ORACLE
x86 SERVERS

Oracle Server X5-8 System Architecture

ORACLE WHITE PAPER | JUNE 2016

ORACLE

Table of Contents

Introduction

Oracle Server X5-8, Oracle's new eight-socket x86 server, is part of a family of Oracle x86 servers that are purpose-built to be best for running Oracle software. Oracle Server X5-8 is a scale-up server optimized for database and enterprise workloads requiring large amounts of memory and compute power. This white paper describes the architecture of the Oracle Server X5-8 system.

Product Overview

Oracle Server X5-8 provides up to eight Intel Xeon E7-8895 v3 processors, each of which contains 18 cores operating at 2.6 GHz with 45 MB of level 3 (L3) cache. Each CPU module houses one Intel Xeon E7-8895 v3 processor along with 24 dual inline memory module (DIMM) slots. The Intel Xeon E7-8895 v3 processor is capable of functioning in many different modes. Building on the success of elastic computing in its previous generation eight-socket server, Oracle Server X5-8 offers customers the ability to modulate between the number of active cores and the maximum Intel Turbo Boost frequency, making it the highest performing x86 eight-socket server in the industry. Further, this system offers unprecedented flash capacity of up to 25.6 TB using high-bandwidth low-latency NVM Express (NVMe) technology.

With up to 192 DDR3 low-voltage DIMMs, Oracle Server X5-8 provides up to 6 TB of memory and is ideal for running medium-to-large databases entirely in memory.

The system's industry-leading 5U enclosure contains 16 PCIe 3.0 slots (eight 16-lane slots and eight 8-lane slots), eight 2.5-inch drive bays, and four USB ports (two front and two internal). The system supports up to eight SAS-3 hard disk drives (HDDs) or solid-state drives (SSDs), allowing for up to 9.6 TB of internal HDD storage or up to 3.2 TB of internal SSD storage with an additional 25.6 TB of flash via eight Oracle Flash Accelerator F320 PCIe Cards.

Figure 1 and Figure 2 show the Oracle Server X5-8 chassis.

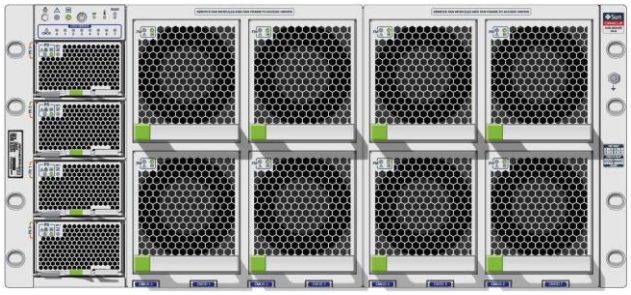

Figure 1. Front chassis view

Figure 2. Rear chassis view

Oracle Server X5-8 is designed for maximum uptime with enterprise-grade availability features. All I/O modules and disks are hot-swappable and support RAID 0, 1, 6, 10, 50, and 60. The power supplies and fans are also redundant and hot-swappable, ensuring that a failure to any single component will not affect the running system. With four power supplies, the system offers N+N power redundancy.

The dual PCIe card carrier (DPCC) is an electro-mechanical subassembly that enables hot-pluggable I/O support to industry-standard low-profile PCIe cards. Each DPCC has one 8-lane and one 16-lane PCIe Gen 3 slot, along with buttons and indicators that allow users to perform I/O hot-plug functions. A number of I/O option cards are supported that enable Ethernet, Fibre Channel, InfiniBand, SAS, and flash connectivity to Oracle Server X5-8.

Each system includes built-in enterprise management tools that reduce operating expenses by finding hardware problems early and making initial setup easy. Oracle Integrated Lights Out Manager (Oracle ILOM) provides remote power control; virtual keyboard, video, mouse (KVM) capability; advanced health monitoring; and remote server configuration.

