ORACLE

SOLARIS

An Oracle White Paper
November 2011

What Oracle Solaris Brings to Oracle Exalogic Elastic Cloud

Disclaimer

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Introduction

The Oracle Exalogic Elastic Cloud is an *Engineered System*, consisting of software, firmware and hardware, on which enterprises may deploy Oracle business applications, Oracle Fusion Middleware or software products provided by Oracle partners. Exalogic is designed to revolutionize datacenter consolidation, enabling enterprises to bring together tens, hundreds, or even thousands of disparate, mission-critical, performance-sensitive workloads with maximum reliability, availability, and security. Exalogic dramatically improves the performance of virtually any standard Linux, Oracle Solaris, or Java application[1] with no code changes required. It also reduces application implementation and ongoing costs versus traditional enterprise application platforms and private clouds assembled from separately sourced components provided by multiple competing vendors.

Oracle Solaris, one of two operating system options for the Exalogic, is a key building block that provides superior reliability, performance, and security, coupled with unique features, such as highly efficient, built-in operating system virtualization for application consolidation. This paper provides a high level overview of the Oracle Exalogic Elastic Cloud and then describes the value that Oracle Solaris brings to the Oracle Exalogic offering.

[1] Any application that supports Oracle Linux (version 5 update 5 or later, Unbreakable Enterprise Kernel, 64-bit), Oracle Solaris 11 (or Oracle Solaris 10 Zone, x86)

Oracle Exalogic System Overview

Oracle Exalogic is an engineered system comprising both hardware and software components, each of which is a strategic technology in Oracle's product portfolio. Figure 1 shows how the software components in Exalogic are layered on top of the Oracle Exalogic Elastic Cloud hardware.

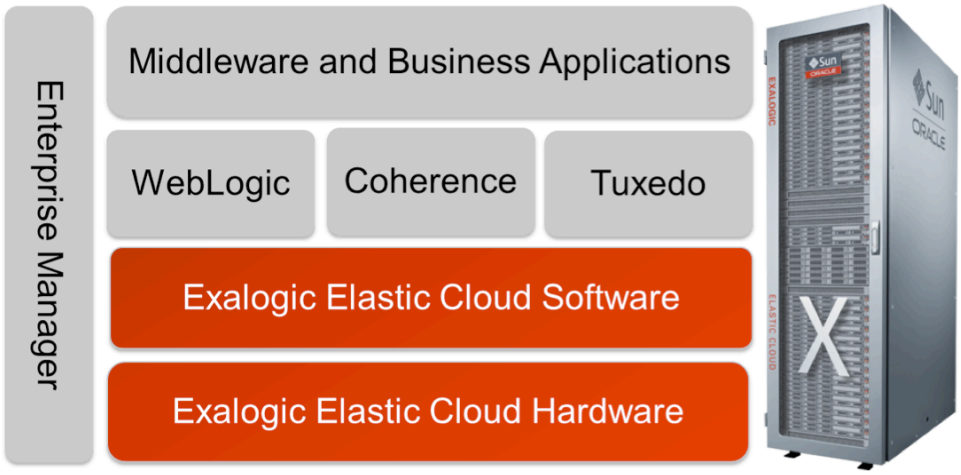

Figure 1. Oracle Exalogic consists of hardware and software engineered together.

Exalogic hardware is pre-assembled and delivered in standard rack configurations. Each Exalogic configuration is a unit of elastic cloud capacity balanced for compute-intensive workloads. Each contains a number of hot-swappable compute nodes, a clustered, high-performance disk storage subsystem, and a high-bandwidth interconnect fabric. The fabric comprises the switches needed to connect every individual component within the configuration as well as to externally connect additional Exalogic or Exadata Database Machine racks. In addition, each configuration includes multiple 10 Gigabit Ethernet ports for integration with the datacenter service network. Gigabit Ethernet ports are used for integration with the datacenter‘s management network. All Exalogic configurations are fully redundant at every level and are designed with no single point of failure.

