

Fujitsu SPARC M12 and Fujitsu M10 Server Architecture

White Paper



Manual Code: C120-E690-14EN
June 2022

Copyright © 2007, 2022, 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 of The Open Group.

Oracle and Java are registered trademarks of Oracle and/or its affiliates.

Fujitsu and the Fujitsu logo are registered trademarks of Fujitsu Limited.

SPARC Enterprise, SPARC64, SPARC64 logo and all SPARC trademarks are trademarks or registered trademarks of SPARC International, Inc. in the United States and other countries and used under license.

Other names may be trademarks of their respective owners.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

Disclaimer: The only warranties granted by Oracle and Fujitsu Limited, and/or any affiliate 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 NONINFRINGEMENT, 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.

Fujitsu SPARCM12 is sold as SPARCM12 by Fujitsu in Japan.

Fujitsu SPARCM12 and SPARCM12 are identical products.

Fujitsu M10 is sold as SPARCM10 by Fujitsu in Japan.

Fujitsu M10 and SPARCM10 are identical products.

Copyright © 2007, 2022, Fujitsu Limited. Tous droits réservés.

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

Oracle et/ou ses affiliés et Fujitsu Limited détiennent et contrôlent chacun 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 droit d'auteur, des brevets, et 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 affiliés et de Fujitsu Limited, et de leurs éventuels concédants de licence. Ce document, bien qu'il vous ait été fourni, ne vous confère aucun droit et aucune licence, exprès 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 le droit d'auteur 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 affiliés ou Fujitsu Limited. Cette distribution peut comprendre des composants développés par des parties tierces. Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie.

UNIX est une marque déposée de The Open Group.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés.

Fujitsu et le logo Fujitsu sont des marques déposées de Fujitsu Limited.

SPARC Enterprise, SPARC64, le logo SPARC64 et 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.

Tout autre nom mentionné peut correspondre à des marques appartenant à leurs propriétaires respectifs.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des États-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des États-Unis, la notice suivante s'applique :

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

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 EXPRESSEMENT STIPULÉE AU DIT CONTRAT, ORACLE OU FUJITSU LIMITED ET/OU

LES SOCIETES AFFILIEES A L'UNE OU L'AUTRE ENTITE DECLINENT TOUT ENGAGEMENT OU GARANTIE, QUELLE QU'EN SOIT LA NATURE (EXPRESSE OU IMPLICITE) CONCERNANT CE PRODUIT, CETTE TECHNOLOGIE OU CE DOCUMENT, LESQUELS SONT FOURNIS EN L'ETAT. EN OUTRE, TOUTES LES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFACON, SONT EXCLUES, DANS LA MESURE AUTORISEE 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'ETAT" ET TOUTE AUTRE CONDITION, DECLARATION ET GARANTIE, EXPRESSE OU TACITE, EST FORMELLEMENT EXCLUE, DANS LA MESURE AUTORISEE PAR LA LOI EN VIGUEUR, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFACON.

Au Japon, Fujitsu SPARCM12 est vendu en tant que SPARCM12 par Fujitsu.

Fujitsu SPARCM12 et SPARC M12 sont des produits identiques.

Au Japon, Fujitsu M10 est vendu en tant que SPARCM10 par Fujitsu.

Fujitsu M10 et SPARCM10 sont des produits identiques.

Contents

Contents	vii
Introduction	1
Fujitsu SPARC M12 and Fujitsu M10 Providing Both High Performance and High Reliability	1
1. Fujitsu SPARC M12	4
Product Lineup	4
Fujitsu SPARC M12-1	6
Fujitsu SPARC M12-2	7
Fujitsu SPARC M12-2S	8
2. Fujitsu M10	9
Product Lineup	9
Fujitsu M10-1	11
Fujitsu M10-4	12
Fujitsu M10-4S	13
3. SPARC64 XII, SPARC64 X+, and SPARC64 X Processors	15
4. System Architecture	17
System Interconnect	17
1. System Bus	17
2. Fujitsu SPARC M12-1 Interconnect Architecture	19
3. Fujitsu SPARC M12-2 Interconnect Architecture	20
4. Fujitsu SPARC M12-2S Interconnect Architecture	21
5. Fujitsu M10-1 Interconnect Architecture	24
6. Fujitsu M10-4 Interconnect Architecture	25
7. Fujitsu M10-4S Interconnect Architecture	26
8. System Interconnect Reliability Features	29
Memory	30
1. Memory Mirroring	31
System Clock	32
I/O Bus	32
1. I/O Subsystem Architecture	32
2. Fujitsu SPARC M12-1 I/O Subsystem	33
3. Fujitsu SPARC M12-2 I/O Subsystem	34
4. Fujitsu SPARC M12-2S I/O Subsystem	35
5. Fujitsu M10-1 I/O Subsystem	36
6. Fujitsu M10-4 I/O Subsystem	37

7. Fujitsu M10-4S I/O Subsystem	38
8. Internal Drives and Peripherals	39
9. PCI Expansion Unit	39
Cooling	41
High-Efficiency Power Supplies	43
5. System Management	44
Reliability, Availability, and Serviceability	44
1. Redundant and Hot-Swappable Components	45
2. Partitioning Feature	46
3. Advanced Reliability Features	47
4. Error Detection, Diagnosis, and Recovery	47
System Management	48
1. eXtended System Control Facility	48
2. Redundant XSCF	49
3. XSCF Control Package	49
4. Role-based System Management	50
5. Oracle Enterprise Manager Ops Center 12c	50
Eco-Friendly Computing	50
6. Oracle Solaris 11 Operating System	52
Oracle Solaris ZFS	52
Oracle VM Server for SPARC	55
Oracle Solaris Zones	56
Image Packaging System	58
Boot Environment	58
7. Technological Enhancements of the SPARC64 X2, SPARC64 X+, and SPARC64 X Processors	59
Microarchitecture	59
1. Chip Configuration	59
2. Core Microarchitecture	60
3. Interface Between Chips	66
Extended Instruction Set Architecture	66
1. HPC-ACE	67
2. SWoC (Software on Chip)	68
Reliability, Availability, and Serviceability Features	69
1. Error Marking	72
2. Internal RAM Reliability and Availability Features	72
3. Internal Register and Execution Unit Reliability Features	73
4. Synchronous Update Method and Instruction Retry	73
5. Increased Serviceability	74
8. Conclusion	75

Introduction

Fujitsu SPARC M12 and Fujitsu M10 Providing Both High Performance and High Reliability

To respond to rapidly changing business environments, IT infrastructures must provide higher speed, reliability, and operational efficiency as well as optimized performance. The Fujitsu SPARC M12 and Fujitsu M10 provide the total platform required in modern IT infrastructures with leading technologies including a high-performance processor, cutting-edge semiconductor advancements, and the high reliability and high quality inherited from Fujitsu's mainframes and SPARC Enterprise M-series. The Fujitsu SPARC M12 and Fujitsu M10 have further enhanced system reliability, flexibility, and scalability to meet the demands of cloud computing environments.

Improved System Performance With the SPARC64™ XII, SPARC64™ X+, and SPARC64™ X

Fujitsu has developed the new SPARC64 XII (twelve) processor, which inherits the high performance and high reliability of the SPARC64 X+ (ten plus) processor and SPARC64 X (ten) processor.

The SPARC64 XII processor mounted in the Fujitsu SPARC M12 provides high core performance, which accelerates the database/DWH processing of customers' mission-critical systems.

The Fujitsu M10 is offered with two kinds of processors, SPARC64 X+ and SPARC64 X. With these processors, customers are able to select the most appropriate engine depending on the workload required.

- System on chip

With the SPARC64 XII, SPARC64 X+, and SPARC64 X processors having been developed for UNIX servers, peripheral LSIs are consolidated into the processors.

- Software on chip

Many processing functions that were traditionally handled by software have been built into the processor hardware, by adding multiple, dedicated instructions. For handling big data

analytics and analysis processing, the number of concurrent processes has increased.

- Cooling technology

Fujitsu's proprietary cooling technologies get the maximum performance out of the processor (Liquid Loop Cooling (LLC) and Vapor and Liquid Loop Cooling (VLLC), a newly introduced cooling technology).

These technologies achieve significant improvements in processing speed (throughput) and overall performance.

Fujitsu SPARC M12 and Fujitsu M10: Most Appropriate for Mission-Critical Systems and Real-Time Information Analysis

The processors are interconnected with the Fujitsu SPARC M12 or Fujitsu M10 using a cutting-edge fast interconnect technology. Moreover, the Fujitsu SPARC M12-2S and Fujitsu M10-4S adopt the Building Block (BB) expansion method (*1), which can significantly enhance performance and functionality. A maximum of 16 servers (*2) can be interconnected to build a single, large system. Up to 16 Fujitsu SPARC M12-2S chassis with up to 32 CPUs/3072 threads (up to 16 Fujitsu M10-4S chassis with up to 64 CPUs/2048 threads) deliver the highest performance in an extremely flexible and scalable system.

From a small-scale configuration at installation, which reduces the initial investment, the system can expand to meet increases in demand, such as to integrate an information analysis system with a database system to run real-time analysis.

These features make the Fujitsu SPARC M12 and Fujitsu M10 the most appropriate servers for datacenters in the cloud computing era.

*1 The Building Block method allows common servers (blocks) containing CPUs, memory, and I/O expansion slots to be connected like children's toy blocks being stacked up.

*2 For the Fujitsu SPARC M12-2S that began selling in April 2022, the maximum is 4 servers.

Fujitsu SPARC M12 and Fujitsu M10: Better Installation

The Fujitsu SPARC M12 and Fujitsu M10 also benefit from increased ease of installation. With peripheral LSIs consolidated into their SPARC64 XII, SPARC64 X+, and SPARC64 X processors, the Fujitsu SPARC M12 and Fujitsu M10 are very densely packaged servers delivering substantial space savings.

The Fujitsu SPARC M12 and Fujitsu M10 housing the SPARC64 XII, SPARC64 X+, and

SPARC64 X processors merge numerous hardware and software technologies to provide customers with the most appropriate solution for their ever-growing IT infrastructures.

1.Fujitsu SPARC M12

Product Lineup

The Fujitsu SPARC M12 is the optimal product for a broad range of needs, like for a database or analytics infrastructure.

- Fujitsu SPARC M12-1: An entry-level server that offers high performance and high reliability in a space-saving 1U chassis
- Fujitsu SPARC M12-2: A mid-range server perfect for system consolidation, such as database or analytics consolidation
- Fujitsu SPARC M12-2S: Capable of covering a broad range of needs from mid-range to high-end




Figure 1-1. Fujitsu SPARC M12 Product Lineup

Table 1-1. Fujitsu SPARC M12 Specifications

		Fujitsu SPARC M12-1	Fujitsu SPARC M12-2	Fujitsu SPARC M12-2S (1BB)	Fujitsu SPARC M12-2S (16BB) (*1)
SPARC64 XII Processor	Frequency	3.2 GHz	3.9 GHz	4.25 GHz	4.25 GHz
	L3 cache	16 MB	32 MB	32 MB	32 MB
	CPUs	1	1 or 2	1 or 2	Up to 32
	Cores	6	12 (per processor)	12 (per processor)	Up to 384
	Threads	48	96 (per processor)	96 (per processor)	Up to 3072
Memory	Maximum size	1 TB	3 TB	3 TB	48 TB
	Mounted memory		(24 memory slots/processor)	(24 memory slots/processor)	(24 memory slots/processor)
			2 TB	2 TB	32 TB
			(16 memory slots/processor)	(16 memory slots/processor)	(16 memory slots/processor)
	modules, maximum	16	48	48	768
Internal Storage			(24 memory slots/processor)	(24 memory slots/processor)	(24 memory slots/processor)
			32	32	512
			(16 memory slots/processor)	(16 memory slots/processor)	(16 memory slots/processor)
PCIe Slot	Interface	SAS	SAS	SAS	SAS
	Built-in disks	8	8	8	128
	Interface	PCI Express 3.0	PCI Express 3.0	PCI Express 3.0	PCI Express 3.0
	Built-in PCIe slots	3	11	8	128
	Built-in PCIe slots (maximum)	33 (*2) With 3 PCI Expansion Units	91 (*2) With 8 PCI Expansion Units	88 (*2) With 8 PCI Expansion Units	1,408 (*2) With 128 PCI Expansion Units
I/O Interface	10 Gb Ethernet ports	4	4	4	64
	SAS ports	1	1	1	16
	USB ports	2	2	2	32
Rack Units		1U	4U	4U	40U x 2 (including XB-Boxes)
Virtualization Function	Physical partitions	1	1	1 (1 partition per BB)	Up to 16
	Logical domains, maximum	48	192	192	256 per physical partition

*1 Up to the 4BB configuration, for the Fujitsu SPARC M12-2S that began selling in April 2022. The resulting maximum value of each item in the specifications becomes 1/4 of that for the 16BB configuration.

*2 When the number of mounted CPUs has reached the maximum and the OS versions on the control domain, root domain, and I/O domain are all Oracle Solaris 11

Fujitsu SPARC M12-1



The Fujitsu SPARC M12-1 is an entry-level server that offers high performance and high reliability in a space-saving chassis measuring one rack unit (1U).

