


*Fujitsu Upgrades Workgroup  
Line with Faster, Reliable  
SPARC64 V-Based Servers*

*June 2003*



*A D.H. Brown Associates, Inc. White Paper Prepared for  
Fujitsu*


This document is copyrighted © by D.H. Brown Associates, Inc. (DHBA) and is protected by U.S. and international copyright laws and conventions. This document may not be copied, reproduced, stored in a retrieval system, transmitted in any form, posted on a public or private website or bulletin board, or sublicensed to a third party without the written consent of DHBA. No copyright may be obscured or removed from the paper. D.H. Brown Associates, Inc. and DHBA are trademarks of D.H. Brown Associates, Inc. All trademarks and registered marks of products and companies referred to in this paper are protected.

This document was developed on the basis of information and sources believed to be reliable. This document is to be used "as is." DHBA makes no guarantees or representations regarding, and shall have no liability for the accuracy of, data, subject matter, quality, or timeliness of the content. The data contained in this document are subject to change. DHBA accepts no responsibility to inform the reader of changes in the data. In addition, DHBA may change its view of the products, services, and companies described in this document.

DHBA accepts no responsibility for decisions made on the basis of information contained herein, nor from the reader's attempts to duplicate performance results or other outcomes. Nor can the paper be used to predict future values or performance levels. This document may not be used to create an endorsement for products and services discussed in the paper or for other products and services offered by the vendors discussed.

# TABLE OF CONTENTS

<b>EXECUTIVE SUMMARY</b> .....	<b>1</b>
<i>Sidebar 1: PRIMEPOWER Server Family Specifications</i> .....	2
<b>NEW PRIMEPOWER SERVERS</b> .....	<b>3</b>
PRIMEPOWER250 – RAISING THE BAR FOR ENTRY SERVERS .....	3
<i>Table 1: PRIMEPOWER250 Specifications</i> .....	4
PRIMEPOWER450 – WORKHORSE WORKGROUP SERVER .....	4
<i>Table 2: PRIMEPOWER450 Specifications</i> .....	5
MULTIPLE FORM FACTORS PROVIDE FLEXIBILITY.....	5
SPARC64 V MICROPROCESSOR AND CROSSBAR ARCHITECTURE.....	6
<i>Figure 1: Crossbar Architecture of the PRIMEPOWER250/PRIMEPOWER450 Servers</i> .....	7
HIGH AVAILABILITY (RAS) .....	7
XSCF – EASED MANAGEMENT.....	8
<i>Figure 2: Remote Monitoring through XSCF</i> .....	9
PERFORMANCE.....	9
<i>Table 3: PRIMEPOWER Performance</i> .....	9
<i>Sidebar 2: Sun Applications/Solaris Interoperability</i> .....	10



# *Fujitsu Upgrades Workgroup Line with Faster, Reliable SPARC64 V-Based Servers*

## **EXECUTIVE SUMMARY**

Today's enterprises increasingly look at their IT staff to build evolving infrastructures to accommodate newer technologies such as web services, which promise better and tighter integration across their distributed systems. CIOs and CTOs are asked to enable these capabilities under an over-arching need to cut costs and gain better return on IT (ROIT) assets. Workgroup servers often form the backbone for such environments as their scale-out capabilities deliver flexibility and performance in an affordable package and a small footprint. Fujitsu's announcement of the much-anticipated addition to its PRIMEPOWER family with the 1.1 GHz 64-bit SPARC64 V chip-based dual-processor capable PRIMEPOWER250 and the quad-processor capable PRIMEPOWER450 running the widely accepted Solaris Operating Environment (OE) provide new levels of reliability and price/performance (economics) for these customers.

The PRIMEPOWER250 and PRIMEPOWER450 servers leverage design fundamentals from Fujitsu's high-end Shared Multi-Processing (SMP) servers stressing Reliability, Availability, and Serviceability (RAS) characteristics throughout the hardware. Featuring standard redundant components, these new servers allow customers to achieve higher levels of availability than previously attainable from the earlier generation of products. Built by using mainframe design points, the SPARC64 V chip introduces unprecedented levels of reliability for a RISC processor. The higher degree of parallelism in the chip coupled with the high-bandwidth crossbar switch architecture and plenty of expansion capabilities distinguish the PRIMEPOWER250 and PRIMEPOWER450 servers as leading systems for all applications. These applications range from serving dynamic web pages to compute-intensive workloads and transaction processing.

