Oracle Database Appliance

Simple. Reliable. Affordable.

Nov 9th 2011

Rhos B. Dyke
Executive Vice President, Cloud Creek Systems

Sohan DeMel
Vice President, Oracle

ORACLE
DATABASE APPLIANCE

ORACLE

Introduction to CCS

Cloud Creek SYSTEMS

- Formed in 1996, over 15 years of Oracle experience
- Single focus on Oracle product expertise
- Over 275 clients served, over 1700 Professional Service Engagements
- 93% repeat & add-on business
- 2007 & 2009 Oracle Titan Award Winner
- 2009 Oracle Global Mid Market Partner of the Year
- Specialized & Certified in:

ORACLE Platinum Partner
Specialized
Oracle Database

ORACLE Platinum Partner
Specialized
Oracle Database 11g Performance Tuning

ORACLE Platinum Partner
Specialized
Oracle Linux

ORACLE Platinum Partner
Specialized
Oracle Business Intelligence Foundation

ORACLE Platinum Partner
Specialized
Oracle Real Application Clusters

ORACLE Platinum Partner
Specialized
Oracle Exadata

ORACLE Platinum Partner
Specialized
Oracle Data Warehousing

ORACLE Platinum Partner
Specialized
Oracle Database Security Options

ORACLE Platinum Partner
Specialized
Oracle GoldenGate

ORACLE

Agenda

- Hardware Architecture
  - integrated server, storage, networking
  - storage architecture
- Software Architecture
  - high availability design
  - appliance manager
- Performance Architecture
  - high performance design
  - i/o and system sizing

DB Appliance Chassis
(Front View) Cluster in a Box

Cloud Creek SYSTEMS

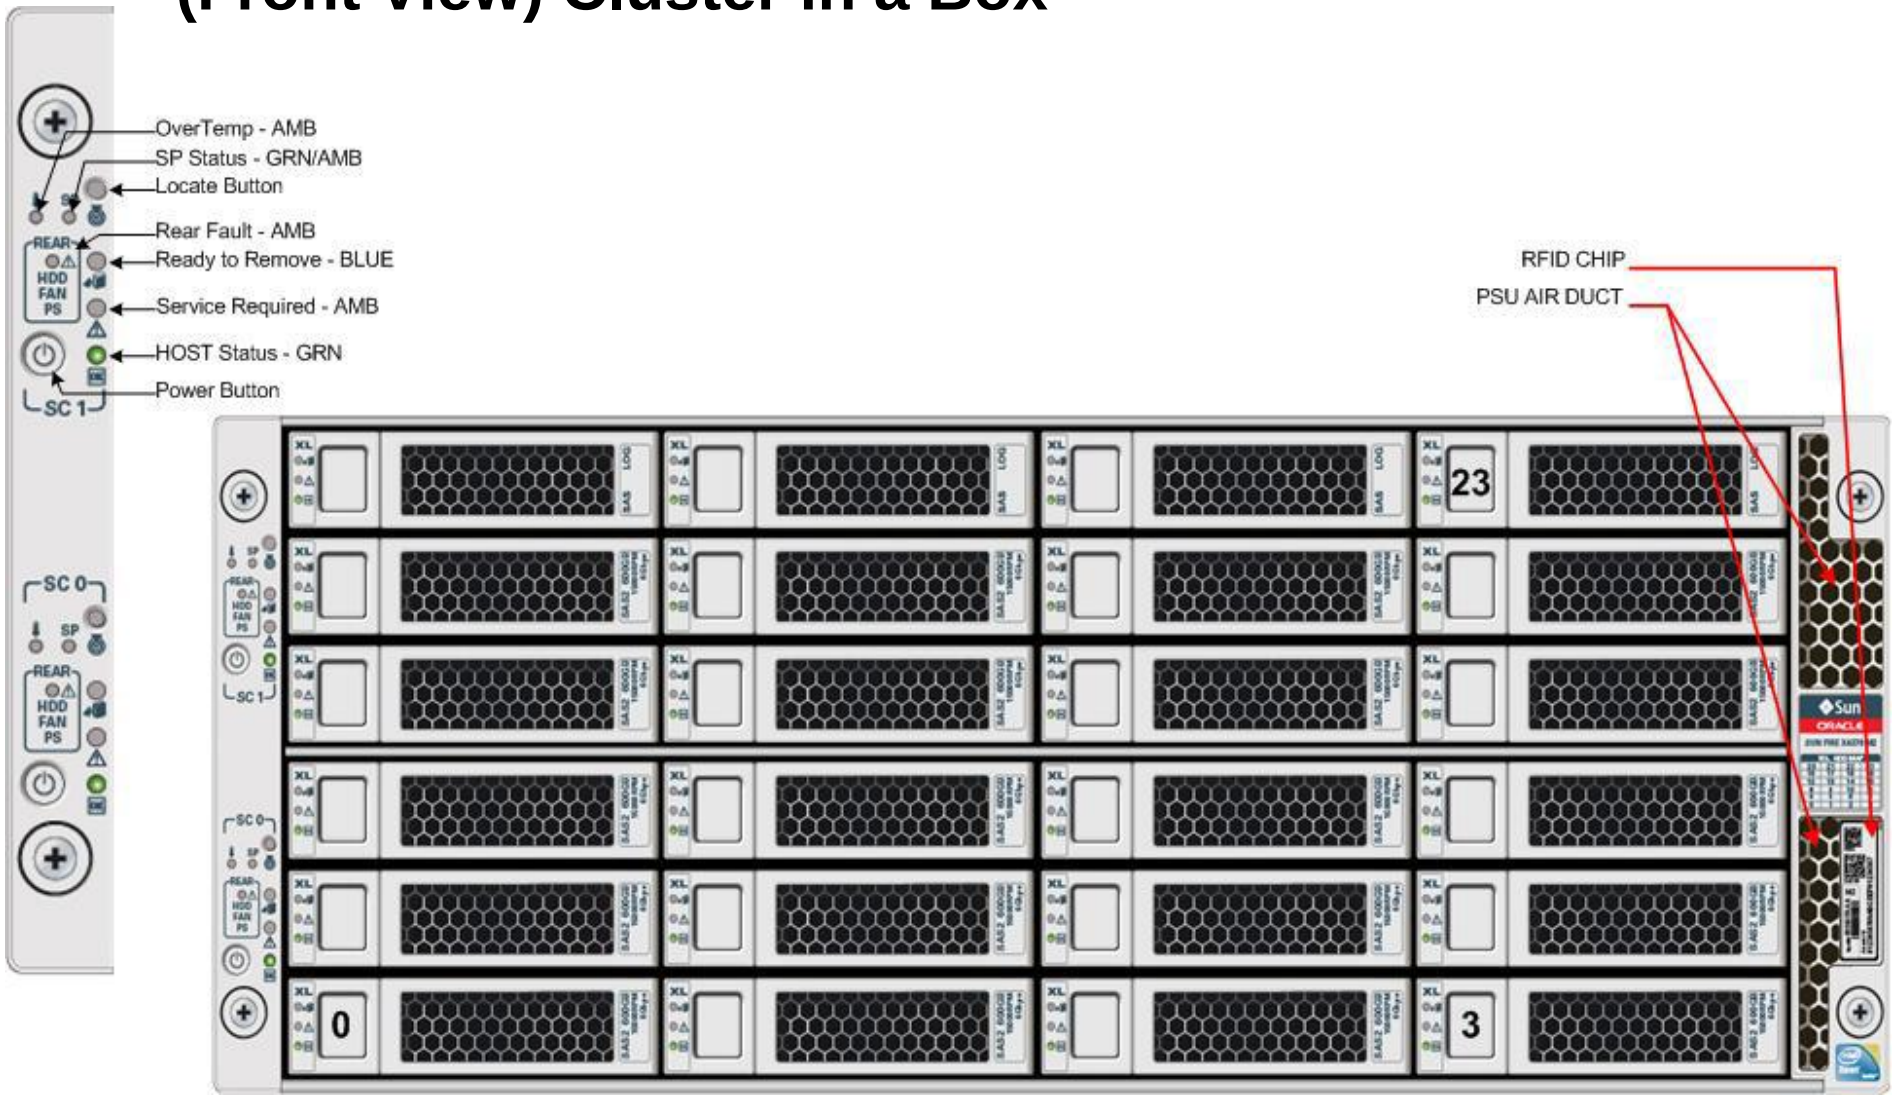

ORACLE

DB Appliance Chassis

(Back View)

