PRIMEPOWER 650 and 850

Technical White Paper

December 2001

PRIMEPOWER650

PRIMEPOWER850

PRIMEPOWER 650 and 850 Contents

Preface

The past several years have seen a significant change in business computing. We are seeing a fast evolution into the broadband era and the expansion of e-commerce, e-governance, and public trunk systems. We are also seeing a rapid expansion of contents business. For a middle-range server to meet such needs, the following conditions have become more and more important:

(1) Performance (Processing power corresponding to broadband requirements)
(2) Reliability and Availability (High reliability and availability in public trunk systems)
(3) Future expandability (Investment protection across new processors and technology)
(4) Space-saving (Ability to set up many devices in limited spaces with good extensibility)

PRIMEPOWER 650 and 850 are being developed as the best solution to these requirements.

PRIMEPOWER650 and 850 Outline

PRIMEPOWER 650 Features

PRIMEPOWER 650 is a mid-range, rack-mounted server model that can accommodate up to eight 675 MHz SPARC64™ GP CPUs.

A source-synchronous method that enables an ultra high-speed transmission frequency of 540MHz on the core system bus is employed to achieve significant expansion of system bus performance to 13.8GB/s. The best possible mounting structure has been adopted to improve memory access performance. This high-density mounting, where eight CPUs can be installed on the same system board, enables short-distance transmission between CPUs resulting in reduced latency between all eight CPUs. In addition high-speed system bus and memory design with reduced latency significantly improves transaction performance.

To meet the demands for high-density installation in computer rooms with limited space, high installation ability has been achieved by the use of high-density mounting. We have created the most compact eight CPU system in the world at 8U. The extremely compact cabinet accommodates a DAT drive, DVD drive, two hard disk bays, plus eight PCI slots for the external connections required to support a significant networked environment. In addition, taking the importance of reliability into consideration, redundant PSU and dual power feed options are supported and can be mounted in the cabinet. Its compactness makes it possible to mount it in the same rack as a disk array and the various other I/O units required for a fully operating solution. Limited installation space in an office or computer center environment can therefore be effectively used.

In order to prepare for and manage the occurrence of a processor, memory, or I/O bus error, a degradation function is supported that isolates the faulty location at system restart. In addition, essential components such as the disks, power supplies, and fans are available in redundant configuration; with hot swappable components (components that can be exchanged during operation) supported as well. By employing these measures, high system availability has been achieved.

For flexibility in expansion of the number and type of networks, a 19-inch rack compatible, 4U PCIBOX (PCI slot expansion unit) system configuration is employed. By connection of these additional units, a further 12 PCI slots can be added, so that PCI cards can be mounted in up to 20 slots in total. This enables the user to flexibly configure the required I/O configurations. . PRIMEPOWER 650, with its ability to be easily installed in either an office or computer center environment, supports the widest range of use for many applications.

- Up to eight SPARC64™ GP processors @ 675 MHz, 8 MB secondary cache
- Height of 8U (19-inch rack)
- Up to 32 GB, memory subsystem
- Up to 20 PCI slots
- 13.8 GB/s data bus width
- N + 1 redundant power supply subsystem and dual power feed option
- Redundant cooling subsystem
- Hot swappable power supplies and fans
- 64-bit Solaris operating environment

PRIMEPOWER 850 Features

PRIMEPOWER 850 is a larger mid-range, rack-mounted server model that can accommodate up to sixteen 675 MHz SPARC64™ GP CPUs.

A source-synchronous method that enables ultra high-speed transmission frequency of 540MHz on the core system bus is employed to achieve excellent expansion of system bus performance to 41.5GB/s. Best possible mounting structure is also adopted to improve memory access performance.

High-density mounting, where eight CPUs can be installed on the same system board, enables short-distance transmission between CPUs resulting in reduced latency between all eight CPUs. In addition, the use of a back-panel passage to connect the two system boards results in further reductions in latency.

The high-speed system bus and memory design with reduced latency also significantly improves transaction performance.

To meet the demands for high-density installation in computer rooms with limited space, high installation ability has been achieved by employing high-density mounting. In doing so we have created the most compact 16 CPU unit in the world at only 17U. This extremely compact cabinet accommodates a DAT drive, DVD drive, hard disk units, plus 16 PCI slots for external connections required by the server in supporting a significant networked environment. In addition, taking the importance of reliability into consideration, redundant PSU and dual power feed options are supported and can be mounted in the cabinet. It's compactness allows it to be mounted it in the same rack as a disk array and the various other I/O units required for a fully operating solution. Limited installation space in an office or computer center environment can therefore be effectively used.

In order to prepare for and manage the occurrence of processor, memory, or I/O bus errors, a degradation function is supported that isolates the faulty location at system restart. In addition, essential components such as the disks, power supplies, and fans are available in redundant configuration, with hot swappable components (components that can be exchanged during operation) supported as well. By employing these measures, high system availability has been achieved.

For flexibility in expansion of the number and types of networks, the same 19-inch rack compatible, 4U PCIBOX (PCI slot expansion unit) system configuration as available with the 650 model, can be used. In this case up to two of these units can be connected, adding up to 24 additional PCI slots. This allows 40 PCI cards to be mounted in total. This enables the user to flexibly configure the I/O units as required. PRIMEPOWER 850, with its broad power capability and ability to be easily installed in either office or computer center environments supports the widest possible range of use for many applications.