The servers feature the new Extended System Control Facility (XSCF) with an independent service processor for complete out-of-band systems management and remote monitoring. The XSCF gathers information from the various error checkers (beyond the CPU) in the system and helps lower operational cost by providing precise and granular information to administrators. It alerts them on potential failures and assists in diagnosing these errors and recovering the system to a full operational state.

Fujitsu's support of the Solaris OE for its PRIMEPOWER family gives customers access to over 12,000 supported applications and compatibility with numerous Independent Hardware Vendors (IHVs) by means of its standards-based servers. Fujitsu also maintains strong partnerships with leading software players such as Oracle, SAP, and VERITAS, and is flexible in building customized solutions addressing specific customer requirements. Its offerings in

the services arena, including the 24x7 business critical environment and global support capabilities earn Fujitsu a reputation for providing end-to-end integration capabilities. This holistic approach towards delivering higher business value combined with the PRIMEPOWER250 and PRIMEPOWER450 system features and the Solaris OE allows customers to develop and operate an efficient infrastructure for their diverse workloads. Customers also gain a faster ROIT without having to deal with a sprawling architecture. The approach also gives Fujitsu a competitive edge over rival vendors who either focus on a “box” approach or necessitate a “beyond the box” investment to gain better ROIT.

**SIDEBAR 1**  
**PRIMEPOWER**  
**Server Family**  
**Specifications**

The PRIMEPOWER server family offers a highly scalable platform for commercial applications from the dual-CPU capable PRIMEPOWER250 server to the 128-CPU capable PRIMEPOWER2500 server. Built around an open standards-based architecture developed on the SPARC64 V microprocessor (SPARC-compatible implementation) and the Solaris OE, the binary compatible PRIMEPOWER family provides leading reliability, performance, manageability and availability characteristics.

	Model 250	Model 450	Model 650	Model 850	Model 900	Model 1500	Model 2500
<b>CPUs</b>	1 – 2	1 – 4	1 – 8	2 – 16	4 – 16	4 – 32	8 – 128
<b>Clock Speed</b>	1.1 GHz	1.1 GHz	1.08, 1.35 GHz	1.08, 1.35 GHz	1.35 GHz	1.35 GHz	1.3GHz
<b>Memory</b>	8 GB	16 GB	32 GB	64 GB	64 GB	128 GB	512 GB
<b>Max. Storage Bays</b>	6	6	2	4	16	32	128
<b>Max. PCI Slots</b>	6	9	20	40	36	72	320
<b>Redundant Components</b>							
– Disk	Yes	Yes	Yes	Yes	Yes	Yes	Yes
– Power Supply	Yes	Yes	Yes	Yes	Yes	Yes	Yes
– Fan	Yes	Yes	Yes	Yes	Yes	Yes	Yes
– Power Feed			Yes	Yes	Yes	Yes	Yes
– System Control Facility						Yes	Yes
– Crossbar							Yes
– System Clock							Yes
<b>Hot-Swap Components</b>							
– Disk	Yes	Yes	Yes	Yes	Yes	Yes	Yes
– Power Supply	Yes	Yes	Yes	Yes	Yes	Yes	Yes
– Fan	Yes	Yes	Yes	Yes	Yes	Yes	Yes
– Tape	Yes	Yes	Yes	Yes			
– PCI Card		Yes			Yes	Yes	Yes
– System Board					Yes	Yes	Yes
– System Control Facility						Yes	Yes
<b>Other Features</b>					8 XPAR or 2 PPAR	15 XPAR or 4 PPAR	15 XPAR or 15 PPAR

## NEW PRIMEPOWER SERVERS

Fujitsu's approach for its PRIMEPOWER250 and PRIMEPOWER450 servers balances high performance, reliability, and availability within a competitively priced package. Customers will find this a compelling solution as most implementations in this space focus on one of these key customer requirements that call for additional investments to bridge the gap. The sections below include an overview of the SPARC64 V microprocessor and discuss the key design points that make the PRIMEPOWER suitable for the workgroup server space.

