

Sun Fire 15K/12K Systems

Sun Fire™ 12K System & Sun Fire™ 15K System

JUST THE FACTS

Sun Microsystems, Inc.

901 San Antonio Road
Palo Alto, CA 94303
U.S.A. 650-960-1300

November 12th 2002

Sun Proprietary – Internal Use Only

Copyrights

© 2002 Sun Microsystems, Inc., 901 San Antonio Road • Palo Alto, CA 94303 USA. All Rights Reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Sun, Sun Microsystems, the Sun logo, Sun Fire, UltraSPARC, Solaris, Sun Fireplane, Sun GigabitEthernet, Sun HIPPI/P1.0, Sun Enterprise Systems Interface, Sun Management Center 3.0, Sun StorEdge, Sun StorEdge Volume Manager, SunATM, Java, Sun HPC ClusterTools, ONC/NFS, SunNet, Solstice Site Manager, Solstice Domain Manager, Solstice DiskSuite, Solstice Backup, Sun StorEdge, Sun Quad FastEthernet, SunSolve, SunVIP, Sun Enterprise, ServerStart, SunReady, Sun Professional Services, SunSpectrum, StorEdge S1, and SunSpectrum Platinum 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. All other trademarks are trademarks of their respective owners.

Table of Contents - "Sun Fire™ 12K/15K Just The Facts"


Chapter 1 - "Sun Fire 12K & 15K Server Positioning"	
Product Overview: Sun Fire 12K/15K Servers.....	1
Target Markets and Users for the Sun fire 12K and 15K Systems.....	2
Technical Computing Customers.....	3
Performance.....	4
Markets and Applications.....	7
Sun Fire 12K One Minute Summary.....	9
A few Sun Fire 12K Q&As.....	10
Chapter 2 - "Selling Highlights"	
Channels and Support.....	1
Performance - Key Selling Factors.....	1
Chapter 3 - "Reliability, Availability, Serviceability Features (RAS)"	
Definitions.....	1
- Reliability - Rate of failure of a system	
- Availability - Quality of being accessible and ready for use	
- Serviceability - Maintenance and Repair Effectiveness	
SPARC processor error protection.....	3
Data cache error protection.....	3
System interface error protection.....	4
Address interconnect error protection.....	4
Data interconnect error isolation.....	4
Console bus error protection.....	4
Redundant Components.....	4
Redundant CPU/Memory boards.....	5
Redundant expander boards.....	5
Redundant hsPCI assemblies.....	5
Redundant PCI cards.....	5
Redundant MaxCPU boards.....	5
Redundant system controllers.....	6
Redundant system clocks.....	6
Redundant power.....	6
Redundant fans.....	7
Reconfigurable Logic Centerplane (Fireplane).....	8
Automatic System Recovery.....	9
System Controller.....	9
Console Bus.....	10
Environmental Monitoring.....	10
Temperature.....	10
Power Subsystem.....	10
Field Replaceable Unit Identification.....	11
Concurrent Serviceability.....	11
Dynamci Reconfiguration of System Boards.....	12
System controller removal and replacement.....	12
Remote Service.....	13
RAS features.....	13
Chapter 4 - "Enabling Technology"	
Technology.....	1

Chapter 5 - "System Architecture"	
Introduction.....	1
Logic Centerplane.....	1
Address Crossbar.....	2
Response Crossbar.....	2
Data Crossbar.....	2
Address Crossbar.....	2
Boardsets (Expanders).....	2
Expander board.....	3
Slot-0 (Top) boards (Uniboards).....	3
Slot-1 (Bottom) boards.....	4
Control Board set	4
Peripherals.....	6
System Interconnect (Fireplane).....	6
Centerplane Interconnect Levels.....	8
Interconnect Performance.....	9
Bandwidth.....	9
Same-board Peak Bandwidth.....	10
Off-board Bandwidth.....	10
Latency.....	10
Chapter 6 - "System Components"	
System Components.....	1
Component Numbering – FRONT F12K F15K.....	3
Component Numbering – BACK Sun Fire F15K	4
Component Numbering – BACK Sun Fire 12K	5
Sun Fire 12K/15K System Cabinet.....	6
System Power.....	7
System Cooling.....	7
Chapter 7 - "System Management"	
Dynamic System Domains.....	1
Domain Configurability.....	2
Domain Protection.....	2
Domain Fault Isolation.....	3
Domain Errors.....	3
System Domain Implementations.....	3
System Controller (SC).....	5
Chapter 8 - "Software"	
Sun Fire™ 12K/15K Server Software.....	1
The Solaris™ 8 (10/01 or later) Operating Environment.....	2
Software Compatibility Matrix.....	2
Sun™ Cluster 3.0 and the Solaris™ 8 Operating Environment.....	2
Sun Management Center (SunMC 3.0).....	3
Console.....	4
Automated Dynamic Reconfiguration (ADR).....	4
InterDomain Networking (IDN).....	5
Solaris Resource Manager (SRM).....	5
Solaris Bandwidth Manager (SBM).....	6
Chapter 9 – "Capacity on Demand"	
COD 2.0	1

Chapter 10 - "Service and Support"	
Sun Professional Services.....	1
Sun Educational Services.....	5
Sun Support Services.....	7
Chart for Platinum/Gold/Silver/Bronze.....	7
Warranty Offering.....	8
Installation Services.....	9
System Availability Guarantee.....	10
Sun Remote Services (SRS).....	10
Service Facts.....	10
Chapter 11 – "Sun Fire 12K/15K Cluster Platform"	
Highlights.....	1
Cluster Platform 15K/9960 - Specifications.....	1
Chapter 12 – "Additional Resources"	
Configuration Information.....	1
Training.....	2
White Papers.....	3
Sun Fire 12K and 15K Success Stories.....	4
Sun Fire 12K and 15K Documents.....	5
Sun Fire 15K BluePrints.....	6
Related Products.....	7
Glossary – Technical Terminology	Glossary - 1

Sun Fire™ 12K and Sun Fire™ 15K System

CHAPTER 1 Sun Fire 12K & 15K Server Positioning



Product Overview: Sun Fire 12K/15K Servers

The Sun Fire™ 12K and 15K systems are Sun's next-generation servers that go beyond the mainframe. The Sun Fire 15K and 12K are based on the Sun Fireplane interconnect, the copper-based UltraSPARC™ III Cu processors, scalable high-end symmetric multiprocessing (SMP), and the Solaris™ Operating Environment. They are the ideal general purpose application and data servers for host-based or client-server applications such as e-commerce, web and application hosting, online transaction processing (OLTP), decision support systems (DSS), data warehousing, communications services, or multimedia services. They are also exceptionally well suited for server consolidation projects. Additional expandability, scalability, and availability can be accomplished by clustering up to eight Sun Fire 12K/15K systems together and employing Sun™ Cluster 3.0.

The Sun Fire 15K system has true "big iron" attributes, include the most scalable and powerful server in Sun's extensive server product line. This SMP system supports up to 106 processors at 900 MHz each with 8 MB of secondary cache, over ½ TB of memory, 18 I/O hubs supporting 72 PCI slots, tested support for over 250 TB of online disk storage, and a wide range of UNIX(R) application software. The new plane interconnect implementation achieves up to 172.8 GB/sec

Sun Fire™ 12K and Sun Fire™15K System

peak (up to 43.2 GB/second sustained) with an overall I/O of up to 21.6 GB/second sustained.

Those needing access to the performance and throughput of the Sun Fireplane Interconnect, high speed memory subsystem, and performance of the UltraSPARC II Cu processor, but are unable to justify the cost of Sun's largest server can choose the Sun Fire 12K server. Based on the same technology as the Sun Fire 15K, the Sun Fire 12K server offers blazing performance, a more modest complement of resources, yet can be readily upgraded to a Sun Fire 15K This SMP system supports up to 52 processors at 900 MHz each with 8 MB of secondary cache, over ¼ TB of memory, 9 I/O hubs supporting 36 PCI slots, tested support for over 120 TB of online disk storage, and a wide range of UNIX(R) application software. The new plane interconnect implementation achieves up to 86.4 GB/sec peak (up to 21.6 GB/second sustained) with an overall I/O of up to 10.8 GB/second sustained.

In addition to unleashing outstanding power, speed, bandwidth and capacity, Sun takes system availability to the next level. Sun Fire 12K/15K are designed with many fault tolerant features, including automatic system controller failover, plane interconnect failure isolation, redundant communications to system controllers, error-correcting code (ECC) on datapaths, and dual power grids. The systems are designed to be upgraded without disrupting users or halting the system.

The Sun Fire 12K/15K systems incorporate advanced features not found in other UNIX servers. Examples include 5G, Sun's fifth generation Dynamic System Domains, IP multi-pathing, second generation InterDomain Networking, and the ability to reboot up to all 18 domains at a time (9 domains at a time for the 12K).

Dynamic System Domains partition the Sun Fire 15K into as many as 18 secure, fault-isolated domains, each running its own version of the Solaris Operating Environment. As many as 9 domains can also be created on the Sun Fire 12K.

Dynamic Reconfiguration balances resources by resizing domains on the fly. CPU/Memory and I/O can be independently and dynamically assigned and reassigned where they are most needed. For example on the Sun Fire 15K, each of 18 CPU/memory boards and 18 I/O assemblies may be placed independently, thereby allowing the configuration of domains optimized for workloads that are either compute intensive, I/O intensive, or both. This capability allows CPU/memory boards and I/O assemblies to be serviced separately.

Target Markets and Users for the Sun Fire 12K and 15K Systems

The Sun Enterprise family of servers is targeted at strategic markets: manufacturing, finance, telecommunications, government, education, health care, retail, design automation, and oil and gas. The Internet adds a significant level of complexity to enterprise applications which now support all business operations, including those conducted with external business allies and customers. Applications must support business-to-business and business-to-consumer electronic commerce. These applications require a platform that provides scalability, availability, reliability, and security. The Sun Fire 12K and 15K, in combination with the Solaris Operating Environment, possesses these qualities and capabilities, positioning it as the platform of choice.

The Sun Fire 15K system is the highest performing SMP system on the market. The Sun Fire 12K offers ½ the mainframe equivalent MIPS as the Sun Fire 15K and can be upgraded to the Sun Fire 15K with no downtime. The Sun Fire 12 and 15K offer enhanced scalability and

Sun Fire™ 12K and Sun Fire™15K System

performance in a large-scale, centralized, enterprise server for parallel processing of commercial and technical applications. Commercial parallel and technical applications take advantage of the scalability of the Sun Fire 12K and 15K systems along with its standard operating environment and commodity hardware components. Technical parallel applications rely heavily on the floating point performance of the Sun Fire 12K and 15K systems. Commercial parallel applications include large-scale data warehousing, high-volume OLTP, server consolidation, and financial, analytical and e-commerce applications.

OLTP (On-Line Transaction Processing) customers face high-volume issues associated with providing short response times and high availability for thousands of users. The Sun Fire 12K and 15K systems addresses this by providing mainframe-like RAS capabilities and the ability to handle high transaction volumes. Sun has published benchmarks for the Sun Fire 15K that show that it can support 4100 concurrent users with low response times while running SAP R/3 And this result was made with a Sun Fire 15K running with just 76 of its possible 106 processors.

Data warehousing customers appreciate the Sun Fire 12K and 15K system's ability to provide greater levels of delivered bandwidth where fast query performance is desired. SMP architecture is employed, making partitioning of data unnecessary. Additionally, the Sun Fire 15K system's large data volume, commodity RDBMS solutions, and mission-critical high-availability make it an even more attractive solution to their needs. No matter which industry or application type, customers can always find a wide range of available commercial off-the-shelf applications.

The Sun Fire 12K and 15K system supports a larger number of processors, memory, I/O and dynamic system domains than other systems in the Sun Fire Server product line. In general The Sun Fire 12K system is recommended for those customers who require 16 or more processors at the time of purchase, but who will upgrade the system to 24 or more processors within 24 months of the time of purchase, but who do not plan on upgrading the system to a Sun Fire 15K in two years or less. The Sun Fire 15K is recommended for those customers that need 24 or more processors and who are likely to need more then 36 processors in 18 months from the time of purchase. Both systems are also recommended for customers that expect to have a large number of simultaneous I/Os, or plan to partition their system into more than 4 domains.

Technical Computing Customers

Technical computing customers seeking high performance compute servers are typically not divided by markets or applications, but by response time, room for growth, and cost.

The Sun Fire 15K system has a peak performance of up to 106 CPUs x 900 MHz x 2 flops/clock = 190.8 Gigaflops. The Sun Fire 12K has ½ the number of Gigaflops as the Sun Fire 15K which means that it is an ideal entry high end technical computing server. Computationally intensive applications, where the Sun Fire 15K system is preferred, are those that are highly parallelized or those where large numbers of users are accessing particular applications. Common technical vertical markets include CAD, EDA, petroleum, and computational chemistry.

For technical customers, the Sun server family is positioned as listed below. Refer to the High Performance Computing 3.0 Just The Facts for further information.

- Sun Fire™ 3800 server: A flexible workgroup server™ delivering exceptional performance at an affordable price.
- Sun Fire™ 4800 server: A powerful mid-range server with exceptional availability.

Sun Fire™ 12K and Sun Fire™15K System

- Sun Fire™ 4810 server: A highly expandable mid-range server with a compact design.
- Sun Fire™ 6800 server: Expandable, high-performance systems with mission-critical availability and integrated storage.
- Sun Enterprise™ 10000 server: Sun's previous generation UltraSparc™ II based sever. It has very high power and is highly availability for high-performance computing. It incorporates great RAS capabilities, like dynamic system domains, dramatically increasing system availability for compute-intensive applications.

Performance

The Sun Fire 12K/15K architecture is designed to offer balanced system performance. These systems feature outstanding integer and floating-point performance, supporting up to 52 or 106 UltraSPARC™ CPUs processors respectively running at 900 MHz each with 8 MB of secondary cache, . The Fireplane interconnect, the coherent shared-memory protocol used by the UltraSPARC™-III/IV processor generations, supports a sustained centerplane interconnect throughput bandwidth of 43.2GB/second for the Sun Fire 15K and 21.6GB/sec for the Sun Fire 12K. High-speed networking is supported by 10/100/1000 MB Ethernet cards. Fast I/O capability is supported through up to thirty-six 33 MHz and up to thirty-six 66 MHz hot swappable PCI adapters including fast/wide SCSI, UltraSCSI, and fibre channel arbitrated loop (FC-AL) interfaces on the Sun Fire 15K, and up to eighteen 33Mhz and up to eighteen 66Mhz hot swap PCI adapters on the Sun Fire 12K. SBus adapters are not supported on the Sun Fire family of servers.

Benchmarks

The Sun Fire 15K set a world record in the TPC-H with Oracle benchmark. TPC-H simulates the work load that is found in a big data warehouse or in a big DSS Database. On this TPC-H SF1000 benchmark the Sun Fire 15K set a new Best performance and best \$/Perf World records. As can be seen from the following comparison the Sun Fire 15K had roughly a 2x better price/performance than HP Superdome.

System	#CPUs	QphH@1000GB	\$/QphH@1000GB	Avail
Sun Fire 15K	72	18,802.1	\$287	07/17/02
NCR WorldMark	128	18,542.9	\$638	07/27/01
HP Superdome	64	13,160.1	\$563	09/05/01

Note:IBM promised p690 results by Dec-01, but have not published as of March 26, 2002.

The Sun Fire 15K also performed a general DSS (Decision Support Systems) benchmark that shows how well it did with query scaling. In this benchmark the Sun Fire 15K demonstrated a near linear scalability result of 71 times faster on 72 CPUs. This is a very important measure of how well a system performs. When a system scales this well there no diminishing returns, each CPU added to the domain pulls its own weight. It is very hard to design a system that scales this well because everything must be balanced. SMP systems from other vendors in the past have had such a dismal record that it was common to see that doubling the CPU count from say 4-8 CPUs would only raise performance by 34% not the 100% you would like to see. For some additional information problems in scaling with SMP systems check out the following link.

<http://www.win2000mag.com/Articles/Index.cfm?ArticleID=3781&pg=2>

Sun Fire™ 12K and Sun Fire™15K System

The Sun Fire 15K, also set another world record in the SPECjbb2000 benchmark by achieving 404,472 JBB operations per second. In this benchmark the Sun Fire 15K beat IBM's 32-way p690 result of 339,484 JBB operations/sec. The SPECjbb2000 benchmark is designed to simulate eCommerce types of applications.

In other benchmarks the Sun Fire 15K has demonstrated a delivered IO bandwidth that is twice the IBM p690's peak IO bandwidth. Sunfire 15K with 72 CPUs delivered and I/O Bandwidth of 10.3GB/sec with a current Peak I/O Bandwidth of 14.1GB/sec. IBM's P690 with its maximum complement of 32 CPUs has not commented yet on what its Delivered I/O Bandwidth number is yet but its Current Peak I/O Bandwidth is 4.8GB/sec. The current peak bandwidth on the Sun Fire 15K was calculated at 8 FCALs each at 100 MB/s for 18 Hot swap PCI (IO) assemblies: $8 \times 100 \times 18 = 14400$ MB/s or 14.1 GB/s.

High-end IO performance is important for data warehouse environments. The test consisted of one Sun Fire 15K using a total of 188 StorEdge T3+ partner-pair arrays with 18 72GB drives each. PCI FC Host Adapter cards were connected with 71 IO slots. This Sun Fire 15K configuration delivered 10.3 GB/s using 3384 drives. on an IO test designed by Sun's Strategic Application Engineering (SAE) group. This IO test simulated a DSS or HPC - High performance computing (HPC) IO workload by doing large sequential reads from raw disk partitions. This test was designed to simulate how Oracle 9i or 8i behaves when running DSS queries. 10.3 GB/s is a huge amount of delivered IO bandwidth.

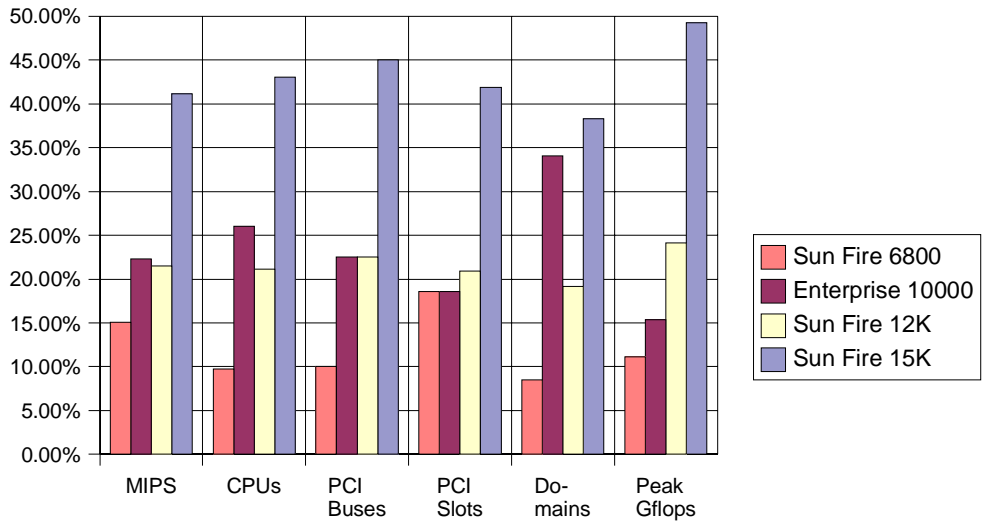
Since the Sun Fire 15K and 12K compete with IBM's Mainframe z series of servers Sun's SAE group has determined what the "Mainframe Equivalent MIPS" ratings of its various servers. These numbers are useful as a guide in comparing the various systems.

MIPS Rating of various Sun Servers			
	CPUs	CPU speed	MIPS
Sun Fire 12K	36-way	900Mhz	3350
IBM Z900	16-way		2694
Sun Fire 6800	24-way	900Mhz	2300
IBM z900	10-way		1934
IBM S/390 G6	12-way		1644
IBM S/390 RY26 G5	10-way		1069

The following table and its associated Chart are useful determining at a glance the various strengths of Sun's high end servers.

	MIPS	CPUs	PCI Buses	PCI Slots	Domains	Peak Gflops
Sun Fire 6800	2300	24	8	32	4	21600
Enterprise 10000	3465	64	18	32	16	29824
Sun Fire 12K	3350	52	18	36	9	46800
Sun Fire 15K	6400	106	36	72	18	95400

Systems Comparison



Performance Type	Sun Fire 6800 Server	Sun Fire 12K/15K System
Processor performance • SPECint_rate2000 • SPECfp_rate2000	24 CPUs @ 900Mhz ● 118 Base 134 Peak ● 114 Base 131 Peak	72 CPUs @ 900Mhz 354 Base 402 Peak 15K 177 Base 201 Peak 12K 342 Base 393 Peak 15K 171 Base 197 Peak 12K
TPC-H benchmark (1000 GB)	4735.7 QphH @ 1000GB \$581 \$/Qph @ 1000GB DB2 with UltraSPARC III 750MHz processors.	F15K - 18,802.1 QphH@1000GB \$287 \$/QphH@1000GB On 72 UltraSPARC III Cu 900Mhz processors
Sustained system bus throughput	9.6 GB per second	43.2 GB per second F15K 21.6 GB per second F12K
Pin-to-pin latency	approximately 200 ns	approximately 313 ns
Current Networking performance	Up to 1000 Mb per second	Up to 14.1 GB per second F15K Up to 7 GB/sec F12K
I/O performance	up to 8 33 MHz PCI buses + 8 66 MHz PCI buses	F15K -up to 36 33MHz PCI buses + 36 66MHz PCI buses F12K – up to 18 33MHz and 18 66MHz PCI buses

Markets and Applications

The following chart illustrates how the Sun Fire 15K system fits into the current line of Sun server products.

Product	Positioning	Applications	Markets
Sun Fire 12K/15K Systems	Enhanced scalability, availability, and performance in a large-scale, mission-critical, centralized, enterprise server for commercial and technical parallel processing applications.	Server consolidation Mainframe consolidation E-commerce Data warehousing Data mining Business applications Customer management systems High-volume OLTP Engineering Design automation Analytics/commercial compute intensive Inter/Intranet	Manufacturing Finance Telecommunications Government Education Health care Retail Oil and gas Pharmaceuticals Chemical Internet commerce

Sun Fire 6800	High-end scalable and expandable Sun server, offering the performance and availability required for mainframe-class, mission-critical applications	Data warehousing Data mining Business applications Customer management systems OLTP NFStm software Design automation Analysis and simulation Video	Manufacturing Finance Telecommunications Government Education Health care Retail Oil and gas Pharmaceuticals Chemical Internet commerce
Sun Fire 4800	Excellent 12-way server for databases, business applications, high-performance computing or Server consolidations.	High performance computing, Internet Data Centers	Telecommunications Government Education Retail Internet commerce
Sun Fire 4810	Affordable data center system designed to deliver high performance and high availability for enterprise-wide applications supporting thousands of users. Front access to all components makes this ideal for situations where rear or side access is not possible.	Internet Data Centers	Telecommunications Government Education Retail Internet commerce
Sun Fire 3800	Workgroup	Internet applications Small application Server	Telecommunications Internet commerce

Specifications: Sun Fire 3800, 4800, 6800, 12K and 15K Systems

	Sun Fire 3800	Sun Fire 4800	Sun Fire 4810	Sun Fire 6800	Sun Fire 12K	Sun Fire 15K
CPU/Memory boards	2	3		6	9	18
Maximum CPUs	8	12		24	36 w/o I/O tradeoff 52 with I/O tradeoff	72 w/o I/O tradeoff 106 with I/O tradeoff
Processor speed	750-900 MHz	750-900 MHz		750-900 MHz	900 MHz	900 MHz
Max Number of DIMMs	64	96		192	288	576
Max Memory capacity (with 1 GB DIMMs)	64 GB	96 GB		192 GB	288 GB	576 GB
Centerplane	Active	Passive	Passive	Passive	Active	Active
Repeater boards	0	2		4	NA	NA
Expander boards	NA				9	18

Domains	2	4	9	18		
I/O assemblies (assemblies)	2	4	19	18		
PCI assembly types	hot-swap CompactPCI	PCI and hot-swap CompactPCI	hot-swap PCI	hot-swap PCI		
PCI slots/assembly	6	8 per PCI, 4 per cPCI	4	4		
Max total PCI slot	12	16	32	36	72	
Bulk power supplies	2	3	6	6	6	
Power requirements	100-120 or 220-240 VAC	220-240 VAC				
System controller boards	2					
Redundant cooling	Yes					
Redundant AC input	No		Yes			
Enclosure	Rackmount	Deskside	Rackmount	Sun Fire 6800 cabinet	Sun Fire 12/15K cabinet	Sun Fire 12K/15K cabinet
Room in rack for peripherals?	Yes			No	No	

Sun Fire 12K One Minute Summary:

The following section attempts to briefly contrast the Sun Fire 12K against the SunFire 15K.

The Sun Fire 12K is a new Entry model using proven SunFire 15K technologies:

- The Sun Fire 12K has all the Data Center Reliability and Quality features of the Sun Fire 15K
- The Sun Fire 12K is a new Entry Point in the High End Servers, it has Up to 9 Dynamic System Domains; up to 52 USIII Cu CPUs; up to 288 GB memory
- The Sun Fire 12K has a new Entry Pricing point.
- The Sun Fire 12K has the ability to do a "On-the-fly"/no-downtime upgrade to Sun Fire 15K
- The Sun Fire 12K uses the CPU/Memory for high Uniboard performance at low cost
- The Sun Fire 12K has the same Datacenter-class warranty, support, and service offerings as the Sun Fire 15k.

