

**ORACLE®**

ORACLE®

# Exadata Technical Overview

Last updated – March 19, 2014



# Exadata Database Machine



- **The ultimate platform for all database workloads**
  - OLTP, Warehousing, Database as a Service
- **Most advanced hardware**
  - Fully scale-out servers and intelligent storage with unified InfiniBand connectivity and PCI flash
- **Most advanced software**
  - Database optimized compute, storage, and networking algorithms dramatically improve performance and cost
- **Standardized, optimized, hardened end-to-end**

# 1000s of Deployments at Leading Companies

Half are Warehouses, Half are OLTP or Mixed Workloads



- Petabyte Warehouses
- Online Financial Trading
- E-Commerce Sites
- Consolidation of 100s of Databases

# Exadata for Packaged Applications

GLENCORE  
INTERNATIONAL AG



CommonwealthBank



NETAS

SWN  
Southwestern Energy



AFG  
AUSTRALIAN FINANCE GROUP



MCF  
Mangalore Chemicals  
& Fertilizers Limited



KIA MOTORS



Santander



HUAWEI

Lyreco  
You're our number One!



CARLSON



Radio Systems  
CORPORATION



IQNavigator



Alpha Natural Resources



Haier

# Exadata Architecture

Complete | Optimized | Standardized | Hardened Database Platform



Fully Redundant

## Standard Database Servers

- 8x 2-socket servers → 192 cores, 2TB DRAM or
- 2x 8-socket servers → 160 cores, 4TB DRAM



## Unified Ultra-Fast Network

- 40 Gb InfiniBand internal connectivity → all ports active
- 10 Gb or 1 Gb Ethernet data center connectivity

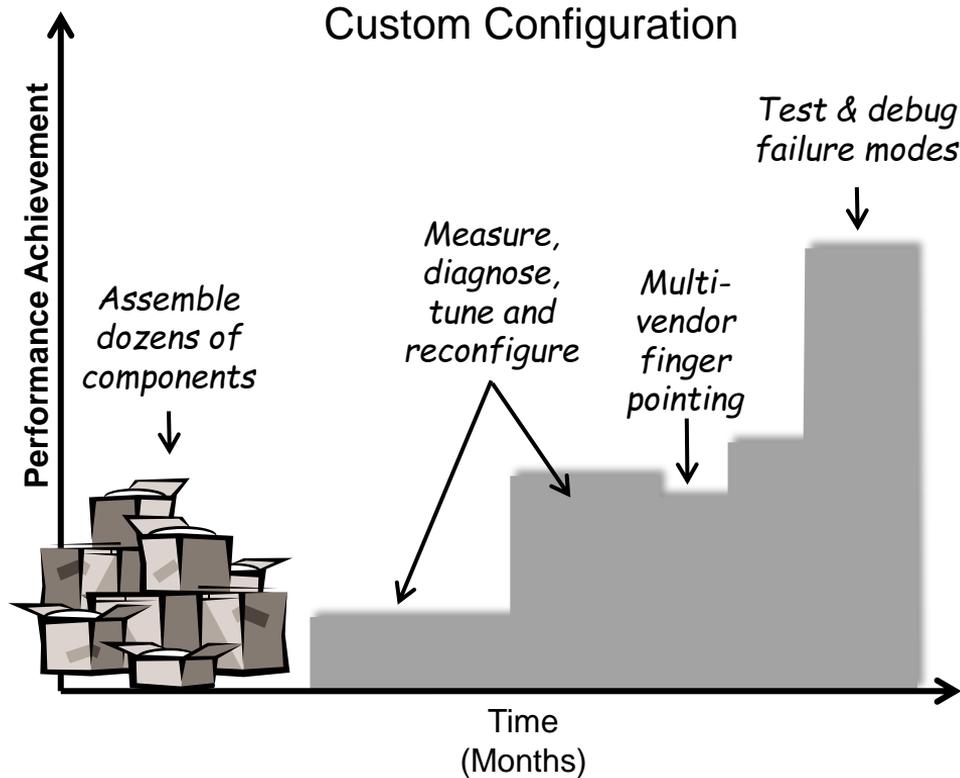
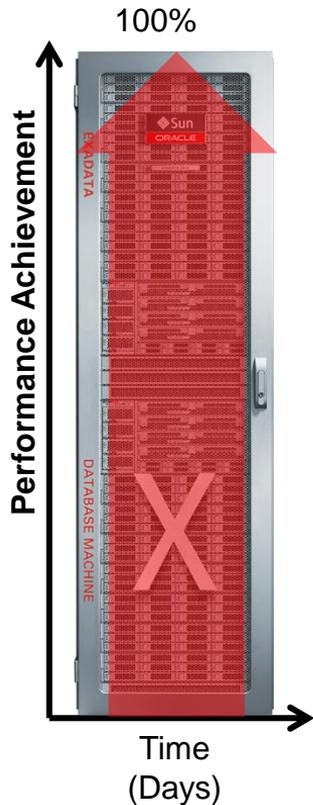
## Scale-out Intelligent Storage Servers

- 14x 2-socket servers → 168 cores in storage
- 168 SAS disk drives → 672 TB HC or 200 TB HP
- 56 Flash PCI cards → 44 TB Flash + compression



ORACLE

# Pre-built and Optimized Out-of-the-Box



# Exadata Engineered System Transformation

Less Risk, Better Results



- Hundreds of engineer years spent optimizing and hardening the system end-to-end
  - Frees I/T talent to focus on business needs
- Standard platform improves support experience
- **Runs all existing Oracle Database workloads**
- Building block of the Oracle Cloud

# Exadata Hardware

# Exadata X4 Compared with X3

Much More Performance and Capacity – Same Price



## X4 Storage

**2X Larger Physical Flash Memory**

**44 TB of Flash Memory**

**Up to 4X Larger Logical Flash Memory**

**88 TB using Flash Cache Compression**

**77% More Flash IOs/sec on X4-2**

**2.66M Reads, 1.96M Writes from SQL**

**33% Larger High Capacity Disks**

**672 TB using 4TB Disks**

**2X Larger High Performance Disks**

**200 TB using 1.2 TB Disks**

Per DB  
Machine  
Full  
Rack

## X4-2 Compute

**50% More Database Cores**

**192 Cores using 12-Core Xeon® CPUs**

**2X Larger DB Server Local Storage**

**2.4 TB per server using 600GB Disks**

**2X Faster InfiniBand**

**InfiniBand PCI-3 Card. All Ports Active**

ORACLE

# X4-2 Database Server

New 12-core “IvyBridge” CPUs, Faster InfiniBand Card, Larger Disks

Processors	<b>2 Twelve-Core Intel® Xeon® E5-2697 v2 Processors (2.7GHz)</b>
Memory	256 GB (16 x 16GB) – Expandable to 512GB (16 X 32GB) via memory kits
Local Disks	4 x <b>600GB</b> 10K RPM SAS Disks (Hot-Swappable)
Disk Controller	Disk Controller HBA with 512MB Cache – <b>Battery Online Replaceable</b>
Network	2 x InfiniBand 4X QDR (40Gb/s) Ports ( <b>PCIe 3.0</b> ) – <b>Both Ports Active</b> 4 x 1GbE/10GbE Base-T Ethernet Ports 2 x 10GbE Ethernet SFP+ Ports (1 Dual-port 10GbE PCIe 2.0 network card based on the Intel 82599 10GbE Controller technology)
Remote Management	1 Ethernet port (ILOM)
Power Supplies	Redundant Hot-Swappable power supplies and fans

# X4-2 Storage Server

6-core IvyBridge CPUs, Larger disks, Larger Flash Cards, Flash Compression

Processors	<b>2 Six-Core Intel® Xeon® E5-2630 v2 Processors (2.6 GHz) - Faster clock</b>
Memory	<b>96 GB (4 x 8GB + 4 x 16GB) - More memory needed to manage larger flash</b>
Disks	12 x <b>1.2 TB 10K RPM</b> High Performance SAS (hot-swap) – <b>2.5” disk size</b> OR 12 x <b>4 TB</b> 7.2K RPM High Capacity SAS (hot-swap) – 3.5” disk size
Flash	4 x <b>800 GB</b> Sun Flash Accelerator <b>F80</b> PCIe Cards – <b>Hardware Compression</b>
Disk Controller	Disk Controller HBA with 512MB Cache - <b>Battery Online Replaceable</b>
Network	2 InfiniBand 4X QDR (40Gb/s) Ports ( <b>PCIe 3.0</b> ) – <b>Both Ports Active</b> Embedded Gigabit Ethernet Ports for management connectivity
Remote Management	1 Ethernet port (ILOM)
Power Supplies	Redundant Hot-Swappable power supplies and fans

# X3-8 Database Server

Processors	8 x Ten-Core Intel® Xeon® E7-8870 Processors (2.40 GHz)
Memory	2 TB (128 x16 GB)
Local Disks	8 x 300GB 10K RPM SAS Disks (Hot-Swappable)
Disk Controller	Disk Controller HBA with 512MB Battery Backed Cache
Network	8 x InfiniBand 4X QDR (40Gb/s) Ports (4 Dual-port PCE 2.0 Express Modules) Two Network Express Modules (NEM), providing a total of <ul style="list-style-type: none"><li>• 8 x 1GbE Ethernet Ports</li><li>• 8 x 10 GbE Ethernet SFP+ Ports (via 4 Fabric Express Modules (FEM) based Intel 82599 10GbE Controller technology)</li></ul>
Remote Management	1 Ethernet port (ILOM)
Power Supplies	Redundant Hot-Swappable power supplies and fans

# Exadata Database Machine X4-2 Full Rack

## Pre-Configured for Extreme Performance

- 8 Xeon-based Dual-processor Database Servers
  - 192 cores (24 per server)
  - 2 TB memory expandable to 4 TB (256 GB per server expandable to 512GB)
  - 10 Gig E-connectivity to Data Center
    - 40 x 10Gb E-ports (5 per server)
- 44.8 TB High Speed Flash
- 14 Exadata Storage Servers X4-2
  - All with High Performance 12 x 1.2 TB SAS disksOR
  - All with High Capacity 12 x 4 TB SAS disks
- 2 Sun Datacenter InfiniBand Switch 36
  - 36-port Managed QDR (40Gb/s) switch
- 1 “Admin” Ethernet switch
- Redundant Power Distributions Units (PDUs)

**Available in Smaller Configurations**



ORACLE

# Exadata Database Machine X3-8 Full Rack

## Larger Database Servers with X4-2 Storage Servers

- 2 Xeon-based Eight-processor Database Servers
  - 160 CPU cores (80 per server)
  - 4 TB memory (2TB per server)
  - 10 GigE connectivity to Data Center
    - 16 x 10Gb E-ports (8 per server)
- 44.8 TB High Speed Flash
- 14 Exadata Storage Servers X4-2
  - All with High Performance 12 x 1.2 TB SAS disks
- OR
  - All with High Capacity 12 x 4 TB SAS disks
- 2 Sun Datacenter InfiniBand Switch 36
  - 36-port Managed QDR (40Gb/s) switch
- 1 “Admin” Ethernet switch
- Redundant Power Distributions Units (PDUs)