The server basically comes with a single high-performance CPU with up to six cores, and supports up to 1 TB of memory and up to eight internal storage slots (supporting hard disk drives and SSDs). The server includes three PCI Express slots. For further expansion, PCI Expansion Units can be added to provide up to 33 internal I/O slots. Thus, the server provides mid-range class scalability. The server has a redundant configuration with two built-in power supplies and seven built-in fan units.

A system can be purchased with the minimum required CPU resources. Then, as needed, the CPU resources can be expanded in a step-by-step manner with the purchase of CPU Activations (CPU Rights to Use). In addition, server virtualization and system consolidation are realized through both Oracle Solaris Zones and Oracle VM Server for SPARC. The Fujitsu SPARC M12-1 can build on the initial investment reductions and step-by-step processing expansion, so TCO can be optimized even by this entry-level server.

Fujitsu SPARC M12-2



The Fujitsu SPARC M12-2 is a high-performance and highly reliable mid-range server that is appropriate for database integration and consolidation of databases and analytics. Its mainframe-class reliability is suitable for customers' mission-critical business.

The Fujitsu SPARC M12-2 occupies 4 rack units (4U) and supports up to 2 CPUs (24 cores), up to 3 TB (24 memory slots per processor) of memory, and up to 8 internal storage slots supporting hard disk drives and SSDs). The Fujitsu SPARC M12-2 also includes 11 PCI Express slots. To support further I/O expansion, PCI Expansion Units can be connected to provide up to 91 PCI Express slots. Four power supply units (two per group) provide redundant power to the server. The Vapor and Liquid Loop Cooling (VLLC) system, a redundant high-efficiency cooling technology developed by Fujitsu, and eight fan units cool the server.

A system can be purchased with the minimum required CPU resources. Then, as needed, the CPU resources can be expanded in a step-by-step manner with the purchase of CPU Activations (CPU Rights to Use). In addition, both Oracle Solaris Zones and Oracle VM Server for SPARC allow flexible changes to a virtualized configuration according to variation requests for the system. Processing expansion is realized in the Fujitsu SPARC M12-2 by expanding CPU capacity from one CPU to two CPUs.

Fujitsu SPARC M12-2S



The Fujitsu SPARC M12-2S provides the world's highest level scalability and flexibility, covering a broad range of needs, from mid-range to high-end. This model is also capable of large-scale consolidation including mission-critical databases.

In addition to high performance and high reliability, the Fujitsu SPARC M12-2S provides flexible scalability by virtue of the Building Block expansion method. With an expansion rack installed, a customer can freely combine scale-up configurations with scale-out configurations best suited for distributed parallel processing while continuing business operations. With 2 expansion racks, the customer can scale up to 16 interconnected Building Blocks without shutting down the entire system.

Fujitsu high-speed interconnect technology provides the connection between Fujitsu SPARC M12-2S Building Block chassis and achieves linear performance improvement up to the maximum 16BB configuration. A system can be purchased with the minimum required CPU, memory, PCI Express slot, and disk bay resources. Then, as needed, the customer can increase capacity by adding more Fujitsu SPARC M12-2S chassis. The CPU resources can also be expanded in a step-by-step manner with the purchase of CPU Activations (CPU Rights to Use) to achieve finer granularity expansion than by adding whole CPU chips.

2. Fujitsu M10

Product Lineup

The Fujitsu M10 covers a product lineup that is applicable to a broad range of industries.

- Fujitsu M10-1: An entry-level server in a space-saving 1U chassis
- Fujitsu M10-4: A mid-range server perfect for system consolidation
- Fujitsu M10-4S: Capable of covering a broad range of needs, from mid-range to high-end




Figure 2-1. Fujitsu M10 Product Lineup

Table 2-1. Fujitsu M10 Specifications

		Fujitsu M10-1	Fujitsu M10-4	Fujitsu M10-4S (1BB)	Fujitsu M10-4S (16BB)
SPARC64 X+ Processor (*1)	Frequency	2.8 GHz/3.2 GHz/3.7 GHz	3.4 GHz/3.7 GHz	3.7 GHz	3.7 GHz
	L2 cache	22 MB/22 MB/24 MB	24 MB	24 MB	24 MB
	CPUs	1	2 or 4	2 or 4	Up to 64
	Cores	16/16/8	16/8 (per processor)	16 (per processor)	Up to 1024
	Threads	32/32/16	32/16 (per processor)	32 (per processor)	Up to 2048
SPARC64 X Processor (*1)	Frequency	2.8 GHz	2.8 GHz	3.0 GHz	3.0 GHz
	L2 cache	22 MB	24 MB	24 MB	24 MB
	CPUs	1	2 or 4	2 or 4	Up to 64
	Cores	16	16 (per processor)	16 (per processor)	Up to 1024
	Threads	32	32 (per processor)	32 (per processor)	Up to 2048
Memory	Maximum size	1 TB	4 TB	4 TB	64 TB
	Mounted memory modules, maximum	16	64	64	1024
Internal Storage	Interface	SAS	SAS	SAS	SAS
	Built-in disks	8	8	8	128
PCIe Slot	Interface	PCI Express 3.0	PCI Express 3.0	PCI Express 3.0	PCI Express 3.0
	Built-in PCIe slots	3	11	8	128
	Built-in PCIe slots (maximum)	23 With 2 PCI Expansion Units	71 (*2) With 6 PCI Expansion Units	58 (*2) With 5 PCI Expansion Units	928 (*2) With 80 PCI Expansion Units
I/O Interface	1 Gb Ethernet ports	4	4	4	64
	SAS ports	1	1	1	16
	USB ports	2	2	2	32
Rack Units		1U	4U	4U	40U x 2 (including XB-Boxes)
Virtualization Function	Physical partitions	1	1	1 (1 per BB)	Up to 16
	Logical domains, maximum	32	128	128	256 per physical partition

*1 Either the SPARC64 X+ or SPARC64 X processor can be mounted. A SPARC64 X Fujitsu M10-4S Building Block can be connected with a SPARC64 X+ Fujitsu M10-4S Building Block.

*2 When the number of mounted CPUs has reached the maximum

Fujitsu M10-1



The Fujitsu M10-1 is an entry-level server, in a space-saving chassis measuring one rack unit (1U), providing excellent performance and reliability appropriate for datacenter consolidation and virtualization.

The server basically comes with a single high-performance CPU with up to 16 cores, and supports up to 1 TB of memory and up to 8 internal storage slots (supporting hard disk drives and SSDs). The Fujitsu SPARC M12-2 also includes 3 PCI Express slots. To support further I/O expansion, PCI Expansion Units can be connected to provide up to 23 PCI Express slots. Two power supplies and seven fan units power and cool the server with built-in redundancy.

A system can be purchased with the minimum required CPU resources. Then, as needed, the CPU resources can be expanded in a step-by-step manner with the purchase of CPU Activations (CPU Rights to Use). In addition, server virtualization and system consolidation are realized through both Oracle Solaris Zones and Oracle VM Server for SPARC. The Fujitsu M10-1 can build on the initial investment reductions and step-by-step processing expansion, so TCO can be optimized even by this entry-level server.

Fujitsu M10-4



The Fujitsu M10-4 is a high-performance and highly reliable mid-range server that is appropriate for datacenter consolidation and virtualization tasks requiring more processors, memory, and I/O capacity than are available in the Fujitsu M10-1 model. It provides the flexibility, scalability, and reliability required to support customers' mission-critical businesses.

The Fujitsu SPARC M12-2 occupies 4 rack units (4U) and supports up to 4 CPUs (up to 64 cores total), up to 4 TB of memory, and up to 8 internal storage slots (supporting hard disk drives and SSDs). The Fujitsu M10-4 also includes 11 PCI Express slots. To support further I/O expansion, PCI Expansion Units can be connected to provide up to 71 PCI Express slots. Two power supply units provide redundant power to the server. The Liquid Loop Cooling (LLC) system and five fan units cool the server.

A system can be purchased with the minimum required CPU resources. Then, as needed, the CPU resources can be expanded in a step-by-step manner with the purchase of CPU Activations (CPU Rights to Use). In addition, server virtualization and system consolidation are realized through both Oracle Solaris Zones and Oracle VM Server for SPARC. The Fujitsu M10-4 builds on the initial investment reductions and step-by-step processing expansion found in the Fujitsu M10-1. The server expands CPU capacity from two CPUs to up to four CPUs, offering the maximum utilization of memory and I/O.

Fujitsu M10-4S



The Fujitsu M10-4S provides the world's highest scalability and flexibility, and covers a broad range of computing needs, from mid-range to high-end. The Fujitsu M10-4S is the most appropriate for use in cloud computing and big data processing infrastructures that are currently large scale or have the potential to grow significantly over time.

In addition to the high performance and high reliability found in the Fujitsu M10 model, the Fujitsu M10-4S provides flexible scalability by virtue of the Building Block expansion method. With an expansion rack installed, a customer can freely combine scale-up configurations with scale-out configurations best suited for distributed parallel processing without interrupting working partitions. With 2 expansion racks, the customer can scale up to 16 interconnected Building Blocks without shutting down the entire system.

The proprietary Fujitsu interconnect technology that provides the connection between Building Blocks achieves linear performance improvement up to the maximum 16BB configuration. A system can be purchased with the minimum required CPU, memory, PCI Express slot, and disk bay resources. Then, as needed, the customer can increase capacity by adding more Fujitsu M10-4S chassis. The CPU resources can also be expanded in a step-by-step manner with the purchase of CPU Activations (CPU Rights to Use) to achieve finer granularity expansion than

by adding whole CPU chips.

Furthermore, additional SPARC64 X Fujitsu M10-4S Building Blocks can be connected to not only SPARC64 X Fujitsu M10-4S Building Blocks but also SPARC64 X+ Fujitsu M10-4S Building Blocks.

Even a system that already uses a Fujitsu M10-4S with a mounted SPARC64 X+ instead of a Fujitsu M10-4S with a mounted SPARC64 X can be upgraded to a higher-performance system without wasting the Fujitsu M10-4S with the mounted SPARC64 X.

3. SPARC64 XII, SPARC64 X+, and SPARC64 X Processors

The SPARC64 XII is the latest of the processors developed by Fujitsu to combine high performance and high reliability for UNIX servers. The processors feature high core performance, high operating frequencies, multi-core and multithreading features, and high memory throughput. Furthermore, the SPARC64 XII can connect up to 32 CPU chips in a single system, and the SPARC64 X+ or SPARC64 X can connect up to 64 CPU chips, offering extremely high performance and scalability. Fujitsu's heritage of mainframe-class high reliability technology is found throughout the processors and corresponding systems.



The SPARC64 XII, SPARC64 X+, and SPARC64 X inherit the robust high-reliability technologies found in previous SPARC64 processors and, at the same time, offer multiple functional enhancements. The enhancements extend beyond microarchitecture improvements, such as significantly higher core counts and the consolidation of peripheral LSI functions into the processor, to include register extensions and significant instruction set enhancements that improve SIMD instruction-based cryptographic and database processing performance.

The SPARC64 XII, SPARC64 X+, and SPARC64 X consolidate peripheral LSI functions into the processor, leading directly to faster processing and improvement in performance per Watt. Compared to the SPARC64 VII+, the SPARC64 X delivers approximately twice the performance per CPU core, and a 7.5-times improvement in per-processor performance. The SPARC64 X+ achieves a further improvement of 2.4 times the performance per CPU core and 9.5 times the performance per processor. Compared to the SPARC64 X+, the SPARC64 XII provides approximately 2.5 times the performance per CPU core and approximately twice the performance per processor (as measured with SPECint_rate).

The significant technological enhancements of the SPARC64 XII, SPARC64 X+, and SPARC64 X are explained in detail in "Chapter 7. Technological Enhancements of the SPARC64 XII, SPARC64 X+, and SPARC64 X Processors."




Figure 3-1. SPARC64 XII, SPARC64 Xand SPARC64 VII+ Functional Enhancement

4. System Architecture

In the Fujitsu SPARC M12 and Fujitsu M10, the processors, system interconnects, and memory and I/O subsystems work in concert to create a scalable, high-performance platform ready to address a wide range of workloads, from consolidation of general-purpose enterprise computing to the fastest, largest, and most secure database processing.

The design of the Fujitsu SPARC M12 and Fujitsu M10 focuses on high reliability, and places emphasis on maximizing the merits of memory locality in a ccNUMA architecture to deliver outstanding performance. The characteristics and capabilities of every subsystem within the Fujitsu SPARC M12 and Fujitsu M10 work toward this goal. A high-bandwidth system bus, the powerful SPARC64 XII, SPARC64 X+, and SPARC64 X processors, dense memory support, and fast PCI Express combine within the Fujitsu SPARC M12 and Fujitsu M10 to deliver the highest levels of uptime and throughput, as well as dependable scaling for enterprise applications.

System Interconnect

The system interconnect underpins the highest levels of performance, scalability, and reliability for the Fujitsu SPARC M12 and Fujitsu M10. Multiple system controllers and crossbar units within the Fujitsu SPARC M12 and Fujitsu M10 provide point-to-point connections between the CPUs, memory, and I/O subsystems. Offering more than one bus route between components enhances performance and allows system operation to continue in the event of a fault.