Instead of developing point solutions to fit limited application areas, Fujitsu has excelled in its ability to employ technology across its entire product line. Transparent with its new servers, this approach began with the SPARC64 V chip. The theme further extends to the overall architectural approach in terms of scalability (increments), reliability, and availability as seen in Sidebar 1. The higher-end servers, while incorporating additional features, remain a superset of the lower-end servers, as expected, but more importantly, in a predictable fashion.

Consistency is a key issue for customers who wish to lower training costs, both for administrators and maintenance staff. A drastic reduction of operating costs associated with managing the full PRIMEPOWER server family across the enterprise is also a welcome attribute. Customers can develop and use the same knowledge base to operate these systems efficiently. This further lowers development costs for Fujitsu, in-turn improving overall system pricing for customers.

A well-balanced architecture demands both a competitively performing processor and a high-bandwidth low latency subsystem. A faster processor on its own in the absence of an equivalently fast system bus is often left sitting idle, wasting valuable CPU cycles. Both the PRIMEPOWER250 and PRIMEPOWER450 servers are designed to provide this balance to ensure full system utilization and eliminate bottlenecks in the system.

### PRIMEPOWER250 – RAISING THE BAR FOR ENTRY SERVERS

The PRIMEPOWER250 server raises the bar for entry UNIX systems by combining the performance and industry-leading reliability characteristics of the SPARC64 V chip with an advanced, high-performance architecture at a competitive price point. Single or dual SPARC64 V chips operating at a 1.1 GHz frequency provide plenty of horsepower for compute-intensive tasks, while simultaneously adding high levels of availability in the system for a complete 24x7 e-business environment. Adding to the CPU performance, the crossbar interconnect provides ample bandwidth of up to 3.5 GB/sec. (peak) per processor for heavy transaction-oriented environments. These include application serving or small business database serving.

The system's small 2U footprint for dense rack environments packs in all the capabilities noted above providing further enhancement. The pedestal form factor adds to the common features with robust expansion capabilities including six 64-bit PCI slots for additional devices and six internal hot-swappable disk expansion bays for nearly 0.9 TB of high-performance storage capacity. Table 2 below highlights the key features of the PRIMEPOWER250.

**TABLE 1**  
*PRIMEPOWER250*  
Specifications

	Rackmount (2U)	Pedestal
<i>CPU</i>	1.1 GHz SPARC64 V (L1 Cache: 256 KB; L2 Cache: 1 MB)	
<i>Memory</i>	8 GB DDR SDRAM through 8 DIMM Slots (Two-Way Interleaving)	
<i>Internal Disk Capacity</i>	294 GB (2 x 147 GB)	882 GB (6 x 147 GB)
<i>64-Bit PCI Slots</i>	3	6
<i>Network Connectivity</i>	FastEthernet, Gigabit Ethernet	
<i>Redundant Components</i>	Power Supply, Fan, HDD (S/W Mirror)	
<i>Hot-Swappable Components</i>	Power Supply, Fan, HDD	Power Supply, Fan, HDD, DAT
<i>Automatic System Reconfiguration</i>	CPU, Memory, PCI bus	
<i>System Management</i>	XSCF	
<i>Operating System</i>	Solaris 8, 9	

## **PRIMEPOWER450 – WORKHORSE WORKGROUP SERVER**

The PRIMEPOWER450 provides midrange expansion capabilities in the workgroup server space with nine 64-bit PCI expansion slots for I/O consolidation. In addition, the PRIMEPOWER450 comes with six internal hot-swappable disk bays for larger storage capacity. Other comparatively positioned servers in the market require customers to upgrade to higher-end models or add optional expansion units for similar capabilities.

The server supports up to four 1.1 GHz SPARC64 V chips that can be added one at a time. The flexibility achieved allows customers to configure their systems based on the type of application without restriction. Some customers require a large memory-to-CPU ratio while others need simpler configurations.