The Sun Fire 12K server is targeting \$500,000 to \$1,000,000 (US dollars) price band (as measured by IDC) market where Sun has never really had a major product presence before. Sun has always done well in the under \$500k and the greater than \$1M space, and expects similar results with this new product.

An "on-the-fly" upgrade kit from the Sun Fire 12k to the 15k can be purchase. The upgrade kit provides yet another Capacity Management offering from Sun to assist customers in the difficult task of Capacity Planning. Customers can buy what they need, when they need it. Conveniently, this upgrade can be done "on-the-fly", using Sun's Dynamic Reconfiguration technology. The upgrade kit enables additional

systems resources (beyond the Sun Fire 12K limits) by simply DRing them into a running system. Normal operations can continue without interruption or reboot while the upgrade is being performed. DR can also resize and reconfigure individual domains.

Some good uses for the Sun Fire 12K:

- As a Mainframe alternative it has more MIPS than IBM's biggest mainframe at far less cost per MIPS.
- As a Starfire replacement, in particular for those applications that may expand in the future.
- For Server/Mainframe consolidation projects.
- In traditional Sun Vertical markets (Financial, Telco, DB, D/W, etc)
- When a competitive server goes "Casters-Up" and hits the end of its useful life, or when the company that created it phases out its software binary compatibility in future servers.

A few Sun Fire 12K Q&As:

Note: a much more complete FAQ on the Sun Fire 12K/15K can be found at <http://edss.corp/products/servers/datacenter/sunfire15k/faq.html>

Q: Can customers upgrade a Sun Fire 12K to a Sun Fire 15k by system replacement?

A: Optionally, customers who may prefer a complete system replacement can use existing trade-in allowance programs, and receive brand new system equipment. Sometimes customers prefer the security of a new system warranty and latest hardware and software. The costs can be higher with this approach and is not without system interruption.

Q: Is it more expensive to do the "upgrade-on-the-fly" than it would be to buy a Sun Fire 15k at the beginning?

A: Yes, but only marginally. The upgrade from 12k to 15k is about 3%-5% more expensive than if the customer had bought a 15k originally. This 3%-5% difference includes the cost of installation at the customer site. On the other hand, many customers will never need to upgrade the machine, pocketing the purchase price difference and never incurring an upgrade expense. In any case, the time value of money alone makes the difference a "wash". It's a very cost effective insurance policy for SunFire 12K buyers.

Q: Has Sun just "gutted" the SunFire 15K to provide a lower entry price?

A: No. In fact, the Sun Fire 12K has been in planning for well over a year at Sun. A key decision to NOT repackage the system in a smaller footprint was made to preserve key RAS features of the Sun Fire 15K in the 12K system infrastructure. This decision guaranteed a stress free upgrade path for our customers.

CHAPTER 2 Selling Highlights

Channels and Support

The Sun Fire™ 12K and 15K system use the same selling channels as the rest of the Sun server line: direct and indirect worldwide. The principal support provider for warranty or for a SunSpectrum™ support contract is Sun Enterprise Services. The Sun Fire 12K/15K system warranty is one year for the hardware and software which includes 7 x 24 x 365 hardware on-site and telephone support (including holidays) with a 4 hour average hardware response time. Customers are also highly encouraged to obtain a SunSpectrum™ Support Program service contract. This contract goes beyond the Sun Fire 12K/15K warranty and provides for a flexible level of service that allows customers to choose the right amount of service based on their specific needs. Sun™ Remote Services (SRS) monitoring is included as part of the Sun Fire 15K warranty and as part of a Gold or Platinum level service contract. The ServerStart™ system offering is included in the purchase price.

Performance - Key Selling Factors

- Expandability

The Sun Fire 12K and 15K servers expands from entry-level configurations to system configurations that can handle hundreds of terabytes of data and thousands of users. The Sun Fire 12K and 15K systems are configured using the same Ultra Sparc III Cu CPUs and Memory that are found on the Sun Fire™ 3800, 4800, 4810 and 6800 servers. These CPUs run at 900 MHz and above and can be found on the Sun Fire Uniboard, which is often referred to as the Sun Fire CPU/Memory board, or just CPU/Memory Board. The Sun Fire 15K system can have as little as 4GB and up to 576 GB of main memory. The Sun Fire 12 K can have as little as 4GB and as much as 288 GB of main memory. The Sun Fire 15K has been tested with and will support well over 250 TB of online disk storage. The Sun Fire 12K can support over 120 TB of online disk storage. Both the Sun Fire 12K and the Sun Fire 15K systems are designed around an all new high speed interconnect with a sustained bandwidth of 43.2 GB/sec on the Sun Fire 15K, and 21.6GB/sec on the Sun Fire 12K. The I/O subsystem for both servers incorporates 64-bit PCI technology clocked at 33 and 66 MHz.

The CPU/Memory boards in these servers scale independently from the I/O subsystem boards. These I/O boards are called hot swap PCI assemblies which is abbreviated hsPCI assemblies, and sometimes hsPCI or hPCI. This means that a Sun Fire 12K can have as few as 2 and many as 9 CPU/Memory boards and at the same time it can have as little as 1 or as many as 9 hsPCI assemblies. The Sun Fire 15K can have as few as 2 and as many as 18 CPU/Memory boards and as few as 1 and as many as 18 hsPCI assemblies.

Having the I/O scale independently of the CPU/Memory Uniboards allowed Sun to also offer the MaxCPU boards. The MaxCPU board offers 2 additional UltraSPARC III Cu processors and Ecache running at 900Mhz and above, but contain no memory. This allows the Servers to expand from a maximum uniboard count of 36 on the Sun Fire 12K and 72 on the Sun Fire 15K to 52 CPUs on the Sun Fire 12K and up to 106 CPUs on the Sun Fire 15K. These maximum CPU count servers are good for HPC - High performance computing (HPC) style applications that do not require very much I/O but need huge amounts of CPU cycles. Sun

Sun Fire™ 12K and Sun Fire™15K System

always recommends however that the maximum CPU count in the Sun Fire 12K be set at 50 CPUs and 104 CPUs for the Sun Fire 15K to allow for I/O redundancy.

- Scalability

The Sun Fire 12K/15K system is highly modular. Customers can easily configure these systems to meet their application and performance requirements by simply adding UltraSPARC™ Uniboards, memory, hot swap PCI assemblies or PCI cards. The high-throughput Fireplane interconnect technologies and I/O architecture helps eliminate system bottlenecks and provides balanced system performance, even in systems with the maximum number of UltraSPARC™ CPUs and PCI cards. CPUs and hsPCI assemblies scale independently.

- Security

System controllers in the UltraSPARC II high end servers which were external workstations have been moved into the Sun Fire 12K/15K cabinet and are attached directly to the centerplane. Comprehensive security is accomplished through improved integration between the Solaris 8 and above Operating Environment and the Sun Fire 15Ks dual system controllers which can log all operations to a designated log host. All administration is designed to accommodate multiple levels of access, with a clear separation of responsibilities between platform and domain, operators and administrators. Dedicated communication paths on the fireplane between the system controller and every domain make the Sun Fire 12K and 15K much more secure than Sun's previous line of High End servers.

The integration of security functionality into the Sun Fire 12/15K server and the available security best practice resources and services help address today's security requirements. See Services section for Sun Professional Services. See Sun Fire 12K and 15K documents for references to the Sun Fire 15K security white paper and available Sun BluePrint articles on security best practice information.

- Investment protection

The CPU/memory board, UltraSPARC™ III Cu processors, dual inline memory modules (DIMMs), and industry standard PCI cards used in the Sun Fire 3800 -6800 servers are also common to the Sun Fire 12K/15K systems. Therefore, when upgrading to the larger Sun Fire 15K system, customers can move these components from an existing chassis to the new chassis, provided that the CPU speed is 900 MHz or faster, protecting their investment. The Sun Fire 15K system uses the same peripherals in the same expansion cabinets as the rest of the family. Mixed speed UltraSPARC-III processors are allowed between boards in the same domain and they run at their native speed without clocking down to the lowest speed processor. The Sun Fire 12K/15K servers runs the Solaris Operating Environment which protects the customer's software investment. Sun provides full support for existing 32-bit applications; in particular, Sun helps ensure that the 32-bit applications will run without recompilation and therefore will interoperate with 64-bit applications.

For example: One day the Sun Fire 15K is upgraded with 2 new 1.2 GHz Uniboards. A 900 MHz Uniboard is pulled out of that Sun Fire 15K and put in a 6800. A slower Uniboard is pulled out of the 6800 and put into a 4800 which only had 1 Uniboard already. In this case, 3 systems were upgraded with just a board purchase.

The Sun Fire 12K and 15K share all parts in common. Hot Swap PCI assemblies, PCI cards, MaxCPU boards, Uniboards, and power supplies may all be moved between on the fly. Note this feature requires Dynamic Reconfiguration which is targeted to be available on the hsPCI

Sun Fire™ 12K and Sun Fire™15K System

assemblies and MaxCPU boards before the end of Calendar Year 2002.

- Solaris™ Operating Environment Applications

The Sun Fire 12K and 15K require the Solaris™ 8 Operating Environment with a portfolio of over 12,000 available Commercial Off-the-Shelf (COTS) applications. The Sun Fire 12K and 15K servers will run with the Solaris™ 9 Operating Environment when it becomes available. Solaris also is required to run on the two System controllers in each Sun Fire 12K or 15K server.

- Upgrade Program

There is a trade-in program available to move customers to the Sun Fire 12K and 15K systems from Sun's other servers, and from selected servers from Sun's competitors.

- Upgradability

The modular design of the Sun Fire 12K and 15K systems simplifies the task of upgrading to new technologies that will improve performance. The Sun Fire 12K/15K system was designed to support future versions of the UltraSPARC™ processor and hot swap PCI I/O interface technology. The Sun Fire 12K may also be upgraded to the Sun Fire 15K in a procedure that results in no down time to the applications running on the Sun fire 12K as it is being upgraded.
competitors.

- Capacity on Demand

The Sun Fire 15K optionally can allow customers to instantly upgrade their machines via use of the Capacity On Demand 1.1 (COD 1.1) program. COD 1.1 for the Sun Fire 15K server allows customers to purchase a Sun Fire 15K server, with up to eight additional CPUs, which can be activated as needed. COD 1.1 provides rapid access to extra capacity for unpredictable peaks in demand and enables data centers to have "hot spare" CPUs on-site. COD 1.1 is available exclusively on the Sun Fire 15K server.

- Reliability, Availability and Serviceability Features (RAS)

Reliability, availability, and serviceability (RAS) are critical requirements of enterprise customers that deploy mission-critical applications. The reliability, availability and serviceability (RAS) goals for the Sun Fire 12K/15K systems are to protect the integrity of the customer's data while providing uninterrupted service to the end user.

The Sun Fire 12K and 15K RAS focus is on three areas:

- Problem detection and isolation- knowing what went wrong and ensuring the problem is not propagated
- Tolerance and recovery- absorbing abnormal system behavior and fixing it, or dynamically circumventing it
- Redundancy --replicating critical components

To help ensure data integrity at the hardware level, all data and data paths are error correcting code (ECC) protected and all address and response buses are protected by ECC or parity. These checks help ensure that errors are managed appropriately and data is never corrupted.

For tolerance to errors, resilience capabilities are designed into the Sun Fire 12K/15K systems to help ensure that the system continues to operate, even in the unlikely event of an error leading to a degraded mode. As a symmetrical multiprocessing system(SMP), the Sun Fire 12K/15K systems can function with one or more processors disabled. In recovering from a problem, the

Sun Fire™ 12K and Sun Fire™15K System

system is checked quickly to determine the fault to allow minimum downtime.

The systems comes loaded with redundant features, such as redundant fans, power supplies, system controllers and even the Centerplane itself if fully redundant. Optionally a Sun Fire 12K/15K server can be configured with redundant hardware such as Uniboards or hsPCI assemblies to reduce downtime. Most major components of the Sun Fire 12K/15K can be repaired while the system is online and in use.

A detailed presentation is available that covers the RAS improvements in the Sun Fire 12K/15K servers over the previous generation of High End servers and may be downloaded from the Sun Internal Website for delivery to Sun's customers.

Details of Sun Fire 12K and 15K RAS features can be found in Chapter 3 of this document.

• Manageability

- System Management Services (SMS), runs on the system controller in combination with Sun Management Center, provides an advanced graphical interface for management, monitoring and control of multiple Sun Fire 15K servers and other servers in the Sun Fire family.
- SRS supports the customer's systems availability needs by providing the ability for Sun Enterprise Services to react quickly to problems. The service proactively detects error conditions based on sets of defined thresholds. If errors are detected, a corrective approach can be taken in partnership with the customer.
- Solaris WebStart is a browser-based tool used to install a single image of the Solaris Operating Environment and co-packages.
- Solaris Live Upgrade supports installation and reconfiguration of new versions of the operating system while the current system is still running.
- Solaris Resource Manager helps enable the consolidation of multiple applications onto a single Solaris server. It provides the ability to allocate and control major system resources, ensuring service availability for critical enterprise applications, IT-defined groups, and individual users.
- Solaris Bandwidth Manager can be used to manage the bandwidth used by IP traffic. It allows both incoming and outgoing network traffic to be prioritized by different classes of service.
- Automated Dynamic Reconfiguration (ADR) allows one to invoke scripts for tasks such as adding or deleting a system board to or from a domain, moving a board between domains, or for determining the status of a system board.
- RAS features dual system controllers. If one system controller fails, the other automatically takes over. The power supplies are designed to do dual grid power input. This means that the Sun Fire 12K or 15K can run even if one power grid fails. Also, the power supplies are N+1 redundant, so even after a power grid failure one of the six power supplies could fail and the system would be unaffected. The Sun Fire 12K can also be configured with N+3, or Triple Grid power.
- Fans are installed as push-pull partners, so should one fan tray fail, its partner would provide all the cooling necessary to keep the system running. In addition, the fan tray fans are also fully redundant themselves in the fan tray.
- All major parts in the Sun Fire 15K have a feature called Dynamic FRUID (FRU = Field Replaceable Unit). Dynamic FRUID is a small bit of storage that contains information about a part, like MAX temperature or failure information, serial numbers, etc. This allows Sun to diagnose problems much faster when a complex failure

Sun Fire™ 12K and Sun Fire™15K System

analysis must be performed.

CHAPTER 3 Reliability, Availability, Serviceability (RAS)

Reliability, Availability, and Serviceability assess and measure a system's ability to operate continuously and to minimize service interruption. A system's *reliability* reduces failures and insures customer data integrity. A system's *serviceability* provides for short service cycles when component upgrades are necessary or failures occur. High reliability, to avoid failures, and quick serviceability, to recover rapidly from failures together lead to enhanced availability. A system's Availability goes beyond reliability and serviceability though by also including the software that runs on the system. The availability of a system defines continuous accessibility to the functions and applications supported by the system.

Definitions

Reliability

The reliability of the system is characterized by the frequency of system outages. This frequency is typically measured by the Mean Time between System Interruptions (MTBSI). During a system interruption the system is unavailable to run customer applications.

When computers had one of everything, and everything needed to work for the computer to function the reliability of the system initially depended on the reliability of every individual component. Now that systems are constructed from parts that are redundant, an individual component may fail, however, the system will continue to operate. Sun's focus is primarily on providing available systems, and clearly availability is impacted by component reliability, but major gains can be made in availability by pairing up individually reliable components. Clearly we need to build the computer out of reliable individual components.

Availability

Availability is the quality of being accessible and ready for use. The availability of a machine depends not only on the technology, but is also on the environment surrounding the machine. Obtaining optimal levels of availability starts at the core system design and extends to the overall data processing or application architecture.

An environment that promotes availability includes the people who are running the computing platform and the processes surrounding the computing environment. Without taking care of the people and processes, the features built into the Sun Fire family of computers will not lead to the availability a customer can expect. The product dimension of availability is impacted by both the system reliability and the reduction of downtime.

System Availability can be measured in minutes per year of unavailability. The "nine" is also often used. A system with "five nines" availability is expected to be available for 99.999% of the time, which means expected unavailability of 0.001% of the year or a little over five minutes per year.

Serviceability

Serviceability is characterized by the effectiveness of maintenance and repair of the system. The

Sun Fire™ 12K and Sun Fire™15K System

major serviceability metric which attributes to the system availability is the Mean Down Time (MDT) which reflects the average time the system is down due to a system interruption. The Mean Down Time is influenced by the Mean Time to Repair (MTTR) of the individual Field Replaceable Units (FRUs), and the amount of time it takes to reboot the system following a repair.

Given that computer components do break from time to time despite all the efforts made to make them 100% reliable. No matter how much technology improves the reliability, and redundancy techniques improve availability — serviceability is a key component. The ease with which the computer is mended and made functional directly impacts the availability of the machine.

Since its inception, Sun has established a consistent trend in delivering increasingly modular and serviceable systems. Improving everything ranging from the reduction in slot dependencies, to tighter integration and greater environmental tolerances, Sun has made the time required for the replacement of a failed module much shorter. These enhancements, coupled with improved diagnostic capabilities, have significantly reduced the service cycle on systems, and simultaneously increased their reliability and availability.

Highly-available vs fault-tolerant

Fault-tolerant operation means that a system will continue to operate and provide service, despite any single fault. Fault tolerant machines are considerably more complex to design and produce than highly available machines.

The Sun Fire family is designed to be highly available — there will be unavailability during a diagnosis and reconfiguration following the failure of a limited range of hardware component. This can be contrasted with a fault tolerant machine, which should maintain service throughout a fault. Sun Fire will withdraw service during the fault and restart service automatically after deconfiguration of the faulty component from the machine.

The Sun Fire 12K and 15K servers are designed with full redundancy and no single points of hard failure. This means that no hardware failure can cause a “hard” down, i.e. a down that prevents the machine from coming back up without being serviced. For the Sun Fire 12K and Sun Fire 15K should certain hardware elements fail, such as a CPU, the domain that the CPU was in will go down. The System will then test the components in the domain and will black list out any component that fails to pass the verification self tests. The domain will then come back up following the failure. Using properly configured hardware and Sun Cluster software customers can make their applications truly fault tolerant with availability that reaches 5-9's. Customers are strongly encouraged to discuss methods of maximizing the availability of their Sun Fire servers with their Sun Sales Engineers.

Here is an exhaustive list of what will happen with various hardware failures on the Sun Fire 12K and 15K. Please note that most, if not all of these items listed as failures can be avoided with proper configuration of the server or by using clustering software.

Survive component failure

- Instruction cache tag or data single-bit error (parity & retry)
- Data cache tag or data single-bit error (parity & retry)
- External cache tag or data single-bit error (ECC)
- DRAM or data interconnect single-bit error (ECC)
- Fan (redundancy)
- Power Supply (redundancy)
- System Controller board (failover)
- Network Interface Card (with IPMP)
- FC-AL adapter (with MPxI/O)

Sun Fire™ 12K and Sun Fire™15K System

Affects one Dynamic System Domain

- Fireplane address bus single-bit error (detected by parity)
- Memory or data interconnect multiple-bit error (detected by parity)
- DRAM chip complete failure (detected by parity)
- CPU failure not covered above
- CPU/Memory board failure not covered above
- I/O Assembly failure not covered above
- OS panic

Survive component failure

- Coherency Directory cache single-bit error (ECC)
- Single-bit error on address or response centerplane crossbar (ECC)
- Centerplane support board (failover)

Affects one Dynamic System Domain

- Expander board failure
- Steering signal parity error
- Note: does not crash other domains that are doing Interdomain transfers

Affects up to two Dynamic System Domains

[Applies when an expander is split between two domains, i.e. the CPU board is in one domain, and the I/O board in another domain]

- Expander board failure
- Steering signal parity error
- Note: does not crash other domains that are doing Interdomain transfers

Affects all Dynamic System Domains

- Centerplane ASIC failure (reboot in degraded mode),
- Note: ASIC failure must be complete and not partial to affect all Dynamic System Domains

SPARC processor error protection

The processor has ECC protection on its external cache SRAM, and parity protection on the major internal SRAM structures. A parity error on an internal cache structure is corrected by software, ensuring correct operation after a fault.

The external cache data resides on eight high-speed (4 ns) SRAMs. A single-bit error correction and a double-bit error detecting code protects the 64-byte-wide cache lines. Errors during data-cache or instruction-cache fills are recovered by software flushing and invalidation. Errors during system data transactions are corrected by hardware.

The Sun Fire address bus connection between the CPU and the address repeater is protected by parity

The CPU generates both parity and ECC for all outgoing data blocks. The parity is checked by the receiving dual-CPU data switch. The ECC is checked by all data switch units in the path of a transfer. ECC is checked and corrected by the CPU when it receives a data block.

Data cache error protection

The on-chip data cache is a 64 KB, 4-way associative cache with 32-byte lines. The cache contents,

Sun Fire™ 12K and Sun Fire™15K System

physical tags, and snoop tags are separately parity protected. Errors in the data or in the physical tags are corrected in software by invalidating the cache, and retrying the load instruction. Errors in the snoop tags cause the line to be invalidated.

System interface error protection

The system address bus connection between the processor and the address repeater is protected by parity. The processor generates both parity and ECC for all outgoing data blocks. The parity is checked by the receiving dual processor data switch. The ECC is checked by all data switch units in the path of a transfer. ECC is checked and corrected by the processor when it receives a data block.

Address interconnect error protection

The address and response busses and control signals are parity protected to detect single-bit errors. In addition, the address and response crossbars on the Sun Fireplane Interconnect have ECC protection to correct single-bit errors and detect double-bit errors. The ECC corrects single-bit address errors on the fly, and detect double-bit errors. An address parity or uncorrectable ECC error stops execution in the affected domain.

Data interconnect error isolation

The ECC checks done by the Data Switches can identify the source of ECC errors in most cases. A particularly hard case for ECC errors is when a device writes bad ECC into memory. These get detected much later by other devices reading these locations. Since the bad writer may have written bad ECC to many locations and these may be read by many devices, the errors appear to be in many memory locations, while the real culprit may be a single bad writer.

Since the Data Switch ASICs check ECC for all data entering or leaving each device from other devices, the original source of errors can be isolated. For example, a bad writer writing bad ECC to a memory on a different board will result in ECC errors being detected in two Data Switches: the first on the bad writer's board, and the next on the write target's board. The direction and transaction tag information can identify which CPU pair was the source of the error and which device is the target of a bad ECC write.

Console bus error protection

The console bus is a secondary bus to allow access by the System Controller to the inner working of the machine without having to rely on the integrity of the primary data and address busses. This allows the System Controller to operate even when there is a fault preventing main operation to continue. It is common to all domains and is protected by Parity.

Redundant Components

System availability is greatly enhanced by the ability to configure redundant components. All hot-replaceable (swappable) components in the system are, or can be, configured redundantly if the customer desires. Each system board is capable of independent operation. The Sun Fire family is built with multiple system boards and is inherently capable of operating with a subset of the configured boards functional. Redundant system components include:

Sun Fire™ 12K and Sun Fire™15K System

- CPU/memory boards (optional)
- Expander boards
- hsPCI assemblies (optional)
- PCI cards (optional)
- Max CPU boards (optional)
- System controller boards
- System clock boards
- Bulk power supplies
- Fan Trays

Redundant CPU/Memory boards

A Sun Fire 15K system can be configured with up to 18 CPU/Memory boards (Uniboards), and the Sun Fire 12K system can be configured with up to 9 CPU/Memory boards. Each board contains up to four processors and their associated memory banks. Each CPU/Memory board is capable of independent operation. They can be hot-swapped out of running systems, and moved between system domains. The system is inherently capable of operating with a subset of the configured boards functional. The Uniboard may also be moved between any of the Sun Fire 3800-15K systems as desired.

Redundant expander boards

A Sun Fire 15K system is configured with up to 18 Expander boards (Expanders), the Sun Fire 12K with 9 Expanders. Expander boards may be hot swapped out and into the system only if the CPU/Memory board (Slot 0 board) and the bottom hsPCI assembly or MaxCPU board has been removed first.

Redundant hsPCI assemblies

A Sun Fire 15K system can configure up to 18 hsPCI assemblies and up to 9 for the Sun Fire 12K. Each board supports up to four PCI cards. The hsPCI assemblies can be hot-swapped out of running systems, and moved between system domains.

Redundant PCI cards

Standard PCI cards are mounted on the Sun Fire 12K/15K hot swap PCI assemblies by means of a special cassette, which allows them to be safely hot swapped. Systems can be configured with multiple connections to the peripheral devices, enabling redundant controllers and channels. Software maintains the multiple paths and can switch to an alternate path on the failure of the primary. One primary benefit of the Sun Fire 12K and 15K over all other Sun Sun Fire servers is that the PCI cards themselves may be hotswapped. Compact PCI cards however can be hotswapped on the Sun Fire 3800, 4800, 4810 and 6800 servers.

Redundant MaxCPU boards

A Sun Fire 15K system can be configured with up to 17 MaxCPU boards, and the Sun Fire 12K with

Sun Fire™ 12K and Sun Fire™15K System

up to 8 MaxCPU boards. Each MaxCPU board contains two processors and their associated Ecache. Each MaxCPU board is capable of independent operation. They can be hot-swapped out of running systems, and moved between system domains.

Redundant system controllers

There are two system controller boards on the Sun Fire 12K and 15K. The system controller software, SMS, runs over Solaris™. The system controllers via SMS keep check on each other and copy state information between them to allow automatic failover should the active system controller board fail. Two system controllers are standard with every Sun Fire 12K or 15K high end server.

There is a main system controller and a hot-spare system controller. The main system controller is responsible for providing all system-controller resources for the system.

If failures of the hardware or software occur on the main system controller, or failures on any hardware control path (e.g. console bus interface, ethernet interface) from the main system controller to other system devices, then upon detection of these failures, the system controller failover software will automatically trigger a failover to the spare system controller. The spare system controller will assume the role as the main system controller, and take over all the main system controller's responsibilities. The system-controller data, configuration, and log files are replicated on the both system controllers.

