

ORACLE®



ORACLE[®]

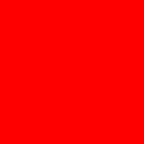
Solaris Engineered Systems

SPARC SuperCluster Introduction

Andy Harrison

andy.harrison@oracle.com

Engineered Systems, Revenue Product Engineering



The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions.

The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Agenda


- Introduction
- Product Overview
- Hardware Components
- Software Components
- System Configurations
- Installation
- Patching and Upgrade
- Q&A

ORACLE®
SPARC SUPERCLUSTER



ORACLE®

Introduction



Introduction

- Who am I?
- Background
- Current Role

ORACLE®
SPARC SUPERCLUSTER



ORACLE®


Product Overview



Motivation for Engineered Systems



Transforming the Enterprise Architecture



Compute, Storage, Network
Building Blocks



Exadata, Exalogic



SPARC Supercluster

**Best of Breed
Traditional Servers**

Engineered Systems

ORACLE®

Engineered Systems & Appliances


Purpose Built



Exadata



Exalogic



ORACLE[®] **11^g**
DATABASE

ORACLE[®]
FUSION MIDDLEWARE

General Purpose



SPARC
SuperCluster

ORACLE[®]
SOLARIS

ORACLE[®] **11^g**
DATABASE

ORACLE[®]
FUSION MIDDLEWARE

ORACLE[®]
OPTIMIZED SOLUTIONS

SPARC SuperCluster – What is it?

- General purpose, flexible, integrated, optimized and scalable server to run database, middleware & applications
- **Compute, Storage, Network**
 - Compute: Scalable compute nodes, flash
 - Storage: DB storage and network connectivity to general purpose storage
 - Network: IB backplane, 10GbE for external connectivity
- **Optimizations and Integration**
 - Exadata Storage Cells & 11gR2 DB
 - Exalogic Elastic Cloud Software
 - Ops Center & Enterprise Manager Grid Control

Standardized and Simple to Deploy

Value Proposition



**Deploy in Days,
Not Months**

All SPARC SuperCluster servers are the same

Delivered Tested and Ready-to-Run

Highly Optimized

Highly Supportable

No unique configuration issues

Identical to config used by Oracle Engineering

Runs existing database, middleware, and custom applications

No SPARC SuperCluster certification required


Full 30 years of Oracle DB capabilities

Leverages Oracle ecosystem

Skills, knowledge base, people, partners


SPARC SuperCluster Architecture

Best infrastructure solution for enterprise applications



- 1,024 CPU threads
 - 4 TB DRAM
 - 97 to 198 TB Hard Disk
-
- 8.66 TB Flash
 - 1.2M IOPS
 - 42 GB/sec Storage Bandwidth
 - 896 Gb/sec InfiniBand Interconnect

Hardware Components



SPARC T4

5x Per Thread Performance



- 3.0 GHz
- 8 Cores, 64 Threads
- Dynamic Threading
- Out of Order Execution
- 2 On Chip Dual-Channel DDR3 Memory Controllers
- 2 On Chip 10 GbE Networking
- 2 On Chip x8 PCIe gen2 I/O Interfaces
- 18 On Chip Crypto functions
- Balanced high-bandwidth interfaces and internals
- Co-engineered with Oracle software


9 World Records and Counting

SPARC T4-4 Systems Key Benefits

Feature	Function	Benefit
3.0 GHz frequency	Faster single threaded processing	Shorter application boot times, rapid batch processing for quicker results
Integrated encryption engines	Up to 3x faster security for data encryption	Reduced cost for secure datacenter operation without a performance penalty
Multithreaded architecture	Preserve T3 system levels of throughput	Time savings due to no application changes for T4 systems
Binary compatibility	Existing SPARC/Solaris applications run unmodified	Rapid time to adoption of new systems and new service deployment
Built-in Virtualization with OVM Server for SPARC	Flexible logical partitioning and live application mobility	Improved uptimes for critical services and higher system utilization rates
Solaris 11	Reliable, secure and streamlined operation	Faster system updates and reboots to improve datacenter operations

SPARC SuperCluster

Hardware Stack – Half Rack



Compute

- 2 * T4-4 nodes, each with:
 - 4 * T4 processors @ 3.0GHz
 - 1 TB memory
 - 6 * 600GB internal SAS disks and 2 * 300GB SSDs
 - 4 * Infiniband HCAs and 4 * 10GbE NICs

Network

- 3 * Sun DataCenter InfiniBand Switch 36-port Switches
- GbE Management Switch

Storage

- 3 * Exadata Storage Servers
- Optional Exadata Storage Server Expansion Rack

Shared Storage


- ZFS Storage Appliance 7320 with 40TB of disk capacity

Data Migration

- Optional FCAL HBA to connect to existing SAN storage

SPARC SuperCluster

Hardware Stack – Full Rack



Compute

- 4 * T4-4 nodes, each with:
 - 4 * T4 processors @ 3.0GHz
 - 1 TB memory
 - 6 * 600GB internal SAS disks and 2 * 300GB SSDs
 - 4 * Infiniband HCAs and 4 * 10GbE NICs