**Add more racks for additional scalability**

ORACLE

# Exadata Storage Expansion X4-2 Full Rack

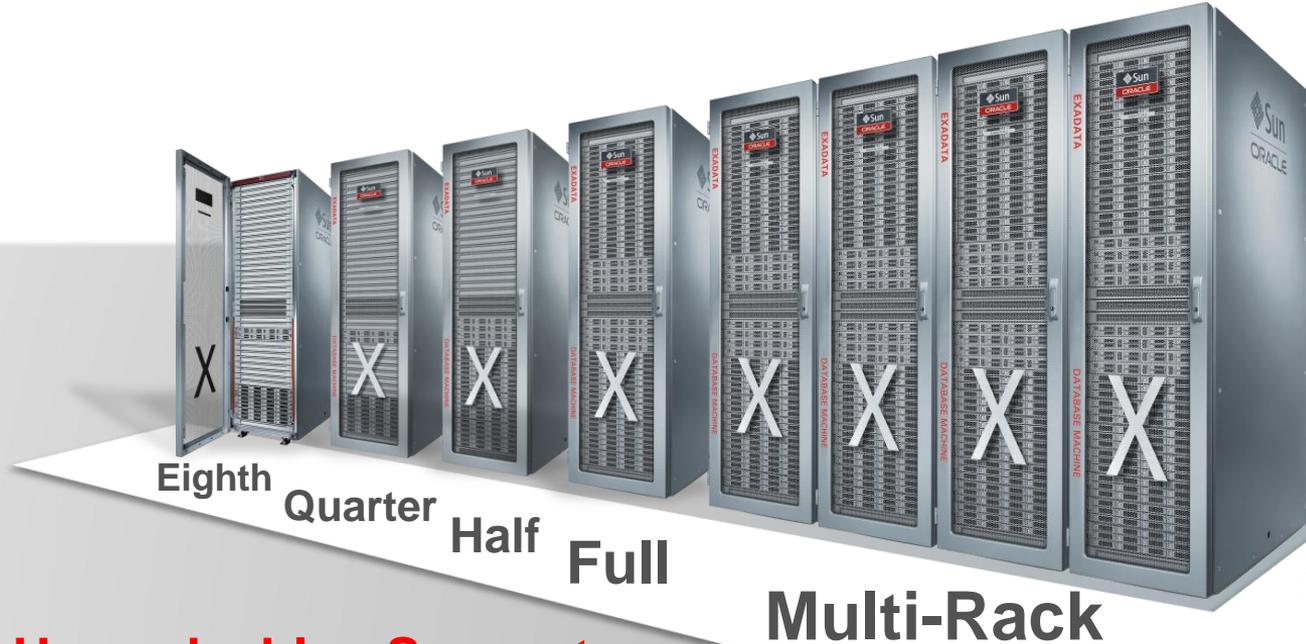
For additional storage – backups, historical data, unstructured data

- 18 Exadata Storage Servers X4-2
  - All with High Performance 12 x 1.2 TB SAS disks
- OR
- All with High Capacity 12 x 4 TB SAS disks
- 57.6 TB High Speed Flash
- 2 Sun Datacenter InfiniBand Switch 36
  - 36-port Managed QDR (40Gb/s) switch
- 1 “Admin” Ethernet switch
- Redundant Power Distributions Units (PDUs)



Add storage capacity to Database Machine online

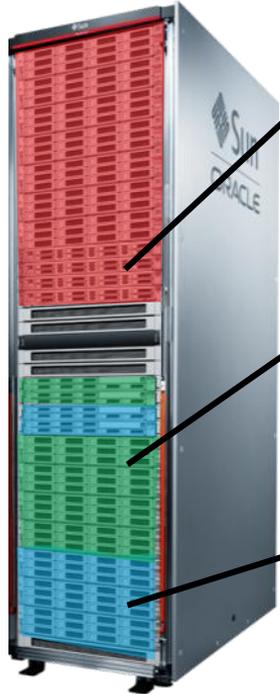
# Scalable from Eighth-Rack to Multi-Rack



**Field Upgradeable - Supports  
Multiple Generations of Hardware**

# Seamless Upgrades and Expansions

## Upgrade Example



**X4-2**

Half to Full Upgrade  
in 2014

**X3-2**

Qtr to Half Upgrade  
in 2013

**X2-2**

Initial Quarter Rack  
deployed in 2011

- **A single Database Machine can have servers from different generations**
- **Databases and Clusters can span across multiple hardware generations**
- **New software runs on older hardware**

# Operating System Updates in 11.2.3.3.0



- Oracle Linux distribution updated to 5.9
  - Including all recent security updates
- Linux kernel updated to UEK2 kernel (2.6.39-400)
  - Same kernel on database servers (2-socket and 8-socket) and the storage servers



- Oracle Solaris updated to S11 Update 1 SRU 9

# X4-2 Database Machine IO Performance from SQL

		X4-2 Full Rack	X4-2 Half Rack	X4-2 Quarter	X4-2 Eighth
<b>Flash Cache SQL Bandwidth<sup>1,3</sup></b>	High Cap Disk	100 GB/s	50 GB/s	21.5 GB/s	10.7 GB/s
	High Perf Disk	100 GB/s	50 GB/s	21.5 GB/s	10.7 GB/s
<b>Flash SQL IOPS<sup>2,3</sup></b>	8K Reads	2,660,000	1,330,000	570,000	285,000
	8K Writes	1,960,000	980,000	420,000	210,000
<b>Disk SQL Bandwidth<sup>1,3</sup></b>	High Cap Disk	20 GB/s	10 GB/s	4.5 G/s	2.25 GB/s
	High Perf Disk	24 GB/s	12 GB/s	5.2 GB/s	2.6 GB/s
<b>Disk SQL IOPS</b>	High Cap Disk	32,000	16,000	7,000	3,500
	High Perf Disk	50,000	25,000	10,800	5,400
<b>Data Load Rate<sup>4</sup></b>		20 TB/hr	10 TB/hr	5 TB/hr	2.5 TB/hr

1 - Bandwidth is peak physical scan bandwidth achieved running SQL, assuming no compression. Effective data bandwidth will be much higher when compression is factored in.

2 - IOPS – Based on read IO requests of size 8K running SQL, typically with sub-millisecond latencies. Note that the IO size greatly effects flash IOPS. Others quote IOPS based on smaller IOs that are not relevant for databases and measure IOs using low level tools instead of SQL.

3- Actual Performance varies by application.

4 –Load rates are typically limited by database server CPU, not IO. Rates vary based on load method, indexes, data types, compression, and partitioning

# Note on Compression and Measurements

- Note that all capacity and performance numbers are shown without including the benefits of compression
  - Compression rates vary greatly by application and compression type
  - Therefore it is inaccurate to quote a single number for all users
  - Capacity and performance will greatly improve over the quoted numbers when compression benefits are factored in
- When comparing Exadata capacity and performance to other vendors note:
  - Many vendors (especially flash array vendors) quote capacity and performance assuming some rate of compression
  - Many vendors quote flash IOs per second running low level IO utilities and using very small blocks. These are several times higher than can be achieved using SQL.
  - All Exadata measurements are from SQL workloads using standard 8K block sizes

# X3-8 Database Machine IO Performance from SQL

		X3-8 Full Rack
<b>Flash Cache SQL Bandwidth<sup>1,3</sup></b>	High Cap Disk	100 GB/s
	High Perf Disk	100 GB/s
<b>Flash SQL IOPS<sup>2,3</sup></b>	8K Reads	1,500,000
	8K Writes	1,000,000
<b>Disk SQL Bandwidth<sup>1,3</sup></b>	High Cap Disk	20 GB/s
	High Perf Disk	24 GB/s
<b>Disk SQL IOPS</b>	High Cap Disk	32,000
	High Perf Disk	50,000
<b>Data Load Rate<sup>4</sup></b>		16TB/hr

1 - Bandwidth is peak physical scan bandwidth achieved running SQL, assuming no compression. Effective data bandwidth will be much higher when compression is factored in.

2 - IOPS – Based on read IO requests of size 8K running SQL, typically with sub-millisecond latencies. Note that the IO size greatly effects flash IOPS. Others quote IOPS based on smaller IOs that are not relevant for databases and measure IOs using low level tools instead of SQL. Exadata Flash read IOPS are so high they are typically limited by database server CPU, not IO.

3- Actual Performance varies by application.

4 – Load rates are typically limited by database server CPU, not IO. Rates vary based on load method, indexes, data types, compression, and partitioning

# X4-2 Storage Expansion IO Performance from SQL

		X4-2 Full Rack	X4-2 Half Rack	X4-2 Quarter	X4-2 Single Cell
<b>Flash Cache SQL Bandwidth<sup>1,3</sup></b>	High Cap Disk	130 GB/s	65 GB/s	29 GB/s	7.25 GB/s
	High Perf Disk	130 GB/s	65 GB/s	29 GB/s	7.25 GB/s
<b>Flash SQL IOPS<sup>2,3</sup></b>	8K Reads	3,420,000	1,710,000	760,000	190,000
	8K Writes	2,520,000	1,260,000	560,000	140,000
<b>Disk SQL Bandwidth<sup>1,3</sup></b>	High Cap Disk	26 GB/s	13 GB/s	6 GB/s	1.5 GB/s
	High Perf Disk	30 GB/s	15 GB/s	7 GB/s	1.75 GB/s
<b>Disk SQL IOPS</b>	High Cap Disk	42,000	21,000	9,500	2,400
	High Perf Disk	64,000	32,000	14,400	3,600

1 - Bandwidth is peak physical scan bandwidth achieved running SQL, assuming no compression. Effective data bandwidth will be much higher when compression is factored in.

2 - IOPS – Based on read IO requests of size 8K running SQL, typically with sub-millisecond latencies. Note that the IO size greatly effects flash IOPS. Others quote IOPS based on smaller IOs that are not relevant for databases and measure IOs using low level tools instead of SQL. Exadata Flash read IOPS are so high they are typically limited by database server CPU, not IO. This is especially true for expansion racks.

3- Actual Performance varies by application.

# Database Machine Capacity (Uncompressed)

		X3-8 or X4-2 Full	X4-2 Half	X4-2 Quarter	X4-2 Eighth
<b>Raw Flash Capacity<sup>1,4</sup></b>		44.8 TB	22.4 TB	9.6 TB	4.8 TB
<b>Effective Flash Capacity Due to Caching<sup>4</sup></b>		440 TB	220TB	96TB	48TB
<b>Raw Disk Capacity<sup>1</sup></b>	High Cap Disk	672 TB	336 TB	144 TB	72 TB
	High Perf Disk	200 TB	100 TB	43.2 TB	21.6 TB
<b>Usable Mirrored Capacity<sup>2,3</sup></b>	High Cap Disk	300 TB	150 TB	63 TB	30 TB
	High Perf Disk	90 TB	45 TB	19 TB	9 TB
<b>Usable Triple Mirrored Capacity<sup>2,3</sup></b>	High Cap Disk	200 TB	100 TB	43 TB	21.5 TB
	High Perf Disk	60 TB	30 TB	13 TB	6.3 TB

1 - Raw Disk Capacity defined using standard disk drive terminology of 1 TB = 1000 \* 1000 \* 1000 \* 1000 bytes.

2 - Capacity calculated using normal space terminology of 1 TB = 1024 \* 1024 \* 1024 \* 1024 bytes.

3 - Actual space available for a database after mirroring (ASM normal or high redundancy). For the ASM normal redundancy case, assume one disk (Quarter and Half) or two disks (Full Rack) of free space to automatically remirror after disk failures.

4 – Effective Flash Capacity is larger than the physical flash capacity and takes into account the high flash hit ratios due to Exadata's intelligent flash caching algorithms, and the size of the underlying disk storage. It is the size of data files that often can be stored in Exadata and be accessed at the speed of flash memory.

# Effective Flash Capacity for Typical Database Workloads

- Exadata flash cache delivers flash performance for data sets that are much bigger than the size of flash – often 10x bigger
  - 440 TB of effective flash on full rack
- Exadata Flash Cache automatically moves frequently accessed data into flash
  - With a very fine granularity and in real-time
  - Flash cache does not need to hold mirror copies
- Typical flash cache hit rates in 95% - 99% range when caching 10x as much disk space
- HCC and OLTP compression further increase effective flash capacity when measured in terms of user data in flash

# Storage Expansion Capacity (Uncompressed)

		X4-2 Full	X4-2 Half	X4-2 Quarter	X4-2 Single Cell
<b>Raw Flash Capacity<sup>1</sup></b>		57.6 TB	28.8 TB	12.8 TB	3.2 TB
<b>Effective Flash Capacity Due to Caching<sup>4</sup></b>		570 TB	280 TB	128 TB	32 TB
<b>Raw Disk Capacity<sup>1</sup></b>	High Cap Disk	864 TB	432 TB	192 TB	48 TB
	High Perf Disk	258 TB	129 TB	57 TB	14.4 TB
<b>Usable Mirrored Capacity<sup>2,3</sup></b>	High Cap Disk	387 TB	194 TB	85 TB	20 TB
	High Perf Disk	116 TB	58 TB	25 TB	6 TB
<b>Usable Triple Mirrored Capacity<sup>2,3</sup></b>	High Cap Disk	260 TB	130TB	58 TB	14.5 TB
	High Perf Disk	78 TB	39 TB	17 TB	4 .25 TB

1 - Raw Disk Capacity defined using standard disk drive terminology of 1 TB = 1000 \* 1000 \* 1000 \* 1000 bytes.

2 - Capacity calculated using normal space terminology of 1 TB = 1024 \* 1024 \* 1024 \* 1024 bytes.

3 - Actual space available for a database after mirroring (ASM normal or high redundancy). For the ASM normal redundancy case, assume one disk (Quarter and Half) or two disks (Full Rack) of free space to automatically remirror after disk failures.