Redundant system clocks

Sun Fire family systems have redundant system clocks. Should the clock system on one system controller board, the consumers of the clock lines will continue drawing their resources from the other system controller until such time as downtime can be arranged to replace the failed system controller board.

Each system controller board generates 75 MHz clocks which are separately distributed to the boards and the other system controller. Each board contains circuitry that selects which incoming clock will be buffered and fanned out to local clock loads. Each system board receives a 75 MHz clock, doubles the frequency to 150 MHz, and sends phase-aligned 75 MHz and 150 MHz clocks to local consumers. One master system controller board is selected to supply clocks. The phase detector on the slave system controller board compares its clock against the clock from the master board., and feeds back the error signal to speed up or slow down the slave system controller board.

Redundant power

The processor cabinet uses six dual-input 4 KW bulk power supplies. All six are standard feature. Two power cables go to each supply, so they can connect to separate power grids. These supplies convert the input power to 48 volt DC. All power in the system is commoned together. The power supplies are also N+1 redundant, such that the system can keep on running with a failed power supply. This means that if the Sun Fire 15K is in its maximum configuration that only 5 of the 12 power connections would be required to be active in order to provide all the power necessary to run the machine. For a Sun Fire 12K only 4 power cords going to unique power supplies are needed to provide N+1 redundancy, and 8 power cords for DualGrid with N+1 redundancy. The additional 2 power supplies that come with the Sun Fire 12K base can be used for Dual Grid N+3 redundancy on the 12K or can just be powered off depending on the owners preference. The power supplies can be replaced while the system is in operation. The Sun Fire 12K can also be configured for triple Grid

Sun Fire™ 12K and Sun Fire™15K System

power if desired.

Power is distributed to the individual system boardsets through separate DC circuit breakers. The boardsets each have their own on-board voltage converters, which transform 48 VDC to the levels required by the on-board logic components. Failure of a DC-to-DC converter will affect only that particular system board.

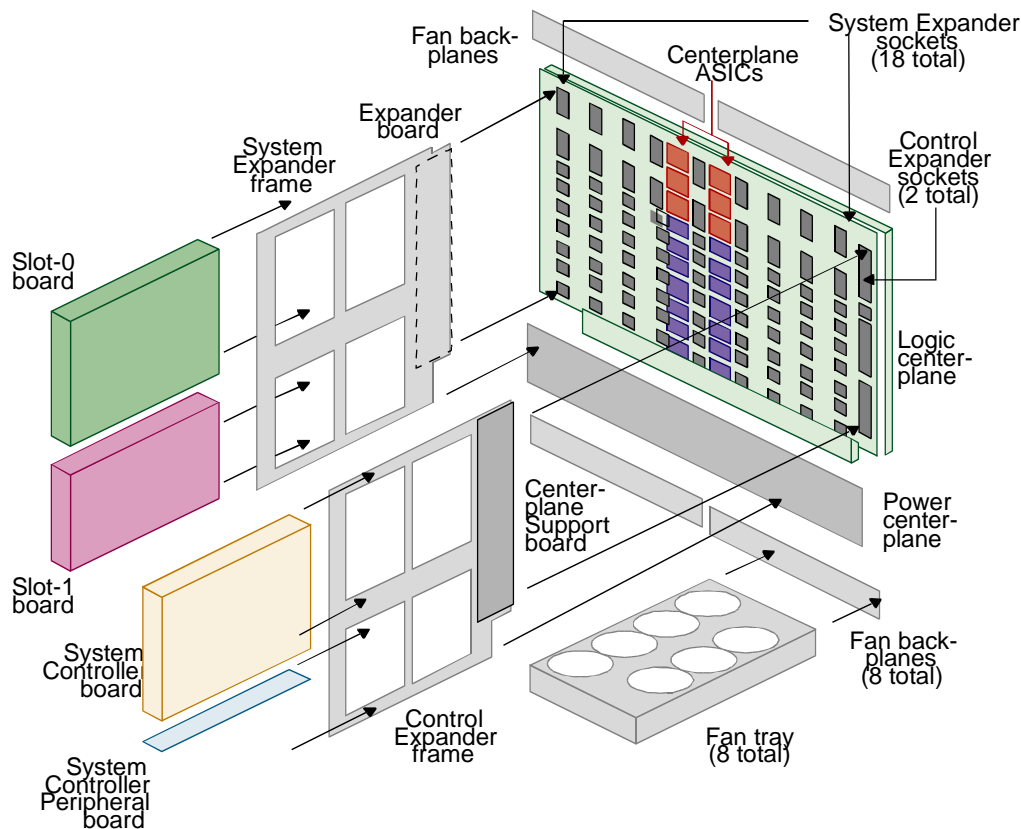
Redundant fans

There are four fan trays above, and four below the system boards. The fans have two speeds, and normally run at low speed. If any of the sensed components get too hot, then the fans are switched to high speed. The fan trays are installed as "push-pull" partners with one on top, and one on the bottom. Should either partner fail, the other partner can provide the full cooling required by the machine from the partner trays. In addition the fans in the fan tray are also fully redundant, such that the fan tray can provide full cooling even if ½ of the fans are disabled in the tray.

Sun Fire™ 12K and Sun Fire™ 15K System

Reconfigurable Logic Centerplane (Fireplane)

The Sun Fire 12K and 15K have three parallel redundant independent crossbars implemented on the logic plane interconnect: one set for addresses, one for responses, and one for data. The plane interconnect contains 20 redundant ASICs, and is the only non-hot-swappable logic component in the system. Since a failed plane interconnect ASIC cannot be removed from a running system, each of the three plane interconnect crossbars (Address, Response, Data) can be independently be moved in and out of degraded mode on the fly. Degraded mode is separately configurable for each system domain. A failed ASIC on the centerplane may result in a all Domain reboot. When this occurs, the machine will reconfigure the centerplane to not use the half of the crossbar that was using the affected ASIC. This means that the machine will go into "double-pump" mode which effectively reduces the machines centerplane speed to ½. However Since the Centerplane has a very large amount of bandwidth for both sustained (43.2GB/sec on the Sun Fire 15K, ½ that on the Sun Fire 12K) and peak loads (172.8 GB/sec on the Sun Fire 15K, ½ that on the Sun Fire 12K) many workloads could be minimally impacted if the system comes back up in double pump mode.



Sun Fire™ 12K and Sun Fire™15K System

Automatic System Recovery

A suitably configured system will always reboot after a failure. The System Controller locates the fault, reconfigures the system excluding the failed processor, memory, or interconnect component, and reboots the operating system.

The System Controller will only configure in parts which have the Fatal Error FRUID Bit clear. Field Replaceable Units (FRUs) that have been already detected as faulty by this or another machine will not be used.

System Controller

The heart of Sun's availability technology is the system controller. It is an Ultra Sparc based stand-alone system running System Management Services (SMS) and the Solaris operating system.

The system controller has access via JTAG to registers in each significant chip in the machine, and continuously monitors the state of the machine. Should there be a problem detected, it is the job of the System Controller to attempt to determine what hardware has misbehaved and then take steps to prevent that hardware from being used until it has been replaced.

The system controller performs the following main functions:

- Configures the system
- Set up the system and coordinate the boot process
- Set up the system partitions and domains
- Generate system clocks
- Monitor the environmental sensors throughout the system
- Handles errors: detection, diagnosis, and recovery
- Provide the platform console functionality and the domain consoles
- Provide routing via syslog of messages to a syslog host

The heart of the system controller (SC) in the Sun Fire 12K and 15K servers is a multifunction, Nordica-based printed circuit board (PCB) which provides critical services and resources required for the operation and control of the Sun Fire system. System boards within the platform can be logically grouped together into separately bootable systems called dynamic system domains, or simply domains. The System Controller can control up to 18 domains simultaneously on a single Sun Fire 15K platform, and 9 domains on the Sun Fire 12K.

The SMS software lets you control and monitor domains, as well as the platform itself. SMS software packages are installed on the SC. In addition, SMS communicates with every Expander board over an Ethernet connection that is etched on to the centerplane.

SMS helps enable the platform administrator to perform the following tasks:

- Administrate domains by logically grouping domain configurable units (DCU) together.
- Dynamically reconfigure a domain
- Perform Automated Dynamic Reconfiguration (ADR)
- Monitor temperatures, currents, and voltage levels
- Monitor and control power to the components
- Execute diagnostic programs

Sun Fire™ 12K and Sun Fire™15K System

- Warns of impending problems
- Notifies when a software error or failure has occurred.
- Monitors a dual SC configuration for single points of failure.
- Automatically reboots a domain after a system software failure
- Keeps logs of interactions between the SC environment and the domains.
- Provides support for the Sun Fire 15K system dual grid power.
- Configurable Administrative Privileges

You perform SMS operations by entering commands on the SC console after remotely logging in to the SC from another workstation on the local area network, or directly via a ASCII terminal session directly with the serial port on the SC. You must log in as a user with the appropriate platform or domain privileges if you want to perform SMS operations (such as monitoring and controlling the platform). Administration tasks on the Sun Fire 12K or 15K system are secured by group privilege requirements. Upon installation, SMS installs 39 UNIX groups to the /etc/group file.

Console Bus

The Console bus is a secondary bus which allows the System Controller to access the inner working of the machine without having to rely on the integrity of the system address and data busses. This allows the System Controller to operate even when there is a fault preventing system operation to continue. It is protected by Parity.

Environmental Monitoring

The System Controller is regularly monitoring the system environmental sensors in order to have enough advance warning of a potential condition so that the machine can be brought gracefully to a halt, thus avoiding physical damage to the system and possible corruption of data.

The environmental items monitored include:

- power state
- voltages
- fan speed
- temperatures
- device failure
- device presence

Temperature

The internal temperature of the system is monitored at key locations as a fail-safe mechanism. Based on temperature readings, the system can notify the administrator of a potential problem, begin an orderly shutdown, or power the system off immediately.

Power Subsystem

DC voltages are monitored at key points within the Sun Fire 12K and Sun Fire 15K. DC current from each power supply is monitored and reported to the System Controller.

Sun Fire™ 12K and Sun Fire™15K System

Field Replaceable Unit Identification

The Field Replaceable Unit IDentification (FRUID) feature on the Sun Fire 12K and Sun Fire 15K is a new and powerful technique Sun has developed to help maximize customer availability.

Each component has an 8 KB SEEPROM chip or area on the component on which is stored 2KB of static, hardware write protected data recording such information as Manufacturing Part Number, Serial Number, Vendor Name, Ethernet Address and Bootbus Timing. The 6 KB of dynamic data holds such information as power-on hours, fatal bit error, hardware level, repair history, temperature log and error log.

The benefits to system availability are:

- The fatal error bit in each FRUID is asserted if the System Controller diagnoses that the FRU is misbehaving. This Fatal Error bit is set until the part has been returned to the repair vendor for investigation or is reset in the field.
- Failure logged twice in succession can be logged, showing that the failure mode was not fixed the first time.
- Trend analysis on serial numbered FRUs allows identification of manufacturing defects.
- Trend analysis on power-on hours, temperature logs etc. enables identification of wearout phenomenon.
- Trend analysis on the vendor that supplied the part to SUN identifies any vendor-specific weaknesses.
- The enhanced configuration information allows selective patching of specific hardware issues by enabling the patching technology to identify the hardware components in the machine and provide only the necessary patches for that configuration.

Concurrent Serviceability

The most significant serviceability feature of the Sun Fire 12K and Sun Fire 15K is to replace system boards online as a concurrent service. Concurrent service is defined as the ability to service various parts of the machine without interfering with a running system. Failing components are identified in the failure logs in such a way that the field replaceable unit (FRU) is clearly identified. With the exception of the plane interconnect, all boards and power supplies in the system can be removed and replaced during system operation without scheduled down time. Replacing the System Controller that is currently active, or switching control to the redundant System Controller can also be done without causing a disruption in the main system operation.

The ability to repair these items without an occurrence of downtime is a significant contributor in achieving higher availability. A by-product of this online repairability of the Sun Fire family concerns upgrades to the on-site hardware. Customers may wish to have additional memory or an extra I/O controller. These operations can be accomplished online with users suffering only a brief (and minor) loss of performance while the system board affected is temporarily taken out of service.

Concurrent service is a function of the following hardware facilities:

- All plane interconnect connections are point-to-point making it possible to logically isolate system boards by dynamically reconfiguring the system.
- The Sun Fire family uses a distributed DC power system, that is, each system board has its own power supply. This type of power system enables each system board to be powered on/off individually.
- All ASICs that connect off-board plane interconnect have a loopback mode that enables the system board to be verified before it is dynamically reconfigured into the system.

Sun Fire™ 12K and Sun Fire™15K System

Dynamic Reconfiguration of System Boards

The online removal and replacement of a system board is called dynamic reconfiguration (DR). Dynamic reconfiguration can be used to remove a troubled board from a running system, or to add new boards into an existing system without requiring a system stop or even a process stoppage. For example, the board can be configured in the system even though one of its processors failed. In order to replace the module without incurring down time, dynamic reconfiguration can isolate the board from the system, hot swap it out, and then allow replacement of the failing hardware parts. Therefore, the dynamic reconfiguration operation has three distinct steps:

- Dynamic detach
- Hot replace
- Dynamic attach

Dynamic reconfiguration enables a board that is not currently being used by the system to provide resources to the system. It can be used in conjunction with hot replace to upgrade a customer system without incurring any down time or to move resources from one domain to another domain on the fly. It can also be used to replace a defective module that was deconfigured by the system and subsequently hot removed and repaired or replaced.

Dynamic deconfiguration and reconfiguration is accomplished by the system administrator (or service provider) working through the system controller or from a domain.

- The first step is the logical detachment of the system board. The Solaris operating system's scheduler is informed, for the board in question, that no new processes should start. Meanwhile, any running processes and I/O operations are enabled to complete, and memory contents are rewritten into other memory banks.
- A switchover to alternate I/O paths then takes place so that when the I/O assembly is removed the system continues to have access to the data.
- The next step in the process is for the system administrator to manually remove the now inert system board from the system — a *hot replace* operation. The removal sequences are controlled by the system controller, and the system administrator follows instructions given by software.
- The removed system board is then repaired, exchanged, or upgraded.
- The second half of hot replace is employed to reinsert the new board into the system.
- Finally, the replaced system board is dynamically configured in by the operating system. The I/O can be switched back, the scheduler assigns new processes and the memory starts to fill.

So with a combination of dynamic reconfiguration and hot replace, the system family can be repaired (or upgraded) with minimal user inconvenience. Hot replace minimizes this interval to minutes by on-site exchange of system boards.

An interesting additional advantage of dynamic reconfiguration and hot replace is that online system upgrades can be performed. For instance, when a customer purchases an additional system board, it too can be added to the system without disturbing operation.

DR is available today on the Sun Fire 12K and Sun Fire 15K. Additionally the ability to DR out MaxCPU boards or hsPCI assemblies is also being completed in Phase 3 of the DR rollout. It should be noted that PCI cards can be "DR'ed out of an hsPCI assembly today.

System controller removal and replacement

The hot-spare system controller board can be removed from a running system. While both system controllers may be removed simultaneously from a running Sun Fire 12K and Sun Fire 15K it is not recommended to do so.

Sun Fire™ 12K and Sun Fire™15K System

Remote Service

The optional capability exists for automatic reporting (to customer service headquarters sites) of unplanned reboots and error log information via email.

Every System Controller has remote access capability that enables remote login to the System Controller. Via this remote connection, all System Controller diagnostics are accessible. Diagnostics can be run remotely or locally on deconfigured system boards while the Solaris is running on the other system boards.

Regulatory Compliances

An important portion of RAS is that it also meets governmental regulatory specifications. These ensure the the product is safe and will not interfere with the proper operation of other machines. The Sun Fire 12K and Sun Fire 15K meet or exceed all of the following Regulation requirements:

Safety	UL/CSA-60950, EN60950, IEC950 CB Scheme with all country deviations, IEC825-1, 2, and CFR21 part 1040
Ergonomics	EK1-ITB-2000
RFI/EMC	EN55022/CISPR22 Class A, FCC CFR47 Part 15 Class A, EN61000-3-2, EN61000-3-3
Immunity	EN55024
Regulatory Markings	CE, FCC A, ICES-003 A, C-tick, VCCI A, GOST-R, BSMI, EK, UL/cUL, TUV-GS

Sun Fire™ 12K and Sun Fire™15K System

RAS features

Reliability Features	Availability Features	Serviceability Features
<ul style="list-style-type: none"> • ECC-protected data path • ECC protected plane interconnect address path • Parity-protected address and control signals within a boardset • Highly reliable distributed power system • Environmental monitors and controls • Connectors, cables, and guides all designed for robustness • Point-to-point routers to maintain bus integrity over multi-drop buses. • ECC error logging 	<ul style="list-style-type: none"> • Highly configurable to increase redundancy (I/O, CPU, system boards, system controller, memory, etc.) • Hot-plug components: CPU, I/O cards, power supplies, and fans • Redundant power subsystem • Redundant power supplies • Redundant cooling subsystem • Auto SC failover • Cluster capability to meet high level availability requirements • Multiple operating system support using dynamic system domains 	<ul style="list-style-type: none"> • Hot -swap components: <ul style="list-style-type: none"> • System boards, I/O, power supplies, fan assemblies, system controller boards • Keyed connectors on hot-swappable components • No slot dependencies except for system controllers • Point-to-point plane interconnect connection enables component isolation, making DR possible • Failed components clearly identified and logged • Hot-swap power/cooling modules • SunVTS software

Sun Fire™ 12K and Sun Fire™15K System

Reliability Features	Availability Features	Serviceability Features
<ul style="list-style-type: none"> • POST • BIST logic in all ASICs <p>PCI Bus parity protected</p>	<ul style="list-style-type: none"> • Compatible with commercial battery backup systems • ECC-protected data path • ECC protected plane interconnect address path • ECC-protected interconnect • Redundant system controllers can be configured for auto-failover high availability • Auto reboot: POST isolation of failed components prior to boot 	<ul style="list-style-type: none"> • Several internal self-tests for error reporting • Dynamic reconfiguration for trouble isolation and online repair • Remote Monitoring available and included in purchase price • Intelligent system controller continuously monitors state of the machine
	<ul style="list-style-type: none"> • IP Multi-pathing increases network availability 	<ul style="list-style-type: none"> • Console bus provides secondary path for system controller to access hardware health information
	<ul style="list-style-type: none"> • System controller monitored via Sun Management Center (SunMC) and can be done remotely 	<ul style="list-style-type: none"> • Electronic serial numbers on all active boards
	<p>Factory-configured hot spares: CPU/memory, I/O assembly, expander board</p>	<ul style="list-style-type: none"> • SunMC simplifies configuration and management for the system and also aids in detecting, identifying, isolating, and correcting problems before they impact the system.
	<ul style="list-style-type: none"> • Redundant power cables/connections 	<ul style="list-style-type: none"> • Sun Validation Test Suite to perform system-level diagnostics
	<ul style="list-style-type: none"> • Distributed dual power grid capability 	
	<ul style="list-style-type: none"> • Dynamic Reconfiguration • Automatic System Recover • Dynamic Multi-Pathing 	
	<ul style="list-style-type: none"> • CPU Power Control 	

CHAPTER 4 Enabling Technology

Technology

Four principal areas of technology used in the design of the Sun Fire™ 12K and 15K systems gives Sun a significant competitive advantage. They are:

- The UltraSPARC™ Microprocessor Family

The Sun Fire 12K and 15K servers features the SPARC Version 9 compliant, 64-bit UltraSPARC™ III Cu 900 processor with up to 8 MB of external cache and operating at 900 MHz. The UltraSPARC III processor provides very high integer and floating-point performance to address the needs of the most computationally demanding applications. With 64-bit data and addressing, the UltraSPARC III Cu processors have a number of important features to improve operating system and application performance:

- Ultra-dense manufacturing process – 16 million transistor design (including cache) implemented using 0.15 micron, 6-layer metal CMOS technology operating at 1.8 volts. Package using a 1200-pin (800 signal) ceramic land gate array
- 4-way associative on-chip 64 KB Data and 32 DB instruction cache, with up to 8 MB of external level-two cache through an integrated controller
- Integrated DRAM controller with support for up to 8 GB of memory that can transfer data at up to 2.4 GB/sec
- Interface to the new Fireplane system interconnect supports peak data rates of 2.4 GB/sec.
- Six-way superscalar issue, no-stall 14-stage pipeline
- Larger cache, improved branch prediction, lower cache latency, and higher clock rates combining to double the performance of the UltraSPARC III Cu.
- High efficiency trap management
- 16 K-entry branch prediction array
- Enhanced error isolation and fault diagnosability
- Enhanced VIS™ instructions set with three new instructions for high performance on multimedia and networking applications
- Binary compatibility across UltraSPARC processor line

- Enormous System Bandwidth

The logic plane interconnect is the heart of the Sun Fire 12K and 15K systems It provides a sustained data bandwidth of 43.2 GB/s on 18 boardsets in the Sun Fire 15K and 21.6 GB/s on 9 boardsets in the Sun Fire 12K. The Fireplane interconnect contains three redundant 18 X 18 crossbar routers. System scalability and low latency are a function of having sufficient internal bandwidth between processors, memory and I/O and the crossbar technology, versus bus technology, was clearly proven in Sun's previous generation Enterprise 10000 server. Sun Fire 12K and 15K employs separate crossbar routers for address, response and data. Bandwidth scales up as system hardware is added.

- The Solaris™ (10/01 or later) Operating Environment

Without a stable and well-proven operating system, the best hardware in the world is useless. The Sun Fire 12K and 15K servers includes the industry's leading enterprise Operating Environment, the Solaris™ Operating Environment. Built on the latest UNIX

Sun Fire™ 12K and Sun Fire™15K System

technology, the Solaris™ Operating Environment delivers unparalleled scalability and performance. The Solaris Operating Environment has been enhanced over the past few years to be able to address very large memories and to scale up to 106 processors and beyond, both important features for the Sun Fire 12K and 15K servers. With enterprise integration by design, the Solaris Operating Environment provides easy access to a wide range of computing environments and network technologies. It delivers a competitive advantage to business through networked computing, scalability, and multi-architecture support. The Solaris Operating Environment provides an advanced, superior solution for all customer IT needs, both technical and business. The Solaris Operating Environment is an industrial-grade solution with the performance, quality, and robustness to deliver mission-critical reliability.

The Solaris Operating Environment delivers a unique advantage for mission-critical environments, providing advanced features and functionality that, combined with built-in networking, gives users a high-performance computing environment enabling faster, and more productive work.

The Solaris Operating Environment delivers the power of the Sun Enterprise servers to users through enhanced networking capabilities and performance, graphics and imaging, increased standards compliance, and key operating system management advancements.

The Sun Fire 12K and 15K systems require the Solaris 8 Operating Environment or beyond.

- Custom ASICs

A number of custom integrated circuits were designed and fabricated for the Sun Fire family of servers. This enabling technology represents a huge engineering investment by Sun and promises improvements in reliability, performance, and overall cost. The following list represents some of the latest technology developed for the Sun Fire generation of compute platforms. For additional detail, please refer to the "The Sun Fire 12K and Sun Fire 15K System Architecture" document.

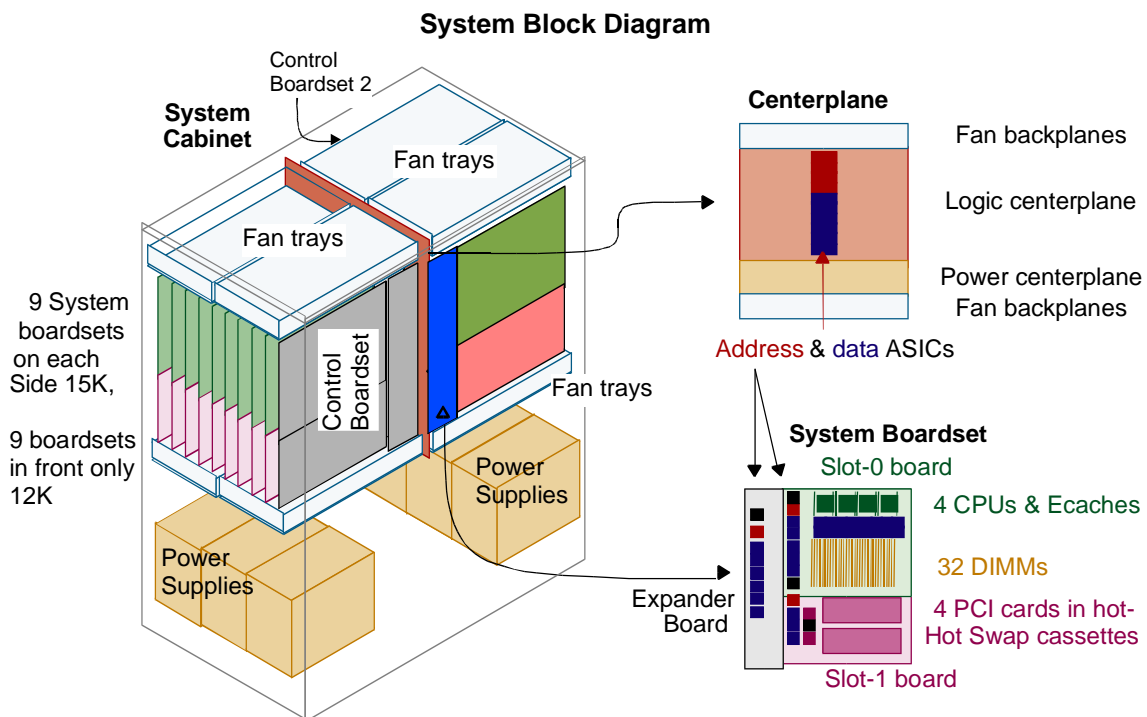
- Slot-0 and Slot-1 Components
 - UltraSPARC™ Processor
 - PCI Controller ASIC (Schizo)
 - Link Controller ASIC (WCI)
 - Address Repeater ASIC (AR)
 - Datapath Controller ASIC (SDC)
 - Dual-Processor Data Switch ASIC (DCDS)
 - Data Switch ASIC (DX)
- Centerplane Components
 - Address Multiplexer ASIC (AMX)
 - Response Multiplexer ASIC (RMX)
 - Data Arbiter ASIC (DARB)
 - Data Multiplexer ASIC (DMX)
- Expander Board Components
 - System Address Controller ASIC (AQX)
 - System Data Interface (SDI) ASIC
- Slot-0, Slot-1, Expander Component
 - Boot Bus Controller ASIC (SBBC)

Sun Fire™ 12K and Sun Fire™ 15K System

CHAPTER 5 System Architecture

Introduction

The Sun Fire™ 12K and 15K systems are comprised of a plane interconnect, system board sets (CPU/memory board, PCI I/O board, plane interconnect expander), and Control Board sets (plane interconnect support board, Control Board, SPARC engine), peripherals, and power and cooling subsystems. These components and their relationships are illustrated in the figure below, and their functions are described in the following pages.



