

An Oracle White Paper

January, 2015

Oracle Exalogic Elastic Cloud System Overview



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.

Table of Contents

Executive Summary	1
Hardware and Software Engineered to Work Together	4
Oracle Standard	4
Performance and Scalability	5
Reliability and Serviceability	5
Exalogic System Architecture	6
Exabus	7
Exabus InfiniBand Gateways and Switches	8
Exabus Software, Firmware and Drivers	9
Traffic Director	9
High Availability	10
High Performance	10
Flexible Routing and Load Control on Back-End Servers	11
Controlling the Request Load and Quality of Service	11
Security	11
Monitoring	11
Integration with Oracle WebLogic Server	12
Virtual Firewall (InfiniBand Partitions)	12
Exalogic Storage	12

Exalogic Compute Nodes.....	14
Exalogic Management Switch	14
Exalogic Operating Systems	15
Oracle Linux Unbreakable Enterprise Kernel (UEK)	15
Oracle Solaris	15
Oracle VM	16
Exalogic Control	17
Exalogic's Infrastructure-as-a-Service User Interface	17
Exalogic System Utilities	19
Exalogic Integrated Lights Out Manager (ILOM)	19
Deploying Oracle Exalogic with Oracle Exadata	20
Exalogic Maintenance, Expansion, Upgrade	21
Datacenter Integration.....	22
Conclusion	23

Executive Summary

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 meet the highest standards of reliability, availability, scalability and performance for a wide variety of time-sensitive, mission-critical workloads.

Exalogic dramatically improves performance of standard Linux, Solaris and Java applications¹ without requiring code changes and reduces costs across the application lifecycle, from initial set-up to on-going maintenance, as compared to conventional enterprise application platforms and private clouds assembled from disparate components sourced from multiple vendors.

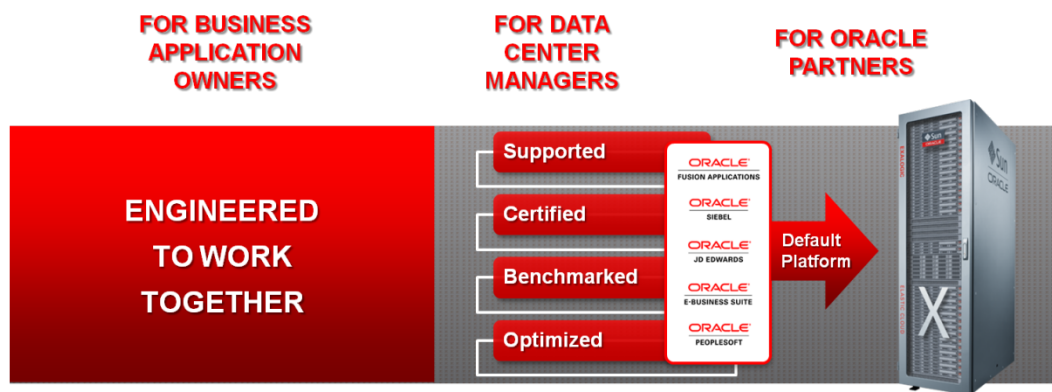


Figure 1: Business Value of Oracle Exalogic Elastic Cloud

Exalogic is an open system, assembled by Oracle from its portfolio of standards-based, best-of-breed component products and technologies. The Exalogic system reflects best-practices acquired and refined from thousands of customer deployments and extensive R&D effort.

¹ Any application that supports Oracle Linux 5 (Unbreakable Enterprise Kernel 2, 64-bit), Oracle Linux 6 (Unbreakable Enterprise Kernel 2, 64-bit), Oracle Solaris 11 (or Solaris 10 Zone, x86)

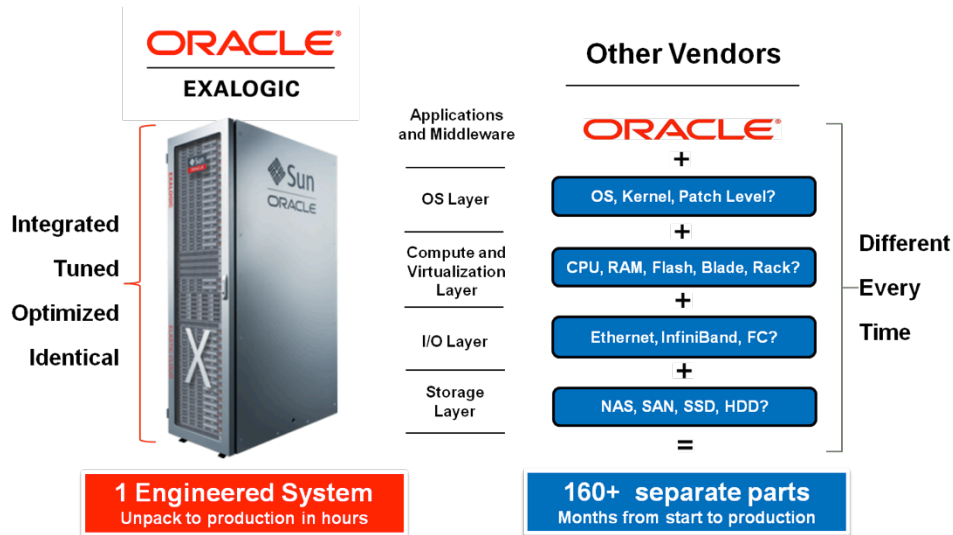


Figure 2: Engineered System Paradigm vs. Conventional "Piece-meal" Approach

With the appropriate investment of resources and time it is possible for customers to aggregate their application requirements, research the available products and technologies, evaluate multiple vendor proposals, select component products, order, receive and install the components, apply firmware and software patches, obtain and install device drivers and finally, test, tune and document the resulting system.