Network

- 3 * Sun DataCenter InfiniBand Switch 36-port Switches
- GbE Management Switch

Storage

- 6 * Exadata Storage Servers
- Optional Exadata Storage Server Expansion Rack

Shared Storage

- ZFS Storage Appliance 7320 with 40TB of disk capacity

Data Migration

- Optional FCAL HBA to connect to existing SAN storage

SPARC SuperCluster

Exadata Storage Servers

Storage Server Hardware

- X4270 M2 with
 - 2 sockets Xeon L5640 6 cores @ 2.26GHz 12MB L3
 - 24GB Memory (6 * 4GB LV 1333Mhz DDR3 DIMMs)
 - SAS-2 RAID HBA
 - 12 disks: 600GB@15Krpm (Hi Perf) or 2TB@7200rpm (Hi Capacity)
 - 4 * F20 (Aura) @ 96GB, total 384GB
 - 1 QDR IB HCA (2 ports)

Half Rack

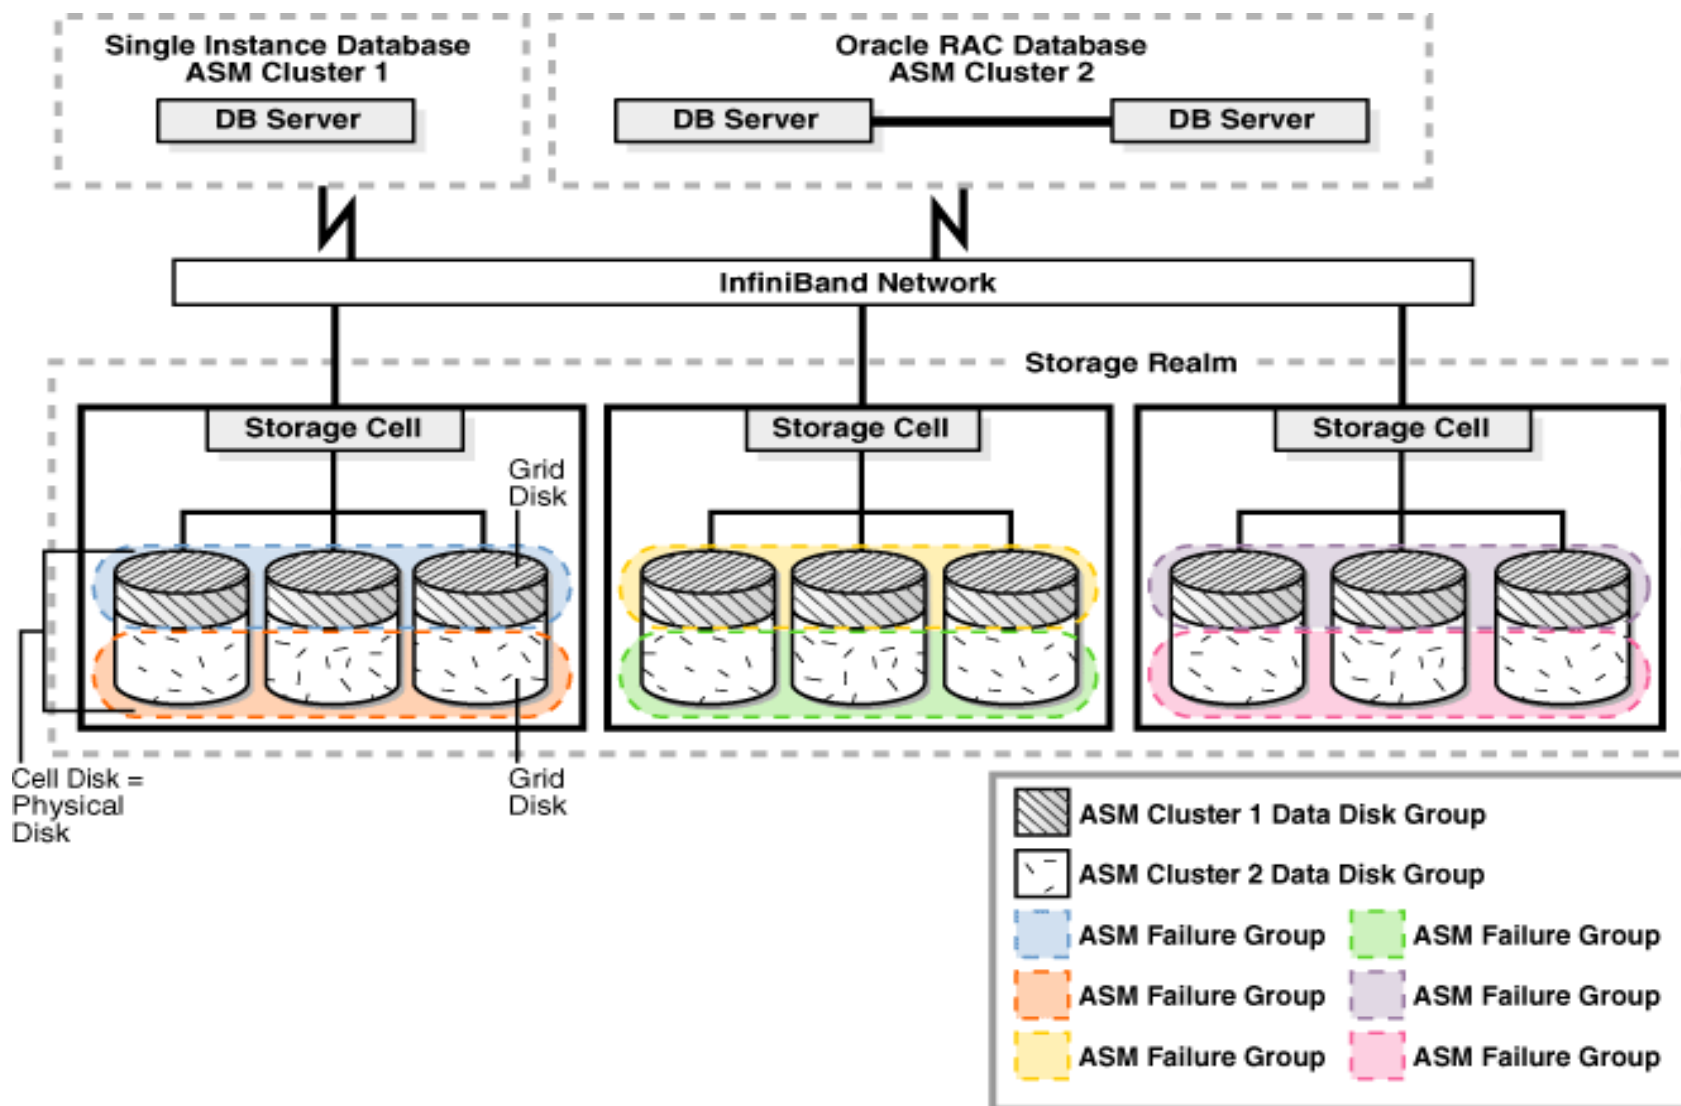
- 3 Exadata Storage Servers (Hi Perf or Hi Capacity)

