

SPARC Enterprise M8000/M9000 Servers

Overview Guide



ORACLE

SPARC

Part No.: E27468-01,
Manual Code: C120-E324-09EN
March 2012

Copyright © 2007, 2012, Fujitsu Limited. All rights reserved.

Oracle and/or its affiliates provided technical input and review on portions of this material.

Oracle and/or its affiliates and Fujitsu Limited each own or control intellectual property rights relating to products and technology described in this document, and such products, technology and this document are protected by copyright laws, patents, and other intellectual property laws and international treaties.

This document and the product and technology to which it pertains are distributed under licenses restricting their use, copying, distribution, and decompilation. No part of such product or technology, or of this document, may be reproduced in any form by any means without prior written authorization of Oracle and/or its affiliates and Fujitsu Limited, and their applicable licensors, if any. The furnishings of this document to you does not give you any rights or licenses, express or implied, with respect to the product or technology to which it pertains, and this document does not contain or represent any commitment of any kind on the part of Oracle or Fujitsu Limited, or any affiliate of either of them.

This document and the product and technology described in this document may incorporate third-party intellectual property copyrighted by and/or licensed from the suppliers to Oracle and/or its affiliates and Fujitsu Limited, including software and font technology.

Per the terms of the GPL or LGPL, a copy of the source code governed by the GPL or LGPL, as applicable, is available upon request by the End User. Please contact Oracle and/or its affiliates or Fujitsu Limited.

This distribution may include materials developed by third parties.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Fujitsu and the Fujitsu logo are registered trademarks of Fujitsu Limited.

All SPARC trademarks are used under license and are registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon architectures developed by Oracle and/or its affiliates. SPARC64 is a trademark of SPARC International, Inc., used under license by Fujitsu Microelectronics, Inc. and Fujitsu Limited. Other names may be trademarks of their respective owners.

United States Government Rights - Commercial use. U.S. Government users are subject to the standard government user license agreements of Oracle and/or its affiliates and Fujitsu Limited and the applicable provisions of the FAR and its supplements.

Disclaimer: The only warranties granted by Oracle and Fujitsu Limited, and/or any affiliate of either of them in connection with this document or any product or technology described herein are those expressly set forth in the license agreement pursuant to which the product or technology is provided. EXCEPT AS EXPRESSLY SET FORTH IN SUCH AGREEMENT, ORACLE OR FUJITSU LIMITED, AND/OR THEIR AFFILIATES MAKE NO REPRESENTATIONS OR WARRANTIES OF ANY KIND (EXPRESS OR IMPLIED) REGARDING SUCH PRODUCT OR TECHNOLOGY OR THIS DOCUMENT, WHICH ARE ALL PROVIDED AS IS, AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. Unless otherwise expressly set forth in such agreement, to the extent allowed by applicable law, in no event shall Oracle or Fujitsu Limited, and/or any of their affiliates have any liability to any third party under any legal theory for any loss of revenues or profits, loss of use or data, or business interruptions, or for any indirect, special, incidental or consequential damages, even if advised of the possibility of such damages.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.



Please
Recycle



Adobe PostScript

Copyright © 2007, 2012, Fujitsu Limited. Tous droits réservés.

Oracle et/ou ses sociétés affiliées ont fourni et vérifié des données techniques de certaines parties de ce composant.

Oracle et/ou ses sociétés affiliées et Fujitsu Limited détiennent et contrôlent chacune des droits de propriété intellectuelle relatifs aux produits et technologies décrits dans ce document. De même, ces produits, technologies et ce document sont protégés par des lois sur le copyright, des brevets, d'autres lois sur la propriété intellectuelle et des traités internationaux.

Ce document, le produit et les technologies afférents sont exclusivement distribués avec des licences qui en restreignent l'utilisation, la copie, la distribution et la décompilation. Aucune partie de ce produit, de ces technologies ou de ce document ne peut être reproduite sous quelque forme que ce soit, par quelque moyen que ce soit, sans l'autorisation écrite préalable d'Oracle et/ou ses sociétés affiliées et de Fujitsu Limited, et de leurs éventuels bailleurs de licence. Ce document, bien qu'il vous ait été fourni, ne vous confère aucun droit et aucune licence, expresses ou tacites, concernant le produit ou la technologie auxquels il se rapporte. Par ailleurs, il ne contient ni ne représente aucun engagement, de quelque type que ce soit, de la part d'Oracle ou de Fujitsu Limited, ou des sociétés affiliées de l'une ou l'autre entité.

Ce document, ainsi que les produits et technologies qu'il décrit, peuvent inclure des droits de propriété intellectuelle de parties tierces protégés par copyright et/ou cédés sous licence par des fournisseurs à Oracle et/ou ses sociétés affiliées et Fujitsu Limited, y compris des logiciels et des technologies relatives aux polices de caractères.

Conformément aux conditions de la licence GPL ou LGPL, une copie du code source régi par la licence GPL ou LGPL, selon le cas, est disponible sur demande par l'utilisateur final. Veuillez contacter Oracle et/ou ses sociétés affiliées ou Fujitsu Limited.

Cette distribution peut comprendre des composants développés par des parties tierces.

Des parties de ce produit peuvent être dérivées des systèmes Berkeley BSD, distribués sous licence par l'Université de Californie. UNIX est une marque déposée aux États-Unis et dans d'autres pays, distribuée exclusivement sous licence par X/Open Company, Ltd.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses sociétés affiliées. Fujitsu et le logo Fujitsu sont des marques déposées de Fujitsu Limited.

Toutes les marques SPARC sont utilisées sous licence et sont des marques déposées de SPARC International, Inc., aux États-Unis et dans d'autres pays. Les produits portant la marque SPARC reposent sur des architectures développées par Oracle et/ou ses sociétés affiliées. SPARC64 est une marque de SPARC International, Inc., utilisée sous licence par Fujitsu Microelectronics, Inc. et Fujitsu Limited. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires.

United States Government Rights - Commercial use. U.S. Government users are subject to the standard government user license agreements of Oracle and/or its affiliates and Fujitsu Limited and the applicable provisions of the FAR and its supplements.

Avis de non-responsabilité : les seules garanties octroyées par Oracle et Fujitsu Limited et/ou toute société affiliée de l'une ou l'autre entité en rapport avec ce document ou tout produit ou toute technologie décrits dans les présentes correspondent aux garanties expressément stipulées dans le contrat de licence régissant le produit ou la technologie fournis. SAUF MENTION CONTRAIRE EXPRESSÉMENT STIPULÉE DANS CE CONTRAT, ORACLE OU FUJITSU LIMITED ET LES SOCIÉTÉS AFFILIÉES À L'UNE OU L'AUTRE ENTITÉ REJETTENT TOUTE REPRÉSENTATION OU TOUTE GARANTIE, QUELLE QU'EN SOIT LA NATURE (EXPRESSE OU IMPLICITE) CONCERNANT CE PRODUIT, CETTE TECHNOLOGIE OU CE DOCUMENT, LESQUELS SONT FOURNIS EN L'ÉTAT. EN OUTRE, TOUTES LES CONDITIONS, REPRÉSENTATIONS ET GARANTIES EXPRESSES OU TACITES, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE À LA QUALITÉ MARCHANDE, À L'APTITUDE À UNE UTILISATION PARTICULIÈRE OU À L'ABSENCE DE CONTREFAÇON, SONT EXCLUES, DANS LA MESURE AUTORISÉE PAR LA LOI APPLICABLE. Sauf mention contraire expressément stipulée dans ce contrat, dans la mesure autorisée par la loi applicable, en aucun cas Oracle ou Fujitsu Limited et/ou l'une ou l'autre de leurs sociétés affiliées ne sauraient être tenues responsables envers une quelconque partie tierce, sous quelque théorie juridique que ce soit, de tout manque à gagner ou de perte de profit, de problèmes d'utilisation ou de perte de données, ou d'interruptions d'activités, ou de tout dommage indirect, spécial, secondaire ou consécutif, même si ces entités ont été préalablement informées d'une telle éventualité.

LA DOCUMENTATION EST FOURNIE « EN L'ÉTAT » ET TOUTE AUTRE CONDITION, DÉCLARATION ET GARANTIE, EXPRESSE OU TACITE, EST FORMELLEMENT EXCLUE, DANS LA MESURE AUTORISÉE PAR LA LOI EN VIGUEUR, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE À LA QUALITÉ MARCHANDE, À L'APTITUDE À UNE UTILISATION PARTICULIÈRE OU À L'ABSENCE DE CONTREFAÇON.

Contents

Preface ix

1. System Overview 1-1

- 1.1 Product Overview 1-1
- 1.2 System Specifications 1-9
 - 1.2.1 Environmental Specifications 1-12
 - 1.2.2 Power Specifications 1-13
 - 1.2.3 M8000 Server Components 1-14
 - 1.2.4 M9000 Server Components (Base Cabinet Only) 1-16
 - 1.2.5 M9000 Server Components (With an Expansion Cabinet) 1-18
 - 1.2.6 Operator Panel Overview 1-19
- 1.3 Server Components 1-22
 - 1.3.1 CPU Module 1-22
 - 1.3.2 CPU/Memory Board Unit 1-23
 - 1.3.3 I/O Unit 1-23
 - 1.3.4 FAN Unit 1-24
 - 1.3.5 Power Supply Unit 1-24
 - 1.3.6 Crossbar Unit 1-24
 - 1.3.7 Clock Control Unit 1-24
 - 1.3.8 Operator Panel 1-24

- 1.3.9 XSCF Unit 1–25
- 1.3.10 Internal Drive Units 1–25
- 1.4 Component Mounting Conditions 1–26
- 1.5 Optional Products 1–27
 - 1.5.1 Power Supply Options 1–27
 - 1.5.2 External I/O Expansion Unit 1–28
 - 1.5.3 M9000 Server (Expansion Cabinet) Option 1–29
- 1.6 Software Features 1–29

- 2. System Features 2–1**
 - 2.1 Hardware Configuration 2–1
 - 2.1.1 CPU 2–1
 - 2.1.1.1 Mounted Processors and CPU Operational Modes 2–2
 - 2.1.2 Memory Subsystem 2–5
 - 2.1.3 I/O Subsystem 2–5
 - 2.1.4 System Bus 2–5
 - 2.1.5 System Control 2–7
 - 2.2 Partitioning 2–7
 - 2.2.1 Features 2–8
 - 2.2.2 Domain Hardware Requirements 2–8
 - 2.2.3 Domain Configuration 2–10
 - 2.3 Resource Management 2–12
 - 2.3.1 Dynamic Reconfiguration 2–12
 - 2.3.2 PCI Hot-plug 2–13
 - 2.3.3 Capacity on Demand 2–13
 - 2.3.4 Oracle Solaris Zones 2–13
 - 2.4 RAS 2–13
 - 2.4.1 Reliability 2–14
 - 2.4.2 Availability 2–14

