

# Oracle's SPARC S7 Server Architecture

## Secure Platform for Scale-Out and Cloud Infrastructures

ORACLE TECHNICAL WHITE PAPER | JULY 2016





## Table of Contents

Introduction	1
SPARC S7 Servers	3
Comparison of Features	4
SPARC S7 Processor	6
SPARC S7 Processor Architecture	7
Processor Core and Cache Architecture	8
Software in Silicon Technology	9
An Open Platform	9
SPARC S7 Server Overview	10
Memory Subsystem	10
I/O Subsystem	10
NVM Express Technology	10
Embedded USB Storage and Oracle Solaris Boot Pool	11
PCIe Adapter Cards	11
SPARC S7 Server Details	12
SPARC S7-2 Server	12
SPARC S7-2L Server	15
Eight 2.5-Inch SFF Disk Drive Bays	16
Twenty-Four Front and Two Rear 2.5-Inch SFF Disk Drive Bays	16
Twelve Front 2.5-inch SFF NVMe Disk Drive Bays	17
Twelve Front 3.5-Inch LFF and Two Rear 2.5-Inch SFF Disk Drive Bays	17
Oracle Solaris	18

Virtualization	20
Systems Management	21
Oracle ILOM and Service Processors	21
Power Management	22
Oracle Enterprise Manager Ops Center	22
Reliability, Availability, and Serviceability	23
Reliability, Availability, and Serviceability Features	23
Advanced Reliability and Availability Features	24
Error Detection, Diagnosis, and Recovery	25
Redundant and Hot-Serviceable Components	25
Conclusion	26
For More Information	27




## Introduction

In today's IT infrastructure, data security can no longer be treated as an afterthought, because billions of dollars are lost each year to computer intrusions and data exposures. This issue is compounded by the aggressive build-out for cloud computing. Big data and machine learning applications that perform tasks such as fraud and intrusion detection, trend detection, and click-stream and social media analysis all require forward-thinking solutions and enough compute power to deliver the performance required in a rapidly evolving digital marketplace. Companies increasingly need to drive the speed of business up, and organizations need to support their customers with real-time data. The task of managing sensitive information while capturing, analyzing, and acting upon massive volumes of data every hour of every day has become critical.

These challenges have dramatically changed the way that IT systems are architected, provisioned, and run compared to the past few decades. Most companies have strategic plans for deploying clouds—on- or off-premises—to run distributed Java applications and databases. However, IT infrastructures that are underutilized, vulnerable, and complex pose a problem by becoming more difficult and expensive to maintain at precisely the moment that organizations are under pressure to increase operating efficiencies, drive down costs, and deliver innovative technologies that can generate new revenue streams.

Oracle's servers that are based on Oracle's SPARC S7 processor extend the technology in Oracle's SPARC T7 and SPARC M7 servers to optimally address the needs of scale-out and cloud infrastructures. SPARC S7 servers leverage Oracle's revolutionary Software in Silicon technology to build the world's most advanced platform for secure computing. The unique hardware and software features built into the SPARC platform for protecting data and enforcing compliance are simpler to implement, are more cost effective, and go far beyond what other solutions in the marketplace can provide. The optimized, balanced design of the SPARC S7 processor and servers provide up to 1.7x better core efficiency than x86 systems for running Java applications and databases, which lowers



deployment and operational costs. Hardware acceleration of data analytics, big data, and machine learning deliver 10x faster time-to-insight and off-load processor cores to run other workloads, enabling real-time analytics to be performed together with on-line transaction processing.

The combination of breakthrough Software in Silicon features and the highest performance in the SPARC platform is the foundation for building the most secure and efficient enterprise clouds.

## SPARC S7 Servers

Oracle's SPARC S7 processor-based servers offer a superior end-to-end security solution. Common hacker exploits and programming errors can be stopped by the Silicon Secured Memory feature of the SPARC S7 processor. Encryption of data at rest and in motion can now be enabled without compromise in performance, through the use of wide-key encryption accelerators in the SPARC S7 processor. Potential intruders cannot create a foothold in SPARC systems because Oracle Solaris offers verified boot from trusted repositories, and unauthorized changes are prevented through immutable zones and virtual machines.

The SPARC S7 processor delivers balanced compute performance with 8 cores per processor, integrated on-chip DDR4 memory interfaces, a PCIe controller, and coherency links. The cores in the SPARC S7 processor are optimized for running key enterprise software, including Java applications and database. The SPARC S7 processor-based servers use very high levels of integration that increase bandwidth, reduce latencies, simplify board design, reduce the number of components, and increase reliability. All this leads to a dramatic increase in system efficiency with a corresponding improvement in the economics of deploying a scale-out infrastructure when compared to other vendor solutions.

Using the same Data Analytics Accelerator (DAX) feature introduced by the SPARC M7 processor, the SPARC S7 processor and servers extend this breakthrough technology into scale-out architectures. The DAX engines accelerates analytics functions, such as scans, filters and joins, by running those in specialized units connected directly to memory interfaces. Applications such as big data, machine learning, risk and fraud analysis, which operate on large data sets in-memory, can run at memory speeds. The SPARC S7 processor offloads data analytics features to the DAX engines, freeing processor cores to run other workloads, dramatically increasing the efficiency of running mixed workloads on the SPARC S7 servers.

SPARC S7 processor-based servers (Figure 1) include single- and dual-processor systems that are complementary to the existing mid-range and high-end systems based on Oracle's SPARC M7 processor. SPARC S7 processor-based servers include two rack-mountable models. The SPARC S7-2 server uses a compact 1U chassis, and the SPARC S7-2L server is implemented in a larger, more expandable 2U chassis. Uniformity of management interfaces and the adoption of standards also help reduce administrative costs, while innovative chassis design provides density, efficiency, and economy for modern data centers.




Figure 1. Oracle's SPARC S7 processor-based server product family.

## Comparison of Features

Table 1 provides a feature comparison of the SPARC S7-2 and S7-2L servers.

**TABLE 1. SPARC S7 PROCESSOR-BASED SERVER FEATURES COMPARISON.**

Feature	SPARC S7-2 Server	SPARC S7-2L Server
Form factor	1U	2U
Height	43 mm / 1.7 in.	88 mm / 3.5 in.
Width	437 mm / 17.2 in.	436 mm / 17.2 in.
Depth	737 mm / 29 in.	737 mm / 29 in.
Maximum weight	19.5 kg / 43 lb.	30 kg / 66 lb.
Processor	8-core 4.27 GHz SPARC S7 processor, up to 64 threads per processor 16 MB Level 3 cache, fully shared and partitioned, 8 MB per core cluster Silicon Secured Memory 16 Data Analytics Accelerator (DAX) engines (query pipes) for In-Memory Query Acceleration and In-Line Decompression Encryption instruction accelerators in each core with direct support for 15 industry-standard cryptographic algorithms plus random number generation: AES, Camellia, CRC32c, DES, 3DES, DH, DSA, ECC, MD5, RSA, SHA-1, SHA-224, SHA-256, SHA-384, SHA-512	
Processor quantity	1 or 2	2
Maximum cores	16	
Maximum threads	128	
Memory	16 GB, 32 GB, or 64 GB DDR4-2400 memory DIMMs, 4 or 8 DIMMs per processor	
Memory capacity	Max. 1,024 GB Min. 64 GB	Max. 1,024 GB Min. 128 GB
Internal disk drive bays	8	Up to 26
SAS support for internal drive bays	One internal SAS3 host bus adapter (HBA) supporting up to 8 SAS drive bays	One internal SAS3 HBA supporting up to 26 SAS drive bays
NVM Express (NVMe) support for internal 2.5-inch disk drive bays	Onboard support for up to 4 NVMe solid-state drives (SSDs)	Onboard support for up to 4 NVMe SSDs, or a total of 12 NVMe SSDs with 3 PCIe switch cards
Max. Oracle Flash Accelerator F320 PCIe Cards	3	6
Management ports	One 1000BASE-T Ethernet port (rear) One serial RJ45 port (rear)	
USB ports	Two USB 2.0 (front) ports	
Ethernet	Four integrated 10GBASE-T ports <sup>1</sup> One onboard Ethernet controller	

1. 10GBASE-T auto negotiates to 100 Mb/sec, 1 Gb/sec, and 10 Gb/sec, full-duplex only. Jumbo frames are supported, up to 9,706 bytes.

**TABLE 1. SPARC S7 PROCESSOR–BASED SERVER FEATURES COMPARISON (CONTINUED).**

Feature	SPARC S7-2 Server	SPARC S7-2L Server
PCIe 3.0 expansion slots	3 slots Supported by 1 or 2 PCIe root complexes	6 slots Supported by 2 PCIe root complexes
	Three x8 slots two of which physically support x16 cards	Six x8 slots two of which physically support x16 cards
Hot-swappable power supplies	Two N+1 redundant 800/1200 W AC power supplies <sup>2, 3</sup>	Two N+1 redundant 1200 W AC power supplies <sup>3</sup>
	Voltage: 100–120 VAC or 200–240 VAC Frequency: 50/60 Hz	Voltage: 200–240 VAC Frequency: 50/60 Hz
N+1 redundant hot-swappable fans	4 quad-fan modules, top loading	4 dual-fan modules, top loading
Operating system	Oracle recommends Oracle Solaris 11.3 or later for enhanced performance and functionality, including features enabled by Software in Silicon technology	
	Control, root, and I/O domains: Oracle Solaris 11.3 SRU 8 or later <sup>4</sup>	
	The following versions are supported within guest domains: » Oracle Solaris 11.3 SRU 8 or later <sup>4</sup> » Oracle Solaris 10 1/13 <sup>5</sup>	
	Applications certified for Oracle Solaris 8 or Oracle Solaris 9 only may run in an Oracle Solaris 8 or Oracle Solaris 9 Branded Zone running within an Oracle Solaris 10 guest domain.	