- Up to 16 SPARC64™ GP processors @ 675 MHz, 8 MB secondary cache
- Height of 17U (19-inch rack)
- Up to 64 GB, memory subsystem
- Up to 40 PCI slots
- 41.5 GB/s data bus width
- N + 1 redundant power supply subsystem and dual power feed option
- Redundant cooling subsystem
- Hot swappable power supplies and fans
- 64-bit Solaris operating environment

PRIMEPOWER Lineup

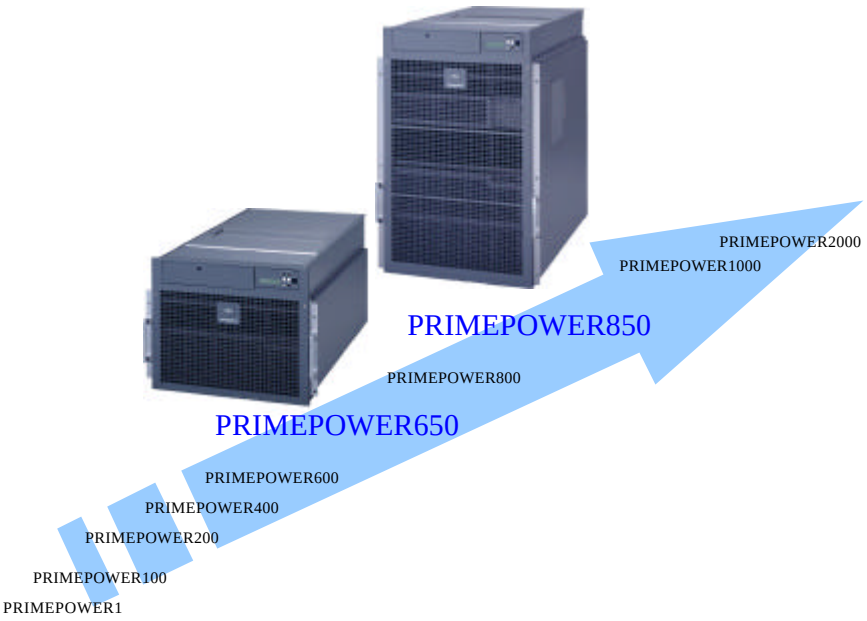

The PRIMEPOWER consists of a comprehensive power lineup starting with the entry-level model 1U single processor PRIMEPOWER 1, under-desk server PRIMEPOWER 200, mid-range servers PRIMEPOWER 400, PRIMEPOWER 600, PRIMEPOWER 650 and PRIMEPOWER 850, the enterprise class PRIMEPOWER 1000 with up to 32 processors, and on to the PRIMEPOWER 2000 that can accommodate either 64 or 128 processors. The PRIMEPOWER thus offers an extremely broad range of servers to create ideal solutions for the widest range of requirements. Each server platform provides complete RAS functionality and the highest in availability when used in cluster configurations.

Performance

The system performance of the PRIMEPOWER 650 and 850 series scales up smoothly as processors are added. This is made possible through an optimum balance of processor, memory, and I/O performance.

- High-speed bus supporting data transfer at 540 MHz
- SPARC64™ GP CPUs with up to 8 MB of secondary cache
- PCI slots capable of 64-bit/66 MHz operation

A high-speed bus called the named Channel Bus connects with the PCI bridge. A PCI bus that conforms to PCI Rev. 2.1 is used for the I/O bus. PCI slots capable of operating at 64-bit/66 MHz are provided enabling PRIMEPOWER 650 and 850 to be used as high-speed network or large-scale file servers. The table below lists the performance values of each key component:

			Maximum configuration	
	Throughput	Notes	PRIMEPOWER650	PRIMEPOWER850
Processor	2.2 GB/s	Performance per processor	8	16
Memory	8.6GB/s	Performance per one-way interleave	2	4
Channel bus	0.8 GB/s x 2(in/out)	Per channel bus	2	
Data bus	13.8 GB/s	Data transfer peak performance of PRIMEPOWER 650	-	-
	41.5 GB/s	Data transfer peak performance of PRIMEPOWER 850	-	-

The table below lists the relative improvements in transaction performance, using a factor of 1 for the performance of the current model type:

Model	PRIMEPOWER600	PRIMEPOWER650
Transaction performance ratio	1	1.42

Model	PRIMEPOWER800	PRIMEPOWER850
Transaction performance ratio	1	1.15

PRIMEPOWER600: CPU 600MHz

PRIMEPOWER800: CPU 675MHz

Compatibility

The PRIMEPOWER architecture follows a consistent design policy. This same hardware design policy is used on both PRIMEPOWER 650 and 850. Solaris, the international operating environment is used. This ensures binary compatibility of applications between PRIMEPOWER server models as well as between the SPARC/Solaris units. As a result, taking into consideration utilization of current resources and future expandability, your investment is protected.

System Architecture

Features

PRIMEPOWER 650 and 850 are rack-mounted, mid-range servers with the following features:

1. High-performance SPARC64™ GP processors with large-capacity secondary cache
- SPARC64™ GP (675 MHz) @ 8 MB E-cache

2. Large-capacity memory subsystem
- 4 GB of memory per processor is supported.

3. SMP architecture
- The architecture supports SMP configurations.
- High-performance system bus using 540 MHz data transfer

4. High-performance I/O interface (maximums shown are with PRIMEPOWER 850)
- As basic slots, up to 16 PCI buses are provided.
- As expansion slots, up to 24 PCI buses are provided to enable flexible configuration.
- High-speed I/O is supported using 64-bit/66 MHz slots.