1. System Bus

High-end systems containing dozens of CPUs provide scalability only if all processors are able to actually contribute to the performance of the application. The ability to deliver near-linear scalability and fast, predictable performance for a broad set of applications rests largely on the capabilities of the system bus. The Fujitsu SPARC M12 and Fujitsu M10 utilize a system interconnect designed to deliver massive bandwidth and consistent, low latency between components. The system bus benefits IT operations by delivering balanced and predictable performance to application workloads.

The interconnect design maximizes the overall performance of the Fujitsu SPARC M12 and Fujitsu M10. Implemented as point-to-point connections that utilize packet-switching technology, this system bus provides fast response times by transmitting multiple data streams.

Packet switching allows the interconnect to operate at much higher system-wide throughput by eliminating "dead" cycles on the bus. All routes are unidirectional, non-contentious paths with multiplexed address, data, and control in each direction.

System controllers within the interconnect architecture of the Fujitsu SPARC M12 and Fujitsu M10 direct traffic between CPUs in each chassis, memory, I/O subsystems, and interconnect paths.




Figure 4-1. Fujitsu SPARC M12 Bandwidth and Data Transfer Rate




Figure 4-2. Fujitsu M10 Bandwidth and Data Transfer Rate

2. Fujitsu SPARC M12-1 Interconnect Architecture

The Fujitsu SPARC M12-1 is implemented on a single physical system board with one SPARC64 XII processor. The memory access, I/O, and system controllers are embedded in the processor. The processor is connected directly to DIMMs and PCI Express switches.

An architecture diagram of the Fujitsu SPARC M12-1 is shown in Figure 4-3.




Figure 4-3. Fujitsu SPARC M12-1 Interconnect Architecture

3. Fujitsu SPARC M12-2 Interconnect Architecture

Each Fujitsu SPARC M12-2 supports one SPARC64 XII processor by default and accommodates up to two processors. The memory access, I/O, and system controllers are embedded in the processor. Each processor connects to DIMMs and PCI Express switches, and both processors are interconnected by the system bus. To increase I/O bus bandwidth and protect against processor failure, each PCI Express switch connects to two SPARC64 X+ or SPARC64 X processors (Figure 4-4).




Figure 4-4. Fujitsu SPARC M12-2 Interconnect Architecture

4. Fujitsu SPARC M12-2S Interconnect Architecture

The Fujitsu SPARC M12-2S implements crossbar units in the server and crossbar boxes in order to connect multiple Fujitsu SPARC M12-2S chassis using the Building Block method. Each Fujitsu SPARC M12-2S supports one SPARC64 XII processor by default and accommodates up to two processors. The memory access, I/O, and system controllers are embedded in the processor. Each processor connects to DIMMs and PCI Express switches, and both processors are interconnected by the system bus. To increase I/O bus bandwidth and protect against processor failure, each PCI Express switch connects to two SPARC64 X+ or SPARC64 X processors (Figure 4-5).




Figure 4-5. Fujitsu SPARC M12-2S Interconnect Architecture

The Fujitsu SPARC M12-2S, adopting the Building Block method, implements crossbar units to connect the system buses between Fujitsu SPARC M12-2S chassis. The system buses support low-latency and high-throughput data transfers.

To improve performance, the physical addressing of memory in the Building Block configuration is evenly spread out across the system controllers of all Fujitsu SPARC M12-2S chassis.

The SPARC64 XII processor has a direct connection to the crossbar units, minimizing the latency of data transfers over multiple Fujitsu SPARC M12-2S chassis. You can use point-to-point connection to directly connect the built-in crossbar units of up to four Fujitsu SPARC M12-2S chassis and expand to up to eight CPU chips (Figure 4-6). With system configurations scalable up to 16 Fujitsu SPARC M12-2S chassis, each chassis is connected to all other chassis via crossbar boxes (Figure 4-7). The system interconnect used in the Fujitsu SPARC M12-2S delivers as much as 5,453 GB/s of peak bandwidth and provides near-linear scalability.





Figure 4-6. System Configuration Directly Connecting Servers (Fujitsu SPARC M12-2S 4BB)



Note: This figure shows only the connection between the BB#0 ports and XB-BOX (0 System).
The BB#1 ports and XB-BOX (1 System) are connected in the same way.

Figure 4-7. System Configuration of Connections via Crossbar Boxes (Fujitsu SPARC M12-2S 16BB)

5. Fujitsu M10-1 Interconnect Architecture

The Fujitsu M10-1 is implemented on a single physical system board. One SPARC64 X or SPARC64 X processor is mounted. The memory access, I/O, and system controllers are embedded in the processor. The SPARC64 X+ or SPARC64 X is connected to DIMMs and PCI Express switches.

An architecture diagram of the Fujitsu M10-1 is shown in Figure 4-8.




Figure 4-8. Fujitsu M10-1 Interconnect Architecture

6. Fujitsu M10-4 Interconnect Architecture

The Fujitsu M10-4 is implemented on two physical system boards, which function as a single logical system board. The Fujitsu M10-4 supports up to either four SPARC64 X+ processors or four SPARC64 X processors; and as with the Fujitsu M10-1, the memory access, I/O, and system controllers are embedded in the CPU. Each controller connects to DIMMs and PCI Express switches, and all four SPARC64 X+ or SPARC64 X processors are interconnected. To increase I/O bus bandwidth and protect against processor failure, each PCI Express switch connects to two SPARC64 X+ or SPARC64 X processors (Figure 4-9).




Figure 4-9. Fujitsu M10-4 Interconnect Architecture

7. Fujitsu M10-4S Interconnect Architecture

The Fujitsu M10-4S implements crossbar units in the server and crossbar boxes in order to connect multiple Fujitsu M10-4S chassis using the Building Block method. In each chassis, as in the Fujitsu M10-4, two physical system boards are implemented and function as a single logical system board. The Fujitsu M10-4 supports up to either four SPARC64 X+ processors or four SPARC64 X processors; and as with the Fujitsu M10-1, the memory access, I/O, and system controllers are embedded in the CPU. Each controller connects to DIMMs and PCI Express switches, and all four SPARC64 X+ or SPARC64 X processors are interconnected. To increase I/O bus bandwidth and protect against processor failure, each PCI Express switch connects to two SPARC64 X+ or SPARC64 X processors (Figure 4-10).




Figure 4-10. Fujitsu M10-4S Interconnect Architecture

The Fujitsu M10-4S, adopting the Building Block method, implements crossbar units to connect the system buses between Fujitsu M10-4S chassis. The system buses support low-latency and high-throughput data transfers.

To improve performance, the physical addressing of memory in the Building Block configuration is evenly spread out across the system controllers of all Fujitsu M10-4S chassis.

The SPARC64 X+ or SPARC64 X processor has a direct connection to the crossbar unit,

minimizing the latency of data transfers over multiple Fujitsu M10-4S chassis. You can use point-to-point connection to directly connect the built-in crossbar units of up to 4 Fujitsu M10-4S chassis and expand to up to 16 CPU chips (Figure 4-11). With system configurations scalable to up to 16 Fujitsu M10-4S chassis, each chassis is connected to all other chassis via crossbar boxes (Figure 4-12). The system interconnect used in the Fujitsu M10-4S delivers as much as 6,553 GB/second of peak bandwidth, offering 9 times more system throughput than Fujitsu's previous generation of high-end servers.

The Fujitsu M10-4S Building Blocks with mounted SPARC64 X processors can be connected to both the Fujitsu M10-4S Building Blocks with SPARC64 X processors and the Fujitsu M10-4S Building Blocks with SPARC64 X+ processors.

This ensures that the system can be upgraded to a higher-performance system without wasting existing investment in the Fujitsu M10-4S Building Blocks with SPARC64 X processors.




Figure 4-11. System Configuration Directly Connecting Servers (Fujitsu M10-4S 4BB)




Figure 4-12. System Configuration of Connections via Crossbar Boxes (Fujitsu M10-4S 16BB)

8. System Interconnect Reliability Features

The built-in redundancy and reliability features of the Fujitsu SPARC M12 and Fujitsu M10 system interconnect enhance the stability of these servers. The interconnect is protected against loss or corruption of both transaction addresses and data by ECC or CRC protection on all system buses. When a single-bit data error is detected in a CPU, memory access controller, or I/O controller, hardware corrects the data and proceeds with the transfer. Also, when a multi-bit data error is detected by CRC on a bus that is connected via Fujitsu high-speed interconnect technology (all crossbar-crossbar, crossbar-processor, and processor-processor buses), the hardware automatically resends the data. If the error cannot be recovered by resending the data, the specific lane is degraded. In the rare event of a hardware failure within the system interconnect, the system uses the surviving bus route on restart, isolating the faulty lane and facilitating the resumption of operation.




Figure 4-13. System Interconnect Reliability Features

Memory

The memory subsystem in the Fujitsu SPARC M12 or Fujitsu M10 increases system scalability and throughput. In a 16 Building Block configuration, the Fujitsu SPARC M12 accommodates up to 3 TB (24 memory slots per processor) or 2 TB (16 memory slots per processor) of memory, and the Fujitsu M10 accommodates up to 64 TB of memory. The Fujitsu SPARC M12 supports DDR4 DIMMs and the Fujitsu M10 supports DDR3 DIMMs, with up to 4-way memory interleaving in both the Fujitsu SPARC M12 and Fujitsu M10 to enhance system performance. The available DIMM capacities include 8 GB, 16 GB, 32 GB, and 64 GB. Further details about the memory subsystem are described in Table 4-1 and Table 4-2.

Table 4-1. Fujitsu SPARC M12 Memory Subsystem Specifications

	Fujitsu SPARC M12-1	Fujitsu SPARC M12-2	Fujitsu SPARC M12-2S (1BB)	Fujitsu SPARC M12-2S (16BB) (*1)
		DDR4-2400/1866 (*2)		
Maximum Memory Capacity	1 TB	3 TB (24 memory slots/processor) 2 TB (16 memory slots/processor)	3 TB (24 memory slots/processor) 2 TB (16 memory slots/processor)	48 TB (24 memory slots/processor) 32 TB (16 memory slots/processor)
DIMM Slots, Maximum	16	48 (24 memory slots/processor) 32 (16 memory slots/processor)	48 (24 memory slots/processor) 32 (16 memory slots/processor)	768 (24 memory slots/processor) 512 (16 memory slots/processor)
Unit of Memory Interleave	4 or 8 DIMMs	8 DIMMs	8 DIMMs	8 DIMMs
Maximum Number of Interleaves	4	4	4	64

*1 Up to the 4BB configuration, for the Fujitsu SPARC M12-2S that began selling in April 2022. The resulting maximum value of each item in the specifications becomes 1/4 of that for the 16BB configuration.

*2 24 memory slots per processor, 24 DIMMs mounted

Table 4-2. Fujitsu M10 Memory Subsystem Specifications

	Fujitsu M10-1	Fujitsu M10-4	Fujitsu M10-4S (1BB)	Fujitsu M10-4S (16BB)
DDR3-1600/1333				
Maximum Memory Capacity	1 TB	4 TB	4 TB	64 TB
DIMM Slots, Maximum	16	64	64	1024
Unit of Memory Interleave	4 or 8 DIMMs	8 DIMMs	8 DIMMs	8 DIMMs
Interleaves, Maximum	4	8	8	128

Beyond performance, the Fujitsu SPARC M12/Fujitsu M10 memory subsystem is built with reliability in mind. ECC protection is implemented for all data stored in main memory, and the following advanced features foster early diagnosis and fault isolation that preserve system integrity and raise application availability.

- Memory patrol

Memory patrol periodically scans memory for errors. This proactive function prevents the use of faulty areas of memory before they can cause system or application errors, thereby improving system reliability.

- Memory Extended ECC

The Memory Extended ECC function provides single-bit error correction, supporting continuous operation even when events caused by memory device failures like burst read errors occur. This feature is similar to IBM's Chipkill technology.

1. Memory Mirroring

The Fujitsu SPARC M12 and Fujitsu M10 support memory mirroring capabilities. Memory mirroring is a high-availability feature appropriate when running applications with the most stringent availability requirements. When memory mirroring mode is enabled on the Fujitsu SPARC M12 or Fujitsu M10, the memory subsystem duplicates the data on write and compares the data on read to each side of the memory mirror. In the event that errors occur at the bus or DIMM level, normal data processing continues through the other memory bus and alternate DIMM set. In the Fujitsu SPARC M12 and Fujitsu M10, memory can be mirrored between memory modules, using the memory access controller (MAC) built into the SPARC64 XII, SPARC64 X+, and SPARC64 X processors (Figure 4-14).




Figure 4-14. Fujitsu SPARC M12 and Fujitsu M10 Memory Mirroring Architecture

System Clock

In the Fujitsu SPARC M12 and Fujitsu M10, the system clock is implemented independently for each CPU. Therefore, even after a system clock failure, the system can be restarted by degrading only the failed processor.