However, such a system, assembled or integrated from disparate components still does not match the capabilities of Exalogic. Oracle has made unique optimizations and enhancements to Exalogic components as well as Oracle's middleware and Oracle's applications that cannot be made by customers or by any 3rd party. These include on-chip network virtualization, high-performance Remote Direct Memory Access (RDMA) at operating system and Java Virtual Machine (JVM) layers and Exalogic-aware workload management in Oracle Weblogic server (Oracle's Java EE application server).

To accelerate production deployments and reduce on-going maintainance and administration costs, the Exalogic system includes a suite of exclusive system-level diagnostic and configuration tools. These tools ensure that Exalogic consistently delivers maximum performance and reliability while requiring less time to install, administer and maintain as compared to competing platforms.

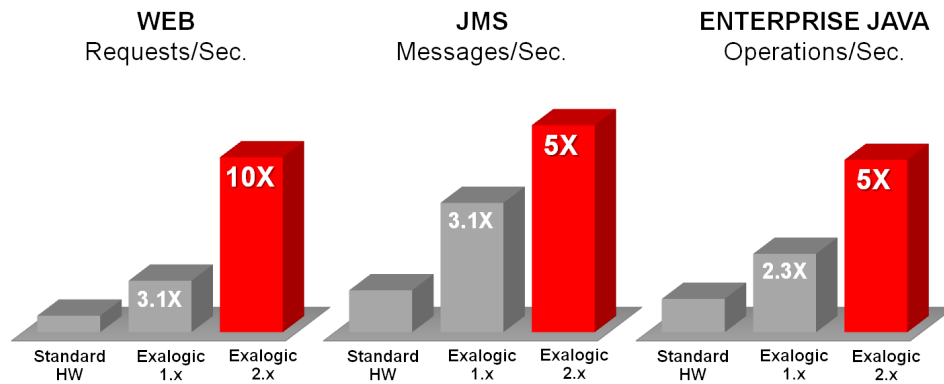


Figure 3: Exalogic Elastic Cloud Software (v 2.X) Performance Benchmark

In real-world settings, Exalogic outperforms other platforms by wide margins, often delivering two to ten times, at times even more, application performance. Not only do applications become more responsive and deliver a greatly improved user experience, they also become more resource efficient. With Exalogic, enterprises are able to support application workloads with less hardware, less power, less heat, less data center space and less software. Being fully pre-integrated by Oracle, Exalogic is also easier to provision, manage and maintain, which further accelerates time to value and lowers TCO for new projects. Exalogic is designed for high availability and zero-down-time maintenance and can be scaled linearly from a single, Eighth-Rack configuration to a large system of multiple Full-Racks without any service disruption or additional external hardware requirements.

Exalogic is at the fore-front of an industry-wide shift from costly and outdated piece-meal application platforms toward *Engineered Systems*. For enterprises looking to get more from their IT investments than merely keeping the “lights on”, Exalogic is the only choice!

Hardware and Software Engineered to Work Together

The Exalogic system consists of two major elements:

- **Exalogic Elastic Cloud Machine:** a high performance hardware system, assembled by Oracle, that integrates storage and compute resources using a high-performance I/O subsystem called Exabus, which is built on Oracle's Quad Data Rate (QDR) InfiniBand
- **Exalogic Elastic Cloud Software:** an essential package of Exalogic-specific software, device drivers and firmware that is pre-integrated with Oracle Linux and Solaris, enabling Exalogic's advanced performance and Infrastructure-as-a-Service (IaaS) capability, server and network virtualization, storage and cloud management capabilities

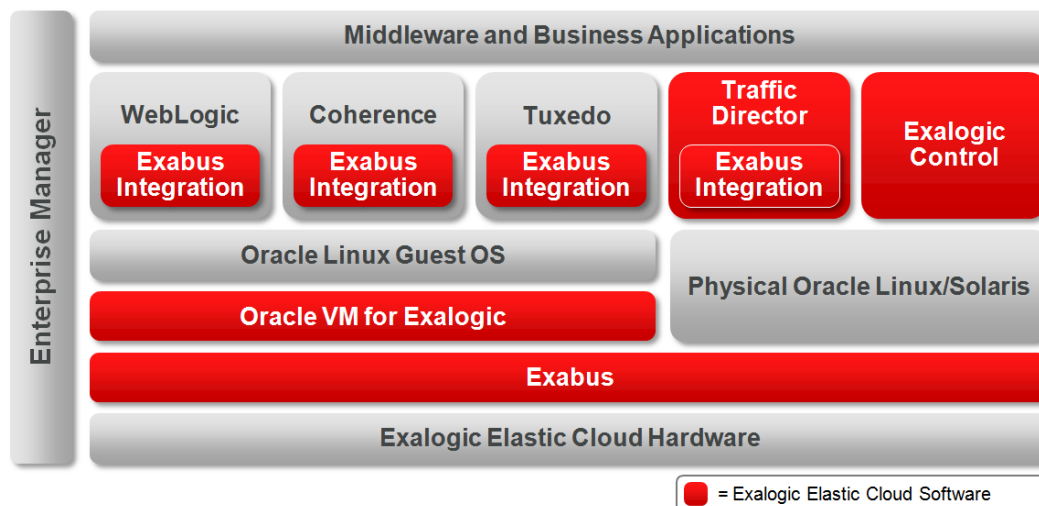


Figure 4: Key Components of Exalogic Elastic Cloud Software

Oracle Standard

It is simply not possible for any vendor to achieve the level of usability, convenience, reliability, and performance that comes from engineering hardware and software to work together if that vendor does not control the development of both the hardware and software. When it comes to providing a platform for Oracle's business applications and middleware, no other vendor is better able to fulfill the promise of an end-to-end engineered system, and no other cloud infrastructure can compete with Oracle Exalogic. Not only is Exalogic designed for Oracle's middleware and business applications, those same middleware products and business applications are being developed and tested on Exalogic.