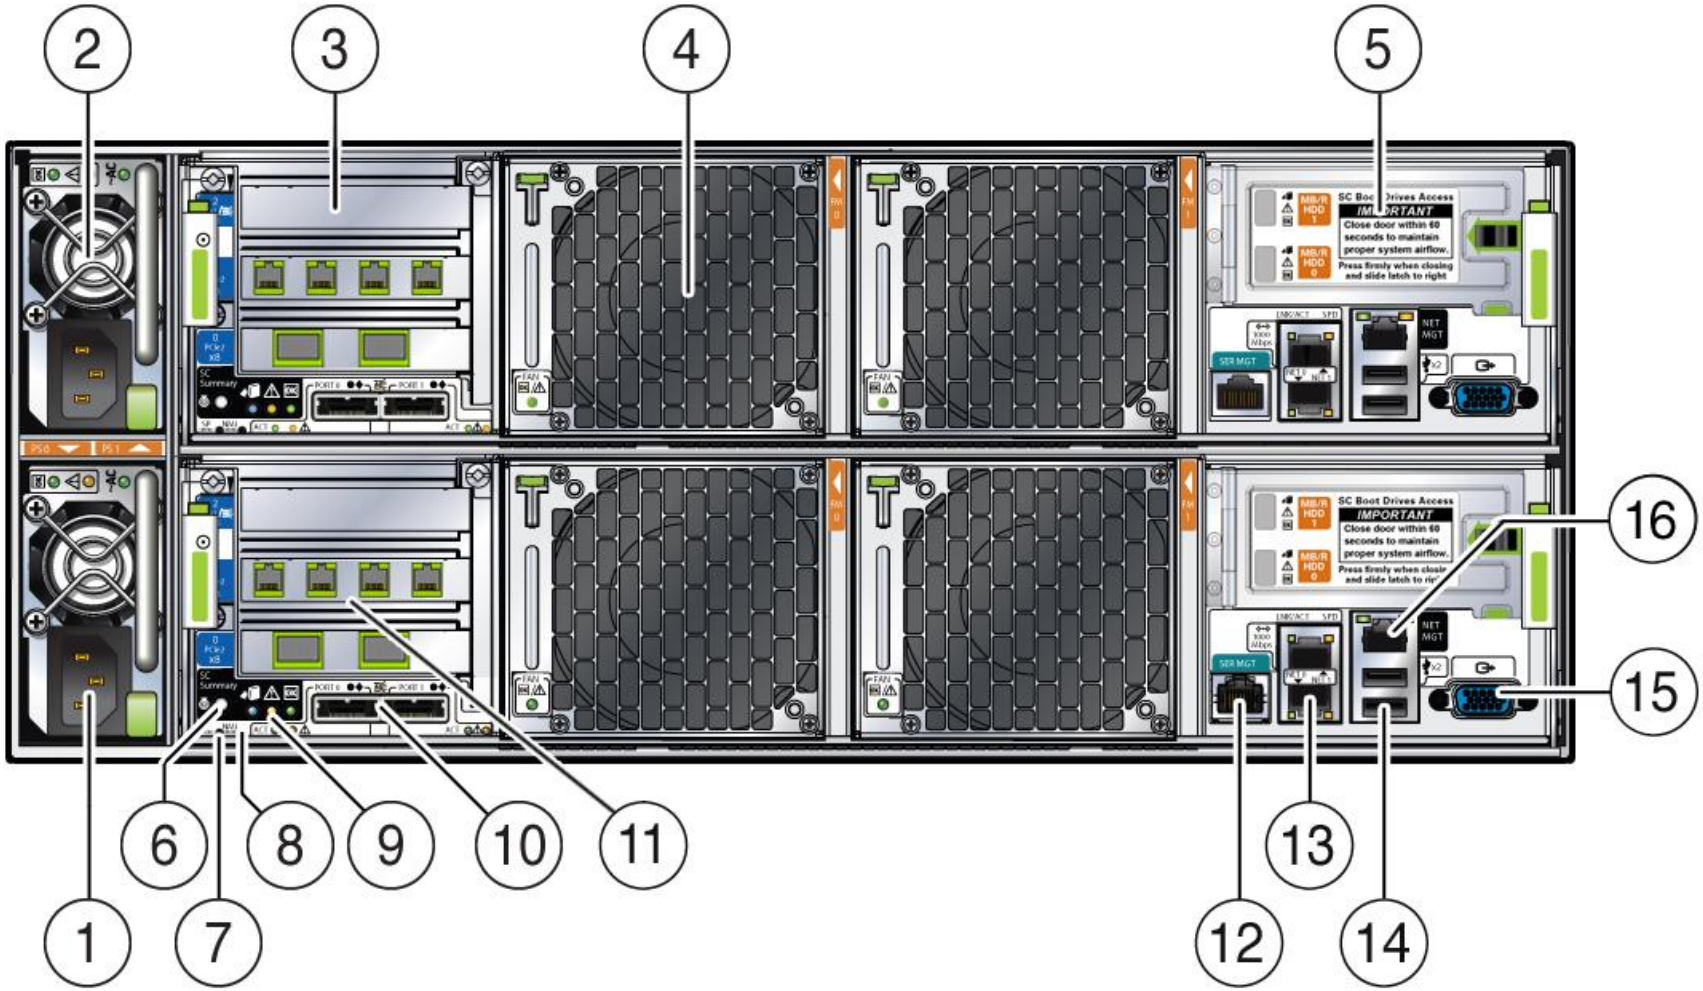

Hardware Architecture

Block Diagram

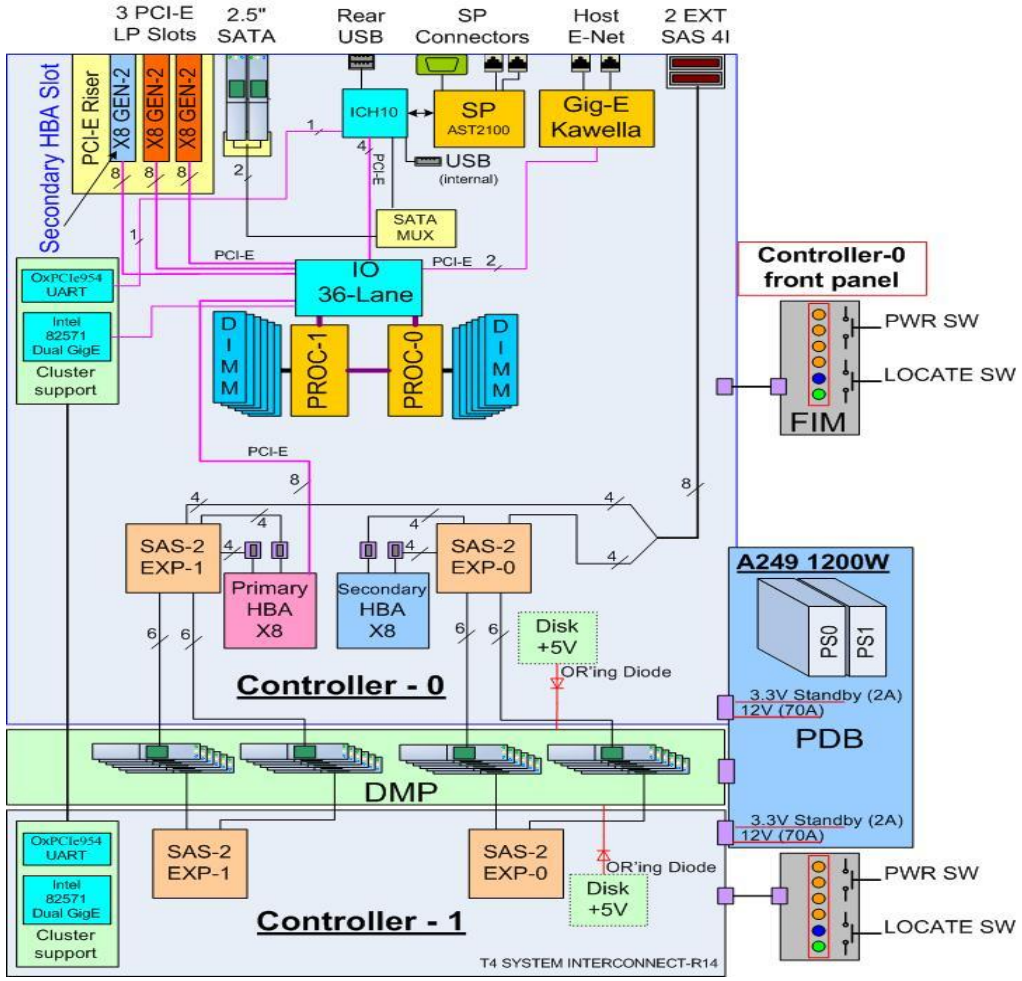

Hardware Architecture

Hardware Elements

Features	Description

Compute	• 2 Compute Nodes
	• On each node
	  • 2 x 6 Core Intel Xeon 5650 3.07GHz CPUs
	  • 12MB L3 cache per socket
	  • 3x full-width bi-directional Intel QuickPath Interconnects per socket
	  • Upto 25.6GB/s per QPI Ports
	  • 96 GB RAM (12 x DDR3-1333 8GB DIMMs)
	  • 3x Channel per socket and 2x DIMMs per channel

I/O	• 4 x 73GB SAS2 SSDs • 20 x 600GB SAS2 15K RPM Disks • 2 x 500GB SATA boot disks • 36 x PCIe Gen2 lanes per node • 3 x PCIe Gen2 8-lanes slots per node
Network	• Intel 82571 Dual GigE as the Cluster Interconnect • 2 x on-board GigE per node • 1 x Quad GigE per node • 1 x 10GigE Dual-ports per node

