

An Oracle White Paper
January, 2011

Oracle Exalogic Elastic Cloud: A Brief Introduction



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.

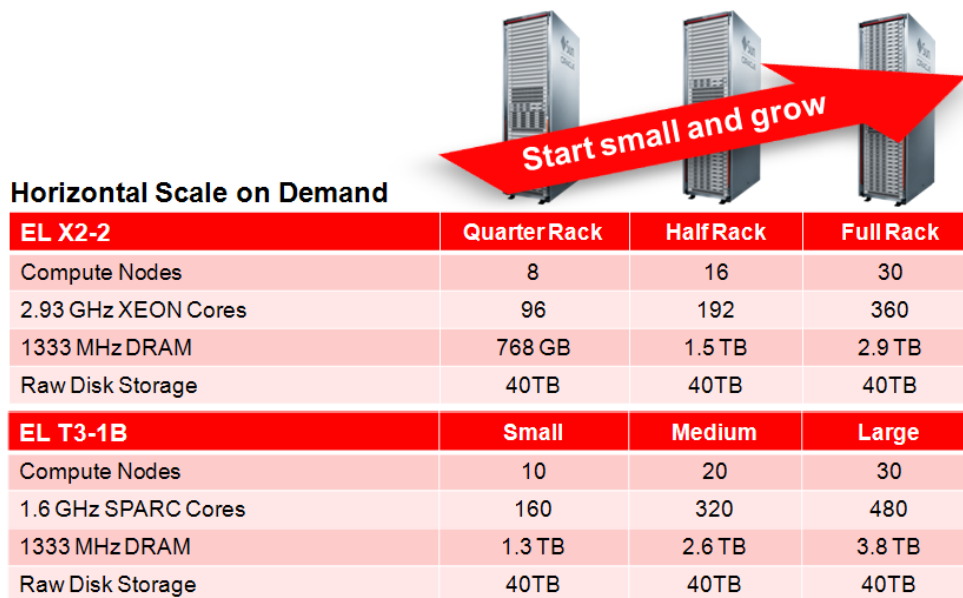


Figure 2. Exalogic is highly scalable.

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 subdivided 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.

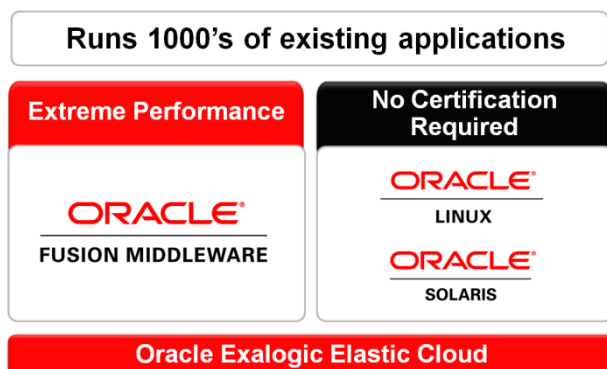


Oracle Exalogic Elastic Cloud Software

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.

