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 remain at the sole discretion of Oracle.



Exadata V2 Sun Oracle Database Machine

Exadata Database Machine



Version 1

- World's Fastest Machine for Data Warehousing
- Extreme Performance for Sequential I/O
- 10x Faster than other Oracle D/W Systems

Version 2

- World's Fastest Machine for OLTP
- Extreme Performance for Random I/O
- 5x Version 1 Data Warehousing Performance
- Dramatic new Exadata Software Capabilities

The Architecture of the Future Massively Parallel Grid



Data Warehousing and OLTP

© 2009 Oracle Corporation

Exadata Storage Server Building Block

- Hardware by Sun
- Software by Oracle



(intel) Xeon* 5500

- High-performance storage server built from industry standard components
- 12 disks 600 GB SAS, or 2TB SATA
- • 2 Xeon quad-core Nehalem



- Dual ported 40 Gb/sec InfiniBand
- __ 4 96 GB Flash Cards
 - Runs at full disk and flash bandwidth

Sun Oracle Database Machine

- Grid is the architecture of the future
- Highest performance, lowest cost, fault tolerant, scalable on demand
- Database Machine is an engineered, optimized, standardized, and tested grid for Oracle database – with intelligent storage



Sun Oracle Database Machine

Extreme Performance

Oracle Database Server Grid

- Millions of transactions per minute
- Tens of millions of queries
 per minute
- Billions of rows per minute

InfiniBand Network

 880 Gb/sec aggregate throughput



Exadata Storage Server Grid

- 21 GB/sec disk bandwidth
- 50 GB/sec flash bandwidth
- 1 million I/Os per second (8K)

Start Small and Grow



Basic	Quarter	Half	Full
System	Rack	Rack	Rack

Scale Performance and Capacity



Scalable

- Scales to 8 rack database machine by just adding wires
 - More with external InfiniBand switches
- Scales to hundreds of storage servers
 - Multi-petabyte databases

- Redundant and Fault Tolerant
 - Failure of any component is tolerated
 - Data is mirrored across storage servers

Drastically Simplified Deployments



- Database Machine eliminates the complexity of deploying database systems
 - Months of configuration, troubleshooting, tuning
- Database Machine is ready on day one
 - Pre-built, tested, standard, supportable configuration
 - Runs existing applications unchanged

Months to Days • Extreme performance out of the box

Best Machine for Data Warehousing



Best Data Warehouse Machine



- Massively parallel high volume hardware to quickly process vast amounts of data
 - Exadata runs data intensive processing directly in storage
- Most complete analytic capabilities
 - OLAP, Statistics, Spatial, Data Mining, Real-time transactional ETL, Efficient point queries
- Powerful warehouse specific optimizations
 - Flexible Partitioning, Bitmap Indexing, Join indexing, Materialized Views, Result Cache



Dramatic new warehousing capabilities

Storage Data Bandwidth Bottleneck



- Today, database performance is limited by storage
 - Storage systems limit data bandwidth from storage to servers
 - Storage Array internal bottlenecks
 - SAN bottlenecks
- Bandwidth limits severely restrict performance for:
 - Data warehouse queries and loads
 - OLTP batch operations and reports
 - Multimedia and document processing

Exadata Smart Storage Breaks Data Bandwidth Bottleneck

- Oracle addresses data bandwidth bottleneck 3 ways
- <u>Massively parallel storage grid</u> of high performance Exadata storage servers (cells)
 - Scales performance and data bandwidth with data volume
 - 40 Gb Infiniband provides maximum bandwidth
- <u>Data intensive processing</u> runs in Exadata storage
 - Queries run in storage as data streams from disk
 - Offloads database server CPUs
- Columnar compression reduces data volume up to 10x
 - Exadata Hybrid Columnar Compression
 - 10x lower cost, 10x higher performance

Exadata Storage Cells



-44				- Ch
	The second		- I have been	
0			II IIIIII	
1			II DOUGON	
1	MARCHINE H	-	- II monorogan	TI MANAGEMENT





Exadata Database Processing in Storage

- Exadata storage servers implement data intensive processing in storage
 - Row filtering based on "where" predicate
 - Column filtering
 - Join filtering
 - Incremental backup filtering
 - Storage Indexing
- New Scans on encrypted data
 - Data Mining model scoring
 - 10x reduction in data sent to DB servers is common
 - No application changes needed
 - Processing is automatic and transparent
 - Even if cell or disk fails during a query







Simple Query Example



Exadata is Smart Storage

Compute and Memory Intensive Processing



- Exadata cell is smart storage, not a database node
 - Storage remains an independent tier

Database Servers

Perform complex database processing such as joins, aggregation, etc.