Hardware Architecture

HDD Storage Subsystem Block diagram

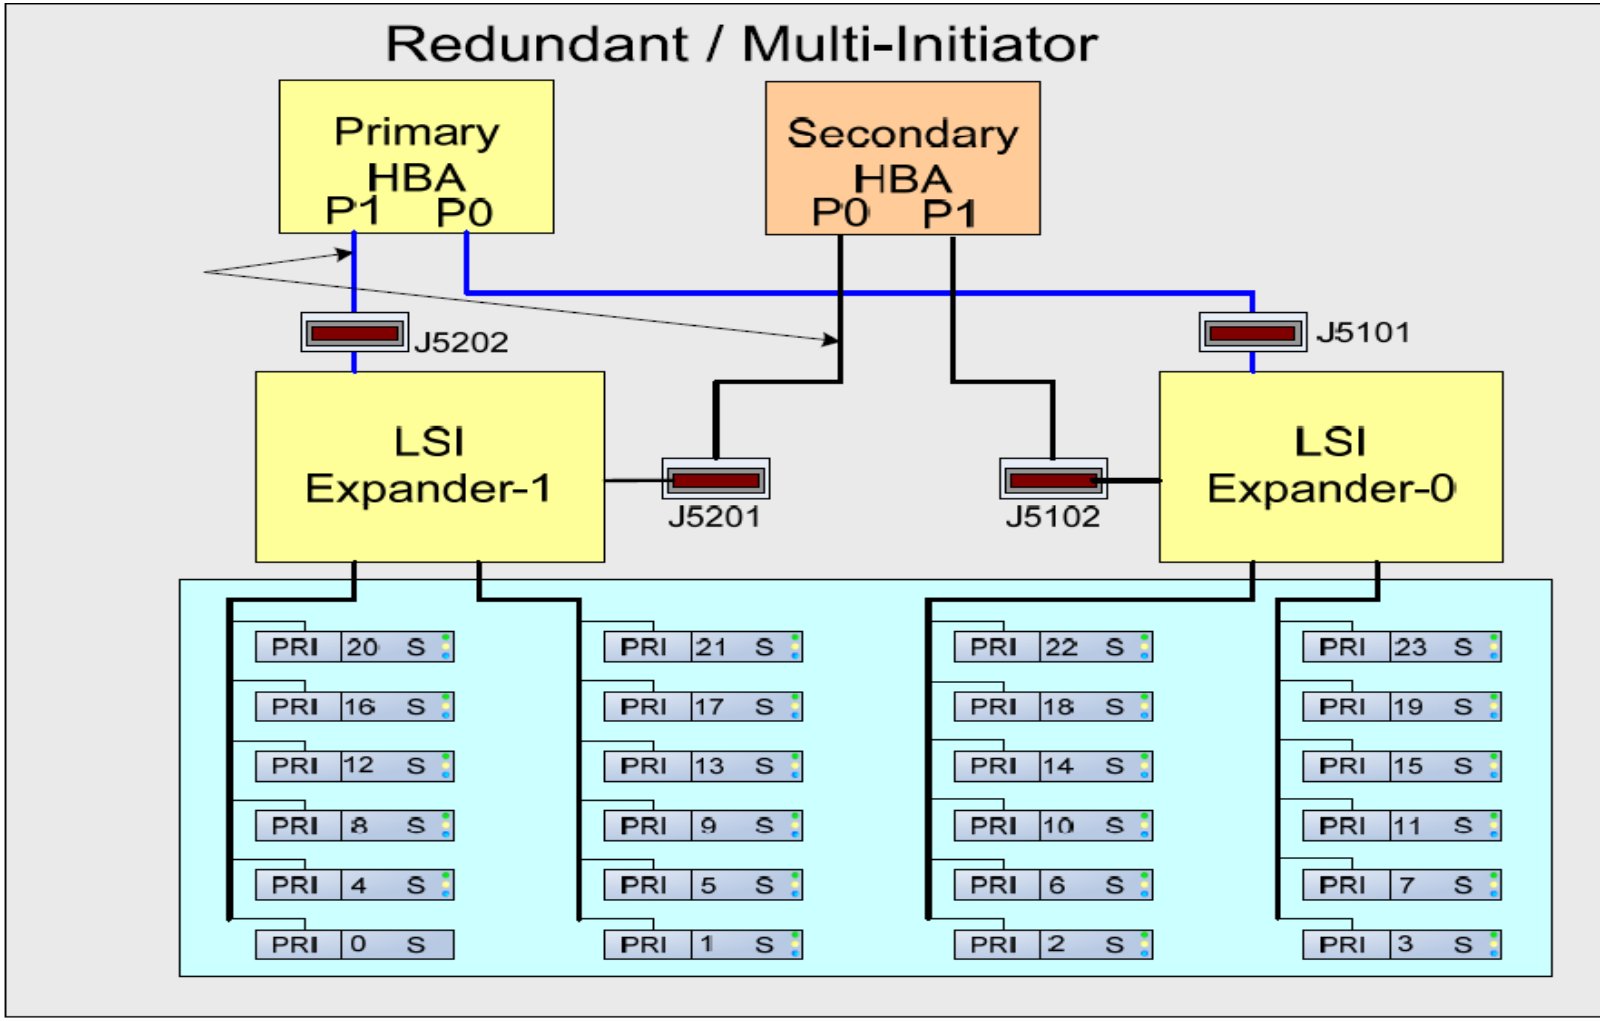

Hardware Architecture

Anatomy of an SSD

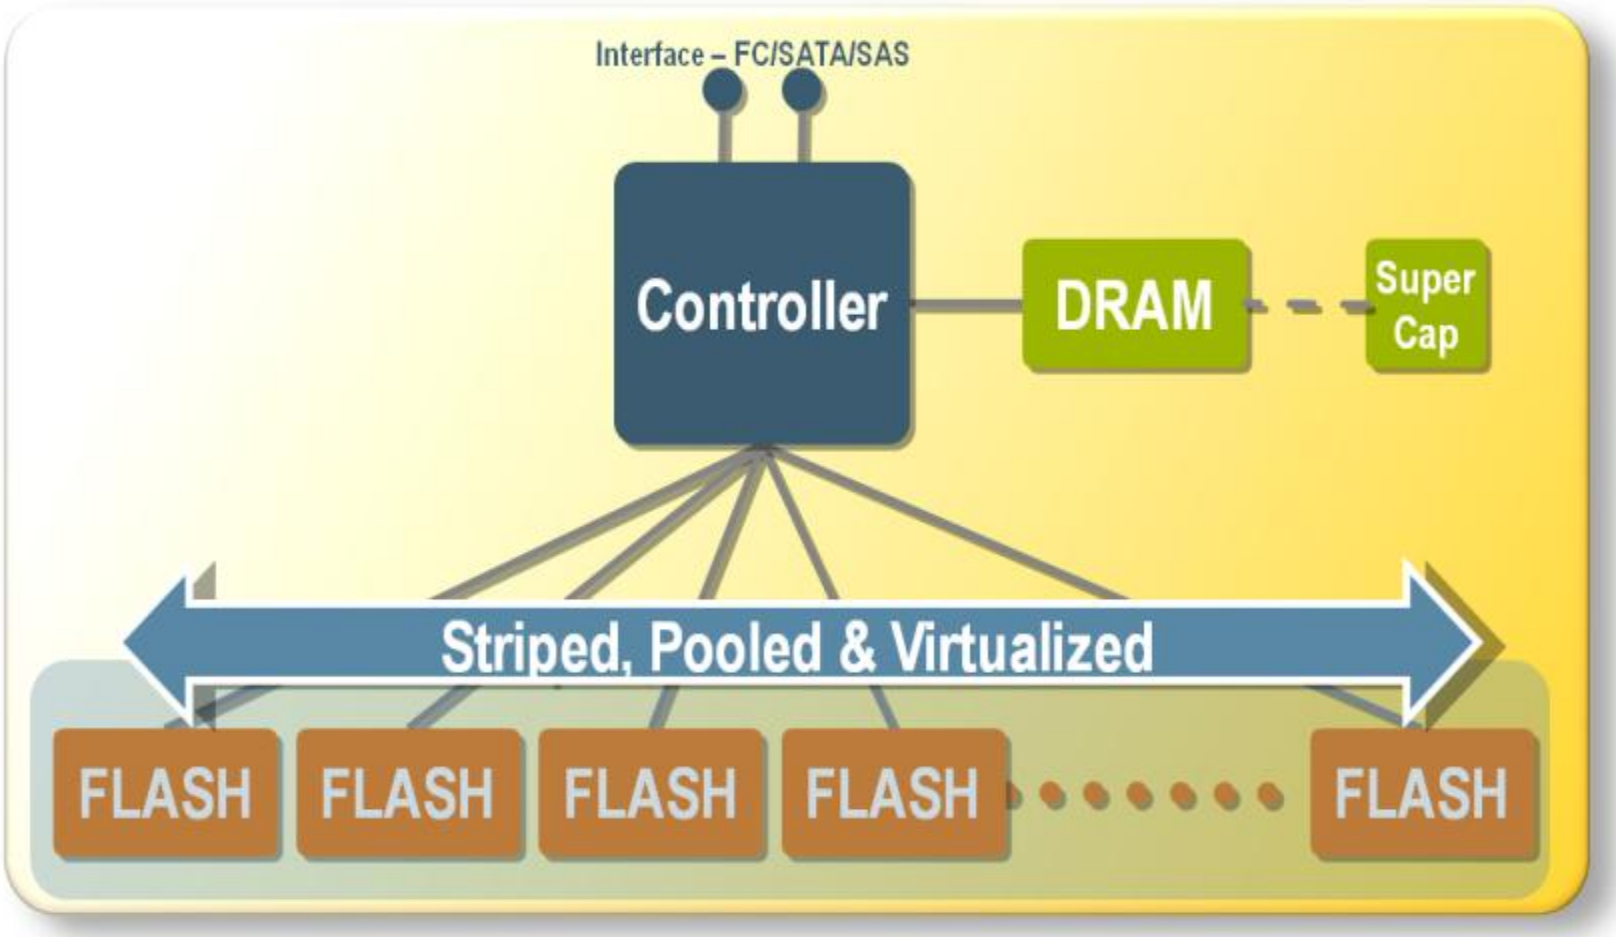

Agenda

- Hardware Architecture
  - integrated server, storage, networking
  - storage architecture
- Software Architecture
  - high availability design
  - appliance manager
- Performance Architecture
  - performance design
  - i/o and system sizing

Software Architecture

Software Stack

- Oracle Database 11*g* Enterprise Edition Release2
  - EE (standalone), RAC One Node, RAC
- Oracle Grid Infrastructure
  - Oracle Clusterware
  - Automatic Storage Management
- Oracle Enterprise Manager Database Control
- Oracle Enterprise Linux 5.5
- Automatic Service Request (phone home)
- Oracle Appliance Manager software

Software Architecture

High Availability Design

- Built on fully redundant hardware – no SPOFs
  - true hot-pluggable disks
  - auto detection, auto correction of hard and soft disk failures
  - disks and server nodes highly serviceable
- Data is triple mirrored
- I/O multipathing is both active/active and HA
- Database is highly available
  - Failover with RAC One Node
  - Active/active with RAC
- Built-in database backups – FRA

Systems Management
Design Objectives

- Easy to deploy
  - Integrate server, storage and networking in ONE chassis
  - Build one button deployment, patching, best practice enforcement
- Easy to manage and maintain
  - Make storage self-discovering, self-configuring, self-managing
  - Build one button patching, upgrade
- Easy to diagnose and support
  - Automate Service Request filing (phone home)
  - Build one button system healthcheck
  - Build one button support diagnostics aggregation and compression
- Easy to scale
  - Enable “pay-as-you-grow” core enablement
  - Adapt (automatically) database to core count

Oracle Appliance Manager

Management Modules

- Configurator Module
- Deployment Module
- Storage Management Module
- Patching Module
- Validation & Diagnostic Tools Module

Oracle Appliance Manager

Configurator Module

- Hides the complexity of setting up the cluster
- GUI-based enquiry collects configuration information
  - Cluster name
  - Domain
  - Networking information
  - Database size
- Validates gathered information
- Configures OS & Networks
- Deploys the Grid Infrastructure & Database