Logic Centerplane

With the rapid movement of processor technology and performance, bus technology has been hard-pressed to keep up. The all-new Fireplane system interconnect reverses this trend by providing superior memory and I/O bandwidth, ensuring balanced and predictable performance under the most demanding loads. The Fireplane system interconnect is the heart of the Sun Fire 12K and Sun Fire 15K. It provides a maximum data bandwidth of 43.2 GB per second between 18 board sets on the Sun Fire 15K, and a maximum data bandwidth of 21.6 GB per second between 9 board sets on the Sun Fire 12K. The logic plane interconnect also delivers a console bus and an Ethernet to each board set.

The plane interconnect contains three redundant 18x18 crossbars:

- Address crossbar

Sun Fire™ 12K and Sun Fire™15K System

- Response crossbar
- Data crossbar

Address Crossbar

The 18x18 redundant Address crossbar provides a path for address transactions between the System Address Controller (AQX) ASIC on each expander board. A pair of unidirectional paths connects to each expander board: one sending, one receiving. Each address transaction requires two clock cycles. The Address crossbar in conjunction with the Response crossbar is one of the major changes over the Enterprise 10000 which uses a bus architecture.

Response Crossbar

The 18x18 Response crossbar provides a reply path between the System Address Controller (AQX) ASIC on each expander board. Each response message takes either one or two clock cycles, depending on the type. The response path is half the width of the address path. A pair of unidirectional paths connects to each expander board: one sending, one receiving. The response path is used for responses from home and slave agents to the original requester, and for completion messages from the requester to the home agent.

Data Crossbar

The 18x18 Data crossbar moves cache line (72-byte wide) packets between the System Data Interface ASICs (SDI) on each expander board. Each connection is a bi-directional 36-byte wide path. The bandwidth is (18 slots) x (32-byte path) x (150 MHz) divided by 2 for bi-directional paths, resulting in a maximum throughput of 43.2 GB per second. To maximize the use of these bi-directional paths, queuing occurs in the Data Multiplexer (DMX) ASICs. The data crossbar is implemented from 12 bit sliced Data Multiplexer (DMX) ASICs, which are under the control of two lockstep data arbiter (DARB) ASICs. If a failure occurs in the data crossbar, it can be put into double-pumped mode using half the datapath. It is possible to move into and out of double-pumped mode during system operation without rebooting.

Board sets

A board set is a combination of three boards that plugs into the plane interconnect.

There are two types of board sets:

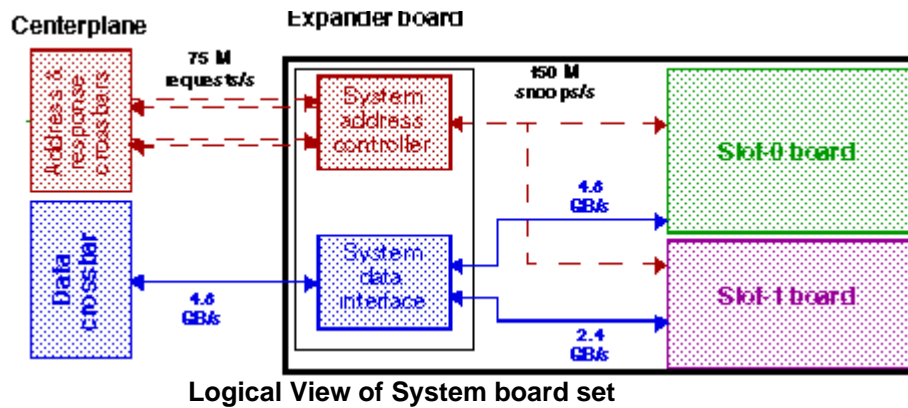
1. System board set. These contains Slot 0 boards with processors and memory, and Slot 1 boards such as the hsPCI assembly or the MaxCPU board.
2. Control Board set is comprised of the Centerplane support board, the SC Peripheral board, the SC CPU board and the Control board.

System board set

A system board set is a combination of three boards, an Expander board, a slot-0 board, and a

Sun Fire™ 12K and Sun Fire™ 15K System

slot-1 board. The board set as a unit is hot-swappable from the plane interconnect, and the slot-0 and slot-1 board are individually hot-swappable from the Expander. Slot 0 boards are CPU/Memory boards. Slot 1 boards are the hsPCI assembly or MaxCPU boards.



Expander board

An expander board expands a plane interconnect slot to accommodate two boards: a slot-0 type and a slot-1 type. It provides a Level-2 Fireplane address bus that can execute 150 million snoops per second. The System Address Controller (AQX) on the expander board recognizes addresses targeted at other board sets, and transmits them across the plane interconnect.

The expander provides a three-port data switch to route data between the slot-0 board, the slot-1 board, and the plane interconnect. This dataport is 36-bytes wide to the plane interconnect and to the slot-0 board, and 18-bytes wide to the slot-1 board. A board set can transfer a maximum rate of 4.8 GB per second in each direction to other board sets.

It is possible to use an expander with only one system board (either slot-0 or slot-1). A system board may be hot plugged into the expander, tested, and configured into a running system, without disturbing the other board. The expander may be hot inserted or removed after its two system boards are removed. The expander also receives a console bus from each of the two system controllers.

18 Expanders come standard in the Sun Fire 15K. 9 Expanders come standard in the Sun Fire 12K. When the upgrade kit is purchased for the Sun Fire 12K to take it to a Sun Fire 15K it will include 9 additional expander boards. These additional expander boards activate the back half of the centerplane in the Sun Fire 12K which is why the Sun Fire 12K may be upgraded to a Sun Fire 15K dynamically.

Slot-0 (Top) boards (Uniboards)

- **System board set: CPU/Memory board (Slot-0)**

Sun Fire™ 12K and Sun Fire™ 15K System

The CPU/Memory board (Uniboard) is the only slot-0 type System board and is common to the Sun Fire 3800, 4800, 4810, 6800, 12K and 15K servers. The CPU/Memory board accommodates four processors. Each processor has an associated memory subsystem of 2 banks of four DIMMs, so memory bandwidth and capacity are both scaled up as processors are added. The memory capacity of the board is 32 GB using 1 GB DIMMs. The maximum memory bandwidth on a board is 9.6 GB per second. The CPU/Memory board has a 4.8 GB (each direction) per second connection to the rest of the system.

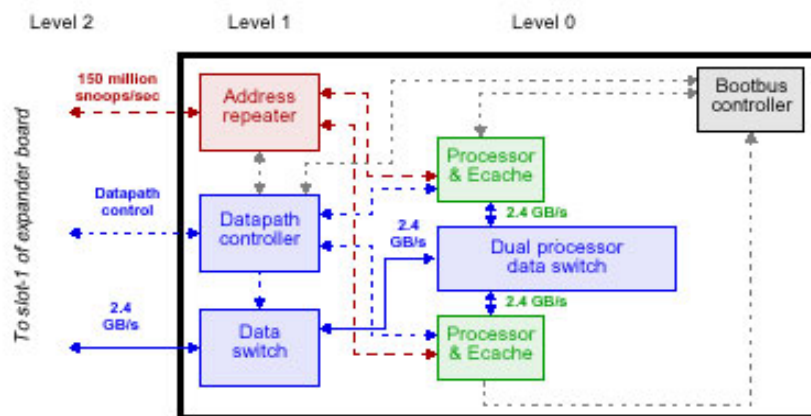
Slot-1 (Bottom) boards

- **System board set: hot swap PCI assembly (hsPCI) (Slot-1)**

The Sun Fire hsPCI assembly has two I/O controllers and is a slot-1 board. It is called the hsPCI assembly since it is hot-swappable. Each controller provides one 66 MHz PCI bus, and one 33 MHz PCI bus, for a total of two of each speed on the I/O board. Each controller has a total bandwidth of 1,200 MB per second. A Sun Fire I/O assembly has a 2.4 GB per second connection to the rest of the system. A cassette is used to provide hot-swap capabilities for industry-standard PCI cards. The cassette is a passive card carrier that adapts the standard PCI edge connector to a pin and socket connector that is suitable for hot-swapping. A PCI card is placed into a hot-swap cassette, and then the cassette is hot-swapped onto the PCI assembly. It relies on software along with the System Controller, for turning power on and off to each PCI slot.

- **System board set: MaxCPU board (Slot-1)**

The MaxCPU board is a slot-1 board. It accommodates two UltraSPARC III Cu processors or better, but does not accommodate main memory. This board allows processors to replace PCI cards in system configurations when more processing power is needed than I/O connectivity or when domains need processors that can be DRed in very quickly.



Logical View of the Dual-CPU board

Control Board Set

The System Control Board set is the heart of the Sun Fire 12K and 15K servers availability and

Sun Fire™ 12K and Sun Fire™15K System

serviceability technology. It provides the critical services and resources required for operation and control of the Sun Fire 15K System. It configures the system, coordinates the boot process, sets up the dynamic system domains, monitors the system environmental sensors, and handles error detection, diagnosis, and recovery. Two System Controller board sets can be configured into the system to provide redundancy and automatic failover in the event that one fails. The operating system used on the Sun Fire 15K is Solaris 8 (10/01) or better, loaded on a hard disk drive.

The Control Board set consists of three boards:

- Centerplane Support board: Plugs into a dedicated plane interconnect slot, and is the same size as an expander board. It provides power/clock/JTAG support for the plane interconnect
- Control Board: Plugs into the Centerplane Support board, and is the same size as a slot-0 board.

Control Board is implemented as a two-board combination:

- An off-the-shelf SPARC engine CP1500 6U high cPCI board with an UltraSPARC-IIi CPU. This runs Solaris and System Management Services (SMS). The SPARC engine cPCI card mounts flat on top of the Control Board, similar to the way PCI cards are mounted onto hsPCI assemblies.
- Control Board, for the Sun Fire 15K-specific logic and connection to the Centerplane Support board.

System Controller: System Clock:

- Driven by 1 of 5 sources.
- Monitors clock input conditions and automatically switches if one is bad.
- Each board comes up running off its own system clocks by default.
- Each board in the system receives a clock from each System Controller .

System Controller: I2C Bus

- The System Controller can access every board, fan, and power supply in the system via an I2C bus.
- Exception is the power plane interconnect which has no I2C components.

System Controller: Console Bus

- The System Controller is the console bus master for the entire system, and either System Controller can be the master.
- Each System Controller uses the CP1500 board for its "brains".
- The CP1500 board uses an UltraSPARC-IIi processor to run the Solaris operating environment and all associated applications needed for bring up, maintenance, and interrogation of the system.
- Two Ethernet ports on the CP1500 are used for the System Management Services (SMS) monitoring software and the other Ethernet port is used for the System Controller-to-System Controller "heartbeat".
- One serial port is used for external access via the front panel, and the other serial port is used for back-door tip access.
- The SCSI port is used to support the boot disk, an optional hard disk drive (HDD), and a DVD drive.

System Controller: Ethernet


- The System Controller has a dedicated Ethernet port for each I/O subsystem.

Sun Fire™ 12K and Sun Fire™ 15K System

- Two “heartbeat” ports exist between the two System Controllers
- Two external ports are provided for the SMS terminal

System Controller: I/O board

- Holds the DVD drive and two HDD drives (one boot device and one optional)
- Provides power to the drives
- I2C will monitor the board’s power supply status
- Provides SCSI port connection from the System Controller to the drives



Logical View of the I/O board

Peripherals

The Sun Fire 3800 through Sun Fire 6800 have room in the same rack with the system enclosure for various peripherals. The Sun Fire 12K/15K cabinet does not currently support additional peripherals inside of the cabinet.

System Interconnect (Fireplane)

The Sun Fire 12K/15K systems uses the Fireplane interconnect which is the coherent shared-

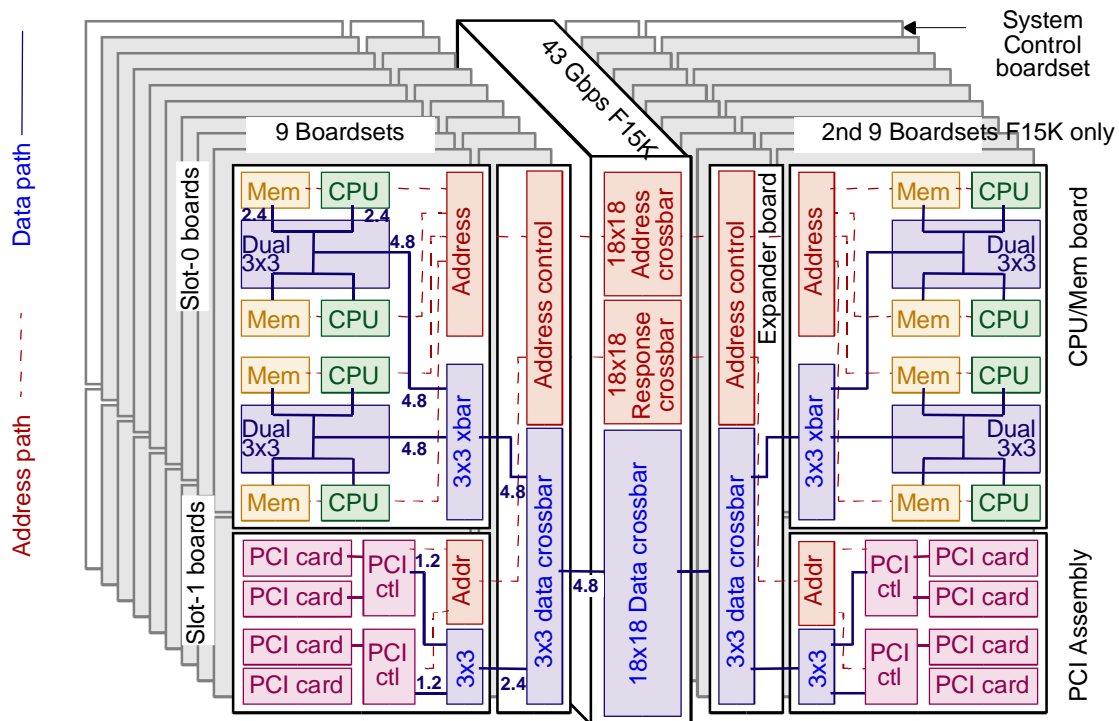
Sun Fire™ 12K and Sun Fire™ 15K System

memory interconnect architecture used in the UltraSPARC-III/IV generation of systems. This is Sun Microsystems' fourth generation of shared-memory interconnect, dating back to the early 1990s. Sun Microsystems uses an improved system interconnect with each new processor generation to keep system performance scaling with CPU performance.

The Sun Fire Fireplane is an evolutionary improvement over the previous generation Ultra Port Architecture (UPA). The system clock rate is increased by 50% from 100 MHz to 150 MHz. The snoops per clock cycle is doubled from one half to one. Taken together, this triples the snooping bandwidth to 150 million addresses per second. The maximum data bandwidth for the Sun Fire 6800 and smaller systems is the snoop rate times the 64-byte coherency block, which is 9.6 GB per second, triple that of the previous generation Sun Enterprise 3500–6500 servers. The bandwidth for the Sun Fire 15K is (18 slots) x (32-byte path) x (150 MHz) divided by 2 for bi-directional paths, resulting in a maximum throughput of 43.2 GB per second, 3.4 times greater than that of the Enterprise 10000. It is ½ that bandwidth for the Sun Fire 12K.

The Sun Fire Fireplane also adds a new layer of point-to-point directory-coherency protocol, for use in systems that require more bandwidth than a single snoopy bus can provide coherency for. This facility allows coherency to be maintained between multiple snoopy buses, and is used in the Sun Fire 12K and 15K.


The small numbers in the Sun Fire 15K logical block diagram below, represent peak bi-directional data bandwidths at each level of centerplane interconnect.



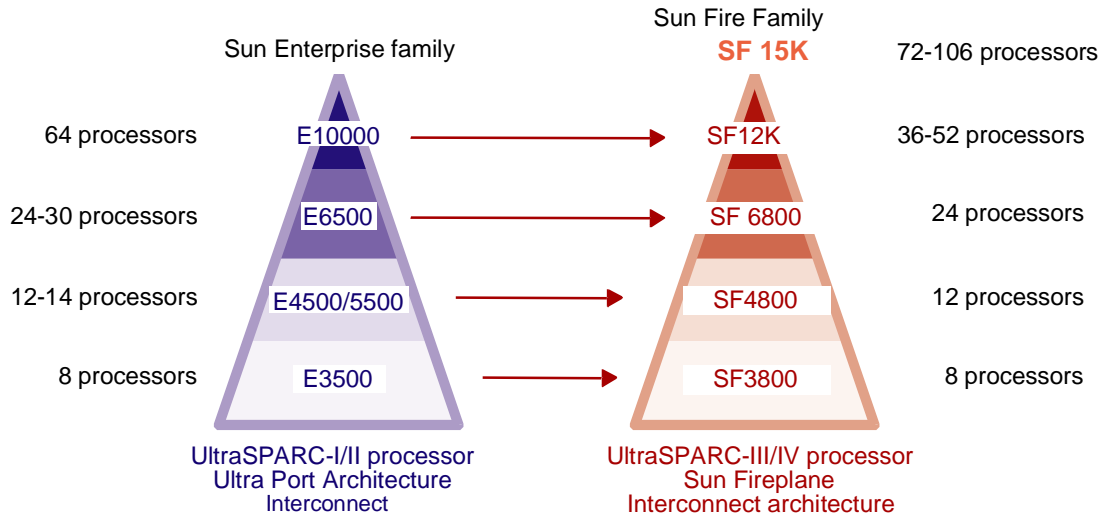
Sun Fire™ 12K and Sun Fire™15K System

Centerplane Interconnect Levels

The Sun Fire 12K/15K interconnect is implemented in several physical layers. The system interconnect of a server is implemented as a hierarchy of levels: chips connect to boards, which plug into a plane interconnect. The latency is lower and the bandwidth is higher between components on the same board, because there are more connections between them than there is to off-board components.



Sun Fire 12K/15K Interconnect Architecture

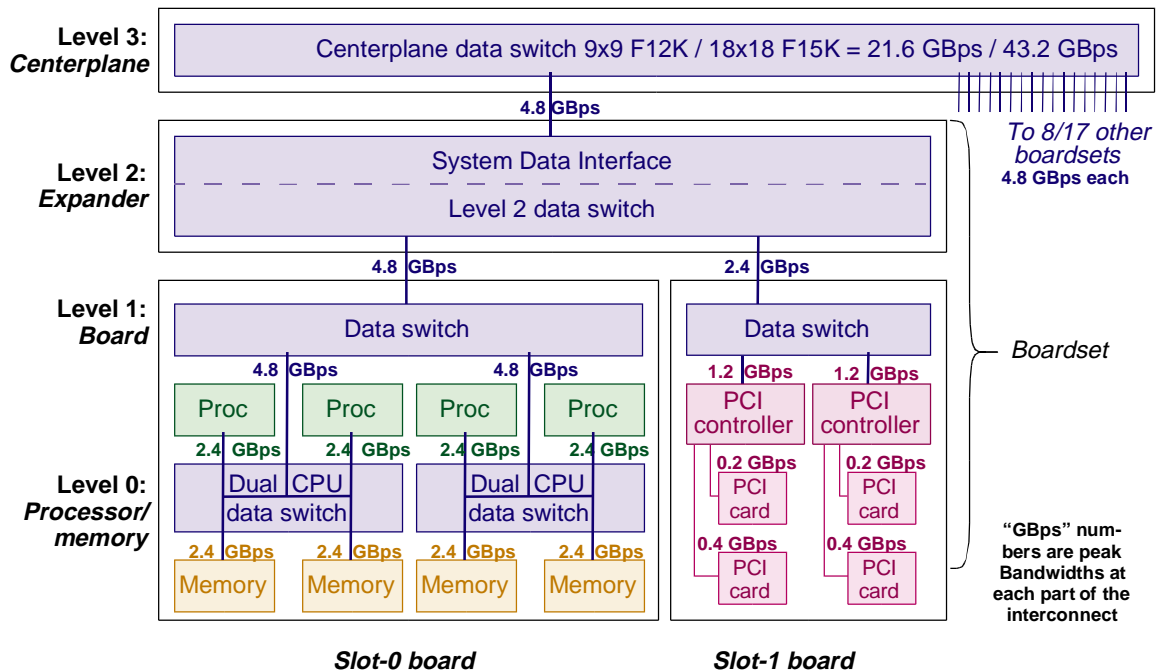


Sun Fire™ 12K and Sun Fire™15K System

Interconnect Performance

This section briefly quantifies the centerplane interconnect latency and bandwidth of the Sun Fire 12K/15K.

Data Interconnect Levels



The numbers in the above figure refer to the peak bandwidth at each level. The Level 3 Centerplane figures represent the peak bisectional bandwidth for the centerplane. All datapaths are bi-directional and the bandwidth on each path is shared between traffic going into and out of a functional unit.

Bandwidth

Bandwidth is the rate at which a stream of data is delivered. These numbers are the peak memory bandwidths, as limited by the centerplane interconnect implementation. Memory is assumed to be interleaved 16-ways across the eight memory banks on one board.

All accesses to memory on:	Sun Fire 12K	Sun Fire 15K
Same CPU as requester	9.6 GB/s x number of board sets, 86.4 GB/s maximum for 9 board sets	9.6 GB/s x number of board sets, 172 GB/s maximum for 18 board sets

Sun Fire™ 12K and Sun Fire™15K System

Same board as requester	6.7 GB/s x number of board sets, 60.3 GB/s maximum for 9board sets	6.7 GB/s x number of board sets, 121 GB/s maximum for 18 board sets
Different board than requester	2.4 GB/s x number of board sets, 21.6 GB/s maximum for 9 board sets	2.4 GB/s x number of board sets, 43 GB/s maximum for 18 board sets
Random data location	23.5 GB/s	47 GB/s

Same-board Peak Bandwidth

The maximum same-board bandwidth is 9.6 GB per second per board. This occurs when: processors all access their own local memory, or all access the memory of the other processor in their pair, or two access their local memory, and two access memory on the other half of the board from themselves. The minimum same-board peak bandwidth is 4.8 GB per second per board. This occurs when all four processors access memory on the other half of the board from themselves. When memory is 16-way interleaved the peak bandwidth is 6.7 GB per second per board. In Sun Fire 3800 through Sun Fire 6800, the total same-board bandwidth is limited to 9.6 GB per second by the address bus snoop rate. Sun Fire 12K and 15K do not have this restriction.

Off-board Bandwidth

The off-board datapath is 32 bytes wide x 150 MHz = 4.8 GB per second. Since this bandwidth has to serve both outgoing requests from this board's CPUs, and incoming requests for memory from other CPUs, the per-board bisection bandwidth is halved to 2.4 GB per second. In Sun Fire 3800 through Sun Fire 6800, the total off-board bandwidth is limited to 9.6 GB per second by the address bus snoop rate.

Latency

Latency is the time for a single data item to be delivered from memory to a processor. There are several kinds of latency that can be calculated or measured. Two are presented below:

1. Pin-to-pin latency: calculated from the centerplane interconnect logic cycles. It is independent of what the processor does with the data.
2. Back-to-back load latency: measured by a kernel of the Imbench benchmark.