2. Power supply output is 800 watts when operating at 100–120 VAC and 1,200 watts when operating at 200–240 VAC.
3. For power consumption, please refer to the Oracle power calculators available at <http://www.oracle.com/goto/power-calculators>
4. Versions of Oracle Solaris 11 prior to 11.3 SRU 8 are not supported on SPARC S7 processor–based servers.
5. Plus required patches.



## SPARC S7 Processor

Oracle's SPARC S7 processor (Figure 2) is optimized for horizontally scalable systems with all the key functionality included in the microprocessor chip. The revolutionary Software in Silicon capabilities, which were introduced with the SPARC M7 processor, are also available in the SPARC S7 processor enabling unprecedented data protection, cryptographic acceleration, and analytics performance. These features include Security in Silicon, which provides Silicon Secured Memory and cryptographic acceleration, and Data Analytics Accelerator (DAX) units, which provide In-Memory Query Acceleration and In-Line Decompression.




Figure 2. The SPARC S7 processor combines eight cores, on-chip memory interfaces, an I/O controller, and networking along with Software in Silicon features to deliver exceptional security protection and analytics performance.

The Silicon Secured Memory feature of the SPARC S7 processor provides real-time data integrity checking to guard against pointer-related software errors and malware. It replaces very costly software instrumentation with low-overhead hardware monitoring. Silicon Secured Memory enables applications to identify erroneous or unauthorized memory access, diagnose the cause, and take appropriate recovery actions, while running in production.

The SPARC S7 processor has cryptographic instruction accelerators integrated directly into each processor core. These accelerators enable high-speed encryption for over a dozen industry-standard ciphers, eliminating the performance and cost barriers typically associated with secure computing.

The SPARC S7 processor also incorporates hardware units that accelerate specific software functions or primitives. The on-chip DAX units offload data analytics functions and database query processing and perform real-time data decompression. In-Memory Query Acceleration, used by Oracle Database In-Memory feature, delivers performance that is up to ten times faster compared to other processors. The In-Line Decompression feature allows more data to be stored in the same memory footprint, without a performance penalty.

The SPARC S7 processor uses the same core cluster design as the high-end SPARC M7 processor. This applies also to the on-chip Level 2 (L2) and Level 3 (L3) cache implementation. The 16 MB L3 cache is partitioned and fully shared, and hot cache lines are migrated to the closest partition to minimize latency and maximize performance. System administration and performance tuning are easier, because the design minimizes interaction between logical domains or between databases. The processor can dynamically trade per-thread performance for throughput by running up to 64 threads, or it can run fewer higher-performance threads by devoting more resources to each

thread. This flexibility allows the system to balance overall throughput versus per-thread performance for optimal results.

The 8-core SPARC S7 processor is binary-compatible with earlier SPARC implementations. It is ideal for virtualized cloud computing environments, supporting a relatively large number of virtual machines and delivering excellent multithreaded performance. This processor, and the server system based on it, enables organizations to rapidly scale the delivery of new network services with maximum efficiency and predictability.

Table 2 provides a comparison between Oracle's SPARC S7 and SPARC M7 processors.

**TABLE 2. SPARC S7 AND SPARC M7 PROCESSOR FEATURE COMPARISON.**

Feature	SPARC S7 Processor	SPARC M7 Processor
CPU frequency	4.27 GHz	4.13 GHz
Out-of-order execution	Yes	Yes
Dual-instruction issue	Yes	Yes
Data/instruction prefetch	Yes	Yes
Cores per processor	8	32
Threads per core	8	8
Threads per processor	64	256
Sockets in systems	1 or 2	Up to 16
Memory per processor	Up to 8 DDR4 DIMMs	Up to 16 DDR4 DIMMs
Memory channels per processor	4	16
Caches	16 KB Level 1 (L1) four-way instruction cache (per core) 16 KB L1 four-way data cache (per core) Shared 256 KB L2 four-way instruction cache (per quad cores) Shared 256 KB L2 eight-way data cache (per core pair) Shared 16 MB (L3) cache	16 KB L1 four-way instruction cache (per core) 16 KB L1 four-way data cache (per core) Shared 256 KB L2 four-way instruction cache (per quad cores) Shared 256 KB L2 eight-way data cache (per core pair) Shared 64 MB (L3) cache
Large page support <sup>1</sup>	16 GB	16 GB
Power management granularity	Entire chip	¼ of chip
Technology	20 nm technology	20 nm technology

1. Large page support with Oracle Solaris 11.3

## SPARC S7 Processor Architecture

Oracle's SPARC S7 processor is designed for scale-out systems with up to two processor sockets. The processor is optimized to include balanced compute capacity and integrated system functions, which eliminates many separate system components such as memory buffer-on-board chips and supporting ASICs.

In addition to the processor cores, the SPARC S7 processor has on-chip DDR4 memory interfaces and a PCIe 3.0 controller. In order to deliver commercial workloads with high levels of throughput, the SPARC S7 processor implements the same cache and memory hierarchy as the SPARC M7 processor. Power management capabilities are also included so dynamic voltage frequency scaling (DVFS) is provided.

Figure 3 illustrates the architecture of the SPARC S7 processor. The processor contains eight cores that are grouped into two core clusters. Two memory controller units (MCUs) are provided, each with two channels connecting directly to DDR4 memory DIMMs. A total of up to eight memory DIMMs is supported per SPARC S7 processor. Two DAX units are associated with each MCU and provide SQL in Silicon functionality. Each of the four total DAX units has 4 pipelines, for a total of 16 Data Analytics Accelerator (DAX) engines (query pipes).

No additional logic is needed to form a single two-way glueless symmetrical multiprocessing (SMP) system. There are four coherency links (CL) per processor for connectivity and coherency between the two SPARC S7 processors in a server. A high-bandwidth, low-latency, on-chip switched data network connects the two L3 cache partitions to each other, to the two MCUs, and to the I/O and coherence gateway. The network maintains coherency both on-chip and off-chip. A PCIe 3.0 root complex is integrated into the processor with a bandwidth of x16 providing connectivity to the I/O devices and networking.

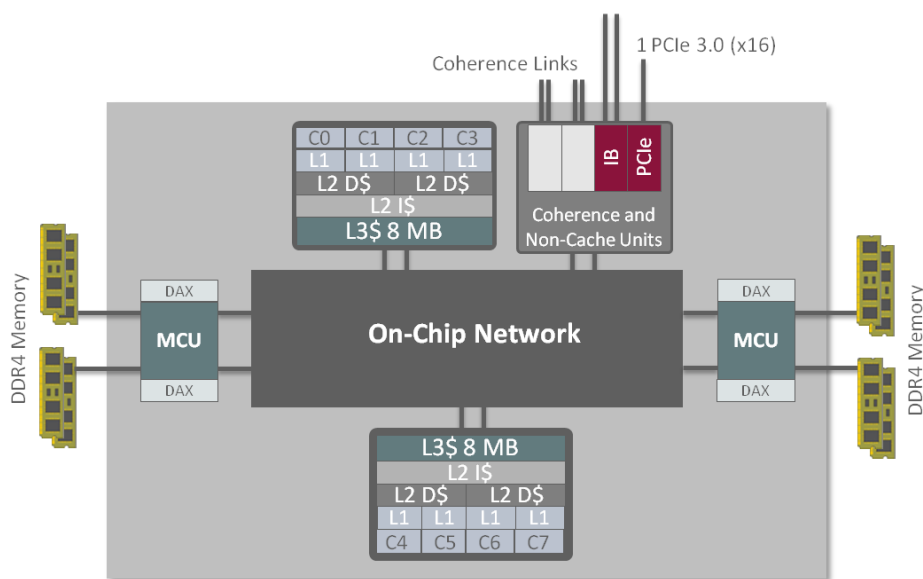


Figure 3. The SPARC S7 processor features eight cores, two core clusters, two MCUs, four DAX units, and a PCIe controller.

## Processor Core and Cache Architecture

The SPARC S7 processor core is dual-issue, out-of-order, and supports up to eight hardware threads. The core provides dynamic threading to optimize for the highest possible per-thread performance. Software can activate up to eight hardware threads (strands) on each core via [Critical Threads Optimization](#). The processor hardware then dynamically and seamlessly allocates core resources among the active strands.

The memory management unit (MMU) in the core provides a hardware table walk (HWTW) and supports 8 KB, 64 KB, 4 MB, 256 MB, 2 GB, and 16 GB pages.

In the SPARC S7 processor, four cores are combined into a core cluster with two core clusters per SPARC S7 processor. Within the SPARC core cluster, each core has its own 16 KB L1 instruction and data cache. Two cores then share a 256 KB L2 data cache with the four cores sharing a 256 KB L2 instruction cache. The L3 cache is fully shared and partitioned. The L3 partition is eight-way set-associative with a 64-byte line size, and is composed of two address-interleaved banks. Any L3 partition may serve a request from any of the eight cores of the SPARC S7 processor. Hot cache lines are migrated to the closest L3 cache partition to optimize performance.