Full Rack

- 6 Exadata Storage Servers (Hi Perf or Hi Capacity)

Expansion Rack

- Optional
- Up to 18 Exadata Storage Servers
- 3 * NanoMagnum2 36-port Infiniband Switches



SPARC SuperCluster

Shared Storage

ZFS Storage Appliance

- Provided with both half- and full-rack configurations
- ZFS Storage Appliance 7320HA (2 controllers)
- Each 7320 controller includes:
 - 2 * Quad core Xeon processors @ 2.4GHz
 - 24GB Memory
 - 4 * 512GB Readzillas (read-optimized SSDs)
 - Infiniband HBA (2 port)
 - GbE Management port

Disk Shelf

- 20 * 2TB disks @ 7200rpm
- 4 * 18GB Logzillas (write-optimized SSDs)

SPARC SuperCluster

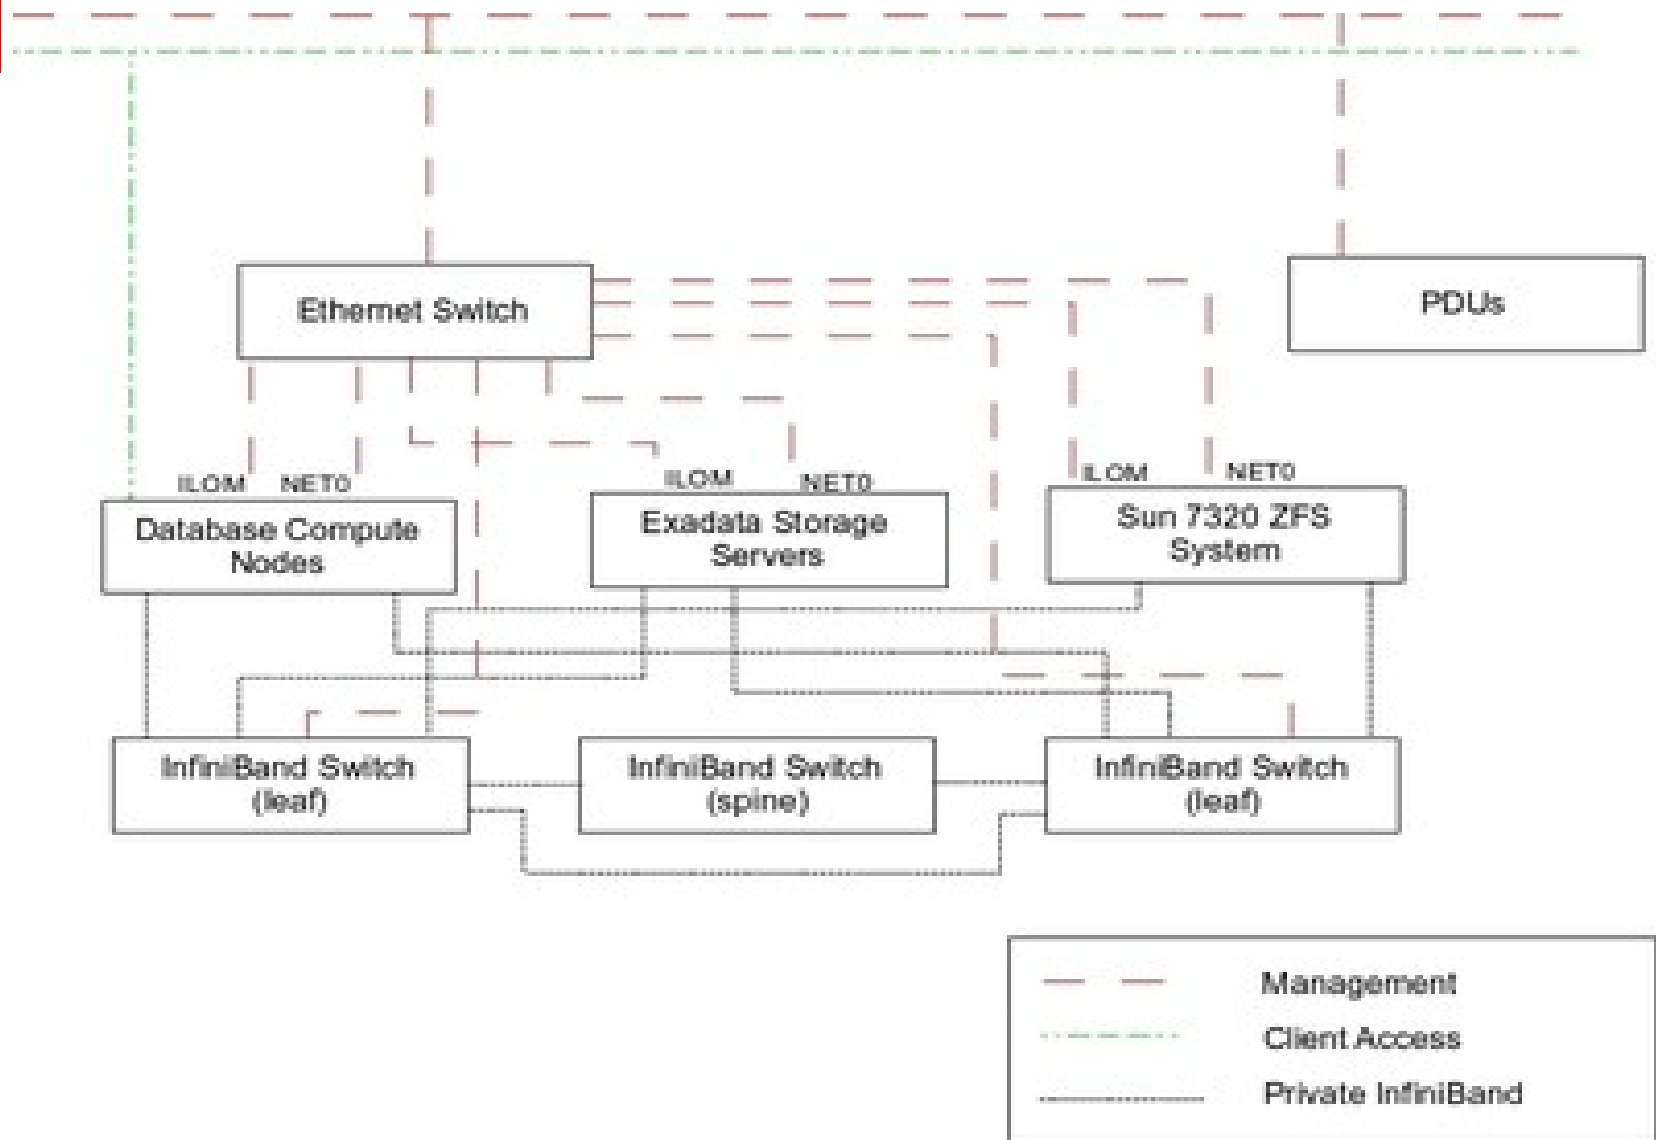
Networking

Cisco Switch

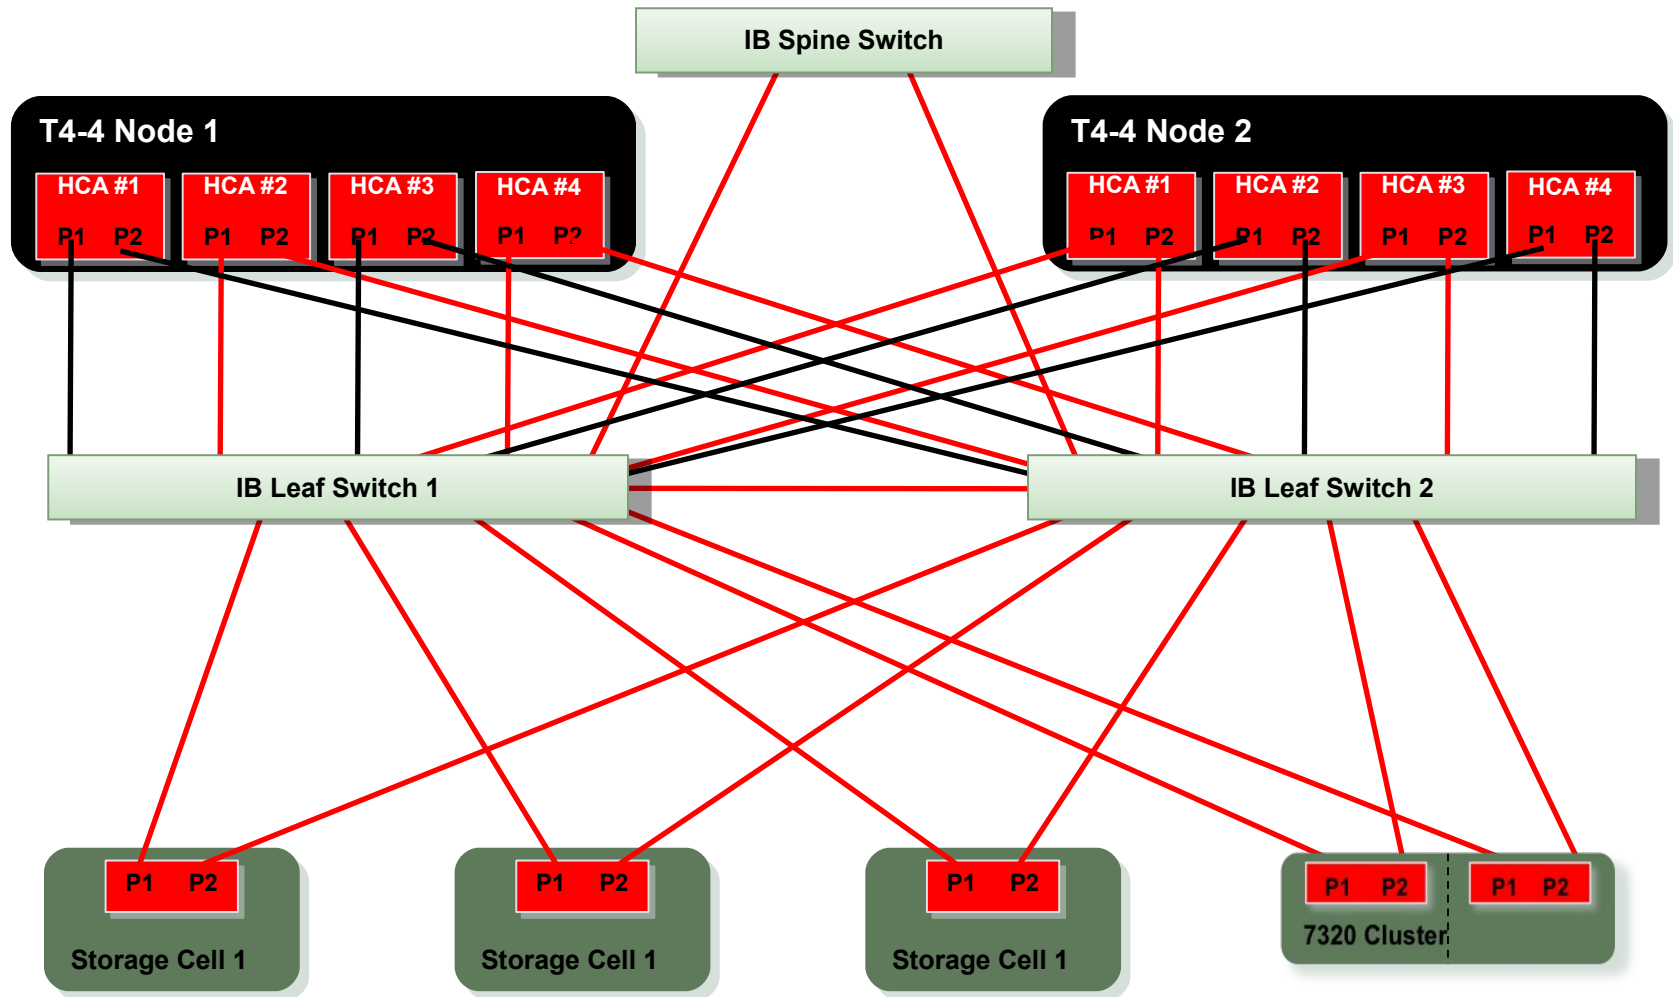
- Cisco 4849 Gigabit Network Switch
- Provided with both half- and full-rack configurations
- Management network
 - iLoms
 - Infiniband switches
 - ZFS SA
 - PDU