Sun Fire™ 12K and Sun Fire™15K System

CHAPTER 6 System Components

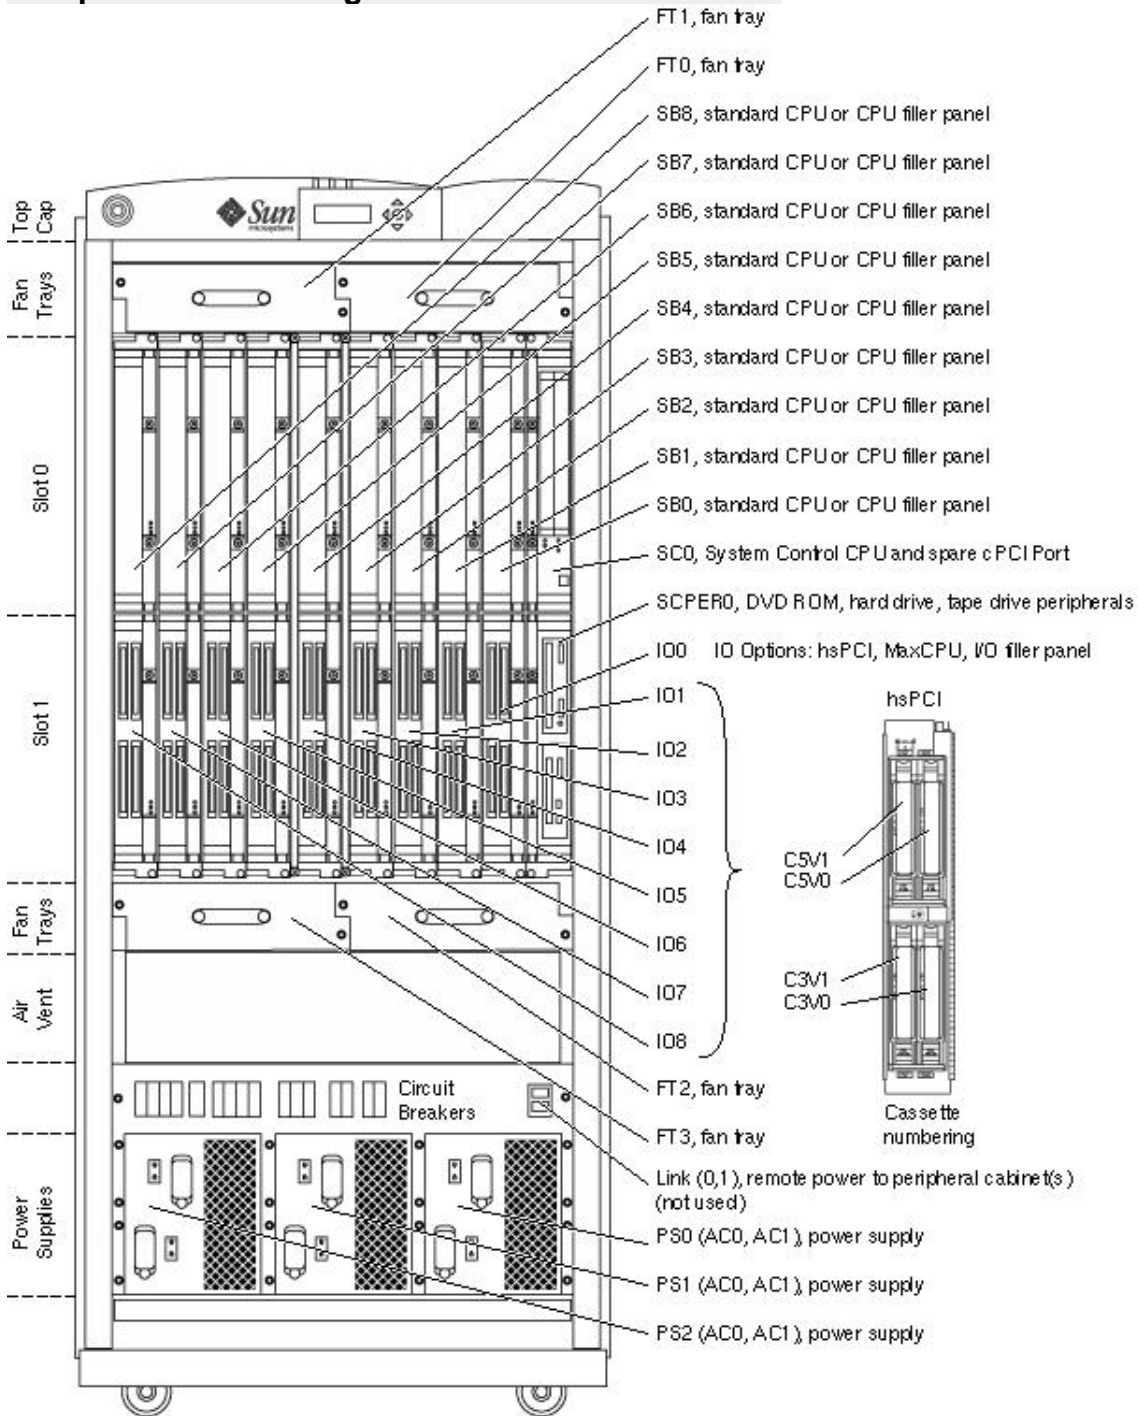
Component	Function	Quantity
Control Board	Set of two boards that are accessible via ethernet or serial interface: <ul style="list-style-type: none"> - SPARC™engine CP1500 6U cPCI board with an UltraSPARC™-Ili embedded system - Control board for Sun Fire 15K specific logic and connection to Centerplane Support Board 	2
SC Peripheral Board	Accommodates a DVD, 2 disk drives, and a 4 millimeter DAT tape drive. Same form factor as a slot-1 board	2
Memory	Four banks of 8 SDRAM DIMMS per CPU/Memory board (1GB DIMMS)	Up to 288 GB F12K and 576GB F15K
I/O	Hot swappable PCI (hsPCI) 2 PCI controllers per slot-1 type PCI Assembly 4 PCI slots total per slot-1 type PCI Assembly <ul style="list-style-type: none"> - 2 at 33MHz - 2 at 66MHz 	Up to 36 F12K Up to 72 F15K
Centerplane (Fireplane Bus)	Fireplane Bus is the coherent shared-memory interconnect architecture which supports: <ul style="list-style-type: none"> - 18 X 18 Data Crossbar Router F15K (9x9 F12K) - 18 X 18 Address Crossbar Router F15K (9x9 F12K) - 18 X 18 Response Crossbar Router F15K (9x9 F12K) - 18 System Boardsets F15K (9 F12K) - 2 System Control Boardsets - Runs at 150 MHz and supports up to 4-52 processors F12K and 4-106 processors F15K 	1
System Cooling	Four fan trays above System and Control Boards (12 fans/tray) Four fan trays below System and Control Boards (12 fans/tray) Base system configured to full fan/cooling capacity	8

Sun Fire™ 12K and Sun Fire™15K System

Component	Function	Quantity
AC power controller	48V bulk power distribution and I ² C distribution	1
Power supply	6 dual-input AC-to-48 volt DC power supplies that can be split across two power girds. Base system configured with full up power	6
Circuit breaker panel	Interrupts power to various components within the system	1
Processor Cabinet	Houses plane interconnect, processors, memory, system controllers 75" high 33.25" wide 65" deep	1
Peripheral Cabinet (Optional)	Houses mass storage devices such as disk and tape drives 75" high 24" wide 42.75" deep	1+
Filters	Cleans the air before it crosses into the fans.	6

Sun Fire™ 12K and Sun Fire™ 15K System


Component Numbering – FRONT F12K F15K



Sun Fire 15K/12K System – Front

Sun Fire™ 12K and Sun Fire™ 15K System


Component Numbering – BACK Sun Fire 15K



Sun Fire 15K System – Rear

Sun Fire™ 12K and Sun Fire™15K System


Component Numbering – BACK Sun Fire 12K



Sun Fire 12K System – Rear

Sun Fire™ 12K and Sun Fire™15K System

Sun Fire 12K/15K System Cabinet



- The Sun Fire 12K/15K system enclosure is a 74.75-inch high, 33.25-inch wide, and 65-inch deep, data center cabinet that is symmetric in the front and the back except for the cap on the top front, and for different doors.
- The processor cabinet is configured with a full complement of eight fan trays, six bulk power supplies, and two system control boardsets which perform RAS services.
- A fully-loaded processor cabinet weighs 2,200 pounds (1000 kg) (F15K). Note: The Sun Fire 12K weighs significantly less however since the Sun Fire 12K can be upgraded to a Sun Fire 15K it is recommended to plan for the maximum possible weight of a Sun Fire 15K when a Sun Fire 12K is purchased.
- Inside the cabinet is an 18-slot card cage for system boards, control boards, and plane interconnect support boards.
- Directly above and below the card cage are the fan trays which draw air up through the cabinet and filters to exhaust out the top. Four trays containing 7 fans are located above the system boards and 4 trays containing 7 fans are located below the system boards.
- The cabinet will house 6 4KW power modules, each with dual AC inputs in order to support dual power grids. 12 power cords will be included to help ensure redundancy.

Sun Fire™ 12K and Sun Fire™15K System

System Power

Site Power

The system runs from 200/240 VAC, single phase power, with a frequency of 47 to 63 Hz. The processor cabinet requires twelve 30-amp circuits, which are usually hooked up to two separate power grids. In North America and Japan, the site power receptacles are NEMA #L6-30R, otherwise they are IEC 309. The power cables that go between the system and the site's power receptacles are supplied with the system.

Processor Cabinet Power System

The processor cabinet uses six dual-input 4 KW bulk power supplies. Two power cables go to each supply. These supplies convert the input power to 48 volt DC, which is commoned together. The system can keep on running with a failed power supply. The power supplies can be replaced while the system is in operation. Power is distributed to the loads through individual DC circuit breakers. The boards each have their own on-board voltage converters, which transform 48 VDC to the levels required by the on-board logic components. Power and cooling components are hot-swappable.

System Cooling


Environmental Requirements

The operating environment limits are - temperature: from 10° to 35° C (50° to 90° F); humidity: from 20% to 80%; and altitude: up to 3,048 m (10,000'). A fully-loaded Sun Fire 15K system draws 24 KW of power (18 KW for the Sun Fire 12K), and has an air conditioning load of approximately 80,000 BTU/hour (54,000 BTU/hour for the Sun Fire 12K). Smaller configurations draw less power. Note: Customers that anticipate that there is a reasonable chance that they will upgrade a Sun Fire 12K to a Sun Fire 15K are encouraged to plan their Environmental requirements for the Sun Fire 12K as if it were a Sun Fire 15K.

Raised Floor Requirements

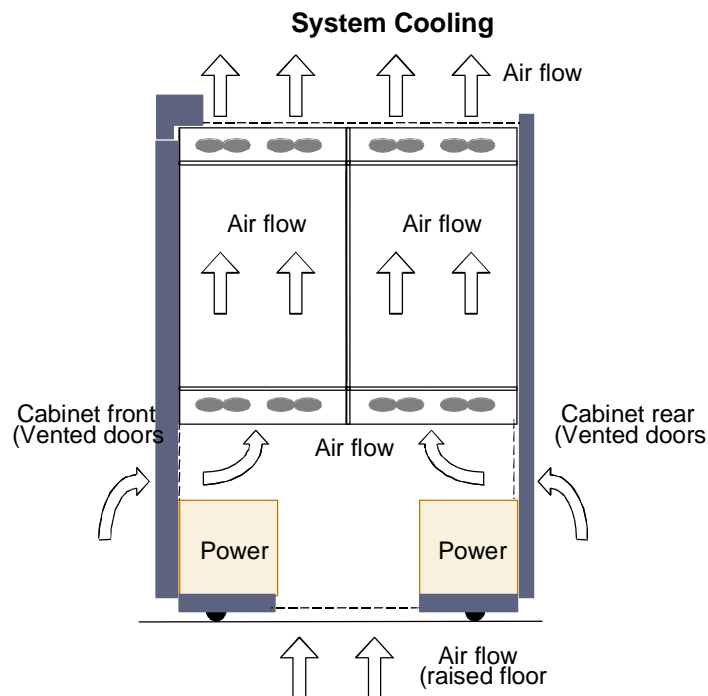
While Sun does not require the Sun Fire 12K or 15K to be installed on a raised floor, it does very strongly, urge its customers to do so. Raised floor environments are more likely to be less subject to air flow blockages, and are generally easier to work with to ensure that hot spots do not occur. When used on a raised floor, a fully-loaded processor cabinet needs seven 2'x2' perforated tiles under it (assuming that the tiles are capable of delivering 600 CFM cooling air at 0.07" of water). Rows of fully-loaded processor cabinets can be located adjacent to each other, with 5' of space per row, as shown below. Though the Sun Fire 12K weighs considerably less than the Sun Fire 15K since it can be upgraded to the Sun Fire 15K customers are strongly urged to plan for the weight and raised floor cooling of the Sun Fire 12K as if it were a Sun Fire 15K.

Sun Fire™ 12K and Sun Fire™15K System



Processor Cabinet Cooling System

- Air flows from bottom to top: in through air inlets in the bottom, front, and back of the processor cabinet, and out through the top.
- There are four fan trays above and four fan trays below the boards.
- The two-speed fans run at medium speed normally. If any of the sensed components exceed temperature thresholds, the fans will switch to high-speed to provide additional cooling.
- The system can run with a failed fan or fan tray
- The fan trays can be swapped while the system is running.



Sun Fire™ 12K and Sun Fire™15K System

The Sun Fire 15K requires boot storage. Sun recommends either the Sun StorEdge S1, Sun StorEdge T3-A, Sun StorEdge T3-B or the Sun StorEdge 5200. These devices need to be installed into a data center rack. The Other storage devices can be used with the Sun Fire 15K, though the ones listed above are the only ones that can be used to boot it.

Complete storage and other I/O device attachment information about the Sun Fire 15K and 12K servers can be found at


http://systems.corp/products/servers/docs/SF_Highend_IO_Support_Matrices_latest.pdf

CHAPTER 7 System Management

Dynamic System Domains

This section describes the dynamic system domains feature of the Sun Fire 12K and 15K systems' interconnect. A Sun Fire 15K system may be dynamically subdivided up into as many as 18 domains, and 9 domains on a Sun Fire 12K. Each domain has a separate boot disk to execute a private instance of the Solaris™ Operating Environment, as well as separate disk storage, network interfaces, and I/O interfaces. CPU boards and hsPCI assemblies may be separately added and removed on the fly from running domains.

Domains are used for server consolidation and to run separate parts of a solution, such as application server, web server, and database server. The domains are hardware-protected from hardware or software faults in other domains. InterDomain Networking (IDN) provides high-bandwidth communication between domains via the plane interconnect, without compromising the isolation between domains. IDN support will be provided in a later release of SMS software.



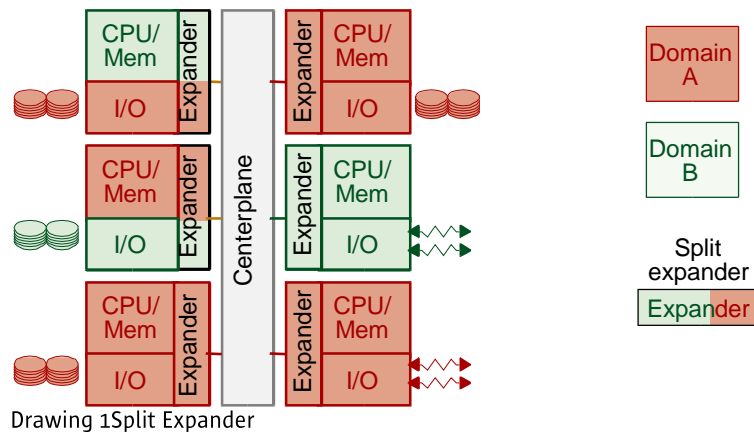
Sun Fire™ 12K and Sun Fire™15K System

Domain Configurability

On the Sun Fire 15K each of the 36 system boards (18 slot-0 boards and 18 slot-1 boards) may be independently added to (or removed from) a running domain. On the Sun Fire 12K each of the 18 system boards (9 slot-0 boards and 9 slot-1 boards) may be independently added to (or removed from) a running domain. This allows processor and memory resources to be moved from one domain to another, without disturbing the disk storage and network connections. Each domain must have an hsPCI assembly, so the maximum number is 18 domains for the Sun Fire 15K and 9 domains for the Sun Fire 12K.

When the two system boards in a boardset are in different domains, it is termed a split expander, since the expander board keeps the transactions separate for each system board. The figure below shows a sample configuration with some of the boardsets split between the two domains. No physical proximity is needed between boards in a domain.

Since split expander hardware is shared between two domains, an expander failure will bring down both domains. Also, memory accesses that go through a split expander take two system clocks (13 ns) longer. If all expanders were split, the load-use latency for accesses to other boardsets would go up about 6%.



Domain Protection

Primary domain protection is done in the system address controller (AXQ) ASICs, by checking each transaction for domain validity when it is first seen. The system data interface chips can also screen data transfer requests for valid destinations (to a granularity of the 36 system boards on the Sun Fire 15K and 18 system boards on the Sun Fire 12K). In addition, each plane interconnect arbiter (data, address, response) screens requests to a granularity of the 18 expanders on the Sun Fire 15K and 9 expanders on the Sun Fire 12K. This is a double-check on the other domain protection mechanisms which are in the system address controller and system data interface chips.

Sun Fire™ 12K and Sun Fire™15K System

If a transgression error is detected in the system address controller, it treats the operation like a request to nonexistent memory. A transgression error in the plane interconnect will cause a domainstop of the transgressing domains, since this must indicate a failure of the primary protection mechanism.

Domain Fault Isolation

Domains are protected against software or hardware faults in other domains. Failures in hardware shared between domains will cause failures only in the domains that share it. Shared hardware includes plane interconnect ASICs, clocks, etc., and the expander board logic serving a CPU/memory board and an I/O board in two different domains.

The steering signals from the system address controllers on the expander boards to the address and response mux chips on the plane interconnect are parity protected. If there is a parity error, the multiple copies of the plane interconnect arbiter could get out of lockstep with each other. So, for a normal transaction, this will cause an immediate domainstop of the domain.

Domain Errors

Nonfatal errors, such as correctable single-bit errors in a datapath, cause a recordstop. History buffers in the ASICs are frozen so that information about the failure can be scanned out via JTAG while the domain continues to run.

Generally, a domainstop shuts down cleanly and quickly only the domain that has encountered the fatal error. When the hardware detects an unrecoverable error, domainstop operates by shutting down the paths in and out of the system address controller (AXQ) and system data interface ASICs. This shutdown is intended to prevent further corruption of data, and to facilitate debugging by not allowing the failure to be masked by continued operation.

The only case where multiple domains are stopped by the same error is when the error is in hardware shared by several domains.

- An error on an expander board shared by two domains will shut down both domains.
- An error in a plane interconnect ASIC, or plane interconnect wiring (connections between two plane interconnect ASICs) may shut down any or all domains.
- Connections between plane interconnect ASICs and expander ASICs are considered a part of that expander's wiring. An error here will shut down only the one or two domains that share the failed expander.

A parity error in the steering signals from an expander board to an arbiter. This error cannot cause the arbiters to so lose lockstep that another domain cannot correctly continue.

The Sun Enterprise 10000 pioneered dynamic system domains (DSD) in the Unix server marketplace. In the Sun Fire generation, domains have been extended to the 3800-15K servers. Each domain is composed of one or more CPU/memory boards and one or more hsPCI assemblies. Each domain runs its own instance of the Solaris™ Operating Environment, and has

Sun Fire™ 12K and Sun Fire™15K System

its own peripherals and network connections. Domains can be reconfigured without interrupting the operation of other domains. Domains can be used for:

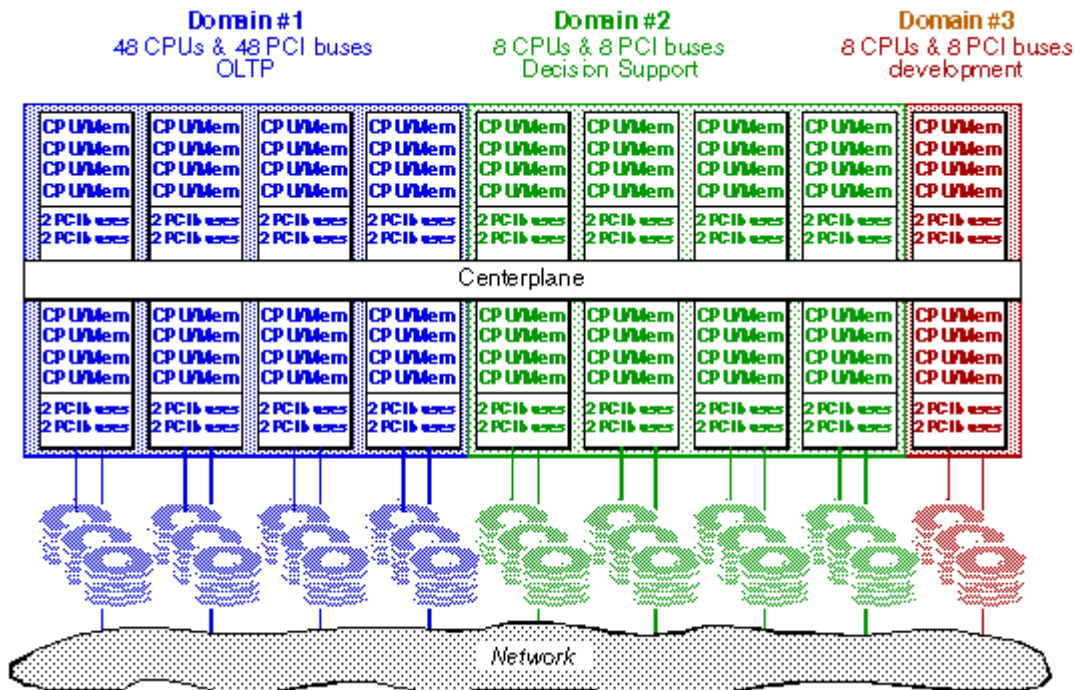
- Testing new applications
- Operating system updates
- Supporting different departments
- Removing and reinstalling boards for repair or upgrade

System Domain Implementations

- **Server consolidation:** One Sun Fire 12K or 15K system can replace two or more smaller servers or mainframes. It is easier to administer (uses a single system controller), more robust (more RAS features), and offers the flexibility to shift resources freely from one “server” to another. This is a benefit as applications grow, or when demand reaches peak levels requiring rapid deployment of additional computing resources.
- **Development, production, and test environments:** In a production environment, most sites require separate development and test facilities. Isolating facilities enables the development work to continue on a regular schedule without impacting production. With a Sun Fire 12K or 15K system, development and test functions can safely coexist on the same platform.
- **Software migration:** Dynamic system domains may be used as a means of migrating systems or application software to updated versions. This applies to the Solaris™ Operating Environment, database applications, new administrative environments, or any type of application.
- **Special I/O or network functions:** A system domain may be established to deal with specific I/O devices or functions. For example, a high-end tape device could be attached to a dedicated system domain, which is alternately merged into other system domains which need to make use of the device for backup or other purposes.
- **Departmental systems:** A single Sun Fire 12K or 15K system may be shared by multiple projects or departments, simplifying cost justification and cost accounting requirements.
- **Configuring for special resource requirements or limitations:** Projects which have resource requirements that might overflow onto other applications may be isolated to their own system domain. For applications that cannot take advantage of all resources (that is, they lack scalability), multiple instances of the application may be run in separate system domains.
- **Data warehouse applications:** Many data warehouses use multiple systems to tier data. The Sun Fire 12K or 15K systems can tier data on the same system and dynamically allocate more resources to individual tiers as needed.

Sun Fire™ 12K and Sun Fire™15K System

Example of a Sun Fire 12K or 15K divided into three domains.



System Controller (SC)

The Sun Fire System Controller (SC) is an embedded computer that resides within each Sun Fire 12K or 15K server. The SC is responsible for many housekeeping and monitoring functions, playing a key role in keeping system availability high and in streamlining management tasks. The SC provides system console access to all domains within the server and the system as a whole. It manages hardware configuration during booting and during dynamic reconfiguration. It is also responsible for system error detection, status logging and reporting, environmental monitoring, identification services, and reporting.

Dual SCs are standard in all Sun Fire 12K or 15K systems, with one configured as a master and one as a slave controller. The slave controller monitors the "heartbeat" of the master to ensure its proper functioning. Should the master controller fail, the slave controller assumes the role of master, with all functions, including the system clock, moving seamlessly from the master controller to the slave. Like other boards in the Sun Fire 12K and 15K servers, the SC boards are hot-swappable, allowing system controllers to be repaired and replaced without requiring any system downtime.

Sun Fire™ 12K and Sun Fire™15K System

Communication with other system boards of the Sun Fire 12K/15K system is handled by the console bus, a dedicated bus used for control and management operations. Communication external to the Sun Fire 12K or 15K server is available via a serial or Ethernet connection.