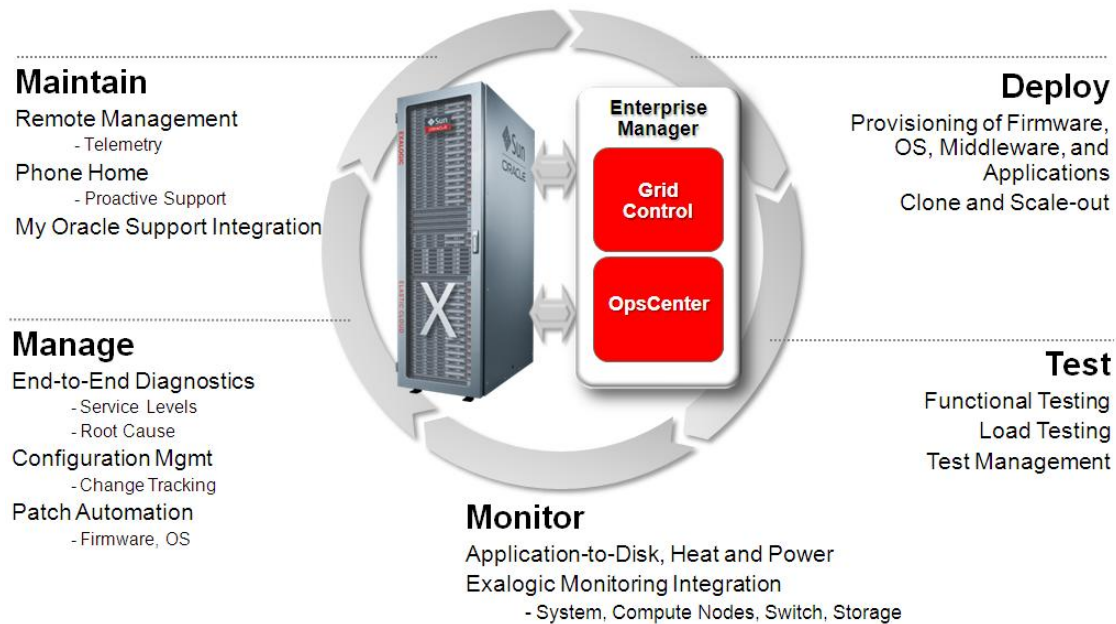


Figure 3. Exalogic software includes Oracle Enterprise Manager for comprehensive system management.

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.

Exalogic Elastic Cloud Software Technical Optimizations

Within the software infrastructure running on Exalogic, there have been specific technical optimizations made to WebLogic Server, Coherence and JRockit and Hotspot that bring dramatic performance benefits. These software optimizations come in two major categories:

1. *Software Optimizations for Exalogic Hardware:* The 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 non-Exalogic hardware infrastructure this can be a challenging tuning exercise given the breadth of configurability available at the operating system, Java VM, the WebLogic Server and the Coherence levels. Within the Exalogic Elastic Cloud not only has Oracle pre-tuned the software infrastructure using existing features, significant new architectural changes have been engineered into the software to dramatically improve performance beyond what is possible on non-Exalogic hardware.

2. *Software Optimizations for Oracle Database:* A critical feature of the Exalogic Elastic Cloud is how well it integrates with Oracle RAC and specifically Oracle RAC running on Oracle Exadata. The performance benefits come in three major areas with RAC and Exadata:
 - a. Exalogic can be directly integrated with Exadata using InfiniBand as a high speed switchless fabric connecting to the two hardware solutions together. As a result of the InfiniBand technology, such configurations effectively become one large high-speed compute fabric running middleware grids and database grids.
 - b. To ensure that the connectivity between Exalogic and Exadata is not only at the InfiniBand fabric level but fully available in the WebLogic runtime infrastructure, the JDBC driver and the WebLogic Server data source infrastructure have been engineered to communicate natively using protocols available on InfiniBand.
 - c. Finally, a significant innovation has been introduced as part of the Exalogic Elastic Cloud Software within WebLogic Server called Active GridLink for RAC. This feature enables WebLogic Server to connect to RAC instances, both on Exadata and independent of Exadata, by augmenting JDBC connections with a new event based channel for RAC instances to share status information back to WebLogic Server.

These optimizations for Exalogic communicating with RAC dramatically improve the performance of OLTP applications by factors of up to 2 to 3 times as compared to WebLogic Server running on standard hardware.

Exalogic Elastic Cloud Software Optimizations for Exalogic Hardware

Perhaps most important to the software infrastructure within Exalogic is the use of InfiniBand as the internal networking fabric. InfiniBand is an industry standard switched fabric communications link widely used in high-performance computing and enterprise data centers.

For Exalogic, Quad Data Rate (QDR) InfiniBand is used, enabling up to 40Gb/s bandwidth with a latency (Message Passing Interface MPI ping) of as low as 1.2 microseconds. The native connection protocol of InfiniBand is Socket Direct Protocol (SDP), but standard IP over InfiniBand (IPoIB) is built in as well, enabling any TCP/IP based solution to run on Exalogic. In addition to higher bandwidth and lower latency, another key characteristic of InfiniBand is the ability to send significantly

larger packets over the network level. With IPoIB the Maximum Transmission Unit (MTU) is fixed at 64K, and with SDP the MTU starting size is 32K and then self-tuned to an arbitrarily larger size for maximum throughput by InfiniBand.

Figure 4 provides a high-level technical view of the InfiniBand infrastructure on Exalogic.