The capabilities of the PRIMEPOWER450 lend themselves nicely to resource-intensive applications such as application server, Supply-Chain Management (SCM), and Product Definition Management (PDM) software. The PRIMEPOWER450 appeals to other areas such as mission-critical database serving, and high-performance technical compute node. It also serves as a consolidation vehicle for older and smaller two- or four-way servers.

**TABLE 2**  
PRIMEPOWER450  
Specifications

	Rackmount (4U)	Rackmount (7U)	Floor-Stand
CPU	1.1 GHz SPARC64 V Chip (L1 Cache: 256 KB; L2 Cache: 1 MB)		
Memory	16 GB DDR SDRAM through 16 DIMM slots (Four-Way Interleaving)		
Internal Disk Capacity	588 GB (4 x 147 GB)	882 GB (6 x 147 GB)	
64-Bit PCI Slots	6	9	
Network Connectivity	FastEthernet, Gigabit Ethernet		
Redundant Components	Power Supply, Fan, HDD (S/W Mirror)		
Hot-Swappable Components	Power Supply, Fan, HDD	Power Supply, Fan, HDD, DAT, PCI Card	
Automatic System Reconfiguration	CPU, Memory, PCI Bus		
System Management	XSCF		
Operating System	Solaris 8, 9		

## MULTIPLE FORM FACTORS PROVIDE FLEXIBILITY

Instead of offering a one-size-fits all approach, Fujitsu's investment in multiple form factors for its PRIMEPOWER250 and PRIMEPOWER450 provides flexibility for a range of users. For example, a datacenter customer dealing with limited space requires densely packaged servers that typically shift data to externally connected network storage devices. In this environment, the rack-optimized 2U-high PRIMEPOWER250 and 4U-high PRIMEPOWER450 servers make an ideal choice. At the other end, small- and medium-sized businesses or branch office locations require most of their storage to be internal to the system as they run applications such as e-mail. The PRIMEPOWER250 in the pedestal form factor packaging offers six internal disk bays and six PCI expansion slots providing plenty of room to grow.

The PRIMEPOWER450 server is also offered in a larger 7U rackmount with a floor-stand option. Besides providing additional headroom for internal storage and I/O expansion with up to nine 64-bit PCI slots, these form factors allow for hot-swap capability on PCI adapters and tape drives, features not supported in the denser rackmount version.

Compared to the existing PRIMEPOWER servers in this space, the new PRIMEPOWER250 and PRIMEPOWER450 servers occupy much smaller volume (up to 40%) with reduced height and depth. Existing customers can upgrade to the new systems and take advantage of the improved performance, RAS, and manageability in a much smaller footprint.

