COLLABORATE 12

TECHNOLOGY AND APPLICATIONS FORUM FOR THE ORACLE COMMUNITY

April 22-26, 2012
Mandalay Bay Convention Center
Las Vegas, Nevada, USA

Rise of the Machines
Session 879

Arup Nanda
Longtime Oracle Technologist

Agenda

- What this is about?
  - The engineered systems
    - Exadata, Exalogic, Exalytics, Sparc Super Cluster, Big Data Appliance, Oracle Database Appliance
  - What are they, where do they fit in, what is best usecase
  - Difference between them
- What will not be covered
  - Greater details on each system
  - Pricing

Systems Covered

- Oracle Database Appliance
- Exadata
- Exalogic
- Super Cluster
- Exalytics
- Big Data Appliance

Oracle Database Appliance

- Traditional database implementation
  - Installation of OS
  - Configuring for Oracle RAC
  - Installation of Oracle, patching
  - Mitigation of issues

ODA is a packaged solution

- 2 nodes
- RAC Database
- Storage Built-in
- 4U Rackmounted Chassis

Specification

- Each node has
  - Two 6-core Intel Xeon processors X5675
  - Cores licensed independently (pay-as-you-grow)
  - 96 GB of RAM
  - The cluster interconnect is via 1GbE (redundant)
  - Six 1GbE and two 10GbE external NIC ports
  - Runs Oracle Enterprise Linux 5.5

The Appliance

- 20X 600 GB 15K SAS drives
  - 12 TB of raw storage
  - (triple mirrored; so 4 TB usable) for database alone.
  - FS on each server are for Linux OS and Oracle software.
- 4X solid state disks of 73GB each for redo logs
- Appliance Manager Software
  - database deployment, patching and health-checks
- Patching the entire appliance I
  - including O/S, firmware, Grid Infrastructure and Database

Good for

- Organizations that don't have a large staff to deploy and administer databases
- Small or medium databases where RAC is essential
- Smaller upfront investment
- It's *not* a mini-Exadata
  - doesn't have all the software Exadata has

Exadata

- Servers, Storage and Network in a Box

A special software runs here – Exadata Storage Server (ESS)