## Software in Silicon Technology

The SPARC S7 processor takes full advantage of Oracle's ability to move in-memory data analytics functions directly onto the processor chip with hard-wired protection for data in memory. By innovating at the processor, system, and application levels, Oracle is in a unique position to optimize application performance through this approach. The SPARC S7 processor capitalizes on the same Software in Silicon functionality that is integrated into the SPARC M7 processor.

The SPARC S7 processor incorporates on-chip accelerators to offload in-memory analytics and database query processing and perform real-time data decompression, and cryptographic instruction accelerators are integrated directly into each processor core. Together, the following Software in Silicon features deliver significant advantages for security, performance, and efficiency:

- » Security in Silicon
  - » Silicon Secured Memory provides real-time data integrity checking to guard against pointer-related software errors and malware, replacing very costly software instrumentation with low-overhead hardware monitoring. Silicon Secured Memory enables applications to identify erroneous or unauthorized memory access, diagnose the cause, and take appropriate recovery actions.
  - » Accelerated cryptography helps eliminate the performance and cost barriers typically associated with secure computing—which is increasingly essential for modern business operations.
- » Data Analytics Acceleration
  - » Analytics and In-Memory Query Acceleration provided by DAX units delivers performance that is up to ten times faster compared to other processors.
  - » The In-Line Decompression feature enables storing up to two times more data in the same memory footprint, without a performance penalty.

In addition to the crypto instruction accelerators that are included in every core, the SPARC S7 processor contains four DAX units, each with four pipelines, or engines. These engines can process a total of 16 independent data streams, offloading the processor cores so they can do other work. The DAX engines can process analytics functions such as decompress, scan, filter, and join.

The DAX engines use very low-overhead interprocess communication and extremely fast atomic operations. For example, DAX engines located on different processors can exchange messages and access remote memory locations, exchanging locks without CPU involvement. Utilizing this functionality requires Oracle Database 12c with the In-Memory option and Oracle Solaris 11.3 or later.

### An Open Platform

Existing applications can be enabled to use Silicon Secured Memory—without recompiling—by linking with the correct Oracle Solaris libraries and being verified in a test environment. Open Oracle Solaris APIs are available for software developers to leverage Silicon Secured Memory and DAX technologies.

For more information on the Software in Silicon features, please refer to the “Get Informed” section of the [Software in Silicon \(Sample Code and Resources\)](#) web page.

## SPARC S7 Server Overview

The SPARC S7-2 and SPARC S7-2L servers are designed for scale-out and cloud infrastructures that required high levels of security, performance, and efficiency. These SPARC servers are ideal for database, Java, middleware, and enterprise applications, and they offer exceptional throughput performance and memory bandwidth. These servers include one or two processors and are complementary to Oracle's larger SPARC T7-1, T7-2, T7-4, M7-8, and M7-16 servers, which scale from one to 16 SPARC M7 processors.

Common hardware features of the servers include the following:

- » SPARC S7 8-core 4.27 GHz processor with Software in Silicon features and system-on-a-chip capabilities
- » 16 GB, 32 GB, and 64 GB DDR4-2400 memory DIMMs
- » Integrated support for internal SAS3 and NVMe Express (NVMe) disk drives
- » Embedded USB (eUSB) storage device that supports booting over the InfiniBand network
- » PCIe 3.0 x8 expansion slots

### Memory Subsystem

Each SPARC S7 processor supports up to 8 DDR4 memory DIMMs connected directly to the on-chip memory controllers with four memory channels. Up to 512 GB of memory is supported per processor with eight 64 GB DIMMs. Raw memory bandwidth is 76.8 GB/sec per SPARC S7 processor. Half-populated and fully populated memory configurations are supported, and they both offer similar memory bandwidth because all four memory channels are used in both cases.

The physical address space provided by the memory DIMMs and controlled by an individual SPARC S7 processor is interleaved to maximize performance. Both half-populated and fully populated memory configurations are four-way interleaved. The SPARC S7 processor also supports a three-way interleaved configuration, which is used if a DIMM or a channel has been deemed unreliable and has been isolated from the active system. The switch from a four-way to three-way configuration is done while the system is offline and requires a reboot.


### I/O Subsystem

The I/O subsystem in the SPARC S7-2 and SPARC S7-2L servers is based on on-chip I/O devices and controllers in the processor. The SPARC S7 processor provides the PCIe bus and communication to the system service processor. The PCIe 3.0 controller in each SPARC S7 processor provides a single x16 root complex, which is deployed using two x8 data ports.

In SPARC S7-2 and SPARC S7-2L servers, two onboard PCIe switches are cross-connected to the on-chip PCIe controllers. Therefore, all PCIe expansion slots and other I/O devices are available in both single- and dual-processor configurations. Dual-processor systems support redundant I/O domains using separate PCIe root complexes. This makes it possible to configure redundant network and storage connections to support virtualized guest domains with I/O and network multipathing. Four integrated 10GBASE-T Ethernet ports are supported by a single PCIe root complex. Redundant network connections are provided with additional Ethernet network adapter cards that are supported by the other root domain. For information on the I/O connectivity to the disk drive bays, please refer to the details for each server provided later in this document.

### NVM Express Technology

SPARC S7 processor-based servers support the emerging flash storage technology known as NVM Express (NVMe). The NVMe specification defines an optimized PCIe-based interface for solid-state drives (SSDs). Utilizing non-volatile memory, NVMe-based SSDs provide both lower latency and better throughput performance relative to



SAS- or SATA-based SSDs. NVMe utilizes PCIe signaling and provides an 8 GT/sec x4 interface per drive, yielding approximately 4 GB/sec full-duplex throughput to the drive.

The SPARC S7 processor-based servers support both 2.5-inch small form factor (SFF) NVMe SSDs and the Oracle Flash Accelerator F320 PCIe Card—an NVMe-based SSD device on a low-profile PCIe card. The support for the NVMe devices varies by model, but most models have at least some NVMe-enabled disk drive bays. The only exception is the SPARC S7-2L server configuration with the 3.5-inch disk drive cage, which supports only SAS drives. For more details, please refer to the server-specific sections later in this document.

NVMe devices are hot-plug capable when the OS-specific hot-plug procedures are followed. An `nvmeadm` command is provided that lets administrators list the drive health and firmware level, check temperatures, get error logs, and access self-monitoring, analysis, and reporting technology (SMART) data, as well as conduct a security data-erasure procedure and perform low-level formatting.

### Embedded USB Storage and Oracle Solaris Boot Pool

SPARC S7 processor-based servers support a boot process that allows booting over InfiniBand. The conventional boot process requires that the boot devices be accessible by the system firmware. Therefore, a conventional network boot over InfiniBand cannot be supported.

The Oracle Solaris boot process includes the concept called a *boot pool*, which stores boot archives on a local device. In SPARC S7 processor-based servers, an embedded USB (eUSB) flash memory device is used to form the boot pool, which is accessible to the OpenBoot PROM firmware. The eUSB device is installed into the system at the factory and is dedicated to the boot process only.

The existence of the local boot pool allows the OpenBoot PROM firmware to load the boot archive and subsequently mount the root file system in the root pool using iSCSI over IP while using the InfiniBand network. SPARC S7 processor-based servers also provide a fallback mechanism for this boot process, which can be used when the eUSB-based boot archives are not available. An alternate boot archive exists in the flash memory of the system service processor (SP). This boot archive is loaded onto the SP at the factory and is intended to be used only when other means of booting are not available.

### PCIe Adapter Cards

SPARC S7-2 and SPARC S7-2L servers feature PCIe 3.0 expansion card slots, which are wired for x8 operation. Select card slots physically support x16 cards, which will operate in up to x8 mode while in these systems. Supported options and requirements vary by server models. At release, available adapter cards from Oracle include the following:

- » Oracle Storage 12 Gb/s SAS PCIe HBA, external
- » Oracle's Sun Storage 16 Gb FC PCIe Universal HBA
- » Oracle's Sun Dual Port GbE PCIe 2.0 Low Profile Adapter, MMF
- » Oracle's Sun Dual Port 10GBASE-T Adapter
- » Oracle Quad 10Gb or Dual 40Gb Ethernet Adapter
- » Oracle Dual Port QDR InfiniBand Adapter M3
- » Oracle Dual Port QDR InfiniBand Adapter M4
- » Oracle Flash Accelerator F320 PCIe Card

## SPARC S7 Server Details

The servers based on the SPARC S7 processor are designed to provide breakthrough security and performance while maximizing reliability and minimizing power consumption and complexity. These systems are ideal for horizontally scaled applications for which high availability is provided by replicating systems with applications that support failover. With common, shared components and subsystems, a separate design is used in each of the SPARC S7 processor-based servers to optimize each system for its particular design point and capacities. The servers feature a robust chassis, component, and subassembly design; enhanced system and component serviceability; and minimized cabling for maximized airflow.

