

**Simple  
Reliable  
Affordable**




ORACLE®

## **The Oracle Database Appliance I/O and Performance Architecture**


Tammy Bednar, Sr. Principal Product Manager, ODA

# Growing Business Critical Services and Data

## High Availability Solutions Desirable



**Costly and Complex**



**Specific Skills  
Required**



**Risk of Failure**

ORACLE

# Oracle Database Appliance

Simple. Reliable. Affordable.

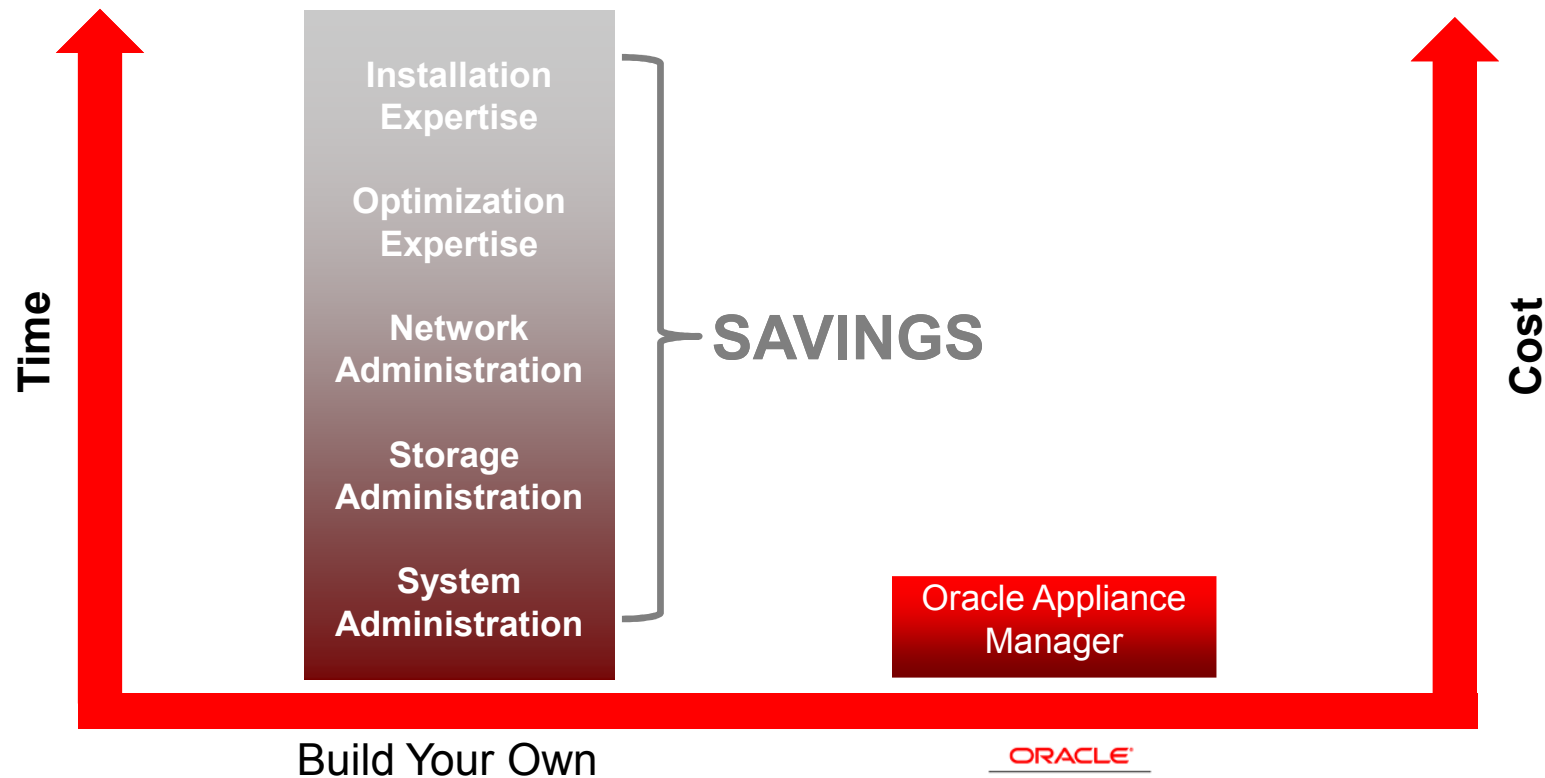
**ORACLE®**  
**DATABASE APPLIANCE**

**Simple.**



**ORACLE®**

# Simple to Install, Manage and Maintain





# Rapidly Deploy A Database Cluster

## Simple to Install

Rack the system

