

An Oracle White Paper
March 2011

Oracle Exalogic Elastic Cloud: A Brief Introduction



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

For most enterprise IT organizations, years of innovation, expansion, and acquisition have resulted in sprawling infrastructure that stretches the limits of manageability. While the individual IT systems and applications in service are often well considered and expertly implemented, the sheer scale of the ongoing IT investment itself has emerged as the dominant concern. Even when best-of-breed technologies, open standards, market-leading vendors, and modern architectural practices like SOA have been employed pervasively, most enterprises now find themselves with too many platforms, too many technologies, too many domains of expertise, and too many vendors to coordinate and manage.

In response, a number of technologies and practices have become staples for large enterprises, ranging from virtualization and centralized storage to enterprise-wide standardization of software and hardware. Recently, however, the discussion of specific cost containment techniques has given way to a larger discussion of the transformation of IT from cost center to profit center. This transformation typically involves adoption of a more centralized, automated, and elastic infrastructure which is commonly referred to as *private cloud*.

Real-world approaches that truly deliver on the promise of private cloud will involve systems engineered together for maximum performance while balancing openness, reliability, cost, flexibility, and resource efficiency. Perhaps most importantly, these engineered systems must not be monolithic and it must be possible for enterprises to implement them over time, at a pace determined by real business needs and prudent investment timelines.

Oracle Exalogic Elastic Cloud is the world's first engineered system specifically designed to provide enterprises with a foundation for 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 Solaris applications widely deployed today. Simply put, Exalogic is a giant step forward in realizing Oracle's vision for the datacenter of the 21st century.

Exalogic System Overview

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. Oracle Exalogic Elastic Cloud consists of hardware and software engineered together.

Oracle Exalogic Elastic Cloud Hardware

Exalogic hardware is pre-assembled and delivered in standard 19" 42U rack configurations. Each Exalogic configuration is a unit of elastic cloud capacity balanced for compute-intensive workloads. Each Exalogic configuration contains a number of hot-swappable compute nodes, a clustered, high-performance disk storage subsystem, and a high-bandwidth interconnect fabric comprising 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 and Gigabit Ethernet ports 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 either multi-core x86 Xeon or SPARC-T3 processors, redundant power supplies, fast ECC DIMM memory, and redundant InfiniBand Host Channel Adapters. Each compute node also contains two solid-state disks (SSDs), which host the operating system images used to boot the node and act as high-performance local swap space and storage for diagnostic data generated by the system during fault management procedures.

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

There are two traditional approaches to growing a datacenter's compute capacity, the most basic approach is vertical scaling. Vertical scaling adds compute capacity to an individual computer, but with major limitations. There are limits to how much a single computer can scale with a balanced configuration. Also, vertical scaling increases the impact of a single system failure. More recently, horizontal scaling has become more common and is accomplished by networking together many individual computers using basic networking technologies like Ethernet. Horizontal scaling increases both compute capacity and the tolerance of individual system failure. However, horizontally expanded systems function as a collection of separate computers that require coordination. This coordination is challenged by throughput, latency and lack of high-end cluster features in basic Ethernet networks.

By contrast, the lossless switched InfiniBand I/O fabric used by the Exalogic system connects all systems together in a way that forms a single large computer.