- » The SPARC S7-2 server is the entry model with one or two processors in a compact 1U enclosure. However, with up to 64 hardware threads, and 1,024 GB of memory, it has formidable compute capacity. It features the full suite of Software in Silicon capabilities as the SPARC M7 processor-based servers, including secure operation with Silicon Secured Memory, in-core cryptographic acceleration, and extreme performance with the SQL in Silicon features.
- » The SPARC S7-2L server offers the same compute and memory capacity as its smaller sibling. However, the 2U enclosure offers twice the number of PCIe 3.0 expansion slots and multiple options for internal storage drives. Up to 96 TB of internal storage with 3.5-inch SAS HDDs or up to 38 TB of NVMe SSD high-performance flash storage is supported. This server supports the same Software in Silicon capabilities as the SPARC S7-2 server.

### SPARC S7-2 Server

The SPARC S7-2 server is the entry model and features single or dual SPARC S7 processors in a 1U enclosure. Standard features include 8 memory DIMM slots per processor—yielding up to 1,024 GB of system memory. The server includes three low-profile PCIe 3.0 expansion slots in the rear of the system. Figure 4 illustrates the front and rear perspectives of the SPARC S7-2 server.

The eight 2.5-inch small form factor (SFF) front-loading drive bays are all supported by an internal 12 Gb/sec SAS controller. As a standard feature, four of the disk drive bays are NVMe-enabled by the onboard PCIe switch. Mixing SAS and NVMe drives is supported. Other standard features of the SPARC S7-2 server include the following:

- » Four 10GBASE-T Ethernet ports (RJ45) are on the back of the system, supported via one onboard controller. Ports auto negotiate to speeds of 100 Mb/sec, 1 Gb/sec, and 10 Gb/sec, full-duplex only. Jumbo frames, up to 9,706 bytes, are supported.
- » Two hot-swappable N+1 redundant power supply units (PSUs) with an output capacity of 1,200 W at 200–240 VAC (800 W at 100–120 VAC) are inserted from the rear of the system.
- » Four top-loading hot-swappable fan modules are provided, each with two sets of dual counter-rotating fans.
- » Dual USB 2.0 ports are in the front.
- » Serial (RJ45) and network (1000BASE-T) management ports for SP connectivity are in the rear.




Figure 4. SPARC S7-2 server front and rear perspectives.



Figure 5 and figure 6 show block diagrams of the SPARC S7-2 server in a dual-processor and single-processor configuration, respectively. Each SPARC S7 processor connects directly to up to eight onboard memory slots. Each SPARC S7 processor provides one PCIe bus (root complex) and connects to both PCIe switches. In a dual-processor configuration, each PCIe switch supports device connections for both PCIe root complexes. The color of the device number next the switch indicates which root complex and which processor supports the device, respectively.

One of the PCIe switches (switch 1) supports the SAS HBA in the internal PCIe slot 4, the four NVMe-enabled disk drive bays, and one of the PCIe expansion slots (slot 3). The other PCIe switch (switch 0) supports the two remaining PCIe expansion slots (slots 1 and 2), the onboard Ethernet controller, and the USB controller.

The four NVMe-enabled disk drive bays and the three available PCIe expansion slots are supported by a combination of the two PCIe root complexes, when available. In a single-processor configuration there is only one PCIe root complex, which supports all devices and expansion slots.

All PCIe card slots are wired for x8 operation. The PCIe expansion slots 1 and 2 have physical connectors that support PCIe 3.0 x16 cards operating in x8 mode. The internal PCIe slot 4 is occupied by the internal SAS HBA card, which supports all eight SAS disk drive bays.

Separate serial (RJ45) and Ethernet (1000BASE-T, RJ45) management ports are provided to interface with the SP. The onboard 10GBASE-T network ports can also be used (with the sideband feature enabled) to connect at speeds of up to 10 Gb/sec to the SP.

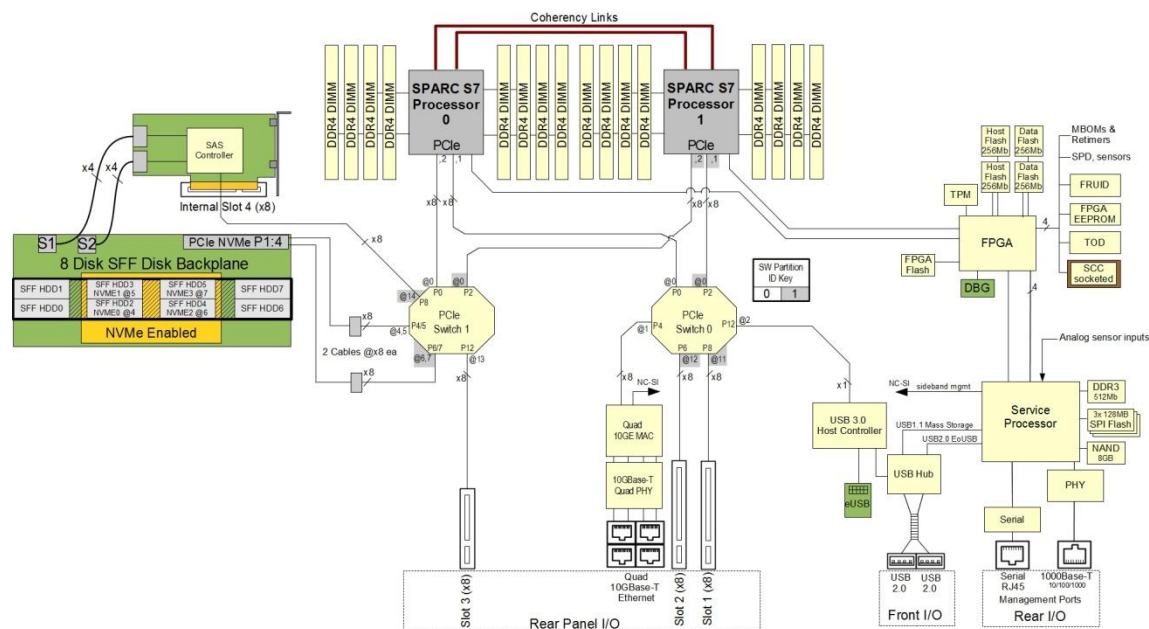


Figure 5. The dual-processor SPARC S7-2 server features cross-over connections to two PCIe switches.



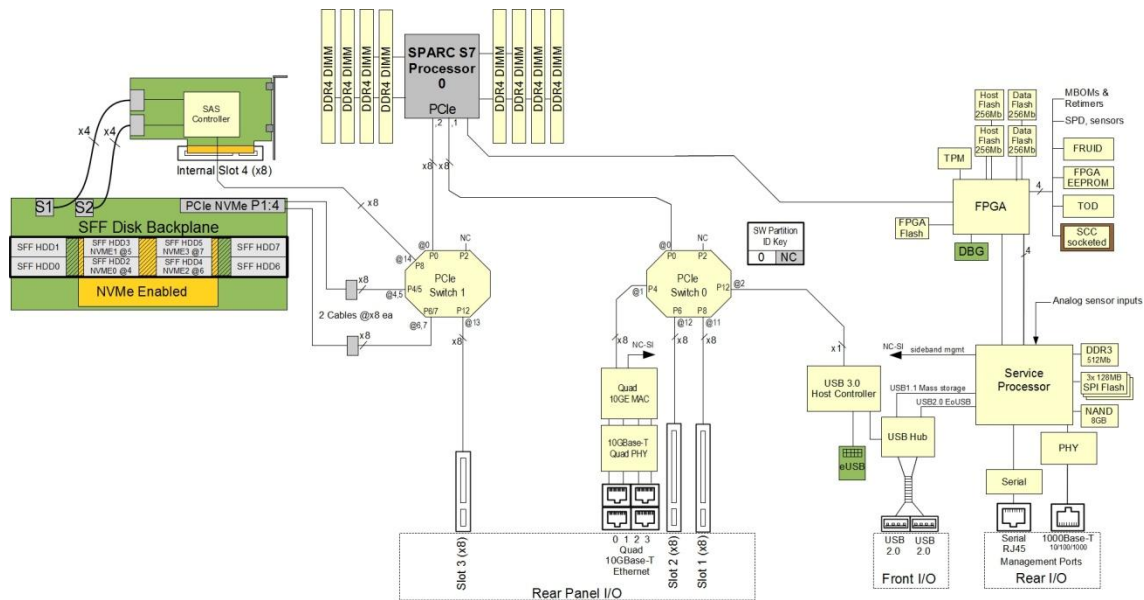


Figure 6. The single-processor SPARC S7-2 server supports all devices with one root complex.

Table 3 describes how the PCIe devices share the root complexes on each of the I/O controllers.

TABLE 3. SPARC S7-2 SERVER ROOT COMPLEX MAPPING.

Processor	PCIe Root Complex	Switch	Device	Target	PCIe Speed
#0	# 0	#0	@1	Quad 10GBASE-T Ethernet ports	x8
			@2	USB controller	x1
		#1	@13	PCIe slot 3	x8
			@4	NVMe disk drive bay #0	x4
#1	# 1	#0	@5	NVMe disk drive bay #1	x4
			@11	PCIe slot 1	x8
		#1	@12	PCIe slot 2	x8
			@14	SAS/SATA I/O controller (PCIe slot 4)	x8
			@6	NVMe disk drive bay #2	x4
			@7	NVMe disk drive bay #3	x4

## SPARC S7-2L Server

The SPARC S7-2L server has similar capabilities as the SPARC S7-2 server but is housed in a larger 2U enclosure with more expandability. Four models are offered with different disk drive options. Standard features include dual SPARC S7 processors with a total of 16 memory DIMM slots—yielding up to 1,024 GB of system memory. The server includes six low-profile PCIe 3.0 expansion slots in the rear of the system. Figure 7 illustrates the front and rear perspectives of the SPARC S7-2L server.