4 – Effective Flash Capacity is larger than the physical flash capacity and takes into account the high flash hit ratios due to Exadata’s intelligent flash caching algorithms, and the size of the underlying disk storage. It is the size of data files that often can be stored in Exadata and be accessed at the speed of flash memory.

# Exadata Database Machine Hardware Summary

		X3-8 Full	X4-2 Full	X4-2 Half	X4-2 Quarter	X4-2 Eighth
<b>Database Servers</b>		2	8	4	2	2
<b>Database Grid Cores</b>		160	192	96	48	24
<b>Database Grid Memory (GB)</b>		4096	2048 (max 4096)	1024 (max 2048)	512 (max 1024)	512 (max 1024)
<b>InfiniBand switches</b>		2	2	2	2	2
<b>Ethernet switch</b>		1	1	1	1	1
<b>Exadata Storage Servers</b>		14	14	7	3	3
<b>Storage Grid CPU Cores</b>		168	168	84	36	18
<b>Raw Flash Capacity</b>		44.8 TB	44.8 TB	22.4 TB	9.6 TB	4.8 TB
<b>Raw Storage Capacity</b>	<b>High Perf</b>	200 TB	200 TB	100 TB	43.2 TB	21.6 TB
	<b>High Cap</b>	672 TB	672 TB	336 TB	144 TB	72 TB
<b>Usable mirrored capacity</b>	<b>High Perf</b>	90 TB	90 TB	45 TB	19 TB	9 TB
	<b>High Cap</b>	300 TB	300 TB	150 TB	63 TB	30 TB
<b>Usable Triple mirrored capacity</b>	<b>High Perf</b>	60 TB	60 TB	30 TB	13 TB	6.3 TB
	<b>High Cap</b>	200 TB	200 TB	100 TB	43 TB	21.5 TB

# Exadata Storage Expansion Rack Summary

		X4-2 Full	X4-2 Half	X4-2 Quarter	Single Cell
<b>InfiniBand switches</b>		3	3	2	-
<b>Ethernet switch</b>		1	1	1	-
<b>Exadata Storage Servers</b>		18	9	4	1
<b>Storage Grid CPU Cores</b>		216	108	48	12
<b>Raw Flash Capacity</b>		57.6 TB	28.8 TB	12.8 TB	3.2 TB
<b>Raw Storage Capacity</b>	<b>High Perf</b>	258 TB	129 TB	57 TB	14.4 TB
	<b>High Cap</b>	864 TB	432 TB	192 TB	48 TB
<b>Usable mirrored capacity</b>	<b>High Perf</b>	116 TB	58 TB	25 TB	6 TB
	<b>High Cap</b>	387 TB	194 TB	85 TB	20 TB
<b>Usable Triple mirrored capacity</b>	<b>High Perf</b>	78 TB	39 TB	17 TB	4.25 TB
	<b>High Cap</b>	260 TB	130TB	58 TB	14.5 TB

# Hardware Generational Advances

V1



2008

V2



2009

X2



2010

X3

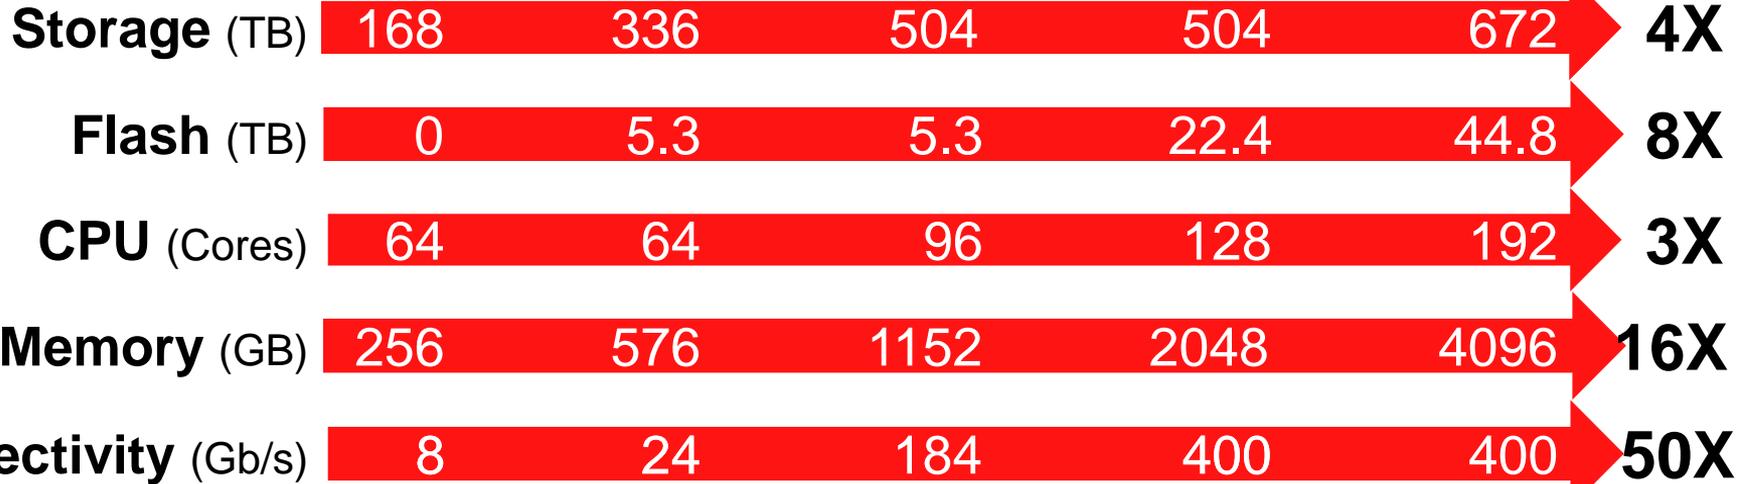


2012

X4



2013



ORACLE

# Exadata Innovations

# Unique Software Optimizes Database Processing

- Query **offload** in storage

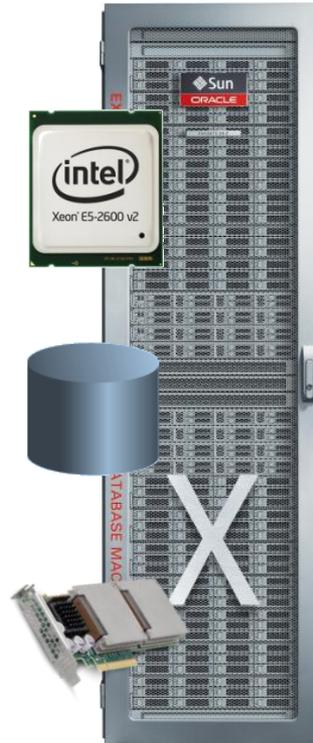
- Data intensive query operations offloaded to storage CPUs
- 100 GB/sec SQL data throughput
- Storage Index data skipping

- Database storage **compression**

- Hybrid Columnar for 10x DB size reduction and faster analytics

- Database optimized **PCI Flash**

- Smart caching of database data
- 2.66 Million Database IOs/sec
- Smart Flash log speeds transactions



- Database optimized **QoS**

- End-to-end prioritization from application to DB and storage

- Database optimized **availability**

- Fastest recovery of failed database, server, storage or switch
- Fastest backup. Incremental offload
- Exachk top-to-bottom validation of hardware, software, settings

- Database optimized **messaging**

- SQL optimized InfiniBand protocol for high throughput low latency SQL

# Innovation Continues: Recent Enhancements

- Query **offload** in storage

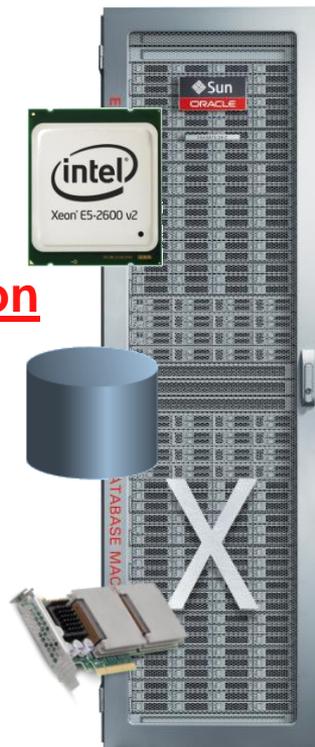
- Offload searches on **LOBs (12c)**
- Offload **joins** for non-parallel queries (11.2.0.4)

- Database optimized **compression**

- Hybrid Columnar enhanced for **OLTP**
- and for **Spatial and Text** data (12c)

- Database optimized **PCI Flash**

- Ultra high speed flash compression (**X3 & X4**) at multi-million IOs/sec
- Automatic caching for **table scans**
- Faster **file initialization**