Oracle Appliance Manager

Configurator Module

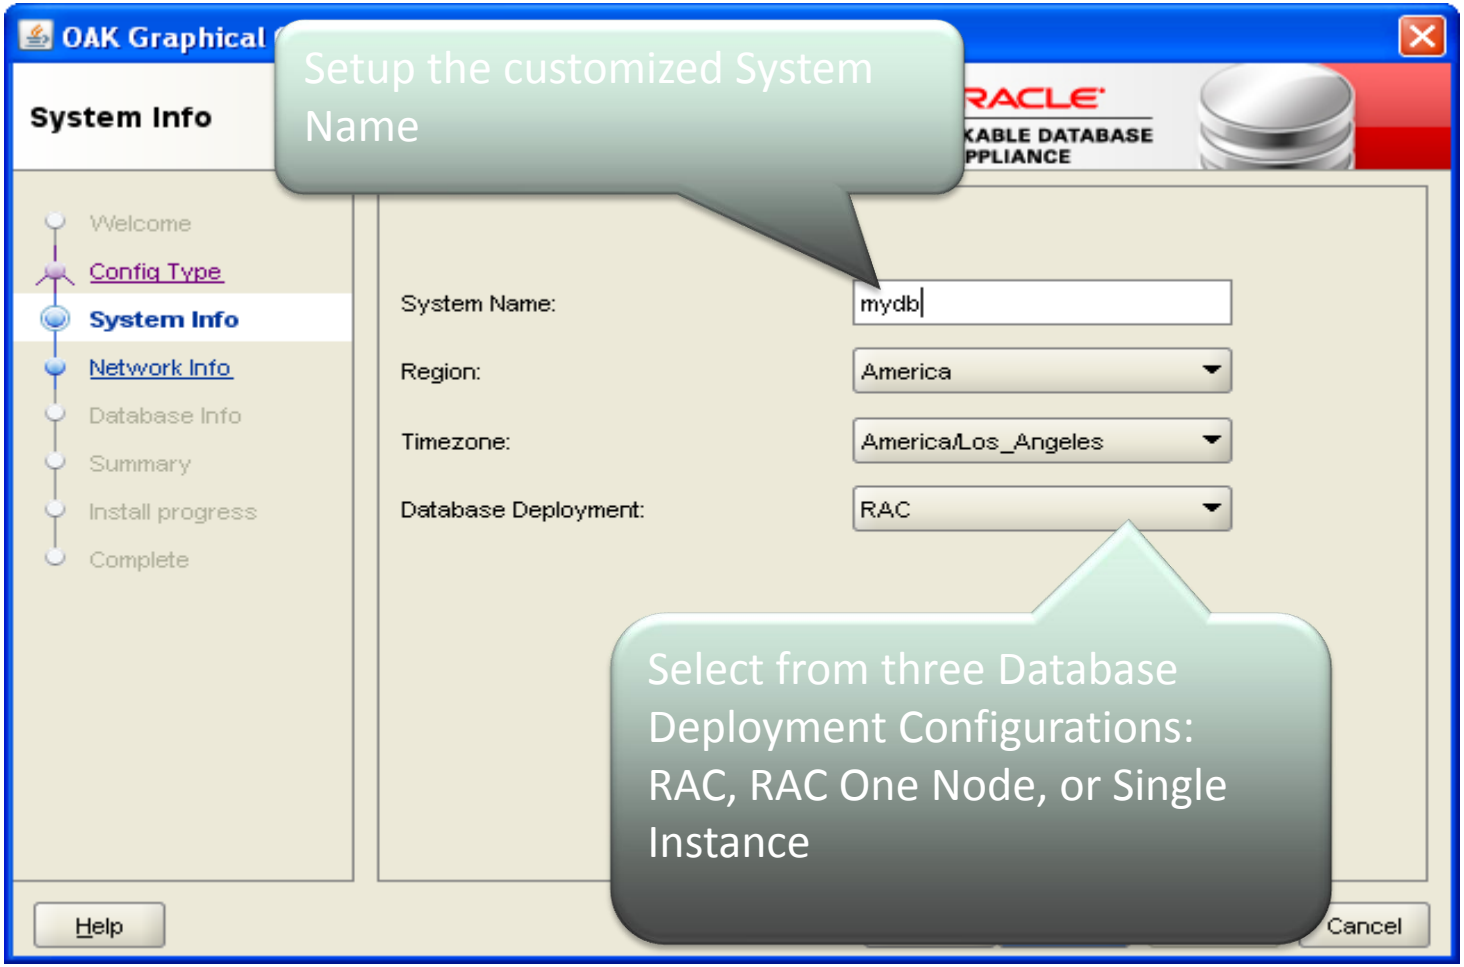

Oracle Appliance Manager
Configurator Module

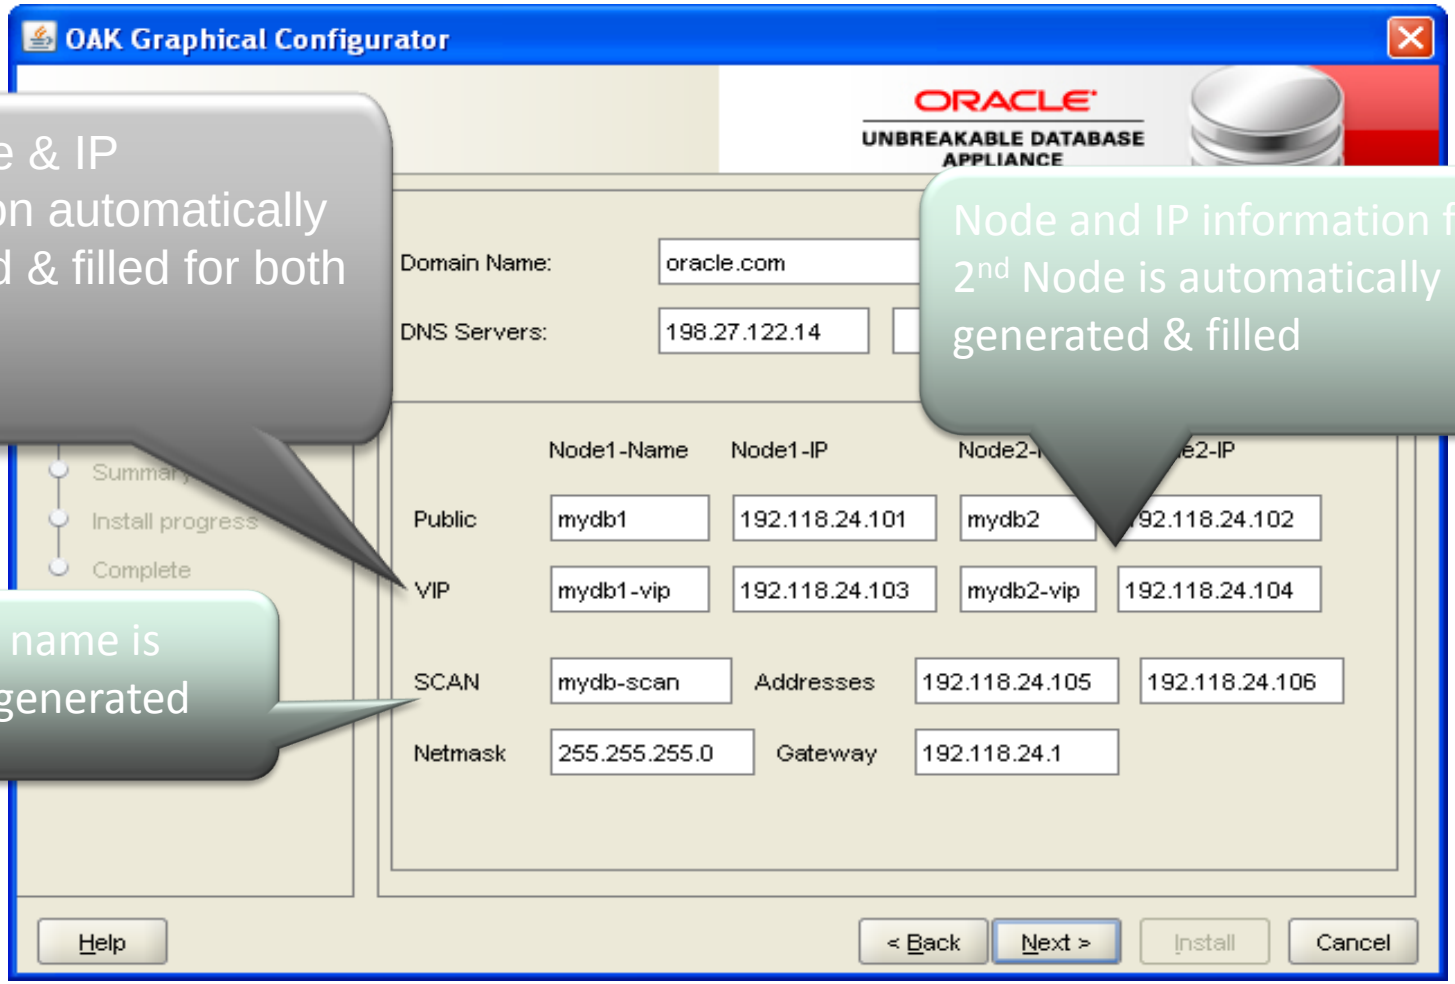

VIP Name & IP information automatically generated & filled for both nodes

Node and IP information for 2nd Node is automatically generated & filled

SCAN name is auto generated

Oracle Appliance Manager

Configurator Module

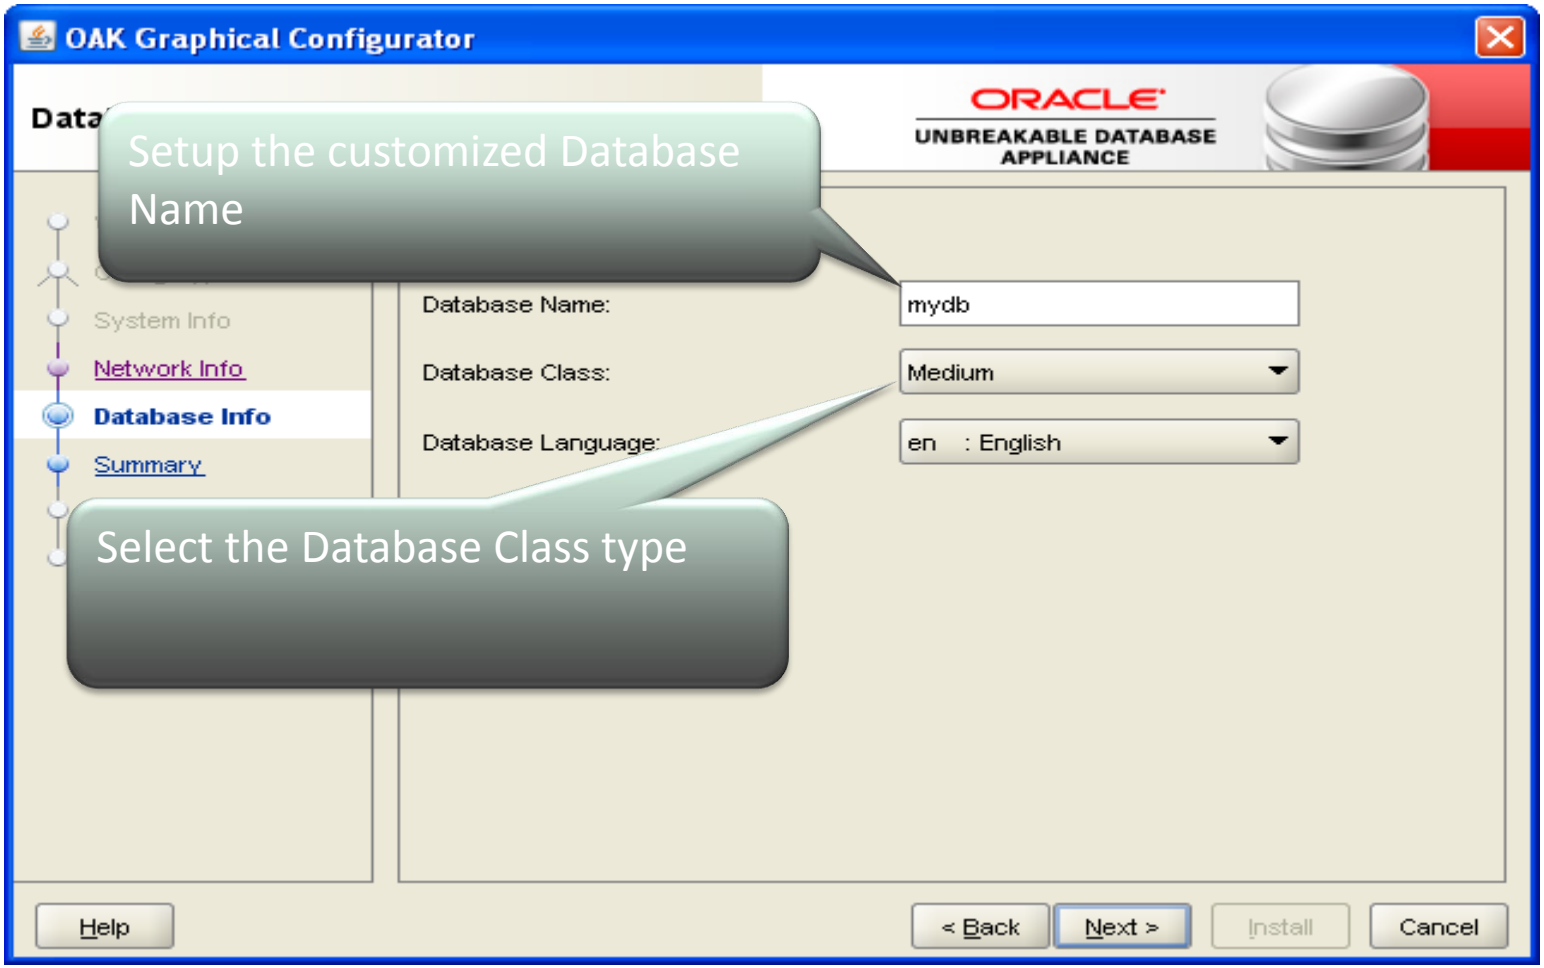

Oracle Appliance Manager
Deployment Module

- Oracle Database Class Options
  - Very Small
    200 Processes, 16 MB log buffer, SGA: 4096-8192 MB,
    PGA: 2048-4096 MB. 1 GB redo log file
  - Small
    400 Processes, 16 MB log buffer, SGA: 8192-16384 MB,
    PGA: 4096-8192 MB, 1 GB redo log file
  - Medium
    800 Processes, 32 MB log buffer, SGA:16384-24576 MB,
    PGA: 8192-12288 MB, 2 GB redo log file
  - Large
    1200 Processes, 64 MB log buffer, SGA: 24576-49152 MB,
    PGA: 12288-24576 MB, 4 GB redo log file
  - Very Large
    2400 Processes, 64 MB log buffer, SGA:49152 MB,
    PGA: 24576 MB, 4 GB redo log file

SGA= System Global Area, PGA= Program Global Area

Oracle Appliance Manager

Deployment Module

- Deploys OS, Oracle Appliance Manager, Grid Infrastructure & Database
- Configures Grid Infrastructure & Oracle Database
- Ensures correct configuration of disks & networks
- Consistent implementation of known best practices