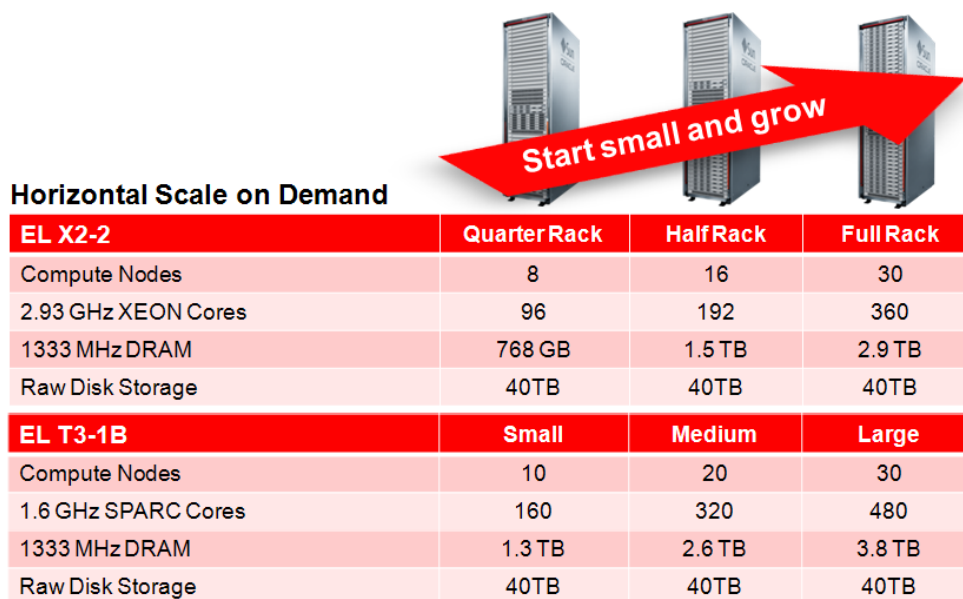


Figure 2. Extreme Scalability with Exalogic

It is possible to start as small or quarter rack, grow to medium or half a rack or a to a large full rack and, beyond this one can connect as many as eight full racks of Exalogic hardware (or any combination of Exalogic and Exadata configurations) together without the need for any external switches. In cases where more than eight racks of Exalogic or Exadata hardware are required, Oracle offers a choice of several high-capacity datacenter switches which allow the creation of Exalogic clouds comprising hundreds of racks and tens of thousands of processors.

Exalogic systems scale horizontally, meaning that there is no degradation of system performance as the size of the cloud increases. Equally importantly, an Exalogic cloud is a resource pool that can be dynamically sub-divided into secure units of capacity. The underpinnings of this capability are in the design of InfiniBand itself, which was designed for precisely this use. InfiniBand supports *partitions*, in which communication between end-points on the I/O fabric is strictly controlled within the fabric switches. Individual compute nodes, or even specific I/O devices, may be grouped into logical partitions, within which communication is allowed. Communication between logical partitions, however, can be controlled at the lowest level.

Augmenting the security of InfiniBand partitions is another feature of InfiniBand called *virtual lanes*. Each I/O end-point communicates over the I/O fabric using one or more of these virtual lanes, each of which is fully independent of the others and may be assigned a priority, thereby ensuring that applications may have guaranteed access to shared resources. In the case, for example, of an extremely mission-critical application, it is possible using these features for an administrator to select a number of compute nodes, assign all of the I/O devices on those nodes to a secure partition, and then assign virtual lanes for exclusive access to shared storage, Exadata Database Machine resources, and external service network ports. More importantly, if capacity requirements change, compute nodes may be added to, or removed from, the desired partition dynamically through simple configuration.

Oracle Exalogic Elastic Cloud Software

Java is the most successful and pervasive application implementation technology in use by enterprises today. Exalogic has been designed from the ground up to provide the ideal environment for enterprise Java applications and Java-based infrastructure. Oracle's entire Fusion Middleware portfolio is designed for deployment on Exalogic. Ultimately, as the Oracle Fusion Applications portfolio is released, they too will be deployable on Exalogic. Oracle Exalogic Elastic Cloud Software includes a number of optimizations and enhancements made to the core products within Oracle WebLogic Suite, the essential Java foundation on which Oracle's next-generation applications are being developed. Oracle WebLogic Suite includes Oracle WebLogic Server, Oracle Coherence, Oracle JRockit, and Oracle HotSpot.

In addition to unique support for Java applications and Oracle Fusion Middleware, Exalogic also provides users with a choice of Oracle Linux or Oracle Solaris operating systems. Exalogic is 100% compatible with all standard Oracle Linux 5 and Solaris 11 applications, and no special certification for Exalogic is required – all Oracle applications that are certified for the appropriate releases of Oracle Linux and Solaris are supported on Exalogic.

Exalogic Elastic Cloud Management

Oracle Enterprise Manager provides application-to-disk management through Grid Control and OpsCenter. Enterprise Manager allows every individual hardware component within an Exalogic deployment to be monitored in real time and, at the customer's option, have system status automatically reported to Oracle Support for proactive system maintenance. Through integration with Oracle Support, Enterprise Manager can apply tested patch bundles tailored for Exalogic that cover

every layer of the system, from device firmware and operating system to JVM, application server, upper-stack Fusion Middleware, and Oracle Applications.