Figure 7. SPARC S7-2L server front and rear perspectives.

Up to 26 disk drive bays are all supported by an internal 12 Gb/sec SAS controller. As a standard feature, select models with 2.5-inch disk drive bays include four drive bays that are NVMe-enabled by the onboard PCIe switch. Mixing SAS and NVMe drives is supported in those models. Other standard features of the SPARC S7-2L server include the following:

- » Four 10GBASE-T Ethernet ports (RJ45) are on the back of the system, supported via one onboard controller. Ports auto negotiate to speeds of 100 Mb/sec, 1 Gb/sec, and 10 Gb/sec, full-duplex only. Jumbo frames, up to 9,706 bytes, are supported.
- » Two hot-swappable N+1 redundant PSUs with an output capacity of 1,200 W at 200–240 VAC are inserted from the rear of the system.
- » Four top-loading hot-swappable fan modules are provided, each with dual counter-rotating fans.
- » Dual USB 2.0 ports are in the front.
- » Serial (RJ45) and network (1000BASE-T) management ports for SP connectivity are in the rear.

Four different models of the SPARC S7-2L server are available. They offer different disk drive configurations as follows:

- » Eight front 2.5-inch small form factor (SFF) SAS disk drive bays, of which four are NVMe-enabled
- » Twenty-four front 2.5-inch SFF SAS disk drive bays, of which four are NVMe-enabled, and two rear 2.5-inch SFF SAS disk drive bays

- » Twelve front 2.5-inch SFF NVMe disk drive bays
- » Twelve front 3.5-inch large form factor (LFF) SAS disk drive bays and two rear 2.5-inch SFF SAS disk drive bays

Figure 8 shows a block diagram of the SPARC S7-2L server with various disk drive cage options. Each SPARC S7 processor connects to eight onboard memory slots. Similar to the SPARC S7-2 server, each SPARC S7 processor provides one PCIe bus (root complex) and connects to both PCIe switches. Each PCIe switch supports device connections for both of the PCIe root complexes.

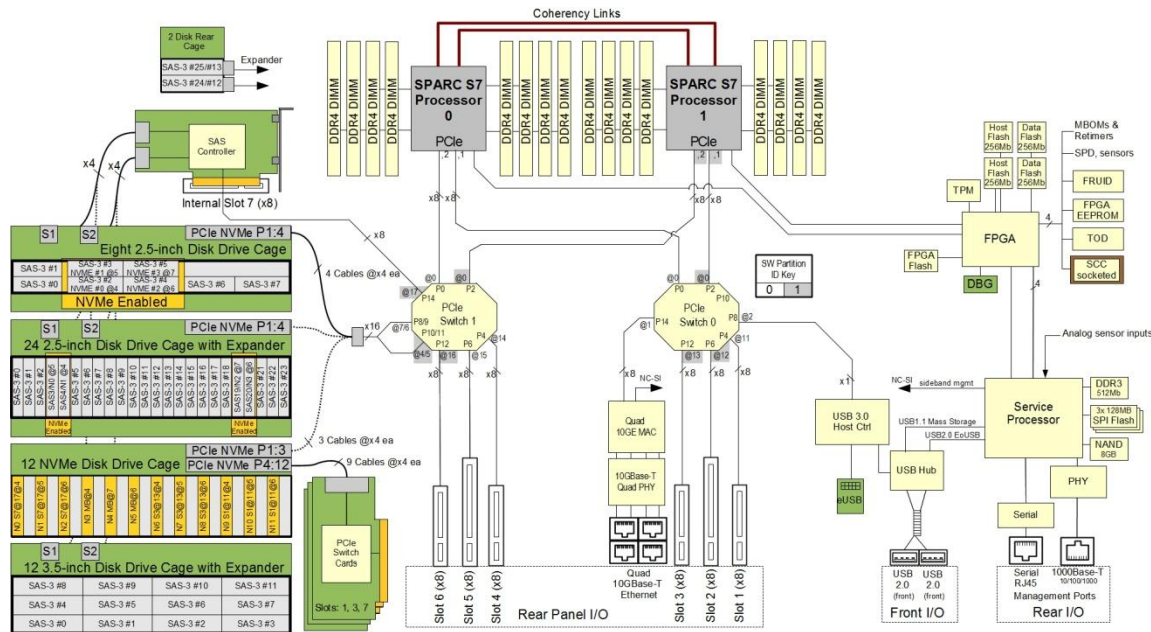


Figure 8. SPARC S7-2L server has dual SPARC S7 processors and a choice of various disk drive cages.

PCIe switch 0 supports half of the PCIe expansion slots (1–3), the onboard Ethernet controller, and the USB controller. PCIe switch 1 supports the internal PCIe slot 7, three or four NVMe-enabled disk drive bays, and the other half of the PCIe expansion slots (4–6).

All PCIe card slots are wired for x8 operation. The PCIe expansion slots 2 and 5 have physical connectors that support PCIe 3.0 x16 cards operating in x8 mode. The internal PCIe slot 7 is occupied either by a SAS HBA or an NVMe PCIe switch card that is installed at the factory.

Separate serial (RJ45) and Ethernet (1000BASE-T, RJ45) management ports are provided to interface with the SP. The onboard 10GBASE-T network ports can also be used (with the sideband feature enabled) to connect at speeds of up to 10 Gb/sec to the SP.

The connections to the internal disk drives go through PCIe switch 1 and vary depending on model. For example, the internal PCIe slot 7 is used either by a SAS HBA or an NVMe PCIe switch card. Following is a brief description of the disk subsystems of each SPARC S7-2L server model:

#### Eight 2.5-Inch SFF Disk Drive Bays

The SAS HBA card in the internal PCIe slot 7 supports all eight SAS disk drive bays. The connections to the four NVMe-enabled disk drive bays are split between the two root complexes with two bays connected to each PCIe bus.

#### Twenty-Four Front and Two Rear 2.5-Inch SFF Disk Drive Bays

In addition to the 24 disk drive bays in the front, this model is equipped with a dual 2.5-inch SFF disk drive cage in the rear. The eight-port SAS HBA in PCIe slot 7 supports all 26 SAS disk drive bays. A SAS expander is included in the twenty-four SFF disk drive cage, which enables support beyond eight SAS devices. The connections to the four NVMe-enabled disk drive bays are split between the two root complexes with two bays connected to each PCIe bus. Note that the drive bay nomenclature and PCIe device numbering are different from the model with the eight 2.5-inch SFF disk drive bays.

#### Twelve Front 2.5-inch SFF NVMe Disk Drive Bays

This model is designed to provide the highest performing storage capacity using NVMe SSDs. It supports only 2.5-inch SFF NVMe disk drives. Three of the disk drive bays (3–5) are enabled by the onboard PCIe switch (switch 1) with a split connection to two different root complexes. Three additional NVMe PCIe switch cards are included at the factory to support the remaining nine remaining disk drive bays. The PCIe switch card in slot 7 (root complex 1) supports disk drive bays 0–2, the switch in PCIe slot 3 (root complex 1) supports bays 6–8, and the switch in slot 1 (root complex 0) supports bays 9–11.

#### Twelve Front 3.5-Inch LFF and Two Rear 2.5-Inch SFF Disk Drive Bays

The 12 LFF drive bays in the front are designed to host high-capacity 3.5-inch disk drives. At release, 96 TB of storage capacity can be supported with the eight TB 3.5-inch disk drives. In addition, a dual 2.5-inch SFF disk drive cage is provided in the rear. The eight-port SAS HBA in PCIe slot 7 supports all 14 disk drive bays. A SAS expander is included in the 12 LFF disk drive cage, which enables support beyond eight SAS devices. There is no support for NVMe disk drives in this model.

Table 4 describes how the PCIe devices share the root complexes on each of the I/O controllers.

**TABLE 4. SPARC S7-2L SERVER ROOT COMPLEX MAPPING.**

Processor	PCIe Root Complex	Switch	Device	Target	PCIe Speed
#0	# 0	#0	@1	Quad 10GBASE-T Ethernet ports	x8
			@2	USB controller	x1
			@11	PCIe slot 1 <sup>1</sup>	x8
		#1	@14	PCIe slot 4	x8
			@15	PCI slot 5	x8
			@6	NVMe disk drive bay #2, #3, or #5	x4
			@7	NVMe disk drive bay #3, #2, or #4	x4
#1	# 1	#0	@12	PCIe slot 2	x8
			@13	PCIe slot 3 <sup>1</sup>	x8
		#1	@16	PCIe slot 6	x8
			@17	PCIe slot 7 (internal) <sup>1, 2</sup>	x8
			@4	NVMe disk drive bay #0, #1, or #3	x4
			@5	NVMe disk drive bay #1 or #0	x4

1. PCIe slots 1, 3, and 7 are occupied by an NVMe PCIe switch card in the model with the 12 SFF NVMe disk drive cage.

2. Internal PCIe slot 7 is occupied by the SAS HBA card in all models that include SAS disk drive cages.

## Oracle Solaris

Oracle Solaris is the world's most advanced enterprise operating system. It provides superior security, virtualization, and high availability, as well as extreme performance for both vertically and horizontally scaled environments. Oracle Solaris runs on a broad range of SPARC- and x86-based systems, and compatibility with existing applications is guaranteed. The Oracle Solaris 11.3 release is specifically designed to take full advantage of the SPARC M7 and S7 processor-based servers, including the complete set of advanced Software in Silicon features. Oracle Solaris 10 is supported in virtualized guest domains.