- Configures optimal disk layout for ASM
- Performs initial configuration of disks & ASM DG(s)

Oracle Appliance Manager

Storage Management Module

Cloud Creek SYSTEMS

- Oracle Appliance Manager Daemon (oakd) is started during boot
- Discovers storage subsystem
- Tracks configuration by storing metadata
- Monitors status of disks
- Generates alerts on failures
- Takes corrective action on appropriate events
- Interacts with ASM for complete automation

ORACLE

Oracle Appliance Manager

Storage Management Module

1. oakd monitors the physical state of disks
2. Monitors disk status in ASM
3. Based on events interacts with ASM for corrective actions
4. ASM takes actions as directed by oakd

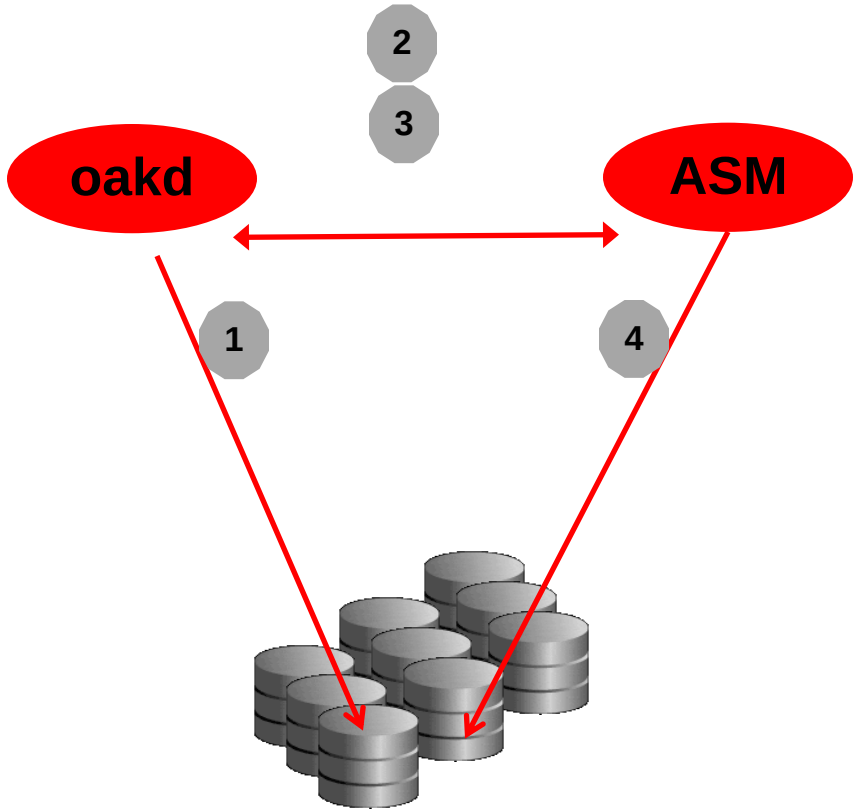

Oracle Appliance Manager

Storage Management Module – command line

- oakcli commands
  - oakcli show - show storage, license, expander, controller, diskgroup, disk
  - oakcli locate - locates a disk
  - oakcli apply - applies the core_configuration_key <key file location>
  - oakcli deploy - deploys the Database Appliance
  - oakcli update - updates the Database Appliance
  - oakcli validate - validates the Database Appliance
  - oakcli manage - manages the oak repository, diagcollect e.t.c
  - oakcli unpack - unpack the given bundle to oak repository
  - oakcli configure - configures the network
  - oakcli copy - copies the deployment config file