5. High availability function support
- Redundant configuration, hot swappable Power Supply Units (PSU), and Fan Trays

Configuration basics

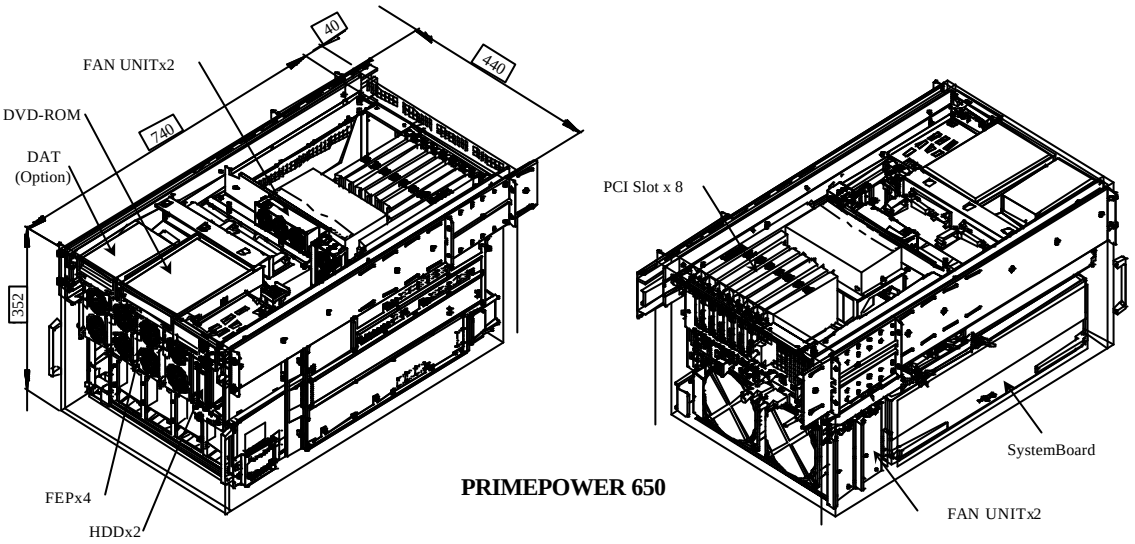

PRIMEPOWER 650

The basic unit of each PRIMEPOWER model is a high-density system board capable of accommodating up to eight processors and a maximum memory capacity of 32 GB. One system board can be mounted in a PRIMEPOWER 650 and two system boards can be mounted in a PRIMEPOWER 850. They support a networked environment in an extremely compact cabinet that can be mounted in general-purpose 19-inch racks. All essential components such as the disk units, the PCI card slots, DVD unit, DAT unit, Ethernet port, RS232C port, and Keyboard/Mouse are mounted in the cabinet. In addition, high reliability is also ensured by the redundant configuration of the power supplies, fans and the built-in dual power feed option.

In the mounting layout of the main unit, the system board, and power supply-I/O section, cooling systems are completely separated for more efficient cooling. This configuration enables two Fan Trays to be installed per system board enabling redundant configuration of the cooling system for the high-density mounted system boards. In turn each of these Fan Trays also contains two small-scale, high-performance fans with adjustable speed control to ensure continued efficient cooling even if a single fan fails.

PRIMEPOWER 650 and 850 are also developed as environmentally friendly products. Their design takes into consideration all laws and regulations including local business standards, market demands, and meets the requirements for the international ecology mark. PRIMEPOWER 650 and 850 are also recognized as green products as prescribed by Fujitsu in its own strict definition. In further consideration of the environment, the radio noise emissions of this equipment do not exceed the EN55022 Class B and VCCI Class B limits.