Designed to Meet Today's Security Challenges

According to the Department of Homeland Security, imported software and electronics are often shipped to the United States with purposely embedded malware, spyware, and security-compromising components by unknown foreign parties. There has been concern about supply-chain security, because computers and IT equipment pass through several suppliers before the final product is deployed. A federal report released on the supply chain between the United States and foreign nations speculated the possibility that somewhere along the line someone could compromise a component and design a "back door entry" capability that could enable cyber attacks.

One example of such a cyber attack has been christened "BIOS Plot." An NSA analyst discovered that a nation state had the intention to destroy computers—via the BIOS—used by US financial institutions.

As recently as December 2015, a security breach was discovered at a major US computer networking equipment manufacturer. US officials worried that hackers working for a foreign government were able to spy on the encrypted

communications of the US government and private companies for years. It is believed that attackers embedded a "back door" into the source code of the communication protocols of the equipment.

These types of attacks are just a few examples of how hackers are becoming more and more sophisticated at attacking multiple layers in the IT stack. It is no longer good enough just to secure applications and the network perimeter of a data center; the enterprise must apply security in depth across hardware, firmware, and software.

Oracle's Approach to Security

Oracle's philosophy on security in-depth is based on the philosophy that "security needs to be built in and 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. The 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. An example alert can be found at:

http://www.oracle.com/technetwork/topics/security/alert-cve-2016-0603-2874360.html

Securing the Foundation—100 Percent In-House Design and Manufacture

The entire x86 server product line is designed 100 percent in-house. No third parties ever touch the motherboard design, ensuring that no components are added to create a "back door entry" into Oracle servers. Additionally Oracle applies strict control over the entire supply chain with all of its servers being manufactured in the United States, thus maximizing supply chain security. This is unique to Oracle.

The firmware installed in Oracle's x86 servers, such as the BIOS and system management stack, are developed and owned by Oracle with no source code ever released to third parties. The Oracle Integrated Lights Out Manager (Oracle ILOM) system management stack is FIPS 140-2 compliant, ensuring the latest cryptography ciphers are supported. Oracle ILOM also includes other security-related features such as fine-grained access control and logging that enable IT administrators to control and monitor access to the infrastructure. Oracle has incorporated technologies into Oracle ILOM that ensure that illegal firmware updates are prevented.

In addition, Oracle ILOM makes sure that newly unpacked and connected Oracle x86 servers are secure "out of the box," because only secure protocols such as HTTPS, SNMP, and IPMI are allowed, while untrustworthy connections are rejected. The Oracle ILOM service processor is ubiquitous across all of Oracle's engineered systems, storage appliances, SPARC servers, and x86 servers—ensuring that common security, reliability, and manageability features are applied across all platforms.

All of these features are embedded within the servers themselves and there are no additional licensing fees to be paid for them.

Best for Oracle Software

Oracle x86 systems are the best x86 platforms for Oracle software. Only Oracle provides customers with an optimized hardware and software stack that comes complete with a choice of operating system, virtualization software, and cloud management tools—all at no extra charge. Oracle's optimized hardware and software stack has enabled a 10x performance gain in its engineered systems and delivered world-record benchmark results. Oracle's

comprehensive, open standards–based x86 systems provide the best platform to run Oracle software with enhanced reliability for data center environments.

Oracle Linux

Oracle Server X5-8 can be combined with one of Oracle's best-in-class operating systems to provide the most reliable, highest performing combination for database and enterprise applications. Oracle offers its own enterprise version of Linux, which is called Oracle Linux. Oracle Linux has been optimized to run best on Oracle hardware through enhancements at all levels including the kernel, which is known as Oracle's Unbreakable Enterprise Kernel.

Specifically, Oracle Linux has been engineered to scale up for large CPU and memory configurations. Further, Oracle Linux has been designed and extensively tested to handle the large number of processor cores and the memory footprint of Oracle Server X5-8. In addition, Oracle Linux has been improved with 10 GbE optimizations and I/O and block-layer enhancements that improve application reliability and performance.

