




x86 SERVERS

Oracle Server X7-2 and Oracle Server X7-2L System Architecture

ORACLE WHITE PAPER | SEPTEMBER 2017





Introduction	1
System Overview	1
Reimagining Database Storage with All-Flash Configurations	1
Extreme Memory and I/O for IaaS	2
Best for Oracle Software	3
The Unsettling Reality of Firmware-Based Attacks	4
Defense-In-Depth Security	4
Security and Oracle Cloud	4
Introducing Oracle Integrated Lights Out Manager 4.0	4
Oracle ILOM 4.0 Development Process and Security Assurance	5
Oracle ILOM Secure Verified Boot	5
Oracle ILOM 4.0 Security Features at a Glance	6
Trusted Operating System Boot	6
Innovative Reliability, Availability, and Serviceability	6
Hardware Designed for Advanced RAS	7
Fault Management and Diagnostics	7
Conclusion	8

Introduction

Oracle Server X7-2 and Oracle Server X7-2L, Oracle's latest two-socket servers, are the newest additions to Oracle's family of x86 servers that are purpose-built to be best for running private cloud solutions. Oracle Server X7-2 is a 1U server that is an ideal platform for infrastructure as a service (IaaS) and for small or medium flash-optimized instances of Oracle Database. Oracle Server X7-2L is a 2U server that is optimized for customers requiring a robust and reliable platform for running instances of Oracle Database with large amounts of direct-attached storage and as a consolidation platform capable of handling the I/O requirements of large numbers of virtual machines (VMs).

System Overview

Oracle Server X7-2 and Oracle Server X7-2L support up to two Platinum, Gold, or Silver Intel® Xeon® Processor Scalable Family processors, each containing up to 24 cores and up to 33 MB of L3 cache. With 24 DDR4-2666 dual inline memory module (DIMM) slots, the servers provide up to 1.5 TB of main memory with 64 GB LRDIMMs. Memory bandwidth increases to 2666 MT/sec from 2400 MT/sec in the previous generation.


In addition, Oracle Server X7-2 has four PCI Express (PCIe) Gen3 slots (2x 16-lane, 2x 8-lane slots), two 10GBase-T or two 25 GbE SFP28 ports, one 1GBase-T port, two USB ports and, optionally, two M.2 flash devices for booting. Oracle Server X7-2 can be configured with up to eight small form factor (SFF) drives in the front of the chassis with hard disk drives (HDDs) or solid-state drives (SSDs). With an optional NVMe Express (NVMe) cable kit, all eight disk bays are capable of supporting NVMe SSDs, each with a flash capacity of 6.4 TB.

Oracle Server X7-2L supports 11 PCIe Gen3 slots (1x 16-lane, 10x 8-lane slots), one 1GBase-T port, two USB ports and, optionally, two M.2 flash devices for booting. Oracle Server X7-2L has a single disk cage that allows for 12 large form factor (LFF) drives. With optional internal NVMe switch cards, Oracle Server X7-2L can be configured with up to 12 low-latency, high-bandwidth NVMe SSDs in the front of the chassis as well up to four NVMe add-in cards for a total of sixteen 6.4 TB NVMe devices. This provides a total flash capacity of 102.4 TB in only 2U of space.

Reimagining Database Storage with All-Flash Configurations

Because NVMe flash has become much less expensive per gigabyte of capacity, it is now possible to affordably configure servers with very large flash configurations. Oracle Server X7-2 allows for up to 51.2 TB of NVMe flash in 1U of space, while Oracle Server X7-2L provides up to 102.4 TB of NVMe flash in 2U of space. With such large amounts of flash, it is no longer necessary to store any part of the database on slower, rotating drives. By storing the entire database on flash, enterprises can remove a major I/O bottleneck and dramatically improve performance. This is especially important for instances of Oracle Database in branch or remote offices away from a corporate SAN.

Both Oracle Server X7-2 and Oracle Server X7-2L offer the option of forgoing conventional storage with a SAS HBA in favor of an all-flash NVMe configuration. Oracle's NVMe devices have a much lower latency and higher bandwidth than standard SAS-3 drives due to the fact that each drive connects directly to four lanes (SFF) or eight lanes (LP-PCIe form factor) of PCIe Gen3 with an aggregate bandwidth of 32 Gb/sec and 64 Gb/sec, respectively, compared to 12 Gb/sec for a traditional SAS-3 SSD.