Exadata Smart Flashcache

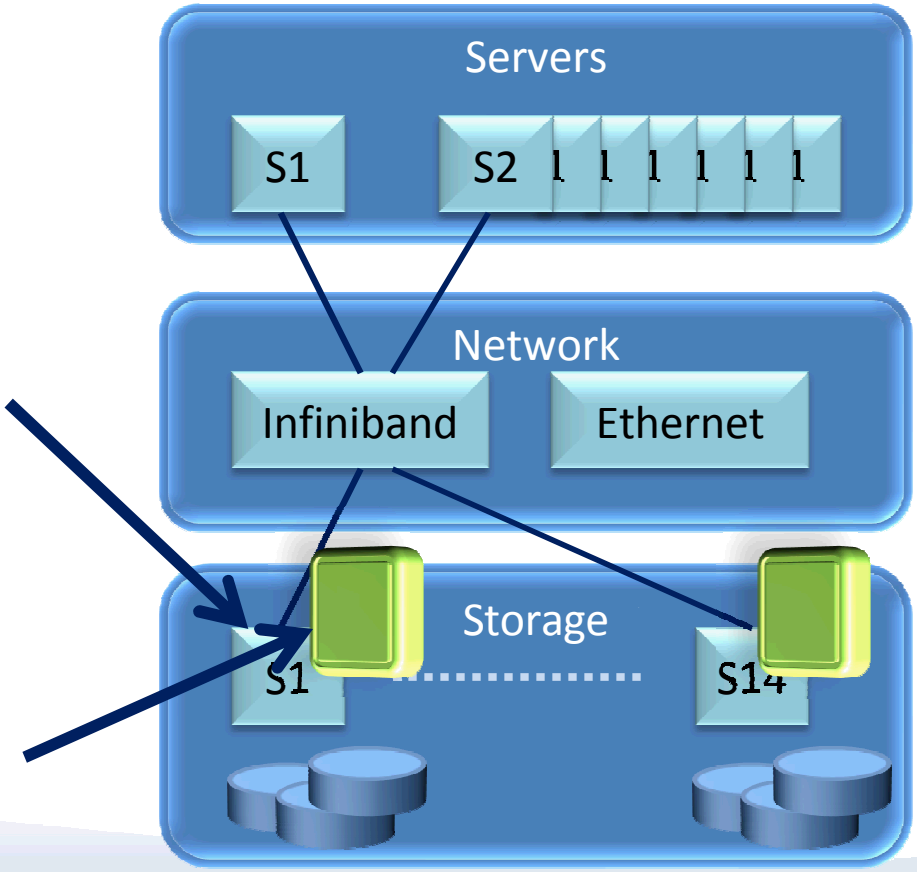

Unique Features

- Smart Scan
  - `select col1 from table1 where col2 = 2`
  - filtering done at storage level
- Storage Indexes
  - Store the min and max values of data on the storage cells
- Flashcache
- Infiniband

Good for

- A super efficient database machine
  - Database alone
  - The filesystem space on compute nodes is very little
  - Usually for other software such as GoldenGate
  - DBFS – a cluster filesystem, for ETL input files
- Great for datawarehouses
  - Smart Scan
- Not so great for OLTP
  - Flashcache helps
- Migration from Oracle based databases is super easy

Specifications

- Comes in 2 models
  - 14 Storage Cells in both models
    - 3 types of disks – 600GB 15K RPM high perf or 2 or 3 TB 7200 RPM high capacity.
    - N0o SAN component or a fiber port to attach an external SAN
  - X2-2
    - 8 Compute Nodes with 96 cores and 768 TB of memory in total
  - X2-8
    - 2 Compute Nodes with 128 cores and 2 TB of memory in total

Three Configurations
- Full Rack
- Half Rack
- Quarter Rack

Storage can be expanded by storage expansion packs
- Full rack
- Half rack
- Quarter rack