Exadata Cells

- Search tables and indexes filtering out data that is not relevant to a query
- Cells serve data to multiple databases <u>enabling</u>
 <u>OLTP and consolidation</u>
- Simplicity, and robustness of storage appliance

Exadata Hybrid Columnar Compression

- Data is grouped by column and then compressed
- <u>Query Mode</u> for data warehousing
 - Optimized for speed
 - 10X compression typical
 - Scans improve proportionally
- <u>Archival Mode</u> for infrequently accessed data
 - Optimized to reduce space
 - 15X compression is typical
 - Up to 50X for some data



Exadata Hybrid Columnar Compression How it works

Compression Unit C C \mathbf{O} olumn olumn olumn N ω 4x to 50x

- Tables are organized into sets of a few thousand rows
 - Compression Units (CUs)
- Within CU, data is organized by <u>column</u>, then compressed
 - Column organization brings similar values close together, enhancing compression
- Useful for data that is bulk loaded and queried
 - Update activity is light
- Compared to best conventional algorithms Gzip, Bzip2
 - Typically 2X the compression, 10X the performance
- Exadata servers offload filtering, projection, etc. for scans on compressed data
 - Indexed accesses return compressed blocks to database so buffer cache benefits from compression

Real-World Compression Ratios Oracle Production E-Business Suite Tables



Hybrid Columnar Comparisons



- Exadata Hybrid Columnar Compression is a second generation columnar technology combining the best of row and column formats
 - Best compression matching full columnar
 - Excellent scan time Compression provides 10x speedup
 - After 10x I/O reduction, most queries become CPU bound
 - Good single row lookup no full columnar "cliff"

Flash

- Flash storage more than doubles scan throughput
 - 50 GB/sec
- Combined with Hybrid Columnar Compression
 - Up to 50 TB of data fits in flash
 - Queries on compressed data run up to
 500 GB/sec



Exadata Data Mining





- All data mining scoring functions are offloaded to Exadata
 - Up to 10x performance gains
 - Reduced CPU utilization on Database Server

select cust_id
from customers
where
prediction_probability(churnmod, 'Y' using *) > 0.8

Exadata Storage Index Transparent I/O Elimination with No Overhead



- Exadata Storage Indexes maintain summary information about table data in memory
 - Store MIN and MAX values of columns
 - Typically one index entry for every MB of disk
- Eliminates disk I/Os if MIN and MAX can never match "where" clause of a query
- Completely <u>automatic and transparent</u>

Select * from Table where B<2 - Only first set of rows can match

Storage Index with Partitions Example

Orders Tak	ole		
Order#	Order_Date Partitioning Column	Ship_Date	ltem
1	2007	2007	
2	2008	2008	
3	2009	2009	

- Queries on Ship_Date do not benefit from Order_Date partitioning
 - However Ship_date and Order# are highly correlated with Order_Date
 - e.g. Ship dates are usually near Order_Dates and are never less
- Storage index provides partition pruning like performance for queries on Ship_Date and Order#
 - Takes advantage of ordering created by partitioning or sorted loading

Benefits Multiply



Data is 10x Smaller, Scans are 2000x faster

DBFS - Scalable Shared File System

- Database Machine comes with DBFS shared Linux file system
 - Shared storage for ETL staging, scripts, reports and other application files
- Files stored as SecureFile LOBs in database tables stored in Exadata
 - Protected like any DB data mirroring, DataGuard, Flashback, etc.
- 5 to 7 GB/sec file system I/O throughput