Infiniband switch (x3)

- 36 Ports
- 3 partitions (RAC, Cluster, Application)




InfiniBand Network Architecture



Scale Racks By Just Adding Cables

Full Bandwidth and Redundancy





SPARC SuperCluster

Cabling requirements for expanding

Cabling requirements

- Interconnecting four (4) racks require the purchase of eight (8) additional cables.
- Interconnecting five (5) racks require the purchase of eighteen (18) additional cables
- Interconnecting six (6) racks require the purchase of twenty six (26) additional cables.
- Interconnecting seven (7) racks require the purchase of forty (40) additional cables.
- Interconnecting eight (8) racks require the purchase of sixty two (62) additional cables.

Software Components



Software Stack

Engineered to Work Together

Database

11g R2, 10g and other DB

Operating System

Oracle Solaris 11

for Exadata and Exalogic nodes
Solaris 10/11 nodes for applications



Virtualization

*Oracle Solaris Zones and Oracle
VM Server for SPARC*

Middleware

*Oracle Fusion Middleware
Applications*

Optional with Exalogic Elastic Cloud



Applications

*E-Business Suite, PeopleSoft, SAP,
Siebel, and much more*

Management

*Oracle Ops Center and Enterprise
Manager Grid Control*

Clustering

*Oracle Solaris Cluster
Oracle Clusterware*

ORACLE®

SPARC SuperCluster

Software Stack

Operating System

- Solaris 11 for Exadata and Exalogic nodes (physical or LDoms)
- Solaris 11 or Solaris 10 nodes for applications (physical or LDoms)

Virtualization

- Zones (including Branded Zones) and LDoms

Management

- Ops Center and Enterprise Manager Grid Control (for DB)

Clustering

- Oracle Solaris Cluster 3.3u1 (Quorum provided by iSCSI ZFS SA lun)
- Oracle Clusterware (for DB)

Database

- 11gR2 to leverage Storage Cells
- Other databases with external storage

Middleware

- WebLogic Server with optional Exalogic Elastic Cloud Software

Applications


- Oracle, ISV and customer applications qualified on Solaris 10 or Solaris 11

SPARC SuperCluster

Virtualization

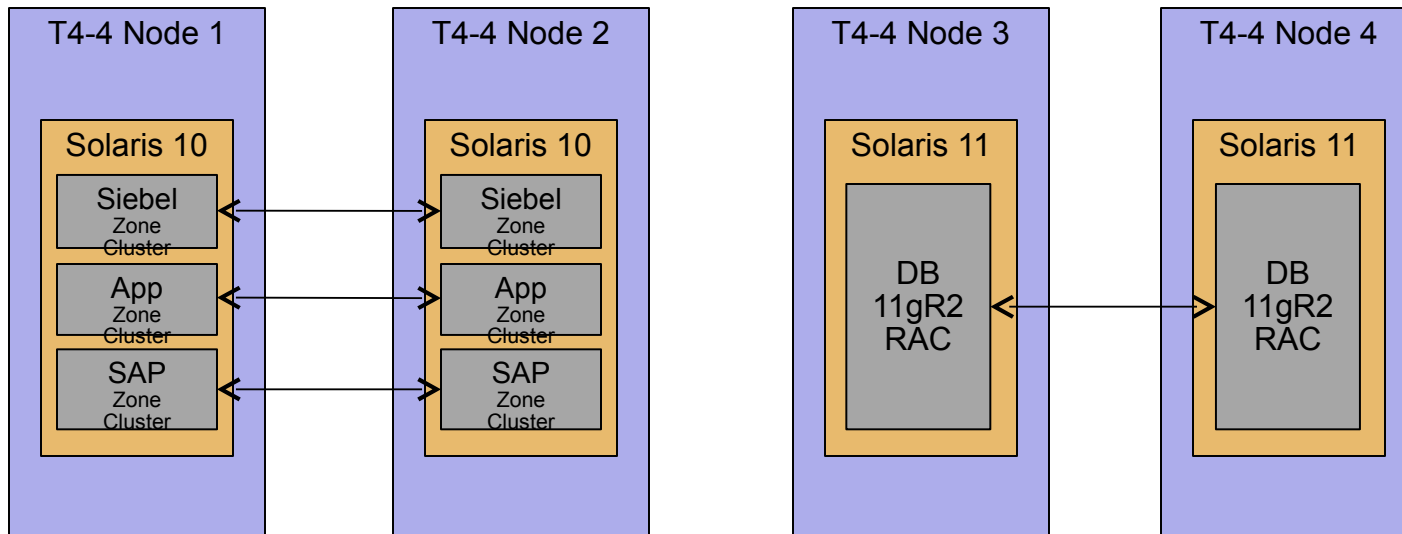
- LDoms for major application stacks
 - Exadata, Exalogic, Solaris 10, Solaris 11
 - Maximum 4 LDoms per node
 - Each with PCI root complex, so no overhead
- Zones
 - Deploy on Solaris 10 and Solaris 11 nodes / domains
 - Full resource management available
 - Minimal overhead