Each Exalogic compute node is a fully self-contained unit of compute capacity with multi-core x86 Intel Xeon processors, redundant power supplies, fast ECC DIMM memory, and redundant InfiniBand Host Channel Adapters. Each compute node also contains two solid-state drives (SSDs), which host the operating system images used to boot the node and act as high-performance local swap space as well as storage for diagnostic data generated by the system during fault management procedures.

InfiniBand is fundamental to the Oracle Exalogic Elastic Cloud system. In addition to providing an extremely fast, high-throughput interconnect between all of the internal hardware units, it also provides extreme scale, application isolation, and elasticity.

The Oracle Exalogic Elastic Cloud hardware platform is a specific combination of compute capacity, memory configuration, networking and storage aimed at middleware workloads. As a result of careful construction of the hardware, the software can be specifically configured and tuned to take advantage

of the numbers of processors and cores on each compute node, the exact networking speed, the locality of flash storage, and the performance characteristics of the storage infrastructure. Even within a hardware infrastructure other than the Exalogic system, this can be a challenging tuning exercise given the breadth of configurability available at the operating system, Java VM, WebLogic Server, and Coherence levels. Within the Oracle Exalogic Elastic Cloud, Oracle has pre-tuned the software infrastructure using existing features and has also engineered significant new architectural changes into the software to dramatically improve performance beyond what is possible on traditional hardware and software infrastructures.

Oracle Exalogic Elastic Cloud is the world's first engineered system specifically designed to provide enterprises with a foundation for a secure, mission-critical private cloud capable of virtually unlimited scale, unbeatable performance, and previously unimagined management simplicity. Exalogic is the ideal platform for applications of all types, from small-scale departmental applications to the largest and most demanding ERP and mainframe applications. While Exalogic is optimized for enterprise Java, Oracle Fusion Middleware, and Oracle's Fusion Applications, it is also an outstanding environment for the thousands of third-party and custom Linux and Oracle Solaris applications widely deployed today. Simply put, Exalogic is a giant step forward in realizing Oracle's vision for the datacenter of the 21rst century.

Oracle Solaris

Oracle Solaris is the strategic platform for today's demanding enterprise, delivering proven results on everything from mission critical enterprise databases to high-performance Web farms. For customers facing challenging business and technical requirements—such as lowering costs, simplifying system administration, and maintaining high service levels—Oracle Solaris is the ideal cross-platform choice.

Oracle Solaris brings all the key advantages of a mission-critical operating system to Exalogic while benefiting from the many advantages of an engineered system. Java applications running on Oracle Solaris or in Oracle Solaris Zones (operating system virtualization discussed below) benefit from all the middleware and InfiniBand tuning, the interconnect fabric itself, the integration, and the extensive testing that go into Exalogic.

Oracle Solaris Zones

Oracle Solaris includes a built-in virtualization capability called Oracle Solaris Zones, which allow users to isolate software applications and services using flexible, software-defined boundaries. Unlike hypervisor-based virtualization, Oracle Solaris Zones provide OS-level virtualization, which gives the appearance of multiple OS instances rather than multiple physical machines. Oracle Solaris Zones enable the creation of many private execution environments from a single instance of the operating system, with full resource management of the overall environment and the individual zones.

The nature of OS virtualization means that Oracle Solaris Zones provide very low-overhead, low-latency environments. This makes it possible to create hundreds, even thousands, of zones on a single system. Full integration with Oracle Solaris ZFS and network virtualization provides low execution and

storage overhead for those areas as well, which can be a problem area for other virtual machine implementations. Oracle Solaris Zones enable close to bare metal performance for I/O, making these software components an excellent match for the outstanding I/O performance throughout of the Exalogic system. Each zone can make full use of the performance characteristics of the InfiniBand infrastructure and take advantage of the InfiniBand system and middleware tuning that went into the Exalogic (for full details see the overview paper on Oracle Exalogic Elastic Cloud[2]).