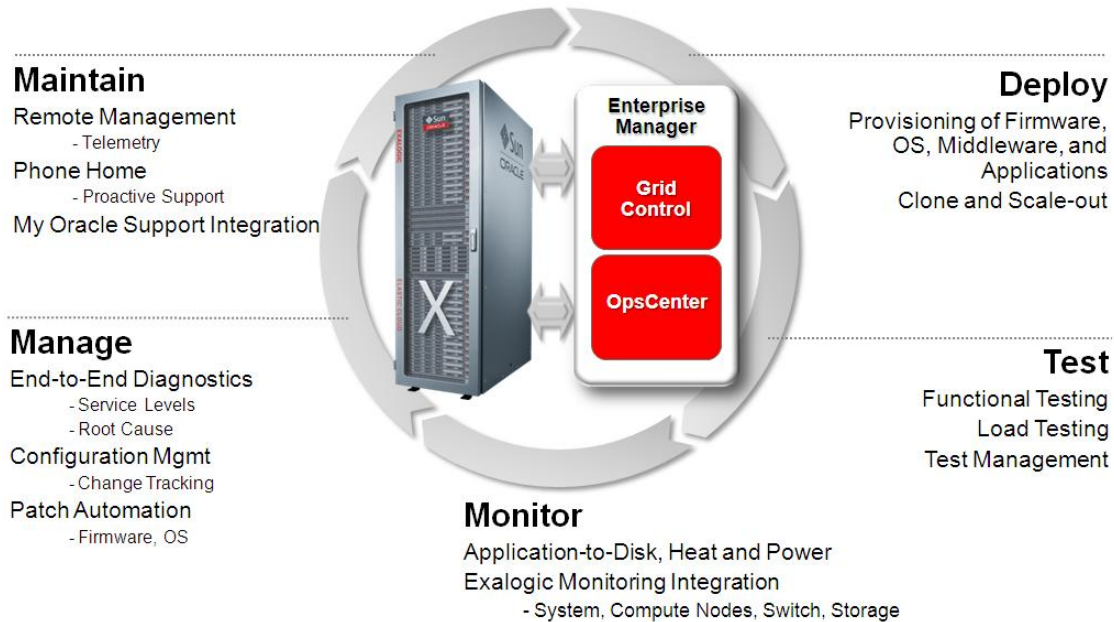


Figure 3. Exalogic software includes Oracle Enterprise Manager for comprehensive system management.

Exalogic Elastic Cloud Software Technical Optimizations

The Exalogic Elastic Cloud Software encapsulates a set of enhancements made to Oracle WebLogic Suite, for optimized performance when running on Exalogic hardware. The WebLogic Suite of products enhanced are the WebLogic Application Server, the JRockit Java Virtual Machine (JVM) and the Coherence In-Memory Data Grid. The software optimizations address performance limitations that only become apparent when the software is running on Exalogic's high-density computing nodes and fast-networking InfiniBand switches. With these software optimizations, the WebLogic Suite of products can utilize the benefits of this high-end hardware to the maximum, resulting in a well balanced hardware-software engineered system.

At a high-level, the key benefits provided by the optimizations can be grouped into 4 main categories:

1. *Increased WebLogic Scalability, Throughput and Responsiveness:* Improvements to WebLogic's networking, request handling and thread management mechanisms, which enable it to scale better on the high multi-core compute nodes that are connected to the fast InfiniBand fabric that ties all the compute nodes together. The net effect is that each WebLogic server can handle more client requests whilst also reducing the time taken to respond to each individual request.

2. *Superior WebLogic Session Replication Performance:*
WebLogic's session replication mechanism for Exalogic is improved to utilize the large InfiniBand bandwidth that is used between WebLogic servers for parallel connections over the network. The net effect is choke free communication between WebLogic servers, resulting in superior performance for Java web applications that require high availability.
3. *Oracle RAC Integration for More Reliable Database Interaction:*
A new technology component called “Active GridLink” has been added to WebLogic on Exalogic that provides optimized WebLogic to Oracle RAC database connectivity and provides intelligent load-balancing across RAC nodes with faster connection failover if a RAC node fails. The net effect is higher throughput and faster response time for enterprise Java applications that are involved in intensive database work.
4. *Reduced Response Times for Exalogic to Exadata Communication:*
In situations where an Exalogic system is directly connected to an Exadata system, using InfiniBand, a “native” InfiniBand networking protocol called SDP is used to interact with the Oracle RAC database on Exadata. This results in low latency for request-response times for calls between WebLogic and the database. The performance gain is most significant when large results sets are transferred from the database. The net effect is applications are able to respond to client requests faster, leading to overall performance gain for enterprise Java applications.

Benefits of Exalogic Elastic Cloud

Extreme Java Performance

The combination of Oracle Exalogic software and hardware results in substantial performance gains for Java-based applications running on WebLogic Server and other Oracle Fusion Middleware technologies. To understand the magnitude of gains achievable by running applications on Exalogic, we compare Exalogic performance to a typical alternative configuration¹ in three representative examples.

¹ A standard reference system has been developed by Oracle based on analysis of typical customer environments in use. This reference system uses the same number and type of Intel processors and the same amount of RAM and local disk storage. The reference system uses standard Ethernet network adapters and switches rather than Exalogic's InfiniBand fabric. Identical versions Oracle Linux and Oracle WebLogic Suite software are deployed to both the Exalogic system and the reference system. For deployment to Exalogic, all of the Oracle Exalogic Elastic Cloud Software enhancements are activated and configured, while the software deployed to the reference system uses default tuning and configuration. The Exalogic software enhancements cannot be used on non-Exalogic hardware.

