Oracle® Fusion Middleware

Enterprise Deployment Guide for Exalogic

Exalogic Release X2-2, X3-2, and X5-2

E64181-02

May 2017

ORACLE

Oracle Fusion Middleware Enterprise Deployment Guide for Exalogic, Exalogic Release X2-2, X3-2, and X5-2

E64181-02

Primary Author: Champa Sampat

Contributing Authors: Michael Rhys, Peter LaQuerre, Venkateswarlu Karnati, Mikael Fransson, Firdaus Fraz

Contents

6 Creating Exalogic Virtual Servers (vServers)

7 Preparing the Host Operating System

Preface

This preface provides supporting information for the *Oracle Fusion Middleware Enterprise Deployment Guide for Exalogic* and includes the following topics:

Audience

This guide is intended for system administrators who are responsible for installing and configuring Oracle Fusion Middleware enterprise deployments on Oracle Exalogic.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info or visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs if you are hearing impaired.

Related Documents

This section identifies additional documents related to Oracle Exalogic. You can access Oracle documentation online from the Oracle Technology Network (OTN) web site at the following URL:

http://docs.oracle.com/

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
italic	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
`monospace`	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Part I

Understanding an Enterprise Deployment on Exalogic

This topic provides a basic overview on Oracle Exalogic and a typical Exalogic enterprise deployment.

1

Understanding Oracle Exalogic

This chapter provides an overview of Oracle Exalogic and how Exalogic functions in an Oracle Fusion Middleware enterprise deployment.

In general, provides information on Exalogic components, architecture, Exalogic networking and deploying Exalogic with Exadata.

1.1 What is Exalogic?

Oracle Exalogic is an integrated hardware and software system designed to provide a complete platform for a wide range of application types and widely varied workloads.

Exalogic is intended for large-scale, performance-sensitive, mission-critical application deployments. It combines Oracle Fusion Middleware software and industry-standard Sun hardware to enable a high degree of isolation between concurrently deployed applications, which have varied security, reliability, and performance requirements. With Exalogic, you can develop a single environment that can support end-to-end consolidation of your applications.

1.2 Understanding Exalogic Components

Oracle Exalogic is delivered as a rack of hardware and in addition to the hardware components, Exalogic can be combined with Oracle Exalogic Elastic Cloud software, which consists of pre-integrated, standard technologies including the operating system, virtualization technology, networking software, device drivers, and firmware.

It consists of the following components:

- Compute Nodes (Servers)
- ZFS Storage (Storage Area Network/SAN)
- Integrated Infiniband Networking

For more information about Exalogic, see 'Introduction to Exalogic Machine' in the *Oracle Exalogic Elastic Cloud Machine Owner's Guide*.

1.3 About the Exalogic Hardware Architecture

This section describes the Oracle Exalogic hardware architecture.

Oracle Exalogic was tested extensively on a wide range of hardware configurations to arrive at the optimal configuration for middleware type deployments. Design considerations included high availability, compute density, state-of-the-art components, balanced system design, field serviceability, centralized storage, and high-performance networking.

Figure 1-1 Exalogic Hardware Architecture

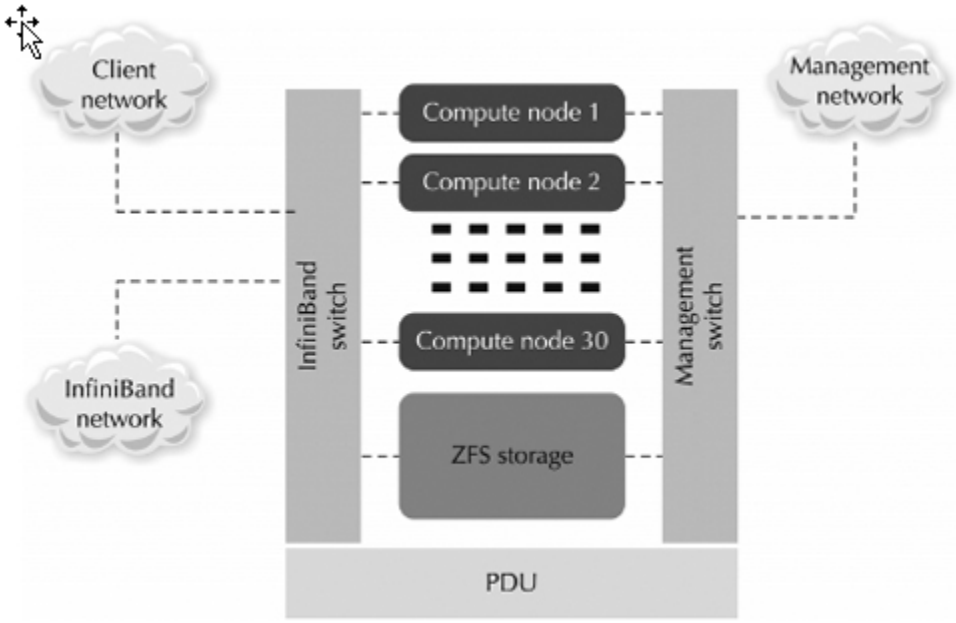

This section contains the following topics:

About Compute Nodes

The compute nodes are much like servers. These compute nodes contain CPUs, networking, and internal flash storage.

About Exalogic Storage
: Shared storage is provided by a Sun ZFS Storage ZS3 appliance, which is accessible by all the compute nodes. ZFS storage features optimized compression, performance and reliability optimizations, and is built in to the Exalogic machine.

About Exalogic Networking
: InfiniBand and Ethernet switches enable network communication in Exalogic.

1.3.1 About Compute Nodes

The compute nodes are much like servers. These compute nodes contain CPUs, networking, and internal flash storage.

Processing is performed by compute nodes. A full rack of Exalogic has 30 compute nodes, a half-rack has 16 compute nodes, a quarter-rack has 8 compute nodes, and a one-eighth rack has 4 compute nodes.

The compute node resembles traditional server hardware and is designed to be a general-purpose processing unit; however, its hardware and software have been specifically constructed and tuned to run Java-based middleware software.

Compute nodes are pre-loaded with the Exalogic Linux base image. They can be re-imaged with either a Solaris or Exalogic Elastic Cloud Software (EECS) server. You can run any type of application you want on a compute node if it is supported on the operating system.

Compute nodes balance high performance with high density. Density is a measure of computing power within a given amount of floor space in a data center. You could have multiple applications deployed on a single compute node. You could configure the compute node to have a backup compute node.

The number of processor cores on an Exalogic compute node depends on the machine version. For example, a compute node on a standard Exalogic X2-2 machine has two 6-core processors (12 cores in total), and a compute node on a standard X3-2 machine has two 8-core processors (16 cores in total).

1.3.2 About Exalogic Storage

Shared storage is provided by a Sun ZFS Storage ZS3 appliance, which is accessible by all the compute nodes. ZFS storage features optimized compression, performance and reliability optimizations, and is built in to the Exalogic machine.

With ZFS, storage has been specifically engineered to hold the binaries and configurations for both middleware and applications therefore reducing the number of installations and simplifying configuration management on the Exalogic system.

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 Quad Date Rate (QDR) 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 through standard network file systems supported by the compute node's operating system. ZFS cluster is made available to compute nodes or virtual machines, depending on the configuration.

1.3.3 About Exalogic Networking

InfiniBand and Ethernet switches enable network communication in Exalogic.

InfiniBand provides reliable delivery, security and quality of service at the physical layer in the networking stack, with a maximum bandwidth of 40Gb/s and latency down to 1 millisecond. The compute and storage nodes include InfiniBand network adapters, which are also referred to as host channel adapters (HCAs). The dual-port infiniband HCA provides a private internal network connecting the compute nodes and storage nodes to the system's I/O fabric.

In addition, the operating system images shipped with Exalogic are bundled with a suite of InfiniBand drivers and utilities called the OpenFabrics Enterprise Distribution (OFED). OFED is a core component of what Oracle refers to as the Exalogic Elastic Cloud Software. The Exalogic Elastic Cloud Software also includes optimizations that have been engineered into Oracle Fusion Middleware and that leverage OFED to provide higher performance over InfiniBand.

IB networking is used for all communications and data transfers within the Exalogic machine and can be used to connect multiple Oracle Engineered Systems together to create a very high performance, multi-purpose computing environment.

Although the hardware within Exalogic utilizes an InfiniBand fabric, the rest of your data center, along with the outside world, still speaks only Ethernet. This includes your application clients, such as web browsers, as well as legacy enterprise information systems, which components running within Exalogic may need to communicate with. Exalogic's switches and nodes enable this communication through the Ethernet over InfiniBand (EoIB) protocol. As the name suggests, EoIB gives InfiniBand devices the ability to emulate an Ethernet connection using IB hardware.

1.4 About Oracle Exalogic Elastic Cloud

Oracle Exalogic Elastic cloud is Oracle's first engineered system for Enterprise Java Applications.

These applications include Oracle Fusion Middleware and any application which can run on one of the Exalogic supported Operating systems namely Linux or Solaris. Hardware and software are engineered together to optimize extreme Java performance.

Oracle Exalogic can also be used for consolidation, by the addition of Oracle Elastic Cloud software and the Exalogic platform can be used to serve virtual server farms.

Figure 1-2 Oracle Exalogic Elastic Cloud

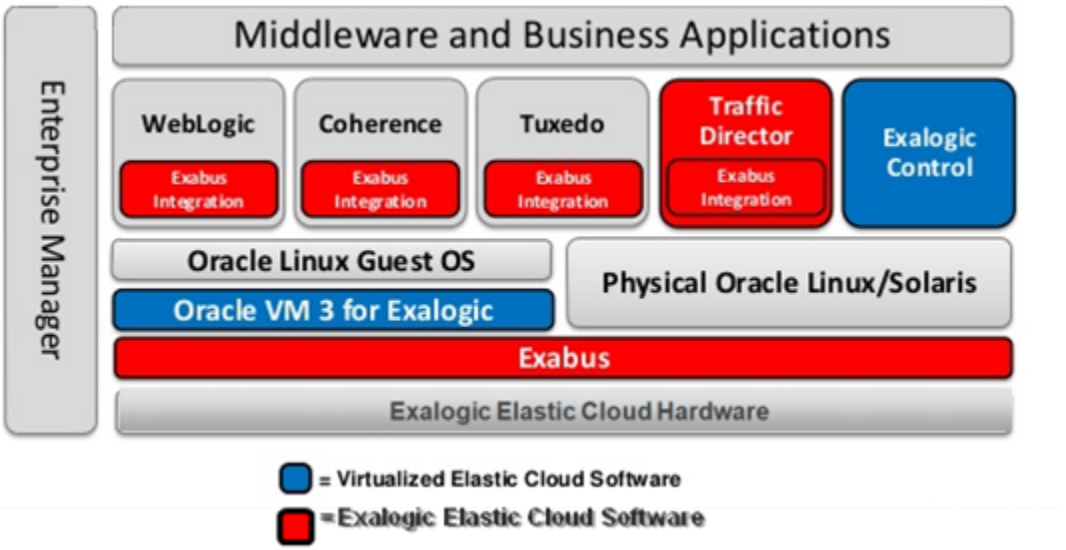

Oracle has made unique optimizations and enhancements to Exalogic components, as well as Oracle's Fusion middleware and Oracle's applications, which includes onchip network virtualization, high performance Remote Direct Memory Access (RDMA) at operating system and Java Virtual Machine (JVM) layers and Exalogicaware workload management in Oracle WebLogic Server (Oracle's Java EE application server), to meet the highest standards of reliability, availability, scalability and performance.

Exalogic Elastic Cloud comprises Exabus, which is a set of hardware, firmware, and software optimizations that enable the operating system, middleware components, and even certain Oracle applications to make full use of the infiniband fabric and the Oracle Traffic Director.

The InfiniBand network fabric offers extremely high bandwidth and low latency, which provides major performance gains with respect to communication between the application server and the database server, and with respect to communication between different application server instances running within the Exalogic system.

Figure 1-3 Exalogic Elastic Cloud Software (v2.X) Performance Benchmark

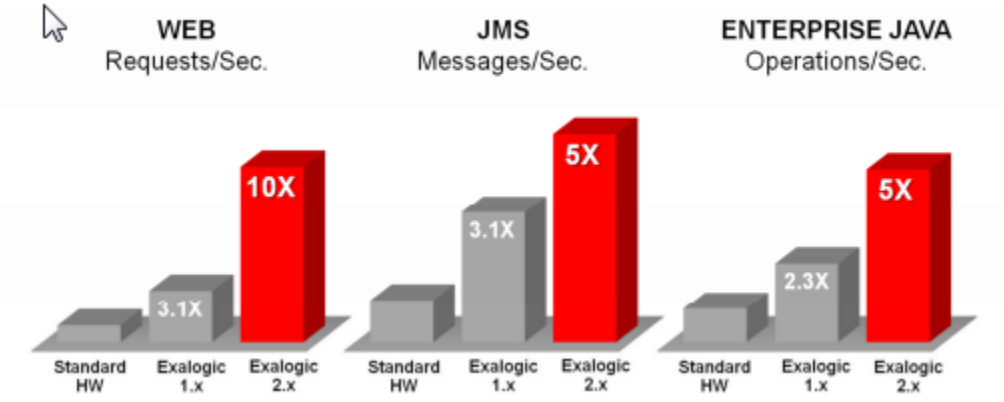

The current release of the Exalogic Elastic Cloud software includes a tightly integrated server virtualization layer with unique capabilities allowing the consolidation of multiple, separate virtual machines containing applications or Middleware on each

server node while introducing essentially no I/O virtualization overhead to the Exabus InfiniBand network and storage fabric.

Physically, Oracle Exalogic Elastic Cloud can be viewed as a rack of physical server machines plus centralized storage, which all have been designed together to cater to typical high-performance Java application use cases.

This section contains the following topics:

1.4.1 Understanding Exalogic Elastic Cloud Architecture

This section describes the Exalogic elastic cloud architecture.

The Exalogic system consists of the following two major elements:

- Exalogic X52 - 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 preintegrated 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.
  - WebLogic Server - Session replication uses the SDP layer of IB networking to maximize performance of large scale data operations as this avoids some of the typical TCP/IP network processing overhead. When processing HTTP requests, WebLogic Server makes native use of the SDP protocol when called by Oracle Traffic Director, or when making HTTP requests to it. Through its Active Gridlink for RAC feature, WebLogic Server JDBC connections and connection pools can be configured to use the low level SDP protocol when communicating natively with Exadata over the IB fabric.
  - Coherence - Cluster communication has been dramatically redesigned to further minimize network latency when processing data sets across caches. Its elastic data feature increases performance in conjunction with the compute nodes built in solid state drives by optimizing both the use of RAM and garbage collection processing to minimize network and memory use. When sending data between caches it uses only a RDMA level IB verb set, thus avoiding nearly all the TCP/IP network processing overhead.
  - Tuxedo - Tuxedo has been similarly enhanced to make increasing use of SDP and RDMA protocols to optimize the performance of inter-process communications within and between compute nodes.

1.4.2 Commissioning an Oracle Exalogic Elastic Cloud

Oracle Fusion Middleware software has been enhanced with performance optimizations for deployment on Exalogic.

You can follow individual software documentation for specific Fusion Middleware Applications.

For example:

- Identity And Access Enterprise Deployment Guide
- Web Center Enterprise Deployment Guide
- SOA Enterprise Deployment Guide

This document does not describe how to commission your Exalogic hardware or how to install the Oracle Exalogic Elastic Cloud software. For information on how to commission your Exalogic hardware refer to Oracle Exalogic Documentation Library Exalogic Release EL X2-2, X3-2, X4-2, and X5–2.

Weblogic Server Exalogic optimizations can be enabled following the "Exalogic Elastic Cloud Software Support" section and refer to Core Server in *What's New in Oracle WebLogic Server guide*.

This document does not describe how to commission your Exalogic hardware or how to install the Oracle Exalogic Elastic Cloud software. For information on how to do this, you should refer to the documentation:

Oracle Exalogic Documentation Library Exalogic Release EL X2-2, X3-2, X4-2, and X5–2.

http://docs.oracle.com/cd/E18476_01/

1.5 Understanding Exalogic Networking

This topic provides information on Exalogic networking.

The following topics describe how an Exalogic machine is networked:

Network Diagram for Exalogic Machine
: This topic provides the information on network diagram for Exalogic machine.

Understanding Network Protocols
: In an Exalogic deployment, all networking is via Infiniband.

1.5.1 Network Diagram for Exalogic Machine

This topic provides the information on network diagram for Exalogic machine.

Figure 4-1 shows the network diagram for an Oracle Exalogic machine.

Figure 1-4 Exalogic Machine Network Overview

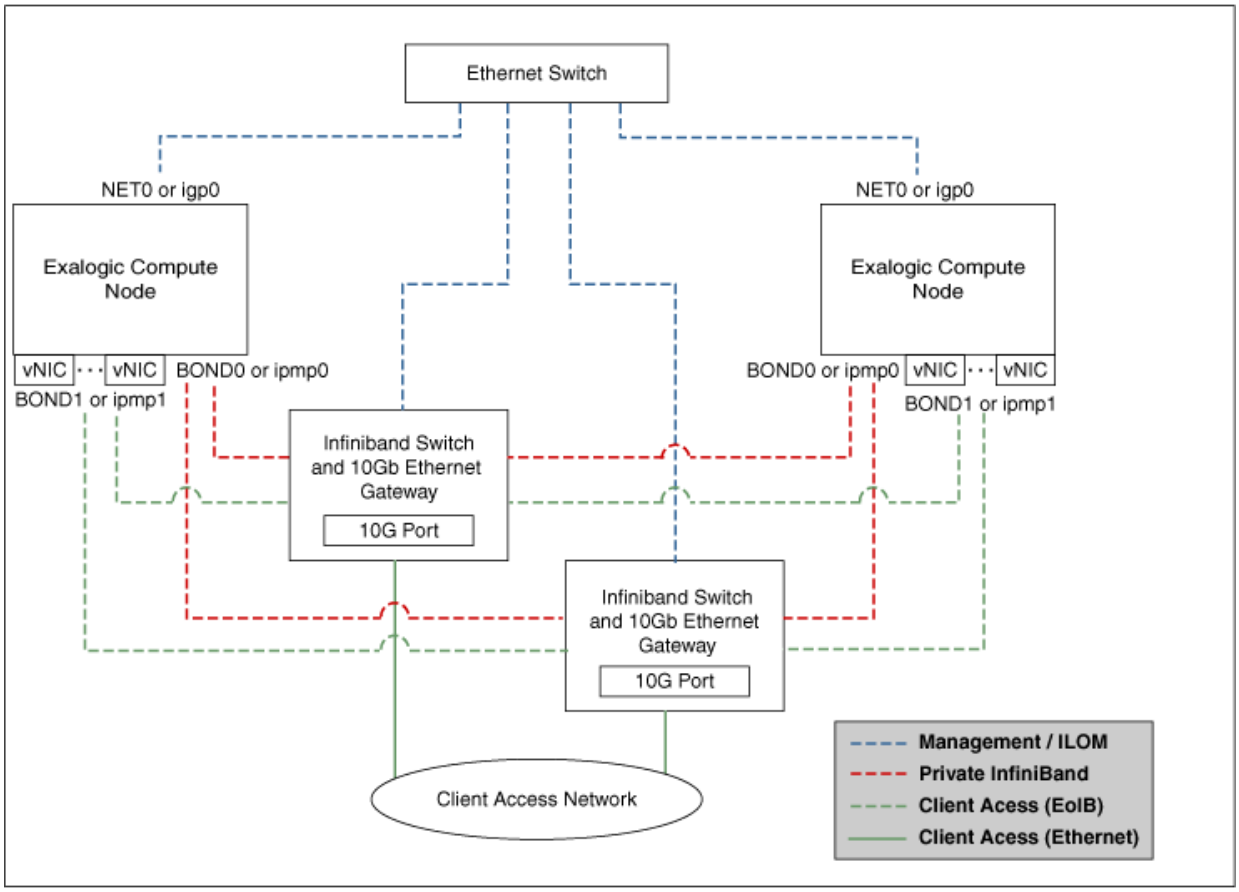

The schematic representation of Oracle Exalogic machine's network connectivity includes the following:

- Default `BOND0` interface, which is the private InfiniBand fabric including the compute nodes connected via Sun Network QDR InfiniBand Gateway Switches

  Typical Uses of this network are:

  - To communicate between compute nodes
  - To access the internal Oracle ZFS Storage Appliance and other Engineered Systems on the fabric
  - To communicate between vServers
  - InfiniBand partitions and memberships provide network isolation and security

Note:

InfiniBand `BOND0` interfaces are the default channel of communication among Exalogic compute nodes and storage server head. IP subnets and additional bonds can be added on top of this default bonded interface.

The device nodes representing the IPoIB network interface for Oracle Linux are referred to as `ib0` and `ib1`. The corresponding logical devices created by Oracle Solaris are referred to as `ibp0` and `ibp1`. The default IPoIB bonded interface `BOND0` or `IPMP0`, configured by the Exalogic Configuration Utility, comprises these Linux-specific interfaces or Solaris-specific interfaces, respectively.

- `BOND1` interface, which is the Ethernet over InfiniBand (EoIB) link

  Typical Uses of this network are:

  – EoIB External Management Network on a vLAN

  – The IP address provided in the ECU spreadsheet created by the ECU configuration process

  – Used for Cloud Administration via Exalogic Control

  – EoIB user access networks on separate vLANs

  – Created by the Exalogic Administrator at post-Exalogic installation

  – Used to access guest vServers and their application services

  Note:

  The device nodes representing the EoIB network interface for Oracle Linux are referred to as `vnic0` and `vnic1`. The Linux kernel creates `eth` device nodes that correspond to the `vnic0` and `vnic1` instances that are created on the Sun Network QDR InfiniBand Gateway Switch.

  The corresponding logical devices created by Oracle Solaris are referred to as `eoib0` and `eoib1`. The EoIB bonded interface `BOND1` or `IPMP1` must be configured manually. When you configure them, choose the network interfaces specific to your operating system.

- `NET0` interface, which is associated with the host Ethernet port 0 IP address for every compute node and storage server head

  Typical Uses of this network are:

  – To access all physical components and ILOMs

  – To perform system administration and life cycle management

  – Used by the Exalogic Control stack

Note:

The device node representing the management network interface for Oracle Linux is referred to as `eth0`. The corresponding logical device created by Oracle Solaris is referred to as `igb0`.

- Client access network for external data center connectivity

1.5.2 Understanding Network Protocols

In an Exalogic deployment, all networking is via Infiniband.

Most likely, your corporate network is Ethernet based. However, you can configure the Infiniband networks so that they understand Ethernet traffic, and then, you can attach your Exalogic machine to the corporate network. This is known as Ethernet over Infiniband (EoIB). The EoIB network communicates with your corporate network using 10 GB Ethernet. This network is known as the client/public/external network.

If you are communicating with other components inside the Exalogic machine, then you do not need to use Ethernet. InfiniBand adapters (HCAs) provide advanced features that can be used via the native "verbs" programming interface:

- Data transfers can be initiated directly from user space to the hardware, bypassing the kernel and avoiding the overhead of a system call.
- The adapter can handle all of the network protocol of breaking a large message (even many megabytes) into packets, generating ACKs, retransmitting lost packets, etc. without using any CPU on either the sender or receiver.
- IPoIB (IP-over-InfiniBand) is a protocol that defines how to send IP packets over IB; for example, Linux has an "ib_ipoib" driver that implements this protocol. This driver creates a network interface for each InfiniBand port on the system, which makes an HCA act like an ordinary NIC.

IPoIB does not make full use of HCA capabilities; network traffic goes through the normal IP stack. This means a system call is required for every message, and the host CPU must handle breaking data up into packets. However, it does mean that applications that use normal IP sockets will work on top of the full speed of the IB link.

IPoIB provides a normal IP NIC interface that can run TCP (or UDP) sockets on top of it.

SDP (Sockets Direct Protocol) is a transport-agnostic protocol to support stream sockets over Remote Direct Memory Access (RDMA) network fabrics. It is specifically designed for Infiniband networks.

The purpose of the Sockets Direct Protocol is to provide a RDMA-accelerated alternative to the TCP protocol on IP. The goal is to do this in a manner that is transparent to the application.

SDP only deals with stream sockets, and if installed in a system, bypasses the operating system resident TCP stack for stream connections between any endpoints on the RDMA fabric. All other socket types (such as datagram, raw, packet, etc.) are supported by the Linux IP stack and operate over standard IP interfaces (that is, IPoIB on InfiniBand fabrics). The IP stack has no dependency on the SDP stack; however, the SDP stack depends on IP drivers for local IP assignments and for IP address resolution for endpoint identifications.

The IPoIB network is known as the internal network in this guide.

Both networks (EoIB and IPoIB) are accessed through an attached IP address. If you want to route traffic through the EoIB network, you send traffic through the IP address associated with that network. Similarly if you want traffic to go through the internal network, you use the IP address associated with that network. For example:

`host1-int` is associated with the internal (IPoIB) network

`host1-ext` is associated with the external (EoIB) network

If you want to communicate with host1 through the internal network, you send traffic to `host1-int`. If you want to use the external network, use `host1-ext`.

1.6 About Deploying Exalogic with Exadata

Most Oracle Fusion Middleware applications will interact with an Oracle Database. This database can reside on external hardware that is connected to the Exalogic machine via the external 10 GB Ethernet connected to the datacenter network.

However, you can obtain maximum performance if your database resides on an Exadata appliance connected directly to the Exalogic machine. If Exalogic is connected to Exadata, you have the option of communicating with the database using IP over Infiniband (IPoIB).

For information about connecting an Oracle Exalogic machine to an Oracle Exadata Database machine, see the *Oracle Fusion Middleware Exalogic Machine Multirack Cabling Guide*.

1.7 Understanding Types of Deployment

This section describes the types of Exalogic deployment.

You can configure Exalogic in physical or virtual deployment.

This section contains the following topics:

1.7.1 About a Physical Exalogic Configuration

In a physical Exalogic configuration, the application software is deployed on compute nodes. Each compute node runs its own single operating system.

All applications, including WebLogic Server, Coherence, and Tuxedo, then share this operating system kernel and the local compute node resources.

The Exalogic compute nodes are engineered servers and thus provide extreme performance to Java-based Middleware software deployed on the compute nodes.

This configuration does not include EECS and Middleware. In addition, applications running on the Exalogic platform are deployed and managed in very much the same way as they are on traditional platforms; new deployments are associated with appropriate physical compute, storage, memory and I/O resources. For more information, see Enterprise Manager in *Manager Cloud Control Managing and Monitoring an Oracle Exalogic Elastic Cloud Machine guide* which is the primary administration tool.

1.7.2 About a Virtual Exalogic Configuration

The purpose of server virtualization is to fundamentally isolate the operating system and applications stack from the constraints and boundaries of the underlying physical servers. By doing this, multiple virtual machines can be presented with the impression that they are each running on their own physical hardware when, in fact, they are sharing a physical server with other virtual machines.

This allows server consolidation in order to maximize the utilization of server hardware, while minimizing costs associated with the proliferation of physical servers-namely hardware, cooling, and real estate expenses.

This hardware isolation is accomplished either through a software based sharing or a direct device assignment (where a I/O device is directly assigned to a VM). Software based sharing is achieved by inserting a very thin layer of software between the operating system in the virtual machine and the underlying hardware to either directly emulate the hardware or to otherwise manage the flow and control of everything from CPU scheduling across the multiple VMs, to I/O management, to error handling.

The challenge with virtualization is to achieve a high enough consolidation ratio to achieve the cost benefits you need while still being able to provide the exceptional, predictable performance required from your core applications.

The Oracle Exalogic Elastic Cloud provides a unique Input/Output subsystem called Exabus. Exabus employs a converged network fabric to provide all Input/Output services to the applications running within an Exalogic system. Applications residing within an Exalogic system can access all network services provided within the datacenter network through Exabus.

In the latest version of Oracle Exalogic, Oracle has virtualized the InfiniBand connectivity in Exabus, using state-of-the-art, standards-based technology to permit the consolidation of multiple virtual machines per physical server with no impact on performance.

Exalogic includes support for a highly optimized version of the EECS 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.

The logical vServers can have specific amounts of physical compute, storage, memory and I/O resources, optionally pre-configured with middleware and applications. This approach allows for maximum levels of resource sharing and agility as vServers can share physical resources and can be provisioned in minutes. Pre-configured OVM templates for Oracle applications are available to download.

EECS has been engineered for tight integration with Exalogic's Exabus I/O backplane using a technique called Single Root I/O Virtualization (SRIOV).

SR-IOV eliminates virtualization overhead to deliver the maximum performance and scalability, while also allowing the same InfiniBand I/O adapter to be shared by up to 63 virtual machines, each with a redundant pair of InfiniBand connections, enabling highly efficient, consolidated operations. SR-IOV's unique ability to nearly eliminate virtualization overhead while still allowing the sharing of hardware permits a much higher server consolidation ratio and higher performance.

1.7.3 About Choosing a Type of Deployment

Both of the Exalogic implementation styles (physical and virtual) can support the creation of a private cloud.

In a virtualized system, Exalogic Control is used to define, manage, and monitor cloud users and services. In a physical system, equivalent functionality is provided by Enterprise Manager with the Cloud Management Pack.

Among the benefits of using virtualized approach is application consolidation, tenant isolation (provision secure Exalogic resources to multiple tenants), deployment simplification, including scaling up or down. With the advent of Exalogic Elastic Cloud technology, the impact of virtualization on application throughput and latency has been minimized to 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.

If you deploy your application using a bare metal (physical) deployment, then you have at your disposal the raw processing power of the compute node. A single compute node is likely to offer more processing power than a single component of your application needs. So, to make the best use of the processing power available, several application components or applications will be installed onto the same compute node.

If you deploy your application using a virtual deployment, you can modularize the components, creating several smaller virtual servers to provide application component isolation. In a virtual deployment, if the underlying hardware fails, then a virtual server can be moved to a different underlying physical host to resume processing. Having a distributed deployment allows you to isolate failures to smaller areas.

2

Understanding a Typical Exalogic Enterprise Deployment

This chapter introduces and describes how an Oracle Fusion Middleware enterprise deployment is typically deployed on Exalogic hardware.

For information regarding specific product deployment architectures on Oracle Exalogic, refer to the appropriate product-specific enterprise deployment guide.

This guide gives an overview of what you need to consider when you are deploying a product on Exalogic. This guide focuses on the steps needed to set up the Exalogic appliance for deploying Oracle Fusion Middleware in an enterprise deployment. Note that this document does not cover how to install Oracle Fusion Middleware products. For information on installing Oracle Fusion Middleware, refer to the enterprise deployment guide specific to the product you are deploying.

2.1 Why Install Oracle Fusion Middleware on Exalogic?

Oracle Exalogic is a highly available, highly performant integrated hardware appliance from Oracle. When deployed onto Exalogic, Oracle Fusion Middleware

components benefit from Exalogic's superior networking, resulting in improved application throughput.

In addition, Oracle WebLogic Server has had a number of optimizations put into place to enable it to run faster on Oracle Exalogic. These optimizations further increase the throughput of the applications deployed.

Deploying Oracle Fusion Middleware on Exalogic ensures that you have a highly available infrastructure that will provide you with maximum availability and performance.

2.2 Understanding the Types of Exalogic Deployment

This guide focuses on two primary reference topologies for deploying Oracle Fusion Middleware on Exalogic. Although the components installed are essentially the same, how the Exalogic appliance is commissioned is different for each topology.

The typical Exalogic deployments are the Physical Exalogic deployment and the Exalogic Elastic Cloud deployment:

- Physical Exalogic Deployment. In this deployment type, the raw processing power of the Exalogic compute node is used directly. Because of the power of the compute node, more products are deployed on a single compute node to make the most of the processing resources available.
- Exalogic Elastic Cloud Deployment. In this deployment, applications do not have direct access to the underlying power of the compute nodes. In an Exalogic Elastic Cloud deployment, the processing power is virtualized. This deployment allows users to create a number of virtual machines, which can be moved between base compute nodes as required. In this deployment, products are compartmentalized. In other words, products are divided into a large number of smaller virtual machines.

Going forward, this guide will refer to the two different types of deployment as physical and virtual.

These topologies are very similar to the standard platform topologies. That is, there is a distributed topology that is more suited to virtual Exalogic deployments and a consolidated topology that is more suited to physical Exalogic deployments. However, there is a third topology that is a hybrid topology, which uses some components installed in Exalogic and others namely the web tier, installed on commodity hardware outside of the Exalogic machine.

The exact Oracle Fusion Middleware topology you install and configure for your organization might vary, but for the three primary topologies, this guide provides the necessary steps to set up your Exalogic environment for an enterprise deployment.

2.3 Diagrams of the Primary Exalogic Enterprise Deployment Topologies

The following sections provide diagrams of the primary Exalogic enterprise deployment topologies.

Note: Each of the diagrams below use arbitrary hostnames such as HOST1 or HOST2. These names are used for illustrative purposes only. In real deployments these names are changed to reflect the host names that are used in the deployment.

This section contains the following topics:

2.3.1 Diagram of a Typical Physical Exalogic Topology

The Figure 2-1 shows a diagram of the physical Exalogic enterprise deployment topology with Oracle Traffic Director.

Figure 2-1 Physical Exalogic Deployment with Oracle Traffic Director

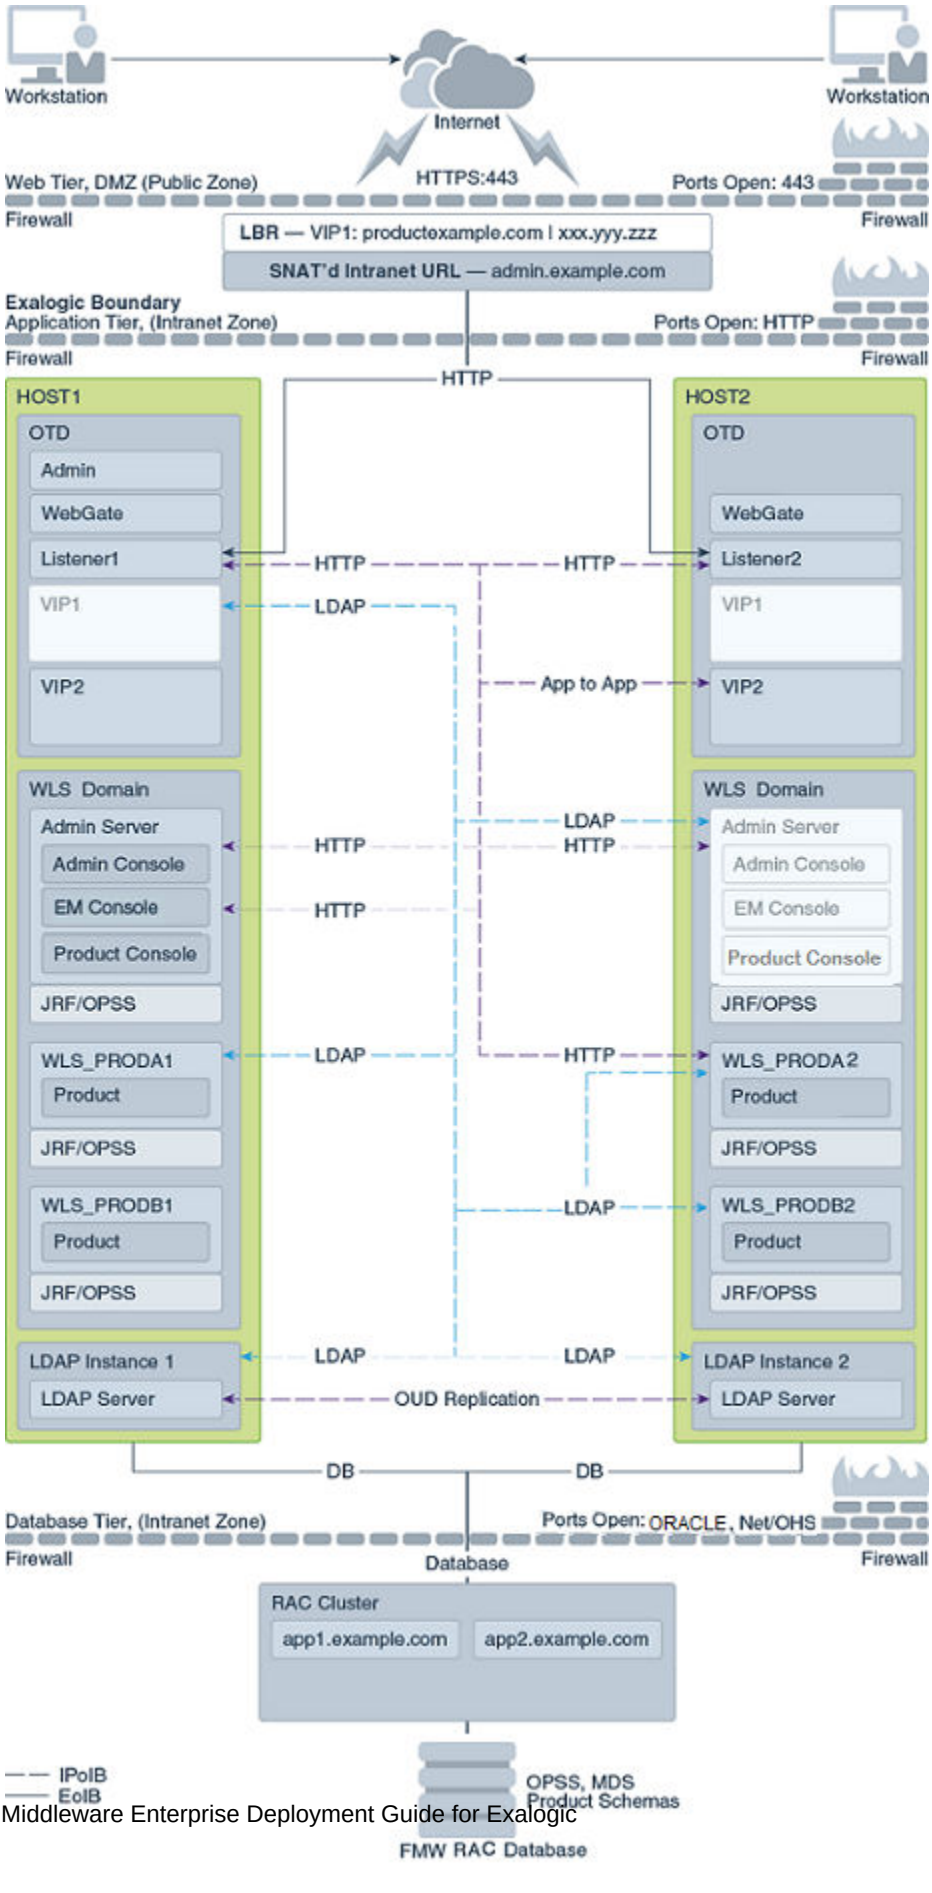

For a description of the standard elements shown in the diagram, see Understanding the Typical Enterprise Deployment Topology Diagrams.