Performance and Scalability

Modern applications, especially those designed in the last few years, have embraced *scale out*² architecture as a solution not only to the challenges of cost, scale and application reliability and serviceability. This approach offers many advantages over legacy approaches which are typically dependent on using increasingly large and costly high-end servers, although it generally suffers from one persistent and challenging limitation: I/O. The performance and efficiency of modern highly distributed systems is primarily constrained by the communication channel that connects all of the system components. The central technological advance in Exalogic, therefore, is the *elimination of I/O bottlenecks at every level* through a networking hardware and software sub-system called *Exabus*. Exabus not only makes applications *run faster*, it also makes them *more efficient*, and it does this consistently and predictably even in extremely large scale deployments with thousands of processor cores and terabytes of memory, for business applications.

Reliability and Serviceability

Exalogic is designed for mission critical applications that must be highly available. Achieving so called *five-nines* (99.999%) availability requires a system that is both fault tolerant and accommodates zero-down-time maintenance and administration. Exalogic is designed for hardware redundancy and automated failover for every major component including power, I/O and cooling (fans). When running scale-out applications, it is possible for every component in the system to be taken out of service in turn, repaired or maintained, and returned to service with no disruption to applications or users. Exalogic features built-in support for block-level storage replication, backup-to-disk and Automated Service Requests³ (ASR). ASR allows Oracle to pro-actively monitor Exalogic systems for actual or impending component failure and pro-actively dispatch replacement parts and service personnel, thereby minimizing or completely avoiding service disruptions.

² “Scale out” refers to so-called horizontal scaling, in which applications scale through the addition of new instances of the application business logic to the system and a combination of traffic shaping, state replication and database sharing.

³ ASR is provided as a part of Oracle’s optional Premier Support for Systems support offering.

Exalogic System Architecture

Exalogic machines are available in four configurations: Eighth Rack, Quarter Rack, Half Rack and Full Rack. Upgrade kits are available which allow for Eighth Rack configurations to be upgraded to Quarter Rack configurations, Quarter Rack configurations to be upgraded to Half Rack configurations, and from Half Rack configurations to Full Rack configurations. It is possible to connect up to eight (8) configurations of any type together using the provided cabling to form a single Exalogic system (sometimes called an *Exalogic fabric*). It is possible to scale an Exalogic system even further, if desired, through the use of separately available datacenter spine switches.




Figure 5: Standard Exalogic Elastic Cloud X5-2 Hardware Configurations

An Exalogic configuration consists of the following major components:

- Exabus
- x86 Compute Nodes
- ZFS Storage Appliance
- Management Switch
- Linux and Solaris
- Oracle VM (hypervisor)
- Exalogic Control
- Exalogic System Utilities

Exabus

The defining architectural feature of Exalogic is the Exabus communication (I/O) fabric that ties all of the system components together and provides the basis for Exalogic's reliability, availability, scalability and performance. Exabus is based on Quad Data Rate (QDR) InfiniBand, and consists of hardware, software and firmware distributed throughout the system and involving every major system component.

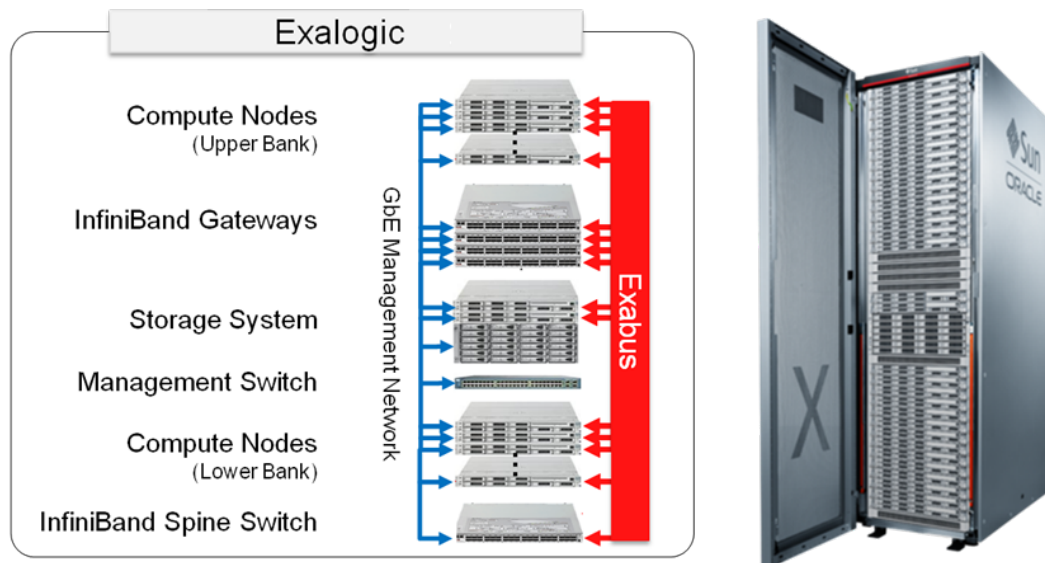


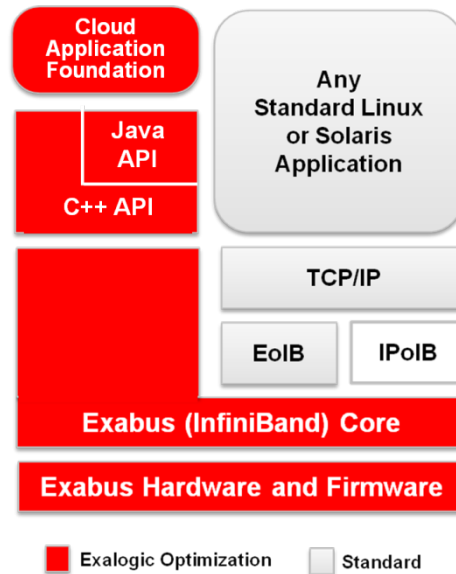
Figure 6: Key Components of Exalogic Elastic Cloud Hardware