Cable the system

Wizard-driven install

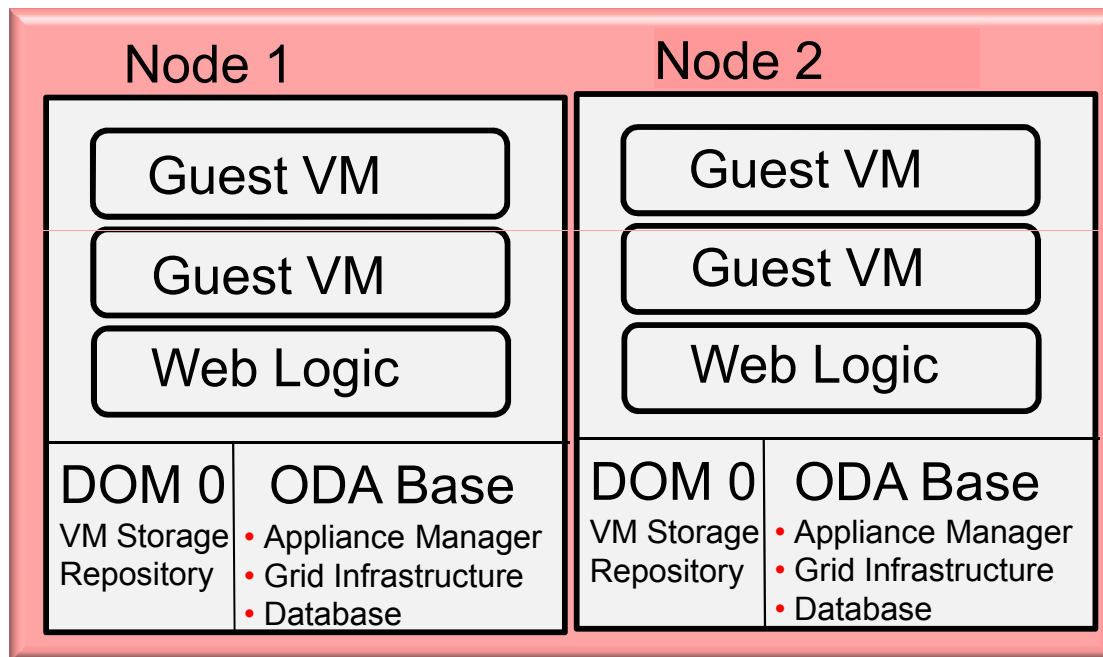


ORACLE

# Bare Metal or Virtualized?

How does the Virtualized Platform work?

## Oracle Database Appliance



- Database runs in the 'ODA Base' domain with access to shared storage
- Oracle Web Logic templates provide application HA
- Domains provide application isolation
- Appliance Manager provides:
  - VM Template and Domain management

ORACLE

# Oracle Database Appliance

Simple. Reliable. Affordable.