2.4.3 Serviceability 2-15

3. About Software 3-1

3.1 Oracle Solaris OS Functions 3-1

3.1.1 Domain Management 3-2

3.1.2 PCI Hot-plug 3-2

3.2 XSCF Firmware Function 3-2

3.2.1 XSCF Features 3-2

3.2.1.1 Command Line-based User Interface (XSCF Shell) 3-3

3.2.1.2 Browser-Based User Interface (XSCF Web) 3-4

3.2.2 XSCF Functional Overview 3-4

3.2.2.1 System Management 3-4

3.2.2.2 Security Management 3-4

3.2.2.3 System Status Management 3-5

3.2.2.4 Error Detection and Management 3-5

3.2.2.5 Remote System Control and Monitoring 3-5

3.2.2.6 Resource Management 3-5

3.2.2.7 Airflow Indicator 3-6

Index Index-1

Preface

This overview guide describes system features, system configurations, hardware functions, and software functions of the SPARC Enterprise M8000/M9000 servers from Oracle and Fujitsu. References herein to the M8000 server or M9000 server are references to the SPARC Enterprise M8000 or SPARC Enterprise M9000 server.

This chapter includes the following sections:

- [“Audience” on page ix](#)
- [“Related Documentation” on page x](#)
- [“Text Conventions” on page xi](#)
- [“Notes on Safety” on page xi](#)
- [“Syntax of the Command-Line Interface \(CLI\)” on page xii](#)
- [“Documentation Feedback” on page xii](#)

Audience

This guide is written for experienced system administrators with working knowledge of computer networks and advanced knowledge of the Oracle Solaris Operating System (Oracle Solaris OS).

Related Documentation

All documents for your sever are available online at the following locations:

Documentation	Link
Sun Oracle software-related manuals (Oracle Solaris OS, and so on)	http://www.oracle.com/documentation
Fujitsu documents	http://www.fujitsu.com/sparcenterprise/manual/
Oracle M-series server documents	http://www.oracle.com/technetwork/documentation/sparc-mseries-servers-252709.html

The following table lists titles of related documents.

Related SPARC Enterprise M8000/M9000 Servers Documents
<i>SPARC Enterprise M8000/M9000 Servers Site Planning Guide</i>
<i>SPARC Enterprise M8000/M9000 Servers Getting Started Guide*</i>
<i>SPARC Enterprise M8000/M9000 Servers Overview Guide</i>
<i>SPARC Enterprise M3000/M4000/M5000/M8000/M9000 Servers Important Legal and Safety Information *</i>
<i>SPARC Enterprise M8000/M9000 Servers Safety and Compliance Guide</i>
<i>External I/O Expansion Unit Safety and Compliance Guide</i>
<i>SPARC Enterprise M8000/M9000 Servers Unpacking Guide*</i>
<i>SPARC Enterprise M8000/M9000 Servers Installation Guide</i>
<i>SPARC Enterprise M8000/M9000 Servers Service Manual</i>
<i>External I/O Expansion Unit Installation and Service Manual</i>
<i>SPARC Enterprise M3000/M4000/M5000/M8000/M9000 Servers Administration Guide</i>
<i>SPARC Enterprise M3000/M4000/M5000/M8000/M9000 Servers XSCF User's Guide</i>
<i>SPARC Enterprise M3000/M4000/M5000/M8000/M9000 Servers XSCF Reference Manual</i>
<i>Dynamic Reconfiguration (DR) User's Guide</i>
<i>SPARC Enterprise M4000/M5000/M8000/M9000 Servers Capacity on Demand (COD) User's Guide</i>
<i>SPARC Enterprise M3000/M4000/M5000/M8000/M9000 Servers Product Notes†</i>
<i>SPARC Enterprise M8000/M9000 Servers Product Notes</i>
<i>External I/O Expansion Unit Product Notes</i>
<i>SPARC Enterprise M3000/M4000/M5000/M8000/M9000 Servers Glossary</i>

* This is a printed document.

† Beginning with the XCP 1100 release.

Text Conventions

This manual uses the following fonts and symbols to express specific types of information.

Font/Symbol	Meaning	Example
AaBbCc123	What you type, when contrasted with on-screen computer output. This font represents the example of command input in the frame.	XSCF> adduser jsmith
AaBbCc123	The names of commands, files, and directories; on-screen computer output. This font represents the example of command output in the frame.	XSCF> showuser -P User Name: jsmith Privileges: useradm auditadm
<i>Italic</i>	Indicates the name of a reference manual, a variable, or user-replaceable text.	See the <i>SPARC Enterprise M3000/M4000/M5000/M8000/M9000 Servers XSCF User's Guide</i> .
" "	Indicates names of chapters, sections, items, buttons, or menus.	See Chapter 2, "System Features."

Notes on Safety

Read the following documents thoroughly before using or handling any SPARC Enterprise M8000/M9000 server.

- *SPARC Enterprise M3000/M4000/M5000/M8000/M9000 Servers Important Legal and Safety Information*
- *SPARC Enterprise M8000/M9000 Servers Safety and Compliance Guide*

Syntax of the Command-Line Interface (CLI)

The command syntax is as follows:

- A variable that requires input of a value must be put in *Italics*.
- An optional element must be enclosed in [].
- A group of options for an optional keyword must be enclosed in [] and delimited by |.

Documentation Feedback

If you have any comments or requests regarding this document, go to the following websites:

- For Oracle users:

<http://www.oracle.com/goto/docfeedback>

Include the title and part number of your document with your feedback:

SPARC Enterprise M8000/M9000 Servers Overview Guide,
part number E27468-01

- For Fujitsu users:

http://www.fujitsu.com/global/contact/computing/sparce_index.html

System Overview

This chapter provides an overview of features, specifications, and configurations of the SPARC Enterprise M8000/M9000 servers.

- [Section 1.1, “Product Overview” on page 1-1](#)
- [Section 1.2, “System Specifications” on page 1-9](#)
- [Section 1.3, “Server Components” on page 1-22](#)
- [Section 1.4, “Component Mounting Conditions” on page 1-26](#)
- [Section 1.5, “Optional Products” on page 1-27](#)
- [Section 1.6, “Software Features” on page 1-29](#)

1.1 Product Overview

This section describes the features and appearances of M8000/M9000 servers.

M8000/M9000 servers have been developed as UNIX servers using a symmetric multi-processing (SMP) architecture. Each of these systems merges mainframe technologies for high reliability, and the associated know-how accumulated over time, with the high-speed technologies of super computers and the openness of UNIX server development.

If a problem occurs during operation, the errors causing them can be corrected or isolated without stopping the system. This feature minimizes problems in many cases, thereby improving job continuity.


Each of the M8000/M9000 servers contains one or more multicore SPARC64 VI, SPARC64 VII, and SPARC64 VII+ processors. They can operate as multiple servers that permit flexible use of resources, including more efficient execution of job operations. In the M8000/M9000 servers, the SPARC64 VI, SPARC64 VII, and SPARC64 VII+ processors can be mounted in combination.

Each server consists of a cabinet containing various mounted components, a front door, rear door, and side panels as parts of the server structure. An operator panel is mounted on the front door and is always accessible. Take special care in handling and storing the dedicated key that is provided for the front door and the operator panel.

[FIGURE 1-1](#) through [FIGURE 1-3](#) show exterior views of the servers.


M8000 Server Appearance

FIGURE 1-1 M8000 Server (Front View)



M9000 Server Appearance (Base Cabinet Only)

FIGURE 1-2 M9000 Server (Base Cabinet Only)



M9000 Server Appearance (With an Expansion Cabinet)

The expansion cabinet is an M9000 option connected to the M9000 (type for the base cabinet only).

FIGURE 1-3 M9000 Server (With an Expansion Cabinet)



M8000/M9000 servers have the following features:

- Multicore SPARC64 VI/SPARC64 VII/SPARC64 VII+ gigahertz processors

These processors provide superior performance, due to their high scalability allowing expansion to up to 64 dual-core CPU modules, and technologies enabling high-speed arithmetic operations and data transfers.

Reliability and availability are enhanced with Error Checking and Correction (ECC) data protection and instruction retry.

As faster and higher-performing CPU modules become available, they can be added to or replace existing installed CPU modules to further improve performance.

The system uses symmetric multiple-processing (SMP), so each CPU can access any part of system memory regardless of its mounting location. Adding more CPUs does not affect memory access to any of the installed CPUs.

- High-speed crossbar-type system bus

The high-speed crossbar-type system bus that provides high-speed wide-band data transfer maximizes the performance of the SPARC64 VI/SPARC64 VII/SPARC64 VII+ processors.

- ECC memory

ECC functionality protects data on all system buses and in memory, so that any errors in data are automatically corrected. In addition to ECC memory, Chipkill memory protection is supported.

- PCI Express (PCIe) is mounted as I/O buses

PCIe, with a maximum bus width of eight lanes is used for the inter-connect bus with the I/O device.

- An optional External I/O Expansion Unit enables I/O slot expansion in the system.

Connect an External I/O Expansion Unit to add more PCIe and PCI-X slots to the server.

An External I/O Expansion Unit is connected by a cable to a link card plugged into a PCIe slot in an I/O unit.

- The main components, power supply units, and FAN units can be configured to be redundant

Redundant configurations can be used for the main components, such as a power supply unit, FAN unit, hard disk drive, and PCI card. Implementation of redundant configurations enables operation to continue without interruption even if one of the units making up part of the system fails.

- The main components, power supply units, and FAN units support active (hot) replacement/addition

Component replacement and addition during system operation supported for the main components, such as a power supply unit, FAN unit, hard disk drive, system control facility (board), system board, and PCI card, with some exceptions.

Dynamic reconfiguration (DR) is used for active replacement and addition of CPU/memory board unit and I/O unit configuring the system board.

PCI hot-plug function enables replacement and addition of PCI cards while the system is running.