Exalogic

- Application Server
- Needs clustering
- Needs storage

Structure

Application Server
- Weblogic
- Fusion Middleware
- Coherence
- Tuxedo
- Hotspot

Storage
- 40 TB SAS
- 4 TB SSD Reads
- 40 GB SSD Writes

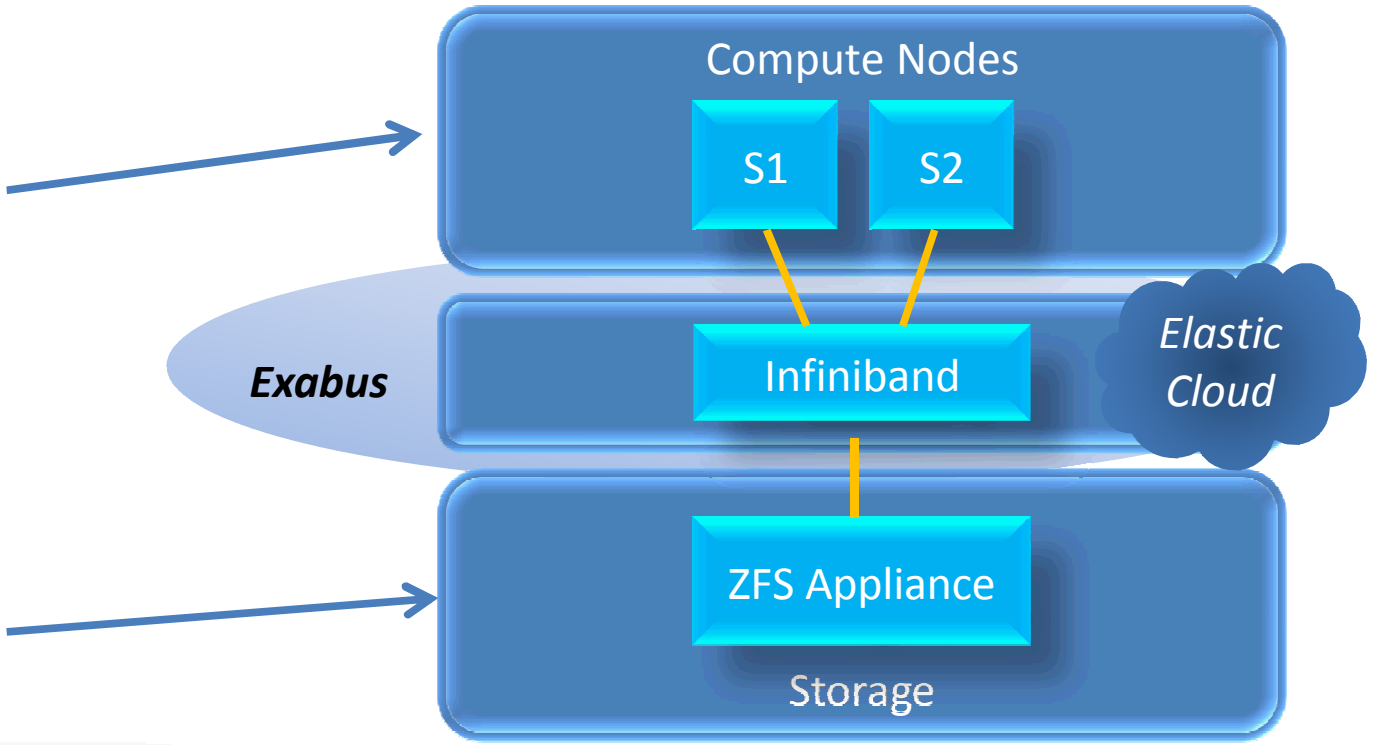

Good for

- It can run
  - any application, e.g. SAP
  - any java applications in the app server
  - coherence caching
- Not for databases
- Backend database can be anything
  - Even non-Oracle
  - Exadata gives it infiniband connectivity.

Typical Use

•Oracle applications
•Java apps
•Weblogic is required
•SAP

Sparc Super Cluster

- The problems with other solutions
  - Two different solutions
    - Exadata – database
    - Exalogic – applications
  - Fiber channel not present
  - Solaris x86 or Linux – less application support