2.3.2 Diagram of a Typical Virtual Exalogic Topology

The Figure 2-2 shows a diagram of the virtual Exalogic enterprise deployment topology.

Figure 2-2 Typical Virtual Exalogic Topology

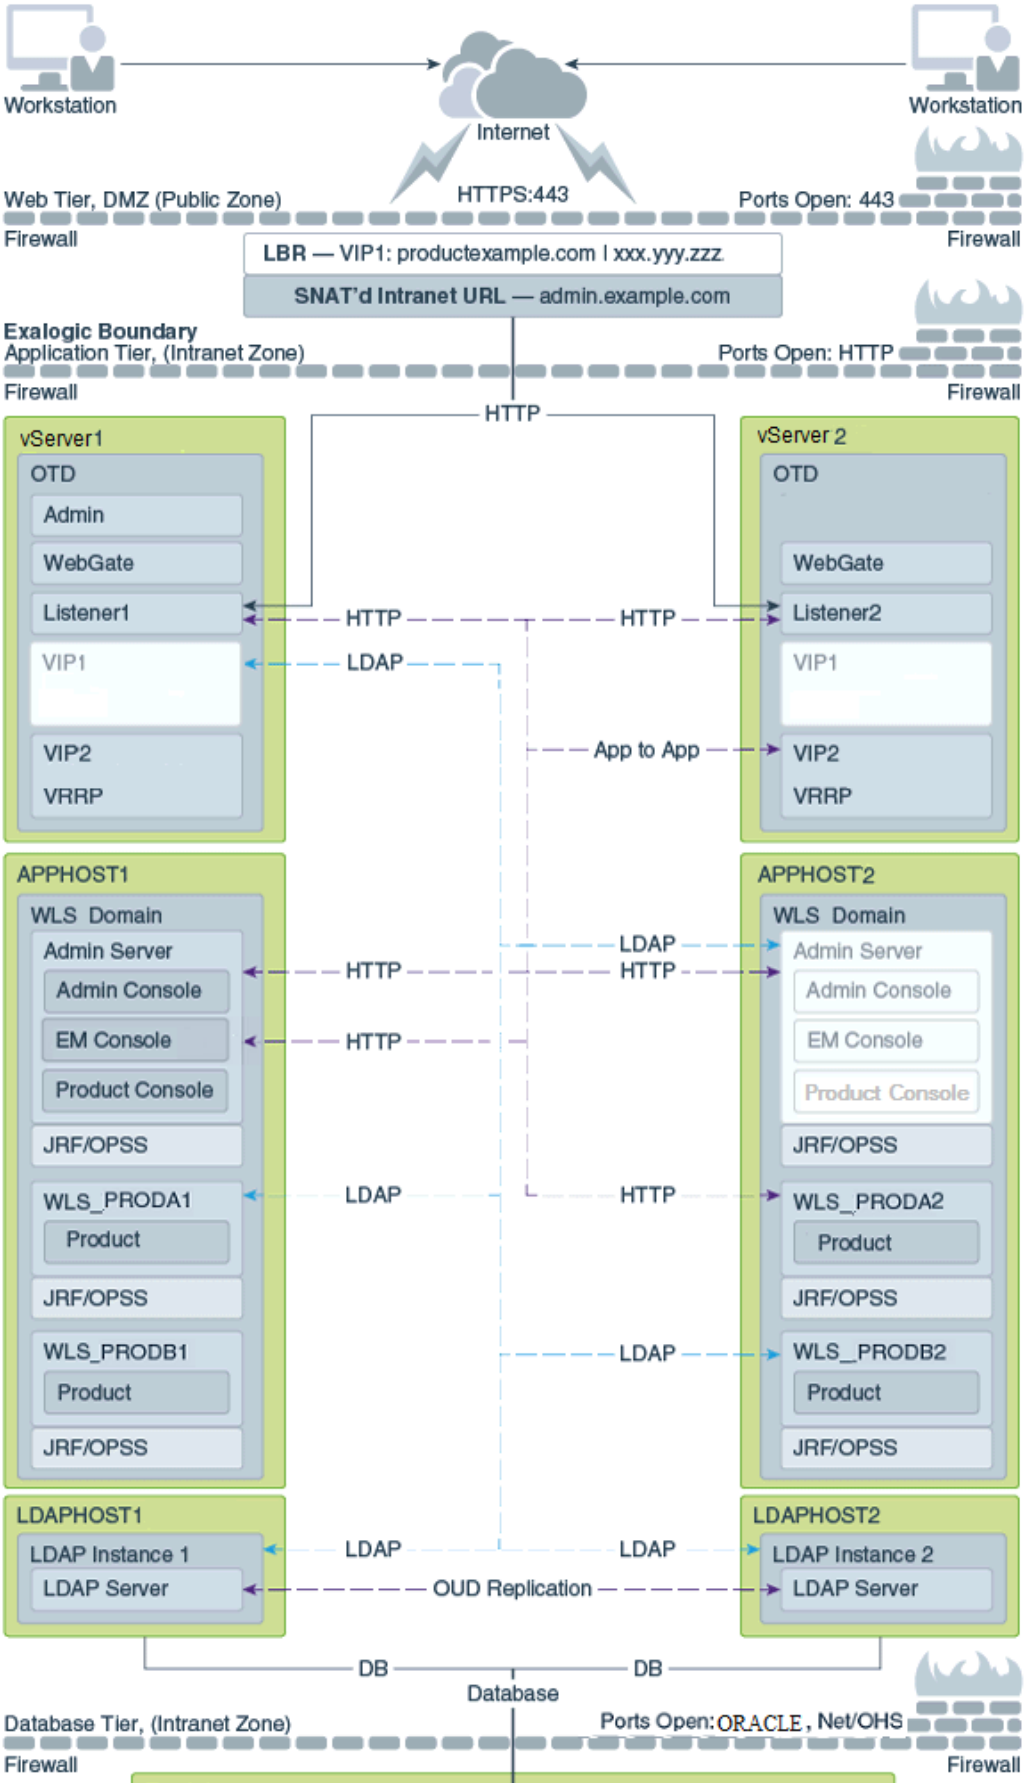

For a description of the standard elements shown in the diagram, see Understanding the Typical Enterprise Deployment Topology Diagrams.

2.3.3 Diagram of an Exalogic Deployment with an External Web Tier

The Figure 2-3 shows a typical enterprise deployment, including the Web tier, application tier and data tier.

Figure 2-3 Typical Enterprise Deployment with Web Tier

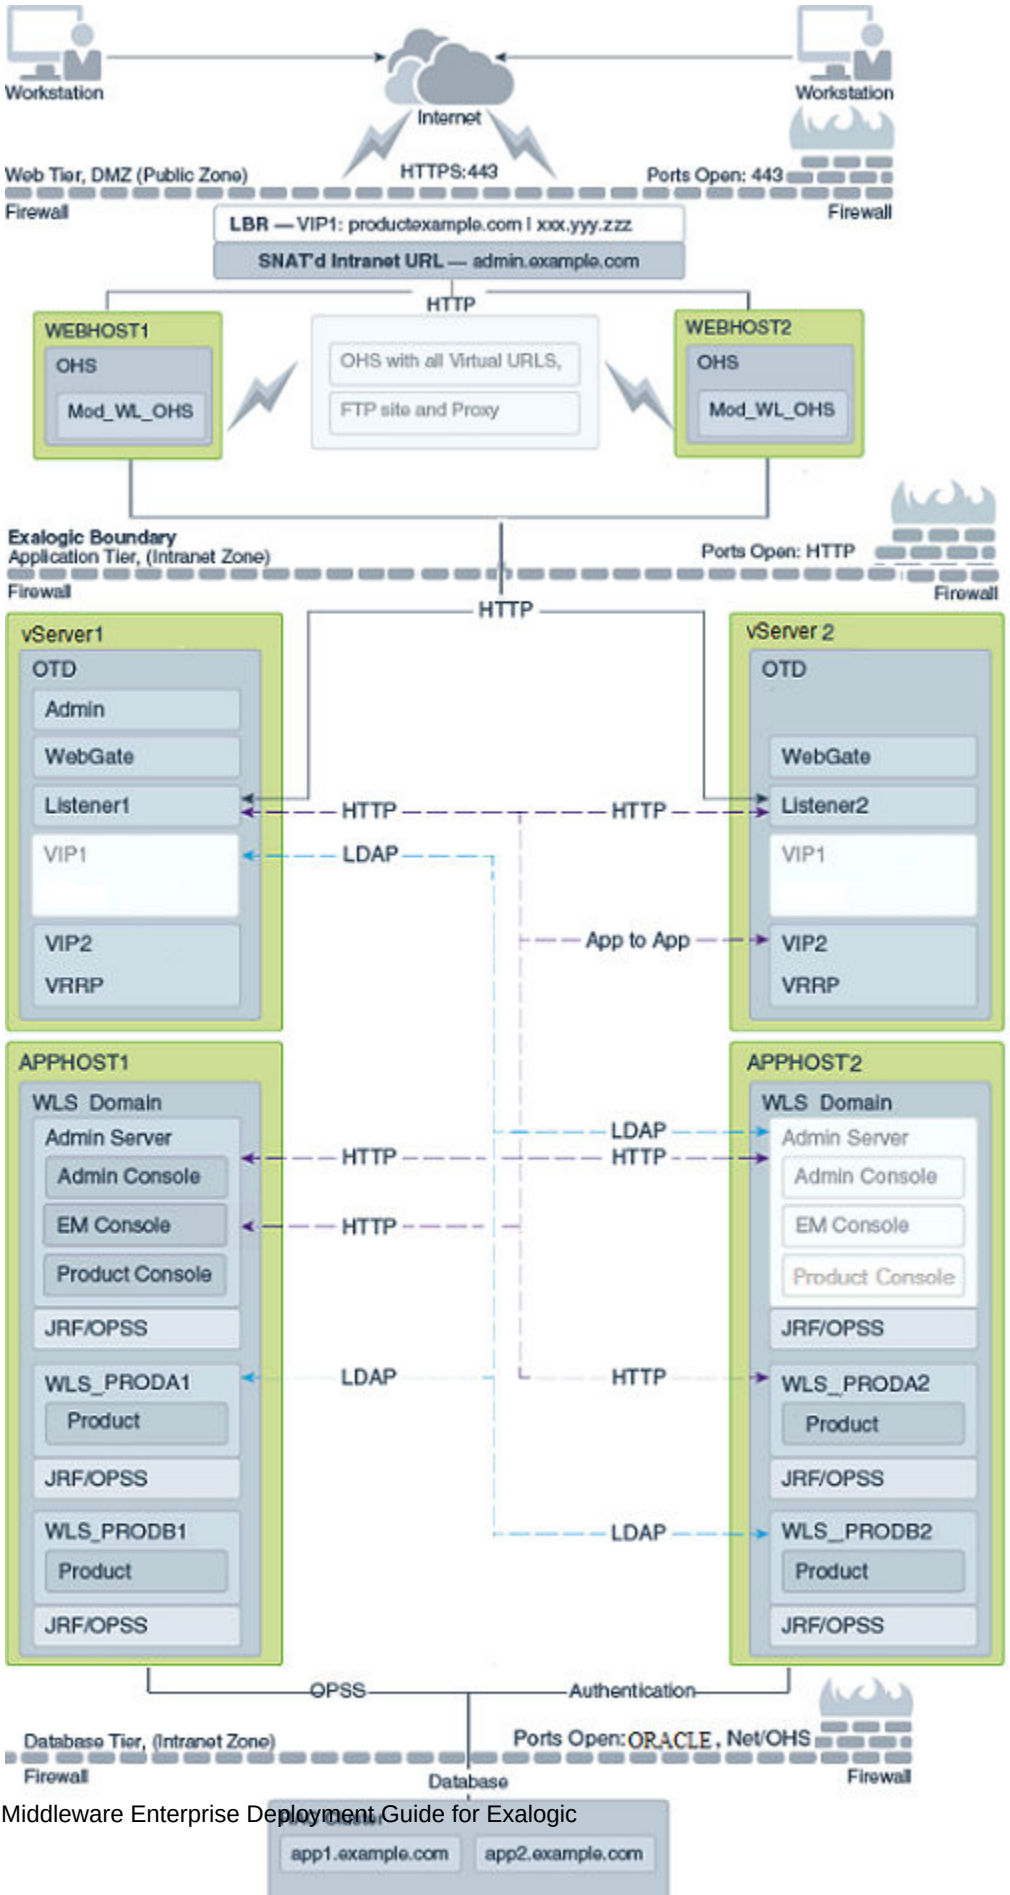

2.4 Understanding the Typical Enterprise Deployment Topology Diagrams

It provides information on Oracle Fusion Middleware and Exalogic networking and elements of a Enterprise deployment topology.

The following topics provide conceptual information about the typical enterprise deployment topology diagrams:

Understanding the Firewalls and Zones of a Typical Exalogic Enterprise Deployment
: When you deploy Oracle Fusion Middleware on Exalogic, most (if not all) of the components are installed inside the Exalogic appliance. A typical platform enterprise deployment is distributed among various tiers. Each tier is separated by a firewall.

Understanding Oracle Fusion Middleware and Exalogic Networking
: One of the core advantages of Exalogic is its fast, flexible network. When you are deploying Oracle Fusion Middleware on Exalogic, you have to consider how you want to use Oracle Exalogic networking.

Understanding the Elements of a Typical Enterprise Deployment Topology
: The enterprise deployment topology consists of a hardware load balancer, web tier, application tier, directory tier and a data tier.

Receiving Requests from the Internet
: The enterprise deployment topology is fronted by a hardware load balancer, which directs incoming HTTP and HTTPS requests from the Internet to the Web tier.

Understanding the Web Tier
: The Web tier of the reference topology consists of two Oracle Traffic Director instances.

Understanding the Application Tier
: The application tier consists of two or more hosts, where Oracle WebLogic Server and the Oracle Fusion Middleware products are installed and configured.

Understanding the Directory Tier
: Oracle Fusion Middleware Products often interact with an LDAP directory. The diagrams above topic depicts how an LDAP directory can be access in such a topology.

Understanding the Data Tier
: In the Data tier, an Oracle RAC database runs on the two hosts (DBHOST1 and DBHOST2). The database contains the schemas required by the Oracle Fusion Middleware components and the Oracle Platform Security Services (OPSS) policy store.

2.4.1 Understanding the Firewalls and Zones of a Typical Exalogic Enterprise Deployment

When you deploy Oracle Fusion Middleware on Exalogic, most (if not all) of the components are installed inside the Exalogic appliance. A typical platform enterprise deployment is distributed among various tiers. Each tier is separated by a firewall.

Since all the hardware components are incorporated into the Exalogic appliance, the number of tiers available is reduced.

There is no need for an application or directory tier. Also, if your Exalogic appliance is linked to an Exadata appliance, then there is no need for a data tier.

In the enterprise deployment topology diagrams, the following two zones are used, which are each separated by a firewall:

- The DMZ, in which the load balancer resides. This zone is accessible only through virtual server names defined on the load balancer.
- The Exalogic zone, in which all application components reside. In this zone, only those components requiring access to external resources are visible on the corporate network.

Note that a common variation on these topologies is to move the web server to the DMZ. This can offer increased security.

2.4.2 Understanding Oracle Fusion Middleware and Exalogic Networking

One of the core advantages of Exalogic is its fast, flexible network. When you are deploying Oracle Fusion Middleware on Exalogic, you have to consider how you want to use Oracle Exalogic networking.

This section contains the following topics:

2.4.2.1 Types of Network

There are three types of network within an Exalogic appliance.

- IP over Infiniband (IPoIB): This is the internal Infiniband network that connects the internal components of the Exalogic appliance. This network is fast, but it cannot be connected to the outside world. The benefit of this network is that it can be used to ensure that network traffic is kept private from the outside world. The downside to using this network is that external components cannot directly access application components inside the Exalogic appliance.
- Ethernet Management Network (eth0): This management network is used for connecting to the Exalogic components through the built-in Ethernet network. This network is only used for management operations and should not be used for production deployments. This network is used to login to the Exalogic components to configure them.
- Ethernet over Infiniband (EoIB): This network also uses the Exalogic Infiniband network, but it is possible to connect this network to the standard corporate network. This allows external components to talk directly to components inside Exalogic. This network is always used for communication between your hardware load balancer and Oracle Traffic Director.

2.4.2.2 Considerations for Choosing your Exalogic Network

By default, some components within Oracle Fusion Middleware only talk on a single network. Other components, such as WebLogic Managed Servers and Oracle Traffic Director, can be configured to talk on both the internal and external networks.

Oracle Traffic Director is the preferred load balancer for traffic once it enters the Exalogic appliance. Oracle Traffic Director is also the preferred Web Server for deployments that reside completely within Exalogic.

When choosing which Exalogic network to use, consider the following:

- If you are using an external Web Tier, then you should configure your components where it listens on EoIB network and routes to IPoIB.
- If you expect that all traffic will come through Oracle Traffic Director and will stay within the Exalogic appliance once it reaches there, then you should choose to configure your components to use the IPoIB network.
- If you expect all of your LDAP traffic to originate within the Exalogic appliance, then you should configure your LDAP server to use the IPoIB network.
- If your database resides in an Exadata appliance that is directly connected to the Exalogic appliance, then you should use the IPoIB network.
- If you are using Exadata to host your databases and these databases are accessed from both Exalogic and Non-Exalogic sources then the Database Listeners need to be configured to Listen on both the IPoIB and EoIB networks.

Note that additional configuration is required if you want to configure WebLogic Managed Servers to listen on multiple networks using different channels. Your usage of the Exalogic network depends on your access requirements. It may be that the solution you adopt encompass elements of both the internal and external network.

There is no mandate rule on which network you use. There are no significant performance gains of one over the other. All components work equally well on either network.

2.4.3 Understanding the Elements of a Typical Enterprise Deployment Topology

The enterprise deployment topology consists of a hardware load balancer, web tier, application tier, directory tier and a data tier.

The enterprise deployment topology consists of the following high-level elements:

- A hardware load balancer, which routes requests from the Internet to the Web servers in the Web tier. In this case, a Web server can either be an external Web server on commodity hardware or Oracle Traffic Director inside the Exalogic host.
- A Web tier, which consists of two or more host computers hosting Web server instances (for load balancing and high availability).
  The Web server instances are configured to authenticate users (via an external identity store and a single sign-on server) and then route the HTTP requests to the Oracle Fusion Middleware products and components running in the Application tier.
  The web tier instances are also used to load balance internal requests without the need for those requests to go through the external hardware load balancer. The Web Server used is typically an Oracle Traffic Director or in the case of an External Web Tier, it is an Oracle HTTP Server.

- An Application tier, which consists of two or more hosts hosting a cluster of Oracle WebLogic Server Managed Servers and the WebLogic Administration Server for the domain. The Managed Servers are configured to run the various Oracle Fusion Middleware products, such as Oracle Identity and Access Management or Oracle SOA Suite.
- A Directory tier, which consists of two or more servers hosting a LDAP-compliant directory, such as Oracle Unified Directory.
- A data tier, which consists of either two or more physical hosts hosting an Oracle RAC Database or an Exadata appliance hosting an Oracle RAC database.

Note:

This topic describes the different tiers of a typical enterprise deployment topology. In this guide, tiers refer to logical separations. In a commodity deployment, tiers are typically separated by firewalls. However, in an Exalogic deployment, this is not necessary as the traffic is confined to the Exalogic machine. In an Exalogic deployment, a firewall typically exists only between the DMZ and the Exalogic appliance.

2.4.4 Receiving Requests from the Internet

The enterprise deployment topology is fronted by a hardware load balancer, which directs incoming HTTP and HTTPS requests from the Internet to the Web tier.

This section contains the following topics:

Purpose of the Hardware Load Balancer
: The hardware load balancer balances the load on the Web tier by receiving requests to a virtual host name and then routing each request to one of the Web server instances, based on a load balancing algorithm.

Specific internal-only communications between the components of the Application tier
: Communications between the Oracle Fusion Middleware components and applications on the application tier are handled by the load balancing capabilities of Oracle Traffic Director (OTD).

Summary of Virtual Server Names
: This section describes the virtual server names recognized by the hardware load balancer and Oracle Traffic Director (OTD) in a typical enterprise deployment.

HTTPS versus HTTP Requests to the External Virtual Server Name
: When you configure the hardware load balancer, a best practice is to assign the main external URL to port 80 and port 443.

2.4.4.1 Purpose of the Hardware Load Balancer

The hardware load balancer balances the load on the Web tier by receiving requests to a virtual host name and then routing each request to one of the Web server instances, based on a load balancing algorithm.

In this way, the load balancer ensures that no one Web server is overloaded with HTTP requests.

For more information about the purpose of specific virtual host names on the hardware load balancer, see Summary of the Typical Load Balancer Virtual Server Names.

Note that in the reference topology, only HTTP requests are routed from the hardware load balancer to the Web tier. Secure Socket Layer (SSL) requests are terminated at the load balancer.

The load balancer provides high availability by ensuring that if one Web server is unavailable, requests will be routed to the remaining Web servers that are up and running.

Further, in a typical highly available configuration, the hardware load balancers are configured such that a hot standby device is ready to resume service in case a failure occurs in the main load balancing appliance. This is important because for many types of services and systems, the hardware load balancer becomes the unique point of access to make invocations and, as a result, becomes a single point of failure (SPOF) for the whole system if it is not protected.

2.4.4.2 Specific internal-only communications between the components of the Application tier

Communications between the Oracle Fusion Middleware components and applications on the application tier are handled by the load balancing capabilities of Oracle Traffic Director (OTD).

The internal-only requests are kept within the Exalogic appliance, using a unique virtual host name that is attached to an OTD failover group.

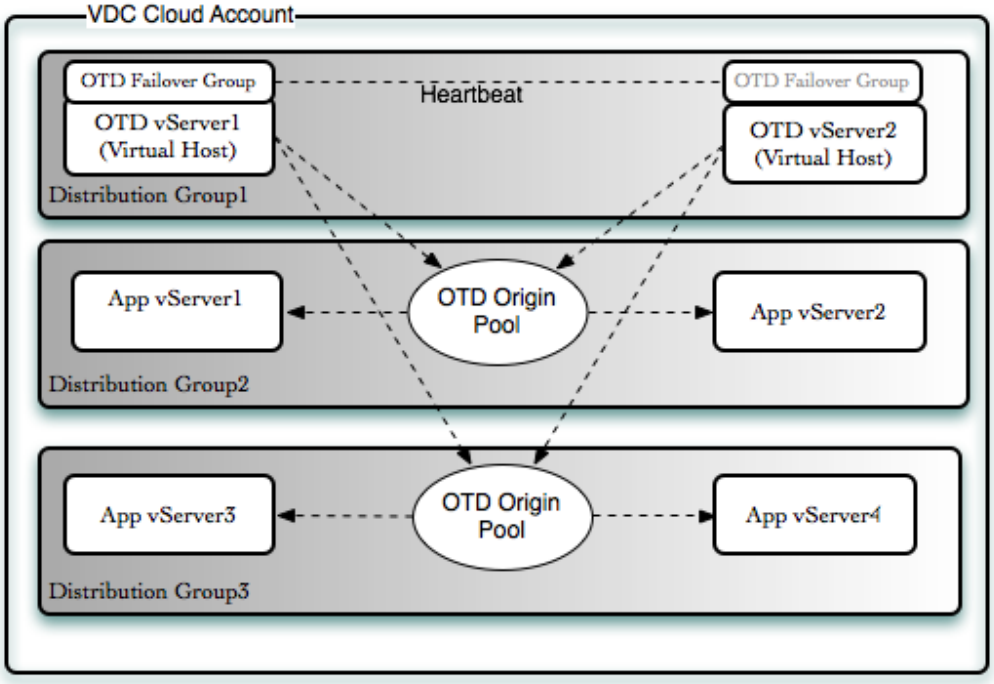

2.4.4.3 Summary of Virtual Server Names

This section describes the virtual server names recognized by the hardware load balancer and Oracle Traffic Director (OTD) in a typical enterprise deployment.

2.4.4.3.1 Summary of the Typical Load Balancer Virtual Server Names

To balance the load on Web hosts and to provide high availability, the hardware load balancer is configured to recognize a set of virtual server names.

As shown in the diagram, the following virtual server names are recognized by the hardware load balancer in this topology:

- *product*`.example.com` - This virtual server name is used for all incoming traffic. Users enter this URL to access the Oracle Fusion Middleware products and custom applications available on this server. The load balancer then routes these requests (using a load balancing algorithm) to one of the servers in the Web tier. In this way, the single virtual server name can be used to route traffic to multiple servers for load balancing and high availability of the Web server instances.
- `admin.example.com` - This virtual server name is for administrators who need to access the Oracle Enterprise Manager Fusion Middleware Control and Oracle WebLogic Server Administration Console interfaces. This URL is known only to internal administrators. It uses the Network Address Translation (NAT) capabilities of the load balancer to route administrators to the active Administration Server in the domain.

These virtual server names are normally defined in your corporate DNS. Intranet facing entry points, such as *product*`.example.com`, are published to the internet. However, management virtual servers, such as `admin.example.com`, are resolvable inside the corporate network only.

For the complete set of virtual server names you must define for your topology, see the chapter in your product enterprise deployment guide that describes the product-specific topology.

2.4.4.3.2 Summary of the Typical OTD Virtual Server Names

To balance the load on internal components and provide high availability, Oracle Traffic Director (OTD) is configured to recognize a set of virtual server names.

As shown in the diagram, the following virtual server names are recognized by OTD in this topology:

- `edginternal.example.com` - This virtual server name is for internal communications only.
  The load balancer uses its Network Address Translation (NAT) capabilities to route any internal communication from the Application tier components that are directed to the `http://edginternal.example.com/` Intranet URL. This URL is not exposed to external customers or users on the Internet.
- `idstore.example.com` – This virtual server name is for internal LDAP communications only.

OTD typically enables these virtual servers on the internal network. They are resolvable only inside the Exalogic appliance using local host entries.

2.4.4.4 HTTPS versus HTTP Requests to the External Virtual Server Name

When you configure the hardware load balancer, a best practice is to assign the main external URL to port 80 and port 443.

Any request on port 80 (non-SSL protocol) should be redirected to port 443 (SSL protocol). Exceptions of this rule include requests from public WSDLs and requests to Oracle Mobile Security Access Server. For more information, see the product-specific Enterprise Deployment Guides.

2.4.5 Understanding the Web Tier

The Web tier of the reference topology consists of two Oracle Traffic Director instances.

This section contains the following topics:

Benefits of Using Oracle Traffic Director Instances to Route Requests
: A Web tier with Oracle Traffic Director (OTD) is not a requirement for many of the Oracle Fusion Middleware products. You can route traffic directly from the hardware load balancer to the WebLogic servers in the Application tier.

Benefits of Using Oracle HTTP Server Instances to Route Requests
: A Web tier with Oracle HTTP Server is not a requirement for many of the Oracle Fusion Middleware products. You can route traffic directly from the hardware load balancer to the WebLogic servers in the Application tier.

Alternatives to Using Oracle Traffic Director Server in the Web Tier
: Oracle Traffic Director (OTD) provides a variety of benefits in an enterprise deployment topology, Oracle also supports routing requests to Oracle HTTP Server or directly from the hardware load balancer to the Managed Servers in the middle tier.

About the Placement and Function of Oracle Traffic Director in a Physical Deployment
: Oracle Traffic Director (OTD) is shown as being deployed onto the same compute nodes as the application. This allows the OTD instances to route specific application traffic to the appropriate application tier and makes the most of the computing power available in the compute node itself.

About the Placement and Function of Oracle Traffic Director in a Virtual Deployment
: In a virtual deployment it is recommended that Oracle Traffic director be placed into dedicated vServers.

About Routing Requests to Oracle Traffic Director
: The load balancer can route requests directly to Oracle Traffic Director (OTD) in the same way that the load balancer would route requests to Oracle HTTP Server. If Oracle Traffic Director is used, you are depending on load balancer monitoring to determine if an Oracle Traffic Director instance is non-responsive.

2.4.5.1 Benefits of Using Oracle Traffic Director Instances to Route Requests

A Web tier with Oracle Traffic Director (OTD) is not a requirement for many of the Oracle Fusion Middleware products. You can route traffic directly from the hardware load balancer to the WebLogic servers in the Application tier.

However, a Web tier does provide several advantages, which is why it is recommended as part of the reference topology:

- If a load balancer routes directly to the WebLogic Server, requests move from the load balancer to the application tier in one single HTTP jump, which can cause security concerns. In the reference topology above, communication with the application tier is confined to the internal Exalogic network to which the load balancer does not have access. Using OTD provides an interface to the application tier with the added advantage of not exposing that tier to the corporate network.
- The Web tier allows the WebLogic Server cluster membership to be reconfigured (new servers added, others removed) without having to change the Web server configuration (as long as at least some of the servers in the configured list remain alive).
- Oracle Traffic Director provides HTTP redirection over and above what WebLogic Server provides. You can use Oracle HTTP Server as a front end against many different WebLogic Server clusters, and in some cases, control the routing via content based routing.
- Oracle Traffic Director provides extra security by acting as an internal load balancer for internal requests.
- Oracle Traffic Director provides the ability to integrate single sign-on capabilities into your enterprise deployment. For example, you can later implement single sign-on for the enterprise deployment, using Oracle Access Manager, which is a component of Oracle Identity and Access Management.

2.4.5.2 Benefits of Using Oracle HTTP Server Instances to Route Requests

A Web tier with Oracle HTTP Server is not a requirement for many of the Oracle Fusion Middleware products. You can route traffic directly from the hardware load balancer to the WebLogic servers in the Application tier.

However, Oracle HTTP Server does provide the following advantages:

- Oracle HTTP Server can be placed on commodity hardware outside of the Exalogic host with a firewall, which adds extra security between the Web server and the Application servers.
- Oracle HTTP Server can be used to host static content.

2.4.5.3 Alternatives to Using Oracle Traffic Director Server in the Web Tier

Oracle Traffic Director (OTD) provides a variety of benefits in an enterprise deployment topology, Oracle also supports routing requests to Oracle HTTP Server or directly from the hardware load balancer to the Managed Servers in the middle tier.

Advantages of Direct Access are:

- Lower configuration and processing overhead than using a front-end Oracle HTTP Server Web tier front-end.

- Monitoring at the application level since the load balancer can be configured to monitor specific URLs for each Managed Server (something that is not possible with OTD).
- You can potentially use this load balancer feature to monitor SOA composite application URLs. Note that this enables routing to the Managed Servers only when all composites are deployed, and you must use the appropriate monitoring software.

Both of the approaches require exposing the internal components to the EoIB network. This would allow the provision of a firewall between the web tier and Exalogic tier. The downside to this approach is that all of the components inside the Exalogic appliance will need to be made visible on the external network.

2.4.5.4 About the Placement and Function of Oracle Traffic Director in a Physical Deployment

Oracle Traffic Director (OTD) is shown as being deployed onto the same compute nodes as the application. This allows the OTD instances to route specific application traffic to the appropriate application tier and makes the most of the computing power available in the compute node itself.

An alternative approach to this is to deploy OTD onto dedicated compute nodes to service the entire Exalogic appliance.

For example, if you have three applications deployed, such as Oracle Identity Management, Oracle SOA Suite, and Oracle WebCenter, you can choose to use common OTD instances to route requests to all of the deployed applications. This is a supported configuration. The individual product enterprise deployment guides use a dedicated OTD deployment for simplicity.

2.4.5.5 About the Placement and Function of Oracle Traffic Director in a Virtual Deployment

In a virtual deployment it is recommended that Oracle Traffic director be placed into dedicated vServers.

In this deployment you can dedicate an OTD cluster to a specific application deployment, this makes management easier. By creating separate distribution groups you can ensure that no 2 OTD instances run on the same underlying compute node.

2.4.5.6 About Routing Requests to Oracle Traffic Director

The load balancer can route requests directly to Oracle Traffic Director (OTD) in the same way that the load balancer would route requests to Oracle HTTP Server. If Oracle Traffic Director is used, you are depending on load balancer monitoring to determine if an Oracle Traffic Director instance is non-responsive.

Internal testing has shown that a faster method of failover detection is available. This involves creating two external OTD failover groups. This method allows the load balancer to direct requests to these external failover groups rather than to the OTD instances themselves. When an OTD failover group is used, then the internal OTD heartbeat is used to detect OTD instance failures. As soon as a failure is detected, a surviving instance will take over processing requests from the failed instance. This failover detection method is typically faster than relying on the failover detection of the load balancer.

2.4.6 Understanding the Application Tier

The application tier consists of two or more hosts, where Oracle WebLogic Server and the Oracle Fusion Middleware products are installed and configured.

The application tier hosts reside inside the Exalogic appliance.

This section contains the following topics:

About the Configuration of the Administration Server and Managed Servers Domain Directories
: Unlike the Managed Servers in the domain, the Administration Server uses an active-passive high availability configuration.

About Using Unicast for Communications Within the Application Tier
: Oracle recommends the unicast communication protocol for communication between the Managed Servers and hosts within the Oracle WebLogic Server clusters in an enterprise deployment.

Understanding OPSS and Requests to the Authentication and Authorization Stores
: Many of the Oracle Fusion Middleware products and components require an Oracle Platform Security Services (OPSS) security store for authentication providers (an identity store), policies, credentials, keystores, and for audit data.

About Coherence Clusters In a Typical Enterprise Deployment
: The standard Oracle Fusion Middleware enterprise deployment includes a Coherence cluster that contains storage-enabled Managed Coherence Servers.

About WebLogic Replication Channels
: By default, WebLogic Server will configure session replication to use the external network. When deploying on Exalogic, you will get better throughput if you configure extra replication channels that utilize the Sockets Direct Protocol (SDP).

2.4.6.1 About the Configuration of the Administration Server and Managed Servers Domain Directories

Unlike the Managed Servers in the domain, the Administration Server uses an active-passive high availability configuration.

This is because only one Administration Server can be running within an Oracle WebLogic Server domain.

In the topology diagrams, the Administration Server on HOST1 is in the active state, and the Administration Server on HOST2 is in the passive (inactive) state.