Table 5 shows the supported releases of Oracle Solaris 11 and Oracle Solaris 10 in various domains of SPARC S7 processor-based servers.

**TABLE 5. SUPPORTED ORACLE SOLARIS RELEASES ON SPARC S7 PROCESSOR-BASED SERVERS.**


Oracle Solaris Version	Control Domain	Root Domain	I/O Domain	Guest Domain
Oracle Solaris 11.3 <sup>1</sup>	✓	✓	✓	✓
Oracle Solaris 10 1/13 <sup>2</sup>				✓

1. Versions of Oracle Solaris 11 prior to 11.3 are not supported on SPARC S7 processor-based servers.

2. Plus required patches.

Oracle Solaris has incorporated many features to improve application performance on Oracle's multicore/multithreaded architectures.


- » **OpenStack cloud management.** Oracle Solaris 11 includes a complete OpenStack distribution, allowing administrators to centrally share and manage data center resources—including infrastructure and virtualization offerings provided by other vendors—through a single management pane. Integrated into the core technology foundations—such as Oracle Solaris Zones, the ZFS file system, Oracle Solaris Unified Archives, and comprehensive software-defined networking—OpenStack on Oracle Solaris provides self-service computing, allowing IT organizations to deliver services in minutes rather than weeks, with enterprise-grade reliability, security, and performance.
- » **Oracle Solaris software-defined networking (SDN).** The Oracle Solaris 11 release enhances the operating system's existing, integrated software-defined networking technologies to provide much greater application agility without the added overhead of expensive network hardware. It now enables application-driven, multitenant, cloud virtual networking across a completely distributed set of systems; decoupling from physical network infrastructure; and application-level network service-level agreements (SLAs)—all built in as part of the platform. Enhancements and new features include the following:
  - » Network virtualization with virtual NICs (VNICs) and virtual switching
  - » Network resource management and integrated quality of service (QoS) to enforce bandwidth limits on VNICs and traffic flows
  - » Cloud-readiness, which is a core feature of the OpenStack distribution included in Oracle Solaris 11
  - » Application-driven, multitenant, cloud virtual networking with Oracle Solaris Elastic Virtual Switch and Virtual Extensible LANs (VXLANs)
  - » Application-level QoS with application-driven SDN
  - » Tight integration with Oracle Solaris Zones and Oracle Solaris 10 Zones
- » **Lifecycle management.** Oracle Solaris 11 includes a complete and integrated set of technologies for managing the software lifecycle of the platform. With support for secure end-to-end provisioning with the Oracle Solaris Automated Installer, failsafe software updates with the Oracle Solaris Image Packaging System, ZFS boot



environments, rapid application deployment using Unified Archives, and a comprehensive compliance framework, Oracle Solaris 11 helps to increase productivity, reduce human error, and greatly reduce IT costs.

- » **Accelerated cryptography.** Accelerated cryptography is supported through the cryptographic framework in Oracle Solaris as well as through the SPARC S7 processor. The SPARC S7 processor permits access to cryptographic cypher hardware implementations. The cyphers are implemented within the appropriate pipeline in the processor core itself rather than as a coprocessor. This means a more efficient implementation as well as no privilege-level changes, resulting in a large increase in efficiency in cryptographic algorithm calculations. In addition, database operations can make much more efficient use of the various cryptographic cyphers that are implemented within the instruction pipeline itself.
- » **Critical-thread optimization.** Oracle Solaris releases 11 and 10 permit either a user or a programmer to allow the Oracle Solaris Scheduler to recognize a critical thread by means of raising its priority to 60 or above through the use of either the command-line interface (CLI) or system calls to a function. The Oracle Solaris Scheduler will usually evenly distribute threads among all the cores on a system. However, when a process is running in FX60 mode, the scheduler will attempt to give that thread exclusive access to a whole core, and allocate the remaining runnable threads to the other available cores.
- » **Multicore/multithreaded awareness.** Oracle Solaris 11 and Oracle Solaris10 are aware of the SPARC S7 processor hierarchy, so the Oracle Solaris Scheduler can effectively balance the load across all available pipelines. Even though it exposes each of these processors as 64 logical processors, Oracle Solaris understands the correlation among core clusters, L2 and L3 cache hierarchies, and cores and the threads they support to provide a fast and efficient thread implementation.
- » **Fine-granularity manageability.** For the SPARC S7 processor, Oracle Solaris 11 and Oracle Solaris 10 have the ability to enable or disable individual cores and threads (logical processors). In addition, standard Oracle Solaris features, such as processor sets, provide the ability to define a group of logical processors and schedule processes or threads on them.
- » **Binding interfaces.** Oracle Solaris allows considerable flexibility in that processes and individual threads can be bound to either a processor or a processor set, as required or desired.
- » **Support for virtualized networking and I/O devices.** Oracle Solaris contains technology to support and virtualize components and subsystems on the SPARC S7 processor. As part of a high-performance network architecture, Oracle multicore/multithreaded-aware device drivers are provided so that applications running within virtualization frameworks can effectively share I/O and network devices.
- » **Non-uniform memory access (NUMA) optimization in Oracle Solaris.** With memory managed by each SPARC S7 processor, these implementations represent a NUMA architecture. In NUMA architectures, the time needed for a processor to access its own memory is slightly shorter than that required to access memory managed by another processor. Oracle Solaris provides the following technology, which can specifically help to decrease the impact of NUMA on applications and improve performance on NUMA architectures:
  - » **Memory placement optimization (MPO).** Oracle Solaris uses MPO to improve the placement of memory across the physical memory of a server, resulting in increased performance. Through MPO, Oracle Solaris helps ensure that memory is as close as possible to the processors that access it, while still maintaining enough balance within the system. As a result, many database applications are able to run considerably faster with MPO.
  - » **Hierarchical Lgroup Support (HLS).** HLS improves the MPO feature in Oracle Solaris by optimizing performance for systems with more-complex memory latency hierarchies. HLS lets Oracle Solaris distinguish between the degrees of memory remoteness, allocating resources with the lowest-possible latency for applications. If local resources are not available by default for a given application, HLS helps Oracle Solaris allocate the nearest remote resources.
- » **Oracle Solaris ZFS.** Oracle Solaris ZFS, the world's only 128-bit file system, offers a dramatic advance in data management, automating and consolidating complicated storage administration concepts and providing unlimited scalability. Oracle Solaris ZFS is based on a transactional object model that removes most of the traditional constraints on I/O issue order, resulting in dramatic performance gains. Oracle Solaris ZFS also provides data integrity, protecting all data with 64-bit checksums that detect and correct silent data corruption.
- » **Multipathing software.** Multipathing software in Oracle Solaris allows organizations to define and control redundant physical paths to I/O devices such as storage devices and network interfaces. If the active path to a





device becomes unavailable, the software can automatically switch or fail over to an alternate path to maintain availability. To take advantage of multipathing capabilities, the server must be configured with redundant hardware, such as redundant network interfaces or two host bus adapters connected to the same dual-ported storage array.

- » **A secure and robust enterprise-class environment.** Existing SPARC applications continue to run unchanged on SPARC S7 processor-based platforms, protecting software investments. Certified multilevel security protects Oracle Solaris environments from intrusion. The Fault Management Architecture in Oracle Solaris enables elements such as Oracle Solaris Predictive Self Healing to communicate directly with the hardware to help reduce both planned and unplanned downtime. Effective tools, such as Oracle Solaris DTrace, help organizations tune their applications to get the most out of the system's resources.

## Virtualization

Virtualization is an essential technology as organizations strive to run workloads more efficiently. SPARC S7 processor-based servers support Oracle's virtualization technologies including Oracle VM Server for SPARC and OS-based virtualization with Oracle Solaris Zones.

Oracle's virtualization technologies are complementary to each other. Indeed, best practices often include layered virtualization where two or three layers are deployed to achieve optimal security, availability, performance, and manageability. Furthermore, Oracle's virtualization technologies are included at no incremental cost. Following is a brief description of Oracle's layered virtualization technologies that are available with SPARC S7 servers:

- » **Logical domains (LDoms)** are created using Oracle VM Server for SPARC and are used to virtualize a server or physical domain to host multiple virtual machines (VMs), each running its own instance of Oracle Solaris. Oracle VM Server for SPARC is a free-of-charge feature that is included in all of Oracle's SPARC servers. Like prior generations of SPARC processors, the SPARC S7 processor supports a hypervisor that is a small firmware layer providing a stable virtual machine architecture that is tightly integrated with the processor. Multithreading is crucial, because the hypervisor interacts directly with the underlying multicore/multithreading processor. Supported in all current SPARC servers from Oracle, Oracle VM Server for SPARC provides full virtual machines that run an independent operating system instance. These VMs can be configured as root domains—in which case, they are assigned PCIe root complexes for direct access to I/O devices—or as guest domains, in which case, they access virtualized I/O devices. Each operating system instance contains processor, memory, storage, console, and cryptographic devices.
- » **Oracle Solaris Zones** enable OS virtualization so that a single instance of Oracle Solaris can securely isolate applications from each other and allocate system resources to each zone. This essentially allows the creation of multiple virtual machines within a single instance of the Oracle Solaris operating system. This isolation helps enhance security and reliability, because processes in one Oracle Solaris Zone are prevented from interfering with processes running in another Oracle Solaris Zone. Virtual CPUs in a multiprocessor system can be logically partitioned into processor sets and bound to a resource pool, which in turn can be assigned to an Oracle Solaris Zone. Resource pools provide the capability to separate workloads so that the consumption of CPU resources does not overlap. They also provide a persistent configuration mechanism for processor sets and scheduling class assignment. In addition, the dynamic features of resource pools enable administrators to adjust system resources in response to changing workload demands. Oracle Solaris 11 provides Immutable Zones, which preserve the zones' configuration by implementing a read-only root file system. Only specific maintenance operations that are done through a trusted path can be performed.