I/O Bus

The Fujitsu SPARC M12 and Fujitsu M10 use PCI Express buses to provide high-speed data transfer within the I/O subsystem. To provide optimal I/O performance for current and future PCI Express cards (such as Fibre Channel, InfiniBand, Gigabit Ethernet, and Flash PCI Express cards), the Fujitsu SPARC M12 and Fujitsu M10 implement the PCI Express 3.0 protocol in each processor. The peak data transfer rate of PCI Express 3.0 reaches 8 GB/s of throughput.

1. I/O Subsystem Architecture

PCI Express switches provide the connections between the SPARC64 XII, SPARC64 X+, or SPARC64 X processor and PCI Express slots, onboard devices, and internal drives. The PCI Express bus can also be extended with PCI Expansion Units to provide a significant increase in PCI Express slots. The Fujitsu SPARC M12 and Fujitsu M10 I/O architecture has been designed to provide the scalability and performance (latest PCI Express 3.0 high throughput) required for server consolidation and virtualization.

To facilitate hot plugging of PCI Express cards, the Fujitsu SPARC M12-2, Fujitsu SPARC M12-2S, Fujitsu M10-4 and Fujitsu M10-4S and the PCI Expansion Unit utilize PCI Express cassettes. PCI Express cards that support PCI hot plug can be mounted in PCI Express cassettes and then inserted into onboard or PCI Expansion Unit PCI Express slots of a running server.

2. Fujitsu SPARC M12-1 I/O Subsystem

A depiction of the I/O subsystem of the Fujitsu SPARC M12-1 is shown in Figure 4-15. The figure shows two PCI Express switches mounted on the physical system board of the Fujitsu SPARC M12-1. These switches connect to all I/O components and the I/O controller in the processor. The I/O subsystem provides three PCI Express slots, one external SAS port, and two USB ports (Front: USB 2.0, Rear: USB 3.0). The external SAS port can be used to connect an SAS tape or storage device, and the USB ports to connect any supported DVD devices. You can increase these slots by connecting a PCI Expansion Unit.




Figure 4-15. Fujitsu SPARC M12-1 I/O Subsystem Architecture

3. Fujitsu SPARC M12-2 I/O Subsystem

Four PCI Express switches are mounted on the physical system board of the Fujitsu SPARC M12-2. These switches connect all I/O components to the I/O controller in the processor. The PCI Express switches and the I/O controller in the CPU are connected one-to-one or one-to-two, depending on the number of CPU chips mounted. The I/O subsystem enables redundant configuration with two internal SAS controllers and two 10 GbE controllers. The I/O subsystem provides 11 PCI Express slots, 1 external SAS port, and 2 USB ports (Front: USB 2.0, Rear: USB 3.0). The external SAS port can be used to connect an SAS tape or storage device, and the USB ports to connect any supported DVD devices. You can increase these slots by connecting a PCI Expansion Unit.




Figure 4-16. Fujitsu SPARC M12-2 I/O Subsystem Architecture

4. Fujitsu SPARC M12-2S I/O Subsystem

Four PCI Express switches are mounted on the physical system board of the Fujitsu SPARC M12-2S. These switches connect all I/O components to the I/O controller in the processor. The PCI Express switches and the I/O controller in the CPU are connected one-to-one or one-to-two, depending on the number of CPU chips mounted. The I/O subsystem enables redundant configuration with two internal SAS controllers and two 10 GbE controllers. The I/O subsystem provides eight PCI Express slots, one external SAS port, and two USB ports (Front: USB 2.0, Rear: USB 3.0). The external SAS port can be used to connect an SAS tape or storage device, and the USB ports to connect any supported DVD devices. You can increase these slots by connecting a PCI Expansion Unit.




Figure 4-17. Fujitsu SPARC M12-2S I/O Subsystem Architecture

5. Fujitsu M10-1 I/O Subsystem

Two PCI Express switches are mounted on the physical system board of the Fujitsu M10-1. These switches connect all I/O components to the I/O controller in the SPARC64 X+ or SPARC64 X processor. The I/O subsystem provides three PCI Express slots, one external SAS port, and two USB 2.0 ports. The external SAS port can be used to connect an SAS tape or storage device, and the USB 2.0 ports to connect any supported DVD devices. You can increase these slots by connecting a PCI Expansion Unit.




Figure 4-18. Fujitsu M10-1 I/O SubsystemArchitecture

6. Fujitsu M10-4 I/O Subsystem

Four PCI Express switches are mounted on the physical system board of the Fujitsu M10-4. These switches connect all I/O components to the I/O controller in the SPARC64 X+ or SPARC64 X processor. The PCI Express switches and the I/O controller in the CPU are connected one-to-one or one-to-two, depending on the number of CPU chips mounted. The I/O subsystem provides 11 PCI Express slots, 1 external SAS port, and 2 USB 2.0 ports. The external SAS port can be used to connect an SAS tape or storage device, and the USB 2.0 ports to connect any supported DVD devices. You can increase these slots by connecting a PCI Expansion Unit.




Figure 4-19. Fujitsu M10-4 I/O Subsystem Architecture

7. Fujitsu M10-4S I/O Subsystem

Four PCI Express switches are mounted on the physical system board of the Fujitsu M10-4S. These switches connect all I/O components to the I/O controller in the SPARC64 X+ or SPARC64 X processor. The PCI Express switches and the I/O controller in the CPU are connected one-to-one or one-to-two, depending on the number of CPU chips mounted in each Fujitsu M10-4S chassis. The I/O subsystem provides eight PCI Express slots, one external SAS port, and two USB 2.0 ports per Fujitsu M10-4S chassis. The external SAS port can be used to connect an SAS tape or storage device, and the USB 2.0 ports to connect any supported DVD devices. You can increase these slots by connecting a PCI Expansion Unit.




Figure 4-20. Fujitsu M10-4S I/O SubsystemArchitecture

8. Internal Drives and Peripherals

The Fujitsu SPARC M12 and Fujitsu M10 support eight hot-swappable internal Serial Attached SCSI (SAS) 2.5-inch hard disk drives and 2.5-inch SSDs. Fujitsu SPARC M12 on-board SAS controllers support RAID 0, RAID 1, RAID 1E, and RAID 10 volumes and hot-swap drives. Fujitsu M10 on-board SAS controller supports RAID 0, RAID 1, and RAID 1E volumes and hot-swap drives. The Fujitsu SPARC M12 and Fujitsu M10 also provide one external SAS port, which can be connected to a SAS tape or storage device. The SAS port offers 4 lanes, supporting up to 24 Gbps total bandwidth. Also, the Fujitsu SPARC M12 supports one USB 3.0 port and one USB 2.0 port, and the Fujitsu M10 supports two USB 2.0 ports, which may be used to connect a supported DVD drive.

9. PCI Expansion Unit

The Fujitsu SPARC M12 and Fujitsu M10 support the attachment of optional PCI Expansion Units to provide additional I/O connectivity. The PCI Expansion Unit is a 2U rack-mountable device that accommodates up to 11 PCI Express slots. By using PCI Express cassettes, the PCI Expansion Unit supports the hot-swapping of PCI Express cards.

A link card mounted in a PCI Express slot in the server provides connectivity to the PCI Expansion Unit and supports management control via sideband signals. The link card is a low-profile card and includes an 8-lane PCI Express bus with 8 GB/s bandwidth. The PCI Expansion Unit architecture of the Fujitsu SPARC M12 and Fujitsu M10 provides high-throughput I/O performance, supporting maximum data rates for various current and future PCI Express cards (such as Fibre Channel, InfiniBand, Gigabit Ethernet, and Flash PCI Express cards).




Figure 4-21. PCI Expansion Unit Architecture

A PCI Expansion Unit connects from the link board in the PCI Expansion Unit to the link card mounted in a PCI Express slot in the server through a cable. Cable options include a copper cable kit and a fibre cable kit. The Fujitsu SPARC M12 and Fujitsu M10 support the connection of multiple PCI expansion units as shown in Table 4-3 and Table 4-4.

Table 4-3. Connectable PCI Expansion Units on Fujitsu SPARC M12

	CPU configuration	OS Version	Connectable PCI Expansion Units (Maximum)
Fujitsu SPARC M12-1	1 CPU	Oracle Solaris 11 (*1)	3
		Oracle Solaris 10 (*2)	2
Fujitsu SPARC M12-2	1 CPU	Oracle Solaris 11 (*1)	4 (*3) or 3 (*4)
		Oracle Solaris 10 (*2)	2
	2 CPU	Oracle Solaris 11 (*1)	8 (*3) or 6 (*4)
		Oracle Solaris 10 (*2)	6
Fujitsu SPARC M12-2S (per Server)	1 CPU	Oracle Solaris 11 (*1)	4 (*3) or 3 (*4)
		Oracle Solaris 10 (*2)	2
	2 CPU	Oracle Solaris 11 (*1)	8 (*3) or 5 (*4)
		Oracle Solaris 10 (*2)	5

*1 Oracle Solaris 11: OS versions on the control domain, root domain, and I/O domain are all Oracle Solaris 11.

*2 Oracle Solaris 10: OS versions on the control domain, root domain, and I/O domain are all Oracle Solaris 10 or a mixture of Oracle Solaris 11 and Oracle Solaris 10.

*3 For firmware version XCP 3040 or later

*4 For firmware version XCP 3030 or earlier

Table 4-4. Connectable PCI Expansion Units on Fujitsu M10

	CPU configuration	Connectable PCI Expansion Units (Maximum)
Fujitsu M10-1	1 CPU	2
Fujitsu M10-4	2 CPUs	3
	4 CPUs	6
Fujitsu M10-4S (per Server)	2 CPUs	3
	4 CPUs	5

Cooling

The Fujitsu SPARC M12-1 and Fujitsu M10-1 use fans for air cooling. The Fujitsu SPARC M12-2 and M12-2S and the Fujitsu M10-4 and M10-4S employ a technology that effectively cools the heat of the processor to increasingly boost its performance.

Liquid Loop Cooling (LLC)

The Fujitsu M10-4 and Fujitsu M10-4S employ Liquid Loop Cooling (LLC). This cooling method combines the advantages of liquid cooling technology (high cooling performance) with the advantages of air cooling technology (ease of installation and operation). LLC consists of the following three functional blocks:

1. Heat transfer block: Pipes and pumps
2. Heat absorber block: Cooling plates
3. Heat dissipation block: Radiator

Refrigerant is circulated in the LLC system by pumps in the heat transfer block. In the heat absorber block, the refrigerant absorbs heat dissipated from the cooling plate mounted to the CPU and power supply components. The refrigerant flows through the pipes of the heat transfer block, and is sent into the radiator where it is cooled by forced air from fans mounted in the front of the chassis. The refrigerant returns through the pipes and pumps back to the cooling plate.

By circulating the refrigerant entirely within the chassis, LLC enhances the ease of server installation and serviceability because there is no need to supply liquid from outside or perform maintenance on the liquid cooling mechanism.




Figure 4-22. Mechanism of Liquid Loop Cooling

Vapor and Liquid Loop Cooling (VLLC)

The Fujitsu SPARC M12-2 and Fujitsu SPARC M12-2S employ Vapor and Liquid Loop Cooling (VLLC), a new cooling technology using evaporative cooling, instead of LLC.

Like LLC, VLLC consists of the following three functional blocks:

1. Heat transfer block: Pipes and pumps
2. Heat absorber block: Cooling plate
3. Heat dissipation block: Radiator

As with LLC, pumps in the heat transfer block circulate refrigerant in the VLLC system, and the radiator dissipates the heat absorbed by the cooling plate into the air.

VLLC uses evaporative cooling, which absorbs heat as the liquid changes to vapor, in addition to the LLC method of hybrid cooling with air and water cooling technologies.

VLLC achieves approximately two times better cooling performance than LLC, while keeping the following LLC features: maintenance-free, ease of server installation, and serviceability.




Figure 4-23. Mechanism of Vapor and Liquid Loop Cooling

High-Efficiency Power Supplies

The Fujitsu SPARC M12 and Fujitsu M10 employ power supply units that are 80 PLUS (*1) certified, indicating the high power conversion efficiency of computer power supply units. Specifically, they are 80 PLUS PLATINUM certified.

The power supply units of the Fujitsu SPARC M12-2 and M12-2S and the Fujitsu M10-4 and M10-4S realize high power conversion efficiency with the following technologies.

- Low-loss components: SiC diodes and Super Junction FETs

Traditional silicon-based power semiconductors in switching power supply circuits have been replaced with low-loss power semiconductors, such as SiC diodes and Super Junction FETs, to lower the power loss.

- High-efficiency circuits: Bridge-less rectifier and standby-less power feed circuits

Bridge-less rectifier circuits, which control rectifier and later-stage PFC (Power Factor Correction) circuits in an integrated fashion, are adopted to achieve high conversion efficiency by reducing the number of stages of power conversion circuits. In addition, Fujitsu's unique standby-less power feed method, which feeds power solely from the main power supply, achieves high efficiency with fewer standby power supply circuits.

- Secondary-side low-loss wiring: Bus bar structure of wiring in the secondary-side output section

The printed-circuit boards previously used for connection between power supply units and secondary-side power components have been replaced with wiring in a copper bus bar structure that reduces direct-current resistance to lower the power loss.