Oracle Database with Oracle Multitenant Option

Oracle Database 12*c* offers a new option called Oracle Multitenant. Historically, IT architects have used server-level virtualization to consolidate small two-socket servers running multiple database instances into fewer physical servers that run each database instances on guest operating systems. While this methodology provided CapEx and OpEx benefits at the system administrator–level, it had quite the opposite effect for database administration by introducing complexities that are inherent to having multiple tenants reside on the same physical server. One of the most notable challenges of this approach is ensuring security and adequate isolation between multiple tenants. Other challenges introduced concern the inability of the database instances to scale with the demands of users. Oracle Multitenant provides a secure, scalable, and cost-effective solution for hosting multiple database instances as pluggable instances in a single container. As shown in Figure 3, the database administrator now has complete control of the operating environment and has a secure and simple way of managing a single database instance for scalability, patching, and data protection for backups and disaster recovery.

The business benefit of running Oracle Multitenant on Oracle Server X5-8 is derived from the fact that there are not only fewer physical servers to manage and patch, but additionally the overall utilization of the infrastructure is greatly enhanced. The physical memory footprint is now used to store business data as opposed to allocating adequate memory to run multiple guest operating systems as well as multiple copies of Oracle Database. This increases the consolidation density within the infrastructure, thereby reducing CapEx.

The large memory footprint of Oracle Server X5-8 enables each pluggable database instance to elastically consume larger chunks of memory based on demand, without being constrained by a guest operating system. Further, the memory management for individual pluggable databases ensures performance scaling based on the demands of each pluggable database. High availability configurations can be easily provisioned and managed, because these can be set and managed at the container level as well as at the level of the individual pluggable database.

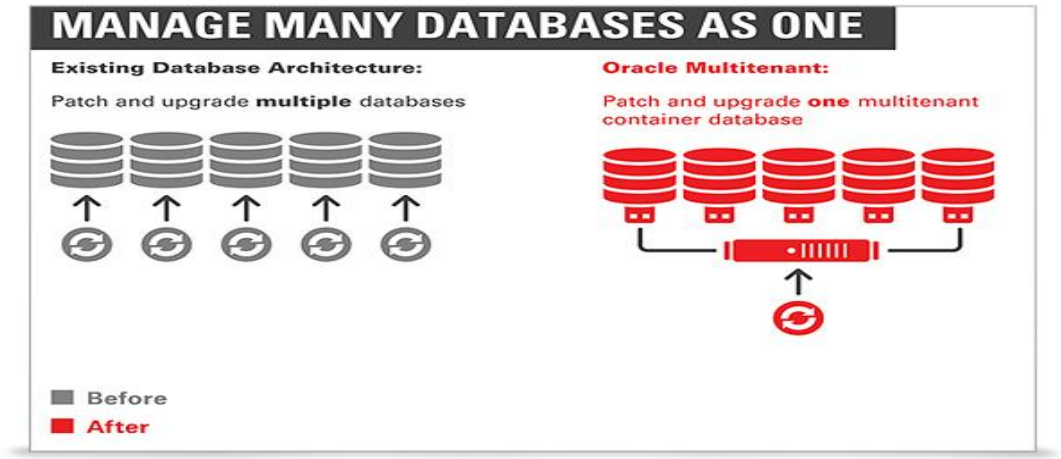

Figure 3. Advantages of an architecture based on Oracle Multitenant

Oracle Database In-Memory Option

Oracle Database runs faster and more reliably when running on Oracle hardware. Oracle Server X5-8 has been engineered to work best when running Oracle Database, with optimizations to both the hardware and firmware. Specifically, the storage controller, BIOS, and network interface card (NIC) firmware have been optimized for Oracle Database. Oracle has built in hundreds of enhancements including improvements to cache recovery, path failover, device failure detection, I/O handling, SCSI recovery, and disk failure and recovery.