- Database optimized **QoS**

- Prioritization of CPU and IO by **multitenant pluggable database (12c)**

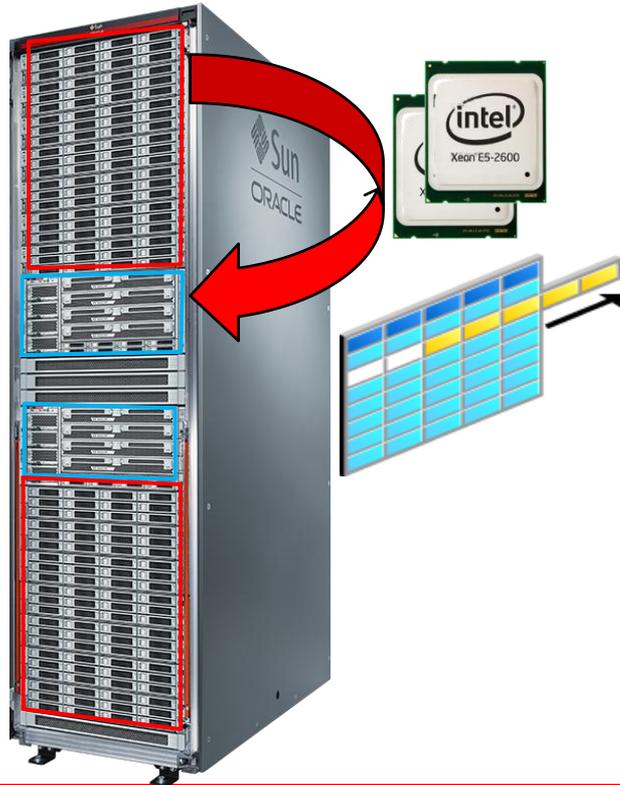
- Database optimized **availability**

- Prioritize **recovery of critical DB files (11.2.0.4)**

- Database optimized **messaging**

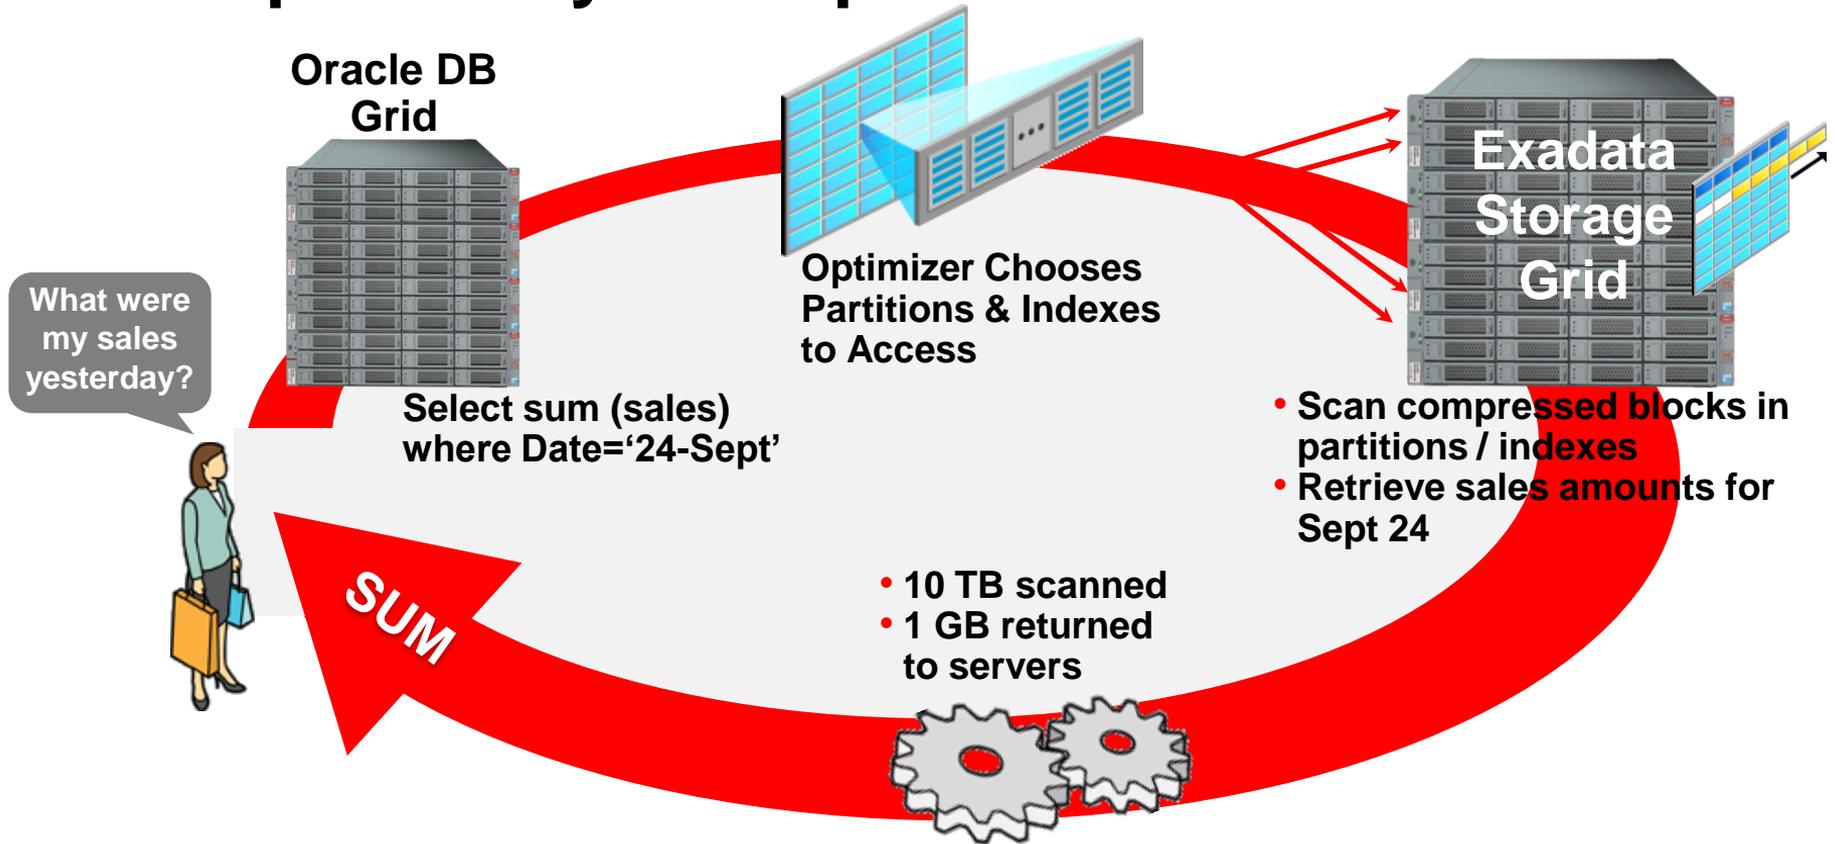
- **End-to-End prioritization of critical database messages (11.2.0.4), including log writes and RAC**

# Exadata Intelligent Storage Grid

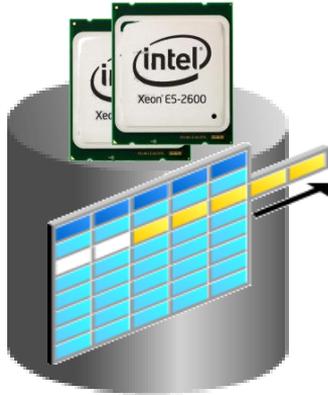


- Data Intensive processing runs in Exadata Storage Grid
  - Filter rows and columns as data streams from disks (168 Intel Cores)
- Example: How much product X sold last quarter
  - Exadata Storage Reads 10TB from disk
  - Exadata Storage Filters rows by Product & Date
  - Sends 100GB of matching data to DB Servers
- Scale-out storage parallelizes execution and removes bottlenecks

# Simple Query Example



# Exadata Intelligent Storage

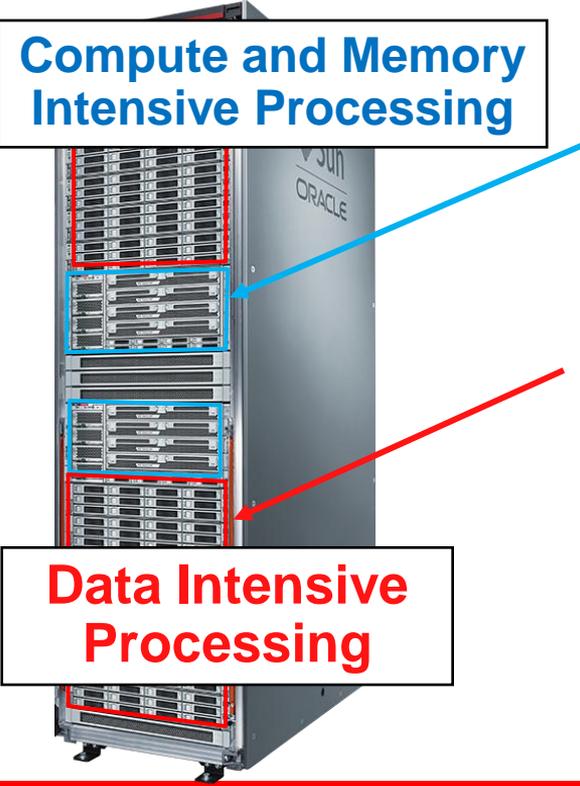


## Exadata Intelligent Storage Grid

- Exadata storage servers also run more complex operations in storage
  - **Join filtering**
  - **Incremental backup filtering**
  - **I/O prioritization**
  - **Storage Indexing**
  - **Database level security**
  - **Offloaded scans on encrypted data**
  - **Data Mining Model Scoring**

- 10x reduction in data sent to DB servers is common

# Exadata is Smart Storage



Compute and Memory  
Intensive Processing

## ▪ Database Servers

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

## ▪ Exadata Storage Servers

- Storage Server is smart storage, not a DB node
- Search tables and indexes filtering out data that is not relevant to a query
- Cells serve data to multiple databases **enabling OLTP and consolidation**
- Simplicity, and robustness of storage appliance

Data Intensive  
Processing

# Exadata Storage Index

## Transparent I/O Elimination with No Overhead

Table				Index
A	B	C	D	
	1			Min B = 1 Max B = 5
	3			
	5			
	5			Min B = 3 Max B = 8
	8			
	3			

- 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 automatic and transparent

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

# Storage Index with Partitions Example

Orders Table			
Order#	Order_Date Partitioning Column	Ship_Date	Item
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

# Storage Index with Joins Example

```
Select count(*) from fact, dim
where fact.m=dim.m and dim.name='Camry'
```

## Dimension

Name	M
Accord	1
Camry	3
Civic	5
Prius	8

Bloom filter constructed  
with min/max for M

Perform IO and  
apply bloom filter

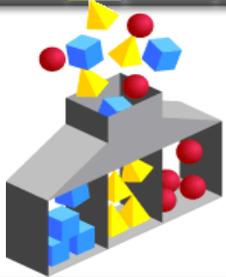
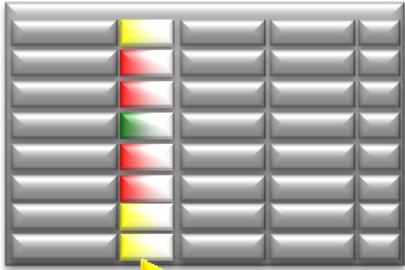
Skip IO  
Due to Storage Index

## Fact

A	M	C	D
	1		
	3		
	5		
	5		
	5		
	5		

# Exadata Hybrid Columnar Compression

Highest Capacity, Lowest Cost



**Faster and Simpler**  
Backup, DR, Caching,  
Reorg, Clone

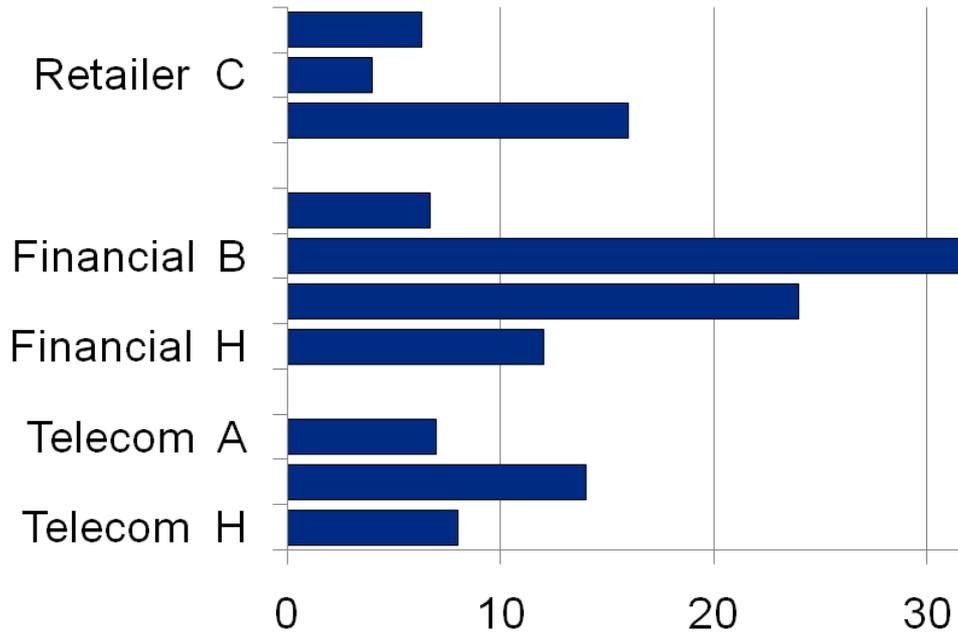
- Data is organized and compressed by column
  - Dramatically better compression
- Speed Optimized **Query Mode** for Data Warehousing
  - 10X compression typical
  - Runs faster because of Exadata offload!
- Space Optimized **Archival Mode** for infrequently accessed data
  - 15X to 50X compression typical



Benefits Multiply

# Compression Ratio of Real-World Data

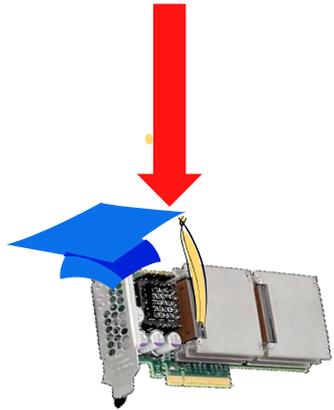
## Query Compression Ratio (Avg = 13x)



- Compression Ratio varies by customer and table
- Trials were run on largest table at 10 ultra large companies
  - Average revenue > \$60 BB
- Average Query Compression ratio was 13x
  - On top of Oracle's already highly efficient format

# Exadata Smart Flash Cache

I/Os



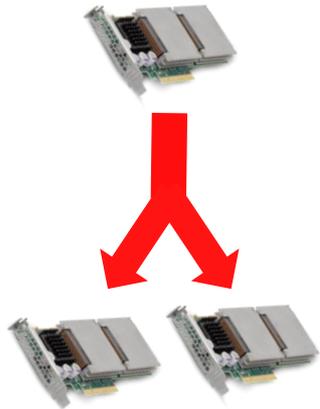
**2.66 Million 8K Read  
1.96 Million 8K Write  
IOPS from SQL**

- Caches Read and Write I/Os in PCI flash
- Transparently accelerates read and write intensive workloads
  - Up to 2.66 million 8K read IOPS from SQL
  - Up to 1.96 million 8K write IOPS from SQL
- Persistent write cache speeds database recovery
- Exadata Flash Cache is much more effective than flash tiering architectures used by others
  - Caches current hot data, not yesterday's
  - Caches data in granules 8x to 16x smaller than tiering
    - Greatly improves the effectiveness of flash

# Exadata Flash Cache Compression

New

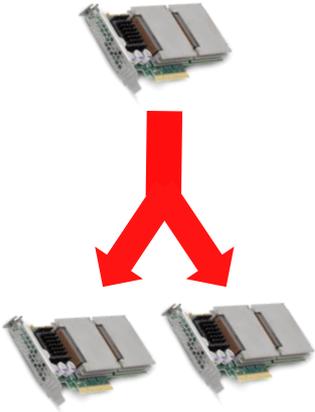
- Exadata uses compression to expand smart flash cache
  - Data automatically compressed as it is written to flash cache
  - Automatically decompressed when it is read out of flash cache
  - Up to 2X more data fits in smart flash cache, so flash hit rates will improve and performance will improve for large data sets
- Flash cache compress/decompress implemented in hardware
  - Performance is same as uncompressed – millions of I/Os per second
    - ZERO performance overhead
  - Supported on X3 or X4 storage servers (requires F40 or F80 cards)
- Note – Flash cache compression does not change DB format, so does not improve backups, buffer cache, network, scan rates, etc.