QDR InfiniBand was selected as the foundation technology for Exabus for several reasons:

- Oracle's InfiniBand products provide the greatest available bandwidth per physical port (40Gb/s) and lowest latency ($\sim 1.07\mu\text{sec}$) of any standard interconnect technology available today⁴, allowing applications to reclaim compute capacity otherwise wasted waiting on slow communication links

⁴ 3.7x the throughput and 1/5 the latency of 10Gigabit Ethernet, the next best option, according to http://www.hpcadvisorycouncil.com/pdf/IB_and_10GigE_in_HPC.pdf

- InfiniBand provides reliable delivery, security and quality of service at the physical layer in the networking stack and natively supports kernel bypass operations, eliminating much of the inefficiency of using system CPU and main memory
- Oracle's InfiniBand products support upper-stack protocols like IP (IPoIB) and Ethernet (EoIB), making it possible for existing applications to run without modification and still benefit from enhanced performance



Exabus InfiniBand Gateways and Switches

One of the most important and differentiating aspects of Exalogic is the ability of the system to present itself to the datacenter network as a single large computer with a large number of 10 Gigabit Ethernet (10GbE) network interfaces. At the physical level, this is achieved through the use of special gateways that connect directly to external 10GbE ports exposed by standard datacenter switches. Exalogic's Ethernet-to-InfiniBand gateways use a unique InfiniBand-to-Ethernet bridge design that allows each Exalogic compute node, through simple software configuration, to have as many as *1000 virtual 10GbE network interface cards*, each of which appears to the datacenter network as if it is a standard Ethernet device.

Figure 7: Exalogic Elastic Cloud Networking Stack

This approach greatly reduces the number of physical network cards, cables, and switches that must be installed, powered, tested, tuned, managed, and maintained while offering improved application performance, security and reliability. This approach eliminates interoperability issues with existing network management tools and guarantees complete interoperability with existing Ethernet-based switches, firewalls and load balancing appliances, since *Exalogic's use of InfiniBand is completely invisible to the rest of the data center*.

Exabus supports the creation of Virtual Local Area Networks (VLANs) within Exalogic as a means of providing application isolation, and transparently isolate inter-cluster communication to non-Ethernet subnets, thereby providing both enhanced security and improved cluster performance.

Exabus Software, Firmware and Drivers

At the software layer, Exabus extends and enhances the Open Fabrics Enterprise Distribution (OFED™). OFED is an industry standard open-source software toolkit for Remote Direct Memory Access (RDMA) and kernel bypass applications. OFED is widely used in high performance InfiniBand-based computing systems that require maximum throughput, minimal latency and a unified infrastructure for storage access, network virtualization and cluster Inter-Process Communication (IPC).

Exabus includes kernel-level drivers, channel-oriented RDMA and send/receive operations, kernel by-passes of the operating system, both kernel and user-level application programming interface (API) and services for parallel message passing, sockets data exchange, storage and file/database system access.

Exabus incorporates a number of reliability, management, and performance features that are not available from any other similar system:

- Unique support for Exalogic's Ethernet-over-InfiniBand (EoIB) gateways
- Simplified management and monitoring with full ILOM and Oracle Enterprise Manager OpsCenter integration
- Quality of Service (QoS) and Partitioning configuration support
- High Availability Sockets Direct Protocol (SDP)
- IPv6 support for SDP, EoIB and IP-over-InfiniBand (IPoIB)
- Automatic disabling of degraded physical links (autonomous port-level failover)
- 100s of separate design fixes and enhancements in the Host Stack and Management Stack, improving compliance, stability, efficiency and performance

Traffic Director

Traffic Director is a fast, reliable, and scalable Application Delivery Controller that may be deployed as the entry point for all HTTP and HTTPS traffic to application servers and Web servers in an Exalogic deployment.

- Traffic Director is fully integrated with the Exabus I/O subsystem and can support

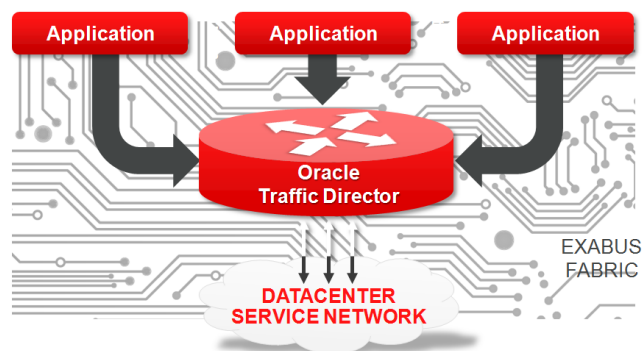


Figure 8: Application Delivery Controller Topology

both extremely high throughput and low latency application traffic workloads.

- Traffic Director supports enterprise-grade high availability via, active-passive or active-active failover.
- Traffic Director may be easily and dynamically scaled in lock-step with varying volume of application traffic.
- Traffic Director may be easily configured to apply multiple, declarative rules when distributing requests to the back-end servers and when forwarding responses to clients.
- Traffic Director is easy to install, configure, and use. It includes a simple, wizard driven graphical interface as well as a robust command-line interface.

High Availability

Health checks for the back-end: If a server in the back-end is no longer available or is fully loaded, Traffic Director detects this automatically through periodic health checks and stops sending client requests to that server. When the failed server becomes available again, Traffic Director detects this automatically and resumes sending requests to the server.

Backup servers in the back-end: When setting up server pools for a Traffic Director instance, you can designate a few servers in the back-end as backup servers. Traffic Director sends requests to the backup servers only when none of the primary servers is available, ensuring continued availability even when some servers in the back-end fail.

Failover for load balancing: To ensure high availability of the load balanced services, you can deploy Traffic Director instances in an active-passive or active-active failover configuration.

Dynamic reconfiguration: Most configuration changes to Traffic Director instances can be deployed dynamically, without restarting the instances.