At Oracle OpenWorld 2013 in San Francisco, Oracle announced the upcoming availability of the Oracle Database In-Memory option—a solution for accelerating database-driven business decision-making to real-time speeds. Unlike specialized approaches that are restricted to particular workloads or applications, the unique approach of the Oracle Database In-Memory option leverages a new in-memory column store format to speed up analytic, data warehousing, and reporting workloads, while also accelerating online transaction processing (OLTP) workloads.

Organizations of all sizes running any application can benefit from real-time analytics performance and acceleration for all workloads. Virtually every existing application that runs on top of Oracle Database runs dramatically faster by simply turning on the new Oracle Database In-Memory feature. By supporting up to 6 TB of memory, Oracle Server X5-8 is ideal for running more of the database in memory and taking advantage of the new in-memory features of Oracle Database.

Oracle's Unique NVMe Design and Database Smart Flash Cache

As depicted in Figure 4 and Figure 5, Oracle's NVMe-based flash cards have a much lower latency and higher bandwidth than standard SAS-3 drives due to the fact that the drive connects directly to four lanes of the PCIe Gen 3 card with an aggregate bandwidth of 32 Gb/sec, as opposed to 12 Gb/sec for a traditional SAS-3 SSD. In addition, the NVMe flash technology is optimized to accelerate Oracle Database using an Oracle Database feature called Database Smart Flash Cache.

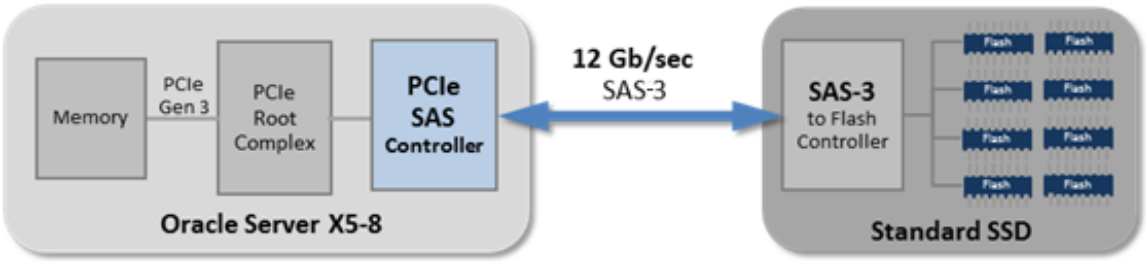

Figure 4. Traditional SAS-3 SSD architecture

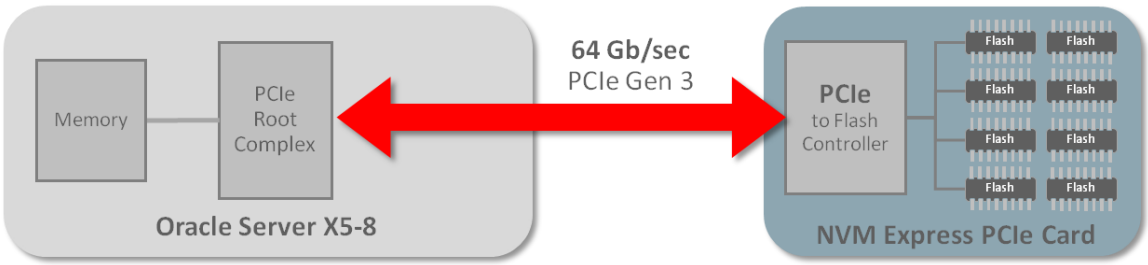

Figure 5. Oracle's NVMe SSD architecture

Oracle Server X5-8 can be configured with up to eight Oracle Flash Accelerator F320 PCIe Cards to support a total of up to 25.6 TB of high-bandwidth low-latency flash.