Oracle Appliance Manager

Patching Module

- Patching module provides tools to patch OS, Oracle Application Manager modules, Grid Infrastructure(GI), DB
- Provides a single interface and command to patch all the components including OS, firmware, BIOS ,GI and DB
- Patching module will update the repository to reflect the newly installed patches and firmware
- Bundle patches for all components to be patched

Oracle Appliance Manager

Patching Module

Cloud Creek SYSTEMS

Phase 1 → Validate

- System Validation → State of System before applying Patch
- Component Validation → Validate Component Connectivity
- Conflict Resolution → Checks for conflicts

Phase 2 → Apply

- Apply patch in order using specific tool for each component
- Start Component after patch application

Phase 3 → Reporting & Clean up

- Report Patch Success and the current Component state
- Clean up → Clean up of Temp areas and reboot as necessary

Oracle Appliance Manager

Validation and Diagnostic Tools Module

Cloud Creek SYSTEMS

- A set of tools for validation & diagnostics
- Validation tool provides detailed information about the components – both HW & SW
- Diagget tool aggregates all the diagnostics information in one tar file
- Healthcheck can be used to check the health of OS, DB, Clusterware and other comet components to ensure they are healthy and functionally optimally

ORACLE

Agenda

- Hardware Architecture
  - integrated server, storage, networking
  - storage architecture
- Software Architecture
  - high availability design
  - appliance manager
- Performance Architecture
  - performance design
  - i/o and system sizing

Performance Architecture

High Performance Design

- Direct-attached storage
  - Exposes the full IOPS of the disks to both servers
    (unlike traditional storage arrays that expose a fraction of the IOPS to a given server)
- Database-aware I/O path optimization
  - REDO writes to flash
  - DATA writes to HDD outer platter
  - RECO writes to HDD inner platter
- Optimal data layout
  - ASM stripes across all available spindles
- Out-of-the-box fully tuned
  - Storage, OS, network, GI, DB

Performance Architecture - 3rd Party Validation

Cloud Creek SYSTEMS

Conclusion drawn in 3rd party paper:

- outer/inner platter placement separation buys you around 50% IOPS at 10ms response time

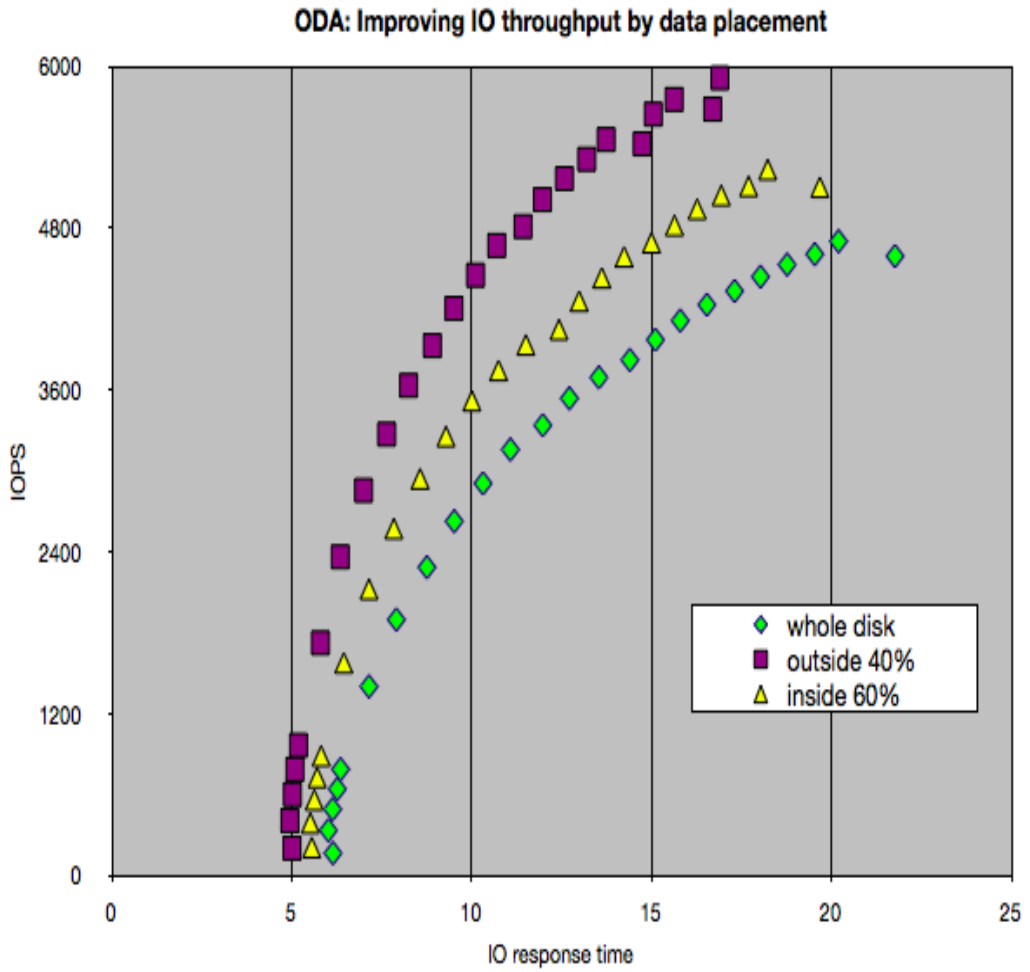

http://www.pythian.com/news/27201/oracle-database-appliance-storage-performance-part-1/

ORACLE

Performance Architecture

HDD Performance - orion

Random	R/W	80/20	60/40	40/60	20/80
8K	IOPS	4,466	4,362	4,310	4,175
	Read(ms)	4.289	4.309	4.271	4.441
	Write(ms)	4.721	4.715	4.717	4.817
16K	IOPS	4,400	4,263	4,258	4,219
	Read(ms)	4.353	4.400	4.352	4.341
	Write(ms)	4.792	4.827	4.769	4.772
32K	IOPS	4,269	4,140	4,114	4,074
	Read(ms)	4.473	4.503	4.490	4.457
	Write(ms)	4.956	4.988	4.938	4.945
64K	IOPS	4,053	3,924	3,882	3,829
	Read(ms)	4.719	4.757	4.739	4.770
	Write(ms)	5.211	5.259	5.237	5.259
Sequential	**R/W**	**MB/s**			
1M	Read	3,329.89			
	Write	3,102.12			

SUMMARY: IOPS = 4,000 for OLTP, BANDWIDTH = 3 GB/s for Data Warehousing

Sizing Model

Design Goals

- Simple, intuitive, easy to use
- Optimal for 80% of workloads
- Flexible and extensible
- Configured out-of-the-box
- Reflects current customer datacenter practices

Sizing Model

Cloud Creek SYSTEMS

	X-Small	Small	Medium	Large	X-Large
CPU(Core) per node	1	2	4	6	12
Mem(GB) per node	8	16	32	48	96
DB Size(GB)	136	273	546	819	1638
Logfile Size(GB)*	1	1	2	4	4
No. DB	12	6	3	2	1
IOPS**	300	600	1300	2000	4000
Throughput(MB/s)***	250	500	1000	1500	3000
Log Generation(MB/s)****	6.83	6.83	13.65	27.30	27.30