- Automatic reboot on failure

If a failure occurs, the faulty component is automatically isolated from the system, and the system is rebooted. If 1-bit errors occur frequently in the cache memory configuring a CPU, the faulty memory can be dynamically isolated without rebooting the Oracle Solaris Operating System (Oracle Solaris OS). This type of graceful degradation function enables the operation of the other resources to continue without interruption, and also provides high fault-tolerance in case of failure.

- Uninterruptible Power Supply (UPS) controller

For measures against commercial power failure, the server is equipped with UPS controller (UPC) ports. Using a UPS enables stable power supply to the system when a power failure or an extensive power interruption occurs.

- Hardware RAID function

Two hard disks which are connected to the IOU onboard device card (IOUA) can be constructed as a single logical volume. The mirrored configuration of the constructed logical volume can secure the data redundancy, as well as achieving the improvement in system fault tolerance.

Note – Hardware RAID is supported only on the M8000/M9000 servers on which a RAID-capable onboard device card (IOUA) is mounted. The RAID-capable IOUA requires the minimum XSCF firmware of XCP 1110 with the applicable Oracle Solaris OS patches. For the latest patch information, see the latest product notes.

Note – When the onboard device card is RAID-capable, the `showhardconf(8)` command displays Type 2 in the output.

```
PCI#0 Name_Property:pci; Card_Type:IOUA;
      + Serial:PP0611T826 ; Type:2;
      + FRU-Part-Number:CA21138-B84X 010AE/371-5000-05 ;
```

- eXtended System Control Facility (XSCF)

This product has a built-in service processor (called the eXtended System Control Facility (XSCF)), which monitors the system temperature, the hardware status of the power supply unit and FAN unit, and the operating status of domains.

You can configure the system to selectively degrade a faulty component for operation if supply failure is detected.

Scheduling is supported to enable automatic power-on and power-off of the server according to the specified operation schedule.

The console of each domain can be controlled from the XSCF via a network.

A browser-based user interface (BUI) and the command line-based interface facilitate operations for making configuration changes and status monitoring in the system.

Note – A console display terminal is required for console control. Prepare it before installation. The devices that can be used as the terminal are listed below.

- PC
- Workstation
- ASCII terminal
- Terminal server (or a touch panel connected to a terminal server)

Note – For the console connection method, see the *SPARC Enterprise M8000/M9000 Servers Installation Guide*.

- Partitioning function

One high-end server can be divided into multiple areas, or domains, for more effective scalability. Each domain manages resources in linkage with the XSCF. A domain may consist of optimized resources depending on its intended use, enabling more efficient system configurations.

Dynamic Reconfiguration (DR) enables adding, deleting, and relocating resources of domains without stopping processing in the domain. This enables dynamic reconfiguration of resources without stopping a job, even when the job load increases suddenly or when a faulty component is replaced.

For details on domain functions, see the *SPARC Enterprise M3000/M4000/M5000/M8000/M9000 Servers Administration Guide*.

For details of the DR function, see the *SPARC Enterprise M4000/M5000/M8000/M9000 Servers Dynamic Reconfiguration (DR) User's Guide*.

- The Oracle Solaris OS (Oracle Solaris 10 or later) is supported.

With an added function for error prediction and self-recovery by the system (Predictive Self-Healing) and enhanced process privilege management and network functions, the Oracle Solaris OS sets new standards for performance, efficiency, availability, and security.

- Capacity on Demand (COD)

The Capacity on Demand (COD) feature allows you to configure spare processing resources on your server in the form of one or more COD CPUs which can be activated at a later date when additional processing power is needed.

For details, see the SPARC *SPARC Enterprise M4000/M5000/M8000/M9000 Servers Capacity on Demand (COD) User's Guide*.

1.2 System Specifications

This section describes the specifications and the environmental conditions of both high-end servers, shows names of server components, and provides an overview of the operator panel.

Note – Contact your sales representative for tape drive unit options on M8000/M9000 servers.

TABLE 1-1 Main Unit Specifications

Item		M9000 Server		
		M8000 Server	Base Cabinet Only	Base Cabinet + Expansion Cabinet
Type		Floor-stand type**		
CPU	Type	SPARC64 VI: 2 cores/1 CPU modules SPARC64 VI compatible mode		
	Number of CPUs	32 cores (maximum 16 CPU modules)	64 cores (maximum 32 CPU modules)	128 cores (maximum 64 CPU modules)
	Type	SPARC64 VII/SPARC64 VII+: 4 cores/1 CPU modules SPARC64 VI compatible mode/SPARC64 VII enhanced mode		
	Number of CPUs	64 cores (maximum 16 CPU modules)	128 cores (maximum 32 CPU modules)	256 cores (maximum 64 CPU modules)
Main storage (memory module)	Maximum memory size	1 TB ^{††}	2 TB ^{††}	4 TB ^{††}
	Error checking function	Error Checking and Correction (ECC)		
PCI slot built into servers (PCI Express)*		Maximum 32 slots	Maximum 64 slots	Maximum 128 slots
External I/O Expansion Unit (maximum number of connections)		8 units (16 boats)	16 units (32 boats)	16 units (32 boats)
Maximum number of slots, with I/O boats mounted		112 slots	224 slots	288 slots
Hard disk drive [†]		16 slots	32 slots	64 slots
CD-RW/DVD-RW drive unit		1 drive		2 drives

TABLE 1-1 Main Unit Specifications (*Continued*)

Item	M9000 Server			
	M8000 Server	Base Cabinet Only	Base Cabinet + Expansion Cabinet	
Tape drive unit	1 drive can be mounted (option)		2 drives can be mounted (option)	
FAN unit	4 units (type A) 8 units (type B)	16 units (type A)	32 units (type A)	
Power supply unit (Maximum number of mounted units) (single phase, one system)	9 units	15 units	30 units	
Redundant configuration	Power supply unit, FAN unit, XSCF, power system (dual power feed option), and clock supply system			
Components that can be active replacement	CPU/memory board unit, I/O unit, XSCF unit, hard disk drive, link card, CD-RW/DVD-RW drive unit, tape drive unit, PCI cassette, power supply unit, FAN unit, DC-DC converter			
Components that can be hot replacement	CPU/memory board unit, I/O unit, XSCF unit, link card, CD-RW/DVD-RW drive unit, tape drive unit, PCI cassette, power supply unit, FAN unit, DC-DC converter			
System control interface	LAN, serial, UPS (Uninterruptible Power Supply) interface, Remote Cabinet Interface (RCI) ^{††} , and USB ^{***}			
Number of domains	16	24	24	
Operating environment	Oracle Solaris OS ^{†††}			
Architecture	Platform group: sun4u Platform name: SUNW, SPARC-Enterprise			
Outside dimensions	Width [mm]	750	850	1674
	Depth [mm]	1260	1260	1260
	Height [mm]	1800	1800	1800
Weight [kg]	700	940	1880	
RFID	RFID standard	ISO/IEC18000-6 TypeC compliant (EPC GEN2 compliant) FSTC RFID tag standard compliant		

TABLE 1-1 Main Unit Specifications (*Continued*)

Item	M9000 Server		
	M8000 Server	Base Cabinet Only	Base Cabinet + Expansion Cabinet
Frequency	860 - 960 MHz ^{‡‡‡}		
Write data	EPCglobal GIAI-96 format ^{****}		
Read range (reference) [‡]	1. 902 - 928 MHz Fixed reader with a maximum output of 4 Watt EIRP: Up to 1.8 m (6 ft) Handheld reader with a maximum output of 2 Watt EIRP: Up to 90 cm (3 ft) 2. 865.6 - 867.6 MHz Fixed reader with a maximum output of 3.2 Watt EIRP: Up to 1.8 m (6 ft) Handheld reader with a maximum output of 1 Watt EIRP: Up to 90 cm (3 ft) 3. 952 - 955 MHz Fixed reader with a maximum output of 4 Watt EIRP: Up to 1.8 m (6 ft) Handheld reader with a maximum output of 0.5 Watt EIRP: Up to 90 cm (3 ft)		

* Up to eight lanes of PCIe bus are connected to each slot.

† A built-in IOU onboard device card (IOUA) is required for using hard disk drive. Active replacement of IOUA is not possible.

‡ The range shows the distance between a RFID tag and a RFID tag reader. The values are for reference measured with the RFID tag reader antenna which is faced toward the RFID tag front. The read range might be shorter or longer than this value for reference, depending on the sending/receiving performance (the size or the method) of the RFID tag reader antenna or the surrounding environment (interference due to the radio wave reflected by the floor, ceiling, or the indoor metallic object).

** The upper part of the M8000 server cabinet has a 12 rack units (RU) space.

†† This is the maximum capacity when 8-GB dual inline memory modules (DIMM) are installed.

‡‡ For information on whether the RCI function is supported for your server, see the *SPARC Enterprise M3000/M4000/M5000/M8000/M9000 Servers Product Notes*.

*** This interface is only used for maintenance by authorized service personnel. It does not support except the maintenance-purpose USB devices.

††† For the latest information on the operating system, visit our website, or contact your sales representative.

‡‡‡ The usable frequency is defined by the government of each country.

**** It is ensured that the written data value is unique, that is, not overlapped with the values written in other RFID tags which comply with the GIAI-96 format. This value is not the serial number of the server.

1.2.1 Environmental Specifications

TABLE 1-2 Environmental Specifications

	Operating Range	Non-Operating Range	Optimum
Ambient temperature	5°C to 32°C (41°F to 89.6°F)	Unpacked: 0°C to 50°C (32°F to 122°F) Packed: -20°C to 60°C (-4°F to 140°F)	21°C to 23°C (70°F to 74°F)
Relative humidity*	20% RH to 80% RH	to 93% RH	45% RH to 50% RH
Altitude restriction†	3,000 m (10,000 ft)	12,000 m (40,000 ft)	
Temperature conditions	5°C to 32°C (41°F to 89.6°F): 0 to less than 1500 m (0 to less than 4921 ft) 5°C to 30°C (41°F to 86°F): 1500 m to less than 2000 m (4921 ft to less than 6562 ft) 5°C to 28°C (41°F to 82.4°F): 2000 m to less than 2500 m (6562 ft to less than 8202 ft) 5°C to 26°C (41°F to 78.8°F): 2500 m to 3000 m (8202 ft to 9843 ft)		

* There is no condensation regardless of the temperature and humidity.

† All altitudes are above sea level.

1.2.2 Power Specifications

Two power input modes are available; single-phase power feed and three-phase power input.