High Performance

SSL/TLS offloading: Traffic Director can be configured as the SSL/TLS termination point for HTTPS requests, reducing the processing overhead on the servers in the back-end.

Content caching: Traffic Director can be configured to cache (in its process memory) content that it receives from origin servers. By caching content, Traffic Director helps reduce the load on servers in the back-end and helps improve performance for clients.

HTTP compression: You can configure Traffic Director instances to compress data received from servers in the back-end and forward the compressed content to the requesting clients. This feature improves the response time for clients connected on slow connections.

Flexible Routing and Load Control on Back-End Servers

Request-based routing: Traffic Director can be configured to route HTTP/S requests to specific servers in the back-end based on information in the request URI: pattern, query string, domain, source and destination IP addresses, and so on.

Request rate acceleration: You can configure the rate at which Traffic Director ramps up the load on specific back-end servers, so that servers that have just been added to the pool or restarted, to perform startup tasks such as loading data and allocating system resources.

Connection limiting: Traffic Director can be configured to limit the number of concurrent connections to a server in the back-end. When the configured connection limit for a server is reached, further requests that require new connections are not sent to that server.

Controlling the Request Load and Quality of Service

Request rate limiting: Traffic Director can be set up to limit the rate of requests from specific clients and request types, enabling optimal utilization of available bandwidth, guaranteeing a certain level of quality of service, and preventing denial-of-service attacks.

Quality of service tuning: To ensure equitable utilization of the available network resources, you can configure Traffic Director virtual servers to limit the number of concurrent connections to clients and the maximum speed at which data can be transferred to clients.

Security

Reverse proxy: By serving as an intermediary between clients outside the network and servers in the back-end, Traffic Director masks the names of servers in the back-end and provides a single point for tracking client access to multiple servers in the back-end.

Support for SSL 3.0 and TLS 1.0: You can configure SSL/TLS-enabled HTTP listeners for Traffic Director instances, using either certificates issued by commercial CAs such as VeriSign or RSA- and ECC-type self-signed certificates with key sizes of up to 4096 bits.

Monitoring

Traffic Director records statistics about server activity at different levels—instances, virtual servers, listeners, connections, and origin servers. For example, for each server instance, Traffic Director collects statistics about the duration for which the instance has been running, number of requests processed, average load, and so on. You can monitor statistics pertaining to the performance of Traffic Director instances through several methods: the administration console, the command-line interface, and a report in XML format.

Integration with Oracle WebLogic Server

Traffic Director is designed to recognize and handle headers that are part of requests to, and responses from, Oracle WebLogic Server instances in the back-end.

Traffic Director can dynamically discover changes in the Oracle WebLogic Server cluster—such as the removal or addition of managed servers, and consider such changes while routing requests.

Virtual Firewall (InfiniBand Partitions)

Exalogic supports a powerful physical I/O traffic separation capability called “InfiniBand partitions”. In an Exalogic system, the QDR InfiniBand fabric is used as the physical foundation on which all other communication networks are virtualized. In Exalogic, the switches in the Exabus I/O backplane of the system enforce the end-point security rules that determine which applications or software components on any given compute node are allowed to send or receive messages (or even see) to any other.




Figure 9: InfiniBand Partitions

This enforcement happens at the lowest possible layer in the system and is highly secure – even securing root level access to compute nodes does not compromise the security of the system as a whole and all applications running on other compute nodes on the system are completely unaffected. In combination with Traffic Director and software firewall technologies like Linux *iptables*, it is possible to implement so-called De-Militarized Zones (DMZ) on the Exalogic system which are as secure as those implemented using traditional hardware firewall appliances.

Exalogic Storage

The Exalogic storage subsystem stores all application binaries, log files and content necessary for the applications to execute. Compute nodes mount configured storage system partitions over the InfiniBand network. This eliminates the need to back up individual compute nodes and in the event of failure of a compute node, the partition can simply be mounted from another




Figure 10: Exalogic Elastic Cloud Storage

compute node while the failed compute node is serviced or replaced.

The Exalogic storage subsystem consists of two physically separate storage heads in an active/standby configuration and large shared disk array. Each of the storage heads is directly attached to the I/O fabric with redundant QRD InfiniBand. The storage subsystem is accelerated with two types of solid state memory that are used as read and write caches, respectively, in order to increase system performance. The storage heads transparently integrate the many Serial Attached SCSI disks in the disk array into a single ZFS cluster which is then made available to Exalogic compute nodes via standard network file systems supported by the compute node's operating system.

The Exalogic storage subsystem provides its own dedicated management interface and offers a number of user-selectable options for security, reliability and quota management. It also offers built-in support for storage replication, allowing each Exalogic configuration to be paired with another geographically remote system as part of a larger disaster recovery strategy.

It is also possible to attach selected Oracle storage appliances directly to the Exalogic I/O backplane using InfiniBand in order to expand the system's storage capacity or implement a high-performance backup solution

Integrated Exalogic Storage Software Features	
File system	Oracle Solaris ZFS (128bit addressability)
File level protocol	NFS v3/v4, FTP/SFTP
Block level protocol	IP over InfiniBand
Monitoring	DTrace Analytics (for system tuning and debugging), dashboard monitoring for key system performance metrics
Automated serviceability	"Phone Home" capability with automatic case creation, configurable alerts
RAID	Striping, mirroring
Snapshots	Read only, Restore
Directory service	NIS, AD, LDA
Data security	Checksum data and metadata, antivirus quarantine
Network services	NTP, DHCP, SMTP
Backup	NDMP v3/v4, ZFS NDMP
Local Replication	Replication within same ZFS Storage appliance (single or cluster)
Clones	Writable snapshots
Remote Replication	Exalogic replication to second remote system. 1:N, N:1, manual, scheduled, continuous

Figure 11: Exalogic Elastic Cloud Software Features

Exalogic Compute Nodes

Each Exalogic compute node (sometimes called a *server*) contains two Intel Xeon processors, each of which has multiple cores. Each compute node also has a large amount of high-speed error correcting RAM, enterprise grade solid state disk drives in a RAID configuration, redundant fans, redundant power supplies and a dual-port InfiniBand Host Channel Adapter that connects it to the system's I/O fabric.