* Four Redo log groups per instance

** 8K Random Read/Write, per Database

*** 1M Sequential Read/Write, per Database

**** 15Min Log Switches, per Database

ORACLE

CUSTOMER PERFORMANCE RESULT

ODA & Exadata Performance Results

Cloud Creek SYSTEMS

CallSource® Aim Higher

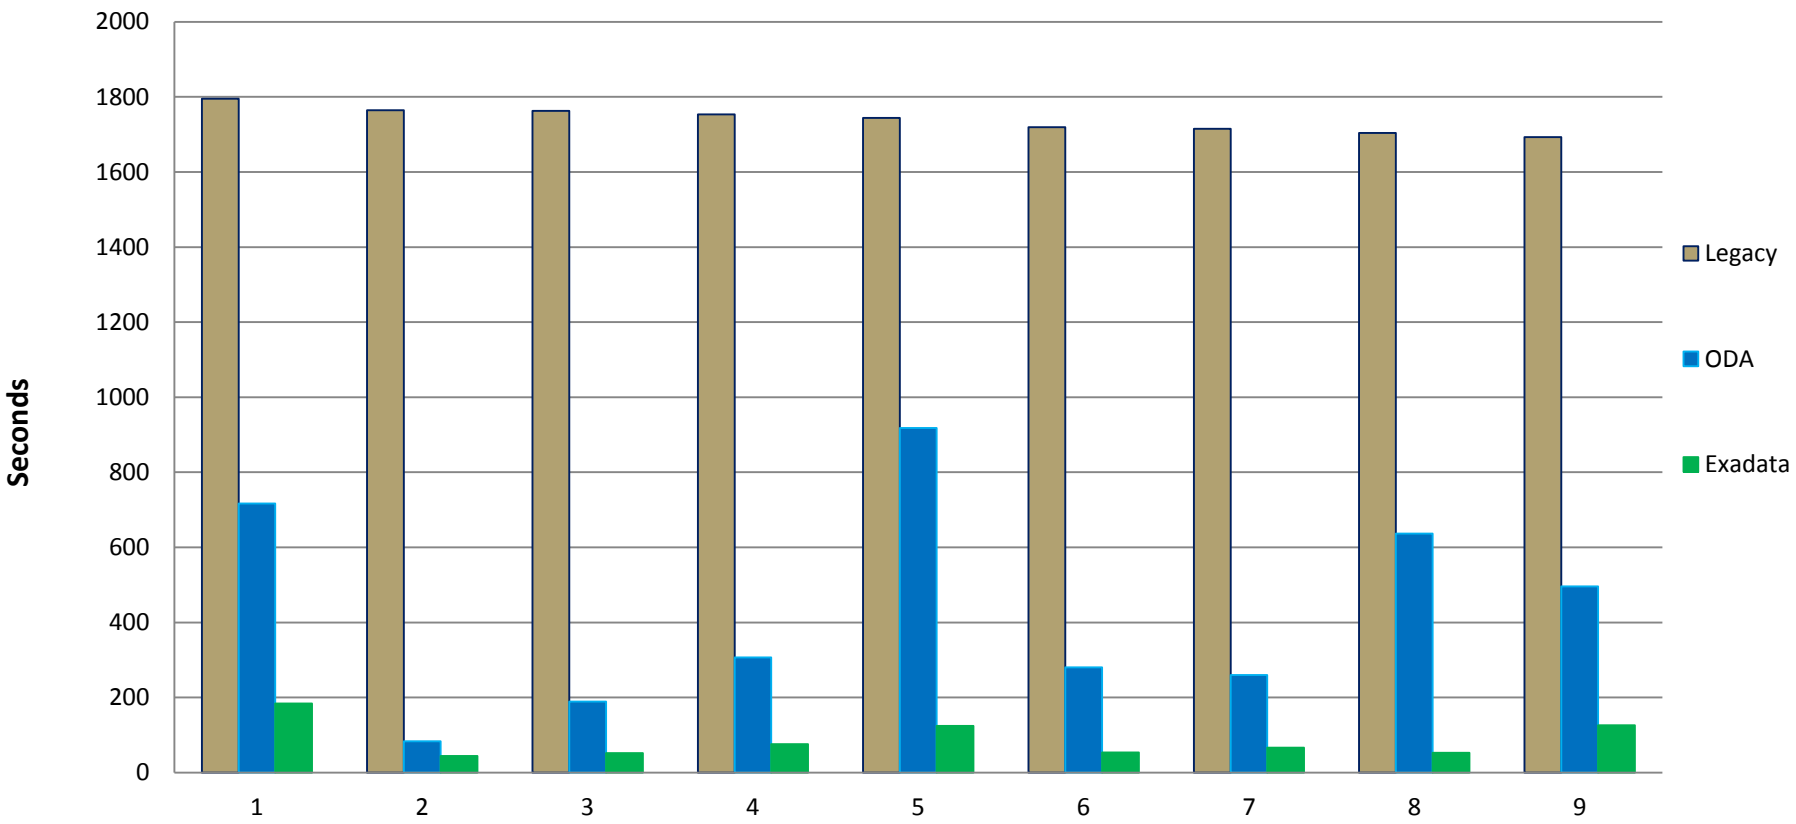

BAR GRAPHS REPRESENT RESPONSE TIMES FOR 9 DIFFERENT CUSTOMER WORKLOADS

ORACLE

HOW DOES THIS PERFORMANCE IMPROVEMENT COMPARE PRICE WISE?

CALLSOURCE
31280 Oak Crest Drive, Suite 3
Westlake Village, CA 91361
(818) 673-4722 fax (888) 261-8990

Purchase Order No. 9192011

Cloud Creek SYSTEMS

PURCHASE ORDER

Vendor

Name	Cloud Creek Systems, Inc.
Attn:	James Yang
email	jyang@cloudcree
Phone	310 740-7977

Ship To

Name	CallSource
Address	31280 Oak Crest Drive, Suite 3
City	Westlake Village St CA ZIP 91361
Phone	(818) 673-4766

Qty	Units	Description	Unit Price	TOTAL
1		Oracle Unbreakable Database Appliance: Sun Fire X4370 M2 server base with 2 server nodes each with 2 Intel® Xeon® X5675 6-core 3.06 GHz processors and twelve 8 GB DDR3-1333 Registered DIMMs	$50,000.00	$50,000.00
1		Power cord: North America and Asia, 2.5 meters, 5-15P plug, C13 connector, 15 A (for factory installation)	$26.00	$26.00
1		Oracle Premier Support	$6,003.12	$6,003.12
1		Installation Service: Servers - Mid - Group I - Exhibit: Standard	$4,200.01	$4,200.01
1		Freight	$59.11	$59.11

SubTotal $60,288.24
Shipping & Handling
Taxes State
TOTAL $60,288.24

Payment Details

- (•) Check
- ( ) Cash
- ( ) Account No.
- ( ) Credit Card

Name
CC #

Shipping Date

Approval

Date 9/23/2011
Order No
Sales Rep
Ship Via

ODA, the consolidated solution, saved this customer an estimated $98,000 in hardware costs alone.

In addition, this customer improved system support by removing more variables from their environment.

By removing these variables… superior performance is attained from a single supplier at almost 1/3rd the price.

ORACLE

ODA Evaluation Program

Cloud Creek SYSTEMS

1. **Data Extraction**
   - Client provides data & test SQL and stored procedures to CCS
2. **CCS Data Load on Evaluation Platform**
   - Oracle sourced, load as is
   - Non-Oracle sourced, migrate data to Oracle
3. **Evaluation Performance Testing & Database Tuning**
   - CCS runs data & SQL processes as is
   - CCS tunes data & SQL processes for performance optimization
4. **Client Accesses Evaluation Platform**
   - Client runs independent tests
5. **Evaluation Tests Completed & Performance Comparison**
   - CCS summarizes results & comparative analysis

*** 2 -4 weeks of Senior Consultant Pro-Services @ $185/hour = $14,800 - $29,600**

ORACLE

Summary

- Hardware Architecture
  - Cluster-in-a-box
  - Resilient storage
- Software Architecture
  - End-to-end high availability design
  - Appliance-like management
- Performance Architecture
  - High performance design
  - Sized right