To support the manual fail over of the Administration Server in the event of a system failure, the typical enterprise deployment topology includes:

- A Virtual IP Address (VIP) for the routing of Administration Server requests
- The configuration of the Administration Server domain directory on the shared ZFS storage device.
- If you have no requirement to access your Administration Server outside of the Exalogic appliance, then the VIP used for the Administration Server should be on

the internal network. If however you wish to use direct external access to the Administration Server for such things as t3 requests or DMS monitoring, then it should be configured on the external network.

In the event of a system failure (for example, a failure of HOST1), you can manually reassign the Administration Server VIP address to another host in the domain, mount the Administration Server domain directory on the new host, and then start the Administration Server on the new host.

However, unlike the Administration Server, there is no benefit to storing the Managed Servers on shared storage. In fact, there is a potential performance impact when Managed Server configuration data is not stored on the local disk of the host computer because of disk contention issues. In an Exalogic deployment, this is achieved by creating a private file system on the ZFS appliance, which is exclusively mounted to each host.

As a result, in the typical enterprise deployment, after you configure the Administration Server domain on shared storage, a copy of the domain configuration is placed on the local storage device of each host computer, and the Managed Servers are started from this copy of the domain configuration. You create this copy using the Oracle WebLogic Server pack and unpack utilities.

The resulting configuration consists of separate domain directories on each host: one for the Administration Server (on shared storage) and one for the Managed Servers (on local storage). Depending upon the action required, you must perform configuration tasks from one domain directory or the other.

For more information about the structure of the Administration Server domain directory and the Managed Server domain directory, as well as the variables used to reference these directories, see Understanding the Enterprise Deployment Directory Structure.

An additional benefit to the multiple domain directory model is that it allows you to isolate the Administration Server from the Managed Servers.

2.4.6.2 About Using Unicast for Communications Within the Application Tier

Oracle recommends the unicast communication protocol for communication between the Managed Servers and hosts within the Oracle WebLogic Server clusters in an enterprise deployment.

Unlike multicast communication, unicast does not require cross-network configuration, and it reduces potential network errors that can occur from multicast address conflicts as well.

When you consider using the multicast or unicast protocol for your own deployment, consider the type of network, the number of members in the cluster, and the reliability requirements for cluster membership. Also, consider the following benefits of each protocol.

Benefits or characteristics of unicast in an enterprise deployment:

- Uses a group leader that every server sends messages directly to. This leader is responsible for retransmitting the message to every other group member and other group leaders, if applicable.
- Works out of the box in most network topologies.
- Requires no additional configuration, regardless of the network topology.

- Uses a single missed heartbeat to remove a server from the cluster membership list.

Benefits or characteristics of multicast in an enterprise deployment:

- Multicast uses a more scalable peer-to-peer model where a server sends each message directly to the network once and the network makes sure that each cluster member receives the message directly from the network.
- Works out of the box in most modern environments where the cluster members are in a single subnet.
- Requires additional configuration in the router(s) and WebLogic Server (that is, Multicast TTL) if the cluster members span more than one subnet.
- Uses three consecutive missed heartbeats to remove a server from the cluster membership list.

Depending on the number of servers in your cluster and on whether the cluster membership is critical for the underlying application (for example, in session-replication intensive applications or clusters with intensive RMI invocations across the cluster), each model may behave better.

Consider whether your topology is going to be part of an Active-Active disaster recovery system or if the cluster is going to traverse multiple subnets. In general, unicast will behave better in those cases.

Most Oracle Fusion Middleware components allow you to use either form of communication. Some, however, are restricted to one form or the other. Refer to the specific product guides as necessary.

2.4.6.3 Understanding OPSS and Requests to the Authentication and Authorization Stores

Many of the Oracle Fusion Middleware products and components require an Oracle Platform Security Services (OPSS) security store for authentication providers (an identity store), policies, credentials, keystores, and for audit data.

As a result, communications must be enabled so the Application tier can send requests to and from the security providers.

For authentication, this communication is to an LDAP directory, such as Oracle Internet Directory (OID) or Oracle Unified Directory (OUD), which typically communicates over port 1389 or 1636. When you configure an Oracle Fusion Middleware domain, the domain is configured by default to use the WebLogic authentication provider. However, for an enterprise deployment, you must use a dedicated, centralized LDAP-compliant authentication provider.

In an Exalogic deployment, the LDAP directory is inside the appliance, and communication with the LDAP servers is typically through the internal Exalogic network.

For authorization (and the policy store), the location of the security store varies, depending up on the tier:

- For the application tier, the authorization store is database-based, so frequent connections from the Oracle WebLogic Server Managed Servers to the database are required for the purpose of retrieving the required OPSS data.
- For the Web tier, the authorization store is file-based, so connections to the database are not required.

For more information about OPSS security stores, see the following sections of *Securing Applications with Oracle Platform Security Services*:

- "Authentication Basics"
- "The OPSS Policy Model"

2.4.6.4 About Coherence Clusters In a Typical Enterprise Deployment

The standard Oracle Fusion Middleware enterprise deployment includes a Coherence cluster that contains storage-enabled Managed Coherence Servers.

This configuration is a good starting point for using Coherence, but depending upon your specific requirements, you can consider tuning and reconfiguring Coherence to improve performance in a production environment or to resolve possible port conflicts.

When reviewing port assignments, note that the Oracle Fusion Middleware products and components default to a Well Known Address (WKA) list that uses the port specified on the Coherence Clusters screen of the Configuration Wizard. The WKA list also uses the listen address of all servers that participate in the Coherence cluster as the listen address for the WKA list.

When configuring Coherence clusters on Exalogic, the coherence clusters are usually configured to use the internal network.

These settings can be customized using the WebLogic Server Administration Console.

For more information, refer to the following resources:

- For information about Coherence clusters, see "Configuring and Managing Coherence Clusters" in *Administering Clusters for Oracle WebLogic Serv*er.
- For information about tuning Coherence, see *Administering Oracle Coherence*.
- For information about storing HTTP session data in Coherence, see "Using Coherence*Web with WebLogic Server" in *Administering HTTP Session Management with Oracle Coherence*Web*.
- For more information about creating and deploying Coherence applications, see *Developing Oracle Coherence Applications for Oracle WebLogic Server*.

2.4.6.5 About WebLogic Replication Channels

By default, WebLogic Server will configure session replication to use the external network. When deploying on Exalogic, you will get better throughput if you configure extra replication channels that utilize the Sockets Direct Protocol (SDP).

For more information, refer to Enabling Cluster-Level Session Replication Enhancements.

2.4.7 Understanding the Directory Tier

Oracle Fusion Middleware Products often interact with an LDAP directory. The diagrams above topic depicts how an LDAP directory can be access in such a topology.

The purpose of including this in the diagrams is to show how the networking would be effected by the inclusion of an LDAP directory inside the Exalogic appliance.

This is shown for example purposes only and by no means assumes that every product will require the installation of an LDAP directory or that every product will

interact with an LDAP directory. Refer to the individual product guides to determine whether that product interacts with LDAP. For details on how to set up an LDAP directory inside Exalogic refer to the *Oracle Identity and Access Management Enterprise Deployment* Guide.

If there is an LDAP directory inside the Exalogic Machine. The directory will typically be configured to listen on the internal IPoIB network. Requests to the LDAP instances will be load balanced through the Oracle Traffic Director. If there is an LDAP directory configured outside of the Exalogic machine then communication with the directory will be through the EoIB network and load balanced through a hardware load balancer.

2.4.8 Understanding the Data Tier

In the Data tier, an Oracle RAC database runs on the two hosts (DBHOST1 and DBHOST2). The database contains the schemas required by the Oracle Fusion Middleware components and the Oracle Platform Security Services (OPSS) policy store.

You can define multiple services for the different products and components in an enterprise deployment to isolate and prioritize throughput and performance. In this guide, one database service is used as an example. In an Exalogic deployment, the database can run on hardware accessible through any standard Ethernet network or on an Exadata system directly connected to Exalogic through Infiniband.

Further, you can use other high availability database solutions to protect the database:

- Oracle Data Guard; for more information, see *Oracle Data Guard Concepts and Administration*
- Oracle RAC One Node; for more information, see "Overview of Oracle RAC One Node" in the *Oracle Real Application Clusters Administration and Deployment Guide*

The above solutions provide protection for the database beyond the information provided in this guide, which focuses on using an Oracle RAC Database, given the scalability and availability requirements that typically apply to an enterprise deployment.

For more information about using Oracle Databases in a high availability environment, see "Database Considerations" in the *High Availability Guide*.

2.5 Understanding the vServers

In an Exalogic virtual deployment physical servers are replaced with virtual servers.

When deploying applications on to virtual servers different deployment options can be considered.

- Create smaller virtual servers with a single instance or JVM running per vServer. The advantages of this approach is it is easier to scale an individual JVM as needed.
- Place dependent applications together in larger vServers. For example placing the components belonging to each domain into a larger vServer. The advantage of this approach is it is easier to scale out an entire domains functionality.

When sizing vServers add up the total memory requirements of the JVM's or instances being hosted in the vServer and add an extra 2GB for the vServer overhead.

2.6 Exalogic Enhancements in Oracle Fusion Middleware

Oracle Weblogic Server includes a number of enhancements to make the Oracle Fusion Middleware products take full advantage of the Exalogic infrastructure.

Enabling these enhancements will enable your Fusion Middleware installation to perform better. Complete list of enhancements are available at Tuning WebLogic Server for Exalogic Environments in the *Oracle Fusion Middleware Performance and Tuning for Oracle WebLogic Server Guide*.

Part II

Preparing Exalogic for an Enterprise Deployment

This topic provides information on preparing the Exalogic appliance, network, storage, host operating system and creating Exalogic virtual servers.

3

Preparing the Exalogic Appliance

This topic provides information on post ECU cloud administration tasks, configuring network information service and Exalogic instrumentation tools.

Post ECU Cloud Administration Tasks
: As an Exalogic Administrator, you must set up the system for cloud users.

Configuring Network Information Service (NIS)
: This topic provides a step-by-step illustration of how to setup and configure Exalogic virtual environment to provide an NIS environment comprised of MASTER/SLAVE servers and NIS clients.

Exalogic Instrumentation Tools
: Exalogic instrumentation refers to tools that help collect, inform, diagnose, or automate configuration or transactional data.

3.1 Post ECU Cloud Administration Tasks

As an Exalogic Administrator, you must set up the system for cloud users.

> **Note:** As part of the Exalogic installation process, the installer is required to validate the newly configured system. Therefore, some of the following steps might have been already executed. For example, Exalogic Guest Base Template (EGBT) may have been downloaded, imported and registered in Enterprise Manager Ops Center.

Also see Task Overviews and Basic Concepts in *Oracle Exalogic Elastic Cloud Administrator's Guide*. Refer to the *Oracle Exalogic Elastic Cloud Administrator's Guide* for the following sections:

1. Create Cloud Admin and Users, using the information available in Creating and Managing Users and Roles.
2. Create User Access EoIB Networks, using the information available in Exalogic vDC Management: Basic Tasks.

Getting Started With Cloud Administration 5-1: Consideration for Creating vServers

3. Create vServer Types using the information available in Exalogic vDC Management: Basic Tasks.
4. Create Account (Cloud Resource Quotas/User Assignment/Tenancies) using the information available in Exalogic vDC Management: Advanced Tasks.
5. Download Exalogic Elastic Cloud Software (EECS) 2.0.6 EGBT templates from Oracle Software Delivery Cloud.

6. Import Server Template using the information available in Exalogic vDC Management Using IaaS CLI: Basic Tasks.
7. Create Custom Template using the information available in Creating Server Templates from vServers.
8. Create Private vNets using the information available in Creating Private vNets.
9. Create Distribution Groups using the information available in Creating Distribution Groups.
10. Create vServers using the information available in Creating vServers.

3.2 Configuring Network Information Service (NIS)

This topic provides a step-by-step illustration of how to setup and configure Exalogic virtual environment to provide an NIS environment comprised of MASTER/SLAVE servers and NIS clients.

Exalogic environment utilizes NFS4 and is connected to a centralised user directory, and it is necessary to configure a network Information System (NIS) environment. NIS allows you to create a master list of users which can then be granted access to various servers. It removes the need to create individual user accounts on servers.

For a user to log into a system through a NIS account, the server has to be configured to allow users to use the NIS service. You can refer to Configuring a Host. If there is another Master NIS server and there are several Slave NIS servers then the failure of one does not prevent users to log in. You have the option of creating a NIS service on the Exalogic appliance. If your Exalogic appliance were to be the master and the rack went down you can be prevented from creating new users.

If the environment allows you to contact an external NIS master from slaves located inside the rack, it is recommended to create the NIS master outside the rack and have the VMs in the rack be slaves to it. If this is not possible, when you bring up the rack you will need to disable the NIS service on the ZFS, bring up the IFS appliance, start the NIS VMs, and then enable the NIS service on the ZFS.

If desired LDAP authentication can be used instead of NIS. For the purposes of this guide we will assume that NIS is being used.

Following sections provide the detailed steps for configuring NIS in an Exalogic virtual environment.

Setting Up the NIS Environment
: Follow the detailed steps for configuring NIS in Exalogic environment.

Configuring LDAP Authentication
: This topic provides basic information on LDAP authentication.

Configure NIS client on ZFS Storage Appliance
: Once you have configured your NIS servers you need to get the ZFS storage appliance to use them. This will enable you to use NFS file systems.

Enabling NFS services on ZFS Storage Appliance
: Once you have configured your NIS servers, you need to get the ZFS storage appliance to use them. Follow the detailed steps to ensure that you can use NFS version 4.

3.2.1 Setting Up the NIS Environment

Follow the detailed steps for configuring NIS in Exalogic environment.

This topic provides a step-by-step illustration of how to setup and configure NIS in an Exalogic environment to provide an NIS environment comprised of MASTER/SLAVE servers and NIS CLIENTS.

Prerequisites for NIS configuration

- If you are placing the NIS server into a virtual server then create a LARGE vServer, which has IPoIB-vServer-shared-storage and Client Access EoIB networks attached.
- If you are configuring NIS for an Exalogic physical environment, then the NIS server should be added to the compute nodes directly.
- Verify that you have the yp rpm's installed on NIS vServers/Compute Nodes.
- In case you do not have the required yp rpm's follow the steps for installing required YP packages on NIS VMs:
  - Create directory `/etc/yum.repos.d`
  - Run the command for downloading the yum repository:

    ```
    wget http://public-yum.oracle.com/public-yum-el5.repo
    ```

  - Run `yum install ypserv` command to install YP serv on NIS VMs.

Creating a NIS Master or Slave

If you are placing your NIS server into vServers then create a large vServer which is connected to the Client EoIB network and the Storage IPoIB network. For details refer to Creating a vServer.

If you are creating a master and slave on the same exalogic host then you should create two vServers and assign them to a dedicated distribution group. If you are not using vServers then you should create the Master and Slave NIS server on different compute nodes.

> **Note:** If you are configuring a corporate NIS, then it is recommended that the NIS Master and Slave reside on different hosts, in different Exalogic racks or split between an external host and the external rack. This ensures that the failure of an Exalogic appliance does not impact the organizations ability to use NIS.

1. Create VMs for running NIS Master & Slave, if required.
2. Create a new distribution group called NIS-Group. Refer to section Creating Distribution Groups under *Oracle Exalogic Elastic Cloud Administrator's* Guide for additional information.
3. Edit `/etc/sysconfig/network file` on each host you are installing a NIS server and add NISDOMAIN and domainname entries defining your NIS domain. Here is an example for additional information:

```
NETWORKING=yes

NETWORKING_IPV6=no
HOSTNAME=nis-server-1
NISDOMAIN=example.com
domainname=example.com
```

4. On each NIS host change the primary IP address to be associated with the ZFS storage network. Do this by modifying the `/etc/hosts` file on both vServers.

   Edit the `/etc/hosts` file on the first vServer: Make the `host-shared-storage` network the primary IP for the machine and add the `host-shared-storage` network for the second vServer.

   Edit the `etc host` file on the second server and add the entry for the first node and modify the entry for the `host-shared-storage-network`. Change the hostname of the host using the hostname command to reflect the hostname of the storage network. For example, `hostname nis-server-1`.

5. Edit `/etc/nsswitch.conf` on both vServers.

   Add NIS to the password, shadow and group lines:

```
# To use db, put the "db" in front of "files" for entries you want to be
# looked up first in the databases
#
# Example:
#passwd: db files nisplus nis
#shadow: db files nisplus nis
#group: db files nisplus nis
passwd: files nis
shadow: files nis
group: files nis
:
```

6. Edit `/etc/yp.conf` file on both hosts.

```
$ cat /etc/sysconfig/network
NETWORKING=yes
NETWORKING_IFV6=no
HOSTNAME=vServerl
NISDOMAIN=example.com
domainname=example.com
[oracle@vServerl bin]$ cat /etc/yp.conf
# /etc/yp.conf - ypbind configuration file
# Valid entries are
# domain NISDOMAIN server HOSTNAME
#    Use server HOSTNAME for the domain NISDOMAIN.
# domain NISDOMAIN broadcast
#    Use broadcast on the local net for domain NISDOMAIN
# domain NISDOMAIN slp
#    Query local SLP server for ypserver supporting NISDOMAIN
# ypserver HOSTNAME
#    Use server HOSTNAME for the local domain. The
#    IP-address of server must be listed in /etc/hosts.
# broadcast
#    If no server for the default domain is specified or
#    none of them is rechable, try a broadcast call to
#    find a server.
domain example.com server vServerl.example.com
domain example.com server vServer2.example.com
```

```
ypserver vServerl.example.com
[oracle@vServerl bin]$
```

The entries in this file should match the entries in the `/etc/hosts` that you added.

7. Edit `/var/yp/Makefile` on the master (first host) and change NOPUSH configuration from true to false.

```
# If we have only one server, we don't have to push the maps to the
# slave servers (NOPUSH=true). If you have slave servers, change this
# to "NOPUSH=false" and put all hostnames of your slave servers in the file
# /var/yp/ypservers.
#
# vvvvvv 21-JUN-20012 (APARKMAN) vvvvvvv
# NOPUSH=true (disable default)
NOPUSH=false
# ^^^^^^ 21-JUN-2012 (APARKMAN) ^^^^^^
```

8. Stop NIS related Services on MASTER and SLAVE nodes as shown below:

```
[root@nis-server-1 yp]# service ypserv stop
Stopping YP server services: [ OK ]
[root@nis-server-1 yp]# service ypbind stop
Shutting down NIS services: [ OK ]
[root@nis-server-1 yp]# service yppasswdd stop
Stopping YP passwd service: [ OK ]
[root@nis-server-1 yp]# service ypxfrd stop
Stopping YP map server: [ OK ]
```

9. Start the YPSERV service on the master (first) vServer.

```
[root@nis-server-1 yp]# service ypserv start
Starting YP server services: [ OK ]
```

10. Execute ypinit -m command to identify SLAVE serves to the NIS MASTER:

```
[root@nis-server-1 yp]# /usr/lib64/yp/ypinit -m
At this point, we have to construct a list of the hosts which will run NIS
servers. nis-server-1 is in the list of NIS server hosts. Please continue to add
the names for the other hosts, one per line. When you are done with the
list, type a control D
next host to add: nis-server-1
 next host to add: nis-server-2.example.com
 next host to add: >>>[CTRL-D]<<<
The current list of NIS servers looks like this:

nis-server-1.example.com
nis-server-2.example.com

Is this correct? [y/n: y] y
We need a few minutes to build the databases...
Building /var/yp/example.com/ypservers...
Running /var/yp/Makefile...
gmake[1]: Entering directory `/var/yp/example.com'
Updating passwd.byname...
Updating passwd.byuid...
Updating group.byname...
Updating group.bygid...
Updating hosts.byname...
Updating hosts.byaddr...
Updating rpc.byname...
Updating rpc.bynumber...
```

```
Updating services.byname...
Updating services.byservicename...
Updating netid.byname...
Updating protocols.bynumber...
Updating protocols.byname...
Updating mail.aliases...
gmake[1]: Leaving directory `/var/yp/example.com'
```

From above snippet you can see that nis-server-1 is set up as a NIS master server. Now you can run `ypinit -s nis-server-1` command on all slave servers.

11. Start the remaining NIS services (ypbind, yppasswdd and ypxfrd) on the MASTER as follows:

```
[root@nis-server-1 yp]# service ypbind start
Binding to the NIS domain: [ OK ]
Listening for an NIS domain server..

[root@nis-server-1 yp]# service yppasswdd start
Starting YP passwd service: [ OK ]

[root@nis-server-1 yp]# service ypxfrd start
Starting YP map server: [ OK ]
```

12. On the Slave host or second virtual machine, start the ypserv service:

```
[root@nis-server-2 yp]# service ypserv start
Starting YP server services: [ OK ]
```

13. Review NIS and YP Configuration. Invoke `nisdomainname` and `domainname` commands to confirm the NISDOMAIN has been appropriately set as follows:

```
[root@nis-server-2 yp]# nisdomainname
example.com

[root@nis-server-2 yp]# domainname
example.com
```

14. Run `ypinit -s` command to initialize NIS configuration on the SLAVE Execute `/usr/lib64/yp/ypinit -s` and provide as it's argument the hostname identified within the output when you ran `/usr/lib64/yp/ypinit -m` command on the SLAVE.

Note: If there are warnings, review to see what went wrong, and fix it accordingly. At this point, make sure that `/etc/passwd` and `/etc/group` files have been edited so that when the NIS is activated, the data bases you have just created will be used, instead of the `/etc ASCII` files.

---

15. Execute following commands for auto restart of NIS services during NIS Master and Slave VMs reboot.

```
chkconfig portmap on
chkconfig ypserv on
chkconfig yppasswdd on
chkconfig ypxfrd on
chkconfig ypbind on
chkconfig nscd on
chkconfig rpcidmapd on
```

3.2.2 Configuring LDAP Authentication

This topic provides basic information on LDAP authentication.

If you want to use the LDAP authentication rather than NIS you can refer to Oracle Exalogic Elastic Cloud - Setting Up LDAP Service for NFSv4.

3.2.3 Configure NIS client on ZFS Storage Appliance

Once you have configured your NIS servers you need to get the ZFS storage appliance to use them. This will enable you to use NFS file systems.

1. Login to the ZFS BUI console as the root user, using the following url:

   `https://exalogicsn01-priv:215`

2. Click on **Configuration**NOT_SUPPORTED**Services**.

   A list of available services is displayed.

3. Click on **NIS** service.

   Enter the following information on the page:

   - *Domain:* This is your corporate domain for example, example.com
   - *Servers:* Select the listed servers. Add an entry for each of the NIS servers you are using, both Master and Slaves. You need to provide the IP addresses of those servers.

4. Click **Apply** and click **Enable** if the NIS service is not enabled.

5. Click **Restart Service** for the changes to take effect.

3.2.4 Enabling NFS services on ZFS Storage Appliance

Once you have configured your NIS servers, you need to get the ZFS storage appliance to use them. Follow the detailed steps to ensure that you can use NFS version 4.

1. Login to the ZFS BUI console as the root user, using the following url:

   `https://exalogicsn01-priv:215`

2. Click on **Configuration**NOT_SUPPORTED**Services**.

   A list of available services is displayed.

3. To edit, click on the **NFS** service.

4. On the NFS page enter the following information:

   Domain: This is your corporate domain for example, example.com

   Servers: Select the listed servers.

   Ensure that the Maximum supported version is set to `NFSv$`

   Enter the Custom `NFSv4` identity Domain and click **Apply**.

Note: This should be the same as the NIS domain.

5. Click **Apply** and click **Enable** if the NIS service is not enabled.
6. Click **Restart Service** for the changes to take effect.

3.3 Exalogic Instrumentation Tools

Exalogic instrumentation refers to tools that help collect, inform, diagnose, or automate configuration or transactional data.

A Master Note On Exalogic Instrumentation , of all of the currently available instrumentation summarizes these tools. Exalogic Kinetic Infrastructure Tools (EKIT) is a collection of tools that simplify, automate and standardize various infrastructure lifecycle management activities on an virtualized Exalogic rack running Linux. These tools can be used to create custom scripts that can automate many Exalogic lifecycle tasks. Refer to the master note on Exalogic Kinetic Infrastructure Tools (EKIT) for additional information.

Exalogic Lifecycle (ELLC) Tools automate lifecycle operations for Oracle Exalogic systems. Refer to the master note on Exalogic Lifecycle Toolkit Releases for additional information on Exalogic Lifecycle (ELLC) Tools.

Note: To open a master note, perform the following steps:

- Select My Oracle Support document ID, and press Ctrl + F9. The Attributes dialog opens.
- In the **Attribute Value** field for the **Url** attribute, enter this URL:

  https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=note_id
- Enter the Note ID or keyword in the search field at the top of the screen.
- Click **Set Value**.

4

Preparing the Network

This chapter provides information about preparing your network for an enterprise deployment on Exalogic.

Overview of Exalogic Networking
: This section provides information about Exalogic networking.

Planning Your Network
: Oracle Traffic Director (OTD) is the preferred load balancer for Fusion Middleware on Exalogic. OTD can route network traffic to both internal IPoIB and external EoIB networks for all Fusion Middleware components.

Understanding How Components use a Network
: Components interact with the network using host names.

Reserving the Required IP Addresses for an Enterprise Deployment
: Before you begin installing and configuring an Oracle Fusion Middleware enterprise topology, you must obtain and reserve a set of IP addresses.

Configuring Exalogic Networking for a Physical Environment
: This section describes the networking in a Physical Exalogic deployment.

Configuring Exalogic Networking for a Virtual Environment
: This section describes virtual Exalogic networking.

Verifying Network Connectivity
: After you have configured the network, it is important to ensure that you can use them to communicate.

4.1 Overview of Exalogic Networking

This section provides information about Exalogic networking.

Types of Network
: There are three types of network within an Exalogic appliance.

Network Diagram for Exalogic Machine
: This topic provides the information on network diagram for Exalogic machine.

4.1.1 Types of Network

There are three types of network within an Exalogic appliance.

- IP over Infiniband (IPoIB): This is the internal Infiniband network that connects the internal components of the Exalogic appliance. This network is fast, but it

cannot be connected to the outside world. The benefit of this network is that it can be used to ensure that network traffic is kept private from the outside world. The downside to using this network is that external components cannot directly access application components inside the Exalogic appliance.

- Ethernet Management Network (eth0): This management network is used for connecting to the Exalogic components through the built-in Ethernet network. This network is only used for management operations and should not be used for production deployments. This network is used to login to the Exalogic components to configure them.
- Ethernet over Infiniband (EoIB): This network also uses the Exalogic Infiniband network, but it is possible to connect this network to the standard corporate network. This allows external components to talk directly to components inside Exalogic. This network is always used for communication between your hardware load balancer and Oracle Traffic Director.

4.1.2 Network Diagram for Exalogic Machine

This topic provides the information on network diagram for Exalogic machine.

Figure 4-1 shows the network diagram for an Oracle Exalogic machine.

Figure 4-1 Exalogic Machine Network Overview

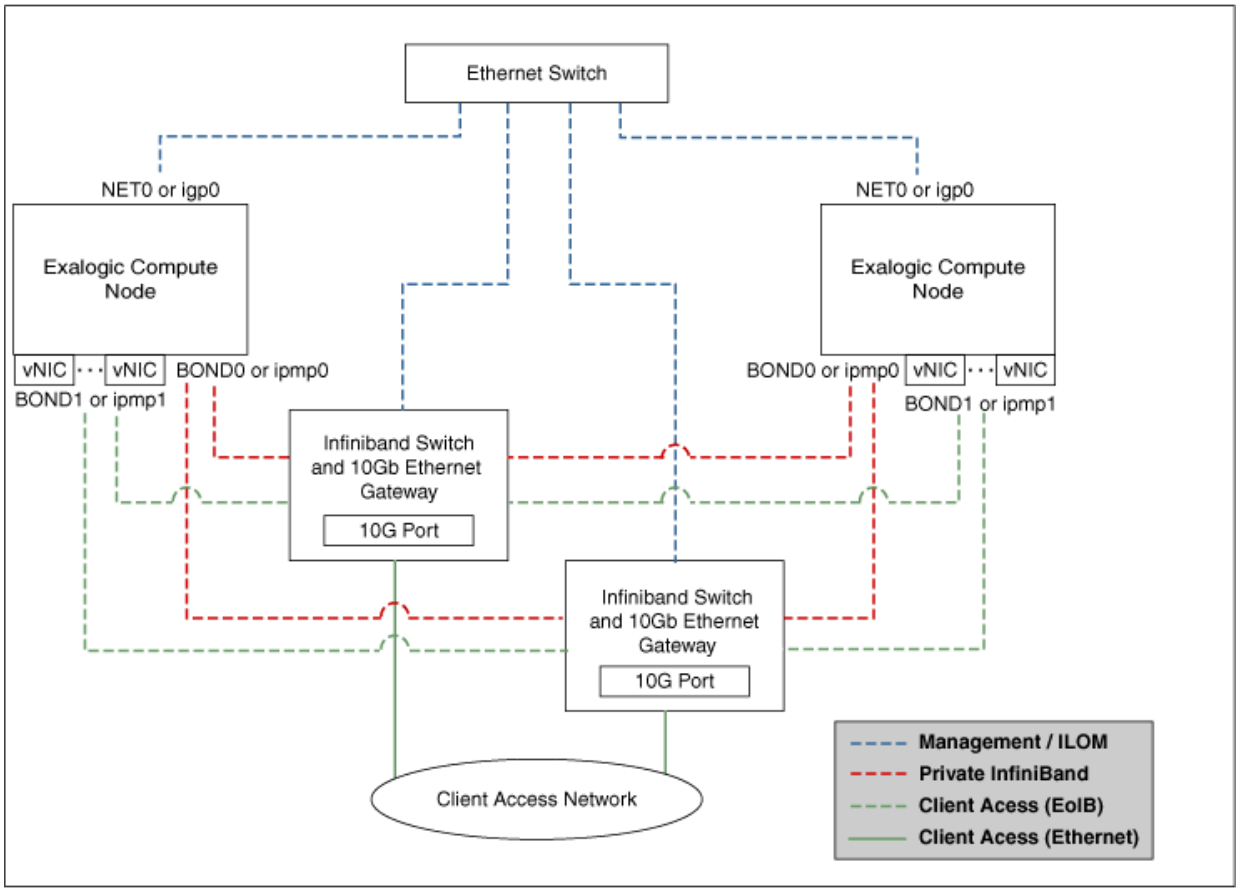

The schematic representation of Oracle Exalogic machine's network connectivity includes the following:

- Default `BOND0` interface, which is the private InfiniBand fabric including the compute nodes connected via Sun Network QDR InfiniBand Gateway Switches

  Typical Uses of this network are:

  – To communicate between compute nodes

  – To access the internal Oracle ZFS Storage Appliance and other Engineered Systems on the fabric

  – To communicate between vServers

  – InfiniBand partitions and memberships provide network isolation and security

  Note:

  InfiniBand `BOND0` interfaces are the default channel of communication among Exalogic compute nodes and storage server head. IP subnets and additional bonds can be added on top of this default bonded interface.

  The device nodes representing the IPoIB network interface for Oracle Linux are referred to as `ib0` and `ib1`. The corresponding logical devices created by Oracle Solaris are referred to as `ibp0` and `ibp1`. The default IPoIB bonded interface `BOND0` or `IPMP0`, configured by the Exalogic Configuration Utility, comprises these Linux-specific interfaces or Solaris-specific interfaces, respectively.

- `BOND1` interface, which is the Ethernet over InfiniBand (EoIB) link

  Typical Uses of this network are:

  – EoIB External Management Network on a vLAN

  – The IP address provided in the ECU spreadsheet created by the ECU configuration process

  – Used for Cloud Administration via Exalogic Control

  – EoIB user access networks on separate vLANs

  – Created by the Exalogic Administrator at post-Exalogic installation

  – Used to access guest vServers and their application services

  Note:

  The device nodes representing the EoIB network interface for Oracle Linux are referred to as `vnic0` and `vnic1`. The Linux kernel creates `eth` device nodes that correspond to the `vnic0` and `vnic1` instances that are created on the Sun Network QDR InfiniBand Gateway Switch.

  The corresponding logical devices created by Oracle Solaris are referred to as `eoib0` and `eoib1`. The EoIB bonded interface `BOND1` or `IPMP1` must be configured manually. When you configure them, choose the network interfaces specific to your operating system.

- `NET0` interface, which is associated with the host Ethernet port 0 IP address for every compute node and storage server head

  Typical Uses of this network are:

  – To access all physical components and ILOMs

  – To perform system administration and life cycle management

  – Used by the Exalogic Control stack

  Note:

  The device node representing the management network interface for Oracle Linux is referred to as `eth0`. The corresponding logical device created by Oracle Solaris is referred to as `igb0`.

- Client access network for external data center connectivity

4.2 Planning Your Network

Oracle Traffic Director (OTD) is the preferred load balancer for Fusion Middleware on Exalogic. OTD can route network traffic to both internal IPoIB and external EoIB networks for all Fusion Middleware components.

OTD provides the common entry point for all traffic. Some components only communicate on a single network.

When you are deciding which Exalogic network to use for these components, consider the following:

- If you will be using an external Web tier, then you should configure your components to use the EoIB network.
- If you expect that all traffic will come through OTD and all traffic will stay within the Exalogic appliance once it reaches there, then you should choose to configure components to use the IPoIB network.
- If you expect all of your LDAP traffic to originate within the Exalogic appliance, then you should configure your LDAP server to use the IPoIB network.
- If your database resides in an Exadata appliance that is connected to the Exalogic appliance via the IB fabric, then you should use the IPoIB network. If the Exadata is connected via Ethernet, then you should use the EoIB network.
- If you are using Oracle Access Manager (OAM), then OAM communications to the proxy port can use both networks (with additional configuration).
- If you plan to access your Administration Servers directly from the corporate network for such things as monitoring, then they should be configured to use the EoIB network. If, however, this is not a requirement, then they can be configured to use the IPoIB network.
- Using the EoIB network is more flexible in that hosts inside the appliance (physical or virtual) can communicate with the corporate network. The downside is it is less secure than the IPoIB network. Traffic accessing components using the IPoIB network must start and terminate in the appliance, which makes man-in-the-middle attacks more difficult.

- Some Oracle Fusion Middleware product Enterprise Deployment Guides may contain specific recommendations on which network to use for various components. Review the appropriate product Enterprise Deployment Guides as part of the planning process.

4.3 Understanding How Components use a Network

Components interact with the network using host names.

For example, if you have three networks for host 1, then you will have three host names as follows:

- `host.example.com` – This will be the default connection to the server via the management network.
- `host-int.example.com` – This will be the host name used to communicate using IPoIB.
- `host-ext.example.com` – This will be the host name used to communicate using the client EoIB network.

Note:

You will also have a host name that is used to connect to the storage appliance (for example, `host-stor.example.com`). This is a private subnet on the IPoIB network that is used to access the ZFS Storage device.

If you are configuring WebLogic Managed Servers to default to the internal network, the listen addresses of the default WebLogic Managed Servers will be set to `host-int.example.com`.

Virtual IP addresses (used in place of the standard IP addresses) will be assigned to an IP address on whichever network they communicate on. The WebLogic Server will use this as the listen address.

This section contains the following topics:

Load Balancers
: In an Exalogic deployment, a hardware load balancer resides outside of the Exalogic machine rack. The load balancer receives external requests for the Oracle Fusion Middleware deployment and sends the requests to the Web hosts.

DMZ
: DMZ is a means of restricting access to components of your infrastructure to those that actually need it.

Firewalls
: This topic provides information on firewalls and its notations.

4.3.1 Load Balancers

In an Exalogic deployment, a hardware load balancer resides outside of the Exalogic machine rack. The load balancer receives external requests for the Oracle Fusion Middleware deployment and sends the requests to the Web hosts.

These Web hosts can either be Oracle HTTP servers or Oracle Traffic Director servers.

The load balancers are configured to receive http and https requests. If a https request is received at the load balancer, the SSL is decrypted at the load balancer and passed on to the Web Servers using the HTTP protocol. This is known as SSL Termination at the load balancer.

The communication from the hardware load balancer to the Web tier is over EoIB.

The load balancer is used to route both application and administrative requests to the Web servers. Administrative requests come from inside the organization intranet. Application requests can be received through the intranet or the internet.

4.3.2 DMZ

DMZ is a means of restricting access to components of your infrastructure to those that actually need it.

In the example mentioned in this guide, there is a public DMZ.