Oracle's layered virtualization technologies can be used together to create resilient high-availability systems. For example, Oracle engineered systems and Oracle Maximum Availability Architectures employ these virtualization technologies and virtualization best practices in order to achieve very high reliability, availability, and serviceability (RAS). For more information on Oracle's layered virtualization technologies please refer to the Oracle white paper [Consolidation Using Oracle's SPARC Virtualization Technologies](#) and other resources listed in the "For More Information" section of this document

## Systems 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 requirements for remote management of systems. One of the many benefits of Oracle servers is the support for lights-out data centers, enabling expensive support staff to work in any location with network access. The design of the SPARC S7 processor-based servers combines with powerful service processors (SPs) running the Oracle Integrated Lights Out Manager (Oracle ILOM) software; this, along with Oracle Enterprise Manager Ops Center software, helps administrators remotely execute and control nearly any task that does not involve physical access to hardware. These management tools and remote functions lower the administrative burden, saving organizations time and reducing operational expenses.

### Oracle ILOM and Service Processors

The Oracle ILOM software on each SP provides the heart of remote monitoring and management capabilities for SPARC servers. Each SP consists of a dedicated processor that is independent of the server system and runs the Oracle ILOM software package. While input power is supplied to the server, each SP constantly monitors the system even if all domains are inactive.

Each SP regularly monitors the environmental sensors, provides advance warning of potential error conditions, and executes proactive system maintenance procedures, as necessary. For example, an SP can initiate a server shutdown in response to temperature conditions that might induce physical damage to the system. The Oracle ILOM software package running on an SP helps administrators to remotely control and monitor physical domains and virtual machines, as well as the hardware platform itself.

Using a network or serial connection to an SP, operators can effectively administer the server from anywhere on the network. Remote connections to an SP run separately from the operating system and provide full control of and authority over a system console. Each Oracle ILOM SP acts as a system controller, facilitating remote management and administration. Each SP is full-featured and is similar in implementation to that used in Oracle's other servers. As a result, SPARC S7 processor-based servers integrate easily with existing management infrastructure. Critical to effective system management, each Oracle ILOM SP does the following:

- » Implements an IPMI 2.0-compliant SP, providing IPMI management functions to the server's firmware, OS, and applications and to IPMI-based management tools accessing the SP via the Oracle ILOM Ethernet management interface. The SP also provides visibility to the environmental sensors on the server module and elsewhere in the chassis.
- » Manages inventory and environmental controls for the server, including processors, DIMMs, fans, and power supplies and provides HTTPS, CLI, and SNMP access to this data.
- » Supplies remote textual console interfaces.
- » Provides a means to download upgrades to all system firmware.

The Oracle ILOM and SPs also allow administrators to remotely manage a server, independent of the operating system running on the platform and without interfering with any system activity. Oracle ILOM can send email alerts about hardware failures, warnings, and other events related to the server. Its circuitry runs independently from the server, using the server's standby power. As a result, Oracle ILOM firmware and software continue to function when the server operating system goes offline or when the server is powered off. Oracle ILOM monitors the following server conditions:

- » CPU temperature conditions
- » Hard drive presence
- » Enclosure thermal conditions
- » Fan speed and status



- » Power supply status
- » Voltage conditions
- » Oracle Solaris Predictive Self Healing, boot timeouts, and automatic server restart events

## Power Management

Power and cooling costs for servers are becoming significant, and lowering these costs is a top challenge in corporate data centers. Limitations in the availability of power and space to expand data centers force organizations to look closely at the power efficiency of servers. Contracts with power providers, which specify penalties for exceeding the stated power consumption, require organizations to be able to cap their servers' power consumption. Power efficiency and carbon footprint have become factors when organizations evaluate servers.

Beyond the inherent efficiencies of Oracle's multicore/multithreaded server design, the SPARC S7 processor incorporates unique power management features at both the core and memory levels of the processor. These features include reduced instruction rates, parking of idle threads and cores, and the ability to turn off clocks in both cores and memory to reduce power consumption.

In addition to the power management support in Oracle ILOM, Oracle Solaris 11.3 or later provides a power manager that supports SPARC S7 processor-based servers. Oracle Solaris can determine which power-saving features to enable based on the `poweradm` settings, which are set by the platform based on the system (Oracle ILOM) policy but can be overridden by an Oracle Solaris administrator. Substantial innovation is present in the following areas:

- » Limiting speculation, such as conditional branches not taken
- » Extensive clock gating in the data path, control blocks, and arrays
- » Power throttling, which allows extra stall cycles to be injected into the decode stage

In a virtualized environment that uses Oracle VM Server for SPARC, the power management manager performs the following tasks when managing LDom guests:

- » Determining which power savings features to enable based on the power management policy.
- » Calling the Power Management engine to initiate power state changes on its resources to achieve a power adjustment or utilization level (for resources not owned by Oracle Solaris 11.3 guests) or telling the hypervisor to enable or disable hypervisor/hardware-managed power states. Only Oracle Solaris 11.3 guests have a power management peer.

## Oracle Enterprise Manager Ops Center

Oracle Enterprise Manager Ops Center delivers a converged hardware management solution for SPARC S7 processor-based servers that integrates management across the infrastructure stack. With advanced virtualization management and reporting capabilities, applications-to-disk management, intelligent configuration management, and more, Oracle Enterprise Manager Ops Center helps IT managers reduce complexity and streamline and simplify infrastructure management. The inclusion of an Oracle Enterprise Manager Ops Center license with every SPARC server enables data center administrators to monitor and manage storage and network devices, servers, Oracle Solaris, and virtualized environments from a single interface. This improves operational efficiency and lowers operational costs.

Oracle Enterprise Manager Ops Center is the most comprehensive management solution for Oracle servers and Oracle engineered systems infrastructure. Offering a single console to manage multiple server architectures and myriad operating systems, Oracle Enterprise Manager Ops Center can manage the components in SPARC S7 processor-based servers using asset discovery, the provisioning of firmware and operating systems, automated patch management, patch and configuration management, virtualization management, and comprehensive compliance reporting (Figure 9).

Oracle Enterprise Manager Ops Center automates workflow and enforces compliance via policy-based management—all through a single, intuitive interface. With Oracle Enterprise Manager Ops Center, IT staff can implement and enforce data center standardization and best practices, regulatory compliance, and security policies while efficiently deploying an infrastructure that meets business requirements.



Figure 9. Oracle Enterprise Manager Ops Center provides detailed management capabilities for SPARC servers.

## Reliability, Availability, and Serviceability

Reducing downtime—both planned and unplanned—is critical for IT services. Data center architectures must include mechanisms that enable fault resilience, quick repair, and even rapid expansion without impacting the availability of key services. SPARC S7 processor-based servers are designed to support complex network computing solutions while recognizing that scale-out systems often rely on server-level redundancy to meet the most stringent high availability (HA) requirements.

SPARC S7 processor-based servers include basic redundancy and hot-serviceable system components, and advanced diagnostic and error recovery features throughout the design, as well as built-in remote management features. The advanced architecture of these reliable servers fosters high levels of application availability and rapid recovery from many types of hardware faults, simplifying system operation and lowering costs for enterprises.

### Reliability, Availability, and Serviceability Features

SPARC servers provide a host of reliability features, extending from the processor, to the hypervisor, to the system as a whole. These features complement and strengthen each other, lending greater resiliency to the system overall. RAS features engineered into SPARC S7 processor-based servers include the following:


- » Oracle's SPARC S7 processor
  - » L1 cache tag, status, and data: Parity protection and retry on error
  - » L2 and L3 cache data: SEC/DED protection, inline correction, and cache-line sparing
  - » L2 and L3 cache status and directory: SEC/DED protection and inline correction
  - » Architectural registers: SEC/DED protection, precise trap, and hypervisor correction/retry
  - » Message retry in hardware
  - » Dynamic processor core deconfiguration

- » Link-level RAS: Automatic frame retry, automatic link retrain, link-level multipathing, and single-lane failover
- » Dynamic voltage and frequency scaling
- » Main memory
  - » SDRAM error protection: Correction of single-device errors and triple-error detection across devices.
  - » Memory channel interconnect: Message retry in hardware and cyclic redundancy check (CRC).
  - » If persistent recoverable errors are localized to a specific address, the system will use “page retire” functionality to prevent hard errors.
- » Fault Management Architecture
  - » Diagnosis engine on SP and Oracle Solaris
  - » Auto-reconfigure on failure
  - » Soft error rate discrimination (SERD)
  - » Bad-page retirement
  - » Operating system and SP watchdogs
- » SPARC S7 hypervisor
  - » Logical domains (LDDs) virtualization and failure containment
  - » Dynamic assignment of CPU, memory, I/O devices
  - » Dynamic PCIe bus assignment
  - » Processor support for error clearing, correction, and collection
- » System
  - » Diagnosis to the level of a field-replaceable unit (FRU) upon first fault
  - » Hot-pluggable internal 2.5-inch or 3.5-inch disk drives
- » System I/O
  - » PCIe end-to-end CRC
  - » PCIe link retry
  - » Dual SPARC S7 processor cross-connected to onboard PCIe switches
  - » Dual PCIe root complexes for redundant root domain and I/O multipathing (dual processor servers)
- » Power and cooling
  - » Advanced power management
  - » Redundant hot-swappable fans and fan modules
  - » Redundant hot-swappable power supplies
  - » Dual-grid system power

## Advanced Reliability and Availability Features

Advanced reliability features included within the components of SPARC S7 processor-based servers increase the overall stability of the platform. Reduced component count and complexity within the server architecture contribute to reliability. In addition, advanced CPU integration and guaranteed data path integrity provide for autonomous error recovery by the SPARC S7 processors, reducing the time to initiate corrective action and subsequently increasing uptime.