Because flash technologies are temperature sensitive, most high-performance flash drives are enabled to throttle down I/O speeds as temperatures rise in order to protect the flash from damage. Oracle's implementation, on the other hand, includes multiple temperature sensors that are monitored by the Oracle Server X5-8 system's Oracle ILOM service processor (SP) to ensure the flash devices maintain an optimum operating temperature. Oracle ILOM modulates the fan speed to ensure sufficient cooling for maximum system performance at all times. The benefit of this design is that the system consistently operates at maximum performance across the full operating temperature range of the server independent of system configuration.

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 X5-8 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.

Extreme CPU Performance

With up to eight Intel Xeon E7-8895 v3 processors, Oracle Server X5-8 provides up to a total of 144 cores in a 5U enclosure. In addition to having the highest core count of any Intel processor, each processor contains 45 MB of shared L3 cache, which is also the highest for any Intel processor. With integrated PCIe 3.0 and Intel Turbo Boost 2.0 technology, the Intel Xeon E7-8895 v3 processors provide extreme performance and elasticity.

Very Large Memory Footprint

With up to 6 TB of memory, Oracle Server X5-8 provides an average of over 51 GB of memory per core. This enables applications to have access to a large amount of memory as well as memory bandwidth. The system's large memory footprint and bandwidth 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.

Glueless Architecture

Oracle designed Oracle Server X5-8 with a "glueless" architecture (shown in Figure 6) 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 (shown in Figure 7). 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.

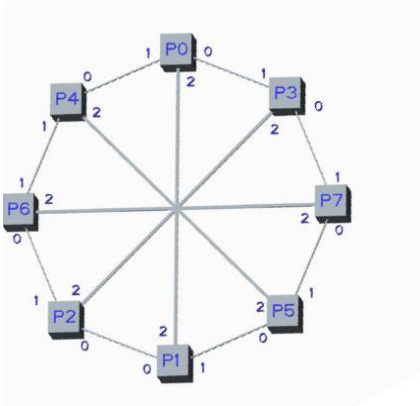

Figure 6. Glueless architecture

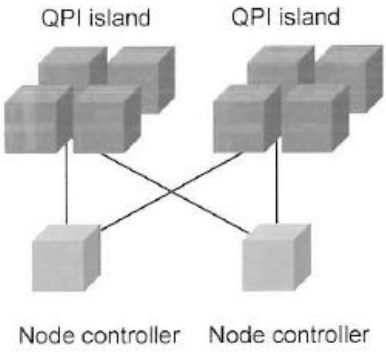

Figure 7. Glued architecture

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 interprocessor communication path.

Node controllers typically need to change from one processor generation to another because of changes to coherency link protocols and other modifications to interprocessor communication. The elimination of the node controller enables Oracle to offer a future-proof chassis that will support three processor releases from Intel.

Elastic Computing with Oracle Server X5-8

Oracle Server X5-8 offers a new elastic computing capability that dramatically reduces operating expenses. While Intel provides a wide range of processor SKUs, each SKU has historically offered only a fixed combination of core count, operational frequency, and power consumption. Customers have, therefore, been forced to make tradeoffs when they select a particular processor SKU, and they've had to base their selection on a particular workload that needs to be known upfront.

Oracle and Intel worked jointly to define a new processor SKU, the Xeon E7-8895 v3 processor, that has unique characteristics and effectively combines the capabilities of three different Xeon processor SKUs into a single processor SKU (see Figure 8). This processor SKU is the only one offered on Oracle Server X5-8.