Oracle's x86 servers provide a major NVMe performance breakthrough by coengineering the NVMe flash with the servers. Because flash technologies are temperature-sensitive, most high-performance flash drives will throttle down their I/O speeds as temperatures rise in order to protect the flash from damage. Oracle's NVMe SSDs, on the other hand, include multiple temperature sensors that are monitored by the servers' Oracle Integrated Lights Out Manager (Oracle ILOM) service processor to ensure the drive maintains optimum operating temperature. Oracle ILOM modulates the fan speed to ensure sufficient cooling for maximum system performance at all times. The benefits of this are that the servers consistently operate at maximum performance across their full operating temperature range independent of the system configuration.

Extreme Memory and I/O for IaaS

As organizations face growing IT expenses while at the same time moving toward modern, private cloud architectures, it is essential to be able to do more with less. Server virtualization is the foundation of private cloud infrastructures and serves as the consolidation mechanism for heterogeneous workloads. Oracle Server X7-2 and Oracle Server X7-2L are the ideal platforms for virtualization, providing the ability to get the most out of each server by simultaneously maximizing memory capacity, I/O, and compute density.

The best virtualization platforms allow for high VM density while providing fast live migration, reliability, and performance. While one important metric for estimating VM density is core density, there are actually many other factors—such as memory capacity, memory bandwidth, and I/O bandwidth—that are equally important in determining how many VMs can be consolidated onto one server. Enterprise-class VM environments rely heavily on I/O bandwidth and low-latency networks to be able to migrate VMs for load balancing as well as failover scenarios. In addition, the two on-board 25 GbE ports on Oracle Server X7-2 provide additional I/O bandwidth without consuming a PCIe slot.

The memory architecture has been dramatically improved with the latest generation of Oracle x86 servers. First, there are now six memory channels per processor rather than four in the previous generation. The additional memory channels, as well as the increased memory bus speed of 2,666 MT/sec, increase the total system memory bandwidth to 256 GB/sec. This allows for significant performance increases due to the additional memory bandwidth. Second, a new motherboard design now supports three Intel® Ultra Path Interconnect (Intel® UPI) links between the two processors, each link operating at up to 10.4 GT/sec. The improved connections between processors significantly reduce the latency when nonlocal memory is accessed and remote I/O operations are performed.

With an optimal balance among core density, memory footprint, and I/O bandwidth, Oracle Server X7-2 and Oracle Server X7-2L can be easily deployed into existing data centers as the building block of a private cloud or IaaS implementation. With up to 106 Gb/sec of raw I/O bandwidth, Oracle Server X7-2L maximizes I/O write operations for virtualized cloud environments to take advantage of large number of cores. The 11 PCIe Gen3 slots allow for a very high VM consolidation factor by eliminating any I/O bottlenecks and providing proportional I/O bandwidth to match the large number of cores (48) that are possible with the latest processors.



Best for Oracle Software

Oracle Server X7-2 and Oracle Server X7-2L systems are ideal x86 platforms for running Oracle software. Only Oracle provides customers with an optimized hardware and software stack that comes complete with choice of operating system, virtualization software, and cloud management tools—all at no extra charge. Running Oracle software on Oracle's x86 systems provides performance, reliability, and supportability advantages not found on third-party x86 servers. The combination of Oracle Database, Oracle Linux, and Oracle's x86 systems represents an integrated stack that has been designed and tested to deliver maximum performance and reliability in enterprise IT environments.

Oracle invests heavily in engineering and quality assurance for its Oracle Linux operating system. While many customers choose Oracle Linux to support their mission-critical applications, Oracle Linux is also the principal development platform for Oracle's own database, middleware, and application software, utilizing more than 175,000 Oracle Linux installations deployed on both physical and virtual servers.

Oracle Linux receives more than 128,000 hours of database and application testing each day, which makes Oracle software more reliable. Even before formal evaluation occurs, Oracle Linux is the base platform on which developers prove functionality, quality, and software viability. In addition, Oracle Linux includes Oracle's Unbreakable Enterprise Kernel, which is specifically optimized for the best performance of Oracle software. Oracle engineers extensively test the Unbreakable Enterprise Kernel across Oracle's database, middleware, and application tiers on Oracle's x86 servers and engineered systems to ensure optimum functionality. This extensive testing ensures that the combination of Oracle Server X7-2L with Oracle Linux provides an extremely reliable, robust, and high-performance server for database and enterprise applications.