The SC is an integral part of the operation of the Sun Fire 12K or 15K server. SC functions include the following:

- **Virtual system clock:** The SC provides time of day to all domains within the system, in addition to providing the system clock signals to system boards.
- **Virtual system console:** The SC provides the system console interface to each domain. A connection to the console of a particular domain may be placed in "advise mode", allowing other read-only domain console connections to view all console activity.
- **Virtual key switch:** Domains do not have physical system key switches. Instead, the system console implements a "virtual key switch" for each domain. The virtual key switch allows the system administrator to secure the domain in the same way a physical key would be used. The virtual key switch supports the same four settings used on all Sun servers: off, normal operation, system diagnostics, and secure operations.
- **Power control:** The SC is used to control the 48V power supplies within the Sun Fire 12K or 15K system and to turn individual boards and field-replaceable units (FRUs) on and off.
- **Environmental monitoring and reporting:** Temperature sensors are placed on or adjacent to all critical parts / FRUs in the Sun Fire 12K and 15K systems, with their measurements transmitted to the SC. The SC monitors these readings and takes appropriate action if an over-temperature condition occurs. Environmental conditions are available via SNMP and vital command-line queries.
- **Error management:** The SC acts on errors reported to it from system components via the console bus. Using information from components along the error path, the SC identifies the component(s) generating the error and marks the failed component for system removal at the next reboot.
- **Hardware configuration management:**
 - Automatic system recovery (ASR) management: In the ASR process, the SC checks each component for errors. If the part is faulty, either by failing a diagnostic test or by a notation of a fault in its EEPROM, the part is removed from the system during boot time.
 - Dynamic reconfiguration (DR) and domain creation responsibilities: The SC is responsible for the configuration within the Fireplane system interconnect to identify which boards will make up each domain. Identification can be done when a domain is created or when resources are added to or deleted from a domain. In either case, the system interconnect must be told which boards make up each domain so that it can isolate its activity to one domain at a time.
- **POST Management:** Like previous Sun servers, the Sun Fire 12K and 15K servers run a series of diagnostic tests when they are powered on. These power-on self-tests (POSTs) are more extensive and complete than in previous generations, since testing of domains and boards is scheduled in parallel whenever possible. POST management features and activity include the following:
 - Exercises the Sun Fire 12K or 15K system logic below the FRU level; with a high degree of accuracy, finds failing components and enables isolation to the FRU
 - Provides a highly-available platform for customer applications, even in the face of hardware failures

Sun Fire™ 12K and Sun Fire™15K System

- Provides low-level, start of day configuration services, including detailed interaction with specific hardware components
 - Records sufficient information about failed and marginal components so that both field replacement and subsequent factory repairs are expedited
 - Remembers which components passed the tests, and will configure only those components into the final system configuration. This is possible by using the JTAG access to each of the key Sun Fire 12K or 15K system ASICs.
 - Can be directed to ignore certain components by looking them up in the blacklist. In this way, components scheduled for service, of questionable functionality, or at a specified revision level can be kept out of a system configuration.
 - Has the responsibility for establishing the final system hardware configuration. If there is a failed or blacklisted component, there are usually a variety of ways in which the final system may be configured.
- **Sun Management Center (SunMC) software:** In combination with the optional software, SunMC, the networked SC gives administrators a powerful tool for system and resource management. The SunMC software provides features such as domain management, a GUI interface to dynamic reconfiguration, alternate pathing, and SC commands. SunMC also provides hardware information, environmental monitoring, and propagation of alarms to associated devices.
 - **Capacity on Demand software:** The Capacity on Demand software beyond 1.1 for the Sun Fire 12K and 15K servers runs on the SC.
 - **Remote service:** All functions on the SC that can be performed remotely are set up for terminal execution via dial-up connections.
 - **SunVTS™ software:** SunVTS™ software (Sun Validation Test Suite) is the replacement product for SunDiag™ software. Like SunDiag software, SunVTS software is run at the UNIX® level and is designed to exercise the entire system. It supports either a graphical or TTY user interface and provides error and information logging. Customers should be encouraged to not accept or use any substitute product then the genuine SunVTS as this is critical to the proper diagnostic capabilities of the Sun Fire 12K and 15K servers.

The key features of SunVTS software are:

- UNIX-level diagnostics - System tests execute real UNIX code under the Solaris™ Operating Environment.
 - Automatic system probing - The system configuration is displayed through the user interface.
 - Two user interfaces - A graphical-based interface and a character-based interface are both available. The SunVTS kernel is cleanly separable from the user interface, such that multiple-user interfaces can communicate with the same SunVTS kernel. The character-based interface permits the writing of shell scripts to control SunVTS software.
 - Application Programming Interface (API) - The API provides a defined interface into the SunVTS kernel from other processes, as well as the user interfaces. A SunVTS execution could be initiated in a cron-like fashion, with no direct user interface at all.
 - Advanced configuration and execution control - Tests can be grouped together based on user requirements, with fine-grained execution control for status and logging information.
- **Network Console (netcon):** The SC provides a service called Network Console (netcon).

Sun Fire™ 12K and Sun Fire™15K System

Netcon provides a "console" for single-user operations. Normally, a SC must be on the same subnetwork as the Sun Fire 12K or 15K system; with netcon, they can be anywhere in the world that has a network attachment. The service is provided such that sessions, similar to login sessions, can be provided to X-Windows clients on the same network as the SC. This helps enable system administrators to access the SC from any location on the same network as the SC.

Netcon:

- Helps enable access to the UNIX console, as with the Sun Enterprise 10000 -- shell prompt, ASCII session
- Requires domain administrator authority
- Private internal Ethernet when the Solaris Operating System is running
- IOSRAM "tunnel" protocol when in OpenBoot PROM (OBP)
- Like the Sun Enterprise 10000, has the ability to log all traffic to a file

Sun Fire™ 12K and Sun Fire™15K System

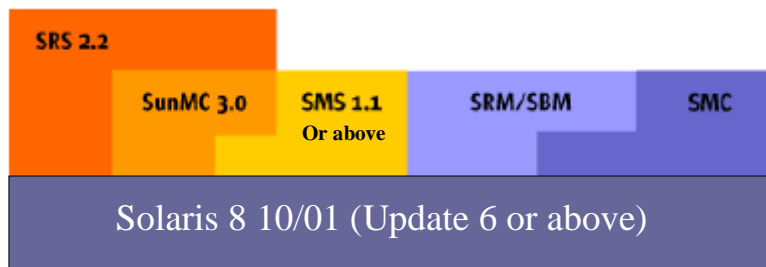
CHAPTER 8 Software

Sun Fire™ 12K and 15K Server Software

Introduction

The software market is constantly evolving, creating new products and services to meet the explosive growth and changing demands of information technology. To manage Sun's products in an enterprise environment, Sun provides software solutions that improve network efficiency, reduce congestion, and manage thousands of system components under a single management solution - satisfying the demands of large corporations.

To operate a Sun Fire 2K or 15K server, users must run the Solaris 8 Operating Environment (10/01 or better) and SMS software. It is also recommended that customers use SunMC and other management services software to enhance server administration and monitoring. The software information is presented as it specifically relates to the Sun Fire 12K and 15K server environments. The following graphic depicts the software stack for the Sun Fire 12K and 15K.



Sun's Sun Fire 12K and 15K software includes:

- **Solaris 8** - provides advanced features and functionality that give users the high-performance operating environment needed for running mission-critical applications.
- **System Management Services** - helps enable users to control and monitor each domain as well as the Sun Fire 12K or 15K system.
- **Sun Management Center** - helps enable users to monitor and manage thousands of Sun systems from a single source. For the Sun Fire 12K or 15K, the software package provides support for the platform and domains.
- **Sun Remote Services** - remote management service that provides proactive problem detection and prompt resolution of system events. This reduces system downtime and provides health and utilization reports that can be an effective IT business planning tool.
- **Solaris Management Console** - client/server application used to manage one or more Solaris domains. It can launch any UNIX application on any Solaris server in the network.
- **Solaris Resource Manager** - helps enable the consolidation of multiple applications onto a single Solaris server. It provides the ability to allocate and control major system resources, allowing service availability for critical enterprise applications, IT-defined groups, and individual users.
- **Solaris Bandwidth Manager** - helps enable control of bandwidth assigned to particular applications, users, and departments that share the same intranet or Internet link. It efficiently

Sun Fire™ 12K and Sun Fire™15K System

distributes network bandwidth and ensures service for mission-critical applications.

The Solaris™ 8 (10/01 or later) Operating Environment

The Solaris™ (10/01 or later) Operating Environment provides an advanced, industrial-grade solution for all customer IT needs, both technical and business; it has the performance, quality, and robustness to deliver mission-critical reliability. The Solaris Operating Environment, with enterprise integration by design, provides easy access to a wide range of computing environments and network technologies while delivering a competitive advantage to business through networked computing, unparalleled scalability, and multi-architecture support. It combines new levels of availability and reliability - to help enable continuous, 24x7 uptime - with massive scalability, sophisticated manageability, and advanced security. With the Solaris Operating Environment, enterprises adapting to the Internet age, as well as dot-com businesses adopting the disciplines of the data center, can increase service levels while at the same time reducing IT risks and lowering service costs.

The Sun Fire 12K and 15K servers includes the Solaris Operating Environment 8, which has been enhanced to address very large memories and scale to support hundreds of processors. The Sun Fire 12K and 15K systems run the Solaris Operating Environment 8 with features and benefits that include:

- Scalability enhancements
- Improved ease-of-use features
- Improved reliability, availability, and serviceability (RAS) features
- Sun Management Center 3.x
- Integrated directory services
- Integrated security
- Hot relief
- Network multipathing solution with IP Multipathing (IPMP) and DR support
- Disk Storage multipathing solution with Sun StorEdge™ Traffic Manager Software and DR
- Clustering with Dynamic Reconfiguration support
- True 64 bit computing to address large memory sets
- 32 and 64 bit support
- 12000 ISV applications
- Live upgrade
 - Installing and reconfiguring new versions of the Operating System while the current system is still running.

Software Compatibility Matrix

The latest Sun Fire 12K/15K Software Compatibility Matrix can always be found at:
<http://edss.corp/products/servers/datacenter/sunfire15k/configuration.html>

Sun Cluster 3.0 and the Solaris 8 Operating Environment

Sun Fire™ 12K and Sun Fire™15K System


"Sun Cluster 3.0 software is a new approach to creating a cluster computing environment for the networked data center. Based on abstracting applications and services, such as data storage and network connectivity, from the physical hardware. Sun Cluster 3.0 extends a high-availability (HA) environment to provide a single, logical view of a commercial computing environment from both an application services and administrative perspective."

Sun Cluster 3.0 extends the Solaris 8 Operating Environment with Sun's cluster framework; it helps enable core Solaris services - devices, file systems, and networks - to operate seamlessly across a tightly coupled cluster while maintaining full Solaris compatibility with existing applications. Sun Cluster 3.0 offers a platform that provides high availability (HA) and scalability to everyday Solaris applications through continuous network and data availability. Services that are written to the simple-to-use Sun Cluster 3.0 API can achieve even higher availability as well as scalability.

Summary of Features and Benefits

- Global devices
- Global file service
- Global network services
- Scalable Services
- Failover services
- Faster failover
- Diskless failover
- Simplified, centralized cluster management
- Cluster-enabling of applications
- Eight-node support
- Dynamic Reconfiguration Support (with SC 3.0 Update 2 or above)

Cluster 3.0 Features



Sun Fire™ 12K and Sun Fire™15K System

Sun Management Center

In combination with the Sun Management Center (SunMC), the networked system controller (SC) gives administrators a powerful tool for system and resource management. SMC software provides features such as domain management, a GUI interface to dynamic reconfiguration, alternate pathing, and SC commands. SunMC also provides hardware information, environmental monitoring and propagation of alarms to associated devices.

- SunMC helps enable flexible integration with the major enterprise management packages
- Critical to SRS (Sun Remote Services)
- Global DR across domains
- Local DR within a domain
- Logical view per domain; tree hierarchy including all hardware and OS components
- Alarms and traps
- Monitor master and standby SCs
- Solaris objects in the SCs
- Environmentals of the system controllers themselves
- Platform Management Information Bus (MIB)
- Platform view and domain view
- Hardware monitoring (presence, state)
- Environmental state and events
- Domain management
- System controller monitoring
- SMS processes on the SC
- Sun Fire 12K and 15K specific physical views "...displaying photo-realistic images of hardware components and pointing to components with an associated event, helping to enable administrators unfamiliar with a particular Sun platform to quickly determine which components need to be replaced"
- Support controlled SC

Console

The System Controller provides a service called console. Console provides a "console" for single-user operations. Normally, a System Controller must be on the same subnet as the Sun Fire 12K and 15K system; with console, they can be anywhere in the world that has a network attachment. The service is provided such that sessions, similar to rlogin sessions, can be provided to X-Windows clients on the same network as the System Controller. This enables system administrators to access the System Controller from any location on the same network as the System Controller.

Console enables:

- Access to UNIX console
- Requires domain administrator authority
- Private internal ethernet when the Solaris™ Operating System is running
- IOSRAM "tunnel" protocol when in OBP

Sun Fire™ 12K and Sun Fire™15K System

Automated Dynamic Reconfiguration (ADR)

Dynamic Reconfiguration (DR) and Automated Dynamic Reconfiguration (ADR) allow resources to be dynamically reallocated between domains. DR helps enable a physical or logical restructuring of the hardware components of a system while the system is running. This high degree of resource flexibility allows administrators to reconfigure the system easily in order to meet load or schedule demands. As a system's resources and requirements change from hour to hour or from month to month, DR optimizes the time and money needed to reallocate those resources.

ADR might be used during those times when it is preferable to reconfigure domains on the Sun Fire 12K/15K server automatically based on predetermined schedules, resource or user load, or other system events. Although domains can still be configured manually as needed, ADR 's high-level, automated UNIX scripts allow fine-grained domain management.

ADR raises domain management resource planning and management, and system-wide management to a new level of flexibility and convenience. Additionally, because ADR reduces the degree of human interaction necessary for domain reconfiguration, errors and time can be significantly reduced -in turn, reducing the overall total cost of operations for the system.

InterDomain Networking (IDN)

Note: IDN is a post-release feature of the Sun Fire 12K and 15K

InterDomain Networking has been completely re-written from the original IDN software now offered on the Sun Enterprise(tm) 10000. When IDN is ready for release, there will be an announcement and a detailed description of how IDN works.

Solaris™ Resource Manager (SRM)

Sun's Solaris Resource Manager software offers the means to allocate, control and monitor system resource usage. Previously, this capability was only available in the mainframe environment. With SRM software, system resources such as CPU, virtual memory, and a number of processes can be allocated to users, groups, and applications. SRM software is a Sun key enabler for server consolidation and increased system resource utilization. Critical application availability is increased because priority can be established over less important applications or functions. It is an effective tool for creating and managing shared service environments. Beyond simple time-sharing schemes, it provides fine-grained, hierarchical control of system resources for users, groups, and applications, enabling an equitable distribution of computational resources within a given Solaris system and promoting server consolidation. Solaris Resource Manager software is particularly effective for use in enterprise servers since it can prevent server resources from being usurped by rogue processes, abusive users, and large computational loads.

Features and benefits include:

- Major system resource controls
- Fair-share allocation

Sun Fire™ 12K and Sun Fire™15K System

- Hierarchical control model
- Policy-based resource administration
- Resource reporting

Solaris™ Bandwidth Manager (SBM)

Solaris™ Bandwidth Manager is a policy-based, directory-enabled software solution regulating IP networks' bandwidth usage in LANs and WANs. SBM is a Sun key enabler for providing differentiated classes of services. It helps enable administrators ability to provide high-quality network service by controlling the bandwidth assigned to applications and users, prioritizing traffic, and building advanced bandwidth management policies. These policies make it possible to mirror the business organization into the information system. The policies can then be stored in a directory to be enforced throughout the organization as required.

Solaris Bandwidth Manager allows monitoring of network traffic, receiving detailed network traffic statistics and collecting of accounting information to feed to the organization's billing system. SBM is customizable and extensible; it provides Java and C programming APIs, which simplifies customization and makes possible the integration of SBM with other applications.

Key features and benefits include:

- Monitoring and allocation of IP traffic priorities and bandwidth
- Directory-enabled
- Policy-based management
- Granular usage reporting

CHAPTER 9 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 12K/15K servers is eight CPUs.

With COD 2.0, F12K/15K 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 12K/15K 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 12K/15K is SMS version 1.3 and above. The SMS 1.3 software will be available on January 31, 2003.

The Sun Fire COD uniboards include two configured options - The first COD uniboard option contains 4 unlicensed CPUs with 900 Mhz/8-MB e-cache and 8 GB of memory (8 1GB memory options/total of 32 DIMMs). The second COD uniboard option contains 4 unlicensed CPUs with 1050Mhz/4x8MB Ecache and 8 banks of 1GB memory options (32x256MB DIMMs, 8 GB total for the board. Also, x-option COD uniboards are also available with 900Mhz and 1050MHz 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):

Sun Fire™ 12K and Sun Fire™15K System

<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 COD 1.0 option is offered on the Sun Enterprise 10000 server. The Sun Enterprise 10000 server does not support COD 2.0. The COD 2.0 option is supported on the Sun Fire 12K/15K, and the Sun Fire 3800-6800 servers.

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.

CHAPTER 10 Services and Support

Sun's extensive service portfolio of Sun Fire™ 12K and 15K service solutions is the "how to" for safely delivering the RAS (Reliability, Availability and Serviceability) features of the Sun Fire 12K or 15K server into the customer's mission critical environment. Based on field-proven methodologies, best practices, and experience implementing thousands of Sun Enterprise™ 10000 servers into data centers worldwide, Sun's approach is to develop a solution that can help customers maintain continuous operations with their Sun Fire 12K and 15K systems.

Sun's experts help customers combine the right people, products, and processes at every stage of the Sun Fire 12K or 15K solution implementation.

The Integrated Services Portfolio includes:

- Professional consulting expertise to help customers architect, design and deploy a solution that can securely handle accelerated growth with high performance and availability.
- Learning solutions that help provide the customer's staff with the skills to assess, architect, build and manage scalable data centers through a curriculum designed specifically for their requirements.
- Comprehensive, flexible support services developed to address the unique needs of the data center and to complement the customer's business model and operational support strategy.

Key Customer Benefits:

- Faster time-to-market
- Maximized return on investment
- Maximized system uptime

Sun Professional Services

Implementing the Sun Fire 12K/15K system requires planning and attention to the customer's entire IT environment. Sun Professional Services understands the issues that may arise. By looking at areas throughout the customer's data center that may affect the Sun Fire 12K/15K server implementation, Sun consultants can help to reduce the risk of system outages and prepare the server for future availability and scalability requirements. The following are the specific services offered from Sun Professional Services to assist with the proper implementation of the Sun Fire 12K/15K server.

- **Application Readiness Service for Sun Fire™ 12K or 15K**
- **Data Center Readiness Service Sun Fire™ 12K and 15K**

Application Readiness Service for Sun Fire™ 12K or 15K

Sun's Application Readiness Service (ARS) for Sun Fire™ 12K or 15K is a critical foundation service, providing essential project management and other services necessary to design, configure, implement, and test the Sun Fire™ 12K or 15K Server to the point of application

Sun Fire™ 12K and Sun Fire™15K System

loading. The paramount objective of ARS is to maximize availability. This service builds on the Server Start Installation Service for the Sun Fire™ 12K or 15K by addressing availability, Solaris security and data storage configuration. All are paramount to ensuring a successful integrated implementation for a mission critical environment.

ARS activities and deliverables

Activities include:

- ◆ **Project Management**— including management of Server Start Installation Service
- ◆ **Architecture Planning**— Platform design consulting, implementation, and testing to ensure:
 - **High Availability:** Dynamic Reconfiguration (DR) , network and storage multi-pathing
 - **Recoverability:** Boot disk recovery procedures
- ◆ **Application Readiness**— Application review; patches for a core application set
- ◆ **Solaris Security**— Hardening of system controllers and domains, user policies, and access-control lists
- ◆ **Storage Consulting**— Application data design, including RAID consulting, file system layout, configuration, and testing of up to 1TB external storage

Deliverables include:

- ◆ **Project schedule**
- ◆ **Build Specification Document**
- ◆ **Test Procedure Plan**
- ◆ **HTML-based Runbook**

ARS Benefits

- ◆ SunPS Project Management coordinates all activities and provides a single point of contact from "kick-off" to "sign-off"
- ◆ Implements and integrates the system right, the first time
- ◆ Faster time-to-production for new Sun Fire™ 12K or 15K Server
- ◆ Customizes Sun Fire™ 12K or 15K configuration and implementation to meet your business needs
- ◆ Optimizes availability to meet mission critical demands
- ◆ Secures the server platform to reduce vulnerability to attacks
- ◆ Implements Sun best practices in security and availability
- ◆ Reduces risk and probability of future outages by leveraging Sun's experience
- ◆ Limits the need to develop Sun Fire-specific technical skills in-house
- ◆ Frees up technical resources so you can focus on higher priority business needs

Sun Fire™ 12K and Sun Fire™15K System

Data Center Readiness (DCRS) for Sun Fire 12K or 15K

NOTE: DCRS for Sun Fire 12K or 15K is expected to be available as a fixed price service in November 2002. In the interim this is available from SunPS as a custom time and material offering. Please contact your local Sun representative

For a more complete integration into a business-critical data center environment, we highly recommend that you utilize the Data Center Readiness Service for the Sun Fire 12K or 15K server. The DCRS builds extensively upon the ARS. It fully leverages SunPS best practices and data center implementation experience. The DCRS really *turns the spotlight* on data center operations and practices, showing the way for a sustainable mission critical environment. The service offers a set of core modules— essential services that address platform, application, and operations-management topics impacting availability. You may also choose from seven additional operations and best practices modules that address data center operations, availability, security, and system and management. You can choose three modules to customize a program that meets your unique needs.

This service adds the following to ARS for Sun Fire 12K or 15K:

CORE SERVICES

Application/Platform Consulting :

- Planning and/or testing assistance on applications such as Oracle, SAP, PeopleSoft.
- Advanced Features Platform Consulting (not included in ARS) such as Inter- Domain Networking (IDN), NFS, fssnap, Solaris iChange Deployment Engine. iChange provides the power to perform software updates while servers continue to run, avoiding interruptions to your software services.

Operations Management Workshop.

Sun Blue print topics included:

- Disaster Recovery,
- Change Management,
- System Performance Management,
- Datacenter Naming schemes,
- Production Environment

Runbook Customization

- 5 customized procedures,
- Incorporate build specification and test procedure plan,
- Resource information

CUSTOMER SELECTABLE SERVICE MODULES (up to 3 from 7 modules)

Data Center Operation Modules

- **i-Status Monitoring Tool:** This service installs a dynamic web/query-based monitoring tool to display hardware and software configurations

Deliverable: Installed monitoring tool incorporated into Runbook

Sun Fire™ 12K and Sun Fire™15K System

- **JumpStart Consulting:** This service designs, implements and tests a customized JumpStart server with core applications and patches.

Deliverable: Customized Jumpstart server.

- **Backup & Recovery Assessment:** This service develops and documents an architecture for an existing NetBackup or Solstice environment.

Deliverable: Backup / Restore Architecture Report

- **Storage Consulting:** This service designs, configures, tests & documents up to 10 TB of external storage; it will also include RAID (Redundant Array of Inexpensive Disks) consulting.

Deliverable: Storage Design

Availability Module

This service includes comprehensive configuration and testing of DR and CADR (automatic DR) features of the Sun Fire 12K or 15K.

Deliverable: DR design spec and test procedures

Security Module

The Security Hardening Service will harden a Sun Fire 12K or 15K domain and include application security considerations for that domain

Deliverable: Build specification and test & verification

Systems & Management Module

This service will configure, implement & test the Sun Management Center (SunMC) network management software.

Deliverable: SunMC implementation report with patches, events & alarms

(Note: ARS for the Sun Fire 12K or 15K server is the prerequisite service for DCRS.)

DCRS benefits

- ◆ Customizable service options meet your unique data center needs
- ◆ Comprehensive DR testing increases system availability
- ◆ Establishes Sun best-practices and proven data center procedures, including backup and recovery, which mitigates the risk of business outages
- ◆ Empowers the IT staff to sustain a mission critical environment
- ◆ Leverages Sun security expertise to assess and reduce unauthorized access and system vulnerabilities
- ◆ Deploys dynamic, web-based tools to proactively monitor, manage, and maintain the ideal platform configuration for your environment
- ◆ Streamlines data center operations while establishing consistent installation standards and reducing IT management costs

Sun Fire™ 12K and Sun Fire™15K System

- ◆ Leverages Sun's expertise to design and run a business-critical data center

Additional information regarding Sun Professional Services can be found at: <http://sunps.Central>

Sun Educational Services

Sun offers Sun Fire™ Server, Solaris™ Operating Environment, Security and Systems Management Learning Paths for customers who want to learn more about the Sun Fire 12K or 15K server hardware and software environment. Sun's Education Consulting Services help companies change the way learning takes place in the corporation by creating custom training solutions that allow people to develop the right skills at the right time.

Who can benefit? Senior system administrators; operations managers; data processing planners; and those individuals responsible for installing, configuring, upgrading, and supporting the Sun Fire 12K or 15K server can benefit from these learning solutions.

ES-421: Sun Fire 12K/15K Server Administration: This five (5) day hands-on course, provides the essential information for administering the Sun Fire 12K or 15K server.

The following key areas are discussed:

- Installing and configuring the Sun Fire 12K/15K system software
- Booting the Sun Fire 12K/15K system
- Monitoring and controlling the Sun Fire 12K/15K server
- Configuring the Sun Fire 12K/15K network environment
- Configuring and using Sun Fire 12K/15K system domains
- Understanding the operation and control mechanisms of the Sun Fire 12K/15K system
- Performing Sun Fire 12K/15K dynamic reconfiguration
- Managing Sun Fire 12K/15K system resources using automated dynamic reconfiguration
- Installing and configuring alternate pathing
- Planning and configuring inter-domain networks

Solaris™ Operating Environment Learning Path

Understanding the Solaris 8 Operating Environment is a prerequisite to understanding and harnessing the full power of the Sun Fire 15K server. Sun offers flexible training options ranging from individual courses to certifications. Successful completion of the Sun Fire 12K/15K Administration courseware and certification requires successful completion of the Solaris courseware listed below.

- Fundamentals of Solaris for Systems Administrators
- Solaris System Administration I
- Solaris System Administration II
- Solaris TCP/IP Network Administration
- Solaris System Performance Management

Security Learning Path

Sun Fire™ 12K and Sun Fire™15K System

To ensure the data stored on the Sun Fire 12K/15K server is implemented and maintained in a secure environment, Sun training can help enterprises understand how to develop and implement solid security strategies to help protect their critical data. Sun's security courses listed below teach corporations how to deploy and manage Sun security products for maximum protection of the massive amounts of corporate data which will reside on their Sun Fire 12K or 15K system.

- Administering Security on Solaris
- Network Security Using SunScreen™ 3.1
- Securing Network Environments
- Trusted Solaris 2.X System Administration
- Developing Secure Commerce Applications

Systems Management using Sun™ Management Center

Systems Management using Sun™ Management Center is a prerequisite course to the Sun Fire 12K/15K Server Administration courseware. This course is designed to help customers realize the full benefits of their Sun Fire 12K or 15K system and prepare for the effective and successful use of Sun Management Center 3.0.