Oracle Solaris 11 provides a fully virtualized networking layer. An entire datacenter network topology can be created within a single OS instance using virtualized nets, routers, firewalls and NICs. These virtualized network components come with high observability, security, flexibility, and resource management. This provides great flexibility while also reducing costs by eliminating the need for some physical networking hardware. The networking virtualization software supports quality of service, which means that appropriate bandwidth can be reserved for key applications.

Oracle Solaris Zones are also integrated with Oracle Solaris DTrace, the Oracle Solaris facility that provides dynamic instrumentation and tracing for both application and kernel activities. Administrators can use DTrace to examine Java application performance throughout the software stack. It provides visibility both within Oracle Solaris Zones and in the global zone, making it easy for administrators to identify and eliminate bottlenecks and optimize performance.

Greater Efficiency Through Virtualization and Consolidation

The Oracle Exalogic Elastic Cloud system is an outstanding consolidation platform by design as discussed above in the Exalogic summary. Oracle Solaris builds on this capability with its combination of low overhead virtualization with Oracle Solaris Zones and the ability to fully benefit from the performance advantages of the Exalogic interconnects and tuning. The result is a system with great flexibility in how applications are deployed and moved within (and between) Exalogic systems. The benefits include improved application performance along with reduced overall hardware, floor space, heating and licensing costs.

Figure 2 shows an example of how Oracle Solaris Zones and the network virtualization can be used to consolidate an application environment to help reduce costs and increase customer satisfaction. A traditional practice has been to provide one dedicated server per application, providing maximum scalability, security, and performance on demand. In this scenario, as shown in the top portion of Figure 2, system utilization is typically very low except for relatively rare periods of peak demand. By consolidating these applications onto a single, larger node within Oracle Exalogic, cost and complexity are greatly reduced and yet each application component still resides in a completely secure environment. Because of the extremely low-overhead of Oracle Solaris Zones, hundreds of zones can

[2] http://www.oracle.com/us/products/middleware/exalogic-elastic-cloud-overview-489031.pdf

run on a single medium-sized system. The application benefits derived from the Exalogic architecture accrue to each of those applications now running on a single node. And there are many more nodes available for other applications.