**ORACLE®**  
**DATABASE APPLIANCE**

**Reliable.**



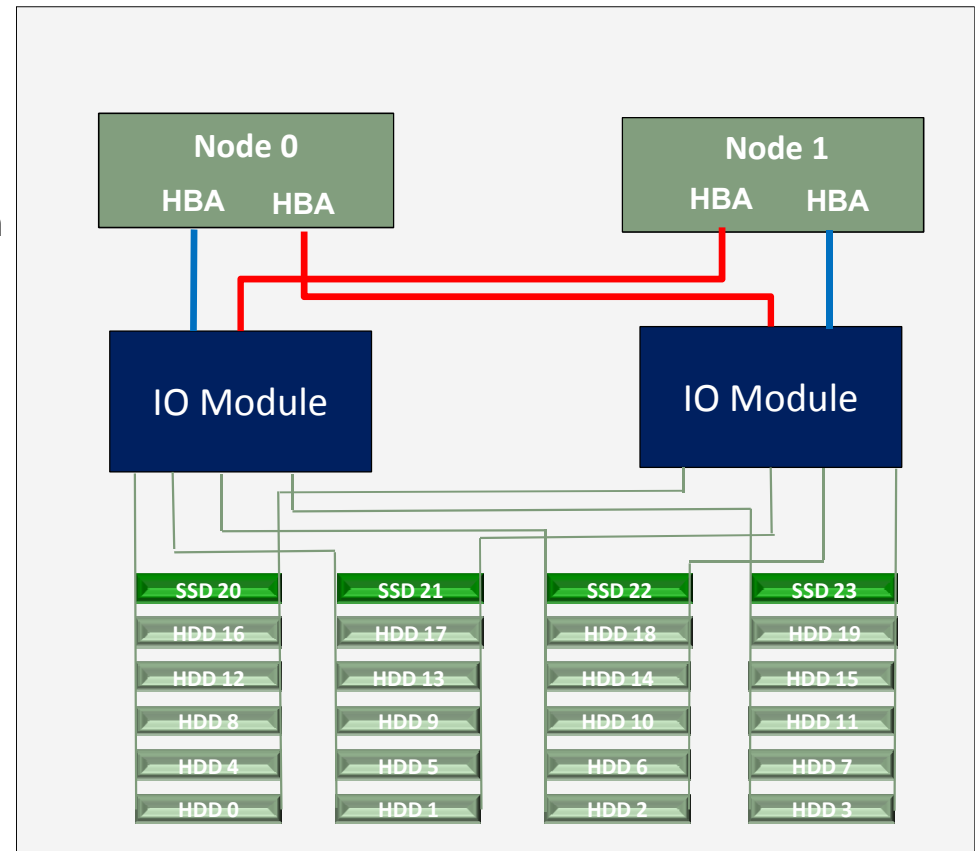
**ORACLE®**

# Fully Redundant Hardware

- Two dual-socket Oracle Linux servers
- Redundant 10GBase-T interconnect
- Double-mirroring or triple-mirroring storage redundancy
- Redundant hot-swappable power, cooling, and fans

# ODA X3-2 Redundant Storage Architecture

- Each Server Node
  - 2x HBA
  - In case of HBA failure
    - Multipath software transparently manages both paths for the database
- Storage Shelf
  - 2x IO Modules (Controllers)
    - Each connects to all 24 HDDs to protect against failure
  - Redundant HDDs and SSDs
    - ASM stripes data across HDDs to protect against failure
- Configuration
  - Load balances IO evenly across adapters and ports



# Highly Reliable Software

- Oracle Database 11g Enterprise Edition
  - Real Application Clusters (RAC)
  - RAC One Node
  - Single Instance
- Oracle Grid Infrastructure
  - Automatic Storage Management (ASM)
  - Oracle Clusterware
- Oracle Linux
- Oracle Appliance Manager

# Oracle Database Appliance

Simple. Reliable. Affordable.


**ORACLE®**  
**DATABASE APPLIANCE**

**Affordable.**



**ORACLE®**

# What Could You Do with 2000 Extra Hours?



**ORC***International*

**The Oracle Database Appliance  
simplifies time-consuming DBA tasks  
and saves:**

**835 hours in first year**

**669 hours each subsequent year**


<http://www.oracle.com/us/products/database/database-appliance-vs-sql-server-1434947.pdf>

ORACLE



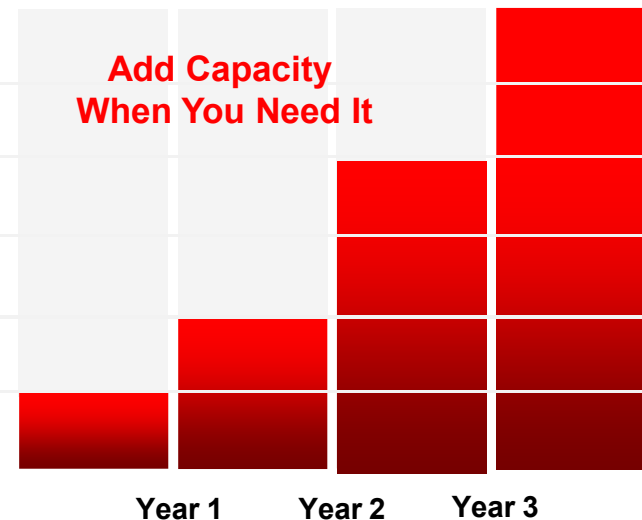
# Capacity On Demand Licensing

**Option 1: Build Your Own**  
*License 32 Cores for Anticipated Growth*



**Purchase Capacity Up Front**

**Option 2: Buy Database Appliance**  
*License as You Grow and Save Significantly*




**Purchase Capacity-on-Demand**

**OR**

ORACLE

# Best-in-Class Performance



ORACLE

# Oracle Database Appliance X3-2

## Up To 2x the Performance

### 2x Processing Power


- 33% more cores: 32 Intel Xeon processor E5-2690 cores
- 2.7x the memory: 512 GB of main memory