The public zone is where the outside world gains access to your systems. You place into this zone only those components that the outside world must access, such as the Load Balancers and Oracle HTTP Servers (if used in the topology). If users from the outside world attempts to access any servers or services below this zone, they are prevented from doing so by firewalls. The public zone is configured so that the servers in this zone can interact with the application servers in the private zone.

The intranet zone is where you place servers that contain core services, such as databases. These services are very tightly controlled by the organization as they contain the most sensitive data.

By using this approach, you restrict access to information to only those components that require it. This approach is useful where you have users coming in from outside of your organization. If, instead of an extranet, you are setting up an intranet, where all communication is from trusted sources, then you might reasonably decide to do away with the public DMZ.

4.3.3 Firewalls

This topic provides information on firewalls and its notations.

Many of the Oracle Fusion Middleware components and services use ports. As an administrator, you must know the port numbers used by these services, and to ensure that the same port number is not used by two services on a host. Most of the port numbers are assigned during installation.

The Firewall notations are:

- FW0 refers to the outermost firewall.
- FW1 refers to the firewall between the web tier and the application tier.
- FW2 refers to the firewall between the application tier and the data tier.

Table 4-1 Ports Used in the Exalogic Reference Topology

Type	Fire wall	Port and Port Range	Protocol/ Application	Inbound/ Outbound	Other Considerations and Time-out Guidelines
Browser request	FW 0	80	HTTP/Load Balancer	Inbound	Time-out depends on all HTML content and the process models used for the Oracle Fusion Middleware products you are using in the Exalogic environment.
Browser request	FW 0	443	HTTP/Load Balancer	Inbound	Time-out depends on all HTML content and the process models used for the Oracle Fusion Middleware products you are using in the Exalogic environment.
Load balancer to Oracle Traffic Director	n/a	7777	HTTP	n/a	Time-out depends on all HTML content and the process models used for the Oracle Fusion Middleware products you are using in the Exalogic environment.
IAM Access Domain Administration Console access	FW 1	7001	HTTP/ Administration Server and Enterprise Manager	Both	You should tune this time-out based on the type of access to the Administration console (whether it is planned to use the Oracle WebLogic Server Administration Console from application tier clients or clients external to the application tier).
IAM Access Domain Administration Console SSL access	FW 1	7002	HTTP/ Administration Server and Enterprise Manager	Both	You should tune this time-out based on the type of access to the Administration console (whether it is planned to use the Oracle WebLogic Server Administration Console from application tier clients or clients external to the application tier).
IAM Governance Domain Administration Console access	FW 1	7101	HTTP/ Administration Server and Enterprise Manager	Both	You should tune this time-out based on the type of access to the Administration console (whether it is planned to use the Oracle WebLogic Server Administration Console from application tier clients or clients external to the application tier).
IAM Governance Domain Administration Console SSL access	FW 1	7102	HTTP/ Administration Server and Enterprise Manager	Both	You should tune this time-out based on the type of access to the Administration console (whether it is planned to use the Oracle WebLogic Server Administration Console from application tier clients or clients external to the application tier).

Table 4-1 (Cont.) Ports Used in the Exalogic Reference Topology

Type	Fire wall	Port and Port Range	Protocol/ Application	Inbound/ Outbound	Other Considerations and Time-out Guidelines
Coherence	n/a	8088 Range: 8080 - 8090		n/a	n/a
Application tier to data tier (Oracle database or RAC outside of Oracle Exalogic machine via Ethernet)	FW 2	1521		n/a	n/a
Managed Server Access (WLS_OAM1, WLS_OAM2, WLS_OIM1. WLS_OIM2, WLS_SOA1, WLS_SOA2)	FW 1	WLS_OAMn: 14100 WLS_OIMn: 14000 WLS_SOAn: 8001	HHTP	Inbound	Managed Servers, which use bond1 floating IP addresses, are accessed via Oracle HTTP Server. This is only required if the topology has external Oracle HTTP Servers.

4.4 Reserving the Required IP Addresses for an Enterprise Deployment

Before you begin installing and configuring an Oracle Fusion Middleware enterprise topology, you must obtain and reserve a set of IP addresses.

- Physical IP (IP) addresses for each of the host computers you have procured for the topology
- Virtual IP (VIP) addresses for each WebLogic Administration Server in the deployment
- Virtual IP (VIP) addresses for each WebLogic Managed Server which will make use of Server Migration

  For information on what VIPs you need for your deployment, refer to the product specific Enterprise Deployment Guide.
- A unique virtual host name to be mapped to each VIP.

This section contains the following topics:

What Is a Virtual IP (VIP) Address?
: A virtual IP address is an unused IP address that belongs to the same subnet as the host's primary IP address on the network you wish to use.

Why Use Virtual Host Names and Virtual IP Addresses?
: For an enterprise deployment, in particular, it is important that a set of VIPs and the virtual host names to which they are mapped are reserved and enabled on the network.

Host Name Resolution
: It is recommended that host IP addresses on the EoIB network are resolvable within the corporate DNS.

Physical and Virtual IP Addresses Required by the Enterprise Topology
Each product will have different requirements for virtual IP addresses. Refer to the relevant product enterprise deployment guide for product-specific information.

4.4.1 What Is a Virtual IP (VIP) Address?

A virtual IP address is an unused IP address that belongs to the same subnet as the host's primary IP address on the network you wish to use.

It is assigned to a host manually. Individual Managed Servers within the Oracle WebLogic Server domain are configured to listen on this IP address. The Virtual IP Addresses will be assigned to the network you wish to default your components to.

If you are configuring your deployment to use the internal IPoIB network, these Virtual IP Addresses will be in the same subnet as host-int.example.com.

If you are configuring your deployment to use the external EoIB network, these Virtual IP Addresses will be in the same subnet as host-ext.example.com.

4.4.2 Why Use Virtual Host Names and Virtual IP Addresses?

For an enterprise deployment, in particular, it is important that a set of VIPs and the virtual host names to which they are mapped are reserved and enabled on the network.

In an Exalogic deployment, these can be on the corporate or the internal IPoIB network.

In the event of the failure of the host computer where the Virtual IP address is assigned, the Virtual IP address can be assigned to another host in the same subnet so that the new host can take responsibility for running the Managed Servers assigned to it.

The reassignment of virtual IP addresses for Managed Servers can be performed automatically using the Server Migration feature of Oracle WebLogic Server. The reassignment of virtual IP address for the Administration Server must be performed manually.

4.4.3 Host Name Resolution

It is recommended that host IP addresses on the EoIB network are resolvable within the corporate DNS.

For IP addresses on the internal IPoIB network, it is recommended that these are resolved through the file `/etc/hosts`.

If desired, all hosts can be resolvable in DNS or in the `/etc/hosts` file.

4.4.4 Physical and Virtual IP Addresses Required by the Enterprise Topology

Each product will have different requirements for virtual IP addresses. Refer to the relevant product enterprise deployment guide for product-specific information.

The figure below shows how this could look in a generic deployment.

In the diagram:

- IP is the physical IP of the host that the Managed Server is running on and is used by the WebLogic Managed Server.
- VIPx are virtual IP addresses currently enabled on the host that the Managed Server is running on.
- The Administration Server is shown twice. One is active with the VIP assigned, and the other is passive. The passive Administration Server is ready to be started on a different host once VIP1 has been transferred.

Before you begin to install and configure the enterprise deployment, reserve a set of host names and IP addresses that correspond to the VIPs in the product enterprise deployment guide

You can assign any unique host name to the VIPs, but this guide references each VIP using the suggested host names in the guide.

4.5 Configuring Exalogic Networking for a Physical Environment

This section describes the networking in a Physical Exalogic deployment.

Enabling Multicast for bond0
: This topic provides steps to enable multicast for bond0.

Verifying Network Connectivity (HOST1-INT and HOST2-INT)
: After having defined the network, ensure that all of the network names are resolvable from each of the compute Nodes/vServers.

Verifying Multicast Connectivity
: Oracle provides a simple command to test that multicast is configured and working correctly. However, this command is available only after the Oracle software has been installed.

4.5.1 Physical Exalogic Network Map

This topic describes about the physical Exalogic network map.

The following image illustrates the IPoIB and EoIB network interfaces needed for an Oracle Fusion Middleware enterprise deployment. The topics that follow provide a detailed description of the image.

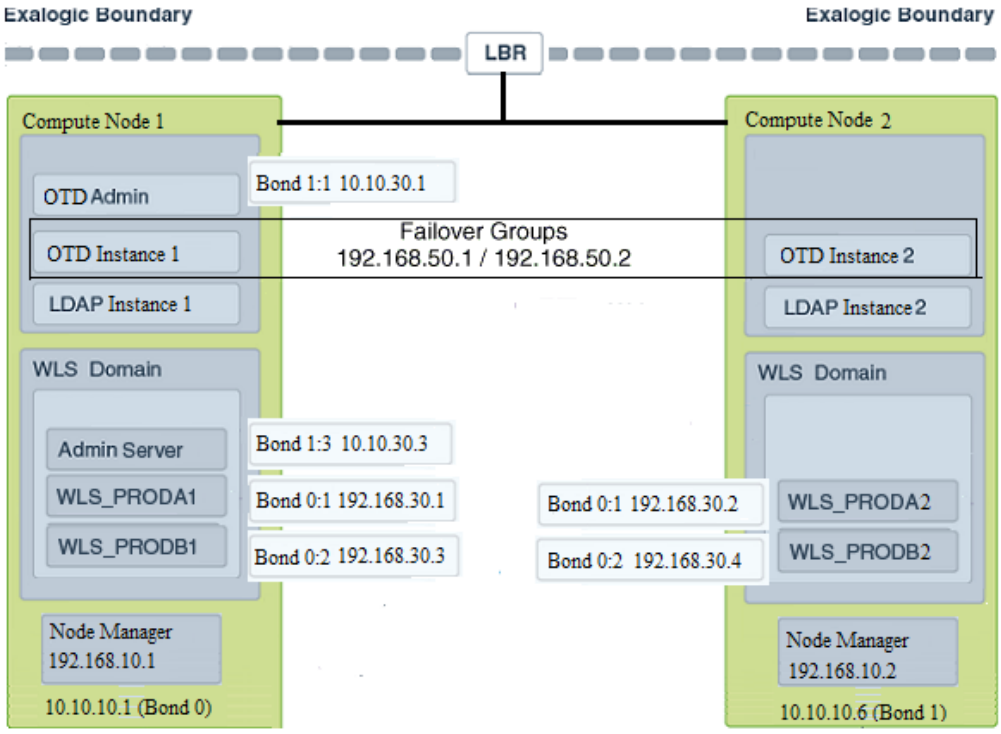

4.5.2 Explanation of the Physical Exalogic Network Interfaces Map

In a physical Exalogic deployment, a hardware load balancer distributes requests to two Oracle Traffic Director instances on the compute nodes in the Exalogic Rack.

All of the Oracle Fusion Middleware and Oracle Traffic Director software is deployed across these two compute nodes in a highly available fashion. The physical Exalogic network map shows how these compute nodes are networked together and to the external corporate network where the load balancer sits.

The diagram above shows how LDAP instances could be configured within the Exalogic appliance. This is important if your application interacts with LDAP. LDAP is shown here for illustrative purposes and may not exist in your deployment.

Note that the diagram of the physical Exalogic network map assumes the following network usage. How you configure your network will depend on your requirements. The example below has the following characteristics:

- The EoIB network is used for communication from the load balancer to Oracle Traffic Director.
- Oracle Traffic Director communicates with the WebLogic Server and LDAP components using the IPoIB network.

This section contains the following topics:

Load Balancer
: An external load balancer sits outside of the Exalogic machine rack.

Physical Network Interface Bonding
: In order to maintain maximum availability, individual network interfaces are bonded together so that the failure of one interface will not affect the availability of the system.

Oracle Traffic Director
: Oracle Traffic Director (OTD) serves several functions within an enterprise deployment on Exalogic.

External Oracle HTTP Servers
: Optionally, you can use Oracle HTTP Servers, which sit on servers outside of the Exalogic machine rack. These servers receive requests from the load balancer and distribute those requests to application virtual servers inside the machine rack using EoIB.

Compute Nodes
: In an Exalogic Physical deployment, the physical compute nodes are used to host services.

4.5.2.1 Load Balancer

An external load balancer sits outside of the Exalogic machine rack.

Its purpose is to receive requests on the public Ethernet network and distribute those requests to the Oracle Traffic Director nodes inside the machine rack using the front end EoIB network or to the external Oracle HTTP Servers.

4.5.2.2 Physical Network Interface Bonding

In order to maintain maximum availability, individual network interfaces are bonded together so that the failure of one interface will not affect the availability of the system.

Note:

The physical Exalogic network map diagram shows a number of bond interfaces. A bond interface is 2 or more physical network interfaces bound together that share a single IP address.

In Physical Exalogic deployments, the following network bonding is assumed:

Network Interface	Network	Purpose
bond0	IPoIB	This is the internal network used for communication between applications and with Exadata (if used).
bond1	EoIB	This is the external client access network

These are the default network interfaces all of which have an associated IP address. Virtual IP addresses can temporarily be added to these network interfaces as required by the deployment. Virtual IP addresses are shown as `:x` where `x` is a sequential number. For example, `bond1:1`.

4.5.2.3 Oracle Traffic Director

Oracle Traffic Director (OTD) serves several functions within an enterprise deployment on Exalogic.

Among these are load balancing, intelligent routing, and SSL termination.

OTD often works in conjunction with external load balancers and external web/HTTP servers.

As a load balancer, Oracle Traffic Director can direct both TCP and HTTP traffic to application components.

Unless External Oracle HTTP Servers are used, then Oracle Traffic Director also functions as an HTTP Server. Oracle Traffic Director listens on the client EoIB network for HTTP requests originating from the external load balancers. If these requests require access to the WebLogic Managed Servers on the compute nodes, then it directs these requests accordingly using the internal IPoIB network.

If you are using the same OTD servers to access multiple applications within the Exalogic Rack then it is better to install the OTD servers onto dedicated compute nodes.

4.5.2.4 External Oracle HTTP Servers

Optionally, you can use Oracle HTTP Servers, which sit on servers outside of the Exalogic machine rack. These servers receive requests from the load balancer and distribute those requests to application virtual servers inside the machine rack using EoIB.

All internal traffic still takes place using IPoIB and Oracle Traffic Director.

External HTTP Servers may be required in organizations that need to place the Web Tier in a separate DMZ than the Exalogic machine or have other applications or policies that require a HTTP Server.

The diagram below shows a typical Network Map where an external Oracle HTTP Server is used and all communication is via the client network.

4.5.2.5 Compute Nodes

In an Exalogic Physical deployment, the physical compute nodes are used to host services.

The following sections describe how a typical enterprise deployment could be configured on physical Exalogic. You might not have all of the components listed, but they are included here for completeness.

The examples below assume you are using the internal IPoIB network for all communication except for the load balancer to Oracle Traffic Director. If you choose to expose everything on the EoIB network, then you must update the sample network paths listed below accordingly.

This section contains the following topics:

ComputeNode1
: ComputeNode1 serves two purposes. It hosts Oracle Traffic Director, which acts as both an internal load balancer and a web server. It also hosts the Oracle Fusion Middleware applications.

ComputeNode2
: ComputeNode2 serves two purposes. It hosts Oracle Traffic Director, which acts as both an internal load balancer and a web server. It also hosts the Oracle Fusion Middleware applications.

4.5.2.5.1 ComputeNode1

ComputeNode1 serves two purposes. It hosts Oracle Traffic Director, which acts as both an internal load balancer and a web server. It also hosts the Oracle Fusion Middleware applications.

- It is configured to use the EoIB client access network. It uses this network to communicate with the external load balancer.
- It is configured to use the IPoIB network for internal communications.
- Oracle Traffic Director can enable an IP address using a failover group to route requests to the LDAP servers using the IPoIB network.
- Oracle Traffic Director acts as a failover node in the event that the IP address used for internal callbacks fails.
- Oracle Traffic Director is used to route application requests to the WebLogic Managed Servers making up the Application tier.
- Node Manager, which is used to start and stop the WebLogic Managed Servers, is configured to accept requests on the internal IPoIB interface.
- This node hosts virtual (floating) IP addresses which are configured on the client access network. These virtual IP addresses are used by the Administration Servers. Although, it is not necessary to use the client access network. The benefit of doing so is that it is possible to monitor the Administration Servers outside of the Exalogic machine.
- Two virtual (floating) IP addresses are attached to the IPoIB interface, which are used by the Managed Servers to facilitate server migration.
- LDAP listens for requests on the internal IPoIB network.

4.5.2.5.2 ComputeNode2

ComputeNode2 serves two purposes. It hosts Oracle Traffic Director, which acts as both an internal load balancer and a web server. It also hosts the Oracle Fusion Middleware applications.

- It is configured to use the EoIB client access network. It uses this network to communicate with the external load balancer.
- It is configured to use the IPoIB network for internal communications.
- Oracle Traffic Director can enable an IP address using a failover group to route internal callback requests using the IPoIB network.
- Oracle Traffic Director acts as a failover node in the event that the IP address used for the LDAP directory fails.
- Oracle Traffic Director is used to route application requests to the WebLogic Managed Servers making up the Application tier.
- Node Manager, which is used to start and stop the local WebLogic Managed Servers, is configured to accept requests on the internal IPoIB interface.
- This node hosts virtual (floating) IP addresses which are configured on the client access network. These virtual IP addresses are used by the Administration servers. Although it is not necessary to use the client access network, the benefit of doing so is that it is possible to monitor the Administration Servers outside of the Exalogic machine.
- Two virtual (floating) IP addresses are attached to the IPoIB interface, which are used by the Managed Servers to facilitate server migration.
- The LDAP directory listens for requests on the internal IPoIB network.

4.5.3 Host Name and Networking Requirements

Networking is a complicated but critical part of any Exalogic deployment.

This guide shows the typical enterprise deployment topology setup that is using the network described in Physical Exalogic Network Map. This guide utilizes the IPoIB network for internal communications and the EoIB network for external communications.

Table 4-2 is a summary of the required networking setup in the Exalogic physical machine rack. The following sections describe in detail how to set up this networking.

A column has been added to the table to allow you to add your own values for easier cross referencing.

Appropriate host name resolution is critical to topology designs that can sustain network changes, system relocation and disaster recovery scenarios. It is important that the required DNS (either`/etc/hosts` or central DNS server) definitions are in place and that WebLogic Servers use host names and virtual host names instead of using IP addresses and virtual IP addresses directly. Additionally, the Exalogic enterprise deployment requires a set of virtual server names for routing requests to the proper server or service within the topology through the external load balancer and the Oracle Traffic Director servers.

These virtual server names must be resolvable in the corporate network. IPoIB addresses must be resolved only inside the rack's name resolution system. If multiple racks are going to be connected, to elude possible IP conflict, it is good practice to place these also in a central DNS server. Network administrators at the corporate level should enable this. Alternatively host names may be resolved through appropriate `/etc/hosts` file propagated through the different nodes. Table 4-2 provides an example of names for the different floating IP addresses used by servers in the typical enterprise deployment.

Table 4-2 Exalogic Physical IP Addresses Worksheet

Hostname Example for This Guide	Interface	IP Address/ Subnet	Customer Value	Type	Host	Bound By	Details
HOST1	bond0	192.168.10.1 / 255.255.224.0		IPoIB/ Fixed	ComputNode1/ HOST1	NA	Access to HOST1 using the internal IPoIB network.
HOST2	bond0	192.168.10.2 / 255.255.224.0		IPoIB/ Fixed	ComputNode2/ HOST2	NA	Access to HOST2 using the internal IPoIB network.
HOST1VHN1	bond0:1	192.168.30.1 / 255.255.240.0		IPoIB/ Floating	ComputNode1/ HOST1	WLS_PRODA1 Default Channel	Initially enabled in HOST1 can be failed over by server migration to HOST2.
HOST2VHN1	bond0:1	192.168.30.2 / 255.255.240.0		IPoIB/ Floating	ComputNode2/ HOST2	WLS_PRODA2 Default Channel	Initially enabled in HOST2 can be failed over by server migration to HOST1.
HOST1VHN2	bond0:2	192.168.30.3 / 255.255.240.0		IPoIB/ Floating	ComputNode1/ HOST1	WLS_PRODB1 default channel	Initially enabled in HOST1 can be failed over by server migration to HOST2.
HOST2VHN2	bond0:2	192.168.30.4 / 255.255.240.0		IPoIB/ Floating	ComputNode2/ HOST2	WLS_PRODB2 default channel	Initially enabled in ComputeNode3 can be failed over by server migration to HOST1.

Table 4-2 (Cont.) Exalogic Physical IP Addresses Worksheet

Hostname Example for This Guide	Interface	IP Address/ Subnet	Customer Value	Type	Host	Bound By	Details
HOST1EXT	bond1	10.10.10.1/255.255.240.0		EoIB/Fixed	ComputNode1/ HOST1	NA	A fixed IP allowing the compute node to be accessed by an External Load balancer
HOST2EXT	bond1	10.10.10.2/255.255.240.0		EoIB/Fixed	ComputNode2/ HOST2	NA	A fixed IP allowing the compute node to be accessed by an External Load balancer
OTDADMINVHN	bond1:1	10.10.30.1/255.255.224.0		EoIB / Floating	ComputNode1/ HOST1	OTD Administration Server	A floating IP address for the Administration Server is recommended, if you want to manually migrate the OTD Administration Server from HOST1 to HOST2.
ADMINVHN	bond1:2	10.10.30.2/255.255.224.0		EoIB / Floating	ComputNode1/ HOST1	WebLogic Domain Administration Server	A floating IP address for the Administration Server is recommended, if you want to manually migrate the Administration Server from HOST1 to HOST2.

Table 4-2 (Cont.) Exalogic Physical IP Addresses Worksheet

Hostname Example for This Guide	Interface	IP Address/ Subnet	Customer Value	Type	Host	Bound By	Details
WEBHOST 1VHN	OTD	10.10.50.1/2 55.255.224.0		EoIB / Floating	ComputNo de1/ HOST1	OTD - HOST1	A floating IP Address managed by OTD. This is the IP Address to which load balancers will connect.
WEBHOST 2VHN	OTD	10.10.50.2/2 55.255.224.0		EoIB / Floating	ComputNo de2/ HOST2	OTD - HOST2	A floating IP Address managed by OTD. This is the IP Address to which load balancers will connect. This is optional
EDGINTER NAL	OTD	192.168.50.1 / 255.255.224. 0		IPoIB/ Floating	ComputNo de1/ HOST1	NA	Oracle Traffic Director failover group.
IDSTORE	OTD	192.168.50.2 / 255.255.224. 0		IPoIB/ Floating	ComputNo de2/ HOST2	NA	Oracle Traffic Director failover group for LDAP Directory

Note: In Table 4-2, where the interface is shown as OTD, means that the IP address is managed by Oracle Traffic Director rather than assigned to a network interface manually. The entries are included in this table for completeness. IDSTORE is shown above for illustrative purposes and may not exist in all topologies.

4.5.4 Additional Requirements for External OHS

This topic provides information on additional requirements for external OHS.

If external Oracle HTTP Servers (OHS) are being used, then the additional host names in the following table apply to an Exalogic Physical configuration. Refer to Exalogic Deployment with an External Web Tier for additional information.

Table 4-3 Exalogic Physical OHS Configuration Worksheet

Hostname Example for This Guide	Interface	IP Address/ Subnet	Customer Value	Type	Host	Bound By	Details
OHSHOST1	eth0	201.19.23.10 / 255.255.255.0		ETH0/ Fixed	External OHSHOST1	Oracle HTTP Server	Fixed IP that Oracle HTTP Server Listens on
HOST1VHN1-EXT	bond1:1	10.10.10.7/255.255.224.0		EoIB/ Floating	ComputeNode1/ HOST1	WLS_PRODA1 Default External Channel	Initially enabled on Compute Node 1, can be failed over to Compute Node 2
HOST1VHN2-EXT	bond1:2	10.10.10.9/255.255.224.0		EoIB/ Floating	ComputeNode1/ HOST1	WLS_PRODB1 Default External Channel	Initially enabled on Compute Node 1
OHSHOST2	eth0	201.19.23.11 / 255.255.255.0		ETH0/ Fixed	External OHSHOST2	Oracle HTTP Server	Fixed IP that Oracle HTTP Server Listens on
HOST2VHN2-EXT	bond1:1	10.10.10.8/255.255.224.0		EoIB/ Floating	ComputeNode2/ HOST2	WLS_PRODA2 Default External Channel	Initially enabled on Compute Node 2, can be failed over to Compute Node 1
HOST2VHN2-EXT	bond1:2	10.10.10.10/ 255.255.224.0		EoIB/ Floating	ComputeNode2/ HOST2	WLS_PRODB2 Default External Channel	Initially enabled on Compute Node 2

Note:

HOSTxVHN-EXT is used in the external Oracle HTTP Server topology instead of the standard HOSTxVHN entries used in the standard topologies. The -EXT is used to show that in an Oracle HTTP Server topology, the HOSTxVHN is bound to the client access network, rather than the internal network as used in the other topologies.

4.5.5 Preparing the Network on Physical Exalogic

By default, compute nodes are not able to communicate outside of the Exalogic machine rack. To do this, you must configure the EoIB network for those hosts that are accessed via external hosts or load balancers.

The compute nodes that require this access are HOST1 and HOST2, which interact with an external load balancer, external database, or Oracle HTTP Server access.

This section contains the following topics:

Summary of the IP Addresses for the EoIB Network Interfaces
: This topic provides a summary of IP addresses you must associate with each EoIB interface on each compute node.

Step 1 - Gather Information
: The following section describes how to gather the information required to create the VLAN and VNICs.

Step 2 - Create a Virtual LAN
: Create a virtual LAN (Local Area Network) on each of the switches.

Step 3 - Create Virtual Network Cards
: Create a virtual network card on the switch to allow compute nodes to recognize it as a network card it can use for communication.

Step 4 - Configure Compute Node Networking and Assign Physical IP Address
: Define a new bonded network interface on the compute node that exposes the VNICs you just created as a single interface to applications, so that the compute node can be accessed using a fixed IP by an external load balancer.

4.5.5.1 Summary of the IP Addresses for the EoIB Network Interfaces

This topic provides a summary of IP addresses you must associate with each EoIB interface on each compute node.

Table 4-4 lists the IP addresses you must associate with each EoIB interface on each compute node. Each of these interfaces is shown in Physical Exalogic Network Map.

Table 4-4 IP Addresses for the EoIB Network and Associated Interfaces

Compute Node	Host Name	Interface Name	External IP Address	Netmask	Used by
ComputeNode1	HOST1–EXT	bond1	10.10.10.1	255.255.224.0	Compute node for external load balancer access

Table 4-4 (Cont.) IP Addresses for the EoIB Network and Associated Interfaces

Compute Node	Host Name	Interface Name	External IP Address	Netmask	Used by
ComputeNode2	HOST2–EXT	bond1	10.10.10.2	255.255.224.0	Compute Node for external load balancer access

Configuring the EoIB network is a multi-stage process:

- Stage 1 - Determine the information required to create the network devices.
- Stage 2 - Create a Virtual LAN (VLAN) on the InfiniBand gateway switches for the compute nodes to communicate.
- Stage 3 - Create Virtual Network Cards (VNIC) on the InfiniBand gateway switches which can be seen by the compute nodes, allowing the compute nodes to utilize the EoIB network.
- Stage 4 - Configure the compute nodes to communicate using the VNICS by assigning IP addresses to them.

4.5.5.2 Step 1 - Gather Information

The following section describes how to gather the information required to create the VLAN and VNICs.

To make things easier, complete the following worksheet as you are progressing:

Table 4-5 VNIC Worksheet

Compute Node	Administrative / External IP Address	Base Lid	GUID	Switch Lid	Switch Name	Connector	Switch GUID	MAC Address
HOST1								
HOST2								

Each compute node is connected to gateway switches, the switches that the compute nodes use must have a VLAN created on them.

Note:

Administrative IP is the IP Address of the compute node as configured on the management LAN at the time of commissioning.

The External IP address is the static IP address you will assign to the EoIB interface.

To determine which switches are connected to the compute nodes:

1. Login to the compute node you wish to expose using the root user.

   For example:

   ssh root@HOST1

2. Retrieve information about the active links on the InfiniBand framework using the following command:

   ```
   iblinkinfo.pl -R | grep hostname
   ```

   For example:

   ```
   # iblinkinfo.pl -R | grep HOST1

   6515[  ] ==( 4X 10.0 Gbps Active/  LinkUp)==>121 2[  ] "el01cn01 EL-C
   192.168.10.3 HCA-1" (Could be 5.0 Gbps)
   6415[  ] ==( 4X 10.0 Gbps Active/  LinkUp)==>120 1[  ] "el01cn01 EL-C
   192.168.10.3 HCA-1" (Could be 5.0 Gbps)
   ```

   The first column shows the Lid id of each of the gateway switches used. In this example, these are lids 65 and 64. The number after the ==> symbol shows the Infiniband Port Base LID: 120 and 121. Make a note of these in Table 4-5.

3. Using the `ibswitches` command, determine the names of the gateway switches to which the compute node is connected.

   ```
   #ibswitches

   Switch  : 0x002128548042c0a0 ports 36 "SUN IB QDR GW switch el01gw03" enhanced
   port 0 lid 63 lmc 0
   Switch  : 0x002128547f22c0a0 ports 36 "SUN IB QDR GW switch el01gw02" enhanced
   port 0 lid 6 lmc 0
   Switch  : 0x00212856d0a2c0a0 ports 36 "SUN IB QDR GW switch el01gw04" enhanced
   port 0 lid 65 lmc 0
   Switch  : 0x00212856d162c0a0 ports 36 "SUN IB QDR GW switch el01gw05" enhanced
   port 0 lid 64 lmc 0
   ```

   The example output shows that:

   - `lid 64` is associated with gateway switch `el01gw05`.
   - `lid 65` is associated with gateway switch `el01gw04`.

   The GUID of the switch is the last 16 characters value after the : For example, the GUID of Switch `el101gw04` is `00212856d0a2c0a0`.

   These are the gateway switches that must have a VLAN and VNICs defined. Make a note of these values in the Table 4-5.

4. Retrieve information about the InfiniBand configuration using the `ibstat` command.

   ```
   # ibstat

   CA 'mlx4_0'
        CA type: MT26428
        Number of ports: 2
        Firmware version: 2.7.8100
        Hardware version: b0
        Node GUID: 0x0021280001a0a364
        System image GUID: 0x0021280001a0a367
   ```

```
Port 1:
     State: Active
     Physical state: LinkUp
     Rate: 40Base lid: 120LMC: 0
     SM lid: 6
     Capability mask: 0x02510868
     Port GUID:0x0021280001a0a365Link layer: IB
Port 2:
     State: Active
     Physical state: LinkUp
     Rate: 40Base lid: 121LMC: 0
     SM lid: 6
     Capability mask: 0x02510868
     Port GUID:0x0021280001a0a366Link layer: IB
```

The output shows that the compute node is connected to two InfiniBand switches, one for each port. The Base Lid links to the value you obtained in Step 2 above.

The `ibstat` command above shows that this compute node has two ports. Each of these ports is associated with a different switch.

Use the base lid to determine the switch to which each port is connected, by comparing the base lid with the output of the `iblinkinfo` command in Step 2. In our example, port 1 has a base lid of 120, which is associated with the switch with a lid of 64. You can now determine the actual switch name by looking at the output of the command `ibswitches` in step 3. In this example, this would be switch `el101gw05`.

To summarize, on this compute node, Port Number 1, which has a GUID of `0x0021280001a0a365` is connected to the switch `el101gw05` whose lid id is `64`.

Make a note of the last 16 characters of each GUID in Table 4-5.

You now have the information about the existing network.

5. Determine a unique MAC address for each of the VNICs you are going to create.

   The MAC address needs to be derived following the rules for a locally administered address. This can be done using the following calculation:

   - The last three octets of the Switch GUID, plus the last three octets of the Internal IP address in hex. For example, the GUID of the switch `el101gw04` is `00212856d0a2c0a0`. The last three octets are: `a2c0a0`.
   - Change the last byte in the last octet to a value which will reflect a locally administered MAC address. The value can be 2, 6, A, or E. For example, choosing E this would make the three octets look like this: `a2c0ae` (The MAC is not case sensitive).
   - Separate each octet with a colon (:), for example, `a2:c0:ae`.
   - The internal (bond0) IP address of the Compute Node HOST1 is: `192.168.10.1`
   - The last three octets are: `168.10.1`. Converted to Hexadecimal and separated by a colon: `a8:0a:01`

   Therefore, you can derive the MAC address as: `a2:c0:a0:a8:0a:01`.

   Make a note of the MAC address in the worksheet.

6. Determine the switch upload connector.

a. Log in to one of the switches as root.

For example:

ssh root@el101gw05

b. At the command prompt, run the following:

```
listlinkup | grep Bridge

Bridge-0 Port 0A-ETH-1 (Bridge-0-2) up (Enabled)
Bridge-0 Port 0A-ETH-2 (Bridge-0-2) down (Enabled)
Bridge-0 Port 0A-ETH-3 (Bridge-0-1) down (Enabled)
Bridge-0 Port 0A-ETH-4 (Bridge-0-1) down (Enabled)
Bridge-1 Port 1A-ETH-1 (Bridge-1-2) down (Enabled)
Bridge-1 Port 1A-ETH-2 (Bridge-1-2) down (Enabled)
Bridge-1 Port 1A-ETH-3 (Bridge-1-1) down (Enabled)
Bridge-1 Port 1A-ETH-4 (Bridge-1-1) down (Enabled)
```

Identify the uplinks which can be used in the gateway. Any uplink that has a value of up can be used. In the example output, only `0A-ETH-1` is available for use.

Using the examples above, the worksheet entries for HOST1 would look as follows:

Table 4-6 Example Worksheet for HOST1

Compute Node	Administrative / External IP Address	Base Lid	GUID	Switch Lid	Switch Name	Connector	Switch GUID	MAC Address
HOST1	10.168.10.1/10.10.10.1	120	0021280001a0a365	64	el01gw05	0A-ETH-1	00212856d162c0a0	62:C0:A0:A8:0A:01
		121	0021280001a0a366	65	el01gw04	0A-ETH-1	00212856d0a2c0a0	A2:C0:A0:A8:0A:01

7. Log in to the InfiniBand switch where the master Subnet Manager is running.

For more information, refer to the *Oracle Exalogic Elastic Cloud Machine Owner's Guide*.

8. Run the following command to start the configuration process:

```
smpartition start
```

9. Run the following command to create a myEoIB partition with the pkey 0x005 with a full membership:

```
smpartition create -n myEoIB -pkey 0x005 -m full
```

10. Use the following command to add the compute node port GUIDs to the IB partition defined by partition key created previously. Use this partition key while creating the VLAN as described in Step 2 - Create a Virtual LAN.

```
smpartition add -pkey pkey -port compute_node_port_GUID -m both
```

Where `pkey` is the partition key of the IB partition created in this step. Use this pkey when creating the VLAN.

`compute_node_port_GUID` is the GUID value from the worksheet.

11. Run the following command to view the changed partition configuration:

```
smpartition list modified
```

This command displays the new partition with its pkey, ports added to the partition, and membership type.

12. Run the following command to confirm the partition configuration:

```
smpartition commit
```

4.5.5.3 Step 2 - Create a Virtual LAN

Create a virtual LAN (Local Area Network) on each of the switches.

1. Log in to the gateway switch that you stored in the worksheet, for example, `el01g04`, as the user `ilom-admin`.

For example:

```
ssh ilom-admin@el01gw04
```

2. Change to the system management framework by entering the following:

```
cd /SYS/Fabric_Mgmt
```

For example:

```
Oracle(R) Integrated Lights Out Manager

Version ILOM 3.0 r47111

Copyright (c) 2010, Oracle and/or its affiliates. All rights reserved.
-> cd /SYS/Fabric_Mgmt
```

3. Launch a restricted shell by entering the `show` command:

```
show
```

4. Run the following command to associate a connector with the VLAN that will be used:

```
createvlan connector -vlan 0 -pkey default
```

Where:

- `connector` is the name of the switch interface from the worksheet.
- `vlan` is the number of the Virtual LAN.
- `pkey` is the partition key.

5. Verify the virtual LAN is working using the following command:

```
showvlan
```

Expected output:

```
Connector/LAG VLN PKEY
   ----------- ---- ----
     0A-ETH-1   125 ffff
     0A-ETH-1     0 ffff
```

6. Repeat once for each switch in the VNIC worksheet.

4.5.5.4 Step 3 - Create Virtual Network Cards

Create a virtual network card on the switch to allow compute nodes to recognize it as a network card it can use for communication.