## **SPARC64 V MICROPROCESSOR AND CROSSBAR ARCHITECTURE**

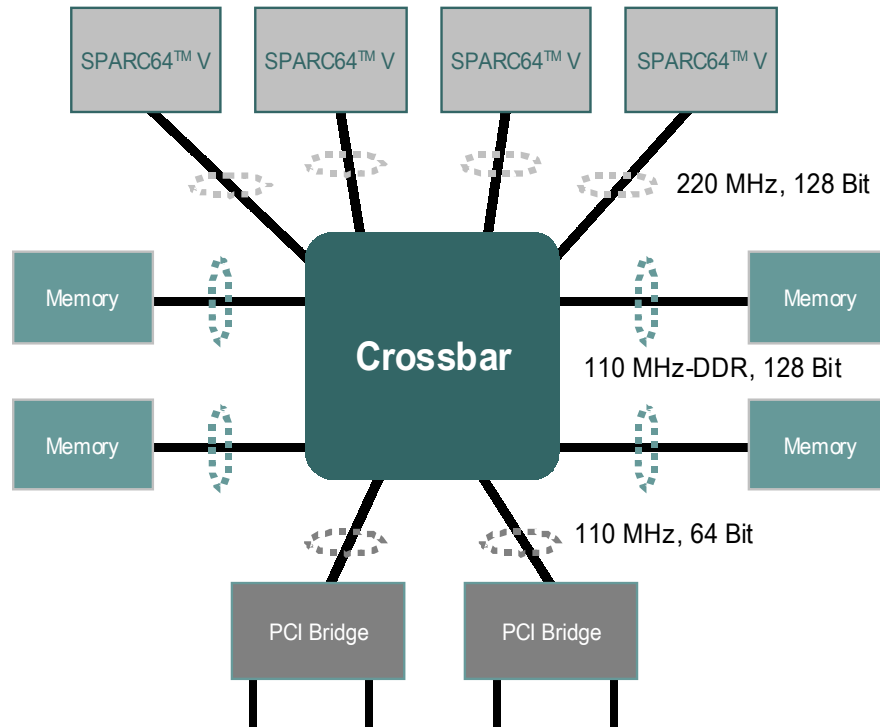
Traditional processor design principles aim to extend the architecture to achieve more instructions per clock cycle as well as increase the number of clock cycles per second. The superscalar out-of-order pipeline capable of executing up to six instructions per clock cycle reveals Fujitsu superiority in this area. In addition, Fujitsu has introduced “mainframe” reliability features never before attempted in a RISC processor (see discussion later in this section). The SPARC64 V chip provides competitive performance, often higher than other processors at similar clock speeds. This ability of the SPARC64 V chip in the PRIMEPOWER250 and PRIMEPOWER450 servers drives higher CPU utilization rates and allows these servers to handle more workload.

Built using the state-of-the-art 0.13-micron fabrication process, Fujitsu engineers have created a processor that consumes only about 50 watts of power. The smaller power consumption carries significant implications on power and cooling costs for a large, dense server installation. The Fujitsu design made maximum use of the large number of circuits available on the die. Some enhance performance; others support error detection and fault isolation.

As chip circuitry gets complex with an ever-increasing number of fast transistors occupying the same footprint as dictated by Moore’s law, the probability of CPU failures due to electrical “noise” poses a real problem. Fujitsu is the first vendor to offer mainframe-class hardware instruction retry in a RISC processor. This feature re-executes a failed instruction, provided proper machine-state has been saved, and avoids crashing the processor. This reliability feature distinguishes the SPARC64 V from other processors. Beyond this, Fujitsu has added extensive error-checking circuitry to the chip (data paths, caches, etc.). In the unlikely event of an error, the chip can deconfigure the failed part dynamically and continue operating in degraded mode. Fujitsu’s SPARC64 V chip offers parity and Error Checking and Correction (ECC) code throughout the processor, including register sets.

The crossbar architecture provides a balance between the CPU capabilities and the rest of the system design. As seen in Figure 1, the crossbar interconnect eliminates a bridge- or a bus-based approach to connect CPU, memory, and I/O, which can cause bottlenecks in the system as more CPUs try to access the same resource. The crossbar design provides separate paths to avoid contention for multiple accesses and operates at a high frequency of 220 MHz. The 128-bit wide data path to memory provides up to 3.5 GB/sec. of bandwidth per CPU, for a total of up to 14 GB/sec. (peak) in the PRIMEPOWER450 server. The memory subsystem supports high-performance Double Data Rate (DDR) SDRAM memory with dual memory controllers to provide bandwidth comparable to the system bus. The 64-bit wide I/O path also operates at a high frequency of 110 MHz to provide sufficient bandwidth and concurrent access to multiple PCI devices.

**FIGURE 1**  
Crossbar  
Architecture of the  
PRIMEPOWER250/  
PRIMEPOWER450  
Servers



### PRIMEPOWER450 (4U) Architecture

The point-to-point nature of the crossbar speeds performance and reduces latency by minimizing the number of data hops. This aspect of the PRIMEPOWER250 and PRIMEPOWER450 servers allows them to easily handle applications such as transaction processing and e-business workloads, which place heavy demands on the entire subsystem.