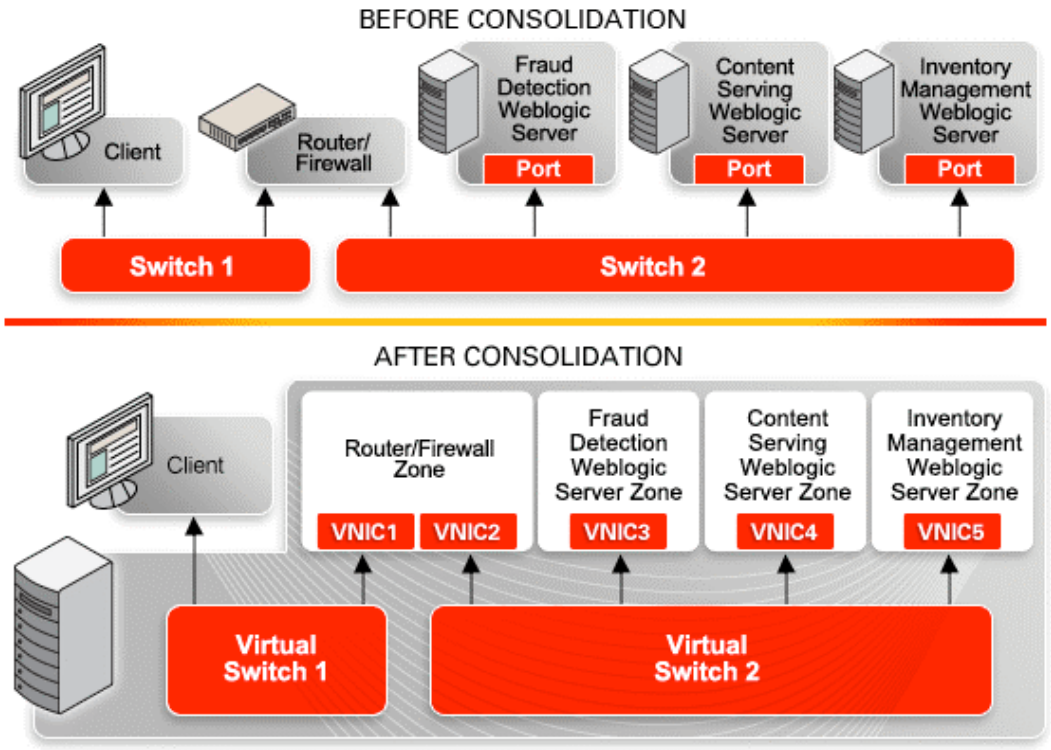

Figure 2. Oracle Solaris Zones consolidation.

Improved Agility

Oracle Solaris Zones also make it easier to move application components from node to node to gain further agility with the application infrastructure. Movement of Oracle Solaris Zones can be done to balance workloads or offload a server that needs maintenance. And the movement can be within an Exalogic system, to an external server, or to a node in another Exalogic system. Movement of zones to a remote Exalogic node can be part of a disaster recovery plan.

The movement of an Oracle Solaris Zone and its application contents will be fastest if the system data is kept on shared disks to minimize copying. Unlike hypervisor-based virtual machines, which have to start up the entire operating system plus the application, the start time for Oracle Solaris Zones is very quick due to their being part of an already running operating system. Thus, the time to deploy a new application is simply the start time of the application itself plus a short start up time for the zone. Oracle Solaris Zones enable administrators to rapidly react to end user demand. If another 10 instances of WebLogic are needed, they can be deployed very quickly with Oracle Solaris Zones.

Moving Existing Application Environments to Oracle Exalogic

Oracle Solaris Zones also greatly simplify application deployment on the Exalogic Elastic Cloud. Application environments currently running on Oracle Solaris 10 can be moved to Oracle Solaris 11 as an Oracle Solaris 10 Zone using physical-to-virtual migration[3]. Existing Oracle Solaris 10 Zones can be moved over directly.

Having made the transition to running as an Oracle Solaris 10 Zone on Oracle Solaris 11, these application environments benefit from overall Exalogic and virtualization advantages including:

- Applications can leverage the engineered system benefits of Exalogic.
- Consolidation can help reduce power, footprint, heat and licensing costs.
- Disaster recovery is simplified by configuring remote Exalogic systems into the architecture.

Performance and Scalability

Oracle Solaris is designed to provide superior performance and scalability. It optimizes the system CPU, memory, I/O, and network resources for a broad range of mission-critical workloads. Built-in technologies that help deliver increased application performance include:

- High-performance networking stack (see InfiniBand discussion below)
- File system performance tuned to take full advantage of solid-state drives, which are the system disks for Exalogic nodes
- Tools and libraries
- Multiple page-size support
- Memory placement optimization

Building on a proven track record, Oracle Solaris is ready to take advantage of the unique performance capabilities inherent in the Exalogic Elastic Cloud. Oracle Solaris DTrace helps developers isolate performance bottlenecks when designing application code, and allows system administrators to safely analyze and resolve a broad range of issues in production environments. Significant performance innovation also comes from Oracle Solaris optimizations for both individual processing cores and the overall multicore microarchitectures, increasing both single-threaded and multithread performance. As a result, the Oracle Solaris kernel and existing single or multithreaded applications run faster with no code changes or recompilation necessary. There are performance improvements in critical areas such as the networking stack and disk subsystems too, providing maximum application throughput.