*1 80 PLUS is a program for power-saving electrical equipment, promoted by the 80 PLUS® Program. It indicates that a power supply unit has an 80% or higher efficiency of power conversion from alternating current to direct current. The efficiency rises in the order shown in the rankings of STANDARD, BRONZE, SILVER, GOLD, PLATINUM, and TITANIUM.

URL: <https://plugloadsolutions.com/80pluspowersupplies.aspx>

5. System Management

Reliability, Availability, and Serviceability

To reduce downtime without impacting the availability of key services, mechanisms that foster fault resilience and quick repair are required. A mechanism for rapid expansion must also be included. Specifically designed to support complex, network computing solutions and stringent high-availability requirements, the Fujitsu SPARC M12 and Fujitsu M10 include redundant and hot-swappable system components, diagnostic and error recovery features throughout the design, and built-in remote management features. The advanced architecture of these reliable servers delivers high levels of application availability and rapid recovery from many types of hardware faults, simplifying system operation and lowering costs for enterprises.




Figure 5-1. Reliability, Availability, and Serviceability of the Fujitsu SPARC M12




Figure 5-2. Reliability, Availability, and Serviceability of the Fujitsu M10

1. Redundant and Hot-Swappable Components

Today's IT organizations are challenged by the pace of non-stop business operations. This is forcing planned downtime windows to shrink and, in some cases, disappear entirely. To meet these demands, the Fujitsu SPARC M12 and Fujitsu M10 employ built-in redundant and hot-swappable hardware to help mitigate the disruptions caused by individual component failures or changes to system configurations. In fact, these systems are able to recover from hardware failures often with no impact to users or system functionality.

The Fujitsu SPARC M12 and Fujitsu M10 feature redundant, hot-swappable power supply and fan units. The Fujitsu SPARC M12-2 and M12-2S and the Fujitsu M10-4 and M10-4S also feature hot-swappable PCI Express cards. Administrators can choose to create redundant internal storage in the Fujitsu SPARC M12 or Fujitsu M10 by combining hot-swappable disk drives with either hardware RAID or disk mirroring software. Furthermore, with the Fujitsu SPARC M12-2S or Fujitsu M10-4S, active maintenance of the Fujitsu SPARC M12-2S or Fujitsu M10-4S in a Building Block configuration allows the CPU or memory configuration to be changed while keeping other physical partitions in operation. In multi-Building Block configurations, the Fujitsu SPARC M12-2S and Fujitsu M10-4S feature redundant service processors. The Fujitsu SPARC M12-2S and Fujitsu M10-4S also include degradable crossbar units. If a fault occurs, these duplicated components support continued operation. Depending upon the component and type of error, the system can continue to operate in degraded mode, or

the system may reboot – with the failure automatically diagnosed and the relevant component automatically configured out of the system. Hot-swappable hardware within the Fujitsu SPARC M12 and Fujitsu M10 speeds up service and allows for simplified replacement or addition of components without a need to stop other physical partitions.

2. Partitioning Feature

In order to reduce costs and administrative burden, many enterprises look to server consolidation. However, organizations require tools that increase the security and effectiveness of hosting multiple applications on a single server. On the Fujitsu SPARC M12-2S and Fujitsu M10-4S (one chassis is referred to below as one Building Block (BB)), a physical partition can be created in units of Building Blocks by connecting the high-speed interconnect of multiple Building Blocks using the Building Block method. When a physical partition is configured across multiple Building Blocks, all of the hardware resources from those Building Blocks are available to the physical partition. The physical partitioning feature provides IT organizations with the ability to construct up to 16 independent physical partitions (with 16 Building Blocks), a single large system spanning all Building Blocks, or any configuration in between. Oracle Solaris operates in each physical partition independently. Oracle VM Server for SPARC allows for the construction of multiple independent logical domains on top of the physical partitions. Oracle Solaris operates on each logical domain independently. With proper configuration, software faults in one partition or domain remain isolated and unable to impact the operation of other partitions or domains.

The Dynamic Reconfiguration (DR) function allows the Fujitsu SPARC M12-2S and Fujitsu M10-4S to expand hardware resources such as CPUs and memory, without interrupting the operating system, by adding a Building Block. The DR function also allows the moving of hardware resources from one partition to the other. By using the DR function, great flexibility and availability can be achieved on the Fujitsu SPARC M12-2S and Fujitsu M10-4S. Some expected use cases are as follows:

- To expand hardware resources without interruption of production, when workload increases
- To reassign hardware resources based on a plan without interruption of production (example: Moving CPUs and memory to a partition, where there is database software for sales management, from another partition at the end of each quarter)
- To disconnect, replace, and connect a broken Building Block without interruption of production in order to repair a faulty CPU/memory

3. Advanced Reliability Features

Advanced reliability features included within the components of the Fujitsu SPARC M12 and Fujitsu M10 increase the overall stability of the platform. By virtue of the consolidated peripheral ASIC functions in the processor, the reduced component count and complexity within the server architecture contribute to reliability. In addition, CRC and ECC protection mechanisms guarantee data path integrity in processors and provide for autonomous error recovery, reducing the time to initiate corrective action and subsequently increasing uptime. The large-capacity data paths use Fujitsu high-speed interconnect technology to provide advanced error detection and correction features. Moreover, each Fujitsu SPARC M12-2S and Fujitsu M10-4S includes multiple degradable crossbar switches that provide redundancy within the system bus; the switches also provide IT organizations with the ability to construct high-availability systems.

The Oracle Solaris Predictive Self-Healing software further enhances the reliability of the Fujitsu SPARC M12 and Fujitsu M10. Implementation of the Oracle Solaris Predictive Self-Healing software provides constant monitoring of CPUs and memory. Depending on the nature of the error, persistent soft CPU errors can be resolved by automatically offlining a single thread, one core, or even all cores in a processor. In addition, the memory page retirement function provides the ability to take memory pages offline proactively in response to multiple correctable errors in a specific memory DIMM.

4. Error Detection, Diagnosis, and Recovery

The Fujitsu SPARC M12 and Fujitsu M10 feature important technologies that correct failures early and keep marginal components from causing repeated downtime. Architectural advances that inherently increase reliability are augmented by error detection and recovery capabilities within the server hardware subsystems, as described below.

- End-to-end data protection detects and corrects errors throughout the system, ensuring complete data integrity.
- State-of-the-art fault isolation helps the Fujitsu SPARC M12 and Fujitsu M10 isolate errors within component boundaries and offline only the relevant device segment instead of the entire component. Isolating errors down to the chip or down to the block in the chip improves stability and provides continued availability of maximum compute power. This feature applies to CPUs, memory, memory access controllers in the CPU, crossbar ASICs, and PCI Express I/O controllers in the CPU.
- Constant environmental monitoring provides a historical log of all pertinent environmental and error conditions.
- The host watchdog feature regularly checks the operation of the processors to ensure

operating system and application operation is proceeding normally. In the event that a watchdog timeout occurs, recovery functions are automatically triggered.

- Periodic component status checks are performed to determine the status of many system devices and detect signs of an impending fault. Recovery mechanisms are triggered to prevent system and application failure.
- Error logging, multistage alerts, electronic FRU identification information, and system fault LED indicators contribute to rapid problem resolution.

System Management

Providing hands-on, local system administration for server systems is no longer realistic for most organizations. Around-the-clock system operation, disaster recovery hot sites, and geographically dispersed organizations lead to a requirement for remote management of systems. One of the many benefits of Fujitsu servers is the support for lights-out datacenters, letting expensive support staff work from any location with network access. The Fujitsu SPARC M12 and Fujitsu M10 have the powerful eXtended System Control Facility (XSCF) and XSCF Control Package to help administrators remotely execute and control nearly any task that does not involve changes to the hardware configuration. These remote functions lower the administrative burden, saving organizations time and reducing operational expenses.

1. eXtended System Control Facility

The eXtended System Control Facility (XSCF) is the heart of remote monitoring and management capabilities in the Fujitsu SPARC M12 and Fujitsu M10. The XSCF consists of a dedicated processor that is independent of the server system and runs XSCF Control Package.

The XSCF regularly monitors environmental sensors throughout the system, provides advance warning of potential error conditions, and executes proactive system maintenance procedures as necessary. For example, the XSCF will initiate a server shutdown in response to temperature conditions that might induce physical system damage. XSCF Control Package running on the service processor helps administrators to remotely control and monitor partitions and domains as well as the platform itself.

Using a network or serial connection to the XSCF, administrators can effectively administer the server from anywhere on the network. Remote connections to the service processor run separately from the operating system and provide the full control and authority of a system console.

2. Redundant XSCF

On the Fujitsu SPARC M12-2S and Fujitsu M10-4S configured with more than one Building Block, one XSCF is configured as active and the other is configured as standby. The XSCF network between the two service processors facilitates the exchange of system management information. In case of an XSCF failure, the service processors are already synchronized and ready to fail over and change roles. In addition, the Fujitsu SPARC M12-2S allows replacement of a failed XSCF by hot-swapping while another XSCF is operating normally.

3. XSCF Control Package

XSCF Control Package helps users to control and monitor the Fujitsu SPARC M12 and Fujitsu M10 and individual partitioning functions quickly and effectively. XSCF Control Package provides a command line interface (CLI) and Web browser user interface that give administrators and operators access to system controller functionality. Secure accounts with specific administration capabilities also provide system security for partition and primary logical domain consoles. Communication between the XSCF and domains supports encrypted connections based on Secure Shell (SSH) and Secure Socket Layer (SSL), supporting secure, remote execution of XSCF Control Package commands.

XSCF Control Package provides the interface for the following key server functions:

- Physical partition administration including the assignment of the Fujitsu SPARC M12-2S and Fujitsu M10-4S to physical partitions;
- Audit administration includes the logging of interactions between the XSCF and partitions and primary logical domains;
- Monitoring and control of power to components in the server;
- Interpretation of presented hardware information and notification of impending problems such as high temperatures or power supply problems;
- Integration with the Oracle Solaris Fault Management Architecture to improve availability through accurate fault diagnosis and predictive fault analysis;
- Execution and monitoring of diagnostic programs, such as the OpenBoot PROM (OBP) and power-on self-test (POST);
- Execution of CPU Activation operations providing the ability to stage and then later activate additional resources; and
- Monitoring of dual XSCF configurations for failures, and performing automatic failover if needed.

4. Role-based System Management

XSCF Control Package facilitates the independent administration of multiple autonomous physical partitions by different system administrators and operators – each utilizing portions of a single Fujitsu SPARC M12 and Fujitsu M10 platform. This management software supports multiple user accounts organized into groups. Different privileges are assigned to each group. Privileges allow a user to perform a specific set of actions on a specific set of hardware, including physical components, physical partitions, or physical components within a physical partition. In addition, a single user can possess multiple, differing privileges on any number of physical partitions.

5. Oracle Enterprise Manager Ops Center 12c

Oracle Enterprise Manager Ops Center 12c is system management software for managing a number of virtual and physical servers in an integrated fashion, on the Fujitsu SPARC M12 and Fujitsu M10. Ops Center can reduce the cost of system management. Operations like assigning hardware resources (CPU, memory, and so on), provisioning physical and virtual servers, updating system firmware, and applying Oracle Solaris 11 patches can be done by Ops Center easily.

Eco-Friendly Computing

The power-saving functions provided in the Fujitsu SPARC M12 and Fujitsu M10 reduce the power being consumed by unused or low-utilization-rate hardware. The power-saving functions are set to invalid at the time of initial installation to give priority to performance over power-saving. However, customers who prioritize lower power consumption over performance can change the setting to the power-saving mode (Elastic Mode). The Elastic Mode is set per physical partition.

In addition to the Elastic Mode, SPARC M12 customers can change the setting to the Performance Mode where performance and power consumption are well-balanced. The Performance Mode can be set per physical partition.

The following Fujitsu SPARC M12 and Fujitsu M10 features promote power-saving.

- Lowered hardware component power consumption

In designing the Fujitsu SPARC M12 and Fujitsu M10, the selection of hardware components was made with considerable attention to lower power consumption.

- **Reduction of power consumption from unused hardware components**
Depending on the version, physical partition configuration, and logical domain configuration, some hardware mounted in the system is unused. For example, if a CPU is not assigned to a partition or domain, the CPU's power state will automatically be lowered to save power. The same also occurs for unassigned memory.
- **Sensor monitoring function**
Using this function, power consumption and airflow are monitored and logged. By collecting and making actual power consumption data available, data center power capacity design can be optimized. In a similar fashion, actual airflow data allows datacenter cooling facilities to be optimized.
- **Power capping function**
The power capping function allows the customer to set an upper threshold for system power consumption. The CPU frequency is automatically controlled so as not to exceed the threshold. This function provides the customer with control of system power consumption to fit the datacenter facilities.

6. Oracle Solaris 11 Operating System

Oracle Solaris is the industry-standard UNIX operating system with superiority in performance, scalability, reliability, and security.

The design of Oracle Solaris is suitable for high-reliability systems. It consists of a small and compact kernel, which lessens the likelihood of operating system failure and the resulting platform downtime. Furthermore, Oracle Solaris clearly distinguishes between the kernel, shared libraries, and applications, limiting the impact caused by application failures. In addition, Oracle Solaris allows enterprises and organizations to install sequentially updated software without rebooting, which increases system uptime and lightens the maintenance load.