Each Exalogic compute node is a physically isolated application environment. While each compute node has built in redundancy, if a compute node does fail or requires servicing, all of the application instances running on that node will fail or must be stopped. For HA, it is therefore recommended that application clusters span at least two compute nodes.

Exalogic compute nodes are balanced for maximum single thread performance: the type of memory used and the number and position of memory modules installed are carefully matched to the PCIe system bus and processor frequency. For latency sensitive or memory intensive applications, such as those built using the Java language, an Exalogic compute node's computational performance is up to 40% greater than similar systems running the same Intel processors. Compute nodes use their internal Solid State Disks only for the operating system bootable images and can be quickly and easily re-imaged using special tools provided with the system.

Exalogic Management Switch

Each Exalogic configuration includes a management switch. Every major component of the Exalogic system is connected directly to this management switch using the standard Gigabit Ethernet and built-in management ports. This management switch is the physical integration point with the data center's secure management network, and allows for complete separation of management traffic from service traffic. In addition to carrying management traffic, the management network is also used to facilitate background replication of data in the Exalogic storage system to paired Exalogic systems at remote sites as part of a disaster recovery solution.

Exalogic Operating Systems

Exalogic provides users with a choice of Oracle Linux or Oracle Solaris operating systems and guarantees 100% compatibility with standard Linux or Solaris applications without requiring any additional certification for Exalogic. All Oracle applications that are certified for the appropriate releases of Oracle Linux and Oracle Solaris are supported on Exalogic.

Each Exalogic configuration ships with both Oracle Linux and Oracle Solaris, Exabus software, drivers and firmware and management tools and utilities already installed and ready for final configuration on the customer's premises. While the versions of Linux and Solaris used with Exalogic are not exclusive to Exalogic, the specific bootable disk images Oracle provides are. These disk images are specifically created for use with Exalogic. No other operating system versions are supported on an Exalogic system, and without the unique software, firmware and device drivers incorporated into these Exalogic 'base images' the Exalogic hardware is unusable.

Oracle Linux Unbreakable Enterprise Kernel (UEK)

The Oracle Linux Unbreakable Enterprise Kernel delivers the best overall Linux performance available today and provides numerous features in the areas of hardware fault management, data integrity and diagnostics, including detecting and logging hardware errors before any affect to OS or application and automatic isolation of defective CPUs and memory.

Taking advantage of these enhancements, require no changes to existing Linux applications. The optimizations provide up to 60% higher workload capacity, 50% lower latency, and 50% greater IPoIB performance⁵.

Oracle Solaris

Solaris 11 is also an option for Exalogic. In addition to many security, network virtualization, management and fault tolerance features, Oracle Solaris supports a kernel-level server virtualization technology known as *Zones*.

Oracle Solaris Zones is one of the most highly adopted, highly used, mature virtualization technologies on the market today and has been a core feature of Solaris since its introduction in

⁵ Based on Oracle internal product testing and benchmarking results

the Solaris 10 release in 2005. Oracle Solaris Zones⁶ includes support for Oracle Solaris 10 Zones, which allows customers to deploy Solaris 10 (x86) applications on Exalogic.

Oracle VM

Exalogic includes support for a highly optimized version of the Oracle VM hypervisor, which can be used to subdivide a physical compute node into multiple virtual servers (vServers), each of which may run a separate Oracle Linux operating system instance and applications.

Oracle VM has been engineered for tight integration with Exalogic's Exabus I/O backplane using a technique called Single Root I/O Virtualization (SR-IOV).

This approach requires all operating systems running in these vServers to have special device drivers for communicating with the outside world via Exabus.

The benefit of this approach is unmatched application performance: in an Exalogic configuration, the impact of virtualization on application throughput and latency is negligible. Applications running in Exalogic vServers perform on par with deployments on bare metal, but retain all of the manageability and efficiency benefits that come with server virtualization.

On Exalogic, Oracle VM significantly outperforms comparable hypervisors from other leading vendors. The substantial performance advantage that Exalogic enjoys results in better application usability, improved hardware utilization and lower software licensing costs. For enterprises with investments in Oracle middleware and business applications, Exalogic is an ideal choice to realize benefits of a high-performance, fully virtualized applications platform and applications-to-disk management.

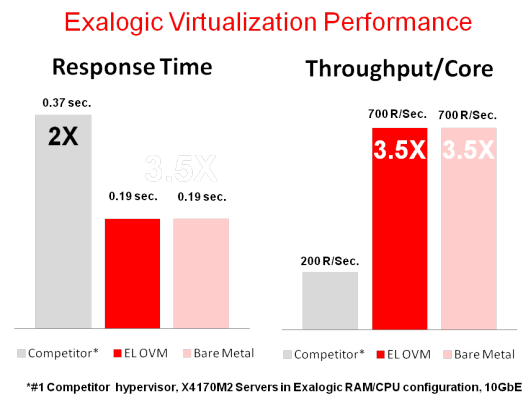


Figure 12: Oracle VM vs. non-Oracle Hypervisor Performance Benchmarks

⁶ See Oracle Solaris Virtualization:

<http://www.oracle.com/technetwork/articles/servers-storage-admin/sol11evirt-186209.pdf>

Exalogic Control

Exalogic is a full-featured, enterprise class private cloud. With Exalogic Control, Exalogic may be deployed as part of an Infrastructure-as-a-Service offering, in which a single IT organization owns and operates the physical system and allows clients (internal or external to the enterprise) to remotely access the Exalogic system to deploy and manage their own applications.