You need to create a VNIC for each port on each switch attached to each externally facing compute node. Refer to Table 4-5 for details.

To create a VNIC:

1. Login to the gateway switch you stored in the worksheet, for example, `el01g04` as the user `ilom-admin`.

   For example:

   ```
   ssh ilom-admin@el01gw04
   ```

2. Change to the system management framework by entering the following:

   ```
   cd /SYS/Fabric_Mgmt
   ```

   For example:

   ```
   Version ILOM 3.0 r47111

   Copyright (c) 2010, Oracle and/or its affiliates. All rights reserved.
   -> cd /SYS/Fabric_Mgmt
   ```

3. Launch a restricted shell by entering the `show` command:

   ```
   show
   ```

4. Run the following command to a VNIC:

   ```
   createvnic connector -guid compute_node_port_GUID -mac unique_mac_address -pkey
   default -vlan 0
   ```

   Where:

   - `connector` is the Connector column in the worksheet.
   - `compute_node_port_GUID` is the GUID column in the worksheet.
   - `unique_mac_address` is the MAC Address in the worksheet.
   - `pkey` and `vlan` are the values you used when you created the VLAN in Step 2 - Create a Virtual LAN.

   For example:

   ```
   createvnic 0A-ETH-1 -guid 0021280001a0a366 -mac A2:C0:A0:A8:0A:01 -pkey default -
   vlan 0
   ```

5. Verify that the VNIC has been created properly by running the following command:

```
showvnics
```

Example output:

```
ID  STATE    FLG IOA_GUID               NODE                              IID
MAC               VLN PKEY GW
--- -------- --- ---------------------- --------------------------------- ----
---------------- --- ---- --------
94  UP       N   0021280001EFA4BF       el01cn01EL-C 192.168.10.1         0000
A2:C0:A0:A8:0A:01 0   ffff 0A-ETH-1
```

Look for the MAC address of the card created and verify that its status is shown as up.

6. If the VNIC is up, repeat Steps 1-5 to add a VNIC to the next interface. If the VNIC is in a WAIT state, follow Steps 7-9 to ensure the Port GUID has been added to the Pkey.

7. Log onto the master switch. This can be obtained by issuing the following command at any switch:

   ```
   getmaster
   ```

8. Log back on to the switch from Step 5 and confirm the VNIC is now up:

   ```
   showvnics
   ```

9. Repeat all the steps, starting with Step 1, to create each of the necessary VNCs.

4.5.5.5 Step 4 - Configure Compute Node Networking and Assign Physical IP Address

Define a new bonded network interface on the compute node that exposes the VNICs you just created as a single interface to applications, so that the compute node can be accessed using a fixed IP by an external load balancer.

Each compute node has two Virtual Network Interface Cards created.

To make configuring the network easier you can use the following worksheet:

Table 4-7 VNIC Worksheet

Compute Node	EoIB IP Address	Netmask	Interface	Network Device	MAC Address	EPORT_ID	IOA_PORT	Interface Device Name	Interface File
HOST1									
HOST2									

To configure the network:

1. Log in to the compute node as the root user.

   For example:

   ssh root@HOST1

2. On the compute node, run the following command to display the list of VNICs available:

```
mlx4_vnic_info -i
```

This command returns the details of the virtual network cards. Make a note of the following in the worksheet:

- `Network Device`
- `MAC Address`
- `EPORT_ID`
- The number following the colon (:) of the `IOA_PORT`.

3. Create interface files for the VNICs on the compute node.

   To ensure correct failover behavior, the name of the VNIC interface file and the value of the DEVICE directive in the interface file must not be based on the kernel-assigned `eth`X interface name (`eth4`, `eth5`, and so on). Instead, Oracle recommends that the interface file name and value of the DEVICE directive in the interface file be derived from the `EPORT_ID` and `IOA_PORT` values:

   Note:

   Any other unique naming scheme is also acceptable.

   a. Determine the interface device name using the following convention:

      ```
      ethEPORT_IOA_PORT
      ```

      For example:

      eth331_1

      Make a note of the interface device name in the worksheet.

   b. Determine the interface file name using the following convention:

      ```
      ifcfg-bond1 ethEPORT_IOA_PORT
      ```

      For example:

      ifcfg-eth331_1

      Make a note of the interface file name in the worksheet.

      Using the examples above for HOST1, the worksheet entry would look as follows:

Table 4-8 VNIC Worksheet

Compute Node	EoIB IP Address	Netmask	Interface	Network Device	MAC Address	EPORT_ID	IOA_PORT	Interface Device Name	Interface File
HOST1	10.10.10.1	255.255.224.0	bond1	eth4	A2:C0:A0:A8:0A:03	331	1	eth331_1	ifcfg-eth331_1

Table 4-8 (Cont.) VNIC Worksheet

Compute Node	EoIB IP Address	Netmask	Interface	Network Device	MAC Address	EPORT_ID	IOA_PORT	Interface Device Name	Interface File
				eth5	62:C0:A0:A8:0A:03	331	2	eth331_2	ifcfg-eth331_2

Note:

Table 4-2 shows the interface (Bond name) for various types of communication. The interface for compute nodes to be accessed by external load balancer using a fixed IP is Bond1.

The MAC address is the value of the MAC address generated in the VNICs worksheet.

c. Create the interface file for the first VNIC, `eth4` in the example, by using a text editor, such as VI, and save the file in the following directory:

```
/etc/sysconfig/network-scripts
```

Name the file `ifcfg-eth331_1` (from the worksheet).

This file will have the following contents:

```
DEVICE=eth331_1
BOOTPROTO=none
ONBOOT=yes
HWADDR=a2:c0:a0:a8:0A:03
MASTER=bond1
SLAVE=yes
MTU=1500
```

Where:

- `DEVICE` is the Derived Name in the worksheet.
- `HWADDR` is the Mac Address in the worksheet.

4. Create a second interface file for the remaining network card.

5. Create a bonded Ethernet Card encompassing each of the network devices by creating a file named `ifcfg`-Interface, for example:

```
ifcfg-bond1
```

The file will have the following contents:

```
DEVICE=bond1
IPADDR=10.10.10.1
NETMASK=255.255.224.0
BOOTPROTO=none
USERCTL=no
TYPE=Ethernet
ONBOOT=yes
IPV6INIT=no
BONDING_OPTS="mode=active-backup miimon=100 downdelay=5000 updelay=5000"
```

```
GATEWAY=10.10.18.1
MTU=1500
```

Where:

- `Device` is the Interface Name.
- `IPADDR` is the external IP address being assigned.
- `NETMASK` is the netmask of the IP Address.
- `GATEWAY` is the IP address of your gateway.

6. Restart networking using the following command:

```
service network restart
```

4.5.6 Enabling Virtual IP Addresses

This topic provides information to enable the virtual IP addresses.

Having completed the network chapter, you need to assign virtual IP addresses to the various network interfaces and hosts as described in Enabling Virtual IP Addresses.

4.5.7 Adjust MTU (maximum transmission units) Value for IPoIB Interface bond0

When the Exalogic rack is commissioned, it sets up the networking configuration for the internal IPoIB interface bond0.

Changed the MTU value created to the value 6400 for improved performance.

To change the value:

1. Open the `ifcfg-bond0` file located in the `/etc/sysconfig/network-scripts` directory.
2. Change the value of `MTU` to 64000. The following example is an ifcfg-bond0 file after editing the MTU value:

```
######## DO NOT EDIT THIS FILE ########
######## GENERATED BY EXALOGIC ########
DEVICE=bond0
IPADDR=192.168.47.67
NETMASK=255.255.240.0
BOOTPROTO=none
USERCTL=no
TYPE=Ethernet
ONBOOT=yes
IPV6INIT=no
BONDING_OPTS="mode=active-backup miimon=100 downdelay=5000 updelay=5000"
MTU=64000
```

3. Save the file and restart the networking using the command service network restart.

Note:

For the latest recommended values, see My Oracle Support document ID *1624434.1*Revised MTU Tuning Recommendations for the IPoIB Related Network Interfaces on Exalogic Physical and Virtual Environments.

To open a master note, perform the following steps:

- Select My Oracle Support document ID *1624434.1*, and press Ctrl + F9. The Attributes dialog opens.
- In the **Attribute Value** field for the **Url** attribute, enter this URL:

  https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=1624434.1
- Enter the Note ID *1624434.1* or keyword in the search field at the top of the screen.
- Click **Set Value**.

4.5.8 Enabling Multicast for bond0

This topic provides steps to enable multicast for bond0.

Even though WebLogic clusters are configured to use unicast for cluster communications, some products, such as Oracle Identity Manager, have an internal dependency on multicast. If you are planning to use one of these components, you must enable multicast on the IPoIB network (assuming that you are using the IPoIB network for communication).

To enable multicast:

1. Create a file called `route-bond0` in the directory `/etc/sysconfig/network-scripts` with the following contents:.

   ```
   224.0.0.0/4 dev bond0
   ```

2. After creating the file, restart the network by executing the command

   ```
   service network restart
   ```

4.5.9 Verifying Network Connectivity (HOST1-INT and HOST2-INT)

After having defined the network, ensure that all of the network names are resolvable from each of the compute Nodes/vServers.

You do this by performing the following command on each compute node/vServer

```
ping -I interface hostname
```

For example:

```
ping -I bond1 IADADMINVHN
```

Perform this test for each entry in Table 4-2 and Table 4-10 , depending on your deployment type.

```
ping -I bond1 ADMINVHN
ping -I bond1 OTDADMINVHN
ping -I bond1 IADDBSCAN
```

```
ping -I bond0 HOST1
ping -I bond0 HOST1VHN1
ping -I bond0 HOST1VHN2
ping -I bond1 HOST1-EXT
ping -I bond1 DBHOST1
ping -I bond0 HOST2
ping -I bond0 HOST2VHN1
ping -I bond0 HOST2VHN2
ping -I bond1 HOST2-EXT
ping -I bond1 DBHOST2
```

In addition, test that the compute nodes allow access to the database servers by pinging the database hosts and the Gridlink scan address. For example:

```
ping -I bond1 DBHOST1
ping -I bond1 DBHOST2
ping -I bond1 IADDBSCAN
```

4.5.10 Verifying Multicast Connectivity

Oracle provides a simple command to test that multicast is configured and working correctly. However, this command is available only after the Oracle software has been installed.

1. Use the following command to verify multicast after the Oracle Fusion Middleware software has been installed:

   ```
   set JAVA_HOME to JAVA_HOME
   ```

2. Change the directory to the following:

   ```
   ORACLE_HOME/coherence_3.7/bin
   ```

3. Run the following command:

   ```
   multicast-test.sh -ttl 0 -local host1-int
   ```

   Where `host1-int` is the name of the host associated with the network interface `bond0`.

   Run the command on each of the `hosts` to perform a multicast test.

   For details about the `multicast-test.sh` program, see "Performing a Multicast Connectivity Test" in the *Oracle Coherence Administrator's Guide*.

4.6 Configuring Exalogic Networking for a Virtual Environment

This section describes virtual Exalogic networking.

Virtual Exalogic Network Map
: This topic describes about the virtual Exalogic network map.

Explanation of the Virtual Network Interfaces Map
: In a virtual Exalogic deployment, a hardware load balancer is used to distribute requests to two vServers in the Exalogic rack that hosts Oracle Traffic Director.

Host Name and Networking Requirements
: Networking is a complicated but critical part of any Exalogic deployment. This guide utilizes the IPoIB network for internal communications and the EoIB network for external communications.

Preparing the Network on Virtual Exalogic
: This topic provides information on creating a private IPoIB network and reserving virtual IP addresses.

Enabling Virtual IP Addresses
: This topic provides information to enable the virtual IP addresses.

4.6.1 Virtual Exalogic Network Map

This topic describes about the virtual Exalogic network map.

The Figure 4-2 shows the Virtual Exalogic Network map.

Figure 4-2 Virtual Network Map

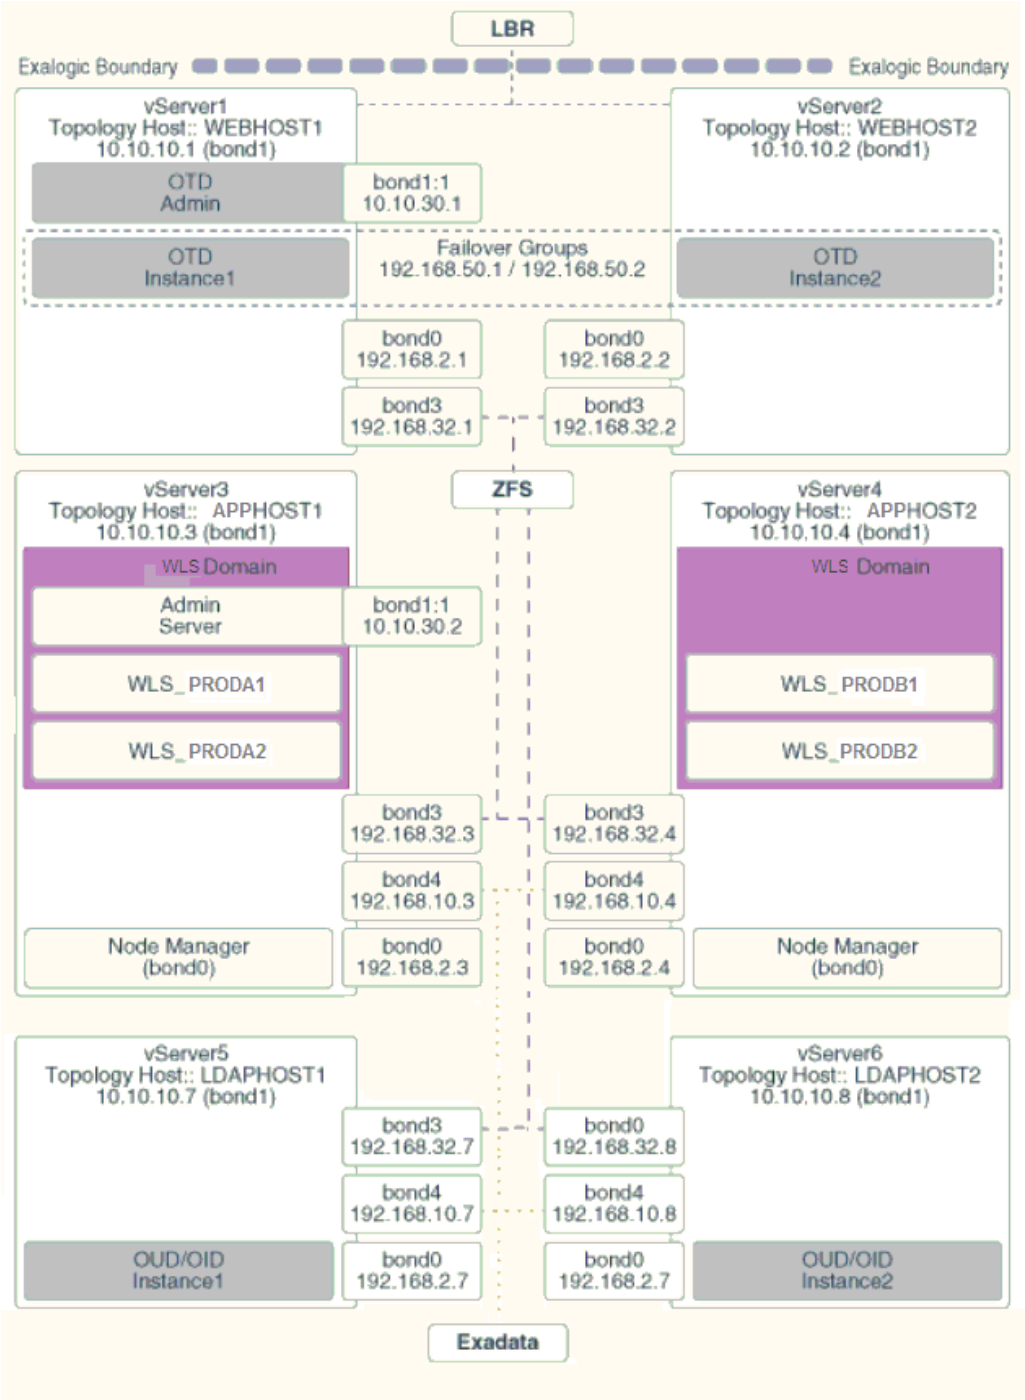

4.6.2 Explanation of the Virtual Network Interfaces Map

In a virtual Exalogic deployment, a hardware load balancer is used to distribute requests to two vServers in the Exalogic rack that hosts Oracle Traffic Director.

The Oracle Traffic Director instances then direct traffic to other vServers, which host the Oracle Fusion Middleware components.

The virtual Exalogic network map diagram shows how these vServers are networked together and to the external corporate network where the load balancer sits.

This section contains the following topics:

Load Balancer
: An external load balancer sits outside of the Exalogic machine rack.

Virtual Network Interface Bonding
: In Virtual Exalogic deployments, these are assigned when the vServer is created. Therefore, the exact bonding is determined by the order the networks are attached to the vServer.

Oracle Traffic Director
: Oracle Traffic Director (OTD) serves several functions within an enterprise deployment on Exalogic.

External Oracle HTTP Servers
: Optionally, you can use Oracle HTTP Servers, which sit on servers outside of the Exalogic machine rack. These servers receive requests from the load balancer and distribute those requests to application virtual servers inside the machine rack using EoIB.

About Virtual Servers
: In an Exalogic Virtual deployment, virtual servers are used instead of physical compute nodes to host services.

4.6.2.1 Load Balancer

An external load balancer sits outside of the Exalogic machine rack.

Its purpose is to receive requests on the public Ethernet network and distribute those requests to the Oracle Traffic Director nodes inside the machine rack using the front end EoIB network or to the external Oracle HTTP Servers.

4.6.2.2 Virtual Network Interface Bonding

In Virtual Exalogic deployments, these are assigned when the vServer is created. Therefore, the exact bonding is determined by the order the networks are attached to the vServer.

For the purposes of this document, assume the network interfaces as shown in Table 4–9:

Table 4-9 Network Interfaces

Purpose	Network	Interface
Management Network	EoIB	eth0
Private Internal Network	IPoIB	bond0
Client Access Network	EoIB	bond1
Internal Administration Network	IPoIB	bond2 (not used)
Internal Storage Network	IPoIB	bond3
Exadata Network	IPoIB	bond4

4.6.2.3 Oracle Traffic Director

Oracle Traffic Director (OTD) serves several functions within an enterprise deployment on Exalogic.

Among these are load balancing, intelligent routing, and SSL termination.

OTD often works in conjunction with external load balancers and external web/HTTP servers.

As a load balancer, Oracle Traffic Director can direct both TCP and HTTP traffic to application components.

Unless External Oracle HTTP Servers are used, then Oracle Traffic Director also functions as an HTTP Server. Oracle Traffic Director listens on the client EoIB network for HTTP requests originating from the external load balancers. If these requests require access to the WebLogic Managed Servers on the compute nodes, then it directs these requests accordingly using the internal IPoIB network.

If you are using the same OTD servers to access multiple applications within the Exalogic Rack then it is better to install the OTD servers onto dedicated compute nodes.

4.6.2.4 External Oracle HTTP Servers

Optionally, you can use Oracle HTTP Servers, which sit on servers outside of the Exalogic machine rack. These servers receive requests from the load balancer and distribute those requests to application virtual servers inside the machine rack using EoIB.

All internal traffic still takes place using IPoIB and Oracle Traffic Director.

External HTTP Servers may be required in organizations that need to place the Web Tier in a separate DMZ than the Exalogic machine or have other applications or policies that require a HTTP Server.

The diagram below shows a typical Network Map where an external Oracle HTTP Server is used and all communication is via the client network.

4.6.2.5 About Virtual Servers

In an Exalogic Virtual deployment, virtual servers are used instead of physical compute nodes to host services.

Virtual hosts are similar to physical servers, but virtual hosts are actually virtual environments managed by Exalogic Control. These are referred to as vServers.

The following sections describe the networking configuration of each of the vServers.

Virtual Server 1 (vServer1)
: vServer1 (WEBHOST1) hosts Oracle Traffic director which acts as both an internal load balancer and a web server.

Virtual Server 2 (vServer2)
: vServer2 (WEBHOST2) serves two purposes. It hosts Oracle Traffic Director, which acts as both an internal load balancer and a web server. It also hosts the Oracle Fusion Middleware applications.

Virtual Server 3 (vServer3)
: vServer3 (HOST1) hosts the Oracle Fusion Middleware applications which comprise the domain.

Virtual Server 4 (vServer4)
: vServer4 (HOST2) hosts the Oracle Fusion Middleware applications which comprise the domain.

Virtual Server 5 (vServer5)
: vServer5 (LDAPHOST1) hosts the Oracle Fusion Middleware applications that comprise the domain.

Virtual Server 6 (vServer6)
: vServer6 (LDAPHOST2) hosts the Oracle Fusion Middleware applications, which comprise the domain.

4.6.2.5.1 Virtual Server 1 (vServer1)

vServer1 (WEBHOST1) hosts Oracle Traffic director which acts as both an internal load balancer and a web server.

- It is configured to use the EoIB client access network. It uses this network to communicate with the external load balancer.
- It is configured to use the IPoIB network for internal communications.
- Oracle Traffic Director acts as a failover node in the event that the IP address used for internal callbacks fails.
- Oracle Traffic Directoryis used to route application requests to the WebLogic Managed Servers making up the Application tier.

4.6.2.5.2 Virtual Server 2 (vServer2)

vServer2 (WEBHOST2) serves two purposes. It hosts Oracle Traffic Director, which acts as both an internal load balancer and a web server. It also hosts the Oracle Fusion Middleware applications.

- It is configured to use the EoIB client access network. It uses this network to communicate with the external load balancer.
- It is configured to use the IPoIB network for internal communications.
- Oracle Traffic Director enables an IP address using a failover group to route internal callback requests to servers using the IPoIB network.
- Oracle Traffic Director is used to route application requests to the WebLogic Managed Servers making up the Application tier.

4.6.2.5.3 Virtual Server 3 (vServer3)

vServer3 (HOST1) hosts the Oracle Fusion Middleware applications which comprise the domain.

- It can be configured to use the EoIB client access network. It uses this network to communicate with external database servers or for external Oracle HTTP servers to communicate with the Weblogic Managed Servers.
- It is configured to use the IPoIB network for internal communications.
- The LDAP instance receives requests from Oracle Traffic Director on the internal IPoIB network.
- Node Manager, which is used to start and stop the WebLogic Managed Servers, is configured to accept requests on the internal IPoIB interface.

- This node hosts a virtual (floating) IP address which is configured on the client access network. This virtual IP address is used by the Administration Server. Although it is not necessary to use the client access network, the benefit of doing so is that it is possible to monitor the Administration Server outside of the Exalogic machine.
- Two virtual (floating) IP addresses attached are used by the Managed Servers to facilitate server migration. The network these VIPs are attached to, is determined by your requirements. See Understanding Oracle Fusion Middleware and Exalogic Networking for additional information. This document assumes you are using the Internal IPoIB network.

4.6.2.5.4 Virtual Server 4 (vServer4)

vServer4 (HOST2) hosts the Oracle Fusion Middleware applications which comprise the domain.

- It can be configured to use the EoIB client access network. It uses this network to communicate with external database servers or for external Oracle HTTP servers to communicate with the Weblogic Managed Servers.
- It is configured to use the IPoIB network for internal communications.
- The LDAP instance receives requests from Oracle Traffic Director on the internal IPoIB network.
- Node Manager, which is used to start and stop the WebLogic Managed Servers, is configured to accept requests on the internal IPoIB interface.
- This node hosts a virtual (floating) IP address which is configured on the client access network. This virtual IP address is used by the Administration Server. Although it is not necessary to use the client access network, the benefit of doing so is that it is possible to monitor the Administration Server outside of the Exalogic machine.
- Two virtual (floating) IP addresses are attached to the IPoIB interface, which are used by the Managed Servers to facilitate server migration.

4.6.2.5.5 Virtual Server 5 (vServer5)

vServer5 (LDAPHOST1) hosts the Oracle Fusion Middleware applications that comprise the domain.

- It can be configured to use the EoIB client access network when external clients need direct access to the LDAP Instance. Normally this takes place via Oracle Traffic Director (OTD). However, if your LDAP Directory is Oracle Internet Directory and the database is on an external host, then EoIB will be required to communicate with that database.
- It is configured to use the IPoIB network for internal communications.
- The LDAP instance receives requests from the Oracle Traffic Director on the internal IPoIB network.
- The Oracle Traffic Director failover group `IDSTORE.example.com` directs requests to the LDAP instance on this and LDAPHOST2.

4.6.2.5.6 Virtual Server 6 (vServer6)

vServer6 (LDAPHOST2) hosts the Oracle Fusion Middleware applications, which comprise the domain.

- Communication with this vServer is via the IPoIB network.
- The Oracle Traffic Director failover group `IDSTORE.example.com` directs requests to the LDAP instance on this and LDAPHOST1.

4.6.3 Host Name and Networking Requirements

Networking is a complicated but critical part of any Exalogic deployment. This guide utilizes the IPoIB network for internal communications and the EoIB network for external communications.

Table 4-10 is a summary of the required networking setup in the Exalogic machine rack. The following sections describe in detail how to set up this networking.

A column has been added to the table to allow you to add your own values for easier cross referencing.

Appropriate host name resolution is critical to topology designs that can sustain network changes, system relocation and disaster recovery scenarios. It is important that the required DNS (either `/etc/hosts` or central DNS server) definitions are in place and that WebLogic Servers use host names and virtual host names instead of using IP addresses and virtual IP addresses directly. Additionally, the Exalogic enterprise deployment requires a set of virtual server names for routing requests to the proper server or service within the topology through the external load balancer and the Oracle Traffic Director servers.

These virtual server names must be resolvable in the corporate network. IPoIB addresses must be resolved only inside the rack's name resolution system. If multiple racks are going to be connected, to elude possible IP conflict, it is good practice to place these also in a central DNS server. Network administrators at the corporate level should enable this. Alternatively host names may be resolved through appropriate `/etc/hosts` file propagated through the different nodes. Table 4-10 provides an example of names for the different floating IP addresses used by servers in a typical enterprise deployment.

Table 4-10 Exalogic IP Addresses Worksheet

Hostname Example for This Guide	Interface	IP Address /Subnet	Customer Value	Type	Host	Bound By	Details
WEBHOST1	bond0	192.168.10.1/255.255.224.0		IPoIB/ Fixed	vServer1/ WEBHOST1	NA	Access to vServer11/ WEBHOST1 using the internal IPoIB network.

Table 4-10 (Cont.) Exalogic IP Addresses Worksheet

Hostname Example for This Guide	Interface	IP Address /Subnet	Customer Value	Type	Host	Bound By	Details
WEBHOST1VHN1	OTD	10.10.50.1/255.255.224.0		EoIB / Floating	vServer1/ WEBHOST1	OTD - WEBHOST1	A floating IP Address managed by OTD. This is the IP Address to which load balancers will connect. This is optional
WEBHOST1-EXT	bond1	10.10.10.1/255.255.240.0		EoIB/ Fixed	Server1/ WEBHOST1	NA	A fixed IP allowing the vServer to be accessed by an External Load balancer
WEBHOST1-STOR	bond3	192.168.32.1/255.255.240.0		IPoIB/ Fixed	vServer1/ WEBHOST1	NA	A fixed IP address allowing the vServer to connect to the ZFS Storage appliance using the internal network.
WEBHOST2	bond0	192.168.10.2/255.255.224.0		IPoIB/ Fixed	vServer2/ WEBHOST2	NA	Access to vServer2/ WEBHOST2 using the internal IPoIB network.
WEBHOST2VHN1	OTD	10.10.50.2/255.255.224.0		EoIB / Floating	vServer2/ WEBHOST2	OTD - WEBHOST2	A floating IP Address managed by OTD. This is the IP Address to which load balancers will connect.
WEBHOST2-EXT	bond1	10.10.10.2/255.255.240.0		EoIB/ Fixed	vServer2/ WEBHOST2	NA	A fixed IP allowing the vServer to be accessed by an External Load balancer

Table 4-10 (Cont.) Exalogic IP Addresses Worksheet

Hostname Example for This Guide	Interface	IP Address /Subnet	Customer Value	Type	Host	Bound By	Details
WEBHOST2-STOR	bond3	192.168.32.2/255.255.240.0		IPoIB/ Fixed	vServer2/ WEBHOST2	NA	A fixed IP address allowing the vServer to connect to the ZFS Storage appliance using the internal network.
HOST1–INT	bond0	192.168.10.5/255.255.224.0		IPoIB/ Fixed	vServer5/ HOST1	NA	Access to vServer5/HOST1 using the internal IPoIB network
HOST1VHN1	bond0:1	192.168.30.5/255.255.240.0		IPoIB/ Floating	vServer5/ HOST1	WLS_PRODA1 Default Channel	Initially enabled in HOST1 and can be failed over by server migration to HOST2
HOST1VHN2	bond0:2	192.168.30.6/255.255.240.0		IPoIB/ Floating	vServer5/ HOST1	WLS_PRODB1 default channel	Initially enabled in HOST1 and can be failed over by server migration to HOST2
HOST1-STOR	bond3	192.168.32.5/255.255.240.0		IPoIB/ Fixed	vServer5/ HOST1	NA	A fixed IP address allowing the vServer to connect to the ZFS Storage appliance using the internal network.
HOST1-DATA	bond4	192.168.10.5/255.255.240.0		IPoIB/ Fixed	vServer5/ HOST1	NA	A fixed IP address allowing the vServer to connect to the Exadata appliance using the default internal network.
HOST2–INT	bond0	192.168.10.6/255.255.224.0		IPoIB/ Fixed	vServer6/ HOST2	NA	Access to vServer6/HOST2 using the internal IPoIB network

Table 4-10 (Cont.) Exalogic IP Addresses Worksheet

Hostname Example for This Guide	Interface	IP Address /Subnet	Customer Value	Type	Host	Bound By	Details
HOST2VHN1	bond0:1	192.168.30.7/255.255.240.0		IPoIB/ Floating	vServer6/ HOST2	WLS_PRODA2 Default Channel	Initially enabled in HOST2 and can be failed over by server migration to HOST1
HOST2VHN2	bond0:2	192.168.30.8/255.255.240.0		IPoIB/ Floating	vServer6/ HOST2	WLS_PRODB2 default channel	Initially enabled in HOST2 and can be failed over by server migration to HOST1.
HOST2-STOR	bond3	192.168.32.6/255.255.240.0		IPoIB/ Fixed	vServer6/ HOST2	NA	A fixed IP address allowing the vServer to connect to the ZFS Storage appliance using the internal network.
HOST2-DATA	bond3	192.168.10.6/255.255.240.0		IPoIB/ Fixed	vServer6/ HOST2	NA	A fixed IP address allowing the vServer to connect to the Exadata appliance using the default internal network.
LDAPHOST1–INT	bond0	192.168.10.7/255.255.224.0		IPoIB/ Fixed	vServer7/ LDAPHOST1	NA	Access to vServer7/ LDAPHOST1 using the internal IPoIB network
LDAPHOST2–INT	bond0	192.168.10.8/255.255.224.0		IPoIB/ Fixed	vServer8/ LDAPHOST2	NA	Access to vServer8/ LDAPHOST2 using the internal IPoIB network

Table 4-10 (Cont.) Exalogic IP Addresses Worksheet

Hostname Example for This Guide	Interface	IP Address /Subnet	Customer Value	Type	Host	Bound By	Details
OTDADMINVHN	bond1:1	10.10.30.1/255.255.224.0		EoIB / Floating	vServer1/ WEBHOST1	OTD Administration Server	A floating IP address for the OTD Administration Server is recommended, if you want to manually migrate the OTD Administration Server from WEBHOST1 to WEBHOST2.
ADMINVHN	bond1:3	10.10.30.3/255.255.224.0		EoIB / Floating	vServer5/ HOST1	Administration Server	A floating IP address for the Administration Server is recommended, if you want to manually migrate the Administration Server from HOST2 to HOST1 If you do not want your Administration Server to communicate with the corporate network, then this address could be on the IPoIB network
EDGINTERNAL	OTD	192.168.50.1/255.255.224.0		IPoIB/ Floating	vServer1/ WEBHOST1	NA	Oracle Traffic Director failover group for internal callbacks
IDSTORE	OTD	192.168.50.2/255.255.224.0		IPoIB/ Floating	vServer2/ WEBHOST2	NA	Oracle Traffic Director failover group for LDAP Directory

Note: The Table 4-10 displays the example network interface names which is used throughout this guide. In an Exalogic virtual environment interface names are assigned when the virtual hosts are created and need not necessarily be the same as those listed in the table.

The IP addresses in this table are examples only. The way IP addresses are assigned in virtual Exalogic, it is highly unlikely that you will be able to use these exact values in your deployment.

This section contains the following topic:

Additional Requirements for External Oracle HTTP Server
If external Oracle HTTP Servers are being used, then the additional host names apply on an Exalogic Virtual configuration.

4.6.3.1 Additional Requirements for External Oracle HTTP Server

If external Oracle HTTP Servers are being used, then the additional host names apply on an Exalogic Virtual configuration.

Table 4-11 Exalogic Virtual OHS Configuration Worksheet

Hostname Example for This Guide	Interface	IP Address/ Subnet	Customer Value	Type	Host	Bound By	Details
OHSHOST1	eth0	201.19.23.10 / 255.255.255.0		ETH0/ Fixed	External OHSHOST1	Oracle HTTP Server	Fixed IP that Oracle HTTP Server Listens on
OHSHOST2	eth0	201.19.23.11 / 255.255.255.0		ETH0/ Fixed	External OHSHOST2	Oracle HTTP Server	Fixed IP that Oracle HTTP Server Listens on
HOST1VHN-EXT	bond1:2	10.10.10.7/255.255.224.0		EoIB/ Floating	ComputeNode1/ HOST1/ vServer3	WLS_PRODA1 Default External Channel	Initially enabled on vServer3, can be failed over by server migration to vServer4
HOST2VHN-EXT	bond1:2	10.10.10.8/255.255.224.0		EoIB/ Floating	ComputeNode2/ HOST2/ vServer4	WLS_PRODA2 Default External Channel	Initially enabled on vServer3, can be failed over by server migration to vServer4

4.6.4 Preparing the Network on Virtual Exalogic

This topic provides information on creating a private IPoIB network and reserving virtual IP addresses.

This section contains the following topics:

4.6.4.1 About Creating the Required Networks

When Exalogic Elastic Cloud is commissioned, a number of networks will be available. You must decide how to assign these networks to the virtual servers.

The following networks are required for the typical enterprise deployment. If these networks do not exist, then you will need to create them. For more information, see the *Oracle Exalogic Elastic Cloud Administrator's Guide*.

Public EoIB Client Access Network

This network is used to communicate with the corporate network. This network will be referred to as EoIB-client.

Private IPoIB FMW Network

This network is used for inter application communication. This network needs to be created as described in Creating a Private IPoIB Network. This network will be referred to as IPoIB-FMW.

IPoIB Storage Network

This is the network that virtual servers will use to connect to the ZFS storage appliance. This network will be referred to as IPoIB-Storage.

IPoIB Data Network

This is the network that virtual servers will use to communicate with a database on an attached Exadata machine. This network will be referred to as IPoIB-Data.

Note that if you are not using Exadata and your database is on an Ethernet-based host, then you might not be able to create this network. If you are unable to create this network, then you will have to use the Public EoIB Client Access Network (EoIB-client).

4.6.4.2 Creating a Private IPoIB Network

You need to create a private network to allow each of the vServers in the deployment to communicate privately with each other.

This network will only be available to assigned vServers and ensures that network communication between the vServers in the deployment is isolated from other network traffic.

To create a private IPoIB network for exclusive communication between the vServers in the deployment, perform the following steps:

1. Log in to Exalogic Control.
2. Expand **vDC Management**.
3. Navigate to **vDCs - Accounts** - *Cloud User Account*.
4. In the Actions window, click **Create Private vNet**.
5. Enter a Name. For example: `IPoIB_EDG`.
6. Click **Next**.
7. Select the **Number of Elements** to reserve on the network.

   Number of IP addresses are seven (HOST1, HOST2, HOST1VHN1, HOST1VHN2, HOST2VHN1,HOST2VHN2, EDGINTERNAL).

   Note: We are considering seven IP addresses in the example. The number of IP addresses you require will consists of 1 per physical host along with 1 per virtual IP you wish to assign on the network.