Oracle Solaris 11 offers cutting-edge new features such as IPS (Image Packaging System), BE (Boot Environment), security enhancements, network virtualization, and Oracle Solaris 10 Zones. It has achieved increases in system installation and operating efficiency, the ability to construct safe and flexible virtual environments, and robust investment protection; all of which offer a strong technological foundation for cloud computing in mission-critical systems.

The Fujitsu SPARC M12 and Fujitsu M10 support bare-metal configurations of Oracle Solaris 10 and Oracle Solaris 11. This maximizes ROI and minimizes investment risks.

Oracle Solaris ZFS

Oracle Solaris ZFS is included by default to provide the storage virtualization function. Oracle Solaris ZFS manages multiple physical storage devices in a storage pool. By allocating the required capacity from the storage pool, a virtualized volume can be created.

Oracle Solaris ZFS enables more efficient and optimized use of storage devices, while dramatically increasing reliability and scalability. Physical storage can be dynamically added and removed from storage pools without interrupting services, providing new levels of flexibility, availability, and performance.

Oracle Solaris ZFS protects all data with 256-bit checksums, resulting in 99.9999999999999999% error detection and correction. If ZFS detects an error in a storage pool with redundancy, the corrupt data is automatically repaired. This contributes to relentless availability by protecting against costly and time-consuming data loss due to hardware or software failures, and by reducing the impact of errors by administrators during file system-related tasks.

Oracle Solaris ZFS is a 128-bit file system, which can manage an infinite capacity, in practical use. Since the metadata used to manage Oracle Solaris ZFS is dynamically allocated as needed, there are no limits on the number of file systems and the number of files. In conventional file systems, the size of the file system has been limited to the size of the physical device. However with Oracle Solaris ZFS, storage pools conceal the physical devices; therefore, size is not limited by a specific physical device. Oracle Solaris ZFS can easily create file system layers without initialization, and automatically extends capacity as disks are allocated to the ZFS storagepool.

The Oracle Solaris ZFS storage pool is a framework that consolidates the management of physical disks. Non-redundant, mirrored, RAID-Z (single parity), RAID-Z2 (double parity), or RAID-Z3 (triple parity) redundant configurations can be selected. All data written to the storage pool is dynamically striped across all available devices. In addition, in a redundant storage pool configuration, when an illegal data block is detected, the correct data is obtained from another redundant copy for self-recovery.




Figure 6-1. ZFS Storage Pool

ZFS deduplication allows for the sharing of duplicated data blocks in the ZFS storage pool, which reduces the amount of data stored and achieves more efficient use of storage. This deduplication function can be used in combination with compression and the cryptographic functionality.

Since the storage pool allows disk addition and replacement online, it is not necessary to secure disk capacity for future use at the time of initial installation.

To achieve significant performance improvements and reduced power consumption, the ZFS

hybrid storage pool was developed. It combines memory, SSDs, and disks together in a single storage pool. The ZFS Intent Log (ZIL) regularly uses the storage pool, and can deliver performance improvements for synchronized data writes by being allocated to high-speed devices such as an SSD. The ZFS cache device (L2ARC: Level 2 Adaptive Replacement Cache) can improve the performance of random reads of static data by adding a high-speed device such as an SSD as the cache between memory and disks.

Oracle Solaris ZFS is a transaction file system. Writing of data does not overwrite the existing data but copies the original data and updates it (Copy-on-Write). After completing a series of data updates, ZFS switches the data pointer for the old and new. This constantly keeps consistency in the file systems. Even with a sudden power loss, the file system is never destroyed.

To prevent data loss even in the case of sudden system interruption, the processing from (1) to (3) in Figure 6-2 makes the old data or updated data always accessible.

- (1) Instead of overwriting, copy the data to a new area and update it.
- (2) After updating the data, copy the administrative information and update it too.
- (3) Rewrite the uberblock with new administrative information, and access the updated data.




Figure 6-2. Copy-on-Write

When a file system is created, Oracle Solaris ZFS automatically generates the mount point. Since it is automatically mounted at server start, mount management is not required. When a ZFS volume is created as a block device, the initial data area size is reserved for the volume and automatically extended as needed.

Oracle Solaris ZFS has end-to-end checksumming for metadata of all user data and administrative information. Since the checksum of the data block is retained in the parent block and iterates up to the top administrative block (uberblock) in order, self-inspection of the entire data tree is possible. When an error is detected, Oracle Solaris ZFS recovers data from the redundant copy. In order to enhance reliability, the Oracle Solaris ZFS metadata is automatically saved across different disks (ditto blocks).

Furthermore, Oracle Solaris ZFS can save multiple copies of user data. Even if the data cannot be redundant across multiple disks, it is possible to recover from disk block read failures.

A ZFS snapshot is a read-only copy of the file system or volume, and can be instantly created without consuming disk capacity. Although the snapshot cannot be directly referred to, it facilitates operations such as clone and backup. A ZFS clone is a writable copy of the file system or volume, and can be created from the snapshot. Just like the snapshot, it can be created instantly without consuming significant disk capacity. The clone only requires enough disk capacity to store the changed data.




Figure 6-3. ZFS Snapshot

To protect against theft of physical storage or outside attack, as well as achieving secure deletion within the storage, data can be encrypted by each ZFS dataset.

Oracle VM Server for SPARC

Oracle VM Server for SPARC can divide the physical server into virtual servers using a firmware layer Hypervisor to configure logical domains in which independent Oracle Solaris

environments operate. CPUs, memory, and I/O devices are flexibly allocated by the Domain Manager, which can run on either Oracle Solaris 11 or 10.

The logical domain concept provides the virtual servers with logically grouped CPUs, memory, and I/O devices. In each logical domain, an independent Solaris OS operates. Each logical domain can be started or stopped individually, and up to 256 domains per physical partition can be created. The logical domains communicate with each other via the Hypervisor's Logical Domain Channel (LDC). Virtual devices such as disks and networks communicate with the virtual services through LDC, to access the physical devices.

Resources such as virtual CPUs, memory, and virtual I/O devices can be dynamically reconfigured while the logical domain is in operation. Furthermore, the number of virtual CPUs in a logical domain can be automatically increased or decreased in accordance with the dynamic resource management policy. This policy can be created by combining conditions such as the number of resources available, usage rate, upper and lower thresholds, and time period. In addition, CPU power management features can reduce power consumption by shutting down the power to a CPU not being used to process the current workload.

In line with changes to workload or server expansion, a guest domain can be migrated to another physical server. Live migration allows for the high-speed transfer of compressed memory contents without suspending guest domain operation.

To consolidate servers, a physical server can also be migrated to a logical domain using the Physical-to-Virtual (P2V) migration tool. This tool collects configuration information from a physical server and creates a file system image. Then, P2V creates the logical domain, based on the collected configuration information and restores the file system image to the virtual disk.

Oracle Solaris Zones

The Oracle Solaris Zones function virtually divides a single OS space and presents multiple OS spaces to the user. Oracle Solaris Zones also include the Oracle Solaris Resource Manager function, which flexibly allocates hardware resources such as CPUs and memory.

Oracle Solaris Zones is a virtualized OS environment that provides a safe and isolated environment suitable for application execution. A process that is executed is isolated by each zone, and does not affect other zones.

The global zone is a single zone that exists in the Oracle Solaris system. The global zone manages the entire system. Operations such as creating and managing non-global zones or allocating physical I/O device can be performed solely in the global zone.

Non-Global zones are the software partitions of the virtual Solaris environment, where

applications can be executed without affecting other zones. Up to 8,191 non-global zones can be created. Each zone can use only the designated file systems and physical I/O devices.

For system files that are used to configure the non-global zones, the necessary packages can be selected and installed when the zone is created. When packages are updated in the global zone, these files are updated and all non-global zone files are also updated in synchronization.

Oracle Solaris Resource Manager periodically monitors resource usage and automatically allocates additional resources, configured for such cases, without stopping the zone. CPU resources to be allocated to the zone are managed by the resource pool. The resource upper limit daemon controls memory. Furthermore, the resource pool consists of CPUs grouped into processor sets and CPU-allocating scheduling classes (in time sharing and fair allocation).

By copying an existing zone, new zones can be created easily. By using a ZFS clone, a zone can be replicated instantly and initially does not occupy additional disk capacity.

When resources such as CPUs or memory are insufficient in a physical server or if the fundamental organization or usage of zones needs to be changed, a zone can be migrated to another server. The zone is detached from the original server and attached to the destination server. Even if there are differences in the environment between the servers, such as the package configuration, the system file used to configure the zone is synchronized with the global zone of the destination server at the time of attachment.

Oracle Solaris 11 provides support for Oracle Solaris 10 Zones by default.

By using the P2V (Physical to Virtual) function or the V2V (Virtual to Virtual) function, an existing Oracle Solaris 10 OS environment can be migrated to a system with Oracle Solaris 11 OS Zones as it is. This allows for the consolidation of Oracle Solaris 10 and Oracle Solaris 11 environments, which contributes to investment protection for existing ICT assets and TCO reduction.

By virtualizing NICs (Network Interface Controllers) using the virtual network function, VNICs (Virtual Network Interface Controllers) can be allocated to multiple Oracle Solaris Zones. An independent network environment consisting of multiple Oracle Solaris Zones can be configured, which can reduce the number of required servers and NICs. In addition, since a virtual switch (ethers tab) is created, a virtual network environment that does not depend on the hardware can be achieved, and the number of switches can be reduced as well. Furthermore, VNIC bandwidth limits can be configured using the resource management functions.

Oracle Solaris 11.3 supports Oracle Solaris Kernel Zones. The Oracle Solaris Kernel Zones feature provides a complete kernel and user environment inside a zone and strengthens the separation of a kernel between a global zone (kernel zone host) and other zones.

Image Packaging System

IPS (Image Packaging System) is a framework that allows the management, installation, update, and deletion of the OS environment in units of packages.

IPS installs an OS environment that is optimized for basic server operation from media and then installs additional packages via a network as required for the specific intended workload. Administrators are not required to independently prepare a network installation server but can use a repository server. IPS automatically navigates through the complicated dependencies of packages. Conventional patch application is now handled by package replacement, which is done without regard to package dependencies and can prevent application failures.

As stated above, IPS aims to improve installation and operation management efficiency, which will lead directly to cost reductions.

Boot Environment

BE (Boot Environment) is a function that enables the management of multiple boot environments and simplifies online upgrade.

The root file system of the Oracle Solaris 11 OS is Oracle Solaris ZFS. To add or update packages, a ZFS snapshot and a ZFS clone enable boot environment duplication in a short period of time. The ZFS tools copy only the data blocks that will be added or updated, reducing the required disk capacity. As well as reboot from the duplicated boot environment, it is also possible to restore the original boot environment should any problems occur during the package update.

BE utilization enables generation management of boot environments, which drastically shortens the downtime for maintenance work.

7. Technological Enhancements of the SPARC64 XII, SPARC64 X+, and SPARC64 X Processors

Microarchitecture

1. Chip Configuration

The SPARC64 X+ and SPARC64 X are semiconductors that adopt the 28 nm CMOS process. Each processor has up to 16 cores and a 24-way shared L2 cache (or 22-way shared L2 cache in the Fujitsu M10-1) of up to 24 MB. The SPARC64 X+ has an operating frequency of up to 3.7 GHz, while the SPARC64 X runs at up to 3.0 GHz.

The SPARC64 XII is a semiconductor that adopts the 20 nm CMOS process. The processor has up to 12 cores and a 16-way shared L3 cache of up to 32 MB. The maximum operating frequency of the SPARC64 XII is 4.25 GHz. (See Table 7-1.)

Table 7-1. SPARC64 XII, SPARC64 X+, and SPARC64 X Specifications

	SPARC64 XII	SPARC64 X+	SPARC64 X
Cores	Up to 12	Up to 16	
L3 Cache	Up to 32 MB (L3 cache)	Up to 24 MB (L2 cache)	
Operating Frequency	Up to 4.25 GHz	Up to 3.7 GHz	Up to 3 GHz
Process Technology	20 nm CMOS	28 nm CMOS	
Die Size	25.8 mm x 30.8 mm	24.0 mm x 25.0 mm	23.5 mm x 25.0 mm
Transistors	5,450 M	2,990 M	2,950 M
Memory Bandwidth (Theoretical Peak Value)	153 GB/s	102 GB/s	

When high-speed mode (*1) of the CPU is enabled on the Fujitsu SPARC M12-2S, equipped with the SPARC64 XII, the CPU operating frequency may be higher than 4.25 GHz (up to 4.35 GHz).

*1 This can be set by an XSCF command (sethsmode). The setting of high-speed mode does not guarantee a maximum CPU operating frequency of 4.35 GHz.

For lower memory access latency and higher throughput, the SPARC64 XII, SPARC64 X+, and

SPARC64 X are equipped with internal memory access controllers (MACs). The memory bandwidth has a theoretical peak of 102 GB/s (SPARC64 X+ and SPARC64 X) or 153 GB/s (SPARC64 XII).