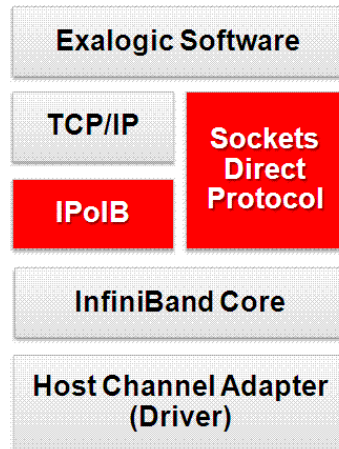


Figure 4. InfiniBand on Exalogic.

Reduced Buffer Copy

Within this highly optimized networking infrastructure the first area necessary to update was the WebLogic Server I/O infrastructure and how network traffic to and from WebLogic Server was read and written to the socket layer in InfiniBand. The goal was to optimize the inflow and outflow of WebLogic traffic onto the InfiniBand fabric. Here the focus was on two areas:

1. Several years ago, with the introduction of Native Input Output (NIO) in Java SE 1.4, WebLogic Server expanded its existing internal native I/O solution to adopt that of Java SE proper. In Exalogic this approach was further optimized by fully taking advantage of Java NIO ByteBuffers to avoid copying the message buffer content to (or from) any intermediate buffer before (or after) each invocation of one of the underlying operating system's native I/O operations. The underlying Java VM of JRockit and Hotspot, have been also been tuned to support this highly streamlined networking architecture.

This optimization was carried through not only to the core kernel infrastructure of WebLogic Server but also into key constructs in the Java EE programming model in areas such as the WebLogic Server servlet runtime, resulting in a dramatically reduced amount of buffer copying at the programming model, application server runtime, Java VM, and ultimately network levels. This directly complements and augments InfiniBand on Exalogic where there also is a reduced level of network copies by the nature of the protocol itself.

To give a sense of the magnitude of this change, within WebLogic Server proper, four buffer copies were eliminated – effectively reducing the internal copying by over 65%, if one were to

follow the Java I/O path from a typical JSP to servlet into the WebLogic Server kernel. Figure 6 provides a pictorial summary of this optimization.

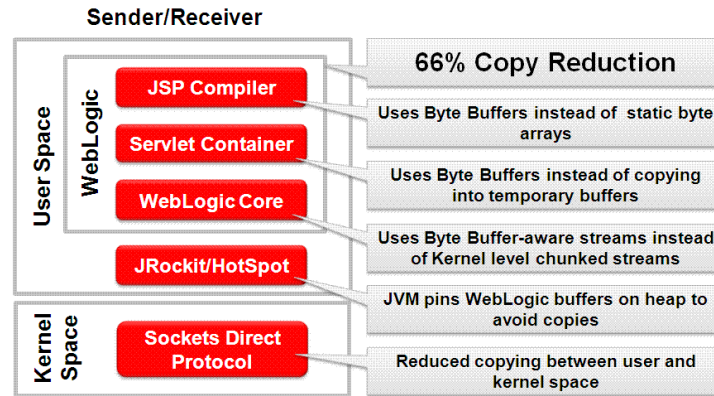


Figure 6. Reduced Buffer Copying

- Given InfiniBand enables inter-process communication at a network level using significantly larger packets than the traditional 1.5K MTU with Ethernet – another architectural change that was introduced into WebLogic Server was the use of vectored (or “scatter/gather”) I/O. Scatter/gather refers to the process of gathering data from, or scattering data into, a given set of buffers or messages in a single operation rather than with multiple operations. Since every read/write operation involves a (costly) call from user space to the kernel, this represents a significant saving in CPU cycles.

With scatter/gather I/O effectively what happens is that WebLogic Server is able to use a single procedure call to sequentially write data from multiple messages (e.g. 4K chunks within WebLogic Server) to a single data stream and correspondingly read data from a single data stream back into multiple messages. With this change there is dramatically less impedance between the internal 4K message size within WebLogic Server and the 64K MTU with IPoIB and particularly the arbitrarily large MTU with SDP. This approach complements the NIO buffer optimizations above and was critical to the construction and deconstruction of larger packets flowing over InfiniBand into, between, and from WebLogic Server instances. Figure 7 provides an overview of scatter/gather I/O with Exalogic.