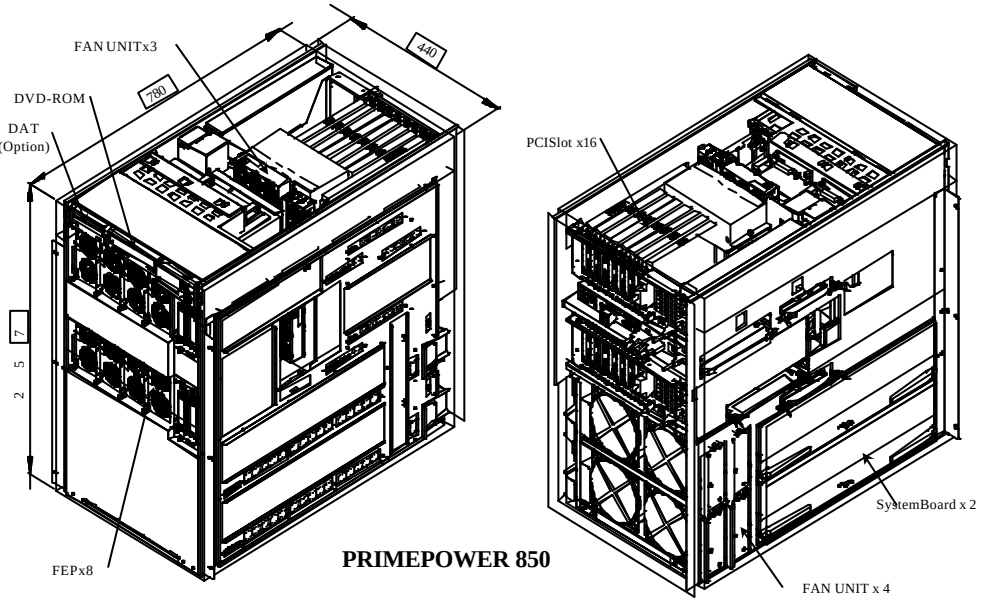

PRIMEPOWER 850

System Board

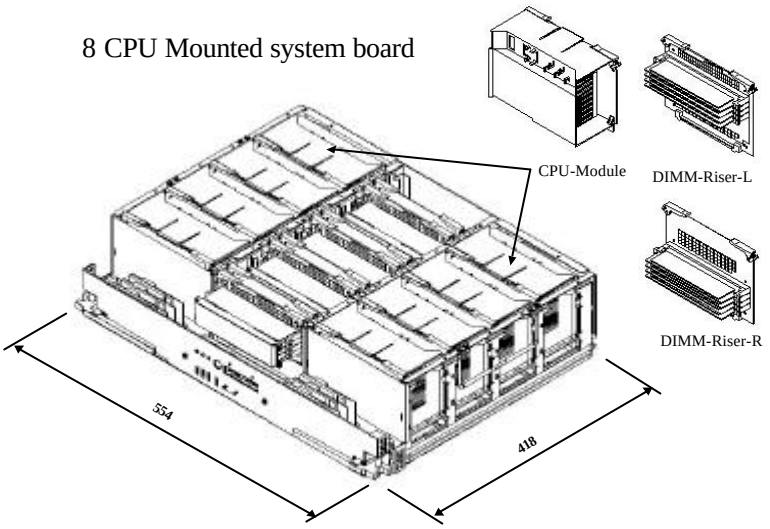

To increase the mounting density, eight CPUs are mounted on one system board. In addition, to fully utilize the space, the DIMMs are mounted three dimensionally using a riser card. 32 GB of memory can be mounted per system board. The DDCON section of the power supply has also been decentralized and mounted on the system board and CPU modules. The result is that mounting density of the power supply has been greatly increased.

SPARC64™ GP Processors

The SPARC64™ GP CPU employs super scalar processing technology and is a highly reliable, highly integrated processor based on the SPARC-V9 64 bit architecture.
PRIMEPOWER 650 and 850 models provide a processor module with a SPARC64™ GP CPU and secondary cache SRAM mounted in high density.

- Full conformity with SPARC Version 9
- Up to four instructions can be issued per cycle
  - Simultaneous execution of up to eight instructions: 4 integer operations (or 2 address operations), 2 floating point operations, and 2 load/store operations
  - Execution of a full range out-of-order instruction for all instructions using the register rename function and reservation station. Except for special instructions like member, etc.
  - High-level dynamic branch forecast function (16K entry forecast table)
  - 64-bit virtual address space
  - Support for multiple page sizes of up to 4 GB. By use of an MMU option with a symmetrical, full-associative 256-entry TLB for all page sizes
  - Large-capacity primary cache of 128 KB (I) + 128 KB (D).
  - Large-capacity 8-MB secondary cache
  - Pipeline L2 cache interface enabling high performance
  - ECC protection for both caches types (primary and secondary)
  - 128-bit data bus and 16-bit ECC
  - High-speed bus (peak bandwidth of 2.2 GB/s)
  - Support for future upgrades

Memory Subsystem

Increasing the data transfer speed has been an important focus area in improving overall system performance. According to the conventional approach, increasing the data transfer speed between the boards is the key to increasing the system clock speed. One of the problems related to data transfer has always been the question of how to reduce the deviation of the clock signals from the data signals (skew). However, a new method (Source-Synchronous method) that involves sending the clock signal together with the data has been adopted here. This has been found to be very effective in significantly reducing skew. By this method, data transfer at 540 MHz has been achieved.

Conventional 16-CPU class models require a board on which a crossbar switch is mounted for the connection between multiple system boards. This requirement causes an increase in memory access latency and an increase in the number of parts. In designing PRIMEPOWER 650 and 850, the need for this crossbar has been eliminated, and the goal of direct connection between the system boards has been achieved.

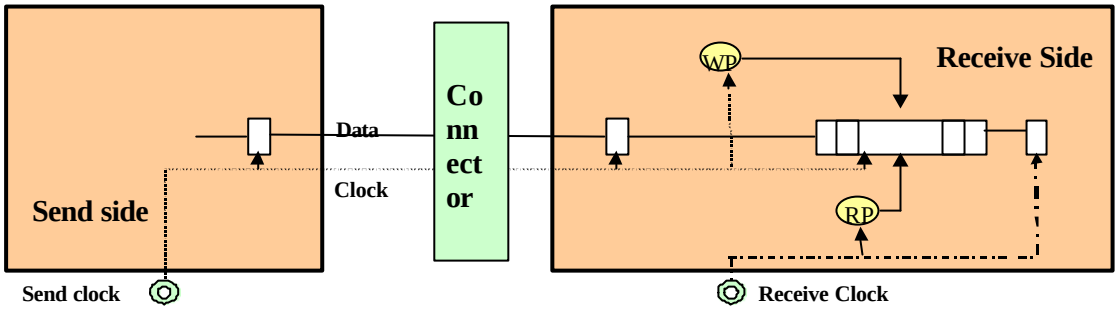

Source Synchronous method

For the address system, a memory controller called a System Controller (SC), where four CPUs and an I/O bridge are combined into a single unit, controls the cache. On these new models, each SC shares the control and maintains cache coherency.