To use a three-phase power source, a three-phase power feed option and a power supply cabinet for mounting the option are required. The three-phase power feed has two connection options: a star connection that connects a neutral line and each phase, and a delta connection that connects each phase.

For details, see the *SPARC Enterprise M8000/M9000 Servers Site Planning Guide*.

TABLE 1-3 shows samples of power consumption of specific configurations and program load. The power consumption of the system varies depending on configuration of the system, characteristics of your running programs and ambient temperature.

TABLE 1-3 Power Consumption Examples

Item		M8000		M9000	
			Base Cabinet Only	Base Cabinet + Expansion Cabinet	
Ambient temperature		25°C	25°C	25°C	25°C
Configuration*	CPU/memory board unit: 2.52GHz CPU x 4, 4GB DIMM x 32	4	8		16
	I/O unit: 73GB HDD x 4, PCIe card x 8	4	8		16
Power consumption†		7.48 kW	14.64 kW		29.96 kW

* 10Watt PCIe cards are installed.

† These power consumption numbers are from a typical workload. You may see higher power consumption values depending on the characteristics of your workload.

1.2.3 M8000 Server Components

FIGURE 1-4 and FIGURE 1-5 show the front and rear views of the M8000 server with a power cabinet connected. The names of server components are shown in each figure.

The dual power feed option and three-phase power feed option can be mounted in the power cabinet. One power cabinet is connected to the M8000 server.

FIGURE 1-4 M8000 Server and Power Cabinet Front View





FIGURE 1-5 M8000 Server and Power Cabinet Rear View



1.2.4 M9000 Server Components (Base Cabinet Only)

FIGURE 1-6 and FIGURE 1-7 show the front and rear views of the M9000 server (base cabinet only) with a power cabinet connected. The names of server components are shown in each figure.

One power cabinet is connected to the M9000 server (base cabinet only).

FIGURE 1-6 M9000 Server (Base Cabinet Only) and Power Cabinet Front View





FIGURE 1-7 M9000 Server (Base Cabinet Only) and Power Cabinet Rear View



1.2.5 M9000 Server Components (With an Expansion Cabinet)

FIGURE 1-8 and FIGURE 1-9 show the front and rear views of the M9000 server with an expansion cabinet and power cabinets connected. The names of server components are shown in each figure.

One power cabinet is connected to each of M9000 server base cabinet and expansion cabinet.

FIGURE 1-8 M9000 Server (With an Expansion Cabinet) and Power Cabinet Front View





FIGURE 1-9 M9000 Server (With an Expansion Cabinet) and Power Cabinet Rear View



1.2.6 Operator Panel Overview

The operator panel has LEDs indicating different states of the M8000 and M9000 servers, a power switch for power control, and a mode switch for setting the operation mode.

The operator panel is mounted on the front panel.


For details about the operator panel, see the *SPARC Enterprise M8000/M9000 Servers Service Manual*.

The following figure shows the operator panel, and its LEDs and switches are described below.

Operator Panel Appearance

FIGURE 1-10 shows the operator panel.




FIGURE 1-10 Operator Panel



Operator Panel LEDs

TABLE 1-4 lists the operating states indicated by the LEDs on the operator panel.

TABLE 1-4 Operator Panel LEDs





LEDs	Name	Light Color	Description of Function and Operating State
	POWER	Green	Indicates whether the main unit power is on. If this LED is on, the power is on. If this LED is blinking, the power-off sequence is in progress.
	STANDBY	Green	Indicates the standby state of the main unit. If this LED is on, the power can be turned on.
	CHECK	Amber	Indicates the main unit operating status. (This is used to indicate a maintenance target, or it indicates that the unit cannot be started.) If this LED is on, a system error has been detected.

Operator Panel Switches

The switches on the operator panel include the mode switch for setting the operation mode and the POWER switch for turning on and off the server. To switch between system operation mode and maintenance mode, insert the dedicated key of the high-end server and change the mode switch setting.

TABLE 1-5 lists functions of the switches on the operator panel.

TABLE 1-5 Operator Panel Switches

Switch	Name	Function
	POWER switch	Controls the main unit power.
	MODE switch	Selects between maintenance and normal operation. Use the dedicated key managed by the customer to switch between normal and maintenance modes.
	Locked	This mode is set for normal operation.
	Service	This mode is set for maintenance.

1.3 Server Components

This section describes the components of both high-end servers.

For details on each, see the *SPARC Enterprise M8000/M9000 Servers Service Manual*.

- CPU Module
- CPU/Memory Board Unit
- I/O Unit
- FAN Unit
- Power Supply Unit
- Crossbar Unit
- Clock Control Unit
- Operator Panel
- XSCF Unit
- Internal Drive Units
 - Hard Disk Drive
 - CD-RW/DVD-RW Drive Unit
 - Tape Drive Unit

1.3.1 CPU Module

The CPU module (CPUM) contains a SPARC64 VI/SPARC64 VII/SPARC64 VII+ processor and a DC-DC converter (DDC). Up to four CPUMs can be mounted on a CPU/memory board unit.

The CPUM has the following features:

- The CPUM contains a SPARC64 VI/SPARC64 VII/SPARC64 VII+, a high-performance multicore CPU that uses the latest LSI process.
- If an unforeseen fault is detected, the SPARC64 VI/SPARC64 VII/SPARC64 VII+ processor enables operation to continue without interruption by using an automatic recovery function, automatic retry function, or automatic degradation function, depending on how the system is configured.
- A redundant DDC configuration enables continuous operation even if a DDC fails.

1.3.2 CPU/Memory Board Unit

The CPU/memory board unit (CMU) contains CPUMs, memory modules, and a DDC. The CMU with an I/O unit can be combined to construct one or more domains.

The CMU has the following features:

- Contains an interconnect LSI module that uses the latest LSI process.
- Uses Double Data Rate (DDR)II DIMM memory.
- Supports the DR function that enables hot maintenance and replacement of CMUs during system operation, and enables addition and deletion of active CMUs during system operation.
- A redundant DDC configuration enables continuous operation even if a DDC fails.

1.3.3 I/O Unit

The I/O unit (IOU) consists of a PCIe bridge control LSI module, a printed circuit board containing a DDC, a hard disk drive (HDD), PCIe slots, and PCI cassettes for the IOU. The IOU and the CMU can be combined to configure domains.

The IOU has the following features:

- It contains eight PCIe slots.
- The IOU onboard device card (IOUA) can be used to connect the built-in hard disk drive (2.5-inch SAS interface), the built-in CD-RW/DVD-RW drive unit, and the tape drive unit. The LAN port (1000BASE-T/100Base-TX/10Base-T) mounted on the card can be used.
- It supports PCI hot-plug for External I/O Expansion Units and PCIe slots.
- Link card can be used to connect IOU to External I/O Expansion Unit.
- It supports the DR function that enables active maintenance and replacement of IOUs during system operation and enables addition and deletion of active IOUs during system operation.
- Insert the PCI card into one of the supplied cassettes before inserting it into a built-in PCIe slot in the IOU. A PCI card whose length is up to 177.8 mm (short size) can be mounted in a slot.
- A redundant DDC configuration enables continuous operation even if a DDC fails.

1.3.4 FAN Unit

The FAN unit is used to cool the server, and has the following features:

- A redundant fan configuration enables continuous system operation even if a fan fails during system operation.
- Active (hot) system maintenance or replacement of a faulty FAN unit can be performed during system operation.

1.3.5 Power Supply Unit

The power supply unit (PSU) feeds power to each unit, and has the following features:

- A redundant configuration enables the system operation to continue without interruption even if a PSU fails during system operation.
- Active (hot) system maintenance or replacement of a faulty PSU can be performed during system operation.

1.3.6 Crossbar Unit

The crossbar unit (XBU) consists of crossbar switches that logically connect CMUs and IOUs.

The XBU has redundant bus routes. If one route fails, the system can be restarted through the other route to continue operation.

1.3.7 Clock Control Unit

The clock control unit (CLKU) contains an LSI module used for the clock.

The CLKU has redundant clock supply routes. If one route fails, the system can be restarted through the other route to continue operation.

1.3.8 Operator Panel

The operator panel can be used to turn on and off the server power, switch between operation modes, and display system status information.

The operations of switches on the operator panel can be limited by switching the operation mode with the dedicated key supplied for the panel.

1.3.9 XSCF Unit

The XSCF unit (XSCFU) includes a dedicated processor, which operates independently from the main unit processors. The XSCFU in the servers adopts a duplicated configuration to increase fault tolerance.

The XSCFU is equipped with hardware interfaces for network connections to remote devices such as personal computers and workstations. A remote device can be connected via a network to the XSCF to control startup, settings, and operation management of the system.

The XSCFU provides the following hardware interfaces for network connections:

- Serial port
- LAN ports (10Base-T/100Base-T(TX))

The XSCF can be accessed through network connections using these interfaces. The command line-based interface (XSCF Shell) and browser-based user interface (XSCF Web) provided by the XSCF enable operation and management of the servers.

For details, see the *SPARC Enterprise M3000/M4000/M5000/M8000/M9000 Servers XSCF User's Guide*.

1.3.10 Internal Drive Units

The M8000/M9000 servers contain the following built-in drive units. These allow active replacement or addition.

- Hard Disk Drive

The hard disk drive is a 2.5-inch hard disk drive with a serial attached SCSI (SAS) interface. It can be mounted in an IOU.

When two hard disk drives are connected via a RAID-capable onboard device card (IOUA), those can be in the mirrored configuration by hardware RAID.

Note – When the onboard device card is RAID-capable, the `showhardconf(8)` command displays Type 2 in the output.

```
PCI#0 Name_Property:pci; Card_Type:IOUA;  
+ Serial:PP0611T826 ; Type:2;  
+ FRU-Part-Number:CA21138-B84X 010AE/371-5000-05 ;
```

- CD-RW/DVD-RW Drive Unit

There are two types of CD-RW/DVD-RW drive units: slot-loading type and tray-loading type.

FIGURE 1-11 Types of CD-RW/DVD-RW Drive Unit

Slot-loading type



Tray-loading type



Note – The locations of LED and button may vary depending on the servers.

The CD-RW/DVD-RW drive unit cannot be directly shared by multiple domains in a server. However, if the multiple domains are connected to one another through a LAN and a certain function of the Oracle Solaris OS is used, the CD-RW/DVD-RW drive unit can be shared by the domains. Adequate consideration of security is necessary for LAN connections between domains.