### HIGH AVAILABILITY (RAS)

As more applications move towards mission-critical status, the IT staff must ensure data integrity and uptime of the entire infrastructure. Customers often have to invest in multiple and often unnecessary “rack-and-stack” servers to act as backups in the event of a system failure. The cumulative effect raises the overall cost of ownership. While the initial cost of these servers may be lower, the overhead associated with managing them negates the initial benefit. Fujitsu’s mainframe-class reliability design aims at making its systems inherently more reliable even in the workgroup space and gives customers maximum benefit and lower cost of ownership with the PRIMEPOWER servers.

Both the PRIMEPOWER250 and PRIMEPOWER450 servers feature numerous self-monitoring and self-diagnosing capabilities, as well as redundant components to ensure the highest levels of availability. There are also a number of error checkers beyond the CPU throughout the system to monitor and log errors. The memory subsystem and data paths are all parity/ECC protected. If the error

threshold exceeds a certain level for key hardware components, the Automatic System Reconfiguration (ASR) function auto-detects the failed component and does a quick reboot without operator intervention. As the system boots up, ASR reconfigures the firmware so it can deconfigure that failed component and resume normal operation. The ASR function is supported on the processor, memory, and PCI bus.

Fujitsu's PRIMECLUSTER software eliminates any single point of failure by creating a highly available cluster environment for customers requiring much higher levels of availability than provided by a single PRIMEPOWER server.

The new workgroup servers feature redundant power supplies and cooling fans and software mirroring allows for hard drive redundancy. This redundancy aspect also allows hot swapping of these components. Besides the power supplies, cooling fans, and hard drives, the PRIMEPOWER450 7U also supports hot swapping of internal tape drives and PCI cards.

The PRIMEPOWER servers feature LED lamps on the front and back of the unit monitoring the hardware condition. They light up on failure of the component and help the maintenance staff identify the failed unit easily. Fujitsu has upgraded the remote management capabilities on its PRIMEPOWER servers through the new XSCF. These improvements build on the previous generation of monitoring tools available from Fujitsu. (XSCF is discussed in greater detail in the following section.)

## **XSCF – EASED MANAGEMENT**

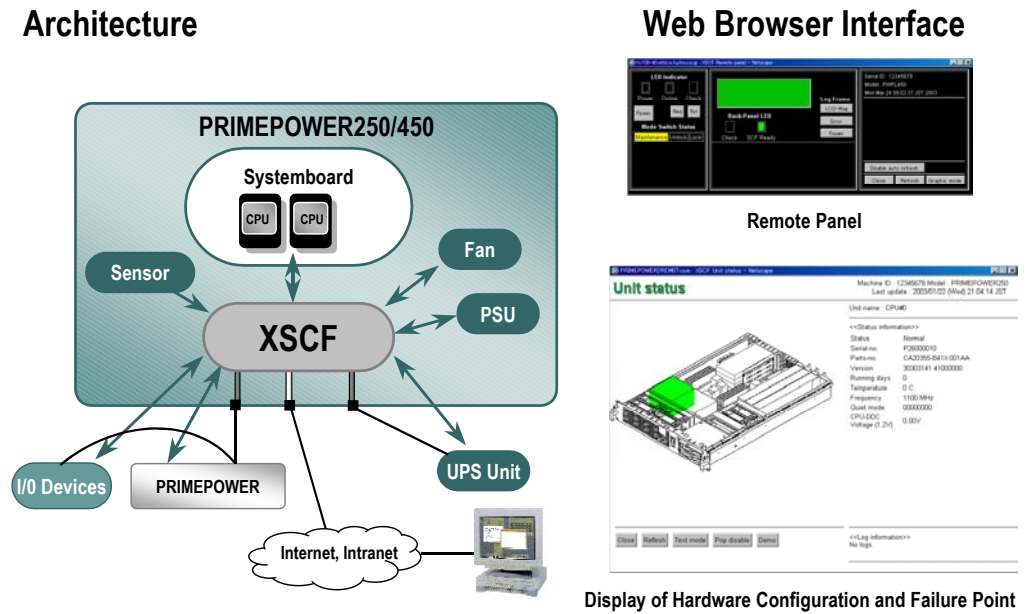
The XSCF includes an independent service processor embedded in the system and monitors all vital components (CPU, fan, power supply, UPS, temperature, etc.) for system operation. This ability of the XSCF improves upon the RAS capabilities of the system by continuously monitoring the hardware and operating system state. A dedicated LAN port alerts administrators of potential failures. The XSCF also provides other features such as remote power on/off and reboot.