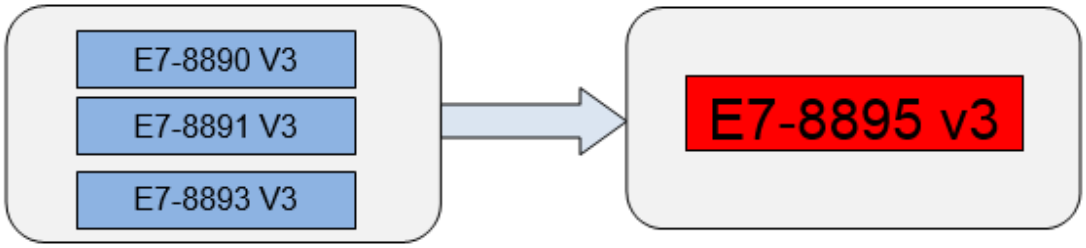

Figure 8. Elastic computing with a single processor SKU

The Xeon E7-8895 v3 processor has been designed to span the performance characteristics of the three processor SKUs shown on the left side of Figure 8. Oracle system design engineers worked closely with the operating system development teams to achieve the ability to transform the core count and operating frequency of the Xeon E7-8895 v3 processor to vary with time without the need for a system-level reboot. This simplifies the process of determining system configuration, streamlines the purchasing process, and decreases the system management overhead through dynamic repurposing of common assets. In addition, power consumption is reduced by maximizing CPU utilization.

Along with the new processor SKU, enhancements have been made to the system BIOS, Oracle Solaris, and Oracle Linux, which allow the processors in the system to dynamically clock up to faster speeds as cores are disabled and to reach higher maximum turbo frequencies for the remaining active cores.

Building on Intel's Turbo Boost technology, Oracle Server X5-4 and Oracle Server X5-8 systems offer customers the ability to modulate between the number of active cores and the maximum Turbo Boost frequency. Customers can choose to run with fewer cores at a higher frequency or with more cores at a lower frequency, all while staying within the same power profile.

This single processor SKU can adapt to the requirements of different workloads based on its runtime configuration. For example, the processor can be configured for transaction processing at a higher speed for one hour and then be switched to higher core counts for the next hour for higher throughput computing.

Figure 9 shows a comparison of the Intel Xeon E7-8890 v3 processor to the Intel Xeon E7-8895 v3 processor.