8. Click **Next**.
9. Click **Finish**.

4.6.4.3 Reserving Virtual IP Addresses

Each enterprise deployment uses a number of virtual IP addresses. To prevent Oracle Elastic Cloud assigning these IP addresses elsewhere, you need to reserve them for your use.

These IP addresses will be taken from one or more of the networks above.

To reserve IP addresses, perform the following steps:

1. Log in to Exalogic Control.
2. Expand **vDC Management**.
3. Navigate to **vDCs - Accounts** - *Cloud User Account*.
4. Click the **Networks** tab.

   The Network Dashboard is displayed.

5. Select the network you want to reserve IP addresses in. For example, **IPoIB_EDG** .

6. Click **Allocate VIP Addresses**.

   The Allocate VIP from vNet window is displayed.

7. Choose the number of virtual IP addresses you wish to reserve, for example 6, and click **Allocate VIP**.

   A window shows what virtual IP addresses have been reserved. Make a note of these.

Note:

This is only for virtual IP addresses on the internal IPoIB network. You will need to allocate virtual IP addresses on the client access network for communication with Oracle Traffic Director and the Administration Servers.

4.6.5 Enabling Virtual IP Addresses

This topic provides information to enable the virtual IP addresses.

Having completed the network chapter, you need to assign virtual IP addresses to the various network interfaces and hosts as described in Enabling Virtual IP Addresses.

4.7 Verifying Network Connectivity

After you have configured the network, it is important to ensure that you can use them to communicate.

To verify network connectivity, ping each IP address from each host.

For example:

```
ping host1
ping host1-stor
ping host1-int
ping host1-ext
ping adminvhn
ping host1vhn1
ping host1vhn2
```

5

Preparing Storage

Each enterprise deployment requires shared storage. In addition, each host requires some private storage.

Private storage can be considered either local disk or disk space on the Storage Area Network (SAN), which is made exclusively available to a host. The advantage of using private storage on a SAN (as opposed to local disk) is you have the built-in storage redundancy and fast backup technologies available within the SAN itself, which makes management easier.

An Exalogic appliance comes with a ZFS Storage appliance built in. This chapter explains the best practices to setup and configure the ZFS storage for a typical enterprise deployment.

The advantage of using local disk is that access is generally faster. It is true in physical deployments where direct access to the underlying local flash disk is available. In a virtual deployment, then local disk is simply a virtual disk on ZFS storage.

This section contains the following topics:

5.1 ZFS Concepts

This section provides conceptual information for setting up and configuring the ZFS storage for a typical enterprise deployment.

About Storage Pools
: Physical disks inside the ZFS storage are allocated to storage pools.

About Projects
: A project defines a common administrative control point for managing shares. All shares within a project can share common settings and can be backed up as a logical unit.

About Shares
: Shares are file systems and LUNs that are exported over supported data protocols to clients of the appliance.

5.1.1 About Storage Pools

Physical disks inside the ZFS storage are allocated to storage pools.

Space within a storage pool is shared between all shares.

5.1.2 About Projects

A project defines a common administrative control point for managing shares. All shares within a project can share common settings and can be backed up as a logical unit.

All file systems and LUNs are grouped into projects. Also, quotas can be enforced at the project level in addition to the share level. Projects can also be used solely for grouping logically related shares together, so their common attributes (such as accumulated space, or write frequency) can be accessed from a single point.

Note: On Exalogic, the use of LUNs or iSCSI are not supported. The only supported use of ZFS is through NFS shares.

5.1.3 About Shares

Shares are file systems and LUNs that are exported over supported data protocols to clients of the appliance.

File systems export a file-based hierarchy and can be accessed using NFS over IPoIB in the case of Exalogic machines. The project/share tuple is a unique identifier for a share within a pool. Multiple projects can contain shares with the same name, but a single project cannot contain shares with the same name. A single project can contain both file systems and LUNs, and they share the same namespace.

File systems can grow or shrink dynamically as needed, though it is also possible to enforce space restrictions on a per-share basis. All shares within a project can be backed up as a logical unit.

5.2 About the Default Storage Configuration

After commissioning an Oracle Exalogic Appliance, the storage will have the following default characteristics.

- Storage disks on the Sun ZFS Storage 7320 appliance are allocated to a single storage pool, such as `exalogic`, by default. Every compute node in an Oracle Exalogic machine can access both of the server heads of the storage appliance. The storage pool uses one of the server heads, which are also referred to as controllers. The server heads use active-passive cluster configuration. Exalogic compute nodes access the host name or IP address of a server head and the mount point based on the distribution of the storage pool.
- If the active server head fails, the passive server head imports the storage pool and starts to offer services. From the compute nodes, the user may experience a pause in service until the storage pool is started to be serviced from the working server head. However, this delay might not affect client activity; it can affect disk I/O only.
- By default, data is mirrored, which yields a highly reliable and high-performing system. The default storage configuration is done at the time of manufacturing, and it includes the following shares:
  - Two exclusive NFS shares for each of the Exalogic compute nodes - one for crash dumps and another for general purposes.
    In this scenario, you can implement access control for these shares based on your requirements.
  - Two common NFS shares to be accessed by all compute nodes - one for patches and another for general purposes.

5.3 Overview of Enterprise Deployment Storage

This topic provides information on enterprise deployment storage which includes sharing requirements, read/write requirements and artifact characteristics.

When deciding how to prepare shared storage, consider the following:

Sharing Requirements

An artifact will either need be shared, not-shared (or private), or available.:

- Shared: The artifact resides on shared storage and can be *simultaneously* viewed and accessed by all client machines.
- Available: The artifact resides on shared storage, but only one client machine can view and access the artifact at a time. If this machine fails, another client machine can then access the artifact.
- Not-shared: The artifact can reside on local storage and never needs to be accessible by other machines.

Read/Write Requirements

An artifact will also have specific read/write requirements:

- Read-Only: This artifact is rarely altered but only read at runtime.
- Read-Write: This artifact is both read and written to at runtime.

Artifact Characteristics

With the above requirements in mind, the artifacts deployed in a typical enterprise deployment can be classified as follows:

Artifact Type	Sharing	Read/Write
Binaries — Application Tier	Shared	Read-Only
Binaries — Web Tier	Private	Read-Only
Managed Server Domain Home	Private	Read-Write
Admin Server Domain Home	Available	Read-Write
Runtime Files	Available	Read-Write
Node Manager Configuration	Private	Read-Write
Application Specific Files	Shared	Read-Write

5.4 Understanding the Enterprise Deployment Directory Structure

Each Oracle Fusion Middleware enterprise deployment is based on a similar directory structure. This directory structure has been designed to separate binaries, configuration, and runtime information.

Binaries are defined as the Oracle software installation. Binaries include such things as the JDK, WebLogic Server, and the Oracle Fusion Middleware software being used (for example, Oracle Identity and Access Management or WebCenter).

The Server Domain Home directories will be read to and written from at runtime.

The Runtime Directories are written at Runtime, and hold information such as JMS queue files. The Server Domain Home directories will be read and written at runtime.

This section contains the following topics:

Shared Binaries
: All machines use the same set of binaries to run processes. Binaries are installed once and then only modified during patches and upgrades.

Private or Shared Managed Server Domain Homes
: Managed server domain homes are used to run managed servers on the local host machine.

A Domain Home for the Administration Server
: The Domain Home from which the Administration Server runs is contained in a separate volume that can be mounted on any machine.

Shared Runtime Files
: Files and directories that might need to be available to all members of a cluster are separated into their own directories.

Local Node Manager Directory
: Each client machine will have its own Node Manager and Node Manager configuration directories.

Shared Application Files
: Any files that must be both shared *and* accessed concurrently (read/write) will be application dependent.

Diagrams of the Typical Enterprise Deployment Directory Structure

The following diagrams show the recommended use of shared and local storage for a typical enterprise deployment:

5.4.1 Shared Binaries

All machines use the same set of binaries to run processes. Binaries are installed once and then only modified during patches and upgrades.

For maximum availability and protection, it is recommended that redundant binaries are created (that is, two different shares are created that mirror each other). One share will be mounted to odd-numbered hosts, and one share will be mounted to even-numbered hosts. This is so that corruption in one set of binaries does not impact the entire system.

5.4.2 Private or Shared Managed Server Domain Homes

Managed server domain homes are used to run managed servers on the local host machine.

For performance reasons, it is recommended to have these domain homes reside on private storage attached to the host machine.

5.4.3 A Domain Home for the Administration Server

The Domain Home from which the Administration Server runs is contained in a separate volume that can be mounted on any machine.

This allows the Administration Server to be readily available if the machine that was hosting it has failed or is otherwise unavailable.

5.4.4 Shared Runtime Files

Files and directories that might need to be available to all members of a cluster are separated into their own directories.

These will include JMS files, transaction logs, and other artifacts that belong to only one member machine of a cluster but might need to be available to other machines in case of failover.

5.4.5 Local Node Manager Directory

Each client machine will have its own Node Manager and Node Manager configuration directories.

A Node Manager is associated with one physical machine, these directories can also optionally reside on private storage.

5.4.6 Shared Application Files

Any files that must be both shared *and* accessed concurrently (read/write) will be application dependent.

No files that are part of the Oracle Fusion Middleware WebLogic Infrastructure have this requirement, but some Fusion Middleware products (such as Oracle WebCenter Content) and custom applications might require shared files on disk.

5.4.7 Diagrams of the Typical Enterprise Deployment Directory Structure

The following diagrams show the recommended use of shared and local storage for a typical enterprise deployment:

Figure 5-1 Recommended Shared Storage Directory Structure

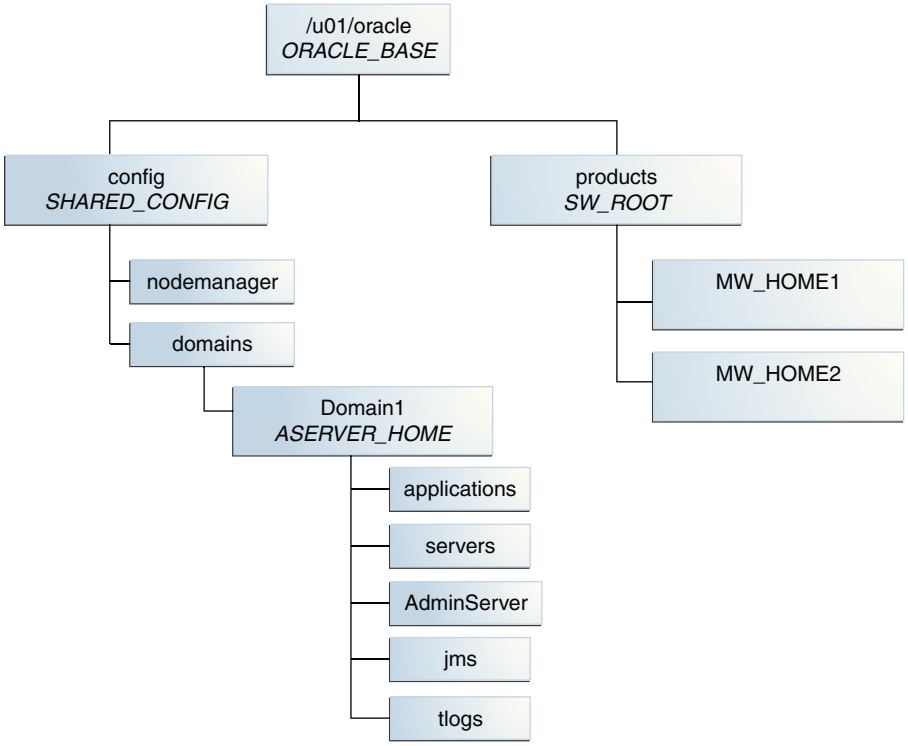

Figure 5-2 Recommended Private Storage Directory Structure

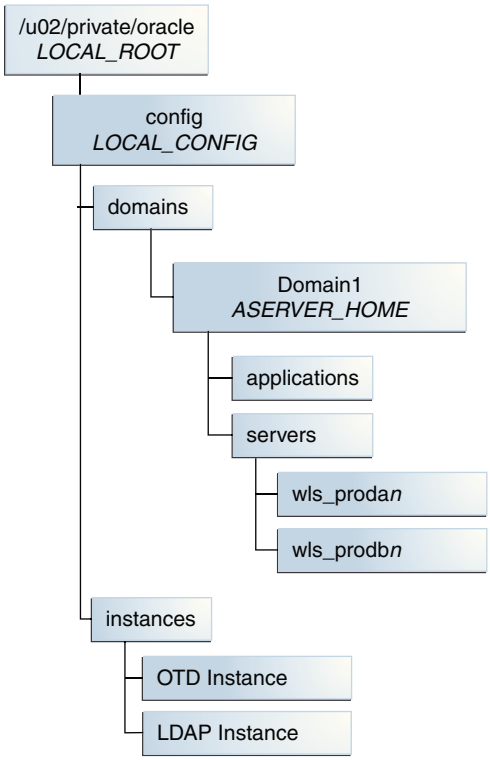

Figure 5-3 Recommended Directory Structure for Private Binary Storage

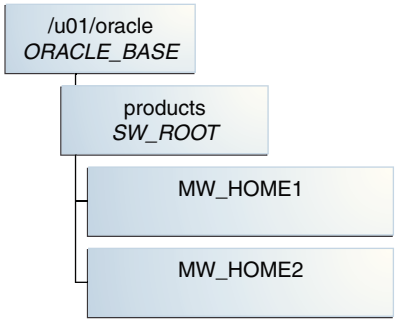

Private binary storage is generally only required when the Web Tier is on a separate host (for example, virtual deployments).

5.5 Shared Storage Concepts

This section describes shared storage concepts.

This section contains the following topics:

Shared Storage Protocols and Devices

ZFS can present storage either as block storage or as NFS shares. This document will describe setting up NFS (Network File System). NFS version 3 (v3) and NFS version 4 (v4) are supported by ZFS.

NFS Version 3

NFS version 3 (v3) is perhaps the most common standard for UNIX clients that requires a remote, shared file system.

NFS Version 4

NFS version 4 (v4) has several advantages over NFS v3. One advantage is support for an authentication protocol, such as NIS, to map user permissions instead of relying on UIDs. The other advantage is support for the expiration of file locks through leasing.

5.5.1 Shared Storage Protocols and Devices

ZFS can present storage either as block storage or as NFS shares. This document will describe setting up NFS (Network File System). NFS version 3 (v3) and NFS version 4 (v4) are supported by ZFS.

A shared file system is both a file system and a set of protocols used to manage multiple concurrent access. The most common of these protocols is NFS for UNIX clients. These solutions consist of a back-end file system, such as ZFS or EXT3, along with server processes (such as NFSD) to manage and mediate access.

5.5.2 NFS Version 3

NFS version 3 (v3) is perhaps the most common standard for UNIX clients that requires a remote, shared file system.

Permissions between the client machine and the storage server are mapped by using the UID of the users. NFS v3 can be configured and exported on most devices. The client machine can then mount the remote device locally with a command such as the following:

```
$ mount -t nfs storage_machine:/export/nfsshare /locmntddir -o
rw,bg,hard,nointr,tcp,vers=3,timeo=300,rsize=32768,wsi ze=32768
```

This command mounts the remote file system at`/locmnntdir`. Most of the values specified above are default, and so an NFS v3 device can also be mounted as:

```
$ mount -t nfs storage_machine:/export/nfsshare /locmntddir -o nointr,timeo=300
```

The right values of `rsize`, `wsize`, and `timeo` will depend on attributes of the network between the client machine and the storage server. For example, a higher timeout value will help to minimize timeout errors.

The attributes `noac` and `actimeo` are not used here. These are used to disable attribute caching and may be necessary when objects such as Oracle data files are placed on NFS. They are not necessary, however, for most WebLogic applications or specific Fusion Middleware applications, such as SOA, WebCenter, or BI. Enabling these parameters will carry performance penalties as well.

Other options include `noatime`, which will provide a small performance boost on read-intensive applications by eliminating the requirement for access-time writes.

5.5.3 NFS Version 4

NFS version 4 (v4) has several advantages over NFS v3. One advantage is support for an authentication protocol, such as NIS, to map user permissions instead of relying on UIDs. The other advantage is support for the expiration of file locks through leasing.

After a specified time, the storage server releases unclaimed file locks. The default lease expiration time is storage server dependent. Typical values are 45 or 120 seconds. In Sun ZFS systems, the default is 90 seconds. In a failover scenario, where a server has failed and another server must claim or access the files owned by the failed server, it is able to do so after the lease expiration time. This value should be set lower than the

failover time so that the new server can successfully acquire locks on the files. For example, this should be set lower than the time for server migration in a cluster using JMS file based persistence.

This differs from NFS v3 where locks are held indefinitely. The locks held by a process that has died have to be released manually in NFS v3.

An NFS v4 file system can be mounted on Linux similarly to NFS v3:

```
$ mount -t nfs4 storage_machine:/export/nfsv4share /locmntddir -o
nointr,timeo=300
```

5.6 Enterprise Deployment Storage Design Considerations

This topic provides information on ztorage design considerations and for storage requirements specific to your deployment type, refer to the appropriate product enterprise deployment guide.

Based on the typical directory structures described in Diagrams of the Typical Enterprise Deployment Directory Structure, storage can be classified as follows:

Table 5-1 Summary of Enterprise Deployment Storage Design Considerations

Project	Shares	Mount point
product_binaries	shared_binaries	`/u01/oracle/products`
	webhost1_local_binaries	`/u01/oracle/products`
	webhost2_local_binaries	`/u01/oracle/products`
product_config	shared_config	`/u01/oracle/config`
	host1_local_config	`/u02/private/oracle/config`
	host2_local_config	`/u02/private/oracle/config`
Runtime	shared-runtime	`/u01/oracle/runtime`

Using this model, all binaries can be set to read-only after installation at the product level. Configuration information can still be read/write.

Backups can be taken at the project level. Disaster recovery can replicate content at the project level, allowing binaries to be replicated only when necessary and for configuration/jms queues to be replicated more frequently.

5.7 Preparing Exalogic Storage for an Enterprise Deployment

Prepare storage for the physical Exalogic deployment by creating users and groups in NIS, creating projects using the storage appliance BUI and shares in a project using the BUI.

This section contains the following topics:

Prerequisite Storage Appliance Configuration Tasks
This topic provides the prerequisites for the storage appliance.

Creating Users and Groups in NIS
: This step is optional. If you want to use the onboard NIS servers, you can create users and groups using the steps in this section.

Creating Projects Using the Storage Appliance Browser User Interface (BUI)
: To configure the appliance for the recommended directory structure, you create a custom project using the Sun ZFS Storage 7320 appliance Browser User Interface (BUI).

Creating the Shares in a Project Using the BUI
: After you have created projects, the next step is to create the required shares within the project.

Allowing Local Root Access to Shares
: In order to allow compute nodes or virtual machines to access ZFS shares, you must add an NFS exception to allow you.

5.7.1 Prerequisite Storage Appliance Configuration Tasks

This topic provides the prerequisites for the storage appliance.

The instructions in this guide assume that the Sun ZFS Storage 7320 appliance is already set up and initially configured. Specifically, it is assumed you have reviewed the following sections in the *Oracle Exalogic Elastic Cloud Machine Owner's Guide*.

The examples below are for the typical enterprise deployment. Refer to the specific product Enterprise Deployment Guide for the list of shares you need to create for the product you are installing.

- "Prerequisites"
- "Getting Started"
- "Sun ZFS Storage 7320 Appliance Overview"
- "Configuration Overview"
- "Naming Service"

5.7.2 Creating Users and Groups in NIS

This step is optional. If you want to use the onboard NIS servers, you can create users and groups using the steps in this section.

First, determine the name of your NIS server by logging into the Storage BUI. For example:

1. Log in to the ZFS Storage Appliance using the following URL:

   ```
   https://exalogicsn01-priv:215
   ```

2. Log in to the BUI using the storage administrator's user name (root) and password.
3. Navigate to **Configuration**, and then **Services**.
4. Click on **NIS**.

   There is a green dot next to it if it is running. If it is not running and you wish to configure NIS, see the *Oracle Exalogic Elastic Cloud Machine Owner's Guide*

5. Click on NIS. You will see the named NIS servers. Make a note of one of the NIS servers.

   Now that you have the name of the NIS server, open a terminal window on the NIS server as root and perform the following steps:

6. Create users on the NIS master server as described in Creating Users and Groups in NIS.

7. Add users to `yp` by performing the following steps:

   a. Navigate to the `/var/yp` directory:

   b. Run the following command:

      ```
      make -C /var/yp
      ```

   c. If required, restart the services using the following commands:

      ```
      service ypserv start
      service yppasswdd start
      service rpcimapd start
      service ypbind start
      ```

   d. Validate that the users and groups appear in NIS by issuing the command:

      ```
      ypcat passwd
      ```

      and

      ```
      ypcat group
      ```

5.7.3 Creating Projects Using the Storage Appliance Browser User Interface (BUI)

To configure the appliance for the recommended directory structure, you create a custom project using the Sun ZFS Storage 7320 appliance Browser User Interface (BUI).

After you set up and configure the Sun ZFS Storage 7320 appliance, the appliance has a set of default projects and shares. For more information, see "Default Storage Configuration" in the *Oracle Exalogic Elastic Cloud Machine Owner's Guide*.

The instructions in this section describe the specific steps for creating a new project for the enterprise deployment. For more general information about creating a custom project using the BUI, see "Creating Custom Projects" in the *Oracle Exalogic Elastic Cloud Machine Owner's Guide*.

To create a new custom project on the Sun ZFS Storage 7320 appliance:

1. Log in to the ZFS Storage Appliance using the URL:

   ```
   https://exalogicsn01-priv:215
   ```

2. Log in to the BUI using the storage administrator's user name (root) and password.

3. Navigate to the **Projects** page by clicking on the **Shares** tab, then the **Projects** sub-tab.

   The BUI displays the Project Panel.

4. Click **Add** next to the Projects title to display the Create Project window.

Enter a name for the project. For example: **EDG_Binaries**

Click **Apply**.

5. Click **Edit Entry** next to the newly created project.
6. Click the **General** tab on the project page to set project properties.

   Update the following values:

   - Mountpoint: Set to /export/EDG_Binaries
   - Under the Default Settings Filesystems section set the following values:
     - User: oracle
     - Group: oinstall
     - Permissions: RWX RWX R_X
7. For the purposes of the enterprise deployment, you can accept the defaults for the remaining project properties.

   For more information about the properties you can set here, see the "Project Settings" table in the *Oracle Exalogic Elastic Cloud Machine Owner's Guide*.
8. Click **Apply** on the **General** tab to create the project.
9. Repeat for each project you are creating.

5.7.4 Creating the Shares in a Project Using the BUI

After you have created projects, the next step is to create the required shares within the project.

For more general information about creating custom shares using the BUI, see "Creating Custom Shares" in the *Oracle Exalogic Elastic Cloud Machine Owner's Guide*.

To create each share, use the following instructions, replacing the name and privileges, as described in Table 5-1:

1. Login to the storage system BUI, using the following URL:

   ```
   https://exalogicsn01-priv:215
   ```
2. Navigate to the Projects page by clicking the **Shares** tab, and then the **Projects** sub-tab.
3. On the **Project** panel, select the project to which you want to add a share. For example: **EDG_Binaries**.
4. Click the plus (+) button next to **Filesystems** to add a file system.

   The Create Filesystems screen is displayed.
5. In the Create Filesystems screen, choose your new project from the **Project** pull-down menu.
6. In the **Name** field, enter the name for the share.

   For example, shared_binaries.

7. From the **Data migration source** pull-down menu, choose **None**.
8. Make sure the user and group displayed are the same as the NIS user and Group you created for owning the filesystems. If not change them accordingly.
9. Select the **Permissions** option and set the permissions for each share.
10. Select the **Inherit Mountpoint** option.

> **Note:** Initially these will be Read/Write but after installation of the product, these may be changed as described in Artefact characteristics table.

11. To enforce UTF-8 encoding for all files and directories in the file system, select the **Reject non UTF-8** option.
12. From the **Case sensitivity** pull-down menu, select **Mixed**.
13. From the **Normalization** pull-down menu, select **None**.
14. Click **Apply** to create the share.

Repeat the procedure for each share listed in Table 5-1.

5.7.5 Allowing Local Root Access to Shares

In order to allow compute nodes or virtual machines to access ZFS shares, you must add an NFS exception to allow you.

You can create exceptions either at the individual, share, or project level. Additionally, if you want to run commands or traverse directories on the share as the root user, there is an option to enable.

To keep things simple, in this example you create the exception at the project level.

To create an exception for NFS at the project level:

1. In the Browser User Interface (BUI), access the Projects user interface by clicking **Shares**, and then **Projects**.

   The Project Panel appears.
2. On the Project Panel, click **Edit** next to the project you wish to change.
3. Select the **Protocols** tab.
4. Click the + sign next to NFS exceptions.
5. Select **Type: network**.
6. In the **Entity** field, enter the IP address of the compute node or virtual machine in CIDR format. For example: `192.168.8.0/22` for the Infiniband CIDR of Exalogic's compute nodes on a physical rack..

   `172.17.0.0/16` for the Infiniband CIDR of Exalogic's IPoIB-vserver-shared-storage network on a virtual rack.

> **Note:** These ranges can be different, based on installation preferences. You may need to analyze the results of `/sbin/ip addr|grep inet`, a completed Exalogic Configurator spreadsheet, and the Configuration, Services or Network screen in the ZFS BUI to determine which NFS exception to define.

7. Set **Access Mode** to **Read/Write** and check **Root Access**.
8. Click **Apply**.
9. Repeat for each compute node or vServer that accesses the ZFS appliance.

6

Creating Exalogic Virtual Servers (vServers)

This chapter describes how to create virtual servers using Enterprise Manager Operations Control (EMOC) and to verify that servers of the same type do not run on the same underlying hardware.

Refer to your product-specific enterprise deployment guide to determine the number and size of vServers you must create.

A virtual server (vServer) is a virtual host created in the Oracle Exalogic Elastic Cloud. A virtual server is similar to a physical server. If you are planning to deploy on Virtual Exalogic, then you need to create a specific number of virtual servers to host the various components of the deployment.

Because a virtual server can run on any physical compute node in the Exalogic machine, for maximum availability, it is recommended that vServers hosting the same component do not run on the same physical server. This ensures continuity of service, even if the underlying compute node fails. For example, both wls_proda1 and wls_proda2 should not run on compute node 1. They should be distributed across two different compute nodes.

This section contains the following topics:

controlled way to do this, create separate volumes for each vServer to mount for the swap and temp space.

vServer Types
: This topic provides information on the vServer types along with memory and swap space.

Creating a vServer
: This topic provides the steps for creating a vServer.

Updating vServers
: This topic provides information on updating the root password, post network configuration and in setting the MTU size on Infiniband interfaces.

Moving Swap and TMP to Separate Volumes
: If you create separate disk volumes for `swap` and `tmp`, update your vServer to use these new volumes.

6.1 Prerequisites

This topic provides the prerequisites for virtual servers.

Before starting an Exalogic deployment, ensure that the following tasks have been performed:

1. Exalogic rack has been commissioned and one-command run.
2. Accounts have been created in Exalogic Control.
3. Private IPoIB network has been created for the account, enabling secure communications between the virtual servers assigned to the account as described in Creating a Private IPoIB Network.
4. You have created and loaded a Server Template for the operating system you wish to deploy.
5. You have created a vServer type that matches the specification of the virtual servers you want to create.
6. A Client Access Network has been created, using a bonded Network Interface for communication between the vServers and an external load balancer.

6.2 Sizing a Virtual Server

Oracle provides a few out of the box templates for virtual servers. These templates have a base operating system of Oracle Linux 5 or Oracle Linux 6.

Virtual Servers come in a number of sizes, each size having an assigned number of virtual CPU's and memory. These templates are only examples and users can create their own with different operating system images and different CPUs and memory.

When you create a Virtual Server, you need to create a Virtual Server which is sufficient to your needs. If you are using this guide in conjunction with a product EDG such as the Identity and Access Management EDG then those guides will have guidelines on the sizes of Virtual Servers to use.

If you do not have this information, you can create a vServer with the same characteristics as a physical server. For example, if you are deploying Oracle HTTP server and it requires 1 CPU 4GB of memory and 10GB of disk space then you need to create a Virtual Server of the same size.

For additional information, refer to the product installation guides and associated support notes such as *Oracle Fusion Middleware System and Requirements and Specifications* document.

6.3 Obtaining a vServer Guest Template

Oracle provides a number of base templates for Oracle Exalogic and are available to download from the Oracle E-delivery.

If the default template provided is not sufficient, download another template from Oracle E-Delivery. It is recommended that Oracle Linux v6 or later is used. For example to download the OL6.5 template:

1. Login to *edelivery.oracle.com*.

   Search for

   - Product Pack: Oracle Fusion Middleware
   - Platform: Linux x86-64

2. Click the link for *Oracle Exalogic Elastic Cloud Software 11g Media Pack*.
3. Click download next to *Oracle Exalogic 2.0.6.1.2 Base Guest Template for Exalogic Linux 6 x86-64(64bit)*.

6.4 Loading the Guest Template into Exalogic Control

Server Templates contain the configuration of an individual vServer with its virtual disk. Templates can be of the format .tgz, .tar or other file types.

You can use HTTP, HTTPS, or FTP protocols to upload Server Templates from any network, including the external EoIB network, that is available to the virtual machine hosting the Enterprise Controller component of Exalogic Control. You can upload a Server Template to Exalogic Control as follows:

1. Log in to the Exalogic Control as a Cloud User.
2. In the navigation pane, click **vDC Management**.
3. Under vDC Accounts, click the name of your account.

   The vDC Account dashboard is displayed.

4. Click **Server Templates** on the top navigation bar.

   The Server Templates available in your account are listed.

5. Under Server Templates, click the **Upload Server Template** icon. Alternatively, click **Upload Server Template** under Operate on the Actions pane.
6. On the Identify Server Template screen, enter a name and description for the Server Template to be uploaded.
7. You can add tags for later identification and search and click **Next**.

   The Specify Server Template Details screen is displayed.

8. Enter the following information:

- Image Sub Type : Template
- Upload Source: Local Host

9. Select the template file you downloaded from edelivery and click **Upload**.

6.5 About Distribution Groups

A distribution group prevents virtual servers assigned to it from running on the same physical nodes. By preventing different vServers of the same type running on the same physical server, you prevent the failure of the underlying physical server from taking out the complete system.

In a Fusion Middleware Exalogic implementation, the following distribution groups are required:

- EDG_OTD: Prevents two Oracle Traffic Director Servers from running on the same physical server
- EDG_PRODA: Prevents two instances of PRODA from running on the same physical server
- EDG_PRODB: Prevents two instances of PRODB from running on the same physical server
- EDG_LDAP: Prevents two LDAP instances from running on the same physical server.

This section contains the following topic:

Creating a Distribution Group

6.5.1 Creating a Distribution Group

To create a distribution group, perform the following steps:

1. Log in to Exalogic Control using the following URL:

   `https://emochost.example.com:9443/emoc/`

   Log in using your assigned administrator account. For more information on accounts, refer to the *Exalogic Elastic Cloud Administrator's Guide*.

2. Expand **vDC Management**.
3. Navigate to **vDCs - Accounts - Cloud Admin Account**.
4. In the actions window, click **Create Distribution Group**.
5. Enter a **Name**, for example: `EDG_OTD`.
6. Click **Next**.
7. Enter **Number of Elements**. This is a number that defines the number of EECS Servers on which the vServers can be placed.
8. Click **Next**.
9. Click **Finish**.

Repeat for each distribution group to be created. Table 6-1 lists distribution groups and the number of elements for each.

Table 6-1 Number of Elements for Distribution Groups

Distribution Group	Number of Elements
EDG_LDAP	2
EDG_OTD	2
EDG_PRODA	2
EDG_PRODB	2

6.6 Creating vServer Volumes

When you create a vServer, by default, it creates one default volume and allocates the space to swap and the root file system. For a more efficient controlled way to do this, create separate volumes for each vServer to mount for the swap and temp space.

To create separate volumes for each vServer:

1. Log in to Exalogic Control using the following URL:

   `https://emochost.example.com:9443/emoc/`

   Log in using your assigned administrator account. For more information on accounts, refer to the Oracle Elastic Cloud Administration Guide.

2. Expand **vDC Management**.
3. Navigate to **vDCs**, **Accounts**, and then **Cloud Admin Account**.
4. Select **Create Volume** from the Actions menu.
5. Give the volume a name, for example **vServer1_tmp**, and a description.
6. Click **Next**.
7. On the Volume Configuration screen, enter a size for the volume.

   Note: Swap space follows the Linux best practice of 2x memory upto 16GB memory.

   Do not select **shared**.
8. Click **Next**.
9. On the Volume Summary screen, click **Finish** to create the volume.
10. Repeat the above steps for the swap volume, example vServer1_swap.

6.7 vServer Types

This topic provides information on the vServer types along with memory and swap space.

Table 6-2 lists the vServer types used in this document. These vServer types can be used as a guide. Refer to the *Oracle Fusion Middleware System Requirements and Specifications* for the latest hardware requirements.

Table 6-2 vServer Types

vServer Type	Memory	Swap Space	Tmp Space
LARGE	8GB	16GB	2GB
EXTRA_LARGE	16GB	16GB	2GB

6.8 Creating a vServer

This topic provides the steps for creating a vServer.

To create a vServer, perform the following steps:

1. Log into Exalogic Control.
2. Expand **vDC Management**.
3. Navigate to **vDCs - Accounts - Cloud Admin Account**.
4. In the **Actions** window, click **Create vServer**.
5. Enter the following:
   - **Name**: For example: `otdhost1`
   - **Number of vServers**:`1`

   Select: **Support High Availability**.
6. Click **Next**.
7. Choose the Server Template you want to deploy.
8. Click **Next**.
9. Choose the vServer type you wish to create, for example: **LARGE**
10. If you have created volumes for `swap` and `tmp` for example, select them here.
11. Click **Next**.
12. Enter all of the virtual networks you want to assign.

Note: Refer to the network diagrams for the networks you wish to assign.

Network Diagrams	Description
Web Servers	EoIB-client, IPoIB-FMW, IPoIB-Storage
Application Servers	IPoIB-FMW, IPoIB-Storage. IPoIB-Data
LDAP Servers	IPoIB-FMW, IPoIB-Storage, IPoIB-Data (*)

LDAP servers will only require access to the IPoIB-Data network, if the directory type is OID.

13. For each chosen network, enter the following:

    - **IP Address Type** - Static or Automatic
    - **IP Address** - Enter the IP address if you have a predetermined IP address to use.

      Note: IPoIB-vserver shared storage typically does not require pre-assigned IP addresses.

    - **Hostname** - Select the fully qualified host name you wish to assign to the IP address.

14. Click **Next**.

15. Enter the Distribution Group to use.

16. Click **Next**.

17. Click **Next** on vServerAccessControl screen.

18. Click **Finish**.

Repeat for each vServer to be created.

6.9 Updating vServers

This topic provides information on updating the root password, post network configuration and in setting the MTU size on Infiniband interfaces.

Now that the vServers have been created, you need to perform the following tasks to make them available for use.

This section contains the following topics:

Updating the root Password
: This topic provides information on updating the root password.

Updating /etc/hosts File
: This topic provides information on updating the host file.

Post Network Configuration
: Now that your vServer has been created, you must configure it as appropriate to your organization.

Set MTU size on InfiniBand Interfaces
: To maintain optimum performance, you must update the MTU size of each of the InfiniBand interfaces on the vServer to 65520.

6.9.1 Updating the root Password

This topic provides information on updating the root password.

When the vServer is created, it has a default password, which is generally `ovsroot`. Change this to a value appropriate for your organization.

6.9.2 Updating /etc/hosts File

This topic provides information on updating the host file.

After configuration, your `hosts` file will look something like:

```
IP Address   Host_Name
```

For example:

```
192.168.32.3 host1-stor
```

Change the `hosts` file so that it contains both fully qualified and short names for each network, for example:

```
192.168.10.3 host1-int.example.com host1-int
192.168.32.3 host1-stor.example.com host1-stor
192.168.10.3 host1-data.example.com host1-data
10.10.10.5   host1-ext.example.com host1-ext
```

Note:

External Network interface names are assumed to be in DNS.

6.9.3 Post Network Configuration

Now that your vServer has been created, you must configure it as appropriate to your organization.

This typically includes the following steps:

Determine vServer Storage IP Address
: When you created your vServer, you added the network IPoIB-Storage. This is the network the vServers use to communicate with the ZFS storage appliance.

Determine Storage Appliance IP Address
: This topic provides information on determining storage appliance IP address.