Application Responsiveness (Latency)

We looked at several frequently-used operations in a typical Web application such as ‘Create Web Service’, ‘Purchase’, ‘Manage’, and ‘Browse’, and have measured improvements of as much as 14x in comparison with a standard platform.

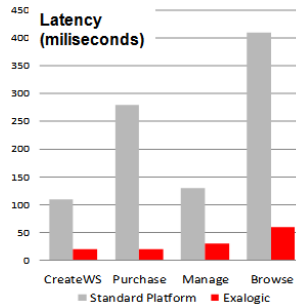


Figure 4. Exalogic improves Web application responsiveness by up to 14x compared with a typical alternative.

Application Capacity (Throughput)

We compared the execution of core Java business logic using a representative benchmark test and saw an increase in application capacity—the number of operations per second—by as much as 60%.




Figure 5. Exalogic improves execution of core Java business logic by 60%.

Database Communication

In a third investigation, we looked at database-intensive application performance—in other words, what can Exalogic do for transactional applications that depend on frequent interactions with a database. Again, when comparing Exalogic with a typical hardware configuration for such an application, we found that Exalogic’s superior processing combined with superior I/O to the database gives a 2-3x improvement in performance for database operations.

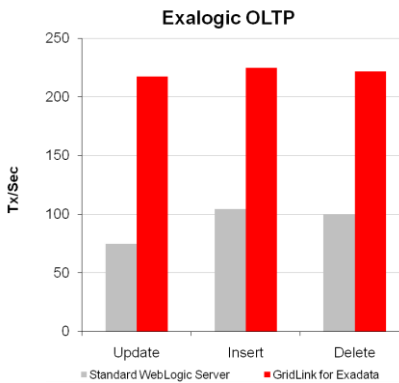


Figure 6. Exalogic accelerates database interactions by 2-3x.

Engineered Together Means Fast Set-Up and Efficient Operation

In addition to extreme performance, the fact that Exalogic’s hardware and software have been engineered together also means that customers are required to do far less in setting up and running Exalogic than for alternative environments. Exalogic is tuned for a wide range of workload types—process- intensive, data-intensive, I/O- intensive, etc.—and thus there is no need to configure parameters for a particular environment. This translates to a reduction in set-up time by as much as 95%. But not only does this mean that there is less work to do to set up, it also means that all customers run the same Exalogic configuration that was tuned and tested at the Oracle factory, substantially reducing risk of errors, simplifying diagnosis, and enabling more efficient, lower-cost operation. Total costs can be reduced by as much as 60%.

Exalogic Elastic Cloud - Foundation for Mission-Critical Cloud

Many enterprise cloud efforts to-date have focused on consolidating non-mission-critical workloads such as testing environments due to lack of experience, trust, or confidence in new technologies and approaches. With Exalogic, enterprises can surpass or leapfrog such steps and bring the benefits of consolidation and elastic capacity to mission-critical, production workloads.

The Exalogic performance benefits are directly relevant for mission-critical applications. The maturity and design excellence of both Exalogic hardware and software confer a reliability that is second to none. The unique delivery guarantees and robust partitioning enabled by the InfiniBand interconnect fabric mean that widely varying workloads can run together on Exalogic with requisite isolation. And the wide scope of workload types supported by Exalogic’s at-the-factory tuning mean that different workloads get Exalogic’s performance benefits without individualized configuration that could compromise other workloads or lead to operational complexity.

Conclusion

The need for enterprise IT organizations to provide next-generation cloud features such as elastic capacity while meeting ever more demanding performance and reliability requirements is driving demand for a new approach to infrastructure. Whether workloads are Web-based or thick-client, whether data-intensive or processing-intensive, whether homogeneous or highly heterogeneous, the key to success is hardware and software engineered together for performance, reliability, and scale. Building or using custom, special-purpose systems for different applications is wasteful and expensive. Oracle Exalogic Elastic Cloud, the world's first and only integrated middleware machine, dramatically surpasses alternatives and provides enterprises the best possible foundation for running applications.

By consolidating applications to Oracle Exalogic Elastic Cloud, enterprises will:

- Accelerate the performance of Java applications by as much as 14
- Improve reliability and scalability beyond even the most mission-critical requirements
- Reduce deployment effort by up to 95% and reduce costs by as much as 60%

We invite you to begin your datacenter transformation with Exalogic today.



Oracle Exalogic Elastic Cloud: A Brief
Introduction
September 2010
Authors: Mike Piech, Mike Palmeter
Mike Lehmann

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

Copyright © 2010, 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.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. 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. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0410

SOFTWARE. HARDWARE. COMPLETE.