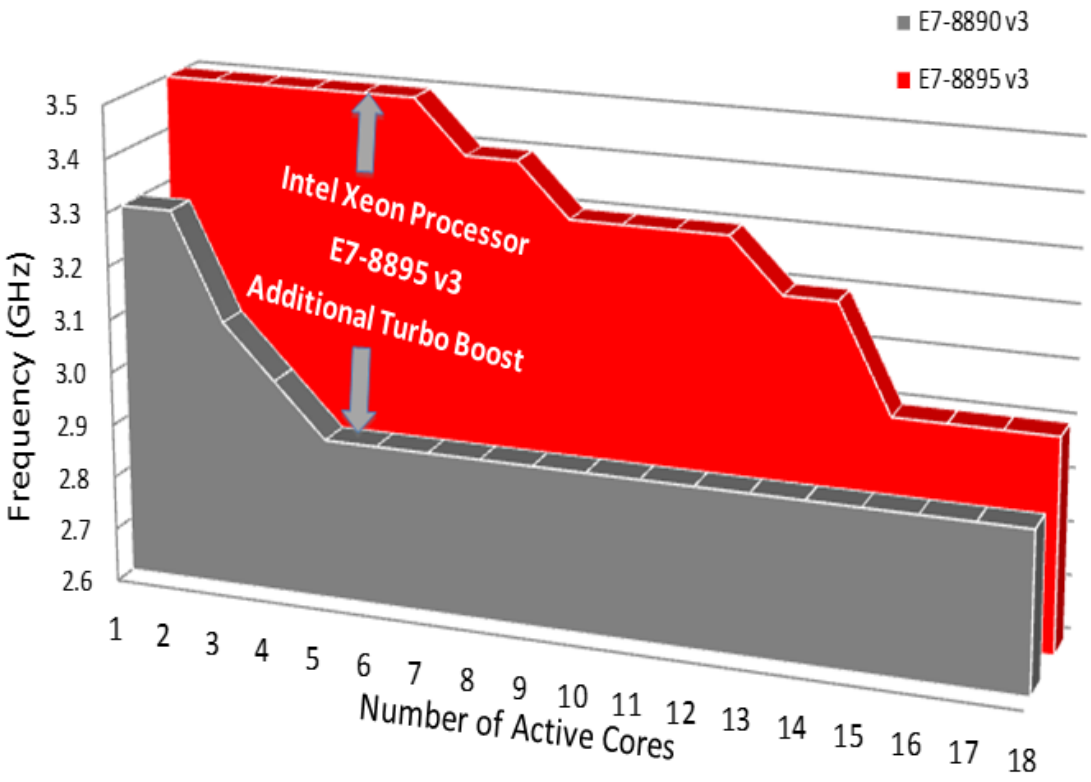

Figure 9. Comparison of Intel Xeon E7-8895 v3 processor to Intel Xeon E7-8890 v3 processor

System Design

Oracle Server X5-8 is a fully modular system, engineered to be easily serviceable. The eight sockets are connected through a passive midplane that maximizes availability and allows for maximum intersocket bandwidth. Oracle Server X5-8 can be configured in either a four-socket or eight-socket configuration, with each socket attached to up to 24 DIMMs.

Each CPU module (CMOD) contains a single Intel Xeon E7-8895 v3 processor and 24 DIMM slots. The processor in each CMOD connects to a total of four memory buffers with 6 DIMM slots attached to each buffer. The CMODs are located just behind the front fans and are front-panel-accessible by removing the fans.

In the rear of the chassis there is a single system module (SMOD) that houses the peripheral control hub functions, the Oracle ILOM service processor, and a PCIe slot for RAID storage HBA connectivity to drive bays. The SMOD also provides connectivity with one VGA, one management, one serial, two USB 2.0, and two GbE ports in the rear of the chassis. Also in the rear are eight hot-swap-capable disk drive bays.

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.

Cooling for Oracle Server X5-8 is front to back using eight variable-speed, dual-counter rotating fan assemblies mounted in the front of the system.

Figure 10 shows the system block diagram and the connections between each CPU to the DIMMs through a memory buffer. Each CPU is connected in a glueless design such that there are never more than two hops between each processor, which allows for the lowest possible latency for a given processor to access memory or I/O that is attached to a different processor.

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. In addition, processor 0 (CPU0) interfaces with the eight disk slots in the SMOD.

Figure 11 shows an exploded view of the chassis.