Tape Drive Unit

Contact your sales representative for tape drive unit options on M8000/M9000 servers.

1.4 Component Mounting Conditions

This section describes the component mounting conditions.

- CPUs can be added in units of two modules.
- DIMMs can be added in units of 16 modules.
- If you add an IOU, a CMU must be mounted for the slot with the same slot number.

- IOU onboard device card (IOUA) can be mounted in PCIe Slot #0, #2, #4, and #6 in the IOU.
- Link cards can be mounted in PCIe Slot #1, #3, #5, and #7 in the IOU.

1.5 Optional Products

The following products are the main options available for the M8000/M9000 servers.

- [Power Supply Options](#)
- [External I/O Expansion Unit](#)
- [M9000 Server \(Expansion Cabinet\) Option](#)

1.5.1 Power Supply Options

The power cabinet and the rack-mountable dual power feed option for the M8000 server are offered as power supply options.

The power cabinet enables dual power feed or three-phase power feed.

The rack-mountable dual power feed option for the M8000 server receives power from two external AC power sources that are independent of each other, and duplicates the input power system.

To use a single-phase dual power feed configuration for the M8000 server, mount the rack-mountable dual power feed option in the rack space itself. This requires a rack space with a height of 6 RUs in the cabinet. For the M9000 server, you must add the power cabinet.

For three-phase power feed in either server, an additional power cabinet is required. Install one power cabinet for each M8000/M9000 server.

For details, see the *SPARC Enterprise M8000/M9000 Servers Site Planning Guide*.

Note – A three-phase power feed option can be installed only at the factory before shipment. A single-phase power feed cannot be changed to a three-phase power feed, or vice versa, after shipment from the factory.

[TABLE 1-6](#) lists specifications of the power cabinet.

TABLE 1-6 Specifications of the Power Cabinet and M8000/M9000 Dual Power Feed Option

Item		Rack-mountable Dual Power Feed Options	Power Cabinet
Outside dimensions	Width [mm]	489	317
	Depth [mm]	1003	1244
	Height [mm]	278(6U)	1800
Weight [kg]		75	350
Input power: Single-phase power input	Voltage [V]	AC200 to 240 \pm 10%	
	Number of phases	Single phase	
	Frequency [Hz]	50/60 +2%, -4%	

Note – For specifications of the three-phase power feed option, see the *SPARC Enterprise M8000/M9000 Servers Site Planning Guide*.

1.5.2 External I/O Expansion Unit

The External I/O Expansion Unit is an optional product used to add PCI slots. The External I/O Expansion Unit has a height of four RUs (rack units), about 18 cm, in an equipment rack.


The External I/O Expansion Unit can accommodate up to two I/O boats by using either six PCIe slots or six PCI-X slots.

- PCIe slots in each I/O boat: short size to long size (to 312 mm)
- PCI-X slots in each I/O boat: short size to long size (to 312 mm)

Also, active addition and replacement is enabled for all slots in the External I/O Expansion Unit.

For details, see the *External I/O Expansion Unit Installation and Service Manual*.

FIGURE 1-12 External I/O Expansion Unit



1.5.3 M9000 Server (Expansion Cabinet) Option

An M9000 server (base cabinet) configuration can contain up to 32 CPU modules (64 cores for SPARC64 VI processors, 128 cores for SPARC64 VII/SPARC64 VII+ processors), up to 2 TB of memory, and up to 224 PCI slots. A configuration containing more components than described above would require the expansion cabinet option of the M9000 server.

A configuration with the M9000 server (expansion cabinet) can contain up to 64 CPU modules (128 cores for SPARC64 VI processors, 256 cores for SPARC64 VII/SPARC64 VII+ processors), up to 4 TB of memory, and up to 288 PCI slots.

For information about connecting the M9000 server (expansion cabinet) and the M9000 server (base cabinet), see the *SPARC Enterprise M8000/M9000 Servers Installation Guide*.

1.6 Software Features

The M8000/M9000 servers use XSCF for system administration and monitoring.

The Oracle Solaris OS can be installed as the operating environment used in a domain.

For details, see [Chapter 3](#).

System Features

This chapter explains the following technical aspects, including features and structures.

- [Section 2.1, “Hardware Configuration” on page 2-1](#)
- [Section 2.2, “Partitioning” on page 2-7](#)
- [Section 2.3, “Resource Management” on page 2-12](#)
- [Section 2.4, “RAS” on page 2-13](#)

2.1 Hardware Configuration

This section explains the hardware configuration, which includes the following items:

- [CPU](#)
- [Memory Subsystem](#)
- [I/O Subsystem](#)
- [System Bus](#)
- [System Control](#)

2.1.1 CPU

The M8000/M9000 servers use the SPARC64 VI/SPARC64 VII/SPARC64 VII+ CPU, a proprietary high-performance multi-core processor. On-chip L2 cache memory minimizes memory latency.

An instruction retry function has been implemented so that operation can be continued by retrying an instruction for which an error has been detected.

The M8000 server, M9000 server, and the M9000 server with expansion cabinet take advantage of system scalability by supporting up to 16, 32, or up to 64 CPU modules, respectively.

CPU modules running at different clock frequencies can be used in a single system. The latest CPUs can therefore be installed when improved processing performance is required.

The SPARC64 VII processor extends the 64-bit integer multiply-accumulate operation function and the hardware barrier function.

The SPARC64 VII+ processor expands the capacity of L2 cache memory to 12MB.

Note – To make maximum use of the 12MB L2 cache memory, it is necessary to use a certain type of CMU (CMU_C) and mount the CPU modules which consist entirely of the SPARC64 VII+ processors. If the CPU modules of different frequencies are mixed on CMU_C, the usable L2 cache memory is 6MB. Also, if you use other types of CMU (CMU_A or CMU_B) and mount the CPU modules which consist entirely of the SPARC64 VII+ processors, the usable L2 cache memory is to 6MB.

The type of CMU which has been mounted on the server can be confirmed by using the `showhardconf` command. For details of the `showhardconf` command, see the *SPARC Enterprise M3000/M4000/M5000/M8000/M9000 Servers XSCF Reference Manual*.


2.1.1.1 Mounted Processors and CPU Operational Modes

The M8000/M9000 servers can mount the SPARC64 VI processors, the SPARC64 VII processors, the SPARC64 VII+ processors, or a mix of those different types of processors. This section applies only to M8000/M9000 servers that run SPARC64 VII or SPARC64 VII+ processors.

Note – Supported firmware and Oracle Solaris OS will vary based on the processor type. For details, see the latest version of the Product Notes (for XCP version 1100 or later) for your server.

FIGURE 2-1 shows an example of a mixed configuration of SPARC64 VI and SPARC64 VII processors.

FIGURE 2-1 CPUs on CPU/Memory Board Unit (CMU) and Domain Configuration Example



Different types of processors can be mounted on a single CMU, as shown in CMU#2 and CMU#3 in FIGURE 2-1. And a single domain can be configured with different types of processors, as shown in Domain 2 in FIGURE 2-1.

An M8000/M9000 server domain runs in one of the following CPU operational modes:

- SPARC64 VI Compatible Mode

All processors in the domain behave like and are treated by the Oracle Solaris OS as SPARC64 VI processors. The new capabilities of SPARC64 VII or SPARC64 VII+ processors are not available in this mode. Domains 1 and 2 in FIGURE 2-1 correspond to this mode.

- SPARC64 VII Enhanced Mode

All boards in the domain must contain only SPARC64 VII or SPARC64 VII+ processors. In this mode, the server utilizes the new capabilities of these processors. Domain 0 in FIGURE 2-1 corresponds to this mode.

For the settings of the CPU operational mode, see the *SPARC Enterprise M3000/M4000/M5000/M8000/M9000 Servers XSCF User's Guide* or the *SPARC Enterprise M3000/M4000/M5000/M8000/M9000 Servers XSCF Reference Manual*.

There are restrictions on the DR operation depending on whether the Oracle Solaris OS operates in the SPARC64 VII enhanced mode or in the SPARC64 VI compatible mode. For DR operation, see the *SPARC Enterprise M4000/M5000/M8000/M9000 Servers Dynamic Reconfiguration (DR) User's Guide*.

Note – If SPARC64 VI processors are intended to be added to a domain which consists only of SPARC64 VII or SPARC64 VII+ processors, we strongly suggest setting the SPARC64 VI compatible mode in advance. Refer to the *SPARC Enterprise M3000/M4000/M5000/M8000/M9000 Servers XSCF User's Guide* or man pages for more information on the `setdomainmode` command.

2.1.2 Memory Subsystem

The memory subsystem controls memory access and cache memory. The M8000/M9000 servers use DDR-II DIMM memory.

Each CMU has thirty-two memory slots.

Also, the M8000 server, M9000 server, and M9000 server with expansion cabinet can mount up to 128, 256, or 512 DIMMs, respectively.

The memory subsystems use up to eight-way interleaving, providing higher-speed memory access.

Memory mirror mode is supported for every pair of memory buses in a CMU. This enables continued operation using the other non-defective bus if an error occurs in one bus. Memory mirror mode can be set up by the system administrator.

2.1.3 I/O Subsystem

The I/O subsystem controls data transfer between the main unit and I/O devices. The M8000/M9000 servers use PCIe as the interconnect bus for I/O devices.

Each IOU contains eight-lane (x8) PCIe slots. Also, eight-lane PCIe slots or 133-MHz 64-bit PCI-X slots can be in a mounted through an External I/O Expansion Unit.

The M8000 server, M9000 server, and the M9000 server with expansion cabinet can mount up to 32, 64, or 128 PCIe-compatible cards, respectively.