App stack –vs– DB stack

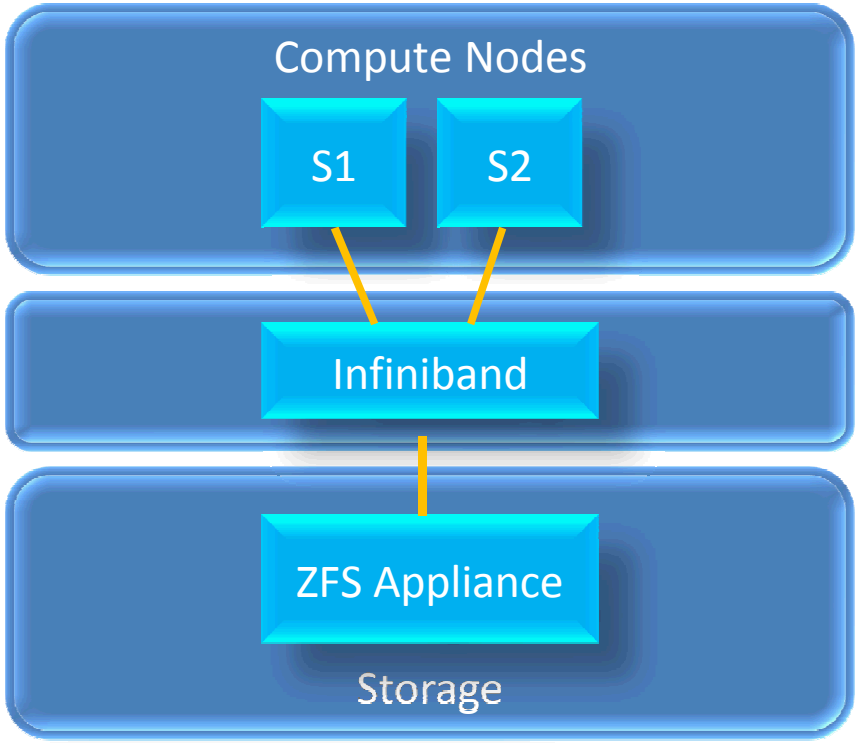

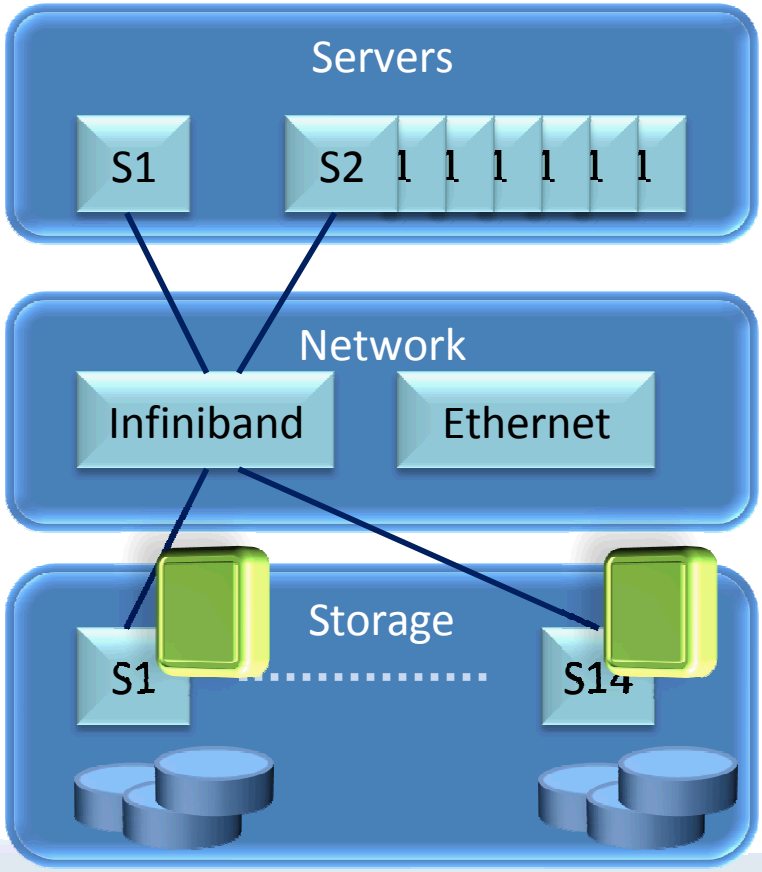

T4 Cluster

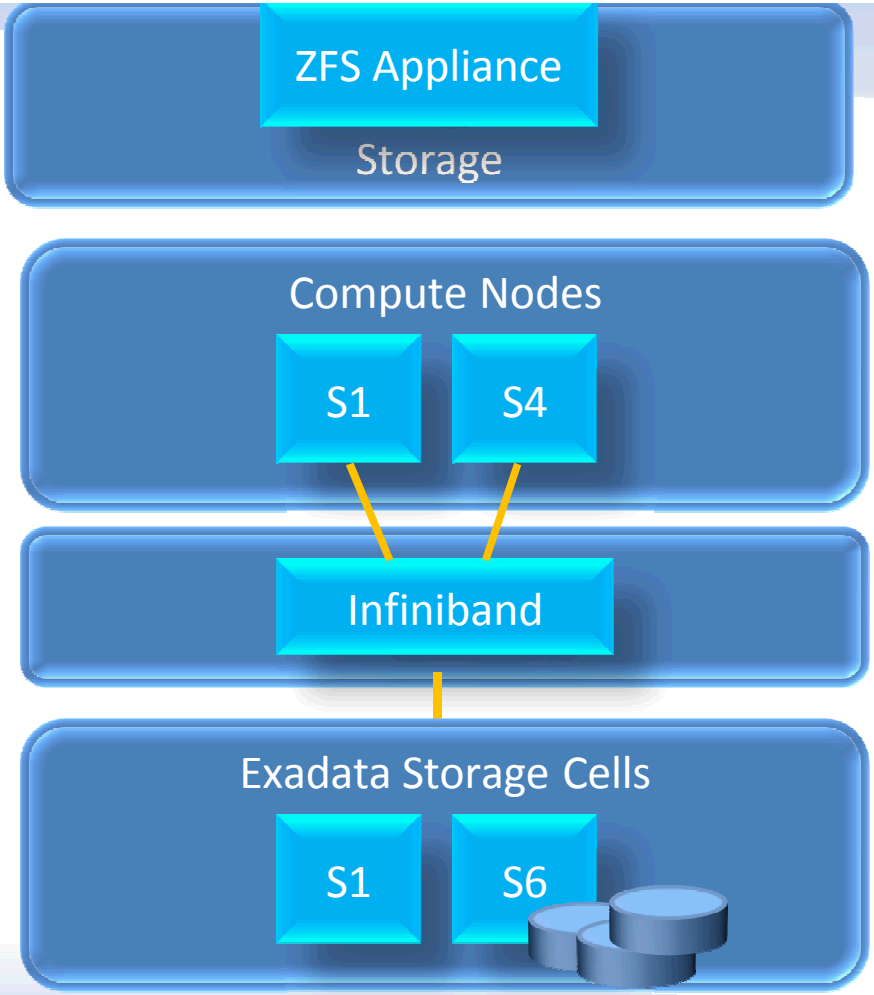
ZFS Appliance
Storage
Compute Nodes
S1
S4
Infiniband
Exadata Storage Cells
S1
S6