As part of Oracle's philosophy of layering security, Oracle Linux is the only Linux distribution today to support Oracle Ksplice, a ground-breaking technology that allows data center managers to apply security patches to the operating system kernel without incurring downtime. This aids IT administrators in staying up to date with the latest security-related patch releases.

Oracle's x86 systems and Oracle Linux are coengineered with Oracle Database so that the combined hardware and software stack, benefiting from numerous optimizations, produces advantages that are greater than the sum of the individual advantages. This yields an increased return on Oracle software investments as well as a system-wide lowering of operating and management costs.

Also, users of an Oracle hardware and software stack realize better quality and higher reliability. Oracle Linux and Oracle's x86 systems are used by Oracle developers and quality assurance teams to write and test the database and application code. With every daily software build by Oracle's product development team, testing provides one more opportunity to examine edge cases and to improve the extensive quality review of Oracle software on Oracle hardware.

Using the combination of Oracle Database, Oracle Linux, and Oracle's standalone x86 servers, customers also realize the benefits from Oracle's highly optimized engineered systems such as Oracle Exadata, which delivers a 10x improvement in large-scale database performance. These engineered systems start with Oracle's x86 standalone servers as core building blocks for compute and storage servers within the overall system architecture.

Oracle's approach to x86 system hardware design is distinctive. Other tier-1 vendors use an off-the-shelf, or a slightly modified off-the-shelf, x86 motherboard design. Only Oracle starts with the standard Intel® Xeon® line of processors and a clean sheet of paper. Oracle's focus is to design a system that runs both Oracle software best and provides capabilities demanded by enterprise and cloud usage. An Oracle-on-Oracle solution provides assurance that the Oracle software is deployed in a thoroughly tested and optimized environment, increasing reliability and performance, and thereby reducing business risk.



The Unsettling Reality of Firmware-Based Attacks

With service processors becoming more powerful and system firmware increasing in complexity, the possibility that firmware can be maliciously attacked and exploited is becoming more likely. As these systems are deployed to the cloud, the fear that attacks originating in a guest operating system could escape past the hypervisor to a privileged domain is becoming a reality. Once it is in a privileged domain, malicious code can then infect the firmware and survive across disk wipes and across tenants and potentially further spread to the control plane of the cloud, wreaking havoc.

Because firmware-based attacks operate at a lower level than can be detected by traditional virus scan software or network-based security tools, the only method for combating these threats is from within the server itself using a defense-in-depth approach. Protective measures in hardware and firmware are necessary to adequately defend against and detect firmware-based attacks. In particular, this requires hardware-based mechanisms to validate the integrity of firmware as well as modern, hardened interfaces across all firmware-based attack surfaces.

Defense-In-Depth Security

Oracle's philosophy on security in depth is based on the philosophy that "security needs to be built in, not bolted on." Oracle has a company-wide initiative to incorporate security features across all of its products, starting with the design and manufacturing of its servers, through the operating systems layers, and extending into the database, middleware, and application layers. Oracle's Global Product Security group is chartered with the goal of setting, auditing, and enforcing security policies across all Oracle products. It also performs periodic security audits and ensures compliance with the latest threat profiles. This organization also publishes regular security alerts to users of Oracle products.

Security and Oracle Cloud

Oracle Cloud is the industry's broadest and most integrated public cloud. It offers best-in-class services across software as a service (SaaS), platform as a service (PaaS), and IaaS, and even lets you put Oracle Cloud in your own data center. Oracle Cloud helps organizations drive innovation and business transformation by increasing business agility, lowering costs, and reducing IT complexity.

Oracle's x86 servers are used as the hardware building blocks of Oracle Cloud solutions and, therefore, must provide the highest levels of security. Because Oracle Cloud runs some of the most mission-critical applications of any cloud, it is essential that it be protected from security threats at all levels. This protection would not be possible using commercial off-the-shelf hardware. Instead, Oracle engineers its own hardware in order to satisfy the unique security requirements of Oracle Cloud.