In addition, the SPARC64 XII, SPARC64 X+, and SPARC64 X are equipped with an internal CPU-to-CPU interface to interconnect the multiple CPUs found within the Fujitsu SPARC M12-2, Fujitsu SPARC M12-2S, Fujitsu M10-4, or Fujitsu M10-4S chassis. This CPU-to-CPU interconnect uses a high-speed serial interface to deliver high throughput when processing spans multiple CPUs.

To complete the System-On-Chip features of the SPARC64 XII, SPARC64 X+, and SPARC64 X, the I/O controller has also been integrated into the silicon. The SPARC64 X+ and SPARC64 X provide two 8-lane PCI Express, 8-GB/s ports per processor, and the SPARC64 XII provides four 8-lane PCI Express, 8-GB/s ports per processor.




Figure 7-1. SPARC64 XII, SPARC64 X+, and SPARC64 X Floor Plan

2. Core Microarchitecture

SPARC64 XII, SPARC64 X+, and SPARC64 X Core

The SPARC64 X+ or SPARC64 X core consists of an instruction fetch block and an instruction execution block, as shown in Figure 7-2. The instruction fetch block contains the L1 instruction cache, and the instruction execution block contains the L1 data cache for operands, execution unit, registers, etc.

The SPARC64 XII core consists of an L2 cache, one instruction fetch block and two instruction execution blocks (referred to below as instruction pipelines) as shown in Figure 7-3.




Figure 7-2. SPARC64 X+ and SPARC64 X Pipelines




Figure 7-3. SPARC64 XII Pipelines

Simultaneous Multithreading (SMT)

The SPARC64 XII, SPARC64 X+, and SPARC64 X carry over the simultaneous multithreading (SMT) technology first deployed in the SPARC64 VII+. With this technology, multiple threads in a single core are simultaneously processed. From the software point of view, each thread appears as an independent CPU. Hardware shares the execution resources such as the instruction buffers, reservation station, pipeline, and caches between the threads. Even when one thread is stalled in a wait for data, another thread continues processing by using the pipeline; thus, the processing performance as a core improves. Also, if one of the threads is idle, it is possible for another thread to utilize all of the execution resources and perform processing.

The SPARC64 X+ or SPARC64 X executes up to two threads per core simultaneously. The SPARC64 XII executes up to four threads per instruction pipeline and up to eight threads per core.

The following sections respectively explain the instruction fetch block and the instruction execution block.

Instruction Fetch

The instruction fetch block reads out a series of instructions from memory or from the L1 instruction cache according to branch prediction. One instruction fetch can read out eight instructions on the SPARC64 X+ or SPARC64 X or eight instructions per instruction pipeline or 16 instructions per core on the SPARC64 XII. The series of instructions that are read out is taken into the instruction buffer all at once. The instruction buffer on the SPARC64 XII is divided for each instruction pipeline. When multiple threads are running in the instruction pipelines, the instruction buffer is divided evenly for each thread. When the instruction buffer is full or when an L1 instruction cache miss occurs, the processor keeps working by executing hardware instruction prefetch to load instructions from the L2 cache or memory to the L1 cache.

Instruction fetch operates independently of the instruction execution block. In the event that the instruction execution is stalled, instruction fetch continues as long as there are available instruction buffers. In contrast, with an event such as a cache miss, instruction supply from the instruction buffer to the instruction execution block continues as long as the instruction buffer includes instructions.

While the throughput of instruction execution is four instructions per cycle for each instruction pipeline (a maximum of six instructions when including the prefix instruction SXAR described below), the throughput of instruction fetch is eight instructions per cycle for each instruction

pipeline. This helps improve system performance by concealing the access latency of the L1 instruction cache.

When using SMT, an instruction fetch of a single thread is performed in the same cycle, and the threads are switched in each cycle.

Furthermore, the SPARC64 XII, SPARC64 X+, and SPARC64 X have an enhanced branch prediction mechanism. When making branch predictions, in addition to the branch history of its own instructions, the taken or not-taken history pattern from the last series of instructions are taken into consideration; thus, prediction accuracy improves.

Instruction Execution

In the SPARC64 X+ and SPARC64 X, instructions from the instruction buffer in the instruction fetch block are supplied to the instruction execution block by up to six instructions per cycle (two of which are the prefix instruction SXAR described below). Within the instruction execution block, these instructions are decoded, issued, executed, and committed.

In the SPARC64 XII, two instruction pipelines (instruction execution blocks) are implemented per core, and up to six instructions are supplied per cycle from the instruction buffer in the instruction fetch block to each instruction pipeline. These instructions are decoded, issued, executed, and committed independently in each instruction pipeline. The subsequent description of instruction execution focuses on the operation of one instruction pipeline.

(1) Instruction Decode and Issue

In the instruction decode and issue stage, when the SXAR prefix instruction is included, up to six instructions (two of which are SXAR) are decoded simultaneously. The prefix instruction SXAR is a new instruction that has the extension information for two subsequent instructions. SXAR performs instruction extension on up to two subsequent instructions in the instruction decode stage. Furthermore, it performs "connection processing" with the subsequent instructions in order to avoid the consumption of execution resources in the pipeline stage later on. After the following processes are performed, instructions extended by SXAR are treated as so-called "valid instructions":

- The instruction that follows the SXAR prefix instruction for instruction extension performs the connection processing.
- The next instruction after the above instruction, if SXAR shows it should be extended, also performs the connection processing.

In the instruction decode stage, up to four valid instructions can be decoded simultaneously.

When using SMT, instruction decode of one of two threads (on the SPARC64 X+ or SPARC64

X) or four threads (on the SPARC64 XII) is performed in the same cycle, and the threads are switched in each cycle.

In the instruction decode stage, the resources required for execution – such as various reservation stations, the fetch port and store port, and the register update buffer – are determined. In the SPARC64 XII, SPARC64 X+, and SPARC64 X, in order to speed up the instruction processing, the resources required for execution such as reservation stations have been enhanced compared to the former SPARC64 VII+ processor.

When performing instruction decode simultaneously, there are no restrictions on instruction type combinations. As long as there are free resources, instructions can be issued. If insufficient space exists for four instructions, as many instructions as possible are issued according to the instruction order in the program. As described above, by eliminating the stall conditions of issued instructions as much as possible, a high multiplicity level is assured for any binary code.

(2) Instruction Execution

Decoded instructions are registered in a reservation station. Among fixed-point instructions, the addition and subtraction operations and the logical operations can be registered in the RSE (Reservation Station for Execution) or RSA (Reservation Station for Address generation). Other fixed-point instructions are registered in the RSE only. In addition, the address calculations of load and store instructions are registered in the RSA. That is, the RSA is shared by the addition and subtract operations, the logical operations of the fixed-point instructions, and the address calculations of load and store instructions. The floating-point arithmetic instructions are registered in the RSF (Reservation Station for Floating-point), and the branch instructions are registered in the RSBR (Reservation Station for BRanch).

The RSE starts the arithmetic pipelines of EXA and EXB. The RSA starts the pipelines of EAGA and EAGB. The RSF starts the pipelines of FLA, FLB, FLC, and FLD. Each instruction stored in a reservation station is dispatched out-of-order to the execution unit corresponding to that reservation station. This prioritizes the older instructions in program order from among those in which input operands are prepared for instructions.

When using SMT, pipelines can be used by multiple threads simultaneously.

The execution block has two fixed-point arithmetic pipelines (EXA/B), two arithmetic pipelines (EAGA/B) that perform the address calculation of load and store or the addition and subtraction operations and logical operation of the fixed-point arithmetic pipelines, and four floating-point arithmetic pipelines (FLA/B/C/D).

Specifically, the floating-point arithmetic pipeline has been significantly extended compared to the former processor. It has adopted the SIMD (Single Instruction Multiple Data) function that

was introduced in the SPARC64 VIIIfx supercomputer processor. With a single instruction, this function performs parallel processing using two pipelines (FLA and FLC, or FLB and FLD). Each pipeline has an FMA (Floating-point Multiply and Add) execution unit as before. A single FMA execution unit can execute floating-point multiplication and addition in every cycle. In each core, eight double-precision floating-point operations can be executed in every cycle.

As described below with SWoC (SoftWare on Chip) enhancements, the SPARC64 XII, SPARC64 X+, and SPARC64 X core also include new execution units to support cryptographic processing and IEEE-754 for decimal floating-point numeric calculations in DPD (Densely Packed Decimal) and Oracle NUMBER formats. Also, in the instruction execution block, an execution unit has been added to the core for the acceleration of business applications, such as databases, through the use of parallel processing.

In the SPARC64 XII, SPARC64 X+, and SPARC64 X, the floating-point registers (FPRs) have been increased by up to four times, which enhances instruction scheduling by software performance (such as loop unrolling or software pipelining).

By using the floating-point arithmetic pipelines for a wider variety of arithmetic operations and the fourfold increase in registers, marked improvement of processing performance in various applications can be achieved.

The L1 data cache block processes the load and store instructions. The data cache can provide the data for the subsequent load instruction without waiting for the address calculation of the store instruction that comes later in the series of instructions. Also, the L1 data cache is in a dual port configuration that is simultaneously accessible by two load instructions. Two 16-byte SIMD load instructions or a single 16-byte SIMD store instruction can be executed. As long as there are no bank conflicts, simultaneous processing of read-out by the load instruction and write by the store instruction is possible, improving the cache throughput.

(3) Instruction Commit

Data resulting from out-of-order executions is stored in different locations, depending on the type of data: GPR Update Buffer (GUB) for fixed-point data, FPR Update Buffer (FUB) for floating-point data, and store port for store data. Next, instruction commit is performed according to the program order to update registers, such as a GPR (General Purpose Register) or FPR, and memory.

The maximum number of valid instructions that can be committed at one time is four. The SPARC64 XII, SPARC64 X+, and SPARC64 X enables four simultaneous writes to a GPR by a fixed-point arithmetic instruction or fixed-point load instruction. The four fixed-point arithmetic and load pipelines deliver increased throughput to maximize performance.

When using SMT, an instruction commit by one of the threads is performed in the same cycle, and threads switch on alternating cycles.

Before an instruction is committed, the execution result remains inaccessible from software. As stated above, control registers including GPR, FPR, and PC (Program Counter) and memory are updated in the commit stage in order and at the same time according to the program order. By using this synchronous update method, precise interrupts are guaranteed, and in-progress processing can be canceled at any time. This method enables the processor to implement instruction retry, which is described below, and contributes to increased reliability.

3. Interface Between Chips

The SPARC64 X, SPARC64 X+, and SPARC64 XII have an on-chip CPU interface function for the connection between CPU sockets. Up to four CPU sockets (SPARC64 X+ or SPARC64 X) or two CPU sockets (SPARC64 XII) can be connected directly. The SPARC64 XII can be expanded beyond 2 sockets, up to 32 CPUs, and the SPARC64 X+ or SPARC64 X can be expanded beyond 4 sockets, up to 64 CPUs, with CPUs connected via a crossbar (XB) chip. By virtue of the increased performance of a single CPU, a fast, high-throughput serial transfer protocol has been implemented for the connection between CPU sockets.

Extended Instruction Set Architecture

The SPARC64 X, SPARC64 X+, and SPARC64 XII have adopted the virtual machine architecture that is fully compatible with the sun4v architecture supported by Oracle SPARC servers.

The SPARC64 X, SPARC64 X+, and SPARC64 XII have also introduced HPC-ACE (High Performance Computing Arithmetic Computational Extensions). This extension of the SPARC-V9 architecture instruction set was first introduced in the SPARC64 VIIIfx supercomputer processor. HPC-ACE enhances the SIMD (Single Instruction Multiple Data) function, enabling parallel processing of operations, and increases the number of floating-point registers (FPRs).

New SWoC (SoftWare on Chip) functionality has also been added to the SPARC64 X, SPARC64 X+, and SPARC64 XII. Using dedicated hardware, execution of cryptographic processing, which has conventionally relied on a combination of general instructions, can now be performed significantly faster. Another enhancement is the ability to perform IEEE-754 decimal floating-point numeric calculations in DPD and Oracle NUMBER formats.

To address trends in business applications, such as database processing, instruction extensions to manage parallel processing of data have also been added to the SPARC64 X, SPARC64 X+,

and SPARC64 XII.

The following sections describe HPC-ACE, SWoC, and the acceleration of business applications through parallel data processing.

1. HPC-ACE

Extension of Floating-Point Registers (FPRs)

The number of floating-point registers (FPRs) in SPARC-V9 is 32. This is not sufficient to fully exploit the performance of many applications. However, to increase the number of registers, the 32-bit fixed instruction length in the SPARC architecture falls short and is difficult to modify. To solve this problem, HPC-ACE has created a new prefix instruction called SXAR (Set eXtended Arithmetic Register). For up to two subsequent instructions, the SXAR instruction performs operations such as register address extension. SXAR has extended the register address with two additional bits, and the number of addressable floating-point registers (FPRs) has been increased to up to 128 – four times that defined on the SPARC-V9 (see Figure 7-4). Compilers use this large-capacity register to perform, among other uses, software pipeline optimization to fully exploit instruction level parallelism in applications.