Education Consulting Services

Sun's Education Consulting Services will develop custom training solutions to enable customers to achieve rapid and efficient integration of their new Sun Fire 15K system into their IT environment. By working with companies on IT organization design, resourcing, and skills development Sun's Education consultants help companies ensure their people have the right skills at the right time to ensure a smooth Sun Fire 15K system integration. Sun helps companies select, retain, and retrain staff to maximize return on IT investments, and reduce total cost of ownership.

Sun Fire Server Skills Package

Sun Educational Services Sun Fire Server Skills Package provides the tool necessary to easily integrate training into the hardware sale. The Sun Fire Server Skills Package is a prepackaged training solution which contains the recommended courseware that will deliver the skills needed to optimize Sun Fire 12K or 15K servers in the computing environment. Once a skills package order is received an education manager will contact the customer to develop a tailored training program. Please contact a local Sun Education representative for details on availability and pricing of these learning solutions. Additional information can be found at <http://education.central>

Sun Fire™ 12K and Sun Fire™15K System

Sun Support Services

Managing a Sun Fire 12K or 15K system requires obtaining the right level of support to help customers improve operational efficiency and increase business and mission-critical system reliability. Sun Support Services offers proactive support, monitoring and sustaining services to help manage the customer's system availability uptime.

SunSpectrum™ Support Programs

	PLATINUM	GOLD	SILVER	BRONZE
	MISSION CRITICAL	BUSINESS CRITICAL	BASIC SUPPORT	SELF SUPPORT
Mission-critical account team				
Availability guarantee				
Patch & FCO management assistance				
Software release planning				
Skills assess. & personalized training plans				
Account support reviews	Monthly	Semi-annual		
Service account manager				
Remote system monitoring				
SunVIP™ Program				
On-site hw service coverage	7x24	8-8, M-F*	8-5, M-F*	
On-site hw service response	2 hours*	4 hours*	4 hours*	
Customer-defined priority				
Parts replacement	On-site	On-site	On-site	By courier
Telephone support coverage	7x24	7x24	8-8, M-F*	8-5, M-F*
Telephone support response	Live trans-P1	Live trans-P2	Live trans-P3	4 hours
Remote dial-in analysis				
SunSolve™ knowledgebase				
Software patches				
Software upgrades				

Features shown are representative of service delivered in the United States.

Program features are subject to local availability. Consult a Sun sales representative for additional information.

*You can, as an option, extend coverage hours and/or expedite response times.

The SunSpectrum™ Support Program is a flexible service that allows customers to choose a service level based on their specific needs. Levels of service range 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, certain Sun-supported peripherals, the Solaris Operating Environment, and Sun-supported software. 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 Services representative for program feature

Sun Fire™ 12K and Sun Fire™15K System

availability in their areas.

Note: On-site hardware service response times are average response times for calls the customer defines as urgent.

Sun RAS Profile Service

Sun RAS (Reliability, Availability, Serviceability) Profile is designed to identify areas, based on best practices and industry standards, where the customers can improve the operational efficiency and increase the mission critical reliability of their Sun system.

The Sun RAS Profile process involves comprehensive system assessment which leverages from both Sun's unique knowledge bank of product configuration data and the expertise of Sun's senior engineers. Sun experts drill deeply into the configuration issues, as well as the potential operational and environmental issues, that can affect the reliability, availability, and serviceability of the customers' Sun system.

RAS Profile Service includes:

- A comprehensive analysis
- Detailed recommendations
- A professional account team to help customers prioritize and develop a related action plan
- Periodic configuration reports and continuous action plan management are provided to help customers proactively manage their systems through out the year

RAS Profile is available as an optional uplift to SunSpectrum Platinum and SunSpectrum Gold.

Warranty Offering

The Sun Fire 12K and 15K server warranty offering includes an increased service level of (7 x 24). This is a significant enhancement over the current Sun Enterprise™ 10000 system service level of (5 x 12).

Sun Fire 12K and 15K server warranty offer is for one (1) year with the following features:

- 7 x 24 x 365 hardware on-site support (including holidays)
- 4 hour average hardware response time
- 7 x 24 x 365 telephone support (including holidays)
- Account Management
- Account Reviews
- Personalized Account Support Plan
- 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

Sun Fire™ 12K and Sun Fire™15K System

- Software upgrade information service
- Software installation and media support for 90 days

Installation Services

ServerStartSM Installation Services

The ServerStart installation service provides Systems Engineers in Support Services to help ensure customers have the proper configuration and the fastest possible time-to-value for high performance systems. A comprehensive system setup through ServerStart installation services includes:

- Installation planning and management
- Hardware Installation
- Installation of up to 18 boot devices (may be mirrored) on the Sun Fire 15K, 9 boot devices on the Sun Fire 12K.
 - Partition boot device
 - Install system software and appropriate patches
- Installation of up to 18 domains on the Sun Fire 15K, 9 domains on the Sun Fire 12K
 - Ensure domain is bootable
 - Ping other systems/domains
- Functional and Store Tools tests all pass
- System administration tutorial

Service for the Sun Fire 15K server includes installation of up to eighteen (18) domains and boot devices, as well as racks of storage devices required for a simple boot. Service for the Sun Fire 12K server includes installation of up to nine (9) domains and boot devices, as well as racks of storage devices required for a simple boot. The boot device may be one (1) device, or two (2), if mirrored. ServerStart does not cover the installation costs for any Sun StorEdge™ T3 Array units, even if they are used as a boot device. These must be installed via the ArrayStartSM offering.

Sun™ HPC Software Installation Services

The Sun HPC installation services provide customers with Sun HPC ClusterTools Software basic setup and configuration. Sun includes the Sun HPC ClusterTools Software installation services for free with the Sun Fire 15K, Sun Fire 6800, and Sun Sun Enterprise 10000 systems. Basic Sun HPC ClusterTools software installation services include:

- Installation planning session
- Install required Solaris patches
- Install NFS software and related patches
- Install Sun HPC ClusterTools software and related patches
- Install license, as required
- Configure as needed to start applications
- Installation and verification of Sun provided interconnects
- Add up to 4 nodes (additional nodes can be installed at T&M)
- Tutorial and reference documentation

Sun Fire™ 12K and Sun Fire™ 15K System

- System turnover

System Availability Agreement (SAA)

The Sun™ Systems Availability program supports a range of high availability percentages on specific Sun Enterprise 10000, Sun Fire 12K and Sun Fire 15K server configurations. Simply defined, the System Availability Agreement is a legal document specifying terms and conditions and each party's commitments to maintaining a target percentage of systems level availability.

Should the Sun systems covered under the SAA fail to meet the target availability percentage during the coverage period, Sun may credit the customer a percentage of its Sun Spectrum Platinum fees for support of the system, provided that the customer is in full compliance with the SAA program contract terms and conditions. More than a contract, the SAA symbolizes a cooperative relationship between Sun and your customer's organization focused on achieving a customer's specific availability goals and represents both the customer's company and Sun's commitment to reduce risk and protect investments by obtaining consistent levels of uptime.

Sun™ Remote Services

Sun Remote Services (SRS) Event Monitoring enhances the Sun Fire 12K and 15K's Warranty and the SunSpectrum GoldSM or SunSpectrum PlatinumSM support coverage by addressing the customer's mission-critical system monitoring needs. SRS Event Monitoring provides customers with 7 x 24 support coverage to help proactively manage event detection and problem resolution within their Sun IT environment. Initial customer infrastructure site setup to use SRS is available for a separate charge. Monthly telco connection charges for using SRS monitoring are the responsibility of the customer. Through proactive notification of potential system problems, a resolution can be identified and implemented to help prevent downtime. The Sun Fire 12K and 15K server's hardware redundancy capabilities combined with Sun's proactive event monitoring service helps customers obtain the highest possible system availability.

Additional information for Sun Support Services can be found at:
http://service.central/Lines_of_Business/Support_Services/index.html

Services Facts

- Sun is the #1 UNIX® platform integrator in the world (IDC, 2001)
- Sun Services supports over 1.8 million systems around the world
- 100% focus on Sun technology and best-of-breed partnerships
- 4,000+ Sun Enterprise Server 10000 servers supported in 63 countries
- Sun Services is one of the top 10 UNIX training providers
- 250,000 students trained annually

Sun Fire™ 12K and Sun Fire™15K System

For more information on Sun Services service offerings, please visit our website at <http://service.central> or at <http://www.sun.com/service>

CHAPTER 11 Sun Fire 12K/15K Cluster Platform

Highlights

For customers requiring a clustered compute infrastructure, the Sun Fire 12K and Sun Fire 15K servers can be delivered via Sun's integrated Cluster Platforms. A combination of Sun's most reliable server, storage, and clustering technologies, Cluster Platforms deliver a highly available architecture inherent in a redundant, failover system design. Sun's compute infrastructure experts pre-select, pre-integrate, and interoperability test the optimal cluster hardware and software combinations to simplify the upfront evaluation and deployment processes.

Key Applications:

The Cluster Platform 15K/9960 is ideal for high-end UNIX® system users who would like to upgrade to next-generation products and mainframe customers who have decided to embrace open systems. The platform provides an ideal infrastructure to run data warehousing and high-volume OLTP applications requiring mission-critical availability. It is also a great fit for server and storage consolidation. Full information about the 15K/9960 Cluster platform can be found at: <http://www.sun.com/architectures-platforms/suntoneplatforms/cp15K-9960/index.html>

Cluster Platform 15K/9960 - Specifications

Primary Components

The Cluster Platform 15K/9960 base configuration includes the following primary components:

- * Redundant Sun Fire[tm] 12K or 15K servers, each with one domain
- * Shared Sun StorEdge[tm] 9960 system
- * 72-inch Sun StorEdge expansion cabinet, which includes:
 - * One Netra[tm] T1 AC200 management server
 - * Four Sun StorEdge S1 arrays as boot devices
 - * Terminal concentrator
 - * Administrative hub
- * Sun[tm] Cluster 3.0 software
- * Solaris[tm] 8.0 Operating Environment

HARDWARE

Sun Fire 12K or 15K servers

Recommended* configuration includes the following for each of two servers:

Each of the two servers includes a minimum of:

- * Six CPU/memory boards
- * 24 900-MHz/8-MB UltraSPARC[tm] III+ processors
- * 48 GB of memory
- * Four hot-swappable I/O boards for a total of 32 PCI slots

Sun Fire™ 12K and Sun Fire™15K System

- * Two Sun Quad FastEthernet[tm] cards; two Gigabit Ethernet cards
- * Two UltraSCSI adapters; two Dual PCI network adapters

Note: Additional Uniboards can be added up to the maximum configuration of the Sun Fire 12K or Sun Fire 15K.

Sun StorEdge 9960 system

Recommended* configuration includes a minimum of:

- * 5.8 TB storage capacity
- * 4.4 TB usable RAID 5 storage capacity
- * 6.1 GB cache
- * Twenty 4x73 GB 10K RPM dual-ported Fibre Channel drives; one spare drive
- * 1.5 GB shared memory
- * Two Fibre Channel four-port adapters for short wavelength; fiber optic cables

* Note: This is the recommended storage configuration for the Cluster Platform 15K/9960. Different sizes of Sun Fire 12K or Sun Fire 15K or Sun StorEdge 9960 or Sun StorEdge 9910 systems can be integrated into the platform based on business needs.
Netra T1 AC200 management server

The management server includes:

- * Single 500-MHz/256-KB UltraSPARC IIe processor
- * 1 GB memory
- * Two 18-GB, 10K RPM SCSI drives
- * Internal AC power supply
- * DVD-ROM

Sun StorEdge S1 thin storage arrays

Each of four boot devices includes:

- * 3x18 GB disk arrays with 10K RPM hot-swappable drives

Terminal concentrator Eight-port 10Base-T Ethernet hub
Administrative hub Twelve-port 10/100Base-T Ethernet hub

SOFTWARE

The following software packages are installed and pre-configured with the platform:

- * Solaris 8 (10/01) Operating Environment
- * Sun Cluster 3.0 software (Update 1)
- * Solstice DiskSuite[tm] software (4.2.1)
- * Sun[tm] Management Center 3.0 software (Update 3)
- * Sun StorEdge Component Manager 2.2 software
- * UNIX® File System (UFS)
- * AnswerBook2[tm] software
- * Apache 1.3 Web server
- * Sun StorEdge Resource Manager software package

The following software packages are optional with the platform:

Sun Fire™ 12K and Sun Fire™15K System

- * VERITAS Volume Manager 3.1.1
- * Sun Management Center 3.0 Advanced Systems Monitoring
- * Sun Management Center 3.0 Reliability Manager
- * Sun Management Center 3.0 Premier Management Applications

The following Sun Cluster software agents are pre-installed on the management server but need to be configured:

- * Sun Cluster 3.0 NFS agent
- * Sun Cluster 3.0 Oracle HA agent

MEDIA KIT

The Cluster Platform 15K/9960 Media Kit includes:

- * Cluster Platform 15K/9960 Getting Started Guide (hard copy)
- * Cluster Platform 15K/9960 Install Guide (hard copy)
- * Cluster Platform 15K/9960 Recovery Guide (hard copy)
- * Sun Cluster 3.0 10/01 Base CD, Release Notes
- * Sun Cluster 3.0 server licenses, one per Sun Fire 12K or 15K server
- * Sun Cluster 3.0 Agent Pack (Agent CD and Agent hard copy material)
- * Sun Cluster 3.0 Software NFS Agents License
- * Recovery DVD

Note: All Cluster Platform 15K/9960 media and documentation can be used with the Sun Fire 12K platform except that you must order a Sun Fire 12K specific server license for Sun Cluster 3.0.

WARRANTIES AND UPGRADES

* Standard warranty terms for each component of the platform are included. SunSpectrum Silver[sm], SunSpectrum Gold[sm], and SunSpectrum Platinum[sm] support upgrades are available upon request. It is highly recommended that all platform components have a minimum of SunSpectrum Gold level of support. Upgrading the Netra T1 AC200 server and the Sun StorEdge S1 boot devices to SunSpectrum Gold level of support will match the Sun StorEdge 9960 and Sun Fire 15K server support levels.

OPTIONAL LICENSES

* Sun Cluster 3.0 software NFS and Oracle HA agents are pre-installed. The NFS HA agent license is provided free of charge; the Oracle HA agent licenses must be purchased separately.

* Sun Cluster 3.0 software HA agents pack is provided in the Media Kit, however agent licenses must be purchased separately.

* Solstice DiskSuite software is provided with the base configuration. VERITAS Volume Manager 3.1.1 software is also available as an option, but its license must be purchased separately.

* Sun Management Center 3.0 software packages are installed and configured with the Cluster Platform 15K/9960. The basic package is available free of charge. Licenses must be purchased in order to use the enhanced versions.

CHAPTER 12 Additional Resources

Configuration Information

Complete and up to date configuration information can be found at the following web sites.

For the Sun Fire 15K:

<http://systems.corp/products/servers/datacenter/sunfire15k/configuration.html>

For the Sun Fire 12K:

<http://systems.corp/products/servers/datacenter/sunfire12k/configuration.html>

Sun Fire 12K and 15K Server Configuration Guide

This document provides Sun Fire 12K and 15K configuration information to technical field operations personnel who will configure systems for customers, and to help System Engineers in sales efforts involved with Sun's UltraSPARC III Cu high-end Enterprise [SunWin Token #312503]

SalesDesk Sales Order Process Information

The Sales Desk website contains critical information required to process Sun Fire 12K and 15K system orders, including the Technical Checklists and Configuration Diagrams required by the MCSO process. Detailed information regarding the complete Sales Process is also available on the SalesDesk website. <http://sales-desk.west>

Sun Fire 12K and 15K Server Ordering Information and Flowcharts

Reference the High-End/Data Center Servers chapter in the online Configuration Guide for step-by-step configuration and ordering instructions for Sun Fire 12K and 15K Full System Bundles, Standard Server Configurations, Building Block Configurations, and associated upgrades.

Sun Fire 12K and 15K Software Compatibility Matrix

The Sun Fire 12K and 15K Software Compatibility Matrix defines the availability and readiness of software and its inter-dependency with other software. Please use the companion Supported Products Matrix to nail down the support status of specific devices or hardware features.

Sun Fire 12K and 15K Server I/O Support Matrix

Includes supported product configuration tables for network interface cards, disk storage host bus adapters, disk and tape storage products.

Product Announcements

Discover the latest information regarding product announcements, system enhancements, price adjustments and more.

Sun Fire™ 12K and Sun Fire™15K System

Training

Two types of training will be available for Sales. Web based training and instructor lead training. Information about web based training for sales on the Sun Fire 12K and 15K can be located at: <http://webtrain.central.sun.com>

Sun Fire 12K and 15K instructor led training can be found at:
<http://webtrain.central.sun.com/gso/>

The following is from the Sales web based training website:

Sun Fire 12K/15K Server Solutions

Welcome to the Sun Fire 15K Server Solutions module. This module provides information about the Sun Fire 12K and 15K servers which are the largest server in the Sun Fire family. It offers enhanced scalability and performance in a large-scale, centralized, enterprise server for parallel processing of commercial and technical applications.

Updated information is posted every Tuesday. Watch for pages to see the latest updates:

- CPU/Memory Board
- System Controller Board
- IBM's Regatta
- IBM's Messaging/Sun's Responses
- Regatta Weaknesses

The following is from the Sales Instructor Led training web site:

Sun Fire 12K/15K Applied Solutions (SASII)

A 4 day intensive look at the Sun Fire 12K and 15K servers covering the overall architecture and design of the server hardware and software, with additional focus on Sun Fire 12K/15K RAS, shared resource domains, and IDN technologies. Exercises and hands-on labs will enable students to effectively distinguish between the advantages of one configuration over another given a particular customer scenario.

Available for both Sun and partner SEs, the web-based training from the Sun Product Learning Portal is prerequisite to attending this instructor-led course. It is available at Geo-specific training hubs worldwide:

- Americas: Milpitas, CA
- Europe/Middle East/Africa: Munich, Germany
- Asia/Pac: Tokyo, Japan

Sun Experts CD - Sun Fire 12K/15K Edition.

The Sun Experts CD series is a pre-sales training and support mechanism, that is self-contained and can be loaded on your laptop, PC, or Solaris workstation and used anytime, anywhere, without a network connection. Comprised of presentation material, reference material like datasheets and white papers, videos, and product demos, the CD is a great way to learn about Sun's products and solutions.

The first Sun Fire 12K/15K edition will be available in December 2001. Topics will include Competition, late-breaking, post-Launch updates, sales strategies, server sizing tips, and much more.

Sun Fire™ 12K and Sun Fire™15K System

Sun Experts CD can also be ordered through Javacart (Sun internal) and via the Partner portal at: <http://reseller.sun.com/training/sunexpert.html>

White Papers

Many fine white papers have already been written about the Sun Fire 12K and 15K servers. These papers may be found on the external Sun website at: <http://www.sun.com/servers/wp/search/search.cgi>

Some of the white papers that can be found there include:

Sun Fire 15K Open System Controller (OpenSC) White Paper

This paper describes how to set up and verify third-party software on system controller (SC) boards, creating an open system controller (OpenSC). The SC board hosts System Management Services (SMS) software that monitors and controls a Sun Fire[tm] 15K system. Applications that run on the OpenSC, other than the SMS software and the Solaris[tm] operating environment, are considered third-party software. Third-party applications are expected to be lightweight, such as monitoring and backup agents, and not to demand intensive system resources. Described here are the SMS resource requirements, the maximum permitted resource consumption for third-party software, and techniques to help ensure that SMS receives the resources it needs to function properly.

Service Portfolio for the Sun Fire 15K Server

This white paper is targeted for current and new Sun customers with high-end data centers who are likely to deploy the Sun Fire 15K server in the near term. Information includes a complete service portfolio from Sun Professional Services, Sun Educational Services, and Sun Support Services.

The Sun Fire 15K Server for High Performance and Technical Markets

This white paper describes today's HPTC markets and the Sun Fire 15K architecture and features that are designed to server these markets. A brief description of the Sun Fire family of systems and Sun's comprehensive HPTC hardware/software stack is also included to provide a better understanding of Sun Fire 15K server's overall value proposition.

Sun Fire 15K Server: System and Resource Management

The Sun Fire 15K server is the most scalable and powerful mainframe class server in Sun Microsystems' extensive server product line, supporting up to 106 900MHz UltraSPARC[tm] III processors, 576 GB of memory, 72 I/O slots, over 250 TB of online disk storage, and over 15,000 software applications for the Solaris[tm] Operating Environment. This paper addresses the complement of management tools offered by the Sun Fire 15K server.

Sun Fire 15K Server: The Ultimate System for Raising Service Levels and Reducing Total Cost of Ownership

This paper provides an overview of the unique capabilities of the Sun Fire 15K server in delivering on the new demands of the networked data center.

Other white papers are available internal to Sun can be found at: http://systems.corp/tools/salestools/datacenter/highend_toolkit/whitepapers.html

Sun Fire™ 12K and Sun Fire™15K System

Sun Fire 15K Server Security Whitepaper

This document provides an overview of the Sun Fire 15K security design characteristics and how these technologies contribute protecting your information and computing resources. (Token #312488)

Sun Fire 15K General Positioning Whitepaper

This paper provides an overview of the Sun Fire 15K server with general technical information, features and benefits. The target audience is IT management. (Token #312485)

Sun Fire 15K Server: Availability for the Net Effect Whitepaper

This white paper details the features of the Sun Fire 15K server and demonstrates that Sun understands highly-available services are the result of the careful interaction of people, process, and product. With the Sun Fire 15K server, Sun Cluster software, Sun Management Services software, Sun Management Center software, and an extensive line of enterprise-class services and training, Sun provides an end-to-end solution for helping customers realize the most demanding availability requirements. (Token #312487)

Service Portfolio for the Sun Fire 15K Server Whitepaper

This document is targeted for current and new Sun customers with high-end data centers who are likely to deploy the Sun Fire 15K server in the near term. Information includes a complete service portfolio from - Sun Professional Services, Sun Educational Services, and Sun Support Services. (Token #314209)

Sun Fire Servers - Lowering the Total Cost of Ownership of Information Technology, PDF

The Sun Fireplane System Interconnect, PDF

Sun Fire 12K and 15K Success Stories

Success Stories are always a useful way to see the creative ways that customers are using Sun's products.

SUN FIRE 15K Success Stories are available at <http://sundoc.central.sun.com>

Success Story: TELUS Mobility

SunWIN #: 339634

Telecommunications: 'The Sun Fire 15K server's Dynamic System Domains, scalability, and computational power help TELUS Mobility in its drastic consolidation.'

Success Story: Uppsala University

SunWIN #: 339632

HPC/Academic Research: 'The Sun Fire 15K server provides Uppsala University with the extra speed, flexibility, and large memory required for its high-performance computing needs.'

Sun Fire™ 12K and Sun Fire™ 15K System

Success Story: Mitthogskolan (Mid-Sweden University)

SunWIN #: 339634

HPC/Academic Research: 'Reliability, scalability, simplification, and computational performance support Mitthogskolan's choice of the Sun Fire 15K server to meet its high-performance computing needs.'

Sun Fire 12K and 15K Documents

The following websites will list the currently available documents from Sun on the Sun Fire 15K server.

<http://docs.sun.com/query.html?Ab2Lang=C&Ab2Enc=iso-8859-1&qt=15K&dc=titl>

And this website will list the currently available Sun Fire 12K documents:

<http://docs.sun.com/query.html?Ab2Lang=C&Ab2Enc=iso-8859-1&qt=12K&dc=titl>

The following link will show the currently available documents covering the Systems Management Services Software which resides on the two System Controllers in the Sun Fire 12K and 15K servers.

<http://docs.sun.com/query.html?Ab2Lang=C&Ab2Enc=iso-8859-1&qt=SMS&dc=titl>

This following list shows some, but not all of the available documents for the Sun Fire 12K and 15K servers.

*Note: All hardware, software and service manuals published for the Sun Fire 15K can be used to manage the Sun Fire 12K except for the **Sun Fire 15K COD 1.1 Installation Guide**.*

Sun Fire 15K System Capacity on Demand (COD) 1.1 Installation Guide

from Sun Fire 15K HW/SW Collection

Sun Fire 15K Product Notes

from Sun Fire 15K HW/SW Collection

Sun Fire 15K Getting Started Guide

from Sun Fire 15K HW/SW Collection

Sun Fire 15K Reference III - FRU Service Document

from Sun Fire 15K HW/SW Collection

Sun Fire 15K Reference I - Nomenclature

from Sun Fire 15K HW/SW Collection

Sun Fire 15K Reference II - Component Numbering

from Sun Fire 15K HW/SW Collection

Sun Fire 15K System Service Manual

from Sun Fire 15K HW/SW Collection

Sun Fire 15K Hardware Installation and De-Installation Guide

from Sun Fire 15K HW/SW Collection

Sun Fire™ 12K and Sun Fire™15K System

Sun Fire 15K Site Planning Guide

from Sun Fire 15K HW/SW Collection

Sun Fire 15K Overview

from Sun Fire 15K HW/SW Collection

Sun Fire 15K Unpacking Guide

from Sun Fire 15K HW/SW Collection

Sun Fire 15K Read Me First

from Sun Fire 15K HW/SW Collection

Sun Management Center 3.0 Release Notes for Sun Fire 15K Systems

from Sun Management Center 3.0 Software

Sun Management Center 3.0 Platform Update 4 Release Notes for Sun Fire 15K Systems

from Sun Management Center 3.0 Software

Sun Management Center 3.0 Supplement for Sun Fire 15K Systems

from Sun Management Center 3.0 Software

Sun Management Center 3.0 Supplement for Sun Fire 15K Systems

from Sun Management Center 3.0 Software

System Management Services (SMS) 1.2 Installation Guide and Release Notes

from System Management Services (SMS) 1.2 Documentation

System Management Services (SMS) 1.2 Reference Manual

from System Management Services (SMS) 1.2 Documentation

System Management Services (SMS) 1.2 Administrator Guide

from System Management Services (SMS) 1.2 Documentation

System Management Services (SMS) 1.1 Installation Guide and Release Notes

from System Management Services (SMS) 1.1 Documentation

System Management Services (SMS) 1.1 Reference Manual

from System Management Services (SMS) 1.1 Documentation

System Management Services (SMS) 1.1 Administrator Guide

from System Management Services (SMS) 1.1 Documentation

Sun Fire 15K BluePrints

Sun realizes how important security is for our customers. Because of this Sun has created a series of papers called BluePrints that define the best way to implement features on Sun's Products. All Security BluePrints can be found from the following weblink:

<http://sun.com/security/blueprints>

The following Sun Fire 12K and 15K Blue Prints are available.