Specifications

- Each full rack machine has 4 Compute Nodes, each with:
  - Four 8-core 3 GHz SPARC T4 Processor
  - 1 TB RAM
  - Six 600 GB 10K RPM SAS Drives
  - Two 300 GB Solid State Disks
  - Four InfiniBand Quad Data Rate Ports
  - Four 10G Ethernet

- ZFS Appliance
  - Two 4-core 2.4 GHz Intel Xeon CPUs
  - 24 GB RAM
  - One dual port InfiniBand HCA
  - Two 500 GB SATA Disk
  - Four 512 GB Solid State Disks (read optimized)
  - Twenty 2 TB 7200 RPM SAS-2 disks
  - Four 18 GB Solid State Disks (write optimized)

Best for

- Database and app server in the same rack
- Database migration from a fiber based SAN
- Existing expertise in Solaris
- Racks
  - Full rack – 4 compute nodes
  - Half rack – 2 compute nodes
  - Both have one Sun ZFS 7329 Storage Appliance

Exalytics

- For business intelligence
- Contains
  - Oracle Business Intelligence Foundation Suite
  - Essbase OLAP Engine (formerly Hyperion)
  - TimesTen In Memory Database

But it's not just that

- TimesTen – an in memory database from Oracle
  - SQL based; not object based
  - Can pull from any datasources
  - Operations occur here
  - Adaptive logic to pull data
- Special Tuning
  - OBIEE is tuned specifically for Exalytics
  - Infiniband

Specifications

- A single 3U Sun Fire X4470 M2 Server
- 4X Intel Xeon E7-4800 CPUs each with 10 cores
- 1TB memory
- Local storage of 3.6 TB
- 2 Quad-rate (40 GB/s) InfiniBand ports

Good for

- Organizations with existing Essbase skillset
- For Exadata customers
  - Infiniband connectivity