6.9.3.1 Determine vServer Storage IP Address

When you created your vServer, you added the network IPoIB-Storage. This is the network the vServers use to communicate with the ZFS storage appliance.

In order for them to communicate properly, you must determine the appropriate IP address of the storage appliance to use.

To determine the IP address, perform the following steps:

1. Log in to Exalogic Control as a Cloud user.
2. From the navigation pane on the left, select **vDC Management**.
3. Under vDC Accounts, expand the name of your account, and select the vServer for which you want to configure access to the storage appliance.

   The vServer dashboard is displayed.

4. Select the **Network** tab, and note the IP address of the vServer for the IPoIB-Storage network. This corresponds with the **-stor** entry in the `/etc/hosts` file.

   For example: `172.17.0.100`

6.9.3.2 Determine Storage Appliance IP Address

This topic provides information on determining storage appliance IP address.

1. Log in to the storage appliance as root.

   For example, type:

   ```
   ssh root@exalogicsn01.example.com
   ```

2. Show the network interfaces using the command:

   ```
   configuration net interfaces show
   ```

3. The output is similar to the following:

   ```
   configuration net interfaces show
   Interfaces:

   INTERFACE   STATE    CLASS LINKS       ADDRS                 LABEL
   igb0        up       ip    igb0        10.244.64.60/21       igb0
   igb1        offline  ip    igb1        10.244.64.61/21       igb1
   ipmp1       up       ipmp  pffff_ibp1  192.168.10.15/24      ipmp1
                             pffff_ibp0
   ipmp2       up       ipmp  p8001_ibp0  192.168.20.9/24       IB_IF_8001
                             p8001_ibp1
   ipmp3       up       ipmp  p8002_ibp0  192.168.21.9/24       IB_IF_8002
                             p8002_ibp1
   ipmp4       up       ipmp  p8005_ibp0  172.17.0.9/16         IB_IF_8005p8005_ibp1
   p8001_ibp0  up       ip    p8001_ibp0  0.0.0.0/8             ibp0.8001
   p8001_ibp1  up       ip    p8001_ibp1  0.0.0.0/8             ibp1.8001
   p8002_ibp0  up       ip    p8002_ibp0  0.0.0.0/8             ibp0.8002
   p8002_ibp1  up       ip    p8002_ibp1  0.0.0.0/8             ibp1.8002
   p8005_ibp0  up       ip    p8005_ibp0  0.0.0.0/8             ibp0.8005
   p8005_ibp1  up       ip    p8005_ibp1  0.0.0.0/8             ibp1.8005
   pffff_ibp0  up       ip    pffff_ibp0  0.0.0.0/8             ibp0
   pffff_ibp1  up       ip    pffff_ibp1  0.0.0.0/8             ibp1
   ```

4. Determine the corresponding IP address by looking for the IP address in the same range as 172.17.0.100.

   In this example, it is the one associated with interface ipmp4, for example: 172.17.0.9.

5. Create an entry in the `/etc/hosts` on all vservers, for example:

   ```
   172.17.0.9  zfsinternal.example.com  zfsinternal
   ```

6.9.4 Set MTU size on InfiniBand Interfaces

To maintain optimum performance, you must update the MTU size of each of the InfiniBand interfaces on the vServer to 65520.

To do this, perform the following steps:

1. Log in to the vServer as the `root` user.

2. Verify the current MTU for `bond2` by running the `ifconfig` command.

   Note:

   The steps in this procedure use `bond2` as an example. This procedure should be repeated for all the InfiniBand interfaces.

```
ifconfig bond2bond2     Link encap:InfiniBand  HWaddr
          80:58:08:CA:FE:80:00:00:00:00:00:00:00:00:00:00:00:00:00:00
          inet addr:192.168.1.12  Bcast:192.168.1.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MASTER MULTICASTMTU:1500Metric:1
          RX packets:9 errors:0 dropped:0 overruns:0 frame:0
          TX packets:7 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:504 (504.0 b)  TX bytes:420 (420.0 b)
```

3. Append the line `MTU=64000` to the `ifcfg` file corresponding to the `bond2` interface:

```
echo MTU=64000 >> /etc/sysconfig/network-scripts/ifcfg-bond2
```

4. Verify whether the `MTU=64000` line was added to the `ifcfg-bond2` file:

```
cat /etc/sysconfig/network-scripts/ifcfg-bond2 | grep MTU
MTU=64000
```

5. Find the slave interfaces for `bond2`:

```
cd /etc/sysconfig/network-scripts
grep "MASTER=bond2" ifcfg-* | awk -F":" '{print $1}'
ifcfg-ib0.8009
ifcfg-ib1.8009
```

6. Set the mode to `connected` for both the slave interfaces of the `bond2` interface:

```
echo connected > /sys/class/net/ib0.8009/mode
echo connected > /sys/class/net/ib1.8009/mode
```

7. Perform Steps 2 through 6 for the other InfiniBand interfaces.

8. Stop and start the vServer as described in Oracle Exalogic Elastic Cloud Machine Owner's Guide.

9. After the vServer starts, log in again to the vServer using SSH as the `root` user.

10. Run the `ifconfig` command for each InfiniBand interface, and verify that the output of the command displays `MTU:64000`, as shown in the following example for `bond2` and its slave interfaces:

```
ifconfig bond2| grep MTU
UP BROADCAST RUNNING MASTER MULTICASTMTU:64000Metric:1

ifconfig ib0.8009| grep MTU
UP BROADCAST RUNNING SLAVE MULTICASTMTU:64000Metric:1

ifconfig ib1.8009| grep MTU
UP BROADCAST RUNNING SLAVE MULTICASTMTU:64000Metric:1
```

6.10 Moving Swap and TMP to Separate Volumes

If you create separate disk volumes for `swap` and `tmp`, update your vServer to use these new volumes.

The disk volumes are added to your virtual server as virtual volumes. They appear in the `/dev` directory as `xvdb/c`.

To determine the exact names, run the following command:

```
fdisk -l
```

Sample command output:

```
255 heads, 63 sectors/track, 1305 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x000c520c

Device Boot Start End Blocks Id System
/dev/xvda1 * 1 32 256000 83 Linux
Partition 1 does not end on cylinder boundary.
/dev/xvda2 32 1305 10223616 8e Linux LVM

Disk /dev/xvdb: 18.3 GB, 18253611008 bytes
255 heads, 63 sectors/track, 2219 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x00000000


Disk /dev/xvdc: 2147 MB, 2147483648 bytes
255 heads, 63 sectors/track, 261 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x00000000
```

As you can see from the output, `/dev/xvda` has a partition created on it so it is in use. This is the default disk.

Disks `/dev/xvdb` and `/dev/xvdc` do not have a partition and are therefore the attached disk volumes. You can determine which is which by the size of the volumes.

To use these volumes in the vServer, create a partition of type LVM. This enables the use of Linux LVM, and make resizing easier if required later. The procedure is the same if you are using the disk for `swap` or `tmp`.

This section contains the following topics:

Creating a LVM partition
: This topic provides detailed steps to create a LVM partition.

Creating Logical Volumes
: Now that you have disk partitions, create logical volumes to use those disks.

Creating a Swap File on the New Logical Volume
: This topic provides instructions to create a sawp file on the new logical volume.

Moving /tmp to the New Logical Volume
: If you have created a logical volume for `/tmp`, you can enable this by first creating a file system on it and then mounting it as a disk.

6.10.1 Creating a LVM partition

This topic provides detailed steps to create a LVM partition.

1. Choose a disk to work on using the following command:

   ```
   fdisk disk_name
   ```

   For example:

   ```
   fdisk /dev/xvdb
   ```

2. When prompted for a command, type `n`.
3. You are asked if you wish to create an extended or primary partition. Select `p` for primary.
4. When promoted for a partition number, enter `1`.
5. You are then asked where on the disk to create the partition. Accept the default value of `1`. Accept the default end value to use the entire disk.
6. You are now asked First cylinder (1-2088, default 1), use default value of `1`. Last cylinder, +cylinders or +size{K,M,G} (1-2088, default 2088): use the default value of `2088`.
7. Now that the partition has been created, give it a type. To do this, when prompted for a command, enter `t`.
8. You can see the list of types available by entering the command `L`.
9. When prompted for the Hex code, enter the code (from the previous list) for the Linux LVM. This is typically `8e`.
10. Save your changes using the command `w`.
11. Validate that the changes are correct using the command `fdisk -l`
12. Repeat the procedure for each disk volume.

Now that you have disk partitions, create logical volumes to use those disks.

6.10.2 Creating Logical Volumes

Now that you have disk partitions, create logical volumes to use those disks.

1. Create a physical volume on the disk partition by using the command:

   ```
   pvcreate disk_partition
   ```

   For example:

   ```
   pvcreate /dev/xvdb1
   ```

Note:

The number `1` at the end of the disk, which denotes the partition number, is the same as the values you saw in the `fdisk -l` command.

Repeat for each disk partition you created above.

2. Verify that the physical volumes have been created correctly using the following command.

   ```
   pvdisplay
   ```

3. Create a volume group, one for each virtual disk. You can create a single volume group for all disks, but this example uses one per disk.

   To create a volume group, use the following command:

   ```
   vgcreate volume_group_name disk partition
   ```

   For example:

   ```
   vgcreate volGroupSwap /dev/xvdb1
   ```

   Repeat for each volume group. For example: `volGroupTemp volGroupSwap`.

4. Validate that the volume groups have been created properly using the following:

   ```
   vgdisplay
   ```

5. Once you've created the volume groups, create a logical volume inside the volume group using the following command:

   ```
   lvcreate --name lvname --size 40G volume_group
   ```

   `size` is the size of space you wish to assign to the volume group. This equates to the size of the file system.

   For example:

   ```
   lvcreate --name Temp1 --size 2G volGroupTemp
   ```

   Repeat for each logical volume to be created.

6. Validate that the logical volumes were created successful using the following command:

   ```
   lvdisplay
   ```

6.10.3 Creating a Swap File on the New Logical Volume

This topic provides instructions to create a sawp file on the new logical volume.

1. Create a swap file using the following command

   ```
   mkswap volume_group
   ```

   For example:

   ```
   mkswap /dev/volGroupSwap/Swap1
   ```

2. Create an entry in the `/etc/fstab` directory for the new swap file. The entry will look similar to the following:

   ```
   /dev/volGroupSwap/Swap1 swap swap defaults 0 0
   ```

   Comment out the original swap entry.

3. Validate that the new swap space is being used by issuing the following commands:

   - `swapoff -a`
   - `swapon -a`
   - `swapon -s`

   You can disable the original swap using the following command:

   ```
   swapoff
   ```

   Note:

   This is not necessary as only your new swap space will be available after a reboot.

6.10.4 Moving /tmp to the New Logical Volume

If you have created a logical volume for `/tmp`, you can enable this by first creating a file system on it and then mounting it as a disk.

You do this by performing the following commands:

1. Create a file system using the command:

   ```
   mkfs.ext3 volume_name
   ```

   For example:

   ```
   mkfs.ext3 /dev/volGroupTemp/Temp1
   ```

2. Add the new file system to `/etc/fstab` so that it is automatically mounted.

   Create an entry similar to:

   ```
   /dev/volGroupTemp/Temp1 /tmp ext3 defaults 1 1
   ```

3. Mount the file system using the following command

   ```
   mount -a
   ```

4. Verify that the file system is created correctly using the following command:

   ```
   df -k
   ```

7

Preparing the Host Operating System

This chapter describes how to set up your operating system on the host, mount file systems and create installation users.

Once the environment is commissioned, you will have a number of compute nodes (physical deployment) or vServers (virtual deployment). In this chapter the vServer or compute node is generically referred to as `Host`.

Verifying Minimum Hardware Requirements for Each Host
: This topic provides information on the minimum hardware requirements required for each host.

Verifying Linux Operating System Requirements
: This topic provides information on verifying the Linux operating system requirements.

Enabling Unicode Support
: Your operating system configuration can influence the behavior of characters supported by Oracle Fusion Middleware products.

Updating DNS Settings
: Configure the host to access your corporate DNS hosts.

Configuring a Host to use a NTP (time) Server
: It is important that all hosts in the deployment have the same time. The best way to achieve this is to use a NTP (Network Time Protocol) server.

Configuring a Host to Use a NIS/YP Host
: If you are using NFS version 4 (v4), configure a directory service or a NIS (Network Information Host).

Network Routing for Multiple Networks
: Now that you have added new interfaces to each host, having only one default gateway might not be sufficient. You might want to have one interface for an Internet connection and another for a corporate WAN, for example.

Enabling Virtual IP Addresses
: The enterprise deployment requires that certain hosts, such as the WebLogic Administration hosts or SOA Managed hosts, use virtual IP addresses.

Configuring Users and Groups
: If you are using NFS 4 then make sure that the users and groups that you enter are in your NIS servers.

Mounting Shared Storage onto the Host
: This topic provides information on mounting shared storage on the host.

7.1 Verifying Minimum Hardware Requirements for Each Host

This topic provides information on the minimum hardware requirements required for each host.

To use a host in an Oracle enterprise deployment, you must verify that it meets the minimum specification described in System Requirements document.

If you are deploying to a virtual host environment, ensure that each of the virtual hosts meets the minimum requirements.

Ensure that you have sufficient local disk and that shared storage is configured as described in Preparing Storage.

Allow sufficient swap and temporary space. Specifically,

- **Swap Space**–The system must have at least 512 MB.
- **Temporary Space**–There must be a minimum of 2 GB of free space in `/tmp`.

7.2 Verifying Linux Operating System Requirements

This topic provides information on verifying the Linux operating system requirements.

Before performing an enterprise deployment, ensure that you have installed a certified operating system and that you have applied all the necessary patches for the operating system as described in *Oracle Fusion Middleware System Requirements and Specifications*.

In addition, review the following sections for information about typical Linux operating system requirements for an enterprise deployment:

Configuring Linux Kernel Parameters
: For production systems, Oracle recommends that you tune these values to optimize the performance of the system. See your operating system documentation for more information about tuning kernel parameters.

Verifying the Open File Limit on UNIX Operating Systems
: On UNIX operating systems, the open file limit is an important system setting, which can affect the overall performance of the software running on the host.

Configuring Local Hosts File
: This topic provides information on configuring the local hosts file.

Setting Huge Page Allocation
: By default, huge pages are enabled in Exalogic compute nodes. It is recommended that the Huge Page allocation be set to `25000`.

7.2.1 Configuring Linux Kernel Parameters

For production systems, Oracle recommends that you tune these values to optimize the performance of the system. See your operating system documentation for more information about tuning kernel parameters.

The kernel parameter and shell limit values shown below are recommended values only.

Kernel parameters must be set to a minimum of those below on all nodes in the topology.

The values in the following table are the current Linux recommendations. For the latest recommendations for Linux and other operating systems, see the *Oracle Fusion Middleware System Requirements and Specifications*.

Table 7-1 UNIX Kernel Parameters

Parameter	Value
kernel.sem	256 32000 100 142
kernel.shmmax	2147483648 or higher
net.ipv4.ip_nonlocal_bind 1	

Note: If the host is used to host an OTD instance, and you are going to create a listener bound to a virtual IP Address (Recommended for faster failover), then you need to set the kernel parameter net.ipv4.ip_nonlocal_bind as described above. Add the Kernel parameters (/etc/sysctl.conf file) if they are missing from your configuration.

To set these parameters:

1. Log in as `root` and add or amend the entries in the file `/etc/sysctl.conf`.
2. Save the file.
3. Activate the changes by issuing the following command:

```
/sbin/sysctl -p
```

7.2.2 Verifying the Open File Limit on UNIX Operating Systems

On UNIX operating systems, the open file limit is an important system setting, which can affect the overall performance of the software running on the host.

On all UNIX operating systems, the minimum open file limit should be 4096.

Note:

The following examples are for Linux operating systems. Consult your operating system documentation to determine the commands to be used on your system.

This section contains the following topics:

Viewing the Number of Currently Open Files
: This topic provides information on viewing the number of currently open files.

Setting the Open File Limit on Linux Operating Systems
: This topic provides information on setting the open file limit for Linux operating systems.

7.2.2.1 Viewing the Number of Currently Open Files

This topic provides information on viewing the number of currently open files.

You can see how many files are open with the following command:

```
/usr/sbin/lsof | wc -l
```

To check your open file limits, use the commands below.

C shell:

```
limit descriptors
```

Bash:

```
ulimit -n
```

7.2.2.2 Setting the Open File Limit on Linux Operating Systems

This topic provides information on setting the open file limit for Linux operating systems.

To change the open file limit on most Linux operating systems:

1. Log in as `root` and edit the following file:

   ```
   /etc/security/limits.conf
   ```

2. Add the following lines to the `limits.conf` file:

   ```
   * soft  nofile  65536
   * hard  nofile  150000
   * soft  nproc   2048
   * hard  nproc   16384
   ```

   Note:

   For information on the most recent suggested values, see the *Oracle Fusion Middleware System Requirements and Specifications*.

3. After editing the file, save your changes and reboot the machine.

7.2.2.3 Setting the Open File Limit on Oracle Linux 6

This topic provides information on setting the open file limit on Oracle Linux 6 systems.

To change the open file limit on Oracle Linux 6:

1. Log in as `root` and edit the following file:

   ```
   /etc/security/limits.conf
   ```

2. Add the following lines to the `limits.conf` file:

   ```
   * soft  nofile  65536
   * hard  nofile  150000
   ```

3. After editing the `limits.conf` file, save your changes.

4. Add the following lines to the `/etc/security/limits.d/90-nproc.conf` file:

```
* soft  nproc   2048
* hard  nproc   16384
```

> **Note:** For information on the most recent suggested values, see the *Oracle Fusion Middleware System Requirements and Specifications*.

5. After editing the `90-nproc.conf` file, save your changes and reboot the machine.

7.2.3 Configuring Local Hosts File

This topic provides information on configuring the local hosts file.

Before you begin the installation of the Oracle software, ensure that your local `/etc/hosts` file is formatted like this:

```
IP_Address       Fully_Qualified_Name       Short_Name
```

For example:

```
# Host Primary Network Interfaces
192.168.10.5 host1-int.example.com host1-int
10.10.10.5   host1-ext.example.com host1-ext
192.168.32.5 host1-stor.example.com host1-sto
192.168.10.5 host1-data.example.com host1-data

#Host Name associated with the ZFS Storage Appliance
172.17.0.9 zfsinternal.example.com zfsinternal

#Virtual Hosts
10.10.30.3   adminvhn.example.com adminvhn
192.168.30.5 host1vhn1.example.com host1vhn1
192.168.30.6 host2vhn1.example.com host2vhn1

#OTD Failover Groups
192.168.50.1 idstore.example.com idstore
192.168.50.2 edginternal.example.com edginternal
```

Each host file should contain the entries for non-DNS registered IP addresses that are used in the deployment topology.

7.2.4 Setting Huge Page Allocation

By default, huge pages are enabled in Exalogic compute nodes. It is recommended that the Huge Page allocation be set to `25000`.

To verify the existing allocation, run the following command as root:

```
grep Huge /proc/meminfo
```

Specify the number of large pages. In the following example 3 GB of a 4 GB system are reserved for large pages (assuming a large page size of 2048k, then 3g = 3 x 1024m = 3072m = 3072 * 1024k = 3145728k, and 3145728k / 2048k = 1536).

To set the Huge Page allocation, run the following command as `root` in the compute node:

```
echo 1536  > /proc/sys/vm/nr_hugepages
```

Note: To make use of huge pages in a Java vm you need to add the following to the arguments field of the web logic managed server: `-XX:+UseLargePages`.

7.3 Enabling Unicode Support

Your operating system configuration can influence the behavior of characters supported by Oracle Fusion Middleware products.

On UNIX operating systems, Oracle highly recommends that you enable Unicode support by setting the `LANG` environment variable to a locale with the UTF-8 character set. For example,

```
LANG=en_GB.UTF-8
```

This enables the operating system to process any character in Unicode. Oracle SOA Suite technologies, for example, are based on Unicode.

If the operating system is configured to use a non-UTF-8 encoding, Oracle Fusion Middleware Suite components might function in an unexpected way. For example, a non-ASCII file name might make the file inaccessible and cause an error. Oracle does not support problems caused by operating system constraints.

7.4 Updating DNS Settings

Configure the host to access your corporate DNS hosts.

To do this, update the DNS settings in the`/etc/resolv.conf` file.

7.5 Configuring a Host to use a NTP (time) Server

It is important that all hosts in the deployment have the same time. The best way to achieve this is to use a NTP (Network Time Protocol) server.

To configure a host to use a NTP server, perform the following steps:

1. Determine the name of the NTP server(s) you want to use.

   Note:

   For security reasons, it is recommended that these servers are inside your organization.

2. Log into the host as the `root` user.
3. Edit the `/etc/ntp.conf` file to include a list of the time servers.

   After editing the `ntp.conf` file, the file should look like the following example:

   ```
   host ntpserver1.example.com
   host ntpserver2.example.com
   ```

4. Run the command to synchronize the system clock with the NTP server:

```
/usr/sbin/ntpdate ntpserver1.example.com
/usr/sbin/ntpdate ntpserver2.example.com
```

5. Start the NTP client by using the following command:

   ```
   service ntpd start
   ```

6. Validate that the time is set correctly using the `date` command.

7. To ensure that the host always uses the NTP server to synchronize time, set the client to start on reboot by using the following command:

   ```
   chkconfig ntpd on
   ```

7.6 Configuring a Host to Use a NIS/YP Host

If you are using NFS version 4 (v4), configure a directory service or a NIS (Network Information Host).

If your organization does not have one already, use the built-in one on the ZFS storage appliance. For more information, see Creating Users and Groups in NIS.

Once you have configured your NIS host, configure each compute node or vserver to use it. If you are using the built-in NIS host on the Exalogic ZFS appliance, perform the following steps:

1. Determine the name of the NIS host by logging into the storage BUI using the following URL:

   ```
   https://exalogicsn01-priv:215
   ```

2. Click **Configuration**, **Services**, and then **NIS**.

3. Make a note of one of the listed NIS hosts.

4. Log into the host as `root`.

5. Open the `/etc/idmapd.conf` configuration file.

6. In the `idmapd.conf` file, set the domain value as shown in the following example:

   ```
   Domain = example.com
   ```

7. Restart the `rpcidmapd` service by running the following command:

   ```
   service rpcidmapd restart
   ```

8. Restart the `rpcbind` service by running the following command:

   ```
   service rpcbind restart
   ```

   Note: If the `rpcbind` service is not started already then start the same with the following command:

   ```
   service rpcbind start
   ```

9. Open the `/etc/yp.conf` configuration file.

10. In the `yp.conf` file, add the following line to set the correct domain value:

    ```
    domain example.com server NIS_Host_hostname_or_IP
    ```

Where *example*.com is the example domain and *NIS_Host_hostname_or_IP* is the host name or IP address of the NIS host. You must replace these sample values with values appropriate for your environment.

11. Set the NIS domain name on the command line:

```
domainname NIS_DOMAIN_NAME
```

For example:

```
domainname example.com
```

12. Open the /etc/nsswitch.conf configuration file, and change the following entries::

```
passwd:     files nis
shadow:     files nis
group:      files nis
automount:  files nis nisplus
aliases:    files nis nisplus
```

13. Restart the rpcidmapd service by running the following command:

```
service rpcidmapd restart
```

14. Restart the ypbind service by running the following command:

```
service ypbind restart
```

15. Check the yp service by running the following command:

```
ypwhich
```

16. Verify if you can access Oracle user accounts:

```
ypcat passwd
```

17. Add ypbind to your boot sequence so that it starts automatically after rebooting.

```
chkconfig ypbind on
```

Note: Check the service rpcbind and rpcidmapd are configured to start at machine boot.

7.7 Network Routing for Multiple Networks

Now that you have added new interfaces to each host, having only one default gateway might not be sufficient. You might want to have one interface for an Internet connection and another for a corporate WAN, for example.

In the example below, the different interfaces are shown, along with example IP addresses and gateway requirements:

Interface	IP Address	Gateway Requirements
eth0	201.19.23.128 / 24	Gateway IP 201.19.23.1
bond0	192.168.10.1 / 24	No Gateway requirements

Interface	IP Address	Gateway Requirements
bond1	10.10.10.101/ 24	Gateway IP 10.10.10.1

As you can see, eth0 and bond1 must have their own respective default gateways.

Bond0, however, does not have any default gateway requirements. It is simply confined to their actual Layer 3 subnet.

To get around this, create rules and tables for routing lookups, as follows.

1. Check the existing table IDs by issuing this command:

   ```
   ip rule list
   ```

2. Choose a unique id that has not already been used. In this example, 224 and 225 will be used.

3. For eth0, create the following two files:

   - The file `/etc/sysconfig/network-scripts/rule-eth0`, which contains:

     ```
     from 201.19.23.128/24 table 224
     to 201.19.23.128 table 224
     ```

   - The file `/etc/sysconfig/network-scripts/route-eth0`, which contains:

     ```
     201.19.23.0/24 dev eth0 table 224
     default via 201.19.23.1 dev eth0 table 224
     ```

4. For bond1, create the following two files:

   - The file `/etc/sysconfig/network-scripts/rule-bond1`, which contains:

     ```
     from 10.10.10.10/24 table 225
     to 10.10.10.10 table 225
     ```

   - The file `/etc/sysconfig/network-scripts/route-bond1`, which contains:

     ```
     10.10.10.0/24 dev bond1 table 225
     default via 10.10.10.1 dev bond1 table 225
     ```

5. Restart the network to make the configuration effective.

   ```
   service network restart
   ```

   The hosts are now accessible from both routers.

7.8 Enabling Virtual IP Addresses

The enterprise deployment requires that certain hosts, such as the WebLogic Administration hosts or SOA Managed hosts, use virtual IP addresses.

You must enable the appropriate IP address on each host. Preparing the Network describes the mapping of IP addresses to the hosts.

The following tables list the assignment of virtual IP addresses to network interfaces for the typical enterprise deployment topology. Managed hosts not using a virtual IP address use the default network interface for communications.

When you configure networking, you will by default create a default network interface that will be assigned to the network card. For example, bond1 for the client EoIB network and bond0 for the internal IPoIB network.

If you want to then assign another IP address to the same network card, then an index number is applied. For example, assigning VIP1 and VIP2 to bond0 would result in bond0:1 and bond0:2.

Note:

Refer to the appropriate product-specific Enterprise Deployment Guide for the list of VIP addresses you need to create for your deployment.

This section contains the following topics:

Summary of Exalogic Virtual IP Addresses
: This topic provides a summary of Exalogic virtual IP address.

Enabling a Virtual IP Address on a Network Interface
: This topic provides the procedure to enable a virtual IP address on a network interface.

7.8.1 Summary of Exalogic Virtual IP Addresses

This topic provides a summary of Exalogic virtual IP address.

Table 7-2 shows the Virtual IP address mapping for a typical enterprise deployment on Exalogic.

For instructions on defining these virtual IP addresses, see Enabling a Virtual IP Address on a Network Interface.

Table 7-2 Virtual IP Addresses Associated with IPoIB and EoIB Network interfaces

Interface	Address Example	Netmask Example	Used By	Virtual Host Name	Default Physical Host	Default Virtual Host
BOND1:1	10.10.30.1	255.255.224.0	OTD Administration Host	OTDADMINVHN	HOST1	WEBHOST1
BOND1:1	10.10.30.2	255.255.224.0	Administration Host	ADMINVHN	HOST1	HOST1
BOND0:1	192.168.30.1	255.255.240.0	WLS_PRODA1	HOST1VHN1	HOST1	HOST1
BOND0:1	192.168.30.2	255.255.240.0	WLS_PRODA2	HOST2VHN1	HOST2	HOST2
BOND0:2	192.168.30.3	255.255.240.0	WLS_PRODB1	HOST1VHN2	HOST1	HOST1

Table 7-2 (Cont.) Virtual IP Addresses Associated with IPoIB and EoIB Network interfaces

Interface	Address Example	Netmask Example	Used By	Virtual Host Name	Default Physical Host	Default Virtual Host
BOND0:2	192.168.30.4	255.255.240.0	WLS_PRODB 2	HOST2VHN2	HOST2	HOST2
	192.168.50.1	255.255.224.0	OTD Failover group for callbacks	EDGINTERN AL	HOST1	WEBHOST1
	192.168.50.2	255.255.224.0	OTD Failover group for LDAP	IDSTORE	HOST2	WEBHOST2

Default Physical Host is the compute node that the Virtual Host is assigned to by default. It will only move in the event of host failure.

Default Virtual Host is the virtual host that the Virtual Host is assigned to by default. It will only move in the event of host failure.

In the example above, the Administration host is listening on the External IPoIB network. It could also be on the internal network.

Note:

The virtual IP addresses used here are examples. You should use the IP addresses you reserved in Reserving Virtual IP Addresses.

7.8.2 Enabling a Virtual IP Address on a Network Interface

This topic provides the procedure to enable a virtual IP address on a network interface.

1. Use the `ifconfig` command to create the virtual IP address.

   ```
   ifconfig subinterface virtual_ip_address netmask netmask_value
   ```

   For example, on HOST1, enter the following:

   ```
   ifconfig bond0:1 192.168.30.2 netmask 255.255.240.0
   ```

2. For each virtual IP address you define, update the ARP caches using the following command:

   ```
   arping -b -A -c 3 -I bond0 192.168.30.2
   ```

   This section contains the following topics:

   Oracle Linux 5
   If you are using Oracle Linux 5, complete the following steps to enable the virtual IP addresses.

   Oracle Linux 6 and Onwards
   In Oracle Linux 6, the `ifconfig` command described above has been deprecated and is replaced with the `ip` command.

7.8.2.1 Oracle Linux 5

If you are using Oracle Linux 5, complete the following steps to enable the virtual IP addresses.

1. Use the `ifconfig` command to create the virtual IP address:

   ```
   ifconfig subinterface virtual_ip_address netmask netmask_value
   ```

   For example, to enable the IP address 192.168.20.3, net mask 255.255.240 on network card bond0, use the following command:

   ifconfig bond0:1 192.168.20.3 netmask 255.255.240.0

   Note: Example in this section is applicable for both physical and virtual Exalogic deployments.

2. For each virtual IP address you define, update the ARP caches using the following command:

   ```
   arping -b -A -c 3 -I bond0 192.168.20.3
   ```

   This command does not return any response.

7.8.2.2 Oracle Linux 6 and Onwards

In Oracle Linux 6, the `ifconfig` command described above has been deprecated and is replaced with the `ip` command.

To enable the virtual IP addresses, complete the following steps:

1. Determine the CIDR notation of the netmask. Each Netmask has a CIDR notation. For example, `255.255.240.0` has a CIDR of 20.

   If the netmask you are adding is the same as the interface, the fastest way to determine this is to examine the existing IP address assigned to the network card. You can do this using the following command:

   ```
   ip addr show dev bond0
   ```

   The following is a sample output:

   ```
   2: bond0: BROADCAST,MULTICAST,UP,LOWER_UP mtu 1500 qdisc pfifo_fast qlen 1000
   link/ether 00:21:f6:03:85:9f brd ff:ff:ff:ff:ff:ff
   int 192.168.20.1/20 brd 10.248.11.255 scope global bond0
   ```

   In this example, the CIDR value is the value after /, that is, 20. If you are unsure of the CIDR value, contact your network administrator.

2. Add the IP address `192.168.20.3`, net mask `255.255.240` (CIDR20) on network card `bond0` using the following command:

   ```
   ip addr add 192.168.20.3/20 dev bond0:1
   ```

3. For each of the virtual IP addresses you define, update the ARP caches using the following command:

   ```
   arping -b -A -c 3 -I bond0 192.168.20.3
   ```

7.9 Configuring Users and Groups

If you are using NFS 4 then make sure that the users and groups that you enter are in your NIS servers.

Do not create them locally. If you are using NFSv3 then you can create your users locally. Create the following users and groups either locally or in your NIS or LDAP host. This user is the Oracle software owner.

The instructions below are for creating users and groups locally. Refer to your NIS documentation for information about creating these users/groups in your NIS host.

This section contains the following topics:

7.9.1 Creating Users and Groups Locally

This topic provides information on creating users and groups.

Refer to creating groups and creating users in the following sections.

7.9.1.1 Creating Groups

This topic provides the procedure to create groups on each node.

You must create the following groups on each node.

- `oinstall`
- `dba`

To create the groups, run the following command as `root`:

```
groupadd groupname
```

For example:

```
groupadd -g 500 oinstall
groupadd -g 501 dba
```

7.9.1.2 Creating Users

This topic provides procedure for creating users.

You must create the following user on each node.

- `oracle` – The owner of the Oracle software. You can use a different name. The primary group for this account must be `oinstall`.

  The account must also be in the `dba` group.

Note:

- The group `oinstall` must have write privileges to all the file systems on shared and local storage that are used by the Oracle software.
- Each group must have the same group ID on every node.
- Each user must have the same user ID on every node.

To create users, run the following command as `root`:

```
useradd -g primary_group -G optional_groups -u userid username
```

For example:

```
useradd -g oinstall -G dba -u 500 oracle
```

7.9.2 Creating Users and Groups in NIS

To create an account for an NIS user on the NIS master server.

- NIS provides maps for password, group, and auto-home.
- The NIS master server uses NFS to export the users home directories.

WARNING: NIS authentication is deprecated as it has security issues, including a lack of protection of authentication data.

1. If the NIS master server does not export the base directory of the users home directories, perform the following steps on the NIS master server:
   a. Create the base directory for user directories, for example `/nethome`:

      ```
      # mkdir /nethome
      ```

   b. Add an entry such as the following to `/etc/exports`:

      ```
      /nethome *(rw,sync)
      ```

      You might prefer to restrict which clients can mount the file system.

      For example, the following entry allows only clients in the 192.168.1.0/24 subnet to mount `/nethome`:

      ```
      /nethome 192.168.1.0/24(rw,sync)
      ```

   c. Use the following command to export the file system:

      ```
      # exportfs -i -o ro,sync *:/nethome
      ```

   d. If you have configured `/var/yp/Makfile` to make the auto-home map available to NIS clients, create the following entry in `/etc/auto.home`:

      ```
      * -rw,sync nissvr:/nethome/&
      ```

      where nissvr is the host name or IP address of the NIS server.
2. Create the user account.

```
# useradd -b /nethome username
```

The command updates the `/etc/passwd` file and creates a home directory on the NIS server.

3. Depending on the type of authentication that you have configured:

   - For Kerberos authentication, on the Kerberos server or a client system with kadmin access, use **kadmin** to create a principal for the user in the Kerberos domain, for example:

     ```
     # kadmin -q "addprinc username@KRBDOMAIN"
     ```

     The command prompts you to set a password for the user, and adds the principal to the Kerberos database.

   - For NIS authentication, use the password command:

     ```
     # passwd username
     ```

     The command updates the `/etc/shadow` file with the hashed password.

4. Update the NIS maps.

   ```
   # make -C /var/yp
   ```

   This command makes the NIS maps that are defined for the all target in `/var/yp/Makefile`. If you have configured NOPUSH=false in `/var/yp/Makefile` and the names of the slave servers in `/var/yp/ypservers`, the command also pushes the updated maps to the slave servers.

Note: A Kerberos-authenticated user can use either kpasswd or passwd to change his or her password. An NIS-authenticated user must use the yppasswd command rather than passwd to change his or her password.

7.10 Mounting Shared Storage onto the Host

This topic provides information on mounting shared storage on the host.

As described in Preparing Storage, you must make shared storage available to each host that will use it.

This section contains the following topics:

7.10.1 Shared Storage Overview

This topic provides information on the shared storage.

Mount the shared storage to the hosts according to the following table.

Table 7-3 Mapping the Shares on the Appliance to Mount Points on Each Compute Node

Volume Mounted	Mounted on Physical Host	Mounted on Virtual Host	Mounted Point	Exclusive
`/export/product_binaries/shared_binaries`	HOST1 HOST2	vServers on HOST1 vServers on HOST2	`/u01/oracle/products`	No
`/export/product_binaries/webhost1_local_binaries`	HOST1	WEBHOST1	`/u01/oracle/products`	Yes
`/export/product_binaries/webhost2_local_binaries`	HOST2	WEBHOST2	`/u01/oracle/products`	Yes
`/export/product_config/shared_config`	HOST1 HOST2	vServers on HOST1 vServers on HOST2	`/u01/oracle/config`	No
`/export/runtime/shared_runtime`	HOST1 HOST2	vServers on HOST1 vServers on HOST2	`/u01/oracle/runtime`	No
`/export/product_config/host1_local_config`	HOST1	vServers on HOST1	`/u02/private/oracle/config`	Yes
`/export/product_config/host2_local_config`	HOST2	vServers on HOST2	`/u02/private/oracle/config`	Yes

Note the following points:

- Each host must have appropriate privileges set within the NAS or SAN so that it can write to the shared storage.
- Temporary mounts are only required during provisioning and patching.
- If WEBHOST1 and WEBHOST2 are in the DMZ, `SW_ROOT` is not shared between those two hosts.
- The mount point should be owned by the user and group created in Configuring Users and Groups.
- Follow the best practices of your organization for mounting shared storage. This section provides an example of how to do this on UNIX or Linux using NFS storage.
- The user ID used to create a shared storage file system owns and has read, write, and execute privileges for those files. Other users in the operating system group can read and execute the files, but they do not have write privileges. For more information about installation and configuration privileges, see the "Understanding Installation and Configuration Privileges and Users" section in the *Oracle Fusion Middleware Installation Planning Guide*.