# Exadata Flash Cache Compression

New

- As always, compression benefits vary based on data



Data Type	Typical Compression
Uncompressed Tables	1.3X to 4X
OLTP Compressed Tables	1.2X to 2X
Indexes	1.3X to 4X
Oracle E-biz uncompressed DB	3x to 5x
HCC Compressed Tables or Compressed LOBs	Minimal

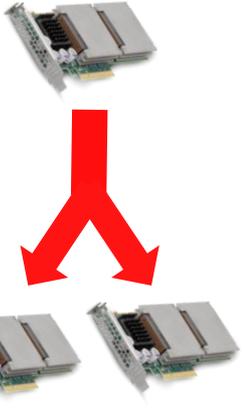
Many OLTP Databases will see 2x Flash Increase

- X4 with flash cache compression stores up to 80TB of data in flash
  - Up to 4X more than X3 (depending on compressibility of data)

ORACLE

# Flash Cache Compression Commands

- Trivial to implement, no management
- Enable using simple cell command:
  - On X4 machines: `alter cell flashCompression=TRUE`
  - On X3 machines also run:
    - `alter cell FlashCacheCompX3Support= TRUE`
- Amount of data cached in Exadata Smart Flash Cache grows and shrinks dynamically and automatically based on data compressibility
- Monitor Flash Cache Compression using cell metric `FC_BY_USED`
  - Reported flash cache size will increase to up to double physical flash size



# Exadata Smart Flash Table Caching

New

- Smarter flash caching for large table scans

- Exadata software understands database table and partition scans and automatically caches them when it makes sense
- Avoids thrashing flash cache when tables are too big or scanned infrequently or scanned by maintenance jobs
- If scanned table is larger than flash, then subset of table is cached
- No need to manually “KEEP” tables that are only scanned



Scan



# Exadata Smart Flash Log

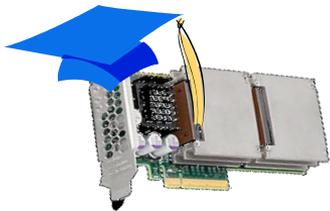
Accelerate Transaction Response Times Using Flash

## Default (on left)

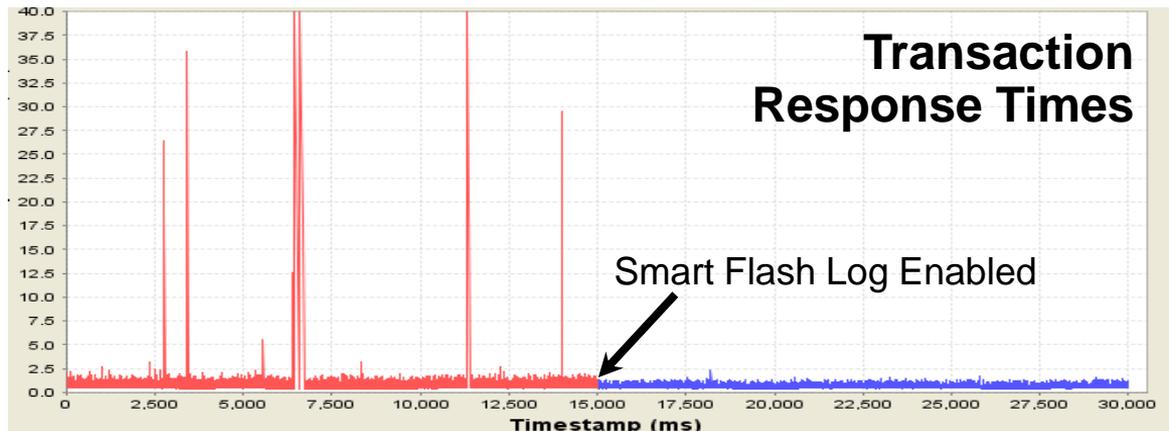
- Choppy Response
- High Outliers

## Smart Flash Log

- 3x faster response
- Much lower outliers



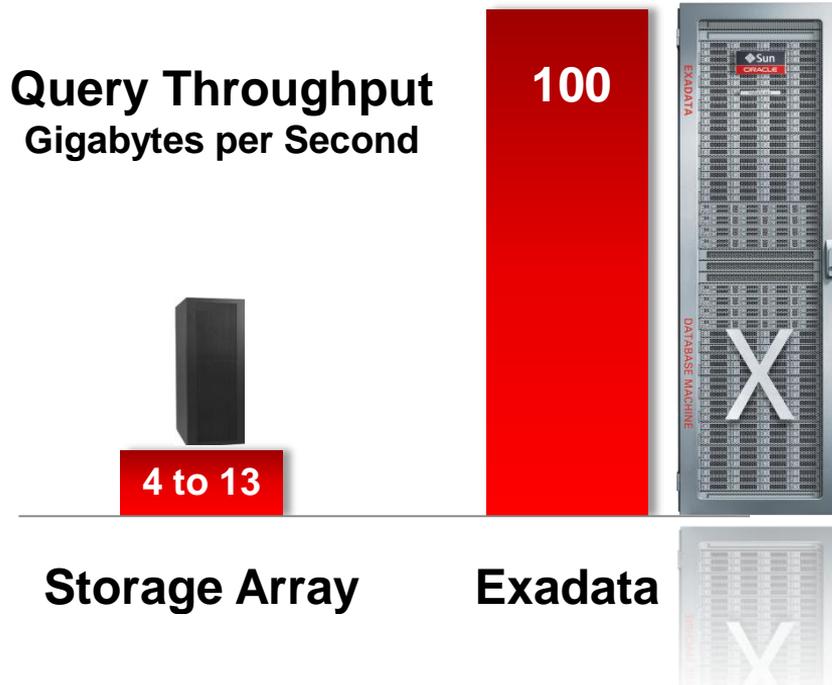
Automatic and  
Transparent



- Uses Flash for Database Logs in a clever way
  - Flash is fast but has slow outliers (Erase cycles, wear leveling, etc)
- Smart Flash Log feature transparently uses Flash as a parallel write cache to disk controller cache
  - Whichever write completes first wins (disk or flash)
- Better response time and more throughput
- Uses almost no flash capacity (0.1% of capacity)

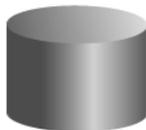
# Exadata Flash Performance Scales Linearly

## Storage Array Architecture Limits Flash Performance



- Exadata scales using
  - True Scale-Out
  - InfiniBand
  - Smart Storage

# Extreme Performance and Low Cost



Disks for Capacity



Flash for I/Os

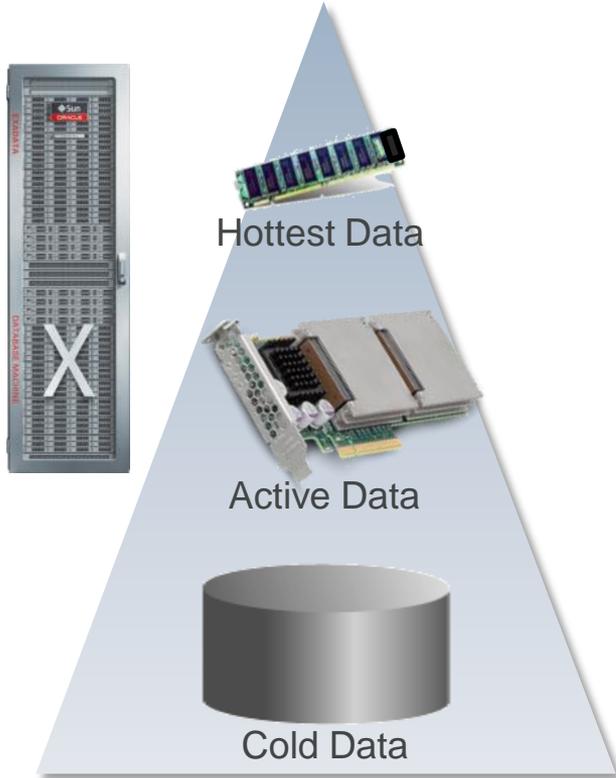


DRAM for Throughput

**Cost of Disk  
I/Os of Flash  
Speed of DRAM**

- **Extreme Performance dramatically lowers cost**
  - Workloads that require huge traditional systems run on small Exadata
  - Hardware needed for an application is often reduced 10x
- **Exadata also delivers Extreme Capacity**
  - Huge disk footprint for large or less active data
    - 672 TB of disk per rack for historical data, low activity data, images, documents, backups, etc.
  - Columnar compression expands disk and memory capacity 10x
  - Blend of tiers gives highest performance at lowest cost
    - Placing all data in DRAM would cost \$20M just for chips

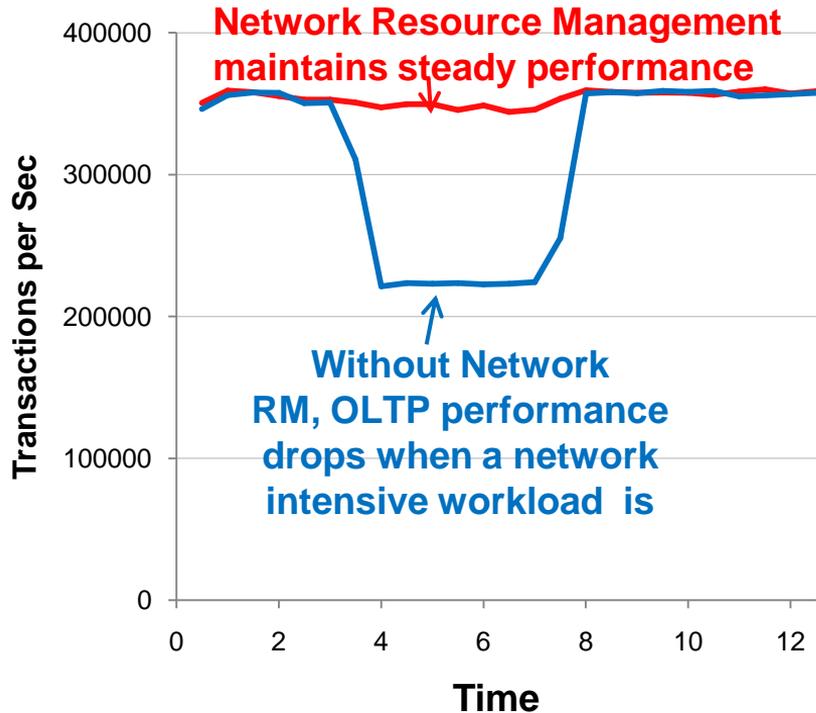
# Business Benefits of X4 Database Machine



- Sub-millisecond latency
  - Interactive OLTP with millions of users
- Sub-second analytics
  - Real-time decision making, instant reports
- 10x faster parallel jobs
  - Quarter close, payroll, supply planning, field inventory, pricing, route planning, sub-ledger accounting

# Exadata Network Resource Management

New



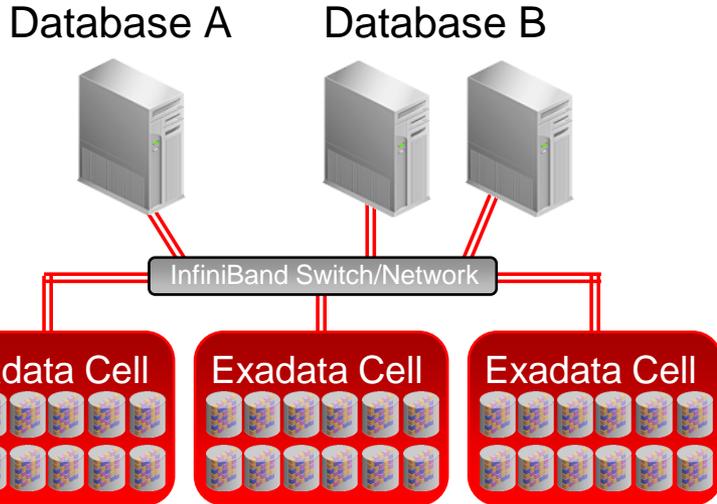
DB Version 11.2.0.4 or 12c, Switch 2.1.3-4