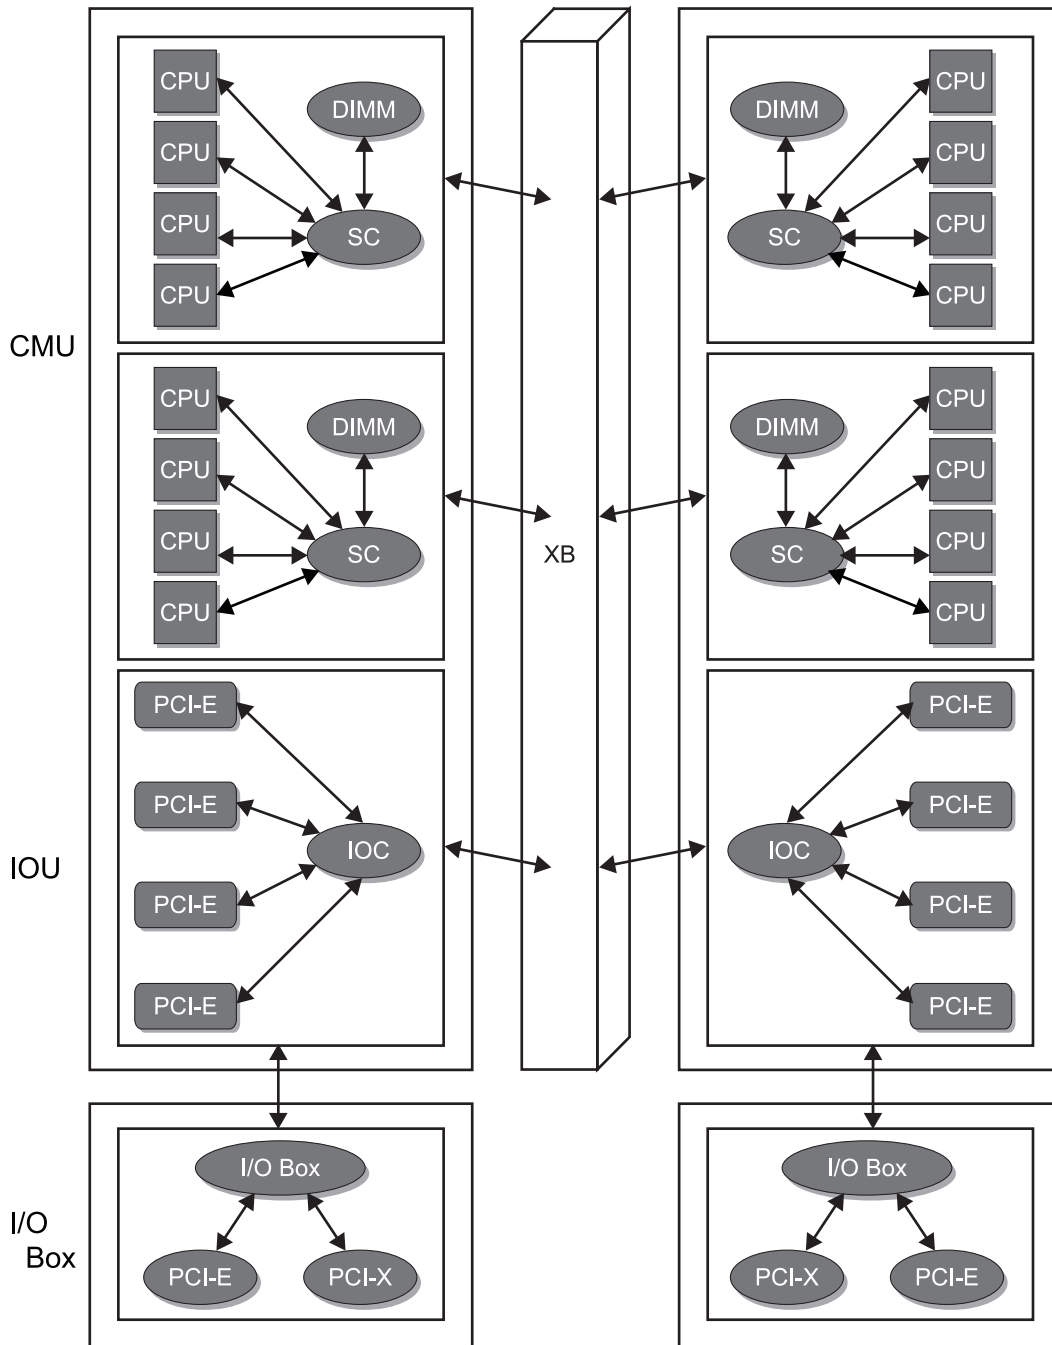
PCI Express slots or PCI-X slots can be added by mounting an External I/O Expansion Unit through a PCI Express slot.

2.1.4 System Bus

The CMU containing a CPU and memory subsystem and each component in an IOU containing an I/O subsystem are used for high-throughput data transfer between all components through a crossbar switch. The crossbar switch has duplicated bus routes. If one crossbar switch has an error, the system can be restarted to isolate the faulty switch, enabling the high-end servers to continue operation.

[FIGURE 2-2](#) shows data transfer in the system.

FIGURE 2-2 Main Component Connections



Note – The SC is the system controller that controls CPUs and memory and handles communication with the XB.

2.1.5 System Control

System control of the M8000/M9000 servers refer to the system control contained within the XSCFU that runs the XSCF and every component controlled by the XSCF.

As long as input power is being supplied to the server, the XSCF constantly monitors the server even if all domains are powered off.

The following functions are provided to increase system availability:

- Configuration management and monitoring
- Cooling unit (FAN unit) monitoring
- Domain status monitoring
- Power-on and power-off of peripheral devices

Note – To use these functions, peripheral devices with dedicated interface (RCI) and the server on which the RCI function is supported are required. For the information whether the RCI function is supported on your server, see the *SPARC Enterprise M3000/M4000/M5000/M8000/M9000 Servers Product Notes*.

- Complete control and monitoring of the server through abnormality monitoring
- Remote partitioning for domain configuration and management
- Server management and monitoring functions by the user through an external network connection
- Notifying the system administrator of fault information on the server
- Remote console input-output

2.2 Partitioning

A single M8000/M9000 server cabinet can be divided into multiple independent systems for operation. This dividing function is called partitioning.

This section describes features of partitioning and system configurations that can be implemented through partitioning.

2.2.1 Features

The individual systems resulting from partitioning can be built in the M8000/M9000 servers. These individual, divided systems are called domains. Domains are sometimes called partitions.

Partitioning enables arbitrary assignment of resources in the server. Partitioning also enables flexible domain configurations to be used according to the job load or processing amount.

An independent Oracle Solaris OS can run in a domain. Each domain is protected by hardware so that it is not affected by other domains. For example, a software-based problem, such as an OS panic, in one domain does not directly affect jobs in the other domains. Furthermore, the Oracle Solaris OS in each domain can be reset and shut down independently.

2.2.2 Domain Hardware Requirements

The basic hardware resources making up a domain are a CMU and an IOU mounted in the high-end servers or a physical system board (PSB) consisting of a CMU.

A PSB can be logically divided into one part (no division) or four parts. The physical unit configuration of each divided part of a PSB is called an extended system board (XSB).

A PSB that is logically divided into one part (no division) is called a Uni-XSB, and a PSB that is logically divided into four parts is called a Quad-XSB.


A domain can be configured with any combination of these XSBs. The XSCF is used to configure a domain and specify the PSB division type.

Note – Although a CMU with two CPUMs can be configured into Quad-XSB mode on an M8000/M9000 server, the server generates a "configuration error" message for those XSBs that do not have a CPUM and memory.

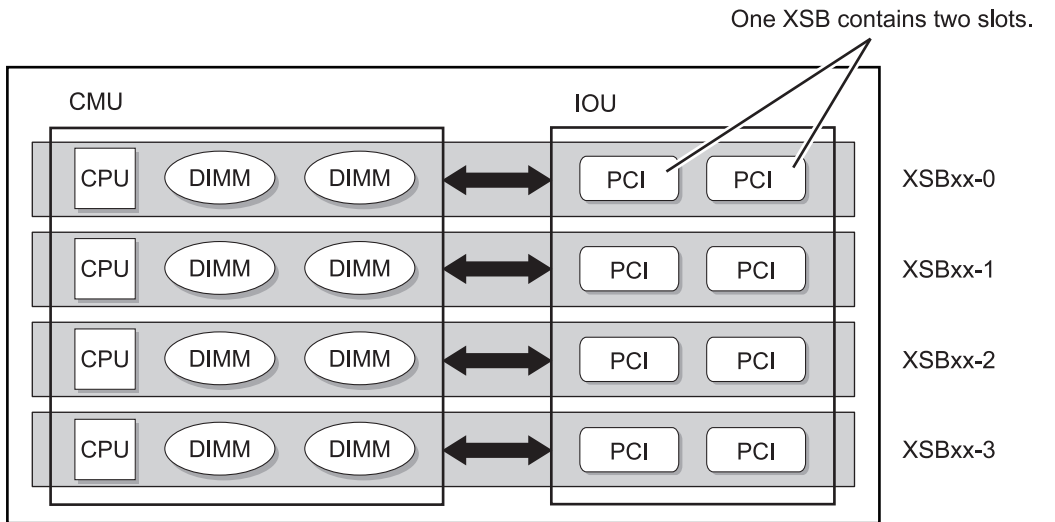
FIGURE 2-3 shows the partitioning division types.

FIGURE 2-3 Partition Division Types of Physical System Board (PSB)

- Uni-XSB types



- Quad-XSB types



2.2.3 Domain Configuration

Any XSBs in the server can be combined to configure a domain, regardless of whether the divided XSB is the Uni-XSBs or Quad-XSBs.

These XSBs can be used in any combination for a flexible domain configuration. Also, the quantity of resources for one XSB can be adjusted according to the division type of a PSB. Thus, a domain can be configured based on the quantity of resources required for job operations.

XSCF user interfaces are used to configure a domain. Each configured domain is managed by the XSCF.

The maximum number of domains that can be configured in the servers depends on the system. Up to 16 domains can be configured in M8000 servers, and up to 24 domains can be configured in M9000 servers.

To configure a domain, an LSB number must first be assigned so that a logical system board (LSB) can function as an LSB of the XSB.

This LSB number is referenced by the Oracle Solaris OS, and it must be a unique number in the domain. However, if one XSB is shared by multiple domains, a common LSB number need not be defined in the domains. An arbitrary LSB number can be assigned for this setting in each domain.

Domain configuration settings are made for each domain. A domain can be configured by specifying an XSB together with this LSB number.

Up to 16 XSBs can be configured in a single domain.