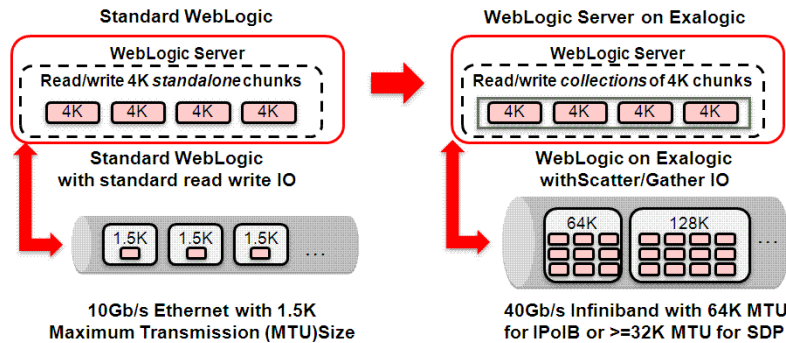


Figure 7: Scatter Gather I/O with WebLogic Server on Exalogic

Parallel Channels over Socket Direct Protocol (SDP)

The next set of changes made was centered on enabling the WebLogic Server network channel infrastructure to natively take advantage of InfiniBand. SDP itself was integrated into the multiple layers of the software stack, specifically into the Oracle JDBC driver and the JRockit and Hotspot, and then surfacing that via standard configurable properties within WebLogic Server. This enabled all network channels within WebLogic Server – HTTP(S), T3(S), IIOP (S) and others – to natively communicate over SDP.

Next, given that WebLogic Server was originally optimized for 10Gb/s Ethernet, it made sense to revisit standard Java EE workloads such as Internet applications, message based applications, and database applications to determine where they could be optimized for InfiniBand. A key area where optimizations became apparent was in the area of socket-based network I/O into, out of, and between WebLogic Server instances.

WebLogic Server uses a concept called network channels to handle such communication. To facilitate multiple protocols – HTTP(S), T3(S), IIOP(S) and others – on a single network channel, WebLogic Server uses a low-level internal construct called a muxer that acts as a multiplexer and demultiplexer for network traffic. In this area two unique changes have been engineered within WebLogic Server:

1. The default Maximum Transmission Unit (MTU) on 10Gb/s Ethernet is 1.5K, and this is the default configured in the standard distribution WebLogic Server. However, because InfiniBand enables a significantly larger MTU with IPoIB, (64K), and an arbitrarily large self-tuned MTU over SDP, WebLogic Server on Exalogic was first optimized via the scatter/gather I/O change to write out larger collections of chunks to the available MTU for the InfiniBand protocol stack – be the choice of protocol IPoIB or SDP.
2. When the larger MTU was combined with the higher InfiniBand bandwidth it was also necessary to rearchitect the WebLogic Server muxer infrastructure to more optimally dispatch the significantly higher rate of packets flowing into, between, and from WebLogic Server instances.

This was accomplished by enabling additional channel selectors executing in parallel, to handle the messages coming across the network channels into WebLogic Server. This newly added parallel muxer capability, tuned to match the networking speed of InfiniBand, effectively removed any lock contention that would have occurred due to the faster incoming flow of packets.

This combination of increasing the MTU and adding more parallelization to the internal dispatching infrastructure of WebLogic Server enabled a dramatic performance increase across all network protocols in WebLogic. Figure 5 illustrates this set of changes with a before and after changes comparison.

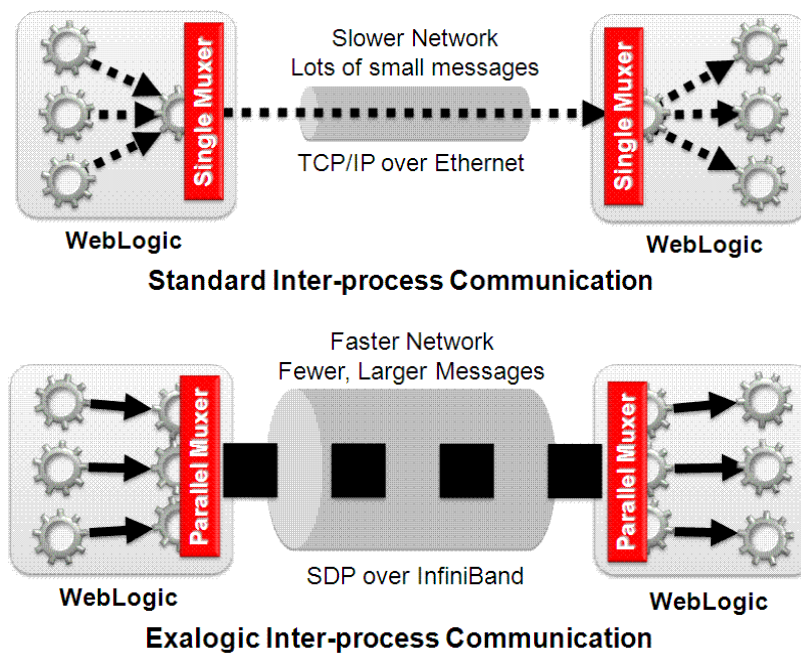


Figure 5. From Standard Inter-Process Communication to Exalogic Interprocess Communication