[3] http://www.oracle.com/technetwork/articles/servers-storage-admin/p2vvirtualizationmigration-170693.pdf

As discussed in some detail in the overview paper on Oracle Exalogic Elastic Cloud[4], InfiniBand has a significant role in delivering application performance. As shown in Figure 3, Oracle Solaris provides the core software layers that support the InfiniBand capabilities and optimizations discussed in this paper. InfiniBand is not only high bandwidth, but also delivers low latency. With an overall well-tuned system and middleware, extreme Java performance is the result.

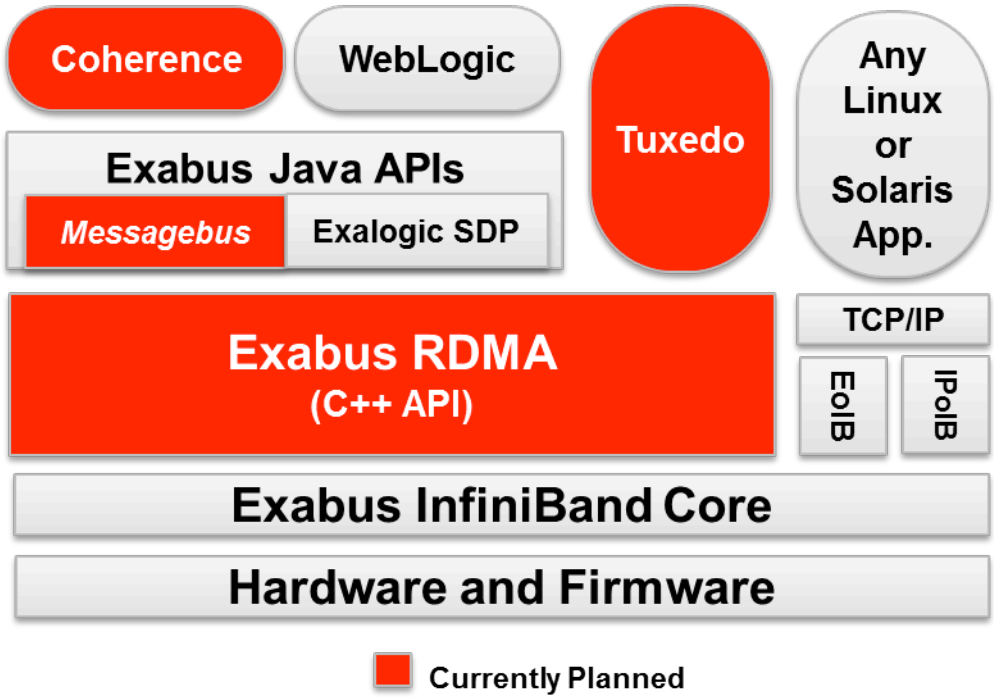

Figure 3. Exalogic InfiniBand infrastructure includes optimized software components in the Oracle Solaris OS.

Unmatched Reliability

Oracle Solaris brings proven datacenter reliability to Exalogic with a long history of industry-leading reliability features and innovation.

Key benefits include:

- **Continuous uptime**—Oracle Solaris **Predictive Self-Healing**[5] diagnoses, isolates, and helps to recover from hardware and application faults. The Oracle Solaris Service Manager Facility manages application services in a uniform way and attempts to recover failed services using dependency

[4] http://www.oracle.com/us/products/middleware/exalogic-elastic-cloud-overview-489031.pdf

[5] http://www.oracle.com/us/products/servers-storage/solaris/solaris-pred-self-healing-ds-075587.pdf

analysis to restart services in the proper order and ensure that all dependent services are addressed. Oracle Solaris Zones isolate the impact of application errors by confining failures to a secure environment.

- **End-to-end data integrity**—Oracle Solaris ZFS[6] delivers dramatic advances in data management with an innovative approach to data integrity, including protection from silent data corruption with Oracle Solaris ZFS check summing.
- **Safely diagnose problems on the fly**—Oracle Solaris DTrace[7] allows administrators to safely analyze, tune, and troubleshoot applications on production systems with little or no performance impact.