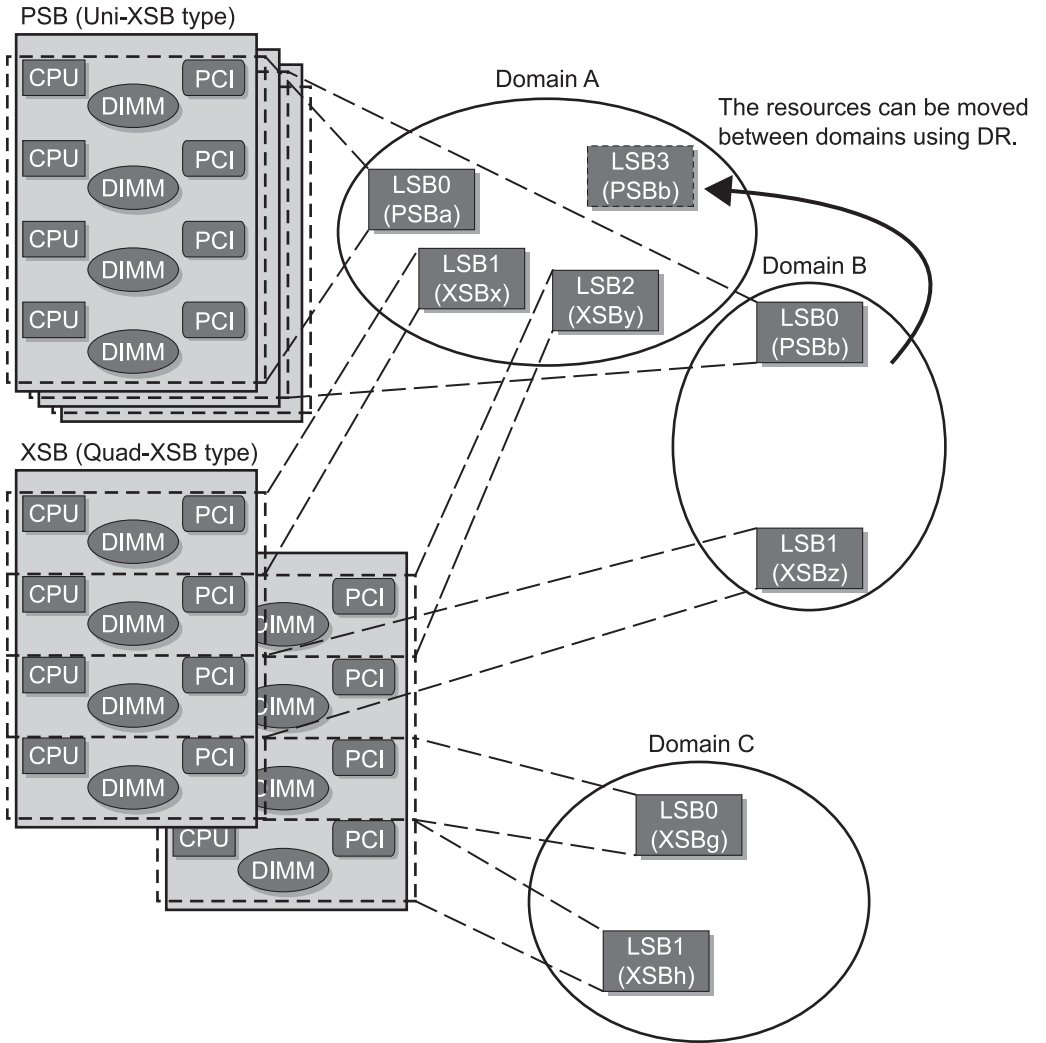
The following as well as the quantity of resources must be considered by the user who is specifying the domain configuration and division type:

- The Uni-XSB type is suitable in the configuration of a domain requiring a large quantity of resources. Also, an XSB of the Uni-XSB type is separated by the physical configuration units of a CMU and IOU. Thus, if a hardware error occurs in a CPU or memory, hardware can easily be replaced without affecting other domains. However, a resource quantity decrease due to an error may be in the range specified by the PSB.
- Quad-XSB type is suitable in the configuration of a small-scale domain, and optimized for flexible resource management. However, since domains are logically separated in a PSB, an error in the hardware shared within the PSB may affect other domains.

In addition, resources of a configured domain can be added to and deleted from individual XSBs, and they can be moved between domains by using DR function.

FIGURE 2-4 shows the domain configuration.

FIGURE 2-4 Domain Configuration



2.3 Resource Management

This section explains the following functions that support dynamic reconfiguration of domain resources during system operation:

- [Dynamic Reconfiguration](#)
- [PCI Hot-plug](#)
- [Capacity on Demand](#)
- [Oracle Solaris Zones](#)

2.3.1 Dynamic Reconfiguration

Dynamic reconfiguration (DR) enables hardware resources on system boards to be added and removed dynamically without stopping system operation. DR thus enables optimal relocation of system resources. Also, if a failure occurs, DR can place the system in a state that enables active replacement of the faulty component.

Using the DR function enables additions or distributions of resources as required for job expansions or new jobs, and it can be used for the following purposes.

- Effective use of system resources

By reserving some resources, the reserved resources can be added according to changes in the work load occurring daily, monthly, or annually. This enables flexible resource allocations on the system that needs to operate 24 hours a day, every day of the year in accordance with changes in the amount of data and the work load.

- Active replacement of system resources

If a failure occurs in a CPU for a domain that has been configured with system resources of multiple system boards, the DR function enables the faulty CPU to be isolated dynamically without stopping the system. The replacement CPU can be configured dynamically in the original domain.

For details on Dynamic Reconfiguration, see the *SPARC Enterprise M4000/M5000/M8000/M9000 Servers Dynamic Reconfiguration (DR) User's Guide*.

2.3.2 PCI Hot-plug

The PCI hot-plug function enables PCI cards to be added or removed under the Oracle Solaris OS without a system reboot.

Examples of uses for the PCI hot-plug function are as follows:

- Replacing or removing a faulty PCI card or one that will probably become faulty, during system operation
- Adding a PCI card during system operation

For details on the PCI hot-plug function, see the *SPARC Enterprise M8000/M9000 Servers Service Manual*.

2.3.3 Capacity on Demand

The Capacity on Demand (COD) feature allows you to configure spare processing resources on your server in the form of one or more COD CPUs which can be activated at a later date when additional processing power is needed. To access each COD CPU, you must purchase a COD hardware activation permit. Under certain conditions, you can use COD resources before purchasing COD permits for them.

For details on COD, see the *SPARC Enterprise M4000/M5000/M8000/M9000 Servers Capacity on Demand (COD) User's Guide*

2.3.4 Oracle Solaris Zones

The Oracle Solaris 10 OS has a function called Oracle Solaris Zones that divides the processing resources and allocates them to applications.

In a domain, resources can be divided into sections called containers, and the processing sections are allocated to each application. The processing resources are managed independently in each container. If a problem occurs in a container, the container can be isolated so that it does not affect other containers. It provides flexible resource allocation that enables optimal resource management with consideration given to the processing load.

2.4 RAS

RAS is an acronym for functions related to Reliability, Availability, and Serviceability.

RAS for M8000/M9000 servers minimize system downtime by providing for error checking at appropriate locations and by providing centralized monitoring and control of error checking.

Also M8000/M9000 servers can be configured with clustering software or centralized management software to enhance the RAS function.

Any scheduled system halt, such as a periodic maintenance or system configuration change can also be performed without affecting operating resources. This can improve service uptime significantly.

2.4.1 Reliability

Reliability represents the length of time the server can operate normally without failure.

Reliability is equally important to both hardware and software.

To improve quality, adequate components must be selected with consideration given to the product service life and the required response in case of a failure. In evaluations such as stress tests that check the service life, components and products are inspected to determine whether they meet the target reliability levels.

Furthermore, software errors are not only triggered by program errors, but also by hardware errors.

M8000/M9000 servers provide the following functions to realize high reliability.

- Monitoring by the XSCF to periodically check whether software such as the Oracle Solaris OS is running in domains (host watchdog monitoring).
- Memory patrol is periodically performed to detect memory software errors and stuck faults, even in memory areas not normally used, to prevent use of faulty memory and thereby prevent system failures caused by faulty memory from occurring.
- Since ECC protects functional data in all routes including a computing unit, a register, cache memory, and a system bus, all 1-bit errors can be automatically corrected by hardware to ensure data integrity.

2.4.2 Availability

Availability is characterized by how easily a server fails and how quickly the user can be recovered from the failure. The amount of time the system is usable is represented as a percentage.

Hardware and software faults in the system cannot be completely eliminated. To provide high availability, the system must include mechanisms that enable continuous system operation even if a failure occurs in hardware, such as components and devices, or in software, such as the OS, or application software.

M8000/M9000 servers provide the functions listed below to obtain high availability. Higher availability can also be obtained by combining the server with clustering software or management software.

- Supporting redundant configurations and active (hot) replacement of power supply units and FAN units
- Supporting redundant configuration of hard disk drive, mirroring by the hardware RAID and active/hot replacement
- Extended range of automatic correction of temporary faults in memory, system buses, and LSI internal data
- Supporting an enhanced retry function and degradation function for detected faults
- Shortening the downtime by using automatic system reboot
- Shortening the time taken for system startup
- XSCF collection of fault information, and preventive maintenance using different types of warnings
- Supporting the Chipkill function in the memory subsystem, which enables single-bit error correction to continue processing in response to continuous burst read errors caused by failures of a memory device
- Supporting the memory mirroring function enables normal data processing through the other memory bus, thereby preventing system failures in response to an error at the bus or device connected to memory bus
- Memory patrol function has no influence on the workload of software operation because it is implemented in hardware

2.4.3 Serviceability

Serviceability is characterized by how easily a server fault can be diagnosed, and how quickly the server can be recovered from the fault or how easily the fault can be corrected.

To achieve high serviceability rates, it must be possible to identify the causes of component or device failure. To facilitate recovery from failure, the system must determine the cause of the failure and isolate the faulty component for replacement. The system must also notify the system administrator and/or field engineer of the event and situation in an easy-to-understand format that prevents misunderstandings.

M8000/M9000 servers provide the following solution to realize high serviceability:

- Status LEDs mounted on the operator panel, indicating the main replaceable components and the components to which active replacement is applicable
- Remote recognition of the device operating status and remote maintenance using the XSC
- LED blinking function for indicating the maintenance target (CHECK LED, which is also called a locator)
- Notes and cautions marked on different types of labels provided for the system administrator and field engineers
- Automatic notification for reporting different types of faults to the system administrator and field engineers
- Centralized systematic monitoring, such as supporting SNMP, of a complex system from a data center

About Software

This chapter explains the following software functions.