Optimized Thread Pool Sizing

The software optimizations based on the Exalogic hardware infrastructure were not only limited to the native adoption of InfiniBand for maximum network performance using SDP and IPoIB and a significant re-thinking of the Java I/O infrastructure within WebLogic Server; they also carried through to the self-tuning thread pool infrastructure.

To understand how the WebLogic Server self-tuning thread pool was optimized it is important to understand how it previously worked. Before Exalogic, the WebLogic Server self-tuning thread pool monitored the overall throughput every two seconds and used the collected data to determine when

the thread count needs to change. A combination of present thread count, measured throughput, and past history were taken into account by the algorithm to determine when the thread count needed to increase or decrease, and new threads were automatically added to the pool or removed, as needed.

To take advantage of the specific Exalogic hardware configuration of 30 compute nodes each with two processors and six cores per processor, the WebLogic Server self-tuning thread pool implementation was altered to be natively aware of the underlying hardware processing capacity and the thread allocation algorithm was changed to adjust based on this information.

Instead of incrementing the total thread count one at a time, WebLogic Server now takes advantage of knowing the higher core density and hyper-threading available on the hardware to increment the thread count by a significantly larger number - 24 at a time. This resulted in a re-architected self-tuning algorithm that allows optimal Java application performance on Exalogic, significantly faster than previous versions of WebLogic Server, particularly as those workloads vary over time.

Figure 8 provides a pictorial overview of the threading model optimizations.

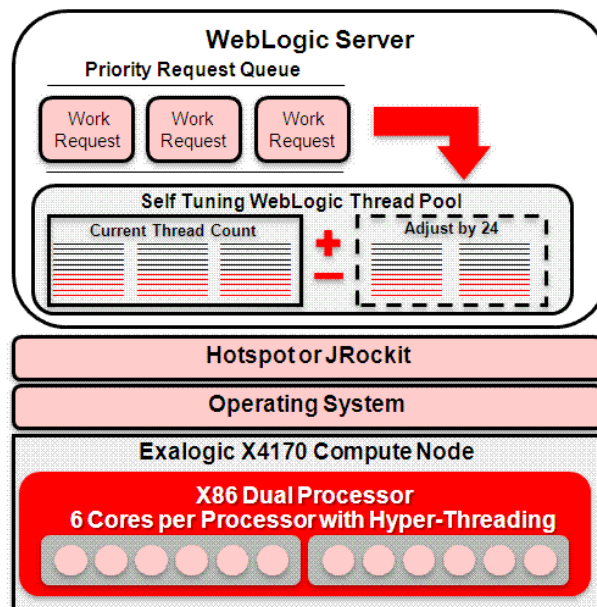


Figure 8. Self-Tuning Thread Pool Optimizations for Optimal Exalogic Performance

Exalogic Elastic Cloud Software Optimizations for Oracle Database

One of the key design goals of Exalogic was to optimize its infrastructure with Oracle RAC and Oracle Exadata. In addition to designing specific InfiniBand connectivity between Exalogic and Exadata racks into the hardware configuration which correspondingly carries into that fabric and the Exalogic Elastic

Cloud Software performance optimizations already covered, WebLogic Server was further optimized with a new feature called Active GridLink for RAC. This feature, part of the Exalogic Elastic Cloud Software, integrates WebLogic Server data sources with Oracle RAC, in a new architectural approach.

Specifically, RAC supports an event-based client notification mechanism to disseminate information about the state of a RAC cluster. The Active GridLink for RAC feature, manifested in WebLogic Server as a GridLink data source type, further optimizes the proven WebLogic Server data source infrastructure by leveraging the information within RAC notifications to provide better overall connectivity in the form of significantly faster response to RAC node failures and maintenance activities, better utilization of RAC cluster resources and improved runtime monitoring and management. Perhaps most importantly, the intelligent runtime processing of this RAC event information has resulted in dramatically improved performance results for WebLogic-based OLTP database workloads on Exalogic and Exadata as compared to non-ExaLogic, non-Exadata infrastructure.