Leading Security

Oracle Solaris includes leading security technologies and is designed to prevent security problems as well as protect against security vulnerabilities. Oracle is committed to software assurance, building security into the design, building, testing, and maintenance of Oracle Solaris and Exalogic (http://www.oracle.com/us/support/assurance/index.html).

Core security features for Oracle Exalogic include:

- Data at rest—Encrypt datasets on Exalogic storage to protect data and meet compliance requirements.
- Data in transport—Build upon the Oracle Solaris Cryptographic Framework or OpenSSL to accelerate Web applications with the integrated hardware cryptographic co-processors.
- Secure by default—Oracle Solaris ships with minimal services running and only one networked service available (ssh).
- User management—Define granular access permissions based on traditional UNIX access control or reduce risk further by using Role Based Access Control (RBAC) to limit the threat of malicious insiders.
- Application/user containment—Guard against unsafe applications or users by granting only the privileges needed for execution.
- Delegated administration for zones—A user or role can be configured with additional abilities to manage specific zones.

[6] http://www.oracle.com/us/products/servers-storage/solaris/solaris-zfs-ds-067320.pdf

[7] http://www.oracle.com/us/products/servers-storage/solaris/solaris-dynamic-tracing-ds-067318.pdf

Investment Protection through Application Compatibility

Oracle Solaris delivers binary compatibility from release to release and source compatibility between Oracle Solaris on the SPARC or x86 architectures.

With a single source code base, Oracle Solaris delivers the same features across both platforms. This means that enterprises can develop and optimize applications on Oracle Solaris on either system and deliver them on both architectures.

Moving Existing Applications to Oracle Exalogic

This application compatibility enables three paths to move applications from existing systems to an Exalogic system:

- Migrate an existing Oracle Solaris 10 server environment over as an Oracle Solaris 10 Zone.
- Move each application separately (or as a group) to an Oracle Solaris 11 Zone or system using binary compatibility.
- Recompile the source code (using source code compatibility) to move from an existing SPARC system to an Exalogic system.

Simplified Management with Unified Software Update Administration

Oracle Solaris 11 introduces a new network-based image packaging system (IPS) to ease the acquisition, installation, and maintenance of Oracle Solaris and additional Oracle, third-party, and in-house applications. IPS offers:

- Dependency-aware tools to update the entire software stack while virtually eliminating opportunities for errors during updates
- Safe reload to a previous boot environment with a simple reboot
- Tight integration with Oracle Solaris Zones to further simply update management
- Minimized downtime by reducing the amount of time a system is offline and out of service
- Simpler and more effective end-to-end experience for system administrators and service providers, thereby helping to reduce costs

Summary

Oracle Exalogic Elastic Cloud is an engineered hardware and software system that offers the IT infrastructure needed to meet today's rising demands for consolidation in the datacenter. Oracle Exalogic Elastic Cloud, the worlds first and only integrated middleware machine, dramatically surpasses alternatives and provides the best possible foundation for running enterprise applications.

Oracle Solaris delivers key datacenter advantages in performance, reliability, scalability, and security for Exalogic systems. Oracle Solaris Zones provide a lightweight virtualization capability that takes full advantage of the performance advantages of Exalogic InfiniBand tuning. Oracle Solaris DTrace provides observability up and down the application through server stack to tune application performance. Customers coming from Oracle Solaris systems today will be able to take advantage of physical to virtual movement to Oracle Solaris 10 Zones and/or source or binary compatibility.

Resources

Oracle Exalogic Elastic Cloud	www.oracle.com/exalogic
Oracle Solaris 11	www.oracle.com/solaris
Oracle Solaris 11 Technical Resources	www.oracle.com/technetwork/server-storage/solaris11/overview/index.html

What Oracle Solaris Brings to Oracle Exalogic Elastic Cloud
November 2011

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com

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

Hardware and Software, Engineered to Work Together