### 10x Networking Bandwidth

- Redundant 10GBase-T Private Interconnect
- Four Public 10GBase-T Ethernet Ports

### 4x Storage Capacity


- 2x Storage Capacity 18 TB Raw Disk Storage
- 2.7 x Flash: 800 GB Raw Solid State Storage
- Storage expansion doubles the capacity



ORACLE

# ODA X3-2 Storage Expansion Shelf

## Zero-Admin/Online Storage Expansion



### Double Available Storage Capacity

- Additional 18 TB HDD, 36 TB Total
- Additional 800GB Flash, 1.6TB total

### Zero Administration

- Automatically integrates when plugged in
- Data automatically distributes to new shelf

### Online Expand Storage

- Hot-plug storage expansion shelf
- No database downtime

ORACLE

# Processor Performance Has Doubled!

- Number of cores has increased by 1.33x
- Per core performance has increased by 1.5x
- Processor performance has doubled –  $1.33 \times 1.5 = 2x$
- Supporting memory has increased by 2.7x


# Overall System Performance

## Is I/O Performance Important?

- I/O performance is probably the most important factor in modern Intel systems performance
- CPU clock speeds (and memory) keep increasing rapidly but the bandwidth to disk is not keeping pace
- How does ODA solve this problem and why is it better than traditional approaches?

# Best-in-Class Performance

- Optimal data layout
  - Ideal disk group configuration
  - Data striping with ASM



# Storage Management

## Initial Configuration

- Discovers disks
- Creates partitions
- Sets up multipath
- Creates diskgroups and lays out data for best performance
  - DATA for database data
  - RECO for archive logs for backup data
  - REDO for redo logs



# Oracle Database Appliance X3-2

## Superior I/O Performance Architecture

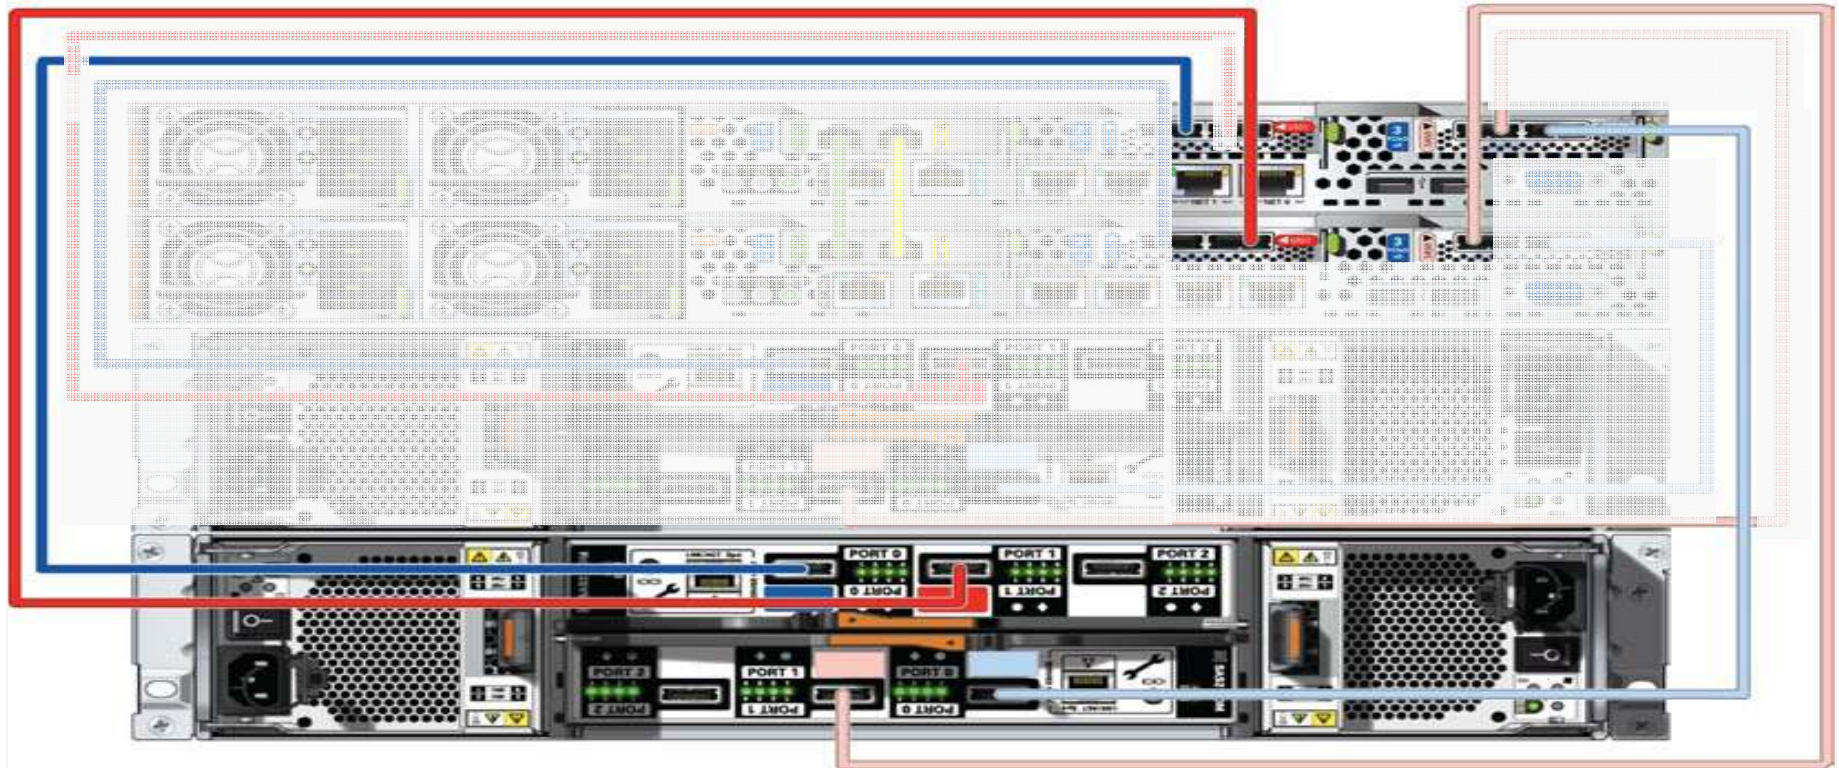
- Superior performance architecture
  - Higher performance architecture to traditional approaches
  - “Database-aware” performance architecture
  - Performance stays predictable over time
  - High performance does not compromise availability
- Linear performance increase with storage expansion
- ODA X3-2 also delivers a superior performance in the Virtualized Platform