Exalogic control is a collection of software components that are deployed directly on the Exalogic system at the time it is installed. These components integrate with the management interfaces of all of the components in the Exalogic system, including the Exabus I/O backplane, the storage system, the power distribution system, the management switch, the physical compute nodes and the Oracle VM hypervisors running on each compute node.

Exalogic control provides users of the system a choice of a Web-based GUI, a command line interface and a Java API, through which they may access Exalogic Control and manage the Exalogic system. These interfaces allow users of the system to access compute resources, network resources and storage resources without having to be aware of the details of the underlying physical system. Exalogic resources are fully virtualized and made available to users in a secure, metered and extremely flexible way.

Exalogic's Infrastructure-as-a-Service User Interface

Virtualized Data Center (vDC)

A vDC is a collection of physical Exalogic compute nodes, storage and switches (usually one or more connected Exalogic configurations). These physical resources are organized into a pool that can then be accessed by self-service users.

Exalogic Systems Administrator

The Systems Administrator is the "super user" with all administrative privileges required for management and administration of all

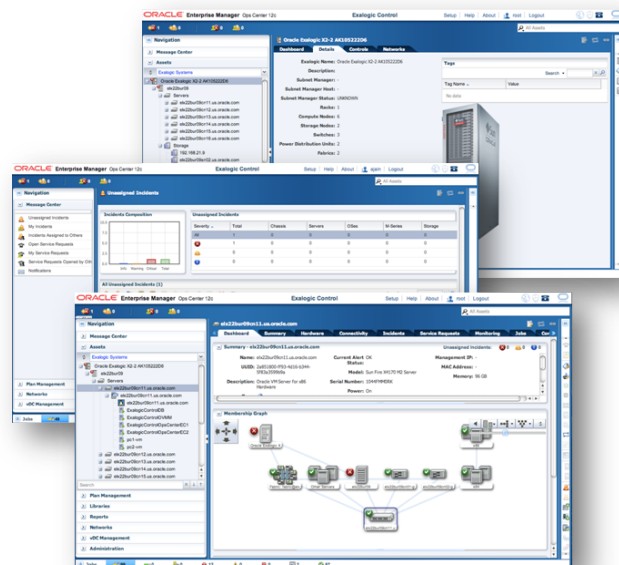


Figure 13: Oracle Exalogic Elastic Cloud Management Console

components and sub-components of the Exalogic system. However, the Exalogic Systems Administrator does not have permission to access the applications that may be running in individual vServers or encrypted application data stored by applications.

Account

Accounts are created by Cloud Administrators. An Account includes quotas for the amount of CPU, memory, storage and networking resources that may be consumed within the context of that Account (within the scope of a single vDC). Usage of resources allocated to accounts may be metered and charged.

Cloud Administrator

A Cloud Administrator is a user who configures the cloud. Cloud Administrators create Accounts and set quotas.

Cloud User

A Cloud User consumes resources in the cloud by creating vServers, creating vNets and Volumes that allow those vServers to communicate with the outside world and each other and store information. Cloud Users are typically responsible for deploying enterprise applications, and are often not aware of the physical location or state of the Exalogic system. A given Cloud User may have access privileges to multiple Accounts within any given vDC in Exalogic.

Virtual Network (vNet)

A virtual network is a logical construct that dictates which servers may communicate with which other servers. Private (within a vDC) vNets are created by Cloud Users. Public vNets (which have connections to the systems outside Exalogic) are created by Cloud Administrators only. The number of vNets that may be created is part of the quota allocated to an Account.

Virtual Server (vServer)

A vServer is a virtual machine (a virtual Exalogic compute node) which is capable of running its own dedicated instance of a guest operating system and all the applications deployed on it. Many individual vServers can co-exist on a single Exalogic compute node at one time, and share all of the physical CPU and memory on that compute node, within the parameters dictated by the account that each vServer belongs to. A vServer can be a member of one or more vNets (public or private) and may access one or more Volumes.

Volume

A Volume is a piece of storage that may be used by applications running in a vServer. Volumes are held within the context of an Account and consume storage quota.

Exalogic System Utilities

Exalogic is manageable using a variety of tools, including the majority of management tools and systems employed in customer data centers today. Irrespective of the specific tools used, managing Exalogic is substantially less complex and error prone than managing traditional systems built from individually sourced components because Exalogic is explicitly designed to be administered and maintained as a single, integrated system.

Exalogic includes a number of very specialized tools that ensure the correct installation and configuration of the Exalogic system and can quickly and easily diagnose critical system-level issues.

- **Exalogic Configuration Utility:** allows configuration of the Exalogic system management and data center service network interfaces and internal subnets from a desktop
- **Exalogic Distributed Command Line Interface:** allows commands to be executed on some or all of the Exalogic nodes simultaneously, at the discretion of the operator.
- **Exalogic Topology Verifier:** verifies the InfiniBand topology of the Exalogic system, ensuring that the correct topology is applied for each given system configuration: Eighth Rack, Quarter Rack, Half Rack or Full Rack.
- **Exalogic InfiniCheck:** verifies the correct operation of every InfiniBand device and port on the fabric, ensuring that all ports and connectors are functioning correctly.
- **Exalogic Hardware & Firmware Profiler:** verifies that all of the hardware devices and firmware versions connected to the Exalogic system fabric are verified and supported, with the correct and compatible device firmware versions.
- **Exalogic Software Profiler:** verifies that all of the Linux or Solaris software packages installed on any of the system's compute nodes are of the correct version and do not jeopardize the Exalogic system's performance, security or stability.
- **Exalogic Boot Manager:** allows system operators to easily re-image individual Exalogic compute nodes, via external PXE servers or network-mounted disk images.