System Configurations



SPARC SuperCluster HA Consolidation

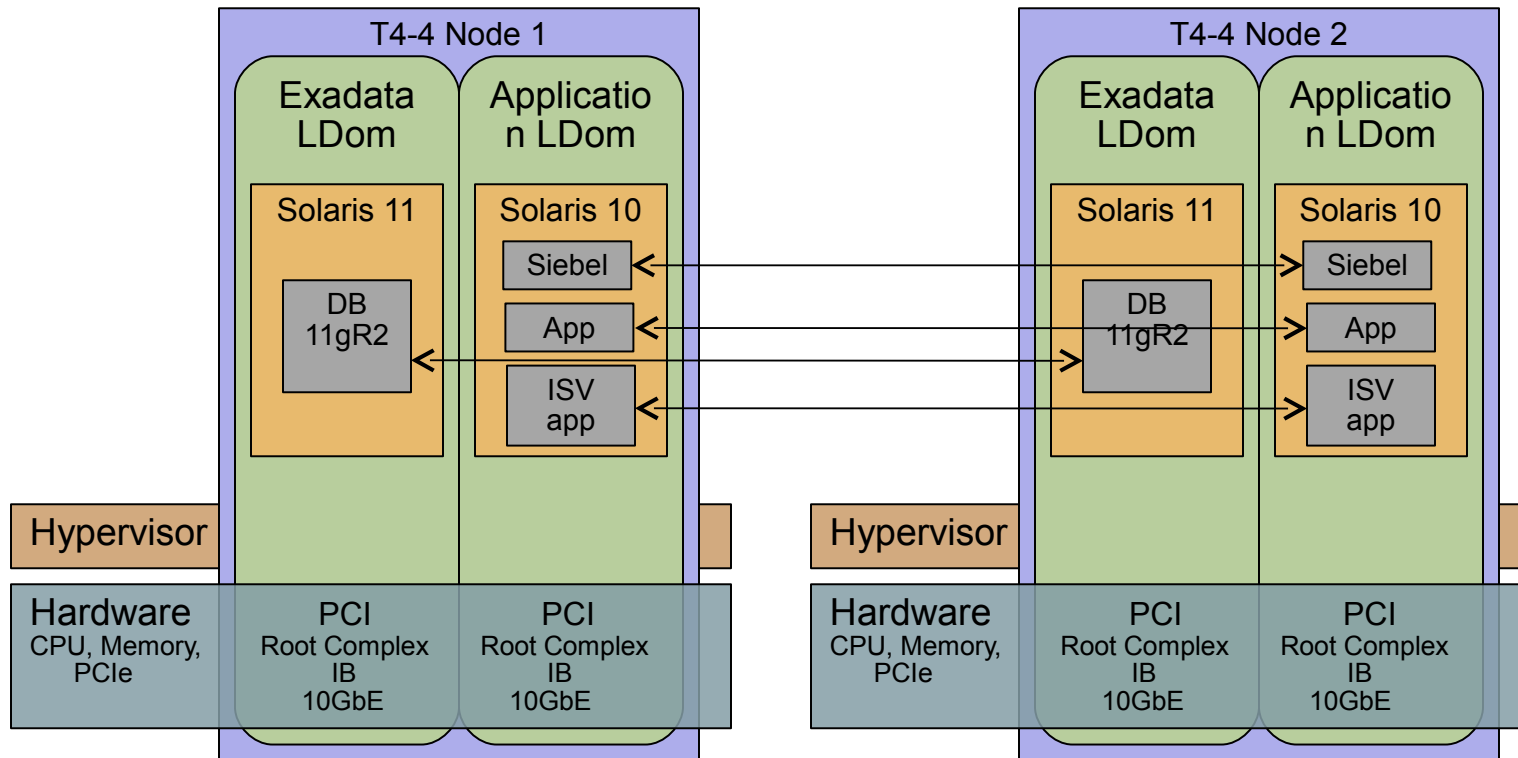
4 Node Configuration Example



SPARC SuperCluster HA Consolidation

Half Rack Configuration Example

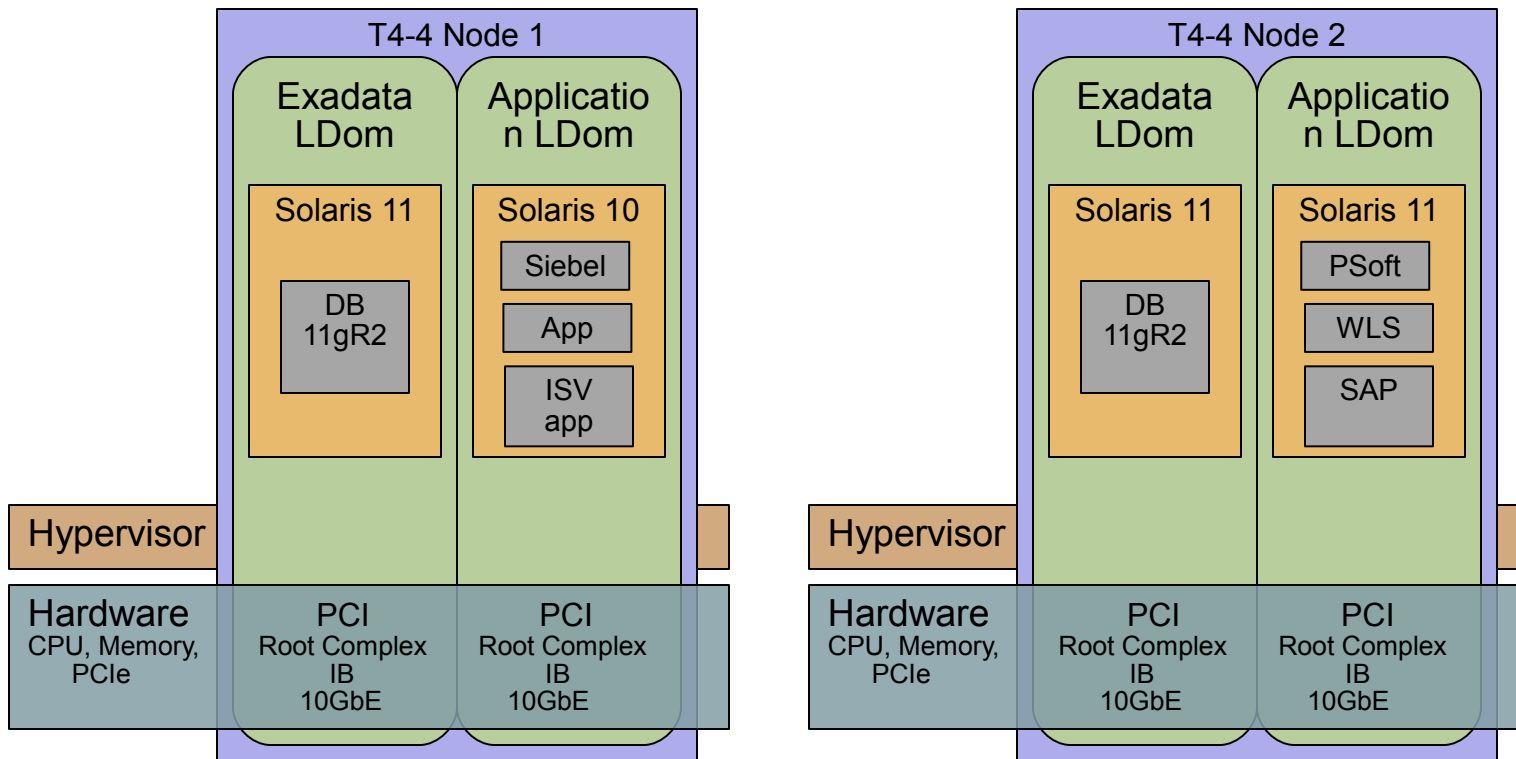
Full failover with
Oracle Solaris Cluster



SPARC SuperCluster Consolidation

Half Rack Configuration Example

Consolidation only;
No failover



Installation



SuperCluster Installation

- Factory install places Exadata/Exalogic Solaris 11 image on all nodes
 - Solaris 10 image can be preloaded (not pre-installed)
 - Ops Center and Solaris Cluster can be preloaded
- At customer site by ACS
 1. Any required LDoms are created
 2. Required OS is loaded on each node/LDom
 3. Exadata onecommand is run on Exadata nodes/LDoms only
 4. Exalogic/Supercluster onecommand is run on other nodes/LDoms

Migration / Backups

- Existing data transferred via 10g interfaces.
- Transferred direct to DBM or via ZFS SA.
 - Import via ZFS on InfiniBand
- Use FC to connect existing storage
 - Migration only / recommend use SuperCluster Storage
- Backups
 - Oracle DataGuard for DBM
 - Geographic Oracle Cluster


Patching and Upgrade



Patching and Updates

- Exadata/Storage Cells updates and patching
 - Follow same process and cadence as for x86 Exadata
 - Only carried out on Exadata nodes/LDoms
- Exalogic updates and patching
 - Follow same process and cadence as for x86 Exalogic
 - Only carried out on Exalogic nodes/LDoms
- Application updates and patching
 - Running on generic Solaris 10 and 11 nodes/LDoms
 - Follow same process and cadence as for same application in other environments

Questions ?



ORACLE®