Big Data Appliance

- Concepts
  - NoSQL Database
  - MapReduce
  - Hadoop
  - Statistical Analysis

Name Value Pairs

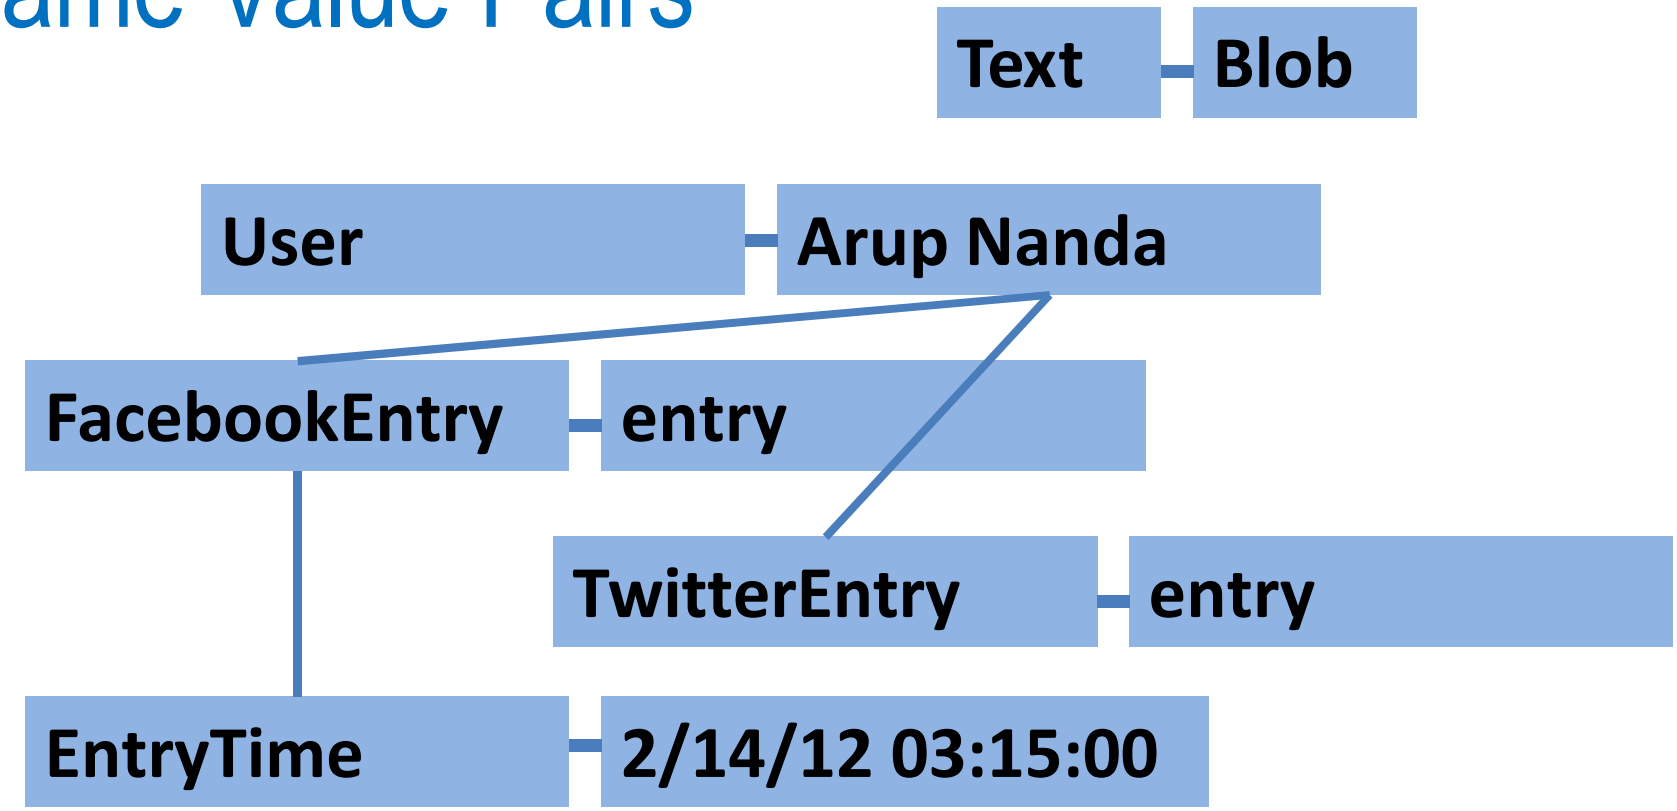

MapReduce

- Map – mapping the values to names
- Reduce – reducing the number of name-value pairs
- Hadoop – an opensource framework
- Analysis – by “R” a statistical package

Big Data Appliance

- Each rack of BDA comes with
  - 18 Nodes – Sun X4270 M2 Servers, each with 48 GB RAM
  - 2 CPUs per node, each with 6 cores (216 cores total)
  - Twelve 2TB disks per node (432 TB raw disk total)
  - Redundant InfiniBand Switches with 10GigE connectivity

- The softwares
  - Oracle Linux and Oracle Hotspot Java VM
  - Open-source distribution of Apache Hadoop
  - Oracle NoSQL Database Enterprise Edition
  - Oracle Loader for Hadoop, that can load data from Hadoop to an Oracle database, for subsequent SQL-based analysis
  - Open-source distribution of R statistical package
  - Data Integrator Application Adapter for Hadoop, for easily specifying MapReduce operations

Best for

- Unstructured or schema-less data
- In-house expertise on NoSQL
- A packaged, engineered all-inclusive solution that can reduce time to market

Conclusion

- Each engineered system is designed for a specific purpose
- Understand the design goals and differences in architecture
- Success depends on the right tool for the job

Thank You!

Session 879

Rise of the Machines

Blog: arup.blogsot.com

Twitter: arupnanda