

PERFORMANCE



GROWTH



INNOVATION



ORACLE DAY 2012

Your Vision, Engineered

# Engineered systems for extreme performance and fast deployment

Dan Stroescu  
Team Leader Infrastructure  
Business Solutions OmniLogic



# Agenda

- Oracle Exadata – extreme performance system
- Unique functions and features
- Components and architecture
- Real environment example



# Oracle Exadata – extreme performance system

## Best platform to run the Oracle database

- Database Machine for Data Warehousing
- Database Machine for OLTP
- Database Machine for database consolidation
- Exadata serves as farm/cloud for databases
  - Large memory enables many databases to be consolidated
  - Extreme performance for complex workloads that mix OLTP, DW, batch, reporting
  - I/O and CPU resource management isolates workloads



# Oracle Exadata – extreme performance system

## What is Exadata?

- Complete Oracle Database system
- Disk storage system
- Processing power, memory, network hardware
- Operating System
- Database software
- Fully installed and configured



# Oracle Exadata Key Features

## What is standard?

- Off-the-shelf hardware
- Intel x86 processors
- Standard disk drives, memory
- Infiniband network adapters and switches
- Linux or Solaris operating system
- Oracle Database software



# Oracle Exadata Key Features

## What is unique?

- Extreme performance
  - 10x to 100x speedup for data warehouses and
  - 20x for OLTP applications are common
- Massively parallel architecture
  - Dynamically scalable
  - Unlimited linear scaling of data bandwidth



# Oracle Exadata Key Features

## What is unique?

- Exadata Smart Scans
  - 10X or greater reduction in data sent to database servers
- Exadata Storage Indexes
  - Eliminate unnecessary I/Os
- Hybrid Columnar Compression
  - Efficient compression increases effective storage capacity and increases user data scan bandwidths by a factor of up to 10X





# Oracle Exadata Key Features

## What is unique?

- Exadata Smart Flash Cache
  - Breaks random I/O bottleneck by increasing IOPs by up to 20X
  - Doubles user data scan bandwidths
- I/O Resource Manager (IORM)
  - Enables storage grid by prioritizing I/Os to ensure predictable performance



# Oracle Exadata – components and architecture

## A new and improved architecture

### The Database Machine:

- Ships less data through pipes
  - Query processing is moved into storage to dramatically reduce data sent to servers while offloading server CPUs
- Has more pipes
  - Modular storage “cell” building blocks are organized into a massively parallel grid
  - Bandwidth scales with capacity
- Has bigger pipes
  - Infiniband interconnect transfers data 5x faster than fiber channel



# Oracle Exadata – components and architecture

## Smart Scan Offload Processing

- Exadata Storage Servers implement smart scans to greatly reduce the data that needs to be processed by database hosts
  - Offload predicate evaluation
  - Only return relevant rows and columns to host
  - Join filtering
- Data reduction is usually very large
  - 10x data reduction is common
- Data transfer is completely transparent
  - Even if a cell or a disk fails during a query



# Oracle Exadata – components and architecture

## I/O Elimination with Storage Index

Storage Indexes in the Database Machine memory contain summary information about data in tables.


- One index entry is typical for every megabyte of disk space
- Each entry holds minimum and maximum values of columns used in the where clauses
- Storage Indexes are created and maintained transparently and automatically

<u>Table</u>				<u>Index</u>
A	B	C	D	
	1			} <b>Min B = 1</b> <b>Max B = 3</b>
	2			
	3			
	5			} <b>Min B = 2</b> <b>Max B = 8</b>
	8			
	2			

# Oracle Exadata – components and architecture

## I/O Resource Management

Coordination and prioritization between different groups or classes of work within a database and between databases



# Oracle Exadata – components and architecture

## Massively Parallel Storage Grid

Exadata Storage Servers are organized into a massively parallel storage grid:

- Scalable
  - The grid is capable of scaling to hundreds of Storage Servers
  - Data is automatically distributed across Storage Servers by Automatic Storage Management (ASM)
  - Data is transparently redistributed when Storage Servers are added or removed
  - Data bandwidth scales linearly with capacity
- Available
  - Data is mirrored across Storage Servers
  - Failure of disk or Storage Server is transparently tolerated
- Simple
  - The grid works transparently, with no application changes necessary


# Oracle Exadata – components and architecture

## Exadata Scalable Architecture



# Oracle Exadata – components and architecture

## Exadata Database Machine Core Hardware Components






# Oracle Exadata – real environment example

## Existing Infrastructure

IBM Power 595

- 2 database nodes (Real Application Cluster)
  - 12 Power 6 64-bit processors (4.2GHz)
  - 100GB RAM memory
  - Oracle Database 10g Enterprise Edition Release 10.2.0.3.0 - 64bit with the Partitioning, OLAP and Data Mining options
  - 11TB Database capacity stored on IBM DS8000 series Fiber Channel storage
- 

# Oracle Exadata – real environment example

## Test Infrastructure

Oracle Exadata X2-2 Quarter Rack

- 2 database servers
- 12 Intel 64-bit processors (3.06GHz)
- 96GB RAM memory
- Oracle Database 10g Enterprise Edition Release 10.2.0.3.0 - 64bit with the Partitioning, OLAP and Data Mining options
- 11TB Database capacity stored on 3 storage servers (36 Intel CPU cores for SQL processing), Infiniband connection
- 1.1 TB Exadata Smart Flash Cache

# Oracle Exadata – real environment example

## Test Results

Test number	Query time (existing)	Query time (Exadata)	Performance increase factor
Query 1	00:03:40	00:00:10	22X
Query 2	00:04:19	00:00:16	16X
Query 3	00:12:34	00:00:22	34X
Query 4	00:01:54	00:00:02	57X
Query 5	00:01:23	00:00:02	42X
Query 6	00:05:17	00:00:01	317X
Query 7	00:47:14	00:02:08	22X
Query 8	00:01:52	00:00:03	37X
Query 9	01:58:16	00:00:03	<b>2365X</b>
Query 10	00:05:03	00:00:03	101X
Query 11	00:38:54	00:03:30	11X

# Q & A