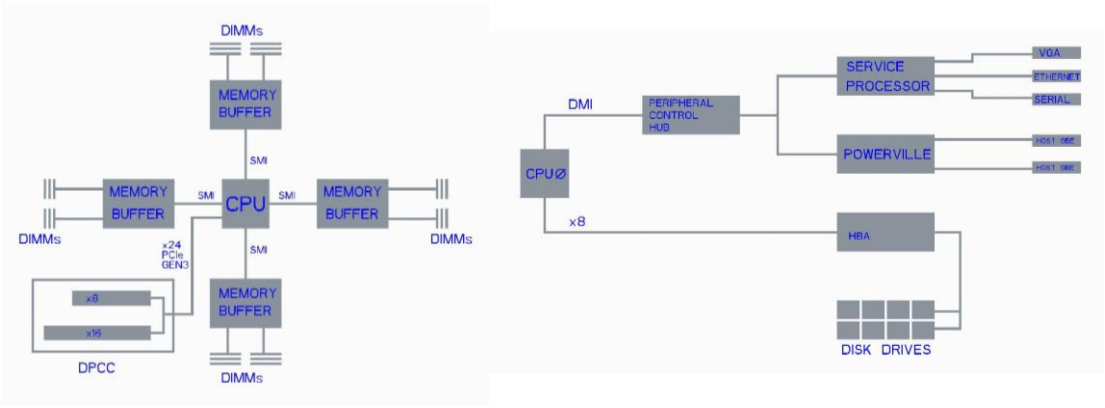

Figure 10. Oracle Server X5-8 System Block Diagram

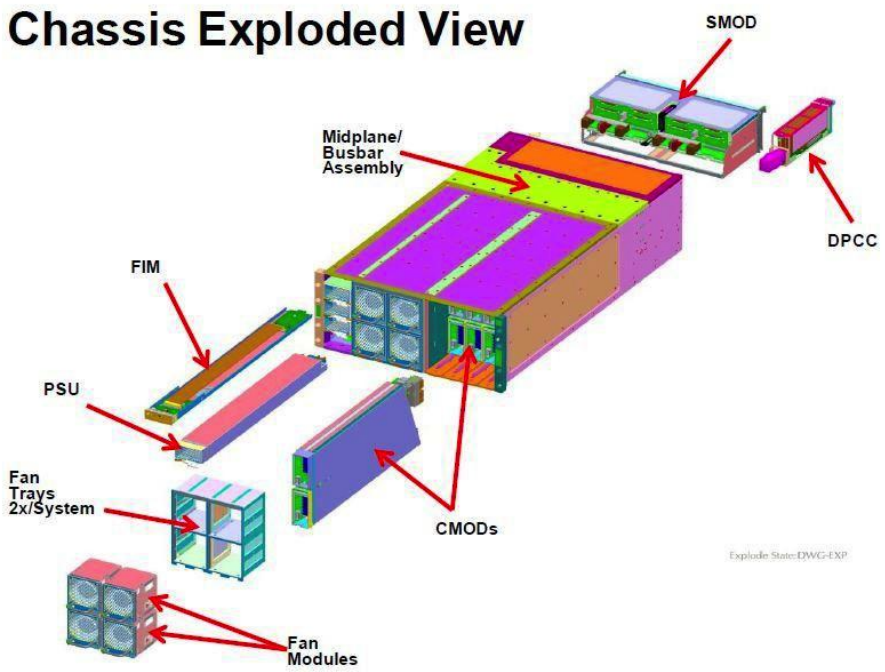

Figure 11. Chassis exploded view

Single-Pane-of-Glass Management

Oracle Enterprise Manager 12*c* is a suite of system management tools that provides a single-pane-of-glass management solution for the entire Oracle stack. This enables organizations to manage their Oracle Server X5-8 systems from the hardware layer all the way up to the database and the applications running on them.

Oracle Enterprise Manager Ops Center 12*c*, part of the Oracle Enterprise Manager family, is an enterprise management tool that allows IT staff to manage all aspects of their servers. In addition to providing detailed hardware monitoring and reporting for hardware problems, Oracle Enterprise Manager Ops Center can provision a bare-metal system with an operating system and also configure virtualization.

Oracle Enterprise Manager Cloud Control 12*c*, also part of the Oracle Enterprise Manager family, can be used to implement a private cloud on Oracle Server X5-8 systems. Oracle Enterprise Manager Cloud Control provides a complete cloud lifecycle management solution enabling customers to quickly set up, manage, and support enterprise clouds and traditional Oracle IT environments from applications to disk.

Conclusion

Oracle continues to deliver products that simplify IT and reduce operating expenses. Oracle Server X5-8 is designed, optimized and pretested for specific Oracle software workloads. The elastic computing capability of Oracle Server X5-8 was codesigned with Intel and provides immediate business value for Oracle customers. Oracle is uniquely positioned to provide a full solution for elastic computing due to its ability to coengineer the system design with Oracle operating systems and Oracle Database. In addition, Oracle's x86 servers are the most reliable, highest performing x86 servers on the market, driving simplification through innovation.

Oracle Corporation, World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries
Phone: +1.650.506.7000
Fax: +1.650.506.7200

ORACLE

CONNECT WITH US

blogs.oracle.com/oracle
facebook.com/oracle
twitter.com/oracle
oracle.com

Integrated Cloud Applications & Platform Services

Oracle Server X5-8 System Architecture
June 2016, Version 1.0
Author: Subban Raghunathan