7.10.2 Mounting Shared Storage

You must create and mount shared storage locations so that each application tier host can see the same location for the binary installation.

You use the following command to mount shared storage from the ZFS storage device to a Linux host.

To mount shared storage on a host, use a command similar to the following:

```
mount -t nfs zfs:volume mountpoint
```

For example:

```
mount -t nfs zfsinternal:/export/product_binaries/shared_binaries /u01/oracle/
products
```

Using the `mount` command mounts the shared storage until the host is rebooted. Once rebooted, the storage must be remounted to the host.

To ensure that storage is made available following a host reboot, place an entry into the file `/etc/fstab` that looks like the following:

```
zfsinternal:/export/product_binaries/shared_binaries /u01/oracle/products nfs4
nointr,timeo=300
```

7.10.3 Validating the Shared Storage Configuration

This topic provides the steps to validate the shared storage.

1. Ensure that you can read and write files to the newly mounted directories by creating a test file in the shared storage location you just configured.

   For example:

   ```
   cd /u01/oracle/products
   touch testfile
   ```

2. Verify that the owner and permissions are correct:

   ```
   ls -l testfile
   ```

3. Remove the file:

   ```
   rm testfile
   ```

Part III

Managing an Exalogic Appliance

This topic provides information on managing a topology and monitoring the topology using Oracle Enterprise Manager cloud control.

8

Managing a Topology

Exalogic Startup and Shutdown Procedure
: It is important to follow the proper sequence in order to startup and shutdown Exalogic and its components.

Maintenance Procedures
: Maintenance procedures provides information about Lifecycle Management Tools, ExaChk, ExaLogs and Patching.

Backup and Recovery Procedures
: This provides guidelines for Exalogic system backup and recovery procedures.

8.1 Exalogic Startup and Shutdown Procedure

It is important to follow the proper sequence in order to startup and shutdown Exalogic and its components.

Refer to the startup sequence and shutdown procedure, ZFS Storage Appliance Power On and Off Procedure and procedures to Start up or Shutdown Exalogic, Control Stack, and Guest vServers.

This section contains the following topics:

Exalogic Startup Sequence
: Startup using the following sequence of steps.

Exalogic Shutdown Sequence
: Shutdown using the following sequence of steps.

ZFS Storage Appliance Power On and Off Procedure
: To power off the ZFS appliance, perform the following sequence on the storage nodes.

Procedures to Start up or Shutdown Exalogic, Control Stack, and Guest vServers
: Refer to the specific sections of startup and shutdown Exalogic machine, control and guest virtual servers for further information.

8.1.1 Exalogic Startup Sequence

Startup using the following sequence of steps.

1. Power on the PDUs of the Exalogic rack
2. Network switches

Note: Ensure that the switches have the power applied for a few minutes to complete the power-on configuration before starting the storage nodes and compute nodes. On the rear side of the InfiniBand Gateway switch (the side on which the InfiniBand cables are plugged in), there are status LEDs. The OK LED on the right bottom (above the USB port) must be steady green. This implies that the gateway is functional without any fault. You can also SSH one of the InfiniBand switches and ensure that ibswitches shows all the InfiniBand Gateway switches.

3. Storage nodes
4. Compute nodes
5. Exalogic Control stack vServers and services (If Using virtual Exalogic)
6. All guest vServers (If Using virtual Exalogic)
7. User application services

8.1.2 Exalogic Shutdown Sequence

Shutdown using the following sequence of steps.

1. All user application services
2. All guest vServers (If Using virtual Exalogic)
3. All control Stack services (If Using virtual Exalogic)
4. Exalogic Control vServers (If Using virtual Exalogic)
5. Power off the host (OVS) of all compute nodes and storage nodes
   - First shutdown and power off the standby node
   - Power off the active storage node
6. Network switches
7. PDUs

8.1.3 ZFS Storage Appliance Power On and Off Procedure

To power off the ZFS appliance, perform the following sequence on the storage nodes.

1. Shutdown and power off the stand-by storage node.
2. Shutdown and power off the active node.

   This avoids unnecessary failover of network and storage back and forth.
3. Power Off:
   - From CLI: maintenance system power off
   - From BUI: Click the Power off appliance icon
4. Power On: To power on the storage nodes:

- Push controller power button
- ILOM: Start /SYS

8.1.4 Procedures to Start up or Shutdown Exalogic, Control Stack, and Guest vServers

Refer to the specific sections of startup and shutdown Exalogic machine, control and guest virtual servers for further information.

References to Startup and Shutdown Exalogic Machine

- Refer to section Operational Procedures for Exalogic Machine in *Oracle Exalogic Elastic Cloud Machine Owner's* Guide at .
  - Non-emergency power procedure
  - Emergency power-off considerations
  - Cautions and warnings
- For more information, see My Oracle Support document ID 1533391.1 Steps To Shut Down or Power Off, and Start Up or Power On an Exalogic Machine.

References to Startup and Shutdown Exalogic Control

- ExaBR provides convenient way to stop and start the control stack. You must first install the Exalogic Lifecycle toolkit. Refer to Section *Lifecycle Management Tools* for more information.
- For manual procedure, see My Oracle Support document ID 1594223.1 How To Stop and Start the Entire Exalogic Control Stack In An Exalogic EECS v2.0.6.0.0 and later Virtual releases.

Note: To open a master note, perform the following steps:

- Select My Oracle Support document ID, and press Ctrl + F9. The Attributes dialog opens.
- In the **Attribute Value** field for the **Url** attribute, enter this URL:

  https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=note_id
- Enter the Note ID or keyword in the search field at the top of the screen.
- Click **Set Value**.

References to Startup and Shutdown Guest Virtual Servers

- Always start and stop the guest vServers via Enterprise Manager Ops Center. Use either the BUI or IaaS CLI as described in *Oracle Exalogic Elastic Cloud Administrator's* Guide at http://docs.oracle.com/cd/E18476_ 01/doc.220/e25258/proc.htm#BABDCBHC

Note: It is important to start and stop vServers under vDC Management rather than from the Assets accordion in the EMOC UI.

- Do not use the xm commands.
- Do not use OS level shutdown command. Otherwise, vServers that are marked HA will be restarted by EMOC automatically.

8.2 Maintenance Procedures

Maintenance procedures provides information about Lifecycle Management Tools, ExaChk, ExaLogs and Patching.

Detailed information on each section is provided in the following topics.

Lifecycle Management Tools
: Oracle Exalogic Lifecycle (ELLC) toolkit is a collection of tools that simplify, automate, and standardize lifecycle management on an Oracle Exalogic Elastic Cloud machine.

ExaChk
: Exachk is a health-check tool that is designed to audit important configuration settings in an Exalogic machine.

ExaLogs
: ExaLogs is a command-line tool for gathering logs, diagnostics, environment and configuration information and other data from key components in an Exalogic physical or virtual configuration.

Patching
: Oracle Elastic Exalogic Cloud Software Recommended Patches are available within My Oracle Support.

Troubleshooting and Action Plan
: The MOS note provides information on common Exalogic outages and restoration steps to recover from those outages for Exalogic Platinum users.

8.2.1 Lifecycle Management Tools

Oracle Exalogic Lifecycle (ELLC) toolkit is a collection of tools that simplify, automate, and standardize lifecycle management on an Oracle Exalogic Elastic Cloud machine.

For more information, see My Oracle Support document ID 1912063.1 Exalogic Lifecycle Toolkit Release 14.2.

Exalogic Tools	New Features and Enhancements
EMAgent PreSetup	A new tool to prepare the Exalogic rack for Enterprise Manager 12c discovery and monitoring.
ExaBR	EECS 2.0.4 Support STIG-hardened Linux Compute Nodes All-ILOM Target
ExaPatch	Improved platform patching

Exalogic Tools	New Features and Enhancements
ExaLogs	Solaris Support Credentials (access) option Network Usage Order
ExaPasswd	A new tool to automate password changes to Exalogic system components
STIGfix	A new tool to make Exalogic guest vServers and Physical Linux Nodes STIG compliant
ModifyLVMImg	A new tool to resize LVM-based vServers (root/swap volumes), and add or remove Linux RPMs
ExaChk	Enhanced Exalogic Health Check tool Support for No DNS Revised scoring Diff comparison

8.2.2 ExaChk

Exachk is a health-check tool that is designed to audit important configuration settings in an Exalogic machine.

- Runs every quarter according to the PSU cycle
- Before and after a maintenance
- Attach Exachk report to the Service Request and save time
- On a scheduled basis for comparison

For more information, see My Oracle Support document ID:

- Exachk Health-Check Tool for Exalogic.
- Exalogic Exachk Diagnostic Information and Suggested Actions.
- Exalogic Exachk Health-Check Tool Known Issues.

Note: To open a master note, perform the following steps:

- Select My Oracle Support document ID, and press Ctrl + F9. The Attributes dialog opens.
- In the **Attribute Value** field for the **Url** attribute, enter this URL:

  https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=note_id
- Enter the Note ID or keyword in the search field at the top of the screen.
- Click **Set Value**.

8.2.3 ExaLogs

ExaLogs is a command-line tool for gathering logs, diagnostics, environment and configuration information and other data from key components in an Exalogic physical or virtual configuration.

- For more information, see My Oracle Support document ID 1912063.1Exalogic Lifecycle Toolkit Release 14.2.
- Run ExaLogs before and after patching (PSUs), upgrades, or when a problem arises.
- ExaLogs are required to open up an SR.

8.2.4 Patching

Oracle Elastic Exalogic Cloud Software Recommended Patches are available within My Oracle Support.

Exalogic Patch Set Updates (PSU) are a collection of Oracle recommended patches. The PSUs are cumulative and released on quarterly schedule.

This section contains the following topic:

8.2.4.1 Patching Recommendation

To ensure that the Exalogic system continues to perform optimally, Oracle periodically provides comprehensive and well-tested patches to the system as a whole.

An Exalogic Patch Set Update (PSU) is released quarterly (January, April, July, and October) with the following features:

- PSU is a single download that contains patches for all Exalogic components (firmware, software and OS) as necessary.
- PSU is a highly recommended update for all Exalogic customers.
- In addition to patches or updates for the Exalogic Infrastructure components (on-node components such as Operating System, ILOM, InfiniBand and RAID controller cards, and off-node components such as InfiniBand switches and ZFS Storage Appliance), patches for Middleware components (WLS, Coherence, JDK) are also included.
- PSU contains optional patches for the guest OS image. They can be applied, when the schedule allows.

Exalogic users should ensure that they align their systems with Oracle's Exalogic releases and recommended patch levels, and should refrain from applying patches which are outside of recommendations for Exalogic. For instance, if a new version of the Oracle ZFS Storage Appliance software is released and is not part of an Exalogic recommended patch or PSU, you must not update the racks with the patch. Applying patches that are not recommended can adversely affect not only the functionality but also the performance of the Exalogic system.

For systems that are in production or in late testing stages before production:

- Plan to periodically adopt more current patch releases
- Not required or necessary to install every new patch release
- A patch should be installed on a production system only after it is validated in a proper test environment
- Systems that are in the early stages of testing before production or proof-of-concept should adopt new releases and patches when they are made available as indicated in the following MOS Note:
  - For more information, see My Oracle Support document ID 1368307.1 Oracle Exalogic Elastic Cloud Supported System Configurations.
  - Oracle Internal Sun EXALOGIC X2-2 , X3-2 , X4-2 Current Product Patches & Firmware provides version numbers of all software and firmware components.

8.2.5 Troubleshooting and Action Plan

The MOS note provides information on common Exalogic outages and restoration steps to recover from those outages for Exalogic Platinum users.

Each outage is categorized as either partial or complete outage. The MOS note also provides information about troubleshooting steps to debug the problem and post issue Root Cause Analysis (RCA) data collection needed for root cause analysis of outage.

For more information, see My Oracle Support document ID 1492461.1 Exalogic Platinum Customer Outage Classifications and Restoration Action Plans.

8.3 Backup and Recovery Procedures

This provides guidelines for Exalogic system backup and recovery procedures.

Whilst hot standby systems are extremely useful for business continuity, they are expensive to maintain and need additional infrastructure. In some circumstances such as simple user errors, it may be quicker to fix the issue than to failover to the DR system especially when DNS needs updating.

Taking regular backups of a system is part of standard operating procedure for most production systems and is done irrespective of whether or not the site has a disaster recovery solution. It allows the flexibility to restore individual files should something happen to the original or the system as whole. Backups can also be stored off site in a secure location. Backups on Exalogic can be within the Exalogic system, disk-to-disk and disk-to-tape.

The data contained in an Oracle Exalogic Machine which needs to be backed up consists of:

- Exalogic Operating System
- Software Binaries
- Configuration Information
- Transactional Data, such as transaction logs as JMS queues
- Switch Configuration

- Other Artifacts which are stored on the disk. These objects can be backed up to:
  - Disk within the same storage appliance
  - Disk on a remote machine, which utilizes the same storage type (ZFS)
  - Disk on a remote machine, which utilizes a different storage type
- Tape

Backup and Recovery Concepts
Volatility

Objects can be grouped by volatility. For example, the operating system changes very infrequently and therefore does not need backing up as frequently as transactional data, which changes on a frequent basis. In a typical Exalogic deployment objects can be grouped into the following categories:

Volatility Groups	Volatility Example Objects
Low	Oracle Binaries
	Operating System
Medium	Configuration Information - WLS Domain
	Oracle Instance
High	File based JMS Queues
	Persistent Stores

Backup Frequency

The volatility of the data can be used to determine the backup frequency. In addition to volatility the following may impact the frequency in which data is backed up:

- Volume of data to be backed up
- Available backup windows
- Regulatory requirements

Using the above volatility groups, the following is a sensible backup frequency.

Table 8-1 Backup Schedule

Volatility Group	Backup Frequency
Low	Monthly
Medium	Weekly
High	Daily

In addition to the scheduled backups, it makes sense to perform ad-hoc backups when major events occur. For example, it is appropriate to take an additional backup of the Oracle binaries after patching or upgrade.

Retention Periods

In determining a backup strategy, you need to factor how long you wish to keep the backups for. This is mainly dependent on your business and regulatory requirements. Using the examples above the following may be appropriate values.

Table 8-2 Retention Periods

Volatility Group	Retention Periods
Low	3 years
Medium	6 months
High	7 days

Backup Types

There are two different types of backups available, full backups and incremental backups. A full backup backs up the entire file system as it is at that moment in time. An incremental backup backs up only the data that has changed since the last backup. Incremental backups can be either cumulative or differential. A cumulative backup backs up all the changes since the last full backup, whilst a differential backup backs up the changes since the last differential – differential backups are not supported on ZFS storage appliances, they are however widely available when backing up to tape.

Incremental backups can be leveled. You can perform a level 0 (full backup) each month, a level 1 (incremental) each Sunday and a level 2 (incremental) each weekday. If this type of strategy is implemented then only the data that has changed since the last level -1 backup is backed up. For example on Tuesday, the data backed up is the data, which has changed since the last level 1 backup that was taken on the previous Sunday. Incremental backups are useful when the volume of data to be backed up is significant.

The advantage of a full backup is that the backup contains all of the information required to perform a restore. In an incremental backup strategy, a restore is likely to use several backups. In the 3 level backup strategy above you would need, the Last level 0 backup, plus the last level 1 backup taken, plus the last level 2 backup taken. If the volume of data to be backed up is small, then it may be easier to perform a full backup each time rather than an incremental one, but this will be determined by the volumes of data being backed up. Incremental backups are supported by Oracle Secure Backup and the operating system dump command.

Recovery Point Objective (RPO) and Recovery Time Objective (RTO) determine the frequency of backups. They are critical factors in an effective business continuity plan. Refer to the following documents for additional information:

- For more information, see My Oracle Support document ID 1546047.1 Backup and Recovery Guide For Exalogic Elastic Cloud Software.
- Backup and Recovery using ExaBR.
- Oracle Exalogic Elastic Cloud Backup and Recovery Guide Release EL X2-2 and X3-2 (E40226).
- ZFS Snapshot for Backup and Recovery of the Exalogic Control Repository and Stack
- ZFS Snapshots for back up of shares on the ZFS Storage Appliance

- Taking manual snapshots (BUI)
- Create a project level snapshot
- Create a share/LUN level snapshot
- Destroying a snapshot (BUI)
- Rolling back to a snapshot (BUI)
- Cloning a snapshot (BUI)
- Scheduled snapshots (BUI)
- Command Line Interface (CLI)
- Listing snapshots (CLI)
- Taking manual snapshots (CLI)
- Renaming a snapshot (CLI)
- Destroying a snapshot (CLI)
- Rolling back to a snapshot (CLI)
- Cloning a snapshot (CLI)
- Scheduled snapshots (CLI)
- Setting the scheduled snapshot label (CLI)

9

Monitoring the Topology Using Oracle Enterprise Manager Cloud Control

Oracle Enterprise Manager Cloud Control 12*c* with Oracle WebLogic Server Management Pack Enterprise Edition's capabilities include Exalogic Elastic Cloud-specific management tools to manage and monitor Oracle software deployed in physical and virtual Exalogic Elastic Cloud environments.

These monitoring capabilities expand on the existing Oracle WebLogic Server management features that span the following:

- Application performance management
- Configuration management
- Service-level management and operations

The Exalogic Elastic Cloud-specific features enable administrators to:

- Access full management capabilities with dedicated dashboards for Oracle Exalogic Elastic Cloud targets to easily monitor overall health, availability, and performance and manage Exalogic hardware and software, including hosts, Oracle WebLogic domains, application deployments, and Coherence clusters running on Oracle Exalogic Elastic Clouds.
- Drill down into application deployments to identify metrics and set thresholds.
- Display alerts, status, and incidents hierarchically across the Exalogic Elastic Cloud environment and provide operations with notifications in relation to service levels and thresholds.
- Drill down from the Exalogic Elastic Cloud dashboard and underlying menus into the detailed component and JVM-level performance metrics, configuration management, and provisioning and cloning that provide end-to-end management for operations and administrators.
- Use the Exalogic Topology View to immediately identify cross-tier relationships between components, middle-tier platforms, and the underlying hosts and hardware that make up an Exalogic Elastic Cloud system.
- Make use of additional Exalogic dashboard enhancements, including:
  - Integrated hardware and software schematics
  - Hardware-software topology views
  - Hardware targets monitoring (Compute, ZFA appliance, Infiniband Fabric, ILOM)
  - Oracle Traffic Director (OTD) monitoring

  - Support for virtual and non-virtual configurations
  - Monitoring of Exalogic vServer guest Virtual Machines (VMs)
  - Health checks
  - Trusted partition/Virtual Central Processing Unit (vCPU) licensing report
  - Performance and filtering optimizations
- Manage hardware and software incidents from a single pane of glass using Enterprise Manager Cloud Control's incident management system, including hardware alarms and alerts from Enterprise Manager Ops Center.

This chapter includes the following topics:

Accessing Oracle Enterprise Manager Cloud Control 12c
: Refer to this topic for information on accessing the Oracle Enterprise manager cloud control 12*c*.

Discovering an Oracle Exalogic Elastic Cloud Target
: To monitor and manage your Oracle Exalogic Elastic Cloud machine using Oracle Enterprise Manager Cloud Control 12*c*, you must prepare the Exalogic Elastic Cloud environment for discovery and then discover the related system targets.

Using Exalogic-Specific Pages in Oracle Enterprise Manager Cloud Control 12c
: Refer to the topic for procedure to navigate to the Exalogic Elastic Cloud-specific pages in Oracle Enterprise Manager Cloud Control 12*c*.

9.1 Accessing Oracle Enterprise Manager Cloud Control 12*c*

Refer to this topic for information on accessing the Oracle Enterprise manager cloud control 12*c*.

If you already have an existing Oracle Enterprise Manager Cloud Control 12*c* implementation, you can access it by navigating to the following URL. Depending upon the configuration of the implementation, the port may or may not need to be provided. See your Enterprise Manager administrator for the URL for your environment — `https://hostname.domain[:port]/em`.

For example:

```
https://exalogicem.mycompany.com:1159/em
```

1. If you do not yet have an Enterprise Manager Cloud Control 12*c* implementation, see the following resources:
   - Oracle Enterprise Manager 12c documentation.
     - *Enterprise Manager Cloud Control Basic Installation Guide*
     - *Enterprise Manager Cloud Control Advanced Installation Guide*
   - Oracle Enterprise Manager resources on Oracle Technology Network (OTN)Oracle Enterprise Manager resources on Oracle Technology Network (OTN).
2. Your ability to use the capabilities of Enterprise Manager Cloud Control 12c to manage Exalogic Elastic Cloud is dependent upon ensuring the availability of the Enterprise Manager system.

See the following resources for additional information:

- Maximum Availability Architecture (MAA) Best Practices for Enterprise Manager.
- EM Operational Considerations and Troubleshooting Whitepaper Master Index.

Note: To open a master note, perform the following steps:

- Select My Oracle Support document ID, and press Ctrl + F9. The Attributes dialog opens.
- In the **Attribute Value** field for the **Url** attribute, enter this URL:

  https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=note_id
- Enter the Note ID or keyword in the search field at the top of the screen.
- Click **Set Value**.

9.2 Discovering an Oracle Exalogic Elastic Cloud Target

To monitor and manage your Oracle Exalogic Elastic Cloud machine using Oracle Enterprise Manager Cloud Control 12*c*, you must prepare the Exalogic Elastic Cloud environment for discovery and then discover the related system targets.

The preparation and discovery steps differ based upon whether your Oracle Exalogic Elastic Cloud machine is installed in a physical or virtual configuration.

1. If your configuration is physical, follow the steps in *Discovering Exalogic Machine - Physical Configuration* in Oracle Enterprise Manager Cloud Control Managing and Monitoring an Exalogic Elastic Cloud Machine
2. If your configuration is virtual, follow the steps in *Discovering Exalogic Machine - Virtual Configuration* in Oracle Enterprise Manager Cloud Control Managing and Monitoring an Exalogic Elastic Cloud Machine.

9.3 Using Exalogic-Specific Pages in Oracle Enterprise Manager Cloud Control 12*c*

Refer to the topic for procedure to navigate to the Exalogic Elastic Cloud-specific pages in Oracle Enterprise Manager Cloud Control 12*c*.

Once you complete the steps to prepare for and discover the Exalogic Elastic Cloud target, do the following to navigate to the Exalogic Elastic Cloud-specific pages in Oracle Enterprise Manager Cloud Control 12*c*.

1. Log in to Oracle Enterprise Manager Cloud Control web interface.

   The home page is displayed.
2. On the home page, select **Targets**NOT_SUPPORTED**Systems** .

   The **Systems** page is displayed.

3. On the **Systems** page, select the target name for your Oracle Exalogic Elastic Cloud machine, such as **Exalogic Enterprise Deployment**. This is the target you have created in Discovering an Oracle Exalogic Elastic Cloud Target.

   The Oracle Exalogic Elastic Cloud Home page is displayed. This is the landing page for all Exalogic Elastic Cloud-specific monitoring and control operations, including configurations, application deployments, WebLogic domains, and metrics pages.

Oracle Enterprise Manager displays comprehensive metrics for Oracle Exalogic. When these metrics are displayed in tabular format, you can sort them in ascending or descending order. This feature enables you to find highest and lowest values easily.

This section contains the following topics:

Management and Monitor Features for Exalogic Configurations
: You can select software components directly from the Exalogic Elastic Cloud menu and can access hardware components by first selecting members.

Management and Monitor Features for Exalogic Virtual Configurations
: Additional monitoring and management features are available in Enterprise Manager Cloud Control 12*c* for an Exalogic Elastic Cloud in a virtual configuration.

9.3.1 Management and Monitor Features for Exalogic Configurations

You can select software components directly from the Exalogic Elastic Cloud menu and can access hardware components by first selecting members.

From the Oracle Exalogic Elastic Cloud home page, you can access the desired functionality, such as the Exalogic Elastic Cloud Dashboard, using the Exalogic Elastic Cloud menu, as shown in figure below.

Figure 9-1 Exalogic Elastic Cloud Menu

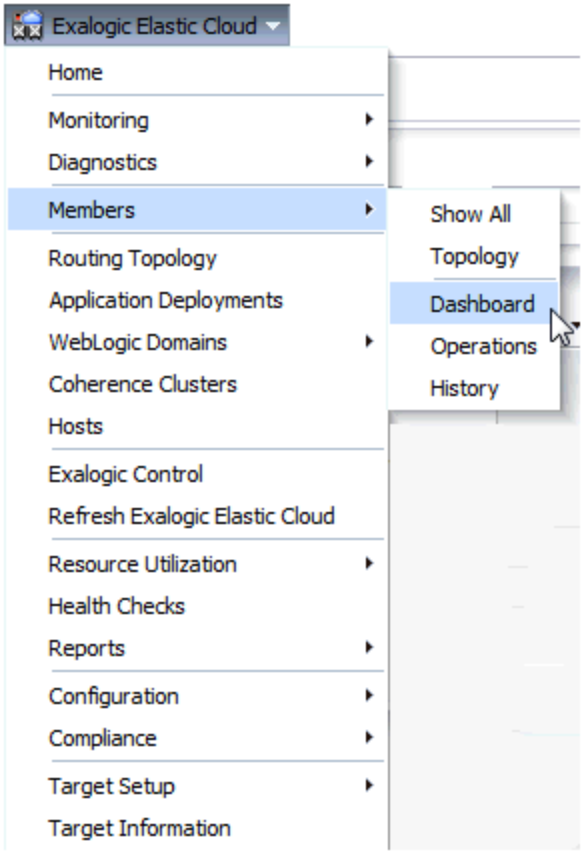

The Exalogic Elastic Cloud Dashboard is comprised of two tabs, Software tab and the Hardware tab.

The Software tab, which is shown in figure below, provides status information including alerts and key performance metrics for Exalogic Elastic Cloud targets, including for example the following, divided into regions:

- Application Deployments
- WebLogic Domains
- Coherence Clusters
- Hosts

Each region contains summary, status, related incidents, and other relevant information, and can be expanded or collapsed. There can be a large number of targets in a given region, each region can be filtered to display only the targets of interest and you can drill down to get further details.

Figure 9-2 Exalogic Elastic Cloud Dashboard Software Tab

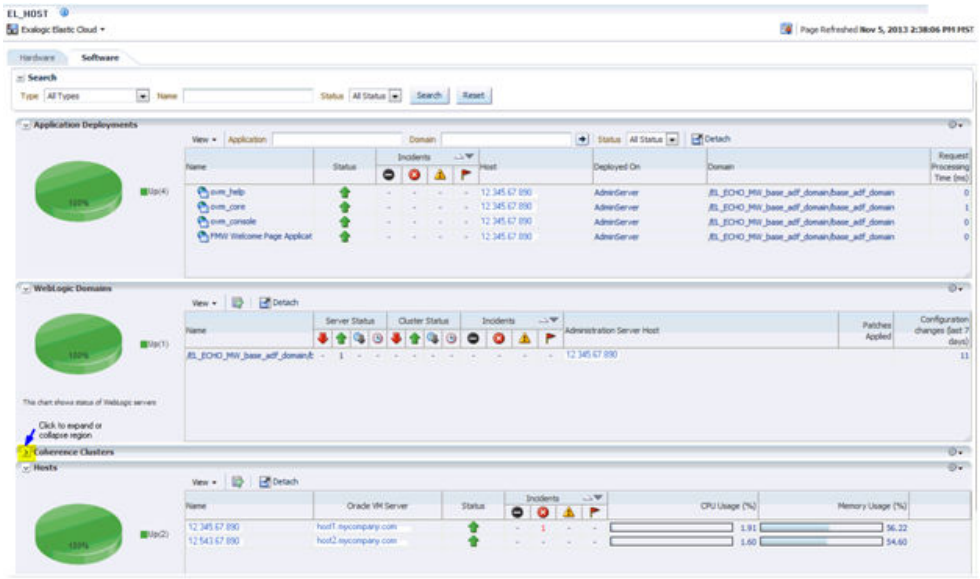

The Hardware tab, which is shown in figure below, provides status, incidents, a schematic diagram of the Exalogic Elastic Cloud, and other information about the hardware and infrastructure of the Exalogic Elastic Cloud, including a schematic diagram.

Figure 9-3 Exalogic Elastic Cloud Dashboard Hardware Tab

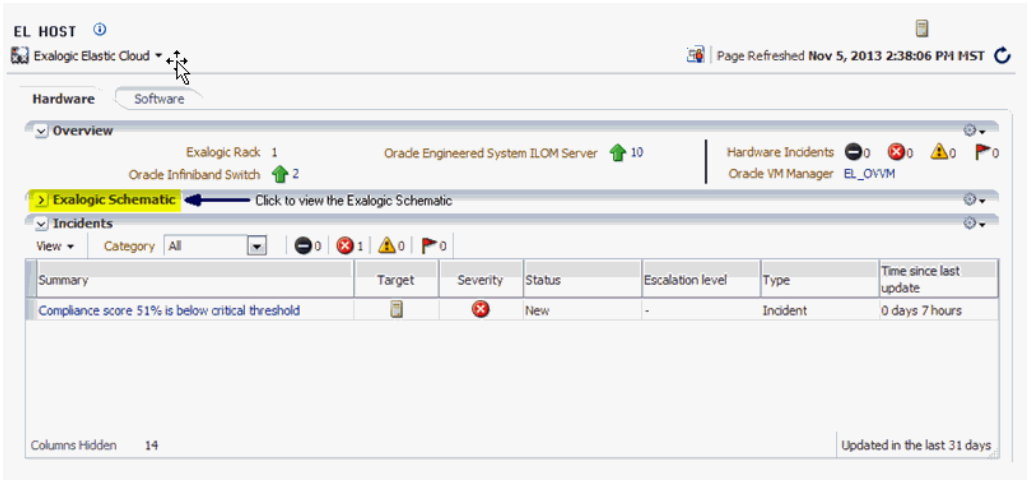

An example of the Exalogic Schematic is shown in the figure below. The information is separated into regions, each of which can be expanded or collapsed. In the Software tab, you can drill down to get further details.

Figure 9-4 Exalogic Elastic Cloud Schematic

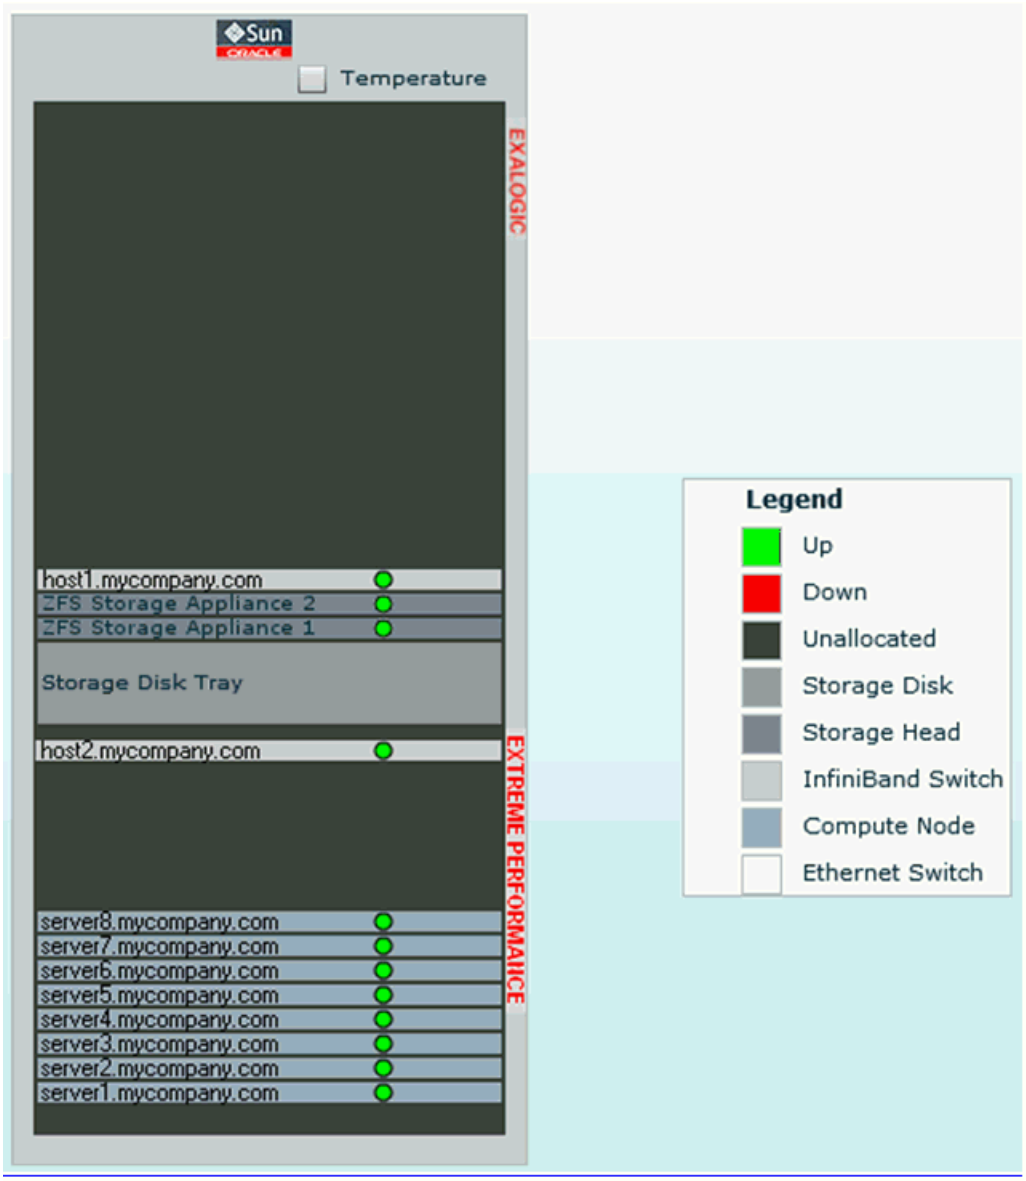

- See Monitoring and Managing Exalogic in *Oracle Enterprise Manager Cloud Control Managing and Monitoring an Oracle Exalogic Elastic Cloud Machine* for further information on the Exalogic Elastic Cloud Dashboard and for details on the following:
  - Monitoring the Hardware Components of Exalogic Elastic Cloud
  - Visualizing Relationships Between Exalogic Software and Hardware Components
  - Analyzing the Impact of Component Failures
- See Monitoring Hosts and Applications in *Oracle Enterprise Manager Cloud Control Managing and Monitoring an Oracle Exalogic Elastic Cloud Machine* for details on the following:
  - Viewing Hosts

  - Viewing Application Deployments
  - Viewing WebLogic Domains
  - Viewing Coherence Clusters
  - Creating Exalogic Reports

9.3.2 Management and Monitor Features for Exalogic Virtual Configurations

Additional monitoring and management features are available in Enterprise Manager Cloud Control 12*c* for an Exalogic Elastic Cloud in a virtual configuration.

- See Monitoring Tasks for Exalogic Virtual Configurations in *Oracle Enterprise Manager Cloud Control Managing and Monitoring an Oracle Exalogic Elastic Cloud Machine* for details on the following:
  - Exalogic Control Stack Monitoring
  - Viewing and Managing Exalogic Consumption Tracking
  - Viewing Incidents and Status Changes Created for an Exalogic System in Ops Center as Incidents in Cloud Control
  - Viewing the vCPU Consumption Report
  - Viewing the Resource Consumption Trend
  - Exporting the vCPU Consumption Report
- See Management Tasks for Exalogic Virtual Configurations in *Oracle Enterprise Manager Cloud Control Managing and Monitoring an Oracle Exalogic Elastic Cloud Machine* for details on the following:
  - Configuring the Exalogic Guest Base Template
  - Creating an Exalogic Control VM
  - Configuring the Exalogic Network