Oracle Server X7-2 and Oracle Server X7-2L offer best-in-breed security features, leveraged from Oracle's own cloud offerings, for customers building their own cloud solutions. This white paper discusses these new features and enhancements.

Introducing Oracle Integrated Lights Out Manager 4.0

Oracle Server X7-2 and Oracle Server X7-2L introduce the next generation of cloud-ready service processor that is designed for today's security challenges. Oracle ILOM 4.0 is a new service processor that uses advanced service processor hardware with built-in hardening and encryption. In order to make server management more secure, Oracle ILOM 4.0 has a new set of modern APIs that are more secure and reduce the attack surface by eliminating older, insecure interfaces. In addition, Oracle ILOM 4.0 uses only industry-leading strong ciphers to encrypt communications and data.

Oracle ILOM 4.0 includes a hardened code base at all levels in order to minimize the potential of malicious code to exploit defects in the kernel, operating system, and other parts of the firmware. In addition, Oracle ILOM 4.0 includes improved firmware image signing to further guarantee the integrity of firmware. This prevents any third-party modification of the firmware and guarantees that only firmware created and signed by Oracle can be used on the servers.

Oracle ILOM 4.0 Development Process and Security Assurance

The advanced security features of Oracle ILOM 4.0 are the result of an industry-leading development and security-assurance process. This process extends throughout the development of Oracle ILOM firmware as well as after the firmware has been released. Oracle's dedicated security team tracks and finds security vulnerabilities and assesses their impact to Oracle ILOM firmware. These vulnerabilities are fixed with regular updates to the kernel and distribution underlying Oracle ILOM. The result is a modern, hardened code base that is maintained through a vigilant software development process that includes fixing security vulnerabilities that are listed in the Common Vulnerabilities and Exposures (CVE) list.

The diagram below shows this process:

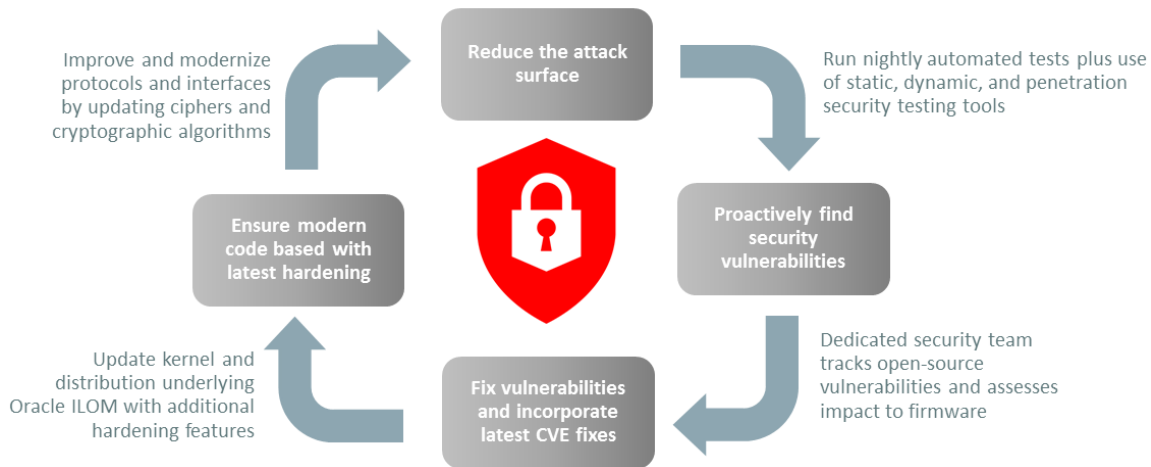


Figure 1. Software development process

Oracle ILOM Secure Verified Boot

One of the most important new security features in Oracle ILOM 4.0 is secure verified boot. This feature anchors validation of the Oracle ILOM firmware in hardware to guarantee its integrity upon every boot. Specifically, Oracle ILOM's firmware signature is validated on every boot by code that is hardware-protected and cannot be modified or compromised by malicious code. By cryptographically preventing the service processor from booting malicious Oracle ILOM firmware, the secure verified boot feature ensures that only firmware created by Oracle can be booted.

The secure verified boot feature is part of a chain of trust that ensures integrity from the hardware up the stack. Specifically, Oracle ILOM 4.0's boot code is validated by hardware-protected code. That validated boot code then validates Oracle ILOM's file system and other components such as the embedded BIOS image to ensure validity.