Figure 7-4. Register Address Extension by SXAR Instruction

SIMD (Single Instruction Multiple Data)

The SIMD enhancement processes data using multiple pipelines in parallel, all with a single instruction. By adopting SIMD, HPC-ACE enables the use of two FMA (Floating-point Multiply and Add) execution units with a single arithmetic instruction.

In conventional scalar operation, there is a 1:1 ratio of data to instruction. Execution of four add operation instructions is required to process four results. With SIMD functionality, the SPARC64 XII, SPARC64 X+, and SPARC64 X is capable of simultaneous multiply-and-add processing of four sets of 64-bit data with two SIMD instructions executed in each instruction pipeline at the same time. In addition, the SPARC64 XII and SPARC64 X+ can compare up to 32 sets of 8-bit data in each instruction pipeline simultaneously by adding new instructions. The SPARC64 XII can compare up to 64 sets of 4-bit data simultaneously.

The SPARC64 XII, SPARC64 X+, and SPARC64 X leverages this SIMD enhancement to accelerate cryptographic processing and business applications, such as a database.

In addition, load and store instructions can use SIMD. For the load instruction especially, processing is done without any penalties even if it steps over the cache line.

By utilizing this function for massive data retrieval and data compression/decompression (encryption/decryption) loads, faster processing of massive data or in-memory databases is achieved.

2. SWoC (Software on Chip)

In the SPARC64 XII, SPARC64 X+, and SPARC64 X, multiple instruction combinations can be replaced by dedicated hardware, leading to much faster computing. This is called "SWoC" (Software on Chip), which accelerates processing as follows.

- Cryptographic Processing

The SPARC64 XII, SPARC64 X+, and SPARC64 X implement cryptographic arithmetic units. The SPARC64 X+ and SPARC64 X implement two cryptographic arithmetic units per core, and the SPARC64 XII implements two per instruction pipeline or four per core.

A dedicated cryptographic instruction executes cryptographic processing (encryption/decryption) at high speed without using an add-on adapter. The supported encryption modes are AES, DES, 3DES, DH, DSA, ECC (*1), RSA, and SHA.

Since the encryption processing is executed by hardware, additional cost and performance deterioration are avoided. Also full database encryption allows the construction of a secure environment.

In addition, OpenSSL and encrypt/decrypt commands are already compliant with the SPARC64 XII, SPARC64 X+, and SPARC64 X cryptographic arithmetic processing functions. Standard libraries such as libpkcs11 also benefit from the cryptographic arithmetic processing enhancements.

*1 For the SPARC64 XII only

- **IEEE 754 for Decimal Floating-Point Numeric Calculations in DPD Format**

The SPARC64 XII, SPARC64 X+, and SPARC64 X inherit this high-speed computing technique from the mainframe.

A decimal floating-point arithmetic unit has been added to execute decimal floating-point arithmetic processing (which conventionally has been done by software) directly by hardware at high speeds.

Conventionally, decimal floating-point arithmetic processing is processed done by software. In the software processing, software converts decimal number data to binary, and then hardware performs arithmetic processing. Computational results are then converted from binary back to decimal number data by software.

In the SPARC64 XII, SPARC64 X+, and SPARC64 X processors, arithmetic processing of the decimal number data can be performed as is by hardware without the above conversion by software. This feature of the processors will accelerate computing across many industries, such as distribution, manufacturing, and finance, by significantly accelerating common processing tasks like sales accounting, cost accounting, rebate accounting, and compound interest calculation.

This SPARC64 XII, SPARC64 X+, and SPARC64 X calculation feature complies with the standard for decimal floating-point arithmetic processing, IEEE 754-2008.

- **Compare and Copy Operations**

With the SPARC64 XII, SPARC64 X+, or SPARC64 X, the width of the memory access bus has been increased. Extended SIMD instructions have also been implemented, allowing for multiple blocks of data to be loaded from memory, compared, and then stored in memory. Extended SIMD instructions allow for 100% utilization of the memory bus capacity, maximizing the computational power of the SPARC64 XII, SPARC64 X+, and SPARC64 X. In this way, memory processing functions such as memcpy(3C) in the standard libc library are accelerated automatically.

Processing that conventionally is executed by application software is now processed by SPARC64 XII, SPARC64 X+, or SPARC64 X hardware, leading directly to improved application performance. Oracle Solaris 11 has been optimized to take full advantage of the SWoC functions in the SPARC64 XII, SPARC64 X+, and SPARC64 X. This translates to performance gains without requiring changes to the applications.

Reliability, Availability, and Serviceability Features

The SPARC64 XII, SPARC64 X+, and SPARC64 X processors increase system reliability by delivering advanced error detection and correction capabilities. In fact, 99% of the SPARC64

XII, SPARC64 X+, and SPARC64 X processor circuitry is protected by error detection and/or data correction mechanisms. RAM units are ECC protected or duplicated, and all 1-bit errors can be corrected. In addition, all latches and execution units are parity protected, and when any of the processors detects a 1-bit error, it retries the instruction. A cache is dynamically degraded either in units of cache-way or units of CPU core, and isolation after reboot is performed for some types of errors. Other reliability features of the SPARC64 XII, SPARC64 X+, and SPARC64 X processors include support for error marking, instruction retry (re-execution) by hardware, and preventive maintenance, which are mentioned below.

As described above, in the SPARC64 XII, SPARC64 X+, and SPARC64 X, RAS functions comparable to mainframe computers have been implemented. With these RAS functions, errors are reliably detected, their impacts are kept within a limited range, recovery processing is attempted, error logs are recorded and reported to software, and so forth. Throughout SPARC64 XII, SPARC64 X+, and SPARC64 X development, RAS functions have been robustly implemented to provide high reliability, high availability, high serviceability, and high data integrity, making the processors a perfect match for mission-critical UNIX servers.

Table 7-2. Error Detection Mechanisms

Unit	Error Detection and Correction Method
Cache (tag)	ECC Duplication + Parity
Cache (data)	ECC Parity
Register	ECC (integer/floating-point) Parity (others)
Execution Unit	Parity/Residue check
Cache dynamic degradation	Implemented
Hardware instruction retry	Implemented
Processor history logging	Implemented




Figure 7-5. RAS Coverage Maps

Unit	Error Detection/ Correction	Error Handling		Logging
		Correction	Deconfiguration	
L1 cache (tag)	Duplication + Parity, ECC	Retry, ECC	Dynamic way deconfiguration (*1)	History
L2 cache (data)	ECC	ECC	Dynamic way deconfiguration (*1)	
L3 cache (data) (*2)	ECC	ECC	Dynamic way deconfiguration (*1)	
Execution Unit	Parity (*3), ECC	ECC, hardware instruction retry	Core deconfiguration	
Register	Parity (*3), ECC	ECC, hardware instruction retry	Core deconfiguration	

*1 "Way" is a unit of cache.

*2 For the SPARC64 XII only

*3 The processor recovers from parity errors by using hardware instruction retry.




Figure 7-6. Error Detection and Recovery Scheme

1. Error Marking

When data read from the main memory or cache is found to contain a multi-bit error, a special value is written into the data by the error marking function to identify the location of the error. In combination with ECC syndrome reporting, the error marking function identifies the original error location that caused the failure, thus aiding accurate degradation and parts replacement.

2. Internal RAM Reliability and Availability Features

The SPARC64 XII, SPARC64 X+, and SPARC64 X processors offer reliability and availability features that support high levels of data integrity. Table 7-3 shows the error detection and correction capabilities of the processors.

Table 7-3. Error Protection of SPARC64 XII, SPARC64 X+, and SPARC64 X Internal RAM

RAM Type	Error Detection/Protection Method	Error Correction Method
L1 Instruction Cache Data	Parity	Invalidation and reread
L1 Instruction Cache Tag	Parity + Duplication	Reread of duplicated data
L1 Data Cache Data	SECDED (*1) ECC	1-bit error correction using ECC
L1 Data Cache Tag	Parity + Duplication	Reread of duplicated data
L2 Cache Data	SECDED (*1) ECC	1-bit error correction using ECC
L2 Cache Tag	SECDED (*1) ECC	1-bit error correction using ECC
L3 Cache Data (*2)	SECDED (*1) ECC	1-bit error correction using ECC
L3 Cache Tag (*2)	SECDED (*1) ECC	1-bit error correction using ECC
Instruction TLB	Parity	Invalidation
Data TLB	Parity	Invalidation
Branch History	Parity	Recovery from branch prediction failure

*1 SECDED: Single Error Detect Double Error Detect

*2 For the SPARC64 XII only

For the L1 cache, L2 cache, and TLB, dynamic degradation can be performed in units of ways. The SPARC64 XII, SPARC64 X+, and SPARC64 X processors employ a set-associative scheme that divides the L1 cache, L2 cache, and TLB into way units. Error occurrence counts are tabulated for each way unit. When the error occurrence count exceeds the upper threshold, degradation is performed and the relevant way is taken out of service. Hardware performs the

dynamic degradation of the way, and software is unaffected except for a decrease in processing speed due to the degradation of the way.

3. Internal Register and Execution Unit Reliability Features

To further increase reliability, the SPARC64 XII, SPARC64 X+, and SPARC64 X processors also provide error protection for registers and execution units. The general purpose registers (GPRs) and the floating-point registers (FPRs) have been enlarged and include ECC protection. When an error occurs, the ECC circuit corrects the error. Other registers are protected by parity. Parity prediction circuitry is implemented within execution units. In multiplication and division execution units a residue check circuit has been implemented to perform further error detection in output results. In the unlikely event of an error, when it is detected, the hardware automatically re-executes the instruction to attempt recovery, as described below.

4. Synchronous Update Method and Instruction Retry

The SPARC64 XII, SPARC64 X+, and SPARC64 X processors employ a synchronous update method at the time of commit. Only instructions that are committed without detecting an error will update the execution results in programmable resources such as a GPR, a FPR, another register, or memory. When an error is detected, all the processing that occurred before the commit is canceled, preventing data integrity issues from occurring. After an error is detected and the processing canceled, hardware can also re-execute the instruction. This is called instruction retry.

As shown in Figure 7-7, instruction retry is triggered by an error and begins automatically. The retry is performed instruction-by-instruction to increase the chance of successful execution. When the execution completes normally, the state automatically returns to the normal execution state. During this period, no software intervention is required. If the instruction retry succeeds, the error has no effect on software. An instruction retry is repeated until the number of retry attempts reaches the threshold. If the threshold is exceeded, the processor logs the source of the error and notifies the operating system of the error occurrence for subsequent operating system processing.




Figure 7-7. Instruction Retry

5. Increased Serviceability

The SPARC64 XII, SPARC64 X+, and SPARC64 X processors, as stated above, provide a variety of error checking mechanisms. Each processor monitors for errors, and sends information to the eXtended System Control Facility (XSCF) if an error occurs. On receipt of this notification, the XSCF firmware collects and analyzes the error log. By taking advantage of extensive error notification features, the Fujitsu SPARC M12 and Fujitsu M10 can identify the location and type of fault quickly and accurately while continuing operation. The system then provides information useful for preventive maintenance, increasing serviceability.

8. Conclusion

As the sheer volume of data explodes in this era of big data and clouds, customers from across the processing spectrum require new levels of scalability and flexibility, reliability and manageability, and above all, high performance. On all counts, the Fujitsu SPARC M12 and Fujitsu M10 deliver on these needs.

The product of a long, stable history in high-performance supercomputing and mission-critical processing, Fujitsu's twelfth-generation SPARC64 XII processor and tenth-generation SPARC64 X+ and SPARC64 X processors provide superior performance and a breakthrough new concept in business computing: Software on Chip. This architecture takes a bold step forward to accelerate key database functions and foreshadow a new blurring of the line between hardware and software.

The Fujitsu SPARC M12 and Fujitsu M10 wrap the processors in highly scalable and flexible systems that boast a unique Building Block architecture, full support of Oracle VM for SPARC, and further flexibility from Oracle Solaris Zones. These technologies, together with per-core activation of resources, allow new levels of granularity in how the Fujitsu SPARC M12 and Fujitsu M10 grow with you. The mainframe-class RAS functionality that Fujitsu prides itself on is found throughout the Fujitsu SPARC M12 and Fujitsu M10 and extends to new levels of protection. Scaling from the entry-priced, but mid-range performing, the Fujitsu M10-1 or Fujitsu SPARC M12-1 single-CPU model all the way to the largest stack of the Fujitsu M10-4S 16BB (64 CPUs) or the Fujitsu SPARC M12-2S 16BB (32 CPUs), the Fujitsu SPARC M12 and Fujitsu M10 deliver a consistent, easy-to-use management interface in the form of the XSCF.

Advanced technical features like Liquid Loop Cooling (LLC) or the more recent Vapor and Liquid Loop Cooling (VLLC), a tradition of reliability, world-beating performance, and a trusted partnership for hardware-software collaboration all combine to make both the Fujitsu SPARC M12 and Fujitsu M10 a great solution for you, our customers and partners, for your ever-expanding needs and growth.

Contact

FUJITSU LIMITED

Shiodome City Center, 5-2, Higashi-shimbashi 1-chome, Minato-ku

Tokyo 105-7123 Japan

Website: <https://www.fujitsu.com/global/products/computing/servers/unix/sparc/>