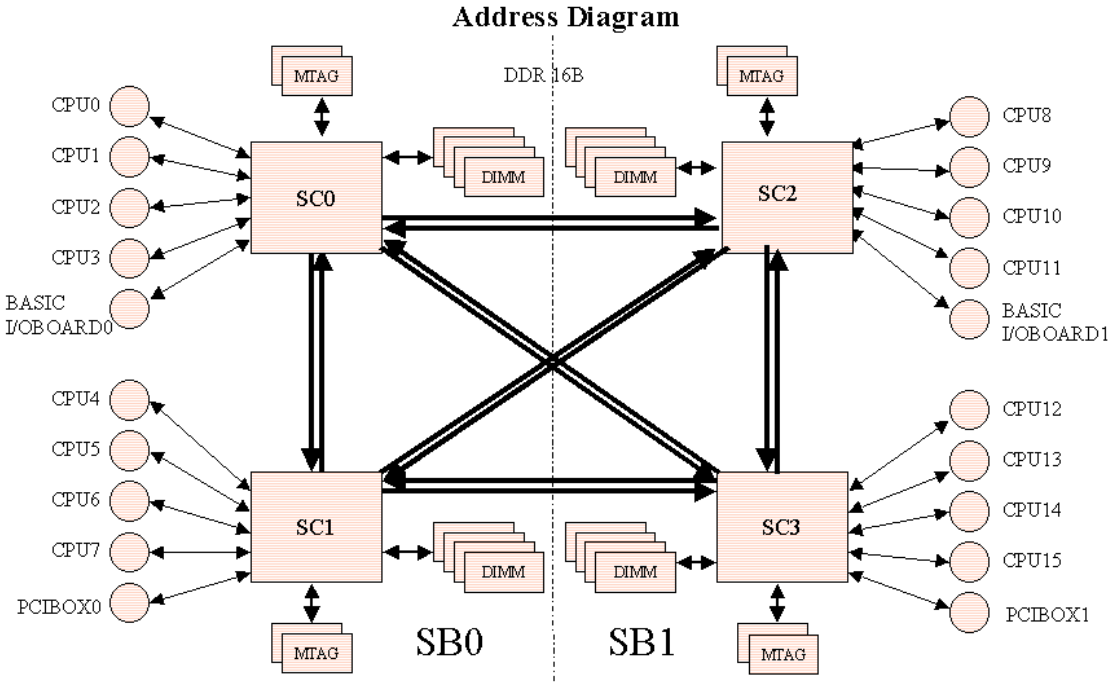

As shown in the figure below, the data system is divided into four groups. Each group connects with four CPUs and an I/O bridge.

Memory is also divided into four parts, with each part controlled by a SC and addresses are interleaved between two SCs on the system board. For applications that frequently access contiguous addresses, the memory is interleaved between the four SCs extending over the system boards. Performance is further improved by interleaving the memory between the four memory modules. In SDRAM, a group of bytes is set to select status for an address specified as the low address. If the low address, at the next memory access, matches, continuous memory access is enabled as the select status is already set. This technique is called low address matching.

The data in memory is protected using ECC, where single-bit errors are corrected automatically. ECC is also employed in the DTAG, used to maintain cache coherency. In addition, cache, memory, and TAG of the processor are all protected using ECC. The memory module of this model uses a configuration that enables memory to be continuously used even if one of the memory elements of the memory module fails completely. This method, previously adopted in high-end PRIMEPOWER 800 1000 and 2000 systems, is now applied to these 8-way class models. As a result high data path reliability is also achieved in these new models.