Best Machine for OLTP



Best OLTP Machine

PEOPLESOFT ENTERPRISE

ORACLE

E-BUSINESS SUITE



- Only Oracle runs real-world business applications "on the Grid"
- Unique fault-tolerant scale-out OLTP database
 - RAC, Data Guard, Online Operations
- Unique fault-tolerant scale-out storage suitable for OLTP
 - ASM, Exadata
- Dramatic New OLTP Capabilities

Exadata Flash Extreme Performance for Random I/O





- Sun Oracle Database Machine has
 5 TB of flash storage
 - 4 high-performance flash cards in every Exadata Storage Server
- Smart Flash Cache caches hot data
 - Not just simple LRU
 - Knows when to avoid caching to avoid flushing cache
 - Allows optimization by application table

Oracle is the First Flash Optimized Database

Exadata Smart Flash Cache Extreme Performance



5X More I/Os than 1000 Disk Enterprise Storage Array

- Running SQL, the Database Machine flash cache achieves:
- 20x more random I/Os
 - Over 1 million per second (8K)
- 10x better I/O response time
 - Sub-millisecond
- Greatly Reduced Cost
 - 10x fewer disks needed for I/O
 - Lower Power

Architected for Flash

- Many systems have added flash
- Exadata fully leverages flash

Scale-Out Storage	No bottlenecks to scaling I/O
InfiniBand	Highest throughput, lowest latency
Queries Offload to Storage	Key to using full flash bandwidth Even InfiniBand can't send 50GB/sec
Flash PCI card	Avoids disk controller bottlenecks. Cards in storage enable HA, RAC
Compression	Multiply flash capacity 3x for OLTP, 10x for Data Warehousing. Also multiplies data scan rates
Flash Cache	Performance of flash, cost of disk Optionally specify placement of tables



Complete, Open, Integrated Availability Maximum Availability Architecture



- Protection from
 - Server Failures
 - Storage Failures
 - Network Failures
 - Site Failures

- Real-time remote standby open for queries
- Human error correction
 - Database, table, row, transaction level
- Online indexing and table redefinition
- Online patching and upgrades

Complete, Open, Integrated Security



Best Machine for Consolidating Databases



Special-Purpose Hardware Configurations



- Biggest driver of ongoing cost is system sprawl
 - Hardware configurations are isolated and customized to specific applications
 - Special systems for high performance, fault tolerance, low cost, data warehousing
- Database machine enables database consolidation
 - Better performance than high-end systems
 - With fault tolerance built in
 - At high volume price
 - Runs any combination of workloads with extreme performance
 - Warehouse oriented bulk data processing
 - OLTP oriented random updates
 - Real time BI against transactional data

Consolidate Database Storage



- Exadata and ASM allow all storage servers to be shared across databases
- Shared Configuration
 - Advanced ASM data striping spreads every database
 across all storage servers
 - Eliminates hot-spots and captive unused space
 - Full storage grid performance available to all databases
 - Database or cluster level storage security
- Predictable Performance
 - Exadata I/O resource manager prioritizes I/Os to ensure predictable performance
 - At user, job, application, or database level
 - No need for isolated storage islands

Consolidate Database Servers



- Many databases can run on Database Machine servers
- Shared Configuration
 - Applications connect to a database service that runs on one or more database servers
 - Services can grow, shrink, & move dynamically
 - Large databases can **span nodes** using RAC
 - Multiple small databases can run on a single node

Predictable performance

- Instance caging provides predictable CPU resources when multiple databases run on the same node
 - Restricts a database to subset of processors

Conclusion



Sun Oracle Database Machine Extreme Performance for all Data Management



- Best for Data Warehousing
 - Grid architecture with scalable storage offload
 - 10x table compression
 - Up to 50 TB of compressed data in super-fast flash
- Best for OLTP
 - Only database that scales real-world applications on grid
 - Smart flash cache runs 1 million IOs per second
 - Compressed for OLTP, up to 15 TB in Flash
 - Up to 50x compression for archival data
 - Secure, fault tolerant
- Best for Database Consolidation
 - Only database machine that runs and scales all workloads
 - Predictable response times in multi-database, multiapplication, multi-user environments

Sun Oracle Database Machine Hardware Improvements

- Same architecture as Exadata V1 Database Machine
 - Same number and type of Servers, CPUs, Disks