- Exadata Network Resource Management uniquely prioritizes critical database messages through the entire fabric
  - From database to InfiniBand card through InfiniBand switches to storage
  - Latency sensitive messages prioritized over batch, reporting, and backup messages
  - Log file writes have highest priority to ensure low latency transactions
- Combines with Exadata CPU and IO Resource management to ensure safe consolidation of workloads and databases
- Completely automatic & transparent

ORACLE

# Exadata I/O Resource Management

## Mixed Workloads and Multi-Database Environment



- Ensure different databases are allocated the correct relative amount of I/O bandwidth
  - Database A: 33% I/O resources
  - Database B: 67% I/O resources
- Ensure different users and tasks within a database are allocated the correct relative amount of I/O bandwidth
  - Database A:
    - Reporting: 60% of I/O resources
    - ETL: 40% of I/O resources
  - Database B:
    - Interactive: 30% of I/O resources
    - Batch: 70% of I/O resources

# Secure Database Machine

## Complete Security Portfolio



**Near-Zero Overhead for fully encrypted database!**

- Application transparent
- Decryption in hardware, not software (5x faster)

### PREVENTIVE

Encryption

Redaction and Masking

Privileged User Controls

### DETECTIVE

Activity Monitoring

Database Firewall

Auditing and Reporting

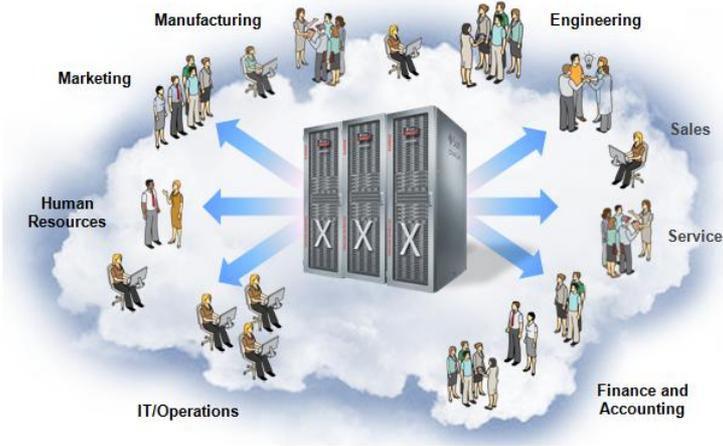
### ADMINISTRATIVE

Privilege Analysis

Sensitive Data Discovery

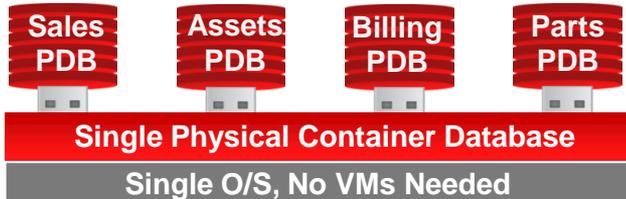
Configuration Management

# Comprehensive Database as a Service Platform



- **Scale-Out Platform optimized for Database**
  - Scale to any size
- **Deploy 100s of databases**
  - Using separate databases or 12c multitenant databases
- **Supports complex & varying mix of workloads**
  - No Performance Bottlenecks
  - Performance Isolation – CPU, I/O, Network

## Multitenant Database



# No Performance Bottlenecks for Consolidation



- Best way to ensure performance is to avoid bottlenecks
- Exadata has unique extreme performance for complex workloads that mix OLTP, DW, batch, reporting
  - Millions of I/Os per second, 100 GB/sec of throughput
  - Sub-millisecond response times
  - Highest bandwidth network
- Unique software optimizations that eliminate bottlenecks
  - e.g. Storage Offload, Smart Flash logging

# Unique Performance Isolation for Consolidation

Application



Database



Network



Storage

**End-to-End  
Prioritization**

- Database Resource Manager provides CPU resource management for normal and pluggable databases
  - Both CPU prioritization and limits on CPU usage
- Exadata uniquely provides I/O resource management by pluggable database, job, user, service, etc.
- Exadata uniquely provides database aware **network resource management**
  - Prioritizes critical DB messages through entire fabric

# Exadata Storage Server Software

12.1.1.1.0

Available on all systems except  
V1



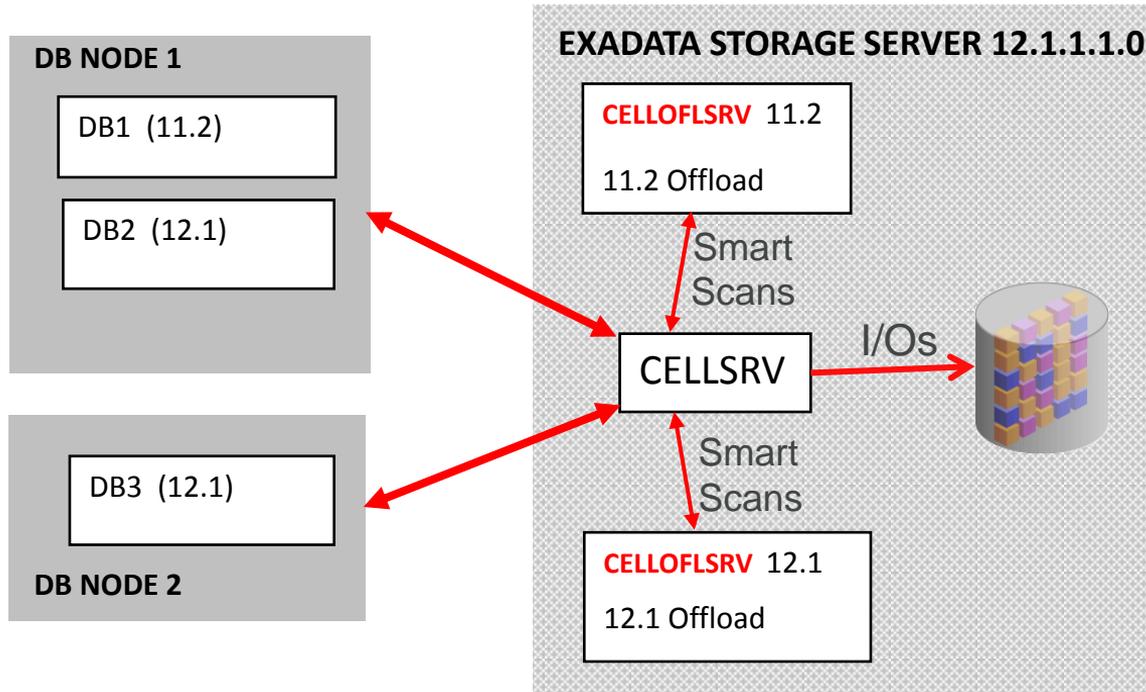
ORACLE

# Exadata Storage Server 12.1.1.1.0



- Supports Database 12c Smart Scan offload
- Supports IORM for Database 12c pluggable database
- Cell-to-cell Data Transfer
  - With Database 12c, offload data movement for ASM operations like resynchronization, resilver, and rebalance
  - Cells transfer data directly to destination cell without going through the database server

# Support for Multiple DB versions on Exadata



- Exadata Systems can run a mix of 11.2 and 12.1 databases
- Different offload server process for each DB version
  - Matches DB libraries
- Offload server startup/shutdown/ restarts automatic
- Offload server communicates only to cellsrv

# Exadata Management

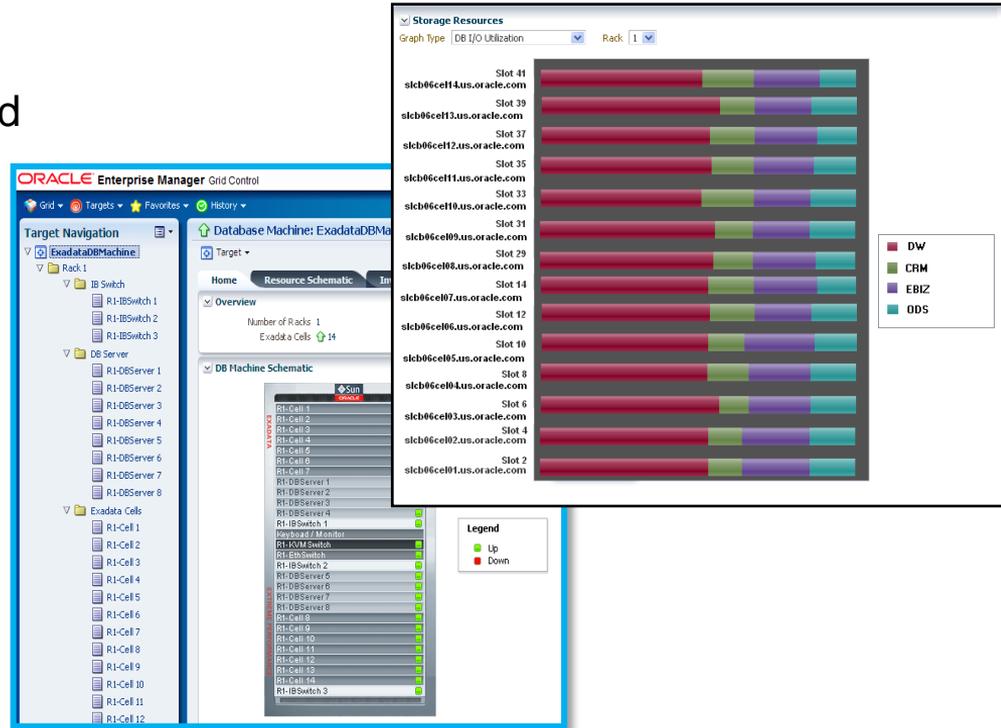
# Exadata Storage Management & Administration

- Enterprise Manager
  - Manage & administer Database and ASM
  - Monitor the Exadata Database Machine Hardware
- Auto Service Request (ASR)
  - File SRs automatically for common hardware faults
- Comprehensive CLI
  - Local Exadata Storage cell management
  - Distributed shell utility to execute CLI across multiple cells
- Embedded Integrated Lights Out Manager (ILOM)
  - Remote management and administration of hardware

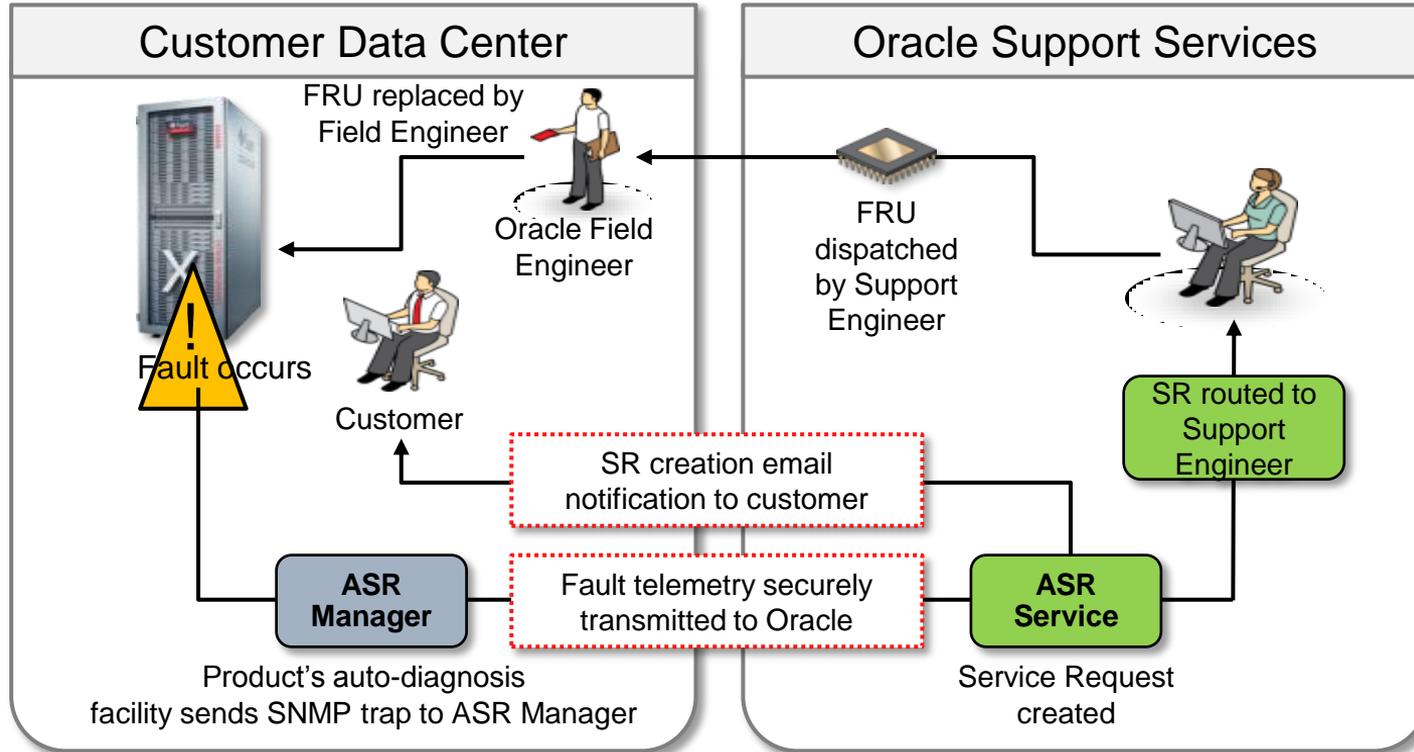
# Enterprise Manager 12c

## Integrated H/W + S/W management for Exadata

- **Hardware view**
  - Schematic of cells, compute nodes and switches
  - Hardware components alerts
- **Software/system view**
  - Performance, availability, usage by databases, services, clusters
  - Software alerts db, cluster, ASM
  - Topology view of DB systems/clusters
- **Configuration view**
  - Version summary of all components along with patch recommendations