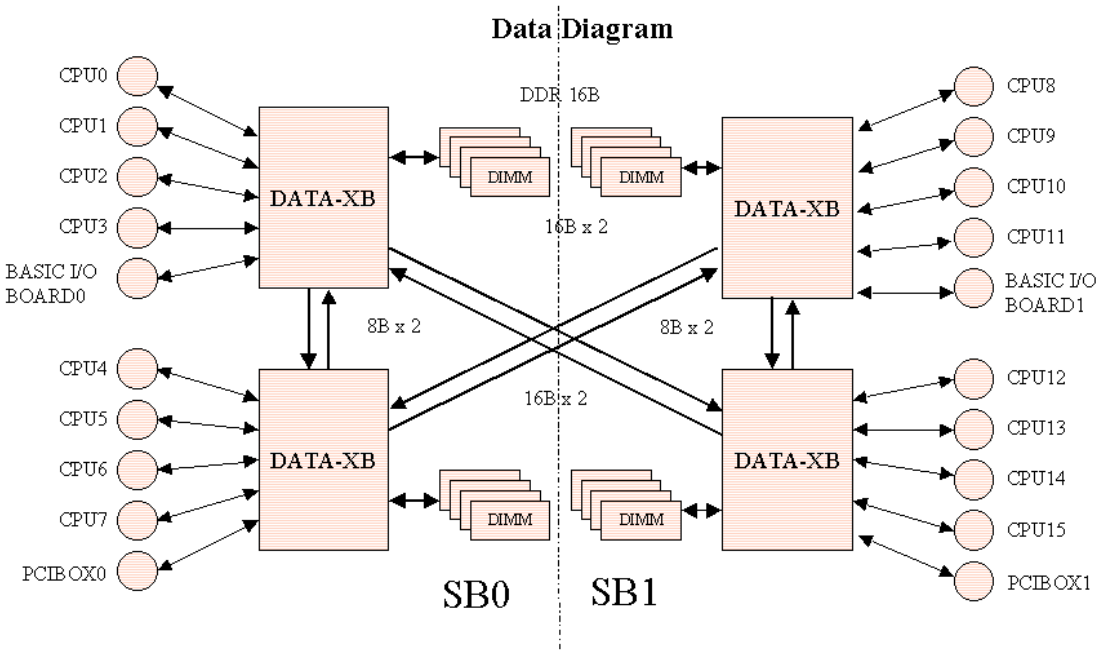

I/O Subsystem and PCIBOX

These models use a PCI bus as the I-O bus. The PCI bus conforms to PCI Rev. 2.1 and can operate in 32/64-bit mode at 33/66 MHz. Eight PCI slots per system board (33 MHz x 6 and 66 MHz x 2) are built into the cabinet. A PCIBOX (33 MHz x 9) has been developed as a PCI slot expansion unit and can be flexible installed into the system to meet customer needs as required. The PCIBOX is connected to the base cabinet using a bus called the CH bus, which provides ultra high-speed transfer at 200MHz. Four of the slots support installation of full-length cards. In addition, for reliability, the power supplies and fans are supplied in redundant configuration. Hot swapping is enabled.

[Basic slots]

Slot No.	PCI bus group	PCI Card Width	PCI Clock Rate	Card Input Voltage	Size
PCI0	A	64bit,/32bit	33 / 66MHz	3.3V or Universal	Short Card
PCI1	B	64bit,/32bit	33MHz	5V or Universal	Short Card
PCI2	B	64bit,/32bit	33MHz	5V or Universal	Short Card
PCI3	B	64bit,/32bit	33MHz	5V or Universal	Short Card
PCI4	C	64bit,/32bit	33 / 66MHz	3.3V or Universal	Short Card
PCI5	D	64bit,/32bit	33MHz	5V or Universal	Short Card
PCI6	D	64bit,/32bit	33MHz	5V or Universal	Short Card
PCI7	D	64bit,/32bit	33MHz	5V or Universal	Short or Long Card

[PCIBOX]

Slot No.	PCI bus group	PCI Card Width	PCI Clock Rate	Card Input Voltage	Size
PCI0	A	64bit,/32bit	33 / 66MHz	3.3V or Universal	Short Card
PCI1	B	64bit,/32bit	33MHz	5V or Universal	Short Card
PCI2	B	64bit,/32bit	33MHz	5V or Universal	Short Card
PCI3	B	64bit,/32bit	33MHz	5V or Universal	Short Card
PCI4	C	64bit,/32bit	33 / 66MHz	3.3V or Universal	Short Card
PCI5	D	64bit,/32bit	33MHz	5V or Universal	Short Card
PCI6	D	64bit,/32bit	33MHz	5V or Universal	Short Card
PCI7	D	64bit,/32bit	33MHz	5V or Universal	Short Card
PCI8	E	64bit,/32bit	33 / 66MHz	3.3V or Universal	Short or Long Card
PCI9	F	64bit,/32bit	33MHz	5V or Universal	Short or Long Card
PCI10	F	64bit,/32bit	33MHz	5V or Universal	Short or Long Card
PCI11	F	64bit,/32bit	33MHz	5V or Universal	Short or Long Card

Power Supply Subsystem

The power supply subsystem of the PRIMEPOWER 650 is capable of handling 100-120 VAC or 200-240 V AC at 50/60 Hz. For 100 V AC, the power supply subsystem is connected using three cables as standard. For 200 V AC, the power supply subsystem is connected using two cables as standard. The power supply subsystem of the PRIMEPOWER 850 is capable of handling 200-240 V AC at 50/60 Hz. For the PRIMEPOWER 850, one AC cable is used to connect the power supply subsystem as standard.

A power supply unit (FEP: Front End Power) is connected to the AC cables of each of the units. The FEP has been shrunk by employing high-density mounting techniques. In addition, a compact, highly efficient DC-DC converter has been developed as conversion unit for each load voltage. The compactness of this DC-DC converter contributes to the reduced design footprint of the overall unit.

An N + 1 redundant configuration (option) is supported for the FEP of each unit. In this case, a system failure can be avoided even if a FEP fails. Because the FEPs are also hot swappable, the system need not be stopped for maintenance.

A dual power feed option (option) is also supported. This option enables operation to continue using the other power supply even if a power supply error or power supply failure occurs. Either of these two options can be accommodated in the base cabinet.

A UPS is supported. By connecting a UPS, the system can be safely shut down when a power failure occurs and data corruption can be prevented.

Cooling system

PRIMEPOWER 650 and 850 both employ the latest cooling technology. In particular these models feature highly capable cooling systems that are compact and especially suitable for high-density mounting. A redundant configuration is used for the cooling fans. High availability is assured through the use of compact, high-performance cooling fans that as well as being extremely reliable are also hot swappable. Importantly the cooling capability of the fans is sufficient to handle future processor enhancements. Plus through the use of rotation speed control, fan noise can be kept below 50 dB, in most customer environments, while simultaneously providing suitable cooling performance even in high than average temperature environments.

To increase cooling efficiency, the system board and power supply-I/O section are cooled separately. Two CPU modules are mounted vertically in a row on the system board. Using the space under the first CPU module as a bypass passage, cold air is blown in at high speed to effectively cool the inherently hard-to-cool second CPU module. In the power supply-I/O section, special airflow control fins are used to achieve effective cooling.