Single Data Source

The Active GridLink for RAC feature also dramatically simplifies and consolidates WebLogic Server data source configuration and administration. It provides a single data source that is enhanced to support the capabilities of RAC. This is significantly different and easier to administer than the traditional Multi Data Source model (still supported), which has a combination of one data source per RAC node and then an additional logical data source as a proxy layer representing the resulting collection of RAC data sources. For example, in a 4-node Oracle cluster on Exadata, the WebLogic Server data source configuration would decline from five distinct data sources with Multi Data Sources (one for each RAC node and one proxy data source) to one with Active GridLink for RAC.

Fast Connection Failover

Oracle Notification Service (ONS) is used by the Oracle RAC database to broadcast events that describe changes in its state. With Active GridLink for RAC, WebLogic Server registers to receive notifications from ONS and quickly becomes aware of any state changes in a RAC database. This approach on Exalogic can be orders of magnitude faster as compared to the multi data source polling approach or in worst case TCP/IP timeouts. Using these state-change notification events, WebLogic Server can intelligently adapt its connection pools at runtime so that it provides continuous and reliable access to the RAC database as changes occur. Figure 9 shows WebLogic Server operating with no downtime due to this feature, despite a RAC node being taken offline.




Figure 9. WebLogic Server with FCF for RAC

Runtime Connection Load Balancing

To provide better throughput and more efficient use of resources, Oracle Database provides a complete runtime load balancing service to distribute connections across the Oracle RAC instance based on performance goals set by the database administrator. The load balancing advisory service issues Fast Application Notification events (“FAN” events) that advise clients on the current state of the cluster, including advice on where to direct connections.

WebLogic Server receives load balancing advisory events issued by the database and distributes connections to the Oracle RAC nodes accordingly. This has been made available as a runtime load balancing configuration option on the Active GridLink for RAC data source. Enabling it dictates that the pool dispatch connections based on the advisories it has received from the database, achieving an outcome where the connections are directed to a specific database node to satisfy a performance goal the database administrator has set. Figure 10 provides an overview of runtime load balancing in action with the WebLogic nodes choosing the least loaded node for routing traffic.




Figure 10. Runtime Load Balancing to Least Loaded RAC Node

XA Affinity

XA affinity is a performance feature that ensures that all database operations performed on an Oracle RAC cluster within the context of a global transaction are directed to the same Oracle RAC instance. Affinity will be established based on the global transaction id, instead of by individual data source, to ensure that connections obtained from different data sources that are configured for the same Oracle RAC cluster are all associated with the same RAC instance. Figure 11 shows a view of how XA affinity enables WebLogic Server to have transactional affinity with RAC nodes.




Figure 11. RAC Node Affinity from WebLogic

Database Graceful Shutdown

A graceful planned shutdown occurs when a database node or service is targeted for a shutdown operation, with the corresponding issuance of a database event indicating the shutdown has been requested. To support graceful planned shutdown of a database, Active GridLink for RAC does not immediately abort connections that are in use when it detects that the database shutdown target is no longer accepting new connections. Instead, the GridLink data source allows any in-progress transactions to complete before closing and recreating the physical connections, while cleaning up idle connections so that new requests for connections are not sent to the database target in active shutdown mode.

Benefits

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 12. 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 13. 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.

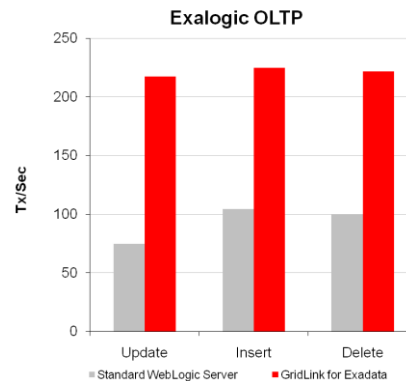


Figure 14. Exalogic accelerates database interactions by 2-3x.

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%.

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.



Figure 15. Exalogic is an ideal platform for consolidating numerous, disparate workloads.

Consistent management across the system with Enterprise Manager provides service-level monitoring, policy-based automation, and robust dynamic capacity adjustment in conjunction with WebLogic Server and Coherence clustering.

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 14x
- 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 Palmetier
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.