Sun Fire™ 12K and Sun Fire™15K System

Securing Sun Fire 15K Domains

This article provides recommendations on how to securely deploy Sun Fire 15k domains. These recommendations apply to environments concerned with security and particularly those where the uptime requirements of the domains and/or the information on the domains is critical to the organization. The configuration documented in this article is supported by Sun.

Securing the Sun Fire 15K System Controller (PDF)

This BluePrint OnLine article provides specific recommendations on how the security of the Sun Fire 15k System Controller (SC) can be improved. These recommendations include specific recommendations on how the Solaris[tm] Operating Environment (Solaris OE) image running on the SC should be configured in secured environments. The configuration documented has been tested by the Sun Fire 15K team and is fully supported.

Related Products

StorEdge™ S1 Thin Storage Array (Boot Device)

I. Product Description

The StorEdge S1 is the next generation of the Netra st D130 thin storage array. Built on the success of the D130, the StorEdge S1 retains many of the key features of the D130 that our customers have embraced, and is enhanced with new features that our customers have requested. The new S1 provides faster data access, increased flexibility of space utilization, and higher scalability at the same price as the D130.

For customers requiring additional storage capacity, a simple boot device, or incremental storage for horizontal-scaling applications in a rack-optimized package, the S1 is the obvious choice. Value to the customers include:

- Low cost storage solution
 - o Less than \$3000 List price.
- Faster data access
 - o With a transfer rate of up to 160MB/s
- Flexibility of space utilization
 - o Long cabling distance of up to 12 meters
- Investment protection
 - o Backward compatible with the Netra st D130 and single ended devices

Sun Management Center (SunMC)

Managing an integrated enterprise IT environment of servers, desktops, storage subsystems, relational databases, transaction monitors, and complex applications like SAP, Baan, and PeopleSoft, is a difficult and time consuming task. While no single application can provide a complete management solution for every platform, organizations can utilize intelligent, flexible enterprise management tools to solve many common problems related to continuous service availability. Enterprise management solutions helps enable organizations to manage systems, middleware, applications, and the networks that connect them, with the overall goal of optimizing service levels while minimizing administrative costs.

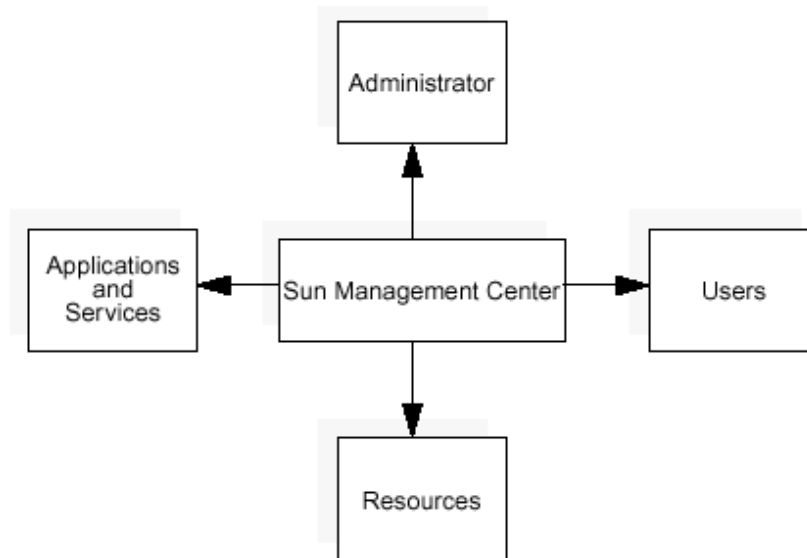
Sun Fire™ 12K and Sun Fire™15K System

One component of enterprise management, and perhaps the most pressing problem organizations face, is systems management. Focusing on the management of the hardware platform, operating system, and storage components, systems management includes the monitoring, performance optimization, and resource management capabilities that are essential to ensuring continuous service availability. Monitoring and adjusting CPU utilization and network bandwidth allocation enables application response time and service-level requirements to be met. In addition, storing data for trend analysis enables effective capacity planning and resource management. When systems such as these operate at peak efficiency, organizations are better equipped to work more productively and make better decisions —factors that can mean the difference between merely surviving and thriving in an increasingly competitive marketplace.

Conventional tools cannot support the move from managing small pockets of disparate systems and services to managing the integrated enterprise. IT organizations need an integrated set of systems management tools—a platform that offers common services for all enterprise management applications. Such solutions must have a consistent look-and-feel, implement enterprise-wide security procedures, integrate with other management tools, and enable new functionality to be added with ease.

Sun Management Center 3.0 is one of the best systems management tools from Sun to date. The Sun Management Center 3.0 three-tiered framework offers a single point of management for all Sun systems, the Solaris™ Operating Environment, applications and services for data center and highly distributed computing environments. With Sun Management Center software, IT organizations can efficiently manage and arbitrate between users, applications, and resources (see figure below) Sun systems, which have long been the preference for both technical and commercial computing users, capitalize on Sun's strengths in reliability, availability, serviceability (RAS), as well as performance and scalability. Enterprise management tools, however, must be designed to understand the systems they intend to manage—and no one knows how to manage Sun hardware and operating systems better than Sun. As a result, Sun Management Center software gives administrators one of the most comprehensive tools for managing Sun systems, as well as a platform that integrates easily with other enterprise management frameworks.

Sun Fire™ 12K and Sun Fire™15K System



Sun Management Center software provides a single point of management for the enterprise
Key Features

Sun Management Center software is an advanced system management tool. Designed to support Sun systems, Sun Management Center technology provides a platform upon which the enterprise can base its administrative and management operations to help ensure all systems and the services they provide are highly available. A powerful tool for managing the enterprise network, Sun Management Center software enables system administrators to perform remote system management, monitor performance, and isolate hardware and software faults for hundreds of Sun systems, all through an easy-to-use Web interface.

Key features of Sun Management Center 3.0 software include the following:

- **Reliability, availability, and serviceability (RAS) features**
 - Proactive and automated management of complex and common administrative tasks, reducing the likelihood of costly errors and helping to ensure availability
 - GUI using Java™ technology, providing heterogeneous GUI support, a common look and feel for all Sun Management Center applications, and the flexibility to manage the enterprise from any platform using Java technology, thereby increasing administrator efficiency; GUI supports dynamic reconfiguration and multi-pathing, allowing higher system availability
 - Predictive failure analysis, helping to enable administrators to predict potential memory and disk hardware failures on a statistical basis, thereby enhancing decision making and increasing availability
- **System administration tools**
 - Health monitoring, a sophisticated set of heuristics that incorporates a large body of administrative knowledge, including an intelligent rules-based health monitor that correlates metrics and gives suggested steps for resolution, resulting in simplified administration
 - Log-file scanning, enabling administrators to search and parse logs and registers for a particular status—the foundation for application health monitoring
 - Physical viewer, displaying photo-realistic images of hardware components and

Sun Fire™ 12K and Sun Fire™15K System

pointing to components with an associated event, guiding administrators unfamiliar with a particular Sun platform to quickly determine which components need to be replaced

- Logical viewer, presenting a tree hierarchy of the managed host or domain, including all hardware and operating system components—if an event is associated with a particular component, the logical viewer will identify its exact location within the hierarchy
- Event and alarm management, providing administrators with the information they need when they need it
- Real-time performance analysis, helping to enable administrators to isolate potential and existing bottlenecks
- ***Integration features***
 - Single event model, enabling information to be shared with multiple consoles or users with ease
 - Standard interfaces and protocols, enabling integration with third-party management tools, including Tivoli TME 10 TEC 3.6, HP OpenView ITOPs, BMC Patrol, and Computer Associates Unicenter TNG, thereby providing a complete enterprise management solution
 - Full SNMP (versions 1, 2,c and V2usec) and RMI connectivity, helping to enable information to be shared with other enterprise management tools
- ***Scalability and customization features***
 - Common management platform, scalable from a single system to thousands of server and desktop systems
 - Configuration flexibility, allowing Sun Management Center technology to be configured out of the box to best fit the needs of the environment, as well as providing easy customization for new rules, scripts, actions, and so on
 - Extensible agent architecture, allowing administrators to add functionality and management features with ease
 - Developer environment, helping to enable organizations to plan, design, develop and integrate Sun Management Center modules as well as third-party applications, tools, and customized solutions based on the Sun Management Center framework
 - Rules writing documentation, helping to enable rules to be created and customized for the Sun Management Center environment
- ***Security features***
 - Enterprise-wide security measures, such as authentication, data integrity, and access control lists for management of data and active management functions
- ***Ease of use features***
 - Single point of management, allowing effective use of administrative resources –
Dynamic agent modules, providing functionality to be added or removed asynchronously and dynamically
 - Domain-aware agents, providing dynamic system domains supported on the Sun Enterprise™ 10000 server and its associated resources to be monitored independently
 - Multiple system support, allowing administrators to monitor and manage all Solaris Operating Environment systems remotely
 - Logical element grouping, helping to enable the grouping of Sun systems by geographical location, server role, administrative responsibility, and so on; the status of the individual systems are summarized by the group, as events and alarms are issued to each system

Sun Fire™ 12K and Sun Fire™15K System

- Hierarchy and topology viewer, a central management application that displays the hierarchy and a topology map of all the objects that are being managed
- Automatic discovery of Sun systems, including IP address, subnet address, hostnames, and OIDs to identify specific types of systems
- Console wrapper, enabling users to access native UNIX® commands with ease

Solaris™ 8 Operating Environment

The industry's premier UNIX® operating environment includes key features such as:

- Hot patching (unique): Apply patches to the operating environment without rebooting.
- Live upgrade: Build a new, upgraded operating environment image while maintaining a full system application service load.
- IP network multipathing: Provides load balancing and failovers with multiple network interface cards.
- Mobile IP: Helps enable users to stay connected to the network regardless of location.
- UNIX file system (UFS) logging: Store transactions in a log before applying to the file system.
- Extensive security capabilities: Including IPSec, Kerberos v5, pluggable authentication modules (PAMs), role-based access control (RBAC), shared-secret and public-key encryption, and smart card support.

Top 10 Reasons for Moving to the Solaris 8 Operating Environment Today

1. **Massive scalability.** Multithreaded design enables it to scale to handle heavy network traffic, huge data sets, and compute-intensive processes.
2. **Reduced cost of ownership.** Offers data center-class reliability, availability, and serviceability at a fraction of the cost of a mainframe.
3. **Reactive scalability.** Helps enable system administrators to write and execute scripts that quickly address peak load situations and dynamically move system resources from one domain to another.
4. **Lower downtime due to planned events.** Live upgrade helps enable system administrators to build a new, upgraded environment image while system is running full application service load. Hot patching framework allows patches to take effect without rebooting, greatly reducing system downtime required to diagnose, test, analyze, and correct problems.
5. **Highly available system and network resources.** Sun Cluster 3.0 helps enable system administrators to manage servers and storage devices as a single entity. Dynamic multipathing technology handles load balancing and failovers with multiple network interface cards.
6. **Secure and predictable operations.** Built-in network protection features include IPSec, Kerberos, RBAC, SunScreen SKIP, SunScreen firewall, and smart card support.
7. **Connectivity to the world.** Supports iPlanet Directory Server, Solaris PC NetLink, mobile IP, XML, Java 2 Platform, and legacy connectivity.
8. **Manageability.** Easily manage systems with Solaris Management Console™ software, diskless client management, and Sun Management Center.
9. **High performance application services.** Solaris Network Cache Accelerator dramatically increases Web server performance. Java technology is tightly integrated into Solaris platform to efficiently leverage its multithreading capabilities.
10. **Universal operating environment.** Delivers technologies that combine billions of users and devices over the Internet and provides extensive support for 37 different languages and 127 locales.

Sun Fire™ 12K and Sun Fire™15K System

Solaris™ Management Console (SMC)

<http://webhome.corp/solarispm/sunmc/>

Sun™ Remote Services (SRS)

<http://srs.ebay/>

Sun StorEdge™

<http://webhome.ebay/networkstorage/software/>

System Management Services (SMS) 1.2

Sun Fire™ 12K and 15K System Management Services 1.2 is the second system management software release specifically designed for use on Sun Fire 12K and 15K System Controllers to manage a Sun Fire 12K and 15K system. An instance of this software is required on each of the main and backup System Controller. While only one System Controller is needed at any time to control the whole system, when configured with a backup System Controllers, the configuration will have automatic failover capability and higher availability.

The Sun Fire™ 12K and 15K System Management Services 1.2 uses command line interface and is a text based management tool. For customers who prefer to use administrative tool with Graphic User Interface (GUI) capability, they should order and install the Sun Management Center3.0 software. Sun Fire™ 12K and 15K System Management Services 1.2 is designed to work with the Sun(TM) Management Center 3.0 seamlessly and is the prerequisite for Sun Management Center 3.0 software.

As the initial formal release of the Sun Fire™ 12K and 15K System Management Services 1.2 software, the Sun Fire™ 12K and 15K System Management Services 1.2 is expected to be distributed only with each newly ordered Sun Fire™ 12K or 15K system.

The Sun Fire™ 12K and 15K System Management Services 1.2 offers following New features for customers' productive use:

- Configuring up to 9 Dynamic Systems Domains on the Sun Fire 12K and 18 Dynamic Systems Domains on the Sun Fire 15K.
- Configuring and managing up to 18 System/Memory boards and up to 18 I/O or MaxCPU boards in combination on the Sun Fire 15K and up to 9 System/Memory boards and up to 9 I/O or MaxCPU boards in combination on the Sun Fire 12K.
- On the Sun Fire 15K each domain can have up to the maximum number (18) of available system boards, up to 18 I/O boards (including up to 17 MaxCPU boards substituting for I/O boards).
- On the Sun Fire 12K each domain can have up to the maximum number (9) of available system boards, up to 9 I/O boards (including up to 8 MaxCPU boards substituting for I/O boards).
- Providing various services to manage up to 18 Dynamic Systems Domains on the Sun Fire 15K and up to 9 Dynamic Systems Domains on the Sun Fire 12K.
- Supporting up to 288 GB of addressing capability for each domain, 576 GB in a single domain on the Sun Fire 15K is a future feature.
- Supporting Solaris 8 10/01 and later as the Operating Environment for each Dynamic Systems Domain
- Supporting Solaris 8 10/01 and later as the Operating Environment for a System Controller
- Allowing CPU/Memory boards running at their rated speeds on a per board basis, once boards with CPU speed greater than 900 MHz are introduced.

Sun Fire™ 12K and Sun Fire™15K System

- Monitor temperatures, currents, and voltage levels
- Monitor and control power to the components
- Execute diagnostic programs
- Warns of impending problems
- Automatically reboots a domain after a system software failure
- Keeps logs of interactions between the SC environment and the domains
- Managing hot plug and swap of an I/O assembly or I/O cards
- Controlling resource management of the platform or domain with access privileges
- Automatic failover of System Controller
- Automatic failover of the main control board
- Automatic failover of clock signal
- Automatic System Recovery
- Rapid booting for up to 6 domains simultaneously. On the Sun Fire 15K all 18 domains can be rebooted automatically in sequence with grouping of 6.
- Dual Power Grid support
- Supporting Sun Management Center 3.0
- Supporting I/O Multipathing (IPMP)
- Supporting Sun StorEdge(TM) Traffic Manager

Post GA features are:

- Providing platform-wide Dynamic Reconfiguration command line interface to initiate DR on each domain - Supporting Dynamic Reconfiguration initiated from domain
- Supporting Reconfiguration Coordination Manager

Further information on SMS may be found at the following URL:

<http://www.sun.com/servers/highend/sunfire15k/>

Relevant URLs

Internal Resources

- Sun Fire 12K Homepage
<http://systems.corp/products/servers/datacenter/sunfire12K>
- Sun Fire 15K Homepage
<http://systems.corp/products/servers/datacenter/sunfire15K>
- Sun Fire 3800, 4800, 4810 and 6800 Servers
<http://systems.corp/products/servers/datacenter/index.html>
- Sun Fire15KProduct Information
<http://onestop/>
- Sun BluePrints Security Documents
<http://sun.com/security/blueprints>

Sun Fire™ 12K and Sun Fire™ 15K System

Glossary – Technical Terminology

ACL

Access Control List. In order to assign a board to a domain, the board name must be listed in the Access Control List (ACL). When a board is listed in the ACL, the system controller software is allowed to process addboard or deleteboard requests for that board..

ADR

Automated Dynamic Reconfiguration. Allows one to invoke scripts for tasks such as adding or deleting a system board to or from a domain, moving a board between domains, or for determining the status of a system board.

alternate pathing

A software feature which gives you the ability to have multiple paths to the same device from one domain. This provides an extra degree of fault tolerance. For disks, this is MPxIO. For networks it is IPMP.

Automatic System Recovery (ASR)

Provides system operation in the event of a hardware failure. Identifies and isolates a failing hardware component, and builds a bootable system configuration without the failed hardware component.

boardset

In the Sun Fire 15K, one expander board plus its attached System board and/or I/O board.

CDC

Coherence Directory Cache. A cache of information associated with each home agent, located in SRAMs on the expander board, which may record a cache line where its owner is or which expanders may have a shared copy.

Cheetah

The internal name for the UltraSPARC-III processor.

CLI

Command Line Interface. A user interface to a computer's operating system or an application in which the user responds to a visual prompt by typing in a command on a specified line, receives a response back from the system, and then enters another command, and so forth.

Coherent

A somewhat overused (and therefore ambiguous) term. *Coherent transactions* may refer only to RTO, RTS, and WB type transactions; or is sometimes used to refer to any transactions to cacheable address space.

COMA

Sun Fire™ 12K and Sun Fire™15K System

Cache-Only Memory Architecture. An architecture in which main memory is treated like a large cache, and memory pages can be present in more than one place. WildCat uses *Simple COMA*, renamed Coherent Memory Replication (see).

control boardset

Plugs into one of two control slots on the plane interconnect. Consists of a plane interconnect support board, a system controller, and a peripheral board.

cPCI

Compact PCI. A form of PCI that allows an individual PCI card to be removed and replaced while the system is running.

CPU

Central Processing Unit. An UltraSPARC III processor in this context.

Domain

A logical grouping of system boards. Each partition contains up to two domains. Domains do not interact with each other. Domains differ from partitions in that they share Repeater boards. A domain is able to run its own copy of the Solaris operating environment and has its own host ID. A domain must contain at least one system board and one I/O board.

domain isolation

Mechanism where by the AXQ confirms that the address transaction destination is in the same domain as the source. Enabling technology for "split slot" configurations.

Double-pumped

A degraded mode in which the information is sent over half the wires using twice the normal number of cycles, thus avoiding a broken wire.

DR

Dynamic Reconfiguration. Adding or removing resources to a domain, such as removing a CPU board from one domain and giving it to another.

DSD

Dynamic System Domain. Mainframe-style partitioning allows a single UNIX server to be divided logically into multiple servers, creating "systems within a system."

dynamic reconfiguration

Software support for the hardware configuration changes made to a domain running the Solaris operating environment. Dynamic Reconfiguration handles the software aspects of dynamically removing a defective board from the system and installing a replacement board without bringing the system down and with minimum disruption to user processes running in the domain.

ecache

Sun Fire™ 12K and Sun Fire™15K System

External cache. In Cheetah, this is the L2 cache, with the data kept in external SRAM chips and the tags kept in the Cheetah chip.

ECC

Error Correcting Code. Extra bits that go along with data bits to detect or correct errors. See SECDED.

expander board

The Sun Fire 15K board that interposes between the CPU and IO boards and the plane interconnect. The expander board contains the AXQ and SDI ASICs.

fault tolerance

The ability of a system to continue to perform its functions, even when one or more components have failed.

fireplane bus

The interconnect architecture used by the Sun Fire servers when communicating between L1 and L2 devices.

FRU

Field replaceable unit or replacement part.

HA

High Availability.

Hot Plug

Physically connecting hardware to a powered system.

Hot Swappable

The ability to physically add or remove a component from the system while the system is up and allowing that hardware to be configured into the system without interruption to system operation. Hot swap is combination of Hot plug and Dynamic Reconfiguration.

IDN

Inter-Domain Networking.

Sun Fire™ 12K and Sun Fire™15K System

IPMP

IP Multipathing. Increases network availability and performance.

keyswitch

See virtual domain keyswitch.

Latency

The time between initiating a request from data and the beginning of the actual data transfer.

LVD

Low Voltage Differential. (Ultra2 and beyond) LVD reduces the amount of power needed to drive the SCSI bus, and increases the cable length from 3 meters to 12 meters.

MaxCPU board

A Dual Proc board, which plugs into slot 1 (the "I/O" slot) of the Expander. Originally called the MacCat board.

MPxIO

MultipathI/O. I/O framework allowing Solaris to represent and manage storage devices which are accessible through multiple host controller interfaces within a single instance of Solaris.

MTBF

Mean Time Between Failure. Is a measure of how reliable a hardware product or component is.

MTBSF

Mean Time Between System Failure. Is a measure of how reliable a hardware product or component is. It is more important than MTBF.

MTTR

Mean Time to Repair. The average time it takes to repair a Failed component.

NAFO

Network Adapter Failover. Exclusively a feature of Sun Cluster.

NetCon

Network Console. System controller service that provides a "console" for single-user operations.

No Single Point of Failure

Implies that for any one component failure (which may cause a system crash), the system can still be rebooted with full data access, though perhaps with degraded performance. When a failure occurs, Sun systems automatically reboot (fast reboot) and reconfigure to recovery.

NUMA

Sun Fire™ 12K and Sun Fire™15K System

Non-Uniform Memory Access. An architecture in which memory and processors are in clumps, and accesses to memory in another clump is possible but slower.

NVRAM

Non-Volatile RAM. Any semiconductor memory device that does not lose its contents when power is turned off.

OLTP

Online Transaction Processing. Is a class of programs that facilitates and manages transaction-oriented applications, typically for data entry and retrieval transactions.

partition

A hardware feature that provides very good hardware failure isolation between two groups of system boards and their associated Repeater boards. Each partition has its own set of Repeater boards. When a system is divided into two partitions, the system logically behaves as two separate systems. Now called a segment.

RAS

Reliability, Availability, and Serviceability.

RCM

Reconfiguration Coordination Manager. A generic framework which allows dynamic reconfiguration (DR) to interact with system management software. The framework enables automated DR removal operations on platforms with proper software and hardware configuration and provides enhanced error reporting on DR operation failures.

SBM

Solaris™ Bandwidth Manager. A policy-based, directory-enabled software solution regulating IP networks' bandwidth usage in LANs and WANs.

SC

System Controller (see).

Scalable Shared Memory (SSM)

A mode of the Fireplane interconnect which allows multiple Fireplane buses to be connected together.

Shared Resource Domain (SRD)

A Sun Fire 12K or 15K domain that is permitted to send and receive a restricted set of transactions to/from other domains (its client domains), thus allowing it to facilitate and control inter-domain communication.

shared expander

In the Sun Fire 12K or 15K, an expander which is either an SRD expander (see) and/or a split expander (see).

Sun Fire™ 12K and Sun Fire™15K System

Slot-0 board

A board that has an off-board bandwidth of 4.8 GBps. One type of slot-0 board is used in the Sun Fire: the System board.

Slot-1 board

A board that has an off-board bandwidth of 2.4 GBps. Three slot-1 board types are used in the in Sun Fire servers: PCI, cPCI, and 3800 cPCI.

SMP

Symmetric Multi-Processor. The Sun Fire servers are SMPs.

SNMP

Simple Network Management Protocol. SNMP is any system listening to SNMP events. This is usually the system with the Sun Management Center software installed.

split expander

When the two system boards in a boardset are in different domains, since the expander board keeps the transactions separate for each system board.

SRD

Shared Resource domain. A Sun Fire domain that is permitted to send and receive a restricted set of transactions to/from other domains (its client domains), thus allowing it to facilitate and control inter-domain communication.

SRS

Sun Remote Services. Continuous system monitoring that utilizes an intelligent agent based architecture to monitor key systems variables 7x24. When a problem is detected, an alarm is automatically generated to notify SRS engineers and initiate resolution.

SSM

Scalable Shared Memory. That part of the Safari architecture that enables point-to-point coherency, thus minimizing system addressing bottlenecks. Coherency messages only get sent to components that need to be involved in a particular transaction.

System Controller

Sun Fire™ 12K and Sun Fire™15K System

The system controller consists of the System Controller board and the system controller software. The system controller provides communication pathways for console traffic and other data that needs to be passed between the system controller and the system. The system controller software monitors and controls the system, manages hardware, and configures domains.

System Controller board

A board containing a microSPARC™ processor, which oversees operation of the system and provides clocks and the console bus.

UltraSPARC-III

The SPARC V9 processor used in all systems comprising the Sun Fire family of systems. Internally known as the Cheetah.

US-III

See UltraSPARC-III.

UltraSPARC IV

Second generation processor model of the Safari-bus generation. Pin compatible with UltraSPARC III.

UPA

Ultra Port Architecture.