# Database Aware Performance

## Optimized for the Oracle Database

- I/O Path optimization on ODA X3-2
  - Redo writes to REDO Disk group made up of SSDs
  - Database writes blocks to DATA Diskgroup residing in Outer HDD Platter
  - Archiving writes blocks to RECO Diskgroup residing in Inner HDD Platter
- ASM stripes data on all spindles

# Storage Expansion



ORACLE

# IO Performance

1 x Storage Shelf	IOPS	Bandwidth
Flash	200,000	2 GB/Sec
HDD	3300	3.5 GB/Sec

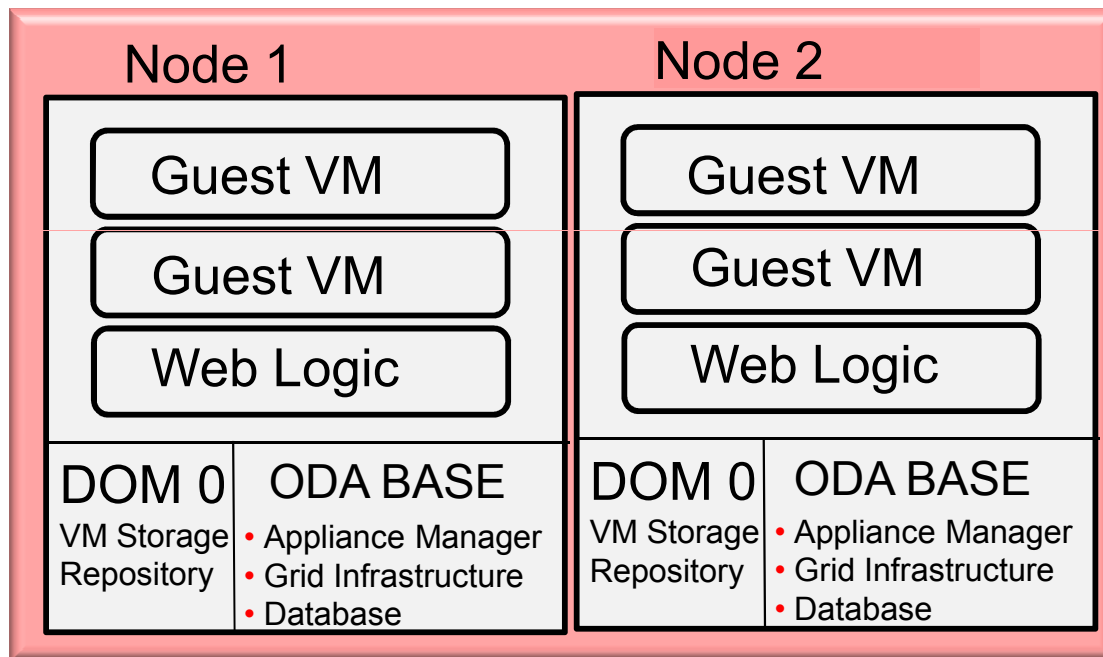
2 x Storage Shelf	IOPS	Bandwidth
Flash	400,000	4 GB/Sec
HDD	6600	5.5 GB/Sec