# Automated Service Request (ASR)



## Comprehensive Fault Coverage

- CPU
- Disk controllers
- Disks
- Flash Cards
- Flash modules
- InfiniBand
- Cards
- Memory
- System Board
- Power supplies
- Fans

# Platinum Support for Exadata

ORACLE<sup>®</sup>  
EXADATA

ORACLE<sup>®</sup>  
DATABASE

ORACLE<sup>®</sup>  
PREMIER SUPPORT

**Complete. Integrated. Proactive.**

- 24/7 support
- Specialized Engineered Systems Support Team
- 2-hour onsite response to hardware issues<sup>1</sup>
- New Updates and Upgrades for Database, Server, Storage, and OS software
- My Oracle Support proactive support portal
- "Phone home" automated service requests (ASR)



**ORACLE PLATINUM SERVICES**

**High Availability. No Additional Cost.**

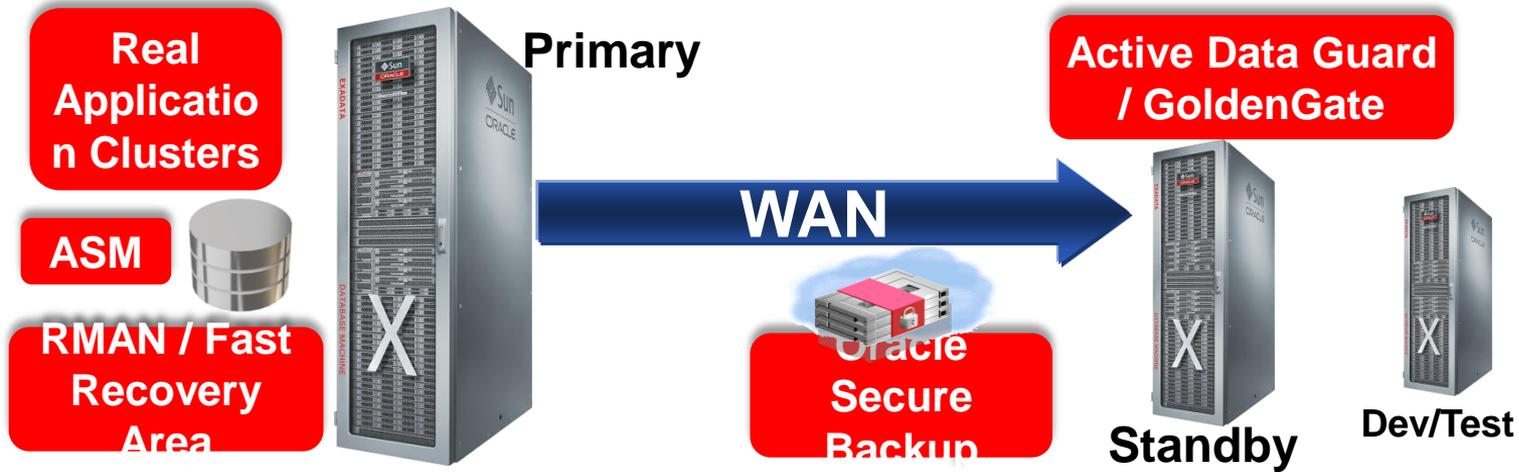
- Better support for the complete Oracle stack
  - Includes higher support levels for Database software
- Proactive remote monitoring for faults
- Industry leading service level response times:
  - 5 Minute Fault Notification
  - 15 Minute Restoration or Escalation to Development
  - 30 Minute Joint Debugging with Development
- Oracle Engineers perform quarterly patching and updates

Available for [certified configurations](#) on Exadata

<sup>1</sup> Covered system must be within an Oracle two-hour service area to receive two-hour response as a standard service.

ORACLE<sup>®</sup>

# Exadata Maximum Availability Architecture



**Comprehensive protection from failures**  
**Server – Storage – Network – Site – Corruptions**

**Active Disaster Recovery: Real-time standby open for query offload**

**Correction from human errors: database, table, row, transaction**

**Online indexing and table redefinition**

**Online patching and upgrades**

# Exadata Comparisons and Customer Use Cases

# PayPal: Real-Time OLTP



## Benefits

*"We chose Oracle Exadata to help keep us nimble by delivering a response time of less than 100ms over petabytes of data."*

- Sehmuz Bayhan, CTO PayPal

Faster Applications

10x Faster Overall



Reduced Risk Exposure

99.99% Analysis



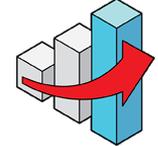
Maximum Availability

99.99%



Increased Payment Volume

200% Increase



## Objectives

- Analyze 100% of payments for risk in < 100 milliseconds
- 99.999% uptime
- 2x increase in transaction volume

## Solution

- 3 Pods, each with two X2-8's plus two X2 Storage Expansion Racks, and Active Data Guard to identical standby configuration

2x Exadata X2-8  
2x Storage Expansion  
Production (Primary Data Center)



2x Exadata X2-8  
2x Storage Expansion  
Standby (DR Data Center)



Active Data Guard  
WAN @ 650+ miles (30 ms)

- 300,000+ executions per second
- 40 ms (avg) response times (99.99%)

24 Exadata Systems in total, across 3 "pods"

120 TB database  
X 3 = 360 TB total

ORACLE

# VocaLink: OLTP Performance

## Benefits

*“Our critical electronic payments service has been live on Exadata since early 2011 with 100% uptime. The service reliably processes the transfer of billions of Euros per week and achieves subsecond response times for online enquires.”* - Martin McGeough, Database Technical Architect

Fast and Scalable  
Performance



Zero Unplanned  
Downtime



More  
Capacity



Zero Data Loss  
DR Protection



## Objectives

- Enable real-time enquiries against transaction data
- Process 20 million financial transactions/day
- Sub-second response time
- Zero data loss DR protection

## Solution

- Two Quarter-rack Exadata V2 systems - Oracle Maximum Availability Architecture

Pre-Exadata  
4 ProdServers, 2 SANs



Exadata V2  
Quarter Rack



- Electronic Payment System
- Mixed workload, high volume OLTP and batch processing

Active Data Guard  
Synchronous Zero Data Loss

Exadata V2  
Quarter Rack



- Performance boost with no application change using Exadata Smart Flash Cache and Smart Scan



Exadata  
test system

- Zero RPO, 15 minute RTO
- Auto block repair
- Standby-first maintenance

# U.S. Customs and Border Protection



## Benefits

*“Oracle Exadata has changed our strategic IT focus from building systems to developing and supporting services.”*

- Ken Ritchhart, Deputy Assistant Commissioner

Faster Applications

10X speedup



32 billion queries/day

Storage Savings

1 Petabyte



Replaced \$\$ SAN

Maximum Availability

99.95%



Cost Savings

75% Cost Reduction



## Objectives

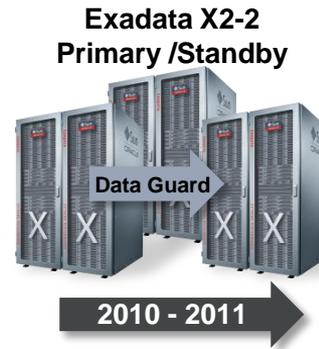
- Real-time processing and monitoring of all cargo and every passenger crossing the US border
- Shrink data center costs

## Solution

- 2010: Implement V2, X2-2
- 2011: Add X2-2, X2-8
- 2012: Add X2-8, Exalogic



*Automated Targeting Systems*



*Automated Export Systems*



*Passenger & Cargo Systems*

*Enterprise Data Warehouse*

# Oracle Beehive: Collaboration

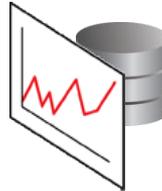


## Benefits

*"Beehive is our largest application in-house. It is Oracle's largest backend database."*

- Campbell Webb, Vice-President IT, Oracle

Capacity for Growth



Faster Response

5x – 60x



100% Uptime

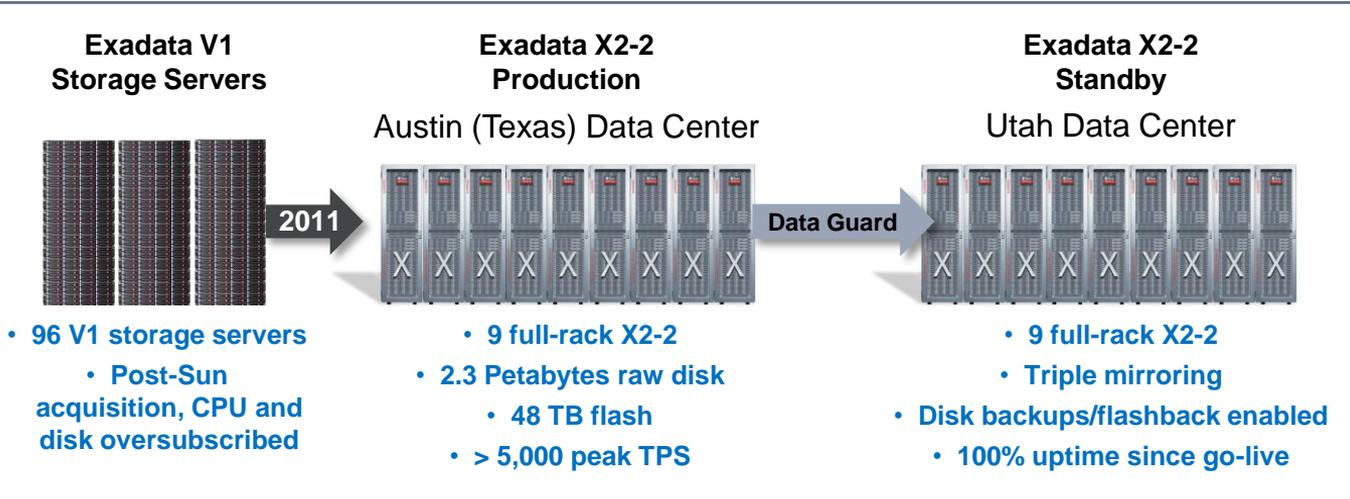


## Objectives

- Company-wide collaboration for > 100K users
- CPU/ storage growth 3+ years
- Improved response times
- Guarantee uptime

## Solution

- 2009: Move Beehive storage to Exadata V1 storage
- 2011: Migrate to Exadata X2-2



# Turkcell: DW and DB Consolidation



## Benefits

*“In a word, Oracle Exadata is fantastic. Almost no report takes more than 10 minutes to run, versus hours before. It sounds unreal, but it’s real.”*

- Power User, Finance Department, Turkcell

Faster Reports

10X



27 min to 3 min  
(avg for 50k rpt)