Oracle ILOM 4.0 Security Features at a Glance

The features described above are only some of the many new Oracle ILOM 4.0 security features. Because security is only as strong as its weakest link, Oracle ILOM 4.0 includes security enhancements in many areas, as outlined in Table 1.

TABLE 1: ORACLE ILOM 4.0 SECURITY FEATURES AT A GLANCE

Security Feature	Description
Hardening	<ul style="list-style-type: none">» Automatically prevents buffer overflow attacks using address space layout randomization» Provides hardware-enforced data execution prevention (DEP)» Fixes relevant vulnerabilities and CVEs» Provides updated kernel and distribution with additional hardening features
FIPS 140-2 Update	<ul style="list-style-type: none">» FIPS 140-2 certification with updated object module version
New Security APIs	<ul style="list-style-type: none">» Consist of new REST APIs designed for secure, cloud deployments» Replace older management interfaces with modern, encrypted web-services APIs» Enable private, secure communication between host and Oracle ILOM
Strong Ciphers	<ul style="list-style-type: none">» Ensure secure code and communications using hardware-based, strong ciphers
Oracle ILOM Secure Verified Boot	<ul style="list-style-type: none">» Validates Oracle ILOM firmware by hardware-protected code prior to each boot» Guarantees Oracle ILOM will not boot malicious firmware
Firmware Integrity	<ul style="list-style-type: none">» Provides Improved firmware image signing with 2048-bit DSA key» Ensures only Oracle authentic firmware is used

Trusted Operating System Boot


In addition to the many firmware and Oracle ILOM 4.0 security enhancements, Oracle Server X7-2 and Oracle Server X7-2L also provide features that secure the host boot process. These features provide built-in security for the BIOS boot loader and operating systems, as described in Table 2.

TABLE 2. TRUSTED OPERATING SYSTEM BOOT FEATURES

Feature	Description
UEFI Secure Boot	<ul style="list-style-type: none">» Validates the authenticity of the host's boot loader» Protects Oracle Server X7-2 and Oracle Server X7-2L against malicious code being loaded and executed early in the boot process
Intel® Trusted Execution Technology (Intel® TXT)	<ul style="list-style-type: none">» Validates that the operating system is authentic and is running in a trusted environment» Hardens the servers from the emerging threats of hypervisor attacks, BIOS, or other firmware attacks, malicious root kit installations, or other software-based attacks

Innovative Reliability, Availability, and Serviceability

Reliability, availability, and serviceability (RAS) are extremely important to customers who demand maximum system uptime when running business-critical applications. If a fault occurs in a server, revenue can be lost.



Oracle Server X7-2 and Oracle Server X7-2L are designed completely in house from the ground up, and they are engineered to be easily serviceable while maximizing reliability. In particular, the chassis design has special features that improve performance while also improving reliability and serviceability. Oracle engineers designed a rigorous testing process for all components of the servers such as memory DIMMs, HDDs, power supplies, and so on. These quality assurance tests are supplementary to those conducted by the component suppliers. All components of the system have to pass these tests prior to the release of the product to market.

Hardware Designed for Advanced RAS

Oracle Server X7-2 and Oracle Server X7-2L are designed for maximum uptime with enterprise-grade availability features. All disks are hot-swappable and support RAID 0, 1, 5, 6, 10, 50, and 60. The RAID controller has 2 GB of onboard SDRAM with write-back flash for cache protection, and it uses an energy storage module to save data in flash upon a server power failure. This energy storage module resides in a location in the servers that guarantees data protection of the write-back cache for all operating conditions of the servers. The power supplies and fans are also redundant and hot-swappable, ensuring that a failure to any single component does not affect the running systems. With two power supplies, the servers offer N+N power redundancy.


The chassis and motherboard are designed to eliminate as many cables as possible; for example, the power supplies mate directly to connectors on the motherboard, eliminating a power distribution cable and a single point of failure. The fans also mount directly to the motherboard, eliminating cables and, hence, improving reliability. All disks are front-accessible and hot-swappable, including the NVMe SFF drives.