\*\*\* Performance measured while the workload is running in both nodes  
4K I/O Size

# Virtualized Platform Performance

Provides Bare Metal I/O Performance

## Oracle Database Appliance




- Virtualized environments typically don't perform well for I/O intensive workloads
- Database is an I/O intensive workload
- With ODA, database I/O is not virtualized
- Database I/O does not incur a VM tax i.e. performs great!

ORACLE

# Best-in-Class Performance

- Optimal data layout
  - Ideal disk group configuration
  - Data striping with ASM
- Direct-attached storage
  - Eliminates SAN/NAS bottlenecks




# SAN/NAS vs. Oracle Database Appliance

Why is the Oracle Database Appliance Performance Architecture Superior to Traditional SAN/NAS ?


ORACLE

# Typical SAN Architecture







# Unbalanced setup equals Unpredictable Performance



# Multiple Queues affects Performance and Diagnosis



## Direct Attached Oracle Database Appliance X3-2




- HDDs are directly attached eliminating multiple queues
- SAS adapters provide 3x bandwidth vis-à-vis FC adapters, improving performance
- Makes expansion easy and predictable
- **Less complexity. Easy to tune.**

# What about Flash Card Vendors ?

- Most prominent vendor is Fusion I/O
- They perform great but aren't redundant
- Customer has to choose – performance OR availability?

# Best-in-Class Performance

- Optimal data layout
  - Ideal disk group configuration
  - Data striping with ASM
- Direct-attached storage
  - Eliminates SAN/NAS bottlenecks
- Out-of-the-box fully tuned
  - Increases throughput
  - Improves response time



# Provisioning – Oracle Database Appliance X3-2

## Database Templates Sized for Performance

Database Class	CPU Cores	Memory	Number of Databases
Very Small	1	8 GB	32
Small	2	16 GB	16
Medium	4	32 GB	8
Large	6	48 GB	4
Very Large	12	96 GB	2
Extra Extra Large	16	128 GB	2


# Provisioning – Oracle Database Appliance X3-2

## Database Templates Sized for Performance

Database Class	System Global Area	Program Global Area	Number of Processes	Log Buffer, Redo Log Size
Very Small	4096-8192 MB	2048-4096 MB	200	16 MB, 1 GB
Small	8192-16384 MB	4096-8192 MB	400	16 MB, 1 GB
Medium	16384-24576 MB	8192-12288 MB	800	32 MB, 2 GB
Large	24576-49152 MB	12288-24576 MB	1200	64 MB, 4 GB
Very Large	49152-65536 MB	24576-32768 MB	2400	64 MB, 4 GB
Extra Extra Large	65536 MB	32768 MB	3600	64 MB, 4 GB

ORACLE

# Summary



- Simple to deploy, manage and maintain your database, application and web tier
- Best-in-class availability
- Best-in-class performance
- Built-in scalability
- Capacity-on-demand licensing for Oracle database and applications
- Solution-in-a-Box

ORACLE



# FOR MORE INFORMATION

**[oracle.com/databaseappliance](http://oracle.com/databaseappliance)**

ORACLE®

# Q&A

ORACLE

# Hardware and Software

The Oracle logo, consisting of the word "ORACLE" in white, sans-serif, uppercase letters, is centered within a solid red rectangular bar.

## Engineered to Work Together

ORACLE