




x86 SERVERS

Oracle Server X8-8 System Architecture

ORACLE WHITE PAPER | APRIL 2019





| | |
|---|---|
| Introduction | 1 |
| System Overview | 1 |
| System Design | 1 |
| Converging Chassis Design | 2 |
| Unique Eight-Socket Glueless Architecture | 2 |
| Best for Scale-Up Enterprise Applications | 2 |
| Extreme CPU Performance and Memory Capacity for Oracle Database | 3 |
| Oracle Database In-Memory on Oracle Server X8-8 | 3 |
| Database Consolidation on Oracle Server X8-8 using Oracle Multitenant | 3 |
| The Unsettling Reality of Firmware-Based Attacks | 4 |
| Defense-In-Depth Security | 4 |
| Security and Oracle Cloud | 4 |
| Oracle Integrated Lights Out Manager 4.0 | 5 |
| Oracle ILOM 4.0 Development Process and Security Assurance | 5 |
| Oracle ILOM Secure Verified Boot | 6 |
| Oracle ILOM 4.0 Security Features at a Glance | 6 |
| Trusted Operating System Boot | 6 |
| Innovative Reliability, Availability, and Serviceability | 7 |
| Fault Management and Diagnostics | 7 |
| Conclusion | 8 |

Introduction

Oracle Server X8-8, Oracle's latest server capable of both four-socket and eight-socket configurations, is the newest addition to Oracle's family of x86 servers that are purpose-built to be best for running Oracle Database. Oracle Server X8-8 is a 5U server that is an ideal platform for Oracle Database In-Memory due to its large memory capacity and bandwidth. With 192 total cores, it is also well suited to Oracle Multitenant, allowing multiple database instances to be consolidated into fewer physical servers. With its unmatched serviceability and reliability features, along with state-of-the-art security in depth, Oracle Server X8-8 is an enterprise-class system capable of even the most extreme workloads.

System Overview

Oracle Server X8-8 supports up to eight Platinum or Gold Intel® Xeon® Scalable Processor Second Generation processors, each containing up to 24 cores and up to 35.75 MB of L3 cache. With ninety-six DDR4-2666 dual inline memory module (DIMM) slots, the server provides up to 6 TB of main memory with 64 GB LRDIMMs in an eight-socket configuration or up to 3 TB of main memory in a four-socket configuration. Memory speed is 2,666 MT/sec, allowing for extreme system-level memory bandwidth. In addition, Oracle Server X8-8 has two PCI Express (PCIe) Gen3 slots (1x 16-lane, 1x 8-lane) per processor for a total of 16 slots.

Oracle Server X8-8 also provides a combination of eight 10GBase-T ports, two SAS-3 disk controllers, eight drive bays for hard disk drives (HDDs) or conventional solid-state drives (SSDs), and two external USB ports. Oracle Server X8-8 can be configured with up to eight 6.4 TB NVMe Express (NVMe) LP-PCIe cards for a total capacity of 51.2 TB.


System Design

Oracle Server X8-8 is a fully modular system, engineered to be easily serviceable. All HDDs and SSDs are hot-swappable. NVMe cards are hot-swappable via I/O slots. The four (N+N) redundant 3,000 watt power supply units (PSUs) and eight redundant fans are also hot-swappable.

The eight sockets are connected through a passive midplane that maximizes availability and allows for maximum inter-socket bandwidth. Oracle Server X8-8 can be configured in a single eight-socket configuration or in a dual four-socket configuration with each socket attached to up to 12 DIMMs. In a single eight-socket configuration, there are three Intel® Ultra Path Interconnect (Intel® UPI) links between each processor, while in a four-socket configuration there are two Intel® UPI links between processors (with the third link disabled in firmware).

Each CPU module (CMOD) contains a single Intel® Xeon Platinum or Gold processor. The processor in each CMOD connects to 12 DIMMs over six memory channels with two DIMMs per channel. The CMODs are located just behind the front fans and are front-panel-accessible by removing the fans. Cooling for Oracle Server X8-8 is front to back using eight variable-speed, dual-counter rotating fan assemblies mounted in the front of the system.

In the rear of the chassis, there are dual system modules (SMODs) that each house peripheral control hub functions, an Oracle Integrated Lights Out Manager (Oracle ILOM) service processor, an internal RAID storage HBA, and four small form factor (SFF) drive bays. Each SMOD also provides connectivity via four 10 GbE ports. In a single eight-socket configuration, the operating system sees dual RAID HBAs that each control four drives. In a dual four-socket configuration, each RAID HBA is dedicated to a single host.



The dual PCIe card carriers (DPCCs) that hold standard low-profile PCIe cards are rear-accessible and, when inserted, sit directly behind a CMOD such that each of the eight CMODs corresponds to one of the eight DPCC modules, respectively. I/O connections are distributed across the eight processors. Because each CMOD has a corresponding DPCC that holds the PCIe slots, each processor is directly connected to one 16-lane and one 8-lane PCIe slot.

Converging Chassis Design

The unique chassis of Oracle Server X8-8 is capable of housing a single eight-socket server or two four-socket servers. In previous generations, the four- and eight-socket x86 servers from Oracle had entirely different designs. By converging the chassis designs of these two different systems into a single chassis capable of multiple configurations, Oracle Server X8-8 provides unparalleled flexibility, serviceability, and reliability. Customers can switch between single eight-socket and dual four-socket modes at any time using BIOS or Oracle ILOM settings as workloads change.

With two four-socket servers in a single 5U chassis, Oracle Server X8-8 reduces data center footprint without compromising reliability and serviceability. All serviceable components have front or rear access to reduce service times and increase asset utilization. A catastrophic failure in one server does not impact the functioning of the other server, because the two four-socket servers are designed to be electrically isolated. Two power supplies are dedicated to each server to deliver power in a 1+1 redundant configuration. In addition, all PCIe slots, HBAs, and service processors, along with the BIOS and operating system, are independent and unique to each server.