Fault Management and Diagnostics

With higher levels of integration among various subsystems in servers, it is becoming more complex to diagnose faults down to the component level. A key element of serviceability that is taken into consideration in Oracle Server X7-2 and Oracle Server X7-2L is automatic fault diagnosis with accurate identification of faulty components. This results in significant reductions in time and effort for debugging problems and waiting on service personnel to replace faulty components.

Oracle Server X7-2 and Oracle Server X7-2L include built-in fault management and diagnostic tools that increase system availability and enable faster service response times that increase server uptime. Oracle Server X7-2 and Oracle Server X7-2L include Oracle ILOM, which performs advanced health monitoring of the server operating environment (power and cooling), CPUs, and memory subsystems. This advanced diagnosis engine is resident in the embedded service processor firmware and constantly monitors the state of these subsystems without interfering with the functionality of the host. Automatic notifications are generated in the event of problems. Building on the fault management infrastructure, Oracle ILOM has the ability to raise automatic service requests (ASRs). This ability enables service requests to be generated automatically with important fields prepopulated for use by Oracle service personnel. The elimination of human intervention in the service-request generation process improves the accuracy of problem notification to Oracle.

On a typical server, the host operating system and the service processor have mutually exclusive (although sometimes partially overlapping) subsystems that they manage. The host operating system has ownership of the CPU, memory, and I/O subsystems while the service processor presides over the fans, power supplies, DIMMs, and other miscellaneous chassis components. For these reasons, data center managers are often forced to monitor the health of the host operating system and the service processor as if they are separate entities.



Oracle Server X7-2 and Oracle Server X7-2L overcome the above limitations by enabling a bidirectional communication path between Oracle ILOM and Oracle Linux or Oracle Solaris. This path facilitates the exchange of critical health information between the host and the service processor. Having a dedicated interconnect between the host operating system and Oracle ILOM allows a holistic and single view of all problems in a system. Data center managers and administrators can depend on this operating system and hardware integration for complete system diagnosis, eliminating the need to connect to multiple management entities.

Oracle Linux and Oracle Solaris include a set of diagnosis engines that process raw error events from the hardware and provide an automated and intelligent method for problem diagnosis and fault isolation. These engines are part of the overall Fault Management Architecture feature of Oracle Linux and Oracle Solaris and include a set of agents that respond to fault events, such as offlining a faulty CPU thread or retiring a memory page on a DIMM. These advanced, self-healing features help reduce unplanned downtime by isolating a problem at runtime and keeping applications running.

Running Oracle Linux or Oracle Solaris on Oracle Server X7-2 and Oracle Server X7-2L ensures maximum system availability by providing early warnings of potential failures, fault visibility, and dynamic offlining of faulty hardware. All of these functions are available at no additional cost.

Conclusion

As enterprises continue to build out their private cloud solutions, IT departments are looking to provide a secure software and hardware infrastructure—one that delivers responsive performance, scalable capacity, and secure “always-on” availability. Oracle Server X7-2 and Oracle Server X7-2L are designed to simplify field installation, reduce cabling, minimize power consumption, maximize system uptime, and improve storage density, making them an ideal choice for system deployment.

Oracle Server X7-2 and Oracle Server X7-2L include many new enhancements that improve cloud implementations. The key features incorporated in Oracle Server X7-2 and Oracle Server X7-2L—such as an improved memory architecture, additional CPU links between processors, and extreme I/O configurations—significantly improve performance for virtualized cloud environments compared to the previous generation. Combining these features with built-in, proactive fault detection and advanced diagnostics ensures Oracle’s new x86 servers provide extreme reliability for enterprise cloud solutions.

Oracle’s x86 systems serve as a key building block for Oracle’s engineered systems, such as Oracle Exadata, which have achieved a 10x performance gain through integration and optimization. These optimizations have been incorporated into the design of Oracle Server X7-2 and Oracle Server X7-2L, further improving their performance and reliability—making them an ideal choice for enterprises that value the quality, system availability, and server efficiency that reduce total cost of ownership.

More information about Oracle Server X7-2 and Oracle Server X7-2L can be found at oracle.com/goto/X7-2 or an Oracle representative can be reached at +1.800.ORACLE1.







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
-  facebook.com/oracle
-  twitter.com/oracle
-  oracle.com

Integrated Cloud Applications & Platform Services

Copyright © 2017, 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. 0817

Oracle Server X7-2 and Oracle Server X7-2L System Architecture
September 2017



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