The PRIMEPOWER250 and PRIMEPOWER450 servers, as with the higher-end SPARC64 V-based servers, provide a number of error checkers inside the system. All checkers send log information to the XSCF, giving administrators detailed information on the errors at a very granular level. Total system availability increases through the extensive error checkers and the ability of the XSCF to gather information and alert administrators to potential failures or detected failures.

The XSCF continues to function even if the operating system on the main unit hangs up. It maintains an error log containing detailed information accessible to the administrators through a browser-based GUI tool for diagnosis and troubleshooting. Fujitsu has designed the XSCF to make it accessible via multiple interfaces by supporting protocols such as telnet, HTTP, and SNMP, eliminating the need for dedicated PC hardware as required by some other vendors.

The PRIMEPOWER servers also come standard with the ESF and WebSysAdmin software tools. The ESF software provides system administration functions, utilities and drivers including the one that controls the XSCF. WebSysAdmin is the primary management software for the PRIMEPOWER servers. Its single point of management feature consolidates views from multiple interfaces into a single GUI screen for easier and quicker navigation, presenting a complete view for administrators to ensure system uptime.

**FIGURE 2**  
Remote Monitoring  
through XSCF



## PERFORMANCE

Fujitsu has always taken an active role in publishing results for its PRIMEPOWER servers across a variety of industry-standard benchmarks for popular application workloads. In this manner, customers may see the potential of the servers by reviewing the benchmark that closely resembles their workload environment. Fujitsu's PRIMEPOWER servers hold records across many of the leading application benchmarks. The SPEC CPU2000 (SPECint2000 and SPECfp2000) benchmarks focus on compute-intensive tasks. The server-side Java Business benchmark, SPECjbb2000, measures the performance of the Java Virtual Machine (JVM) and Just-in-Time (JIT) compiler, in addition to CPUs, caches, and memory subsystems.

**TABLE 3**  
PRIMEPOWER  
Performance

	PRIMEPOWER250 (Two-Way)	PRIMEPOWER450 (Four-Way)
SPECint2000 (peak)	680	682
SPECfp2000 (peak)	1030	1037
SPECjbb2000 (ops/sec.)	30216	58212

The available results for the new PRIMEPOWER servers cover single CPU performance and the results are very comparable for both systems. This lends a proof point to the consistency in the crossbar architecture. As the number of

processors in the system scale, the single CPU performance of a more scalable system is often much lower than that of a limited scaling system using the same exact processor. This is caused by latency issues in the backplane. The SPECjbb2000 result for the PRIMEPOWER450 shows a strong scalability of 93% from dual-CPU to quad-CPU, extending the proof point discussed earlier.

#### **Sidebar 2: Sun Applications/Solaris Interoperability**

The Solaris OE is based on, and is compatible with the open SPARC V9 Level 2 architecture. Like Sun, Fujitsu is a member of the SPARC International Consortium. Its SPARC64 V processor also conforms to the published specification of the SPARC V9 Level 2 architecture. Fujitsu Ltd. has licensed the Solaris OE as the strategic operating system for its PRIMEPOWER servers. Fujitsu has passed the tests as part of the Sun Solaris Hardware Partners (SHWP) Program for OEMs to ensure the hardware's (SPARC64/PRIMEPOWER servers) compatibility with the Solaris OE. In terms of functionality from the Application Binary Interface (ABI) and Application Program Interface (API) perspectives, the Solaris OE running on the Fujitsu SPARC64-based PRIMEPOWER servers is identical to the Solaris OE running on Sun UltraSPARC/Fire servers. This means that applications conforming to the published ABIs and APIs will run on Fujitsu's SPARC64-based PRIMEPOWER servers without any modification.

Standard benchmarks provide a good representation of the system performance, but by nature they are intended to run across multiple systems. This means that they only take advantage of the lowest common denominator of any system. As such it is possible to achieve much higher performance by tuning the application for the specific platform and operating environment.