Unique Eight-Socket Glueless Architecture

Oracle designed the single eight-socket configuration of Oracle Server X8-8 with a “glueless” architecture in which the number of hops from any processor to any other processor is a maximum of two, compared to as many as three hops for designs using a “glued” approach. Using a spoke-and-wheel connectivity pattern, the glueless design enables the fastest possible performance for database and enterprise applications, because it minimizes the latency for shared-memory access.

In addition, a glueless architecture removes the need for a node controller. Eliminating the node controller saves power and improves reliability. The power savings come from reducing the number of components needed to interconnect the processors. The reliability enhancements are realized by eliminating non-Intel components in the inter-processor communication path.

Best for Scale-Up Enterprise Applications

While many organizations rely on scale-out architectures for their web and virtualization tiers, large scale-up architectures remain the optimal choice for organizations that need extreme performance for specific enterprise applications. Because Oracle Server X8-8 provides hundreds of threads and industry-leading memory density, it is the ideal x86 platform for scale-up applications that require large amounts of memory and I/O executing on a single-instance operating system.

The large memory footprint and bandwidth of Oracle Server X8-8 are well suited to run memory-intensive applications such as financial trading applications, batch processing workloads, and applications that require large amounts of simulation or computation.

Extreme CPU Performance and Memory Capacity for Oracle Database

Oracle Server X8-8 in both single eight-socket and dual four-socket configurations is an ideal x86 platform for running Oracle Database. 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.

Oracle's x86 systems 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.

Using the combination of Oracle Database 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.


Oracle Database In-Memory on Oracle Server X8-8

Oracle Database In-Memory adds in-memory database functionality to existing databases, and transparently accelerates analytics by orders of magnitude while simultaneously speeding up mixed-workload online transaction processing (OLTP). This breakthrough capability is enabled by the "dual-format" architecture of Oracle Database In-Memory. Up to now, databases have forced users to store data in either column or row format. Column format is highly efficient for analytics, but imposes very large overheads when used in OLTP environments. Similarly, row format enables extremely fast OLTP, but is less optimized for analytics. The only way to optimize for both OLTP and analytics has been to copy data from OLTP systems to analytic systems using complex extract, transform, and load (ETL) processes that add a great deal of expense and latency.

The dual-format architecture of Oracle Database In-Memory eliminates this trade-off by representing tables simultaneously using traditional row format and a new in-memory column format by taking advantage of systems with large amounts of system memory. Configured with up to 6 TB of memory using 64 GB LR-DIMMs, Oracle Server X8-8 provides an average of 32 GB of memory per core. This makes Oracle Server X8-8 the ideal platform for using memory to optimize database performance using Oracle Database In-Memory.

Database Consolidation on Oracle Server X8-8 using Oracle Multitenant

IT organizations have traditionally used virtualization and clustering technologies to consolidate their databases, and many have embarked on major application redevelopments to consolidate database schemas. This has typically resulted in limited consolidation density, increased management costs and, in many cases, high development costs. Oracle Multitenant simplifies the consolidation process by plugging multiple databases into a multitenant container database without changing applications. In this new architecture, memory and background processes are required only at the multitenant container database level, enabling IT organizations to achieve a greater level of scalability and consolidation density without compromising the security of previously separate database silos.



Oracle Multitenant, an Oracle Database 12c Enterprise Edition option, introduces a new architecture that enables customers to easily consolidate multiple databases without changing their applications. This new architecture delivers all the benefits of managing many databases as one, yet retains the isolation and resource prioritization of separate databases. In addition, Oracle Multitenant enables rapid provisioning and upgrades, and fully complements other options including Oracle Real Application Clusters (Oracle RAC) and Oracle Active Data Guard.

Oracle Server X8-8 is uniquely positioned to be the optimal platform for database consolidation using Oracle Multitenant because it scales I/O, memory, and compute power together, allowing for high consolidation without any drawbacks. The large memory capacity allows for many databases to be consolidated while also taking advantage of the in-memory optimizations of Oracle Database. In addition, the hot-swappable and serviceable I/O components provide better uptime for mission-critical databases.

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 infrastructure as a service (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 X8-8 offers 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.

Oracle Integrated Lights Out Manager 4.0

Oracle Server X8-8 includes the latest generation of cloud-ready service processor that is designed for today's security challenges. Oracle ILOM 4.0 is a modern 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 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 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 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">» 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 X8-8 also provides 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 X8-8 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 server 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 X8-8 is designed completely in house from the ground up and is 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 server such as memory DIMMs, hard disk drives, 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.

Fault Management and Diagnostics

With higher levels of integration among various subsystems in the server, 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 X8-8 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 X8-8 includes built-in fault management and diagnostic tools that increase system availability and enable faster service response times that increase server uptime. Oracle Server X8-8 includes 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 X8-8 overcomes 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 and Oracle Solaris on Oracle Server X8-8 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

Oracle continues to deliver products that simplify IT and reduce operating expenses. Oracle Server X8-8 is designed, optimized, and pretested for specific Oracle software workloads. The chassis design for Oracle Server X8-8 that allows for both eight-socket and dual four-socket configurations delivers extreme performance and flexibility when running any Oracle Database workload. In addition, Oracle's x86 servers are the most reliable, highest performing x86 servers on the market, driving simplification through innovation.

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 X8-8, further improving its performance and reliability—making it 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 X8-8 can be found at oracle.com/servers/x86 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 © 2019, 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 X8-8 System Architecture
March 2019