The Fault Management Architecture (FMA) that is implemented both in firmware and in the Oracle Solaris Predictive Self Healing software further enhances the reliability of SPARC servers. The FMA provides constant monitoring of processors, memory, and I/O devices. Depending upon the nature of the error, persistent CPU soft errors can be resolved by automatically off-lining a thread, a core, or an entire processor. In addition, the memory page retirement function supports the ability to take memory pages offline proactively in response to multiple corrections to data for a specific memory DIMM.



Data center deployments of scale-out servers ensure high service levels—partially by providing hardware redundancy at the server level rather than using advanced RAS features within every server. Oracle's SPARC S7 processor-based servers are optimized to include appropriate levels of RAS within the servers themselves while keeping system cost low. Therefore, redundancy and the ability to perform hot-service operations are limited to key components.

### Error Detection, Diagnosis, and Recovery

SPARC S7 processor-based servers 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. Ultimately, the following features work together to raise application availability:

- » End-to-end data protection detects and corrects errors throughout the system, ensuring complete data integrity.
- » State-of-the-art fault isolation helps the server isolate errors within component boundaries and offline only the relevant item, chip, or component subsection. Isolating errors down to the smallest possible entity improves stability and provides the continued availability of maximum compute power. This feature applies to processors, memory DIMMs, switch ASICs, connectivity links, and SPs.
- » Constant environmental monitoring provides a historical log of pertinent environmental and error conditions.
- » The host watchdog feature periodically checks for the operation of software, including the operating system for each domain. This feature also uses the Oracle ILOM firmware to trigger error notification and recovery functions.
- » The FMA capability of the system and dynamic CPU resource deconfiguration of the SPARC S7 processor enable powerful isolation and recovery. If necessary, the system can dynamically retire processor resources (for example, a core) without interrupting the applications that are running.
- » The system performs periodic component status checks to detect signs of an impending fault. Recovery mechanisms are then 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.

### Redundant and Hot-Serviceable Components

SPARC servers feature redundant, hot-swappable power supply and fan units, as well as the option to configure multiple storage devices, memory DIMMs, and I/O cards. All SPARC S7 processor-based servers are available with dual SPARC S7 processors. Administrators can combine redundant hot-pluggable internal storage with Oracle Solaris ZFS for RAID protection and resiliency. If a fault occurs, these duplicated components can support continued operation or recovery without an immediate need for hardware replacement. Depending upon the component and type of error, the system might continue to operate in a degraded mode or it might reboot—with the failure automatically diagnosed and the relevant component automatically configured out of the system. In addition, hot-serviceable hardware within these servers speeds service and allows for the simplified replacement or addition of components, without a need to interrupt or stop the system.



## Conclusion

Modern technology and business challenges require innovative solutions, significant compute power, and a range of server form factors to deliver on the promise of a rapidly evolving digital marketplace. Based on the strengths of Oracle's innovative SPARC S7 processor, Oracle's SPARC servers provide a new approach that can help organizations efficiently deliver their most demanding computational workloads in a resilient and highly available fashion. Through specialized acceleration engines on each SPARC S7 processor, innovative Software in Silicon technology provides Silicon Secured Memory, in-core cryptographic acceleration, and SQL in Silicon, which can make a dramatic difference for applications.

Oracle's entry-level systems based on the SPARC S7 processor let organizations shift computing costs away from underlying infrastructure and focus on the productive business use of their workloads. For example, new levels of performance and efficiency let organizations run both OLTP and data analytics on the same system. Coupled with Oracle Solaris functionality and Oracle's layered virtualization technologies, organizations can create the most efficient on-premises clouds with superior capacity and scalability, helping to enable accelerated time to market and enhancing the bottom line.

## For More Information

For more information, visit the resources listed in Table 6.

**TABLE 6. RESOURCES.**

Web Resources	URL
Oracle's SPARC servers	<a href="http://oracle.com/us/products/servers-storage/servers/sparc/oracle-sparc">oracle.com/us/products/servers-storage/servers/sparc/oracle-sparc</a> and <a href="http://oracle.com/technetwork/server-storage/sun-sparc-enterprise">oracle.com/technetwork/server-storage/sun-sparc-enterprise</a>
Software in Silicon	<a href="http://community.oracle.com/community/softwareinsilicon/overview">community.oracle.com/community/softwareinsilicon/overview</a>
Oracle Solaris	<a href="http://oracle.com/solaris">oracle.com/solaris</a>
Oracle Optimized Solutions	<a href="http://oracle.com/optimizationsolutions">oracle.com/optimizationsolutions</a>
Oracle benchmark results	<a href="http://oracle.com/benchmarks">oracle.com/benchmarks</a>
Oracle Help Center documentation	<a href="http://docs.oracle.com">docs.oracle.com</a>
Security	<a href="http://oracle.com/technetwork/topics/security">oracle.com/technetwork/topics/security</a>
Virtualization	<a href="http://oracle.com/technetwork/topics/virtualization">oracle.com/technetwork/topics/virtualization</a>
Oracle Integrated Lights Out Manager	<a href="http://docs.oracle.com/en/hardware/?tab=4">docs.oracle.com/en/hardware/?tab=4</a>
Oracle Enterprise Manager Ops Center	<a href="http://oracle.com/technetwork/oem/ops-center">oracle.com/technetwork/oem/ops-center</a>
White Papers and Technical Articles	URL
"Oracle's SPARC T7 and SPARC M7 Server Architecture"	<a href="http://oracle.com/technetwork/server-storage/sun-sparc-enterprise/documentation/sparc-t7-m7-server-architecture-2702877.pdf">oracle.com/technetwork/server-storage/sun-sparc-enterprise/documentation/sparc-t7-m7-server-architecture-2702877.pdf</a>
"Consolidation Using Oracle's SPARC Virtualization Technologies"	<a href="http://oracle.com/technetwork/server-storage/sun-sparc-enterprise/technologies/consolidate-sparc-virtualization-2301718.pdf">oracle.com/technetwork/server-storage/sun-sparc-enterprise/technologies/consolidate-sparc-virtualization-2301718.pdf</a>
"Best Practices for Oracle Solaris Kernel Zones"	<a href="http://oracle.com/technetwork/articles/servers-storage-admin/solaris-kernel-zones-best-practices-2400370.html">oracle.com/technetwork/articles/servers-storage-admin/solaris-kernel-zones-best-practices-2400370.html</a>
"Oracle VM Server for SPARC Best Practices"	<a href="http://oracle.com/technetwork/server-storage/vm/ovmsparc-best-practices-2334546.pdf">oracle.com/technetwork/server-storage/vm/ovmsparc-best-practices-2334546.pdf</a>
"Implementing Root Domains with Oracle VM Server for SPARC"	<a href="http://oracle.com/technetwork/server-storage/vm/ovm-sparc-rootdomains-wp-1914481.pdf">oracle.com/technetwork/server-storage/vm/ovm-sparc-rootdomains-wp-1914481.pdf</a>
"Oracle Multitenant on SPARC Servers and Oracle Solaris"	<a href="http://oracle.com/technetwork/articles/servers-storage-admin/multitenant-on-sparc-solaris-2016889.html">oracle.com/technetwork/articles/servers-storage-admin/multitenant-on-sparc-solaris-2016889.html</a>

**Oracle Corporation, World Headquarters**

500 Oracle Parkway  
Redwood Shores, CA 94065, USA

**Worldwide Inquiries**

Phone: +1.650.506.7000  
Fax: +1.650.506.7200

## CONNECT WITH US

[blogs.oracle.com/oracle](https://blogs.oracle.com/oracle)[facebook.com/oracle](https://facebook.com/oracle)[twitter.com/oracle](https://twitter.com/oracle)[oracle.com](https://oracle.com)**Integrated Cloud Applications & Platform Services**

Copyright © 2016, Oracle and/or its affiliates. All rights reserved. This document is provided *for* information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0615

Oracle's SPARC S7 Server Architecture: Secure Platform for Scale-Out and Cloud Infrastructures  
July 2016



Oracle is committed to developing practices and products that help protect the environment