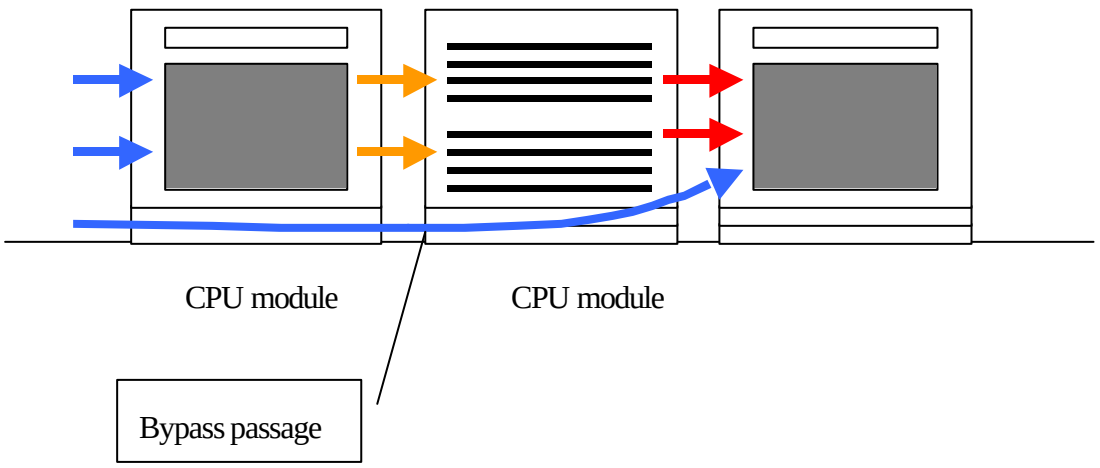

High Availability Functions

RAS/HA Concept

For PRIMEPOWER 650 and 850 models, HA (High Availability) functions, based on mainframe technology, have been employed. To implement HA functions, the RAS (Reliability, Availability, Serviceability) concept has to be guaranteed for each function.

High Availability, or the elimination of job stoppage time, cannot be achieved simply by increasing the reliability of the hardware components. High Availability must also be achieved for the software, applications, and support services. It is therefore essential to provide "system" RAS functions.

Of course to achieve high reliability, the quality of parts must be increased to the maximum. In addition, appropriate parts must be selected taking the product lifespan into consideration. However, there are no parts that can be guaranteed to never break, and it is therefore always necessary to consider the possibility of a failure. This applies to software as well as to hardware. Naturally, it is highly desirable to have software that is free of bugs. However, since there are software bugs that are triggered by hardware failures, it is extremely difficult to completely eliminate all bugs. Still, it goes without saying that all efforts must be made to improve the reliability of the hardware and software.

Fujitsu controls and guarantees the reliability of the parts used. When new parts are used, Fujitsu evaluates them by checking the lifespan using stress tests such as burn-in tests to determine whether the parts provide the reliability that the product aims for.

Availability can be expressed as an index indicating the time when the system is available for job operation. Because the number of errors cannot be kept to zero, mechanisms that ensure high availability must be installed to enable system operation to continue when a hardware failure occurs in a part or unit, an error occurs in the basic software such as the OS, or an error or failure occurs in an application process.

PRIMEPOWER 650 and 850 models incorporate the following basic mechanisms to provide for high availability:

- An expanded automatic error checking and correction range
- Improved retry functions when an error is detected and the provision of a degradation function which isolates failed components and allows a restart using a valid, if reduced, configuration.
- An automatic system restart to reduce down time
- A panel display function for error fault location at system startup
- Reduced system start time
- Redundant configurations for power supplies and fans and the provision of hot swappable components.

Serviceability refers to the functions that are used to quickly and easily recover the system from any problems that may occur during system operation. To achieve this, the cause of any occurring error must be identified, and the component or components that caused the error isolated and replaced. In addition, the event and conditions must be reported to the system administrator and maintenance personnel in a format that is easy to understand.

Machine management software is provided with PRIMEPOWER 650 and 850 to support the isolating of fault locations and the replacement of components without having to stop the system. This software also enables the system administrator and maintenance personnel to clearly identify the operating status of all units and to enable the maintenance personnel to perform the appropriate maintenance work.

Redundant Configuration and Hot Swapping

The power supply and fan units of these modules have a redundant configuration. Storage can also be installed in redundant configuration by using mechanisms such as dual RAID controllers and disk mirroring. This can be achieved for these models by combining Fujitsu's SynfinityDisk and a multipath disk control package. Moreover, SynfinityDisk can be used to mirror the system volumes themselves. Even if a disk error occurs at booting, the boot disk is switched automatically and the OS is restarted without the system process stopping.

Degradation Function Using Initial Diagnosis

At power-on, initial diagnosis is executed before the OS is booted. Initial diagnosis checks first whether the system boards are connected correctly. After the connections have been confirmed, SPARC64™ GP CPU operation is started and memory and I/O units are diagnosed. If initial diagnosis detects an error, the fault location is isolated. By default, fault isolation is performed in such a way that the isolation unit is the smallest possible. This initial diagnostic function enables isolation of faulty components so that operation will be performed using only normal components. Errors detected in the power-on diagnosis are posted to the system administrator and maintenance personnel via the system console.