- [Section 3.1, “Oracle Solaris OS Functions” on page 3-1](#)
- [Section 3.2, “XSCF Firmware Function” on page 3-2](#)

3.1 Oracle Solaris OS Functions

The Oracle Solaris OS has the following features:

- Reliability that has been built over the years
- Affinity that fully brings out the hardware performance of the SPARC architecture.
- A variety of products from ISVs (application software and middleware)
- Resource optimization using the partitioning and DR functions
- Dynamic addition/change of I/O device using PCI hot-plug
- Resource management with Oracle Solaris zones using Oracle Solaris container technology
- Advanced system management in cooperation with XSCF

For details of the Oracle Solaris OS, see the manual on the following URL.

<http://www.oracle.com/technetwork/documentation/index.html>

Irrespective of the capabilities of the software suite, the Oracle Solaris OS has the following functions for communicating with the SPARC Enterprise server hardware:

- [Domain Management](#)
- [PCI Hot-plug](#)

3.1.1 Domain Management

In the M8000/M9000 servers, a physical system board (PSB) can be logically divided into one part (no division) or four parts by a partitioning function unique to the system.

A PSB that is logically divided into one part (no division) is called a Uni-XSB, and a PSB that is logically divided into four parts is called a Quad-XSB.

The physical unit configuration of each divided part of a PSB is called an extended system board (XSB).

In the M8000/M9000 servers, a domain can be configured with any combination of these XSBs.

3.1.2 PCI Hot-plug

M8000/M9000 servers support insertion and removal of PCI cards for specific PCI Express and PCI-X hot-plug controllers. Before removing the PCI card, be sure to cancel the settings and isolate the card by using the Oracle Solaris OS `cfgadm(1M)` command, and make sure that the card is physically removable.

For details of PCI hot-plug, see the *SPARC Enterprise M8000/M9000 Servers Service Manual*.

3.2 XSCF Firmware Function

The XSCF firmware is a System Control Facility including a dedicated processor that is independent from the processors in the server. This section describes XSCF firmware features and provides an overview of XSCF firmware functions.

3.2.1 XSCF Features

The XSCF is firmware that is preinstalled, and operates on the XSCFU as standard equipment. As long as input power is being supplied to the server, the XSCF constantly monitors and manages the server even if the power to domains is turned off. In addition, the XSCF provides a user interface from which users can operate and manage the server.

The XSCFU has a serial port and LAN port as external interfaces. A terminal such as a personal computer or workstation can be connected to the XSCF through a serial connection or Ethernet connection.

The servers can be operated and managed from the command line-based XSCF Shell or browser-based XSCF Web provided by the XSCF.

Only the XSCF Shell can be used through the serial connection. Both the XSCF Shell and XSCF Web can be used through the Ethernet connection.

The XSCFU supports a redundant configuration (duplicate configuration) for high reliability.

The XSCF that is currently in control of the server is called the active XSCFU, and the other XSCF is called the standby XSCF or standby XSCFU since it serves as the backup XSCF to the active one.

The active XSCF and standby XSCF monitor each other to implement a failover mechanism to switch between the active XSCF or standby XSCF when one of them detects an error in the other.

For details on functions provided by the XSCF, see [Section 3.2.2, “XSCF Functional Overview” on page 3-4](#), and the *SPARC Enterprise M3000/M4000/M5000/M8000/M9000 Servers XSCF User’s Guide*.

3.2.1.1 Command Line-based User Interface (XSCF Shell)

The XSCF Shell is a command line-based user interface that can be used from a terminal, such as a personal computer or workstation, connected to the XSCF through a serial connection or Ethernet connection.

With the serial connection, the terminal is connected directly to the server to use shell commands provided by the XSCF. Also, the console redirection function of the XSCF enables the terminal to be used as an OS console.

With the Ethernet connection, the terminal is connected to the XSCF via Secure Shell (SSH) or telnet to use shell commands provided by the XSCF.

The following are the main operations that can be performed with the XSCF Shell:

- Display the server configuration or status and various related settings
- Display the domain configuration or status and various related settings
- Start or shut down a domain
- Make settings for various network services
- Make settings for various security functions
- Make various settings for the remote maintenance service function

3.2.1.2 Browser-Based User Interface (XSCF Web)

XSCF Web is a browser-based user interface that can be used from a terminal, such as a personal computer or workstation, connected to the server through an Ethernet connection.

However, the XSCF Web cannot be used through a serial connection. If the browser function of a terminal is used for a connection to the XSCF, BUI operations can be performed.

3.2.2 XSCF Functional Overview

This section provides an overview of the main functions supported by the XSCF.

For details on each function, see the *SPARC Enterprise M3000/M4000/M5000/M8000/M9000 Servers XSCF User's Guide*.

3.2.2.1 System Management

The main tasks of the XSCF are control and monitoring of the entire system, such as configuration management and monitoring of the server's cooling components (FAN units), domain status monitoring, power-on and power-off of peripheral units, and error monitoring. Also, the XSCF provides a partitioning function for domain configuration and management.

The XSCF constantly monitors the server status so that the system can operate stably.

When any error is detected in the system, the system status monitoring mechanism immediately collects error information about the hardware (hardware log) and analyzes it to identify the error location and check the error status. The XSCF displays the status, and it degrades the relevant component or domain as necessary or it resets the system, thereby preventing the error from occurring again.

The XSCF provides high reliability, high availability, and high serviceability for the entire system.

3.2.2.2 Security Management

The XSCF manages user accounts for the XSCF. The operating range in the XSCF Shell and XSCF Web can be limited based on user account types and settings. Moreover, the XSCF provides an IP address filtering function for permitting access to the XSCF and an encryption function using SSH and SSL. Operator errors and unauthorized access during system operation are recorded in a log. The system administrator can use these functions to investigate the cause of a system problem.

3.2.2.3 System Status Management

The XSCF provides functions as XSCF operations for displaying the system configuration status, creating and changing a domain configuration definition, and starting and stopping domains. Also, the XSCF provides the DR function, which assists in dynamically changing a system board configuration during domain operation. Thus, domain resources can be optimized for job operations. Moreover, the XSCF in linkage with the Oracle Solaris OS manages CPU, memory, and I/O resources.

3.2.2.4 Error Detection and Management

The XSCF constantly monitors the system status so that the system can operate stably. When any error is detected in the system, the XSCF immediately collects error information about the hardware (hardware log) and analyzes it to identify the error location. To continue operation, the XSCF degrades the relevant component or domain as necessary according to error conditions or it resets the system, thereby preventing the problem from occurring again. The user can take prompt action for problems since easy-to-understand and accurate information on hardware errors and fault locations are thus provided to the user.

3.2.2.5 Remote System Control and Monitoring

The XSCF provides functions for monitoring the server through an Ethernet connection so that the user can remotely manage the server. Moreover, a function for reporting error information to the system administrator and a remote console input-output function are supported. System availability is thus increased.

3.2.2.6 Resource Management

The XSCF manages hardware resources on configured domains and system boards. Resource management provides Dynamic Reconfiguration (DR) and Capacity on Demand (COD) functions.

Dynamic Reconfiguration (DR)

DR enables users to add, remove or exchange system boards while the domains that contain these boards remain up and running. It also enables dynamic reconfiguration of domains.

For details on DR, see the *SPARC Enterprise M4000/M5000/M8000/M9000 Servers Dynamic Reconfiguration (DR) User's Guide*.

Capacity on Demand (COD)

The COD feature allows you to configure spare processing resources on your server in the form of one or more COD CPUs which can be activated at a later date when additional processing power is needed.

For details, see the *SPARC Enterprise M4000/M5000/M8000/M9000 Servers Capacity on Demand (COD) User's Guide*

3.2.2.7 Airflow Indicator

The airflow indicator indicates the amount of air exhausted from the server while the M8000/M9000 servers are up and running. The values do not include the peripheral devices.

To display the amount of exhaust air, use the `showenvironment air` command.

```
XSCF> showenvironment air
Air Flow:5810CMH
```

Note – The `showenvironment air` command displays the calculated airflow based on the fan speed such as Low speed or High speed etc. The fan speed is displayed by the `showenvironment Fan` command.

For details of the `showenvironment(8)` command, refer to the man page. For installation details of the SPARC Enterprise M8000/M9000 servers from Oracle and Fujitsu, see the *SPARC Enterprise M8000/M9000 Servers Site Planning Guide* and the *SPARC Enterprise M8000/M9000 Servers Installation Guide*.

You can also obtain the exhaust air data using the SNMP agent function. To obtain the data of exhaust air using the SNMP agent function, install the latest XSCF extension MIB definition file to the SNMP manager. For details on the XSCF extension MIB definition file, see the *SPARC Enterprise M3000/M4000/M5000/M8000/M9000 Servers XSCF User's Guide*.

Index

A

- airflow indicator, 3-6
- appearance
 - M8000, 1-2
 - M9000 (base cabinet only), 1-3
 - M9000 (with an expansion cabinet), 1-4
- availability, 2-14

C

- components
 - CPU module, 1-22
 - CPU/memory unit, 1-23

F

- features
 - high-end server, 1-5
 - partitioning, 2-8
 - software, 3-1
 - system, 2-1
 - XSCF, 3-2

H

- hardware configuration
 - CPU, 2-1
 - I/O subsystem, 2-5
 - memory subsystem, 2-5
 - system bus, 2-5
 - system control, 2-7
- hardware RAID, 1-6, 1-25, 2-15

N

- names of components

- M8000 front view, 1-14
- M8000 rear view, 1-14
- M9000 (base cabinet only) front view, 1-16
- M9000 (base cabinet only) rear view, 1-16
- M9000 (with an expansion cabinet) front view, 1-18
- M9000 (with an expansion cabinet) rear view, 1-18

O

- operator panel
 - LED, 1-19
 - switch, 1-21
- optional products, 1-27
 - external I/O expansion unit, 1-28
 - M9000 server (expansion cabinet) option, 1-29

P

- physical unit, 2-8
- power cabinet, 1-14, 1-16, 1-18

Q

- Quad-XSB, 2-8

R

- reliability, 2-14
- resource management
 - Capacity on Demand (COD), 2-13
 - Oracle Solaris Zone, 2-13
 - PCI hot-plug, 2-13

S

serviceability, 2-15

software

- Oracle Solaris operating system, 3-1

- XSCF firmware, 3-2

specifications

- system, 1-9

system

- specifications, 1-9

U

Uni-XSB, 2-8

X

XSCF firmware, 3-2

XSCF shell, 3-3

XSCF web, 3-4