Exalogic Integrated Lights Out Manager (ILOM)

Exalogic's Integrated Lights Out Manager (ILOM) provides an agent-less means of managing Exalogic's compute nodes, InfiniBand gateways and switches, storage subsystem heads, the storage array and any installed power distribution units. ILOM is implemented in the firmware of each component in the Exalogic system, independent from the Operating System, and is accessible only through the physical management network for added security. ILOM supports

both command line and Web-based direct access to the core management capabilities of each component.

Deploying Oracle Exalogic with Oracle Exadata

Oracle Exalogic and Oracle Exadata are, in many respects, one platform. While Oracle Exalogic is designed to provide a platform for application business logic and compute-intensive workloads, Oracle Exadata is designed for application databases and storage-intensive workloads. Exalogic and Exadata share many common technologies and although they can be, and are, deployed independently from each other there are unique technical and operational benefits to deploying them together.

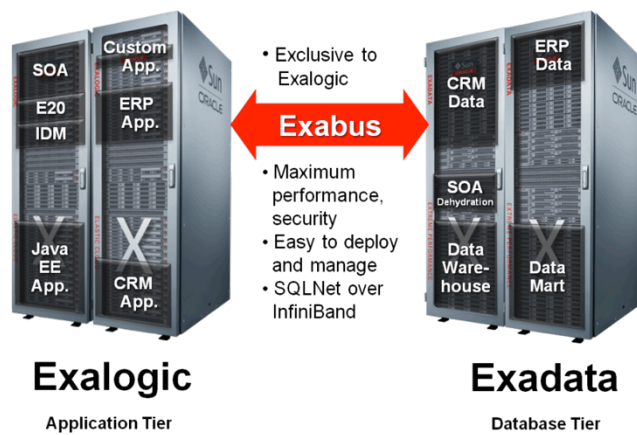


Figure 14: Datacenter Transformation with Oracle Exalogic Elastic Cloud and Oracle Exadata

Exalogic and Exadata may be connected directly via Exabus (QDR InfiniBand), providing a level of overall performance and security that is unsurpassed and exclusive to Oracle. Middle-tier application components deployed to an Exalogic Full Rack configuration, for example, are able to communicate with Oracle Real Application Cluster database nodes in an Exadata Full Rack system with a total of 960 Gigabits per second of total available throughput and microsecond

latency⁷, making the combination of Exalogic and Exadata unbeatable for On Line Transaction Processing (OLTP) and other performance sensitive applications.

Beyond the obvious benefits of such high-performance integration at the physical layer, Exadata and Exalogic:

- are uniquely able to benefit from optimizations in Oracle's JDBC implementation that further improve application performance by taking advantage of a native InfiniBand protocol called Sockets Direct Protocol (SDP)
- are Oracle's standard platforms for the database and middle tier, respectively, for Oracle's Fusion Middleware and business applications products
- share a number of key technologies and components, including compute node designs, InfiniBand switches and host channel adapters, device drivers and operating system enhancements
- can be managed using the same management infrastructure, including Enterprise Manager
- are supported by Oracle using a unique support infrastructure and pool of experts who are expertly trained to diagnose and resolve issues with application deployments that combine both products

Exalogic Maintenance, Expansion, Upgrade

Like all of Oracle's engineered systems, Exalogic is designed to be easy to maintain and upgrade.

Approximately four times per year, Oracle releases a set of device firmware and software patches that have been tested together and are approved as a "known good configuration" for the current and supported Exalogic hardware. Security and stability patches are made available for Exalogic as required.

Exalogic configurations may be upgraded to larger configurations through upgrade kits that are separately available. It is possible to purchase kits that will upgrade Exalogic Eighth Rack

⁷ All 24 available 40 Gb/s QDR InfiniBand ports may be used to connect both systems

configurations to Quarter Rack configurations, Exalogic Quarter Rack configurations to Half Rack configurations, and kits that will upgrade Half Rack configurations to Full Rack configurations. These upgrade kits are installed by Oracle Advanced Customer Services personnel on-premise and may be installed without taking the Exalogic system out of service.

Oracle Advanced Customer Service also offers multi-rack cabling services that allow Exalogic configurations to be integrated as part of an expansion of an existing Exalogic installation or an initial Exalogic deployment. Multi-rack cabling services are also available for customers that wish to connect Exalogic and Exadata systems. It is possible to connect up to eight (8) Exalogic and/or Exadata Full Rack configurations with the built-in expansion ports and provided cabling and external InfiniBand switches are not required.

Datacenter Integration

Integration of Exalogic with external systems, such as hardware firewalls, load balancers, databases, storage systems or other application environments is possible using the 10GbE connectivity provided by the Exalogic gateways.

InfiniBand is an important technology within the Exalogic system and is a key enabler of the Exabus. InfiniBand is used only as an internal communications fabric within the Exalogic system, however, or between a very small number of carefully tested Oracle products, such as Oracle Exadata and certain storage appliances and datacenter InfiniBand switches. At the moment, Oracle does not support connection of Exalogic to any third-party products using InfiniBand.

Installation of Exalogic is very straight-forward and can often be completed by Oracle's Advanced Customer Services professionals in a matter of hours, from start to finish.




Figure 15: Exalogic Datacenter Integration

Conclusion

Oracle Exalogic is an engineered hardware and software system designed to provide the ideal private cloud platform for Oracle's Fusion Middleware and business application platform for Oracle applications and third party applications. By deploying Exalogic Elastic Cloud, enterprises will:

- Accelerate time-to-value by eliminating uncertainty, complexity and effort across the application and infrastructure deployment lifecycle
- Improve ROI of IT investments by reducing CAPEX (in terms of software licenses and hardware outlay) and OPEX (in terms of data-centre power and cooling costs, people intensity and maintenance)
- Increase the performance and efficiency of existing Linux, Solaris and Java applications by up to 10X or more
- Improve application and infrastructure reliability, scalability and availability to levels needed for the most mission-critical systems

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



Oracle Exalogic Elastic Cloud: System Overview

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 © 2015, 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 is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

Published: 1/15