Disk Subsystems

The various RAID and file units are supported as disk subsystems. The disk units mounted in these file units can be used as system volumes. Software such as SynfinityDisk can be used to enable the file units to have a RAID configuration. SynfinityDisk also enables mirroring of the system volumes.
In addition, to maintain High Availability, it is essential that the system can restart as quickly as possible after a system failure occurs. To quickly execute file system recovery (fsck) after a system failure occurs, a journal file system such as Fujitsu's SynfinityFile is required. In addition to faster fsck, installing SynfinityFile enables extension of the file system size that accompanies data expansion and load distribution. SynfinityFile's multi-volume function enables such expansion to be executed without having to stop the system.
When RAID units are used, setting a dual controller configuration can increase the fault tolerance of the RAID units and also enable duplication of the access paths including the PCI cards.

SynfinityDisk
SynfinityDisk is a software product that provides mirroring and hot spare functions between disk units in a single-system configuration and a mirroring function between shared disk units in a SynfinityCluster environment.

SynfinityFile
SynfinityFile is a UFS and API compatible file system. This software product provides a log function (metadata journal) and advanced functions such as multi-volumes and shared access.
- Log function: After a system failure occurs, fsck recovers integrity within seconds from its log information
- Multi-volume function: Multiple disks can be registered in one file system. Because the size can be easily extended without restoring, operating costs can be significantly reduced. In addition, each of the three elements (metadata, data, and log) of the file system size can be allocated to individual disks. This can significantly improve I/O characteristics and performance.
- Server sharing function (cluster file system): Simultaneous access from different servers is enabled by a standard API. Shared access of file systems is enabled even between different partitions.

Multipath Disk Control
This software product supports duplicated configuration of the interface between the disk array units.

Network Subsystem

SynfinityLink
In a network subsystem, SynfinityLink enables the construction of redundant network configurations of network communication components (network adapters, transmission lines, and so on) based on the HA functions of the PRIMEPOWER 650 and 850 models. SynfinityLink thereby provides communication functions with superior reliability and availability.
In addition, switching and system reconfiguration when an error occurs are automatically executed in the network subsystem. As a result, continuity of communication between applications is increased without needing to be aware of these redundant configurations.

Cluster System

In cluster system configurations, PRIMEPOWER 650 and 850 systems provide even more reliability, availability, and scalability. High operating ratio can be achieved with cluster systems using components such as SynfinityCluster, SynfinityDisk, SynfinityFile, and SynfinityLink with these servers. In addition, for the provision of very high performance, parallel processing of up to 16 servers can be used to implement a mode of operation called a "scalable cluster system." This is achieved with a combination of High Availability Solutions and Parallel Database Solutions, including SynfinityCluster, to provide the cluster functions; SynfinityDisk, SynfinityFile, and SynfinityLink for increased availability, and Synfinity-VIA for further increases performance, and network reliability.

Also provides is our new global cluster system "PRIMECLUSTER," an integration of Fujitsu Siemens Computers's Reliant Cluster with Synfinity series products. In addition to the HA cluster function that ensures application failover, PRIMECLUSTER provides a wider range of functions such as a network environment for higher bandwidth capability, network sub-system availability, a file system that allows concurrent direct sharing from multiple servers, volume management for advanced operability, availability of disk sub-systems in a SAN environment, and a variety of disk mirroring options. This ensures high availability in the whole system. PRIMECLUSTER even provides a global network where a parallel database using the Oracle9i Real Application Cluster and IP address is available on two or more nodes.

Basic Software

Software Configuration

The Solaris 64-bit operating environment developed by Sun Microsystems, Inc. is used as the basic software (server media kit) for all PRIMEPOWER models including 650 and 850. Solaris8 is a 64-bit OS that supports 64-bit SPARC processors based on SPARC V9 architecture. Solaris8 maintains high compatibility with conventional 32-bit applications. Solaris8 supports a 64-bit virtual address space and 64-bit file systems.

The software group required for efficient operation of PRIMEPOWER 650 and 850 is included in the Enhanced Support Facility supplied with PRIMEPOWER 650 and 850 as standard.

This software group is required for efficient operation of PRIMEPOWER server models, machine management, the Auto Power Control System (APCS) and System Control Facility (SCF) drivers, etc.

- Machine management

Machine management supports the settings for the hardware environment, status monitoring, and information collection. Machine management accesses the syslog messages output by the OS, plus the hardware, via the System Control Facility (SCF), and collects, analyzes, and displays information related to the hardware.

- Auto Power Control System

The Auto Power Control System (APCS) is a tool for automatically turning on and off the system power supply based on a set operational schedule. The APCS sets the power-on and power-off times for the System Control Facility (SCF).

Browser Based Software (WebSysAdmin)

WSA (WebSysAdmin) is provided as a tool for system administrators to support easier Browser-based system administration. This tool provides an interactive, easy-to-use graphical interface environment that can be used to display the state and mounting locations of PRIMEPOWER processing units (CPU, memory, etc.), display and delete process states, monitor log data, and execute user, software, and task administration.

Summary

Both PRIMEPOWER 650 and 850 offer high performance in a compact cabinet, flexible I-O configuration, plus investment protection, with future processor upgrades taken into consideration in their design. Fujitsu has pursued the High Availability concept with PRIMEPOWER 650 and 850 by employing redundancy and support for hot swapping of all essential components. Moreover, PRIMEPOWER 650 and 850 provide a full range of functions required by Enterprise Servers. This includes integrated operation and maintenance functions. In particular they support greater data sizes, handle higher demands, and flexibly respond to job changes, to mention only a few of their capabilities.