Storage Savings

900 TB



1,000 TB to 100 TB

Reduced Admin

20%



Data Center Cost Savings

80% Less Power  
30 m<sup>2</sup> Less Space

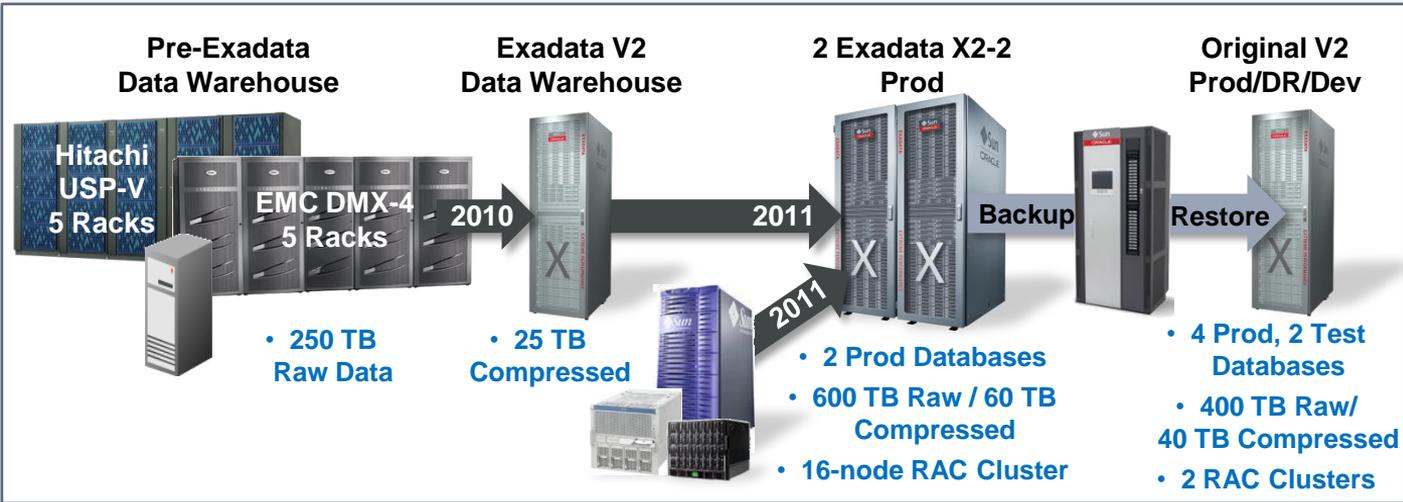


## Objectives

- Speed up BI
- Lean, green data center
- Prepare for big data growth

## Solution

- 2010: Replace 11 racks with 1 full-rack Exadata V2 for DW
- 2011: Add 2 full-rack Exadata X2-2s for DB consolidation



ORACLE

# SK Telecom: Revenue Assurance DW



## Benefits

*“Oracle Exadata Database Machine has overwhelmingly superior performance, with average data processing rates 8 to 20 times greater than other SK Telecom in-house systems.”*

- Jin-hyung Lee, Manager, Network Engineering Department, SK Telecom

Extreme Capacity

18+ TB/day



Storage Savings

1620 TB



1800 TB to 180 TB

Capture Lost Revenue

\$ Millions



## Business Objectives

- Ensure billing accuracy
- Enhance customer service
- Predict data traffic trends

## Solution

- 2009: Create new DW on Exadata
- 2010: Move to full-rack Exadata V2
- 2011: Add full-rack Exadata X2-2
- 2012: Clustered 3 full racks

Exadata V1  
Data Warehouse  
(2009)



Exadata V2  
Data Warehouse



Exadata V2+ two X2-2  
Data Warehouse



Exadata X2-2  
Qtr Rack (2010)



- Billing Analysis System

- Billing Analysis System
- Usage & Roaming Inquiries

- Raw Data 1800 TB
- Compressed to 180 TB
- DB size: 350 TB with indices
- 24-node RAC, 42 Exadata Cells

- Spam Filtering

ORACLE



# P&G: DW and Mixed Workload

## Benefits

*“Performance with Exadata is much better, and getting an out-of-the-box solution dramatically reduced the time and money to build and maintain our DW platform.”*

- Brian Beckman, DW Platform Manager, Procter & Gamble

Faster Applications

2 - 30X



Storage Savings

3x – 10x Compression



Simplified Support



## Objectives

- Improve stability and supportability
- Lower costs
- Boost performance

## Solution

- 2010: Implement 1<sup>ST</sup> DW on V2
- 2011: Move 2<sup>nd</sup> DW to X2-2, Move OLTP/Analytics to X2-2
- 2012: Move 3<sup>rd</sup> DW to X2-8

Exadata V2 Data Warehouse  
Prod Non-Prod



• Trade/Mkt/POS Data

Exadata X2-2 Data Warehouse  
Prod DR Dev UAT



• Shipments Data

Exadata X2-8 Data Warehouse  
Prod Non-Prod



• Fin'I/Planning Data

Exadata X2-2 OLTP/Analytics  
Prod Non-Prod



• Trade Fund Management

# Alpha Natural Resources: OLTP



Alpha Natural Resources

## Benefits

*“Oracle Exadata has enabled us to deliver exceptional service to our business users while reducing the cost involved . Exadata has been a big win for our business and for IT.”*

- Saul Hernandez, CIO

Faster Applications

5X+ Speedup



Reduced TCO

50%



Maximum Availability

99.95%



Data Center Cost Savings

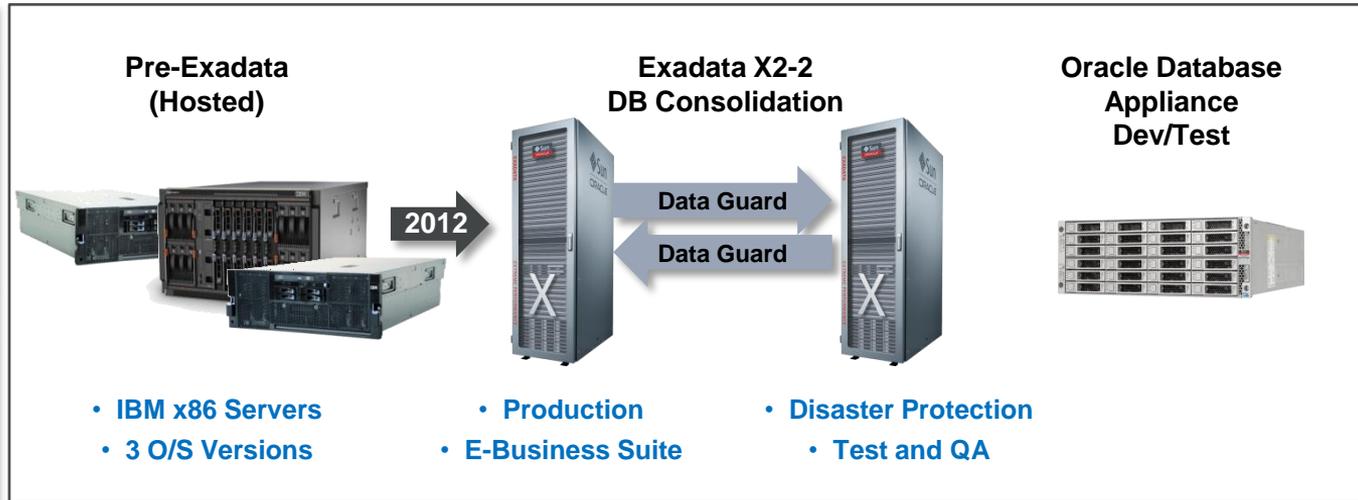


## Objectives

- Reduce cost and complexity
- Deploy Oracle E-Business Suite quickly
- Dramatically improve performance and availability

## Solution

- X2-2 for production E-Bus Suite
- X2-2 for Standby DR and test/QA
- Database Appliance for testing



ORACLE

# Garmin: OLTP Consolidation



## Benefits

*“Consolidation on Exadata reduced costs and eliminated critical performance bottlenecks in our Manufacturing and Planning systems”*

- Ed Link, Vice-President IT, Garmin

Faster  
Month End Reports

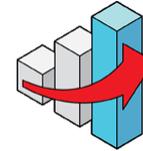
Up to 11X



99.95%  
Uptime



4x Growth



Data Center  
Cost Savings

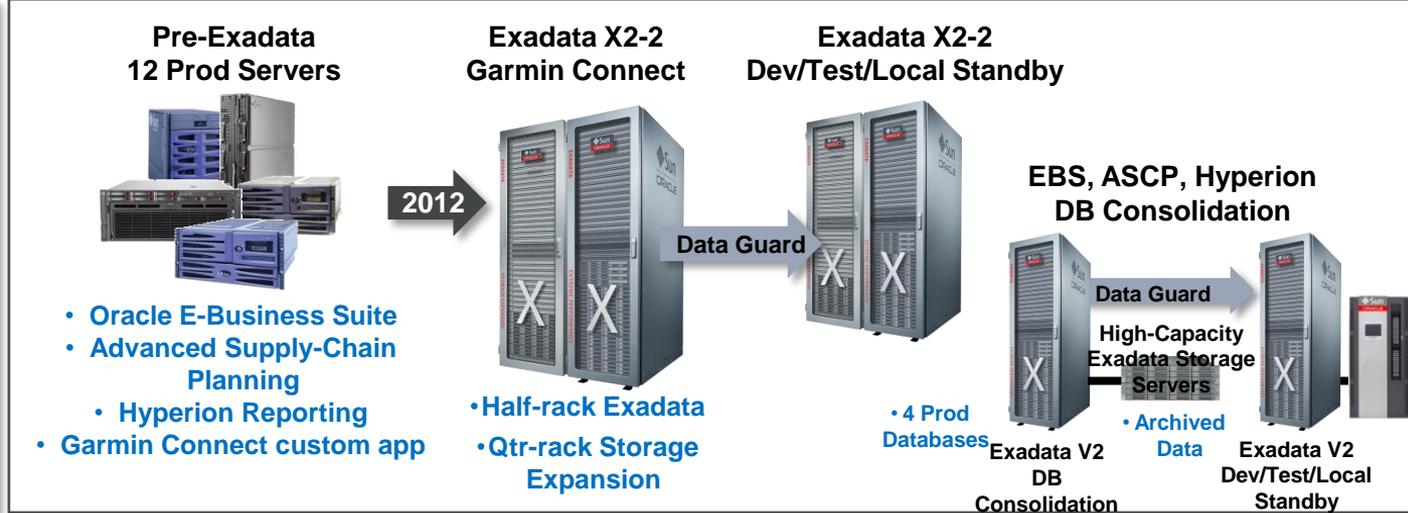


## Objectives

- Support 400% growth in customer facing applications
- Eliminate bottlenecks in Manufacturing and Planning
- Consolidate to reduce costs

## Solution

- 2012: Consolidate 12 servers onto two Half-rack Exadata V2
- 2012: Add two Half-rack X2-2 systems, two Qtr-rack storage expansion



# Lion: SAP on Exadata and Consolidation



## Benefits

*“Exadata has always high performance. Data Compression is also wonderful result. It is best consolidation database platform. Also, thanks to careful planning and verification, it was smooth without any major problems as a project.”*

- Masatoshi Utsunomiya, Director, Integration system department, Lion

### Faster Application

Max. up to 100x



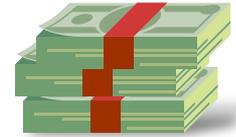
Avg. up to 20x

### Storage Savings

SAP FI: 2x  
Teradata: Over 5x



### Operation Cost Savings



## Business Objectives

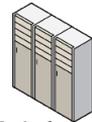
- Reduce TCO
- Improve Operation and Maintenance Cost
  - ✓ Standardize the environment

## Solution

2010: Project start Exadata V2 Quarter Rack  
2012: Software Upgrade X2-2 SAP DB Consolidation

### Mainframe Migration

- Sales/Logistics, Master Data Cost Mgmt, Sales Analytics



Mainframe

Migrated Aug 2012

### SAP Database Migration

- SAP FI (from Linux)
- Consolidate 3 systems
- Compress 1400 GB to 700 GB



IA Server

Migrated May 2012

### Teradata DWH Migration

- Compress 320GB to 60GB



Migrated Mar 2013



✓ Oracle Exadata Quarter Rack

✓ DWH used HCC, SAP used Advanced Compression

✓ Delete useless index of Teradata

# Univ. of Minnesota: DBaaS



## Benefits

*“We consolidated dozens of database servers onto Exadata and freed up many of our admins for more strategic tasks. Standardizing our database services and configurations has yielded benefits across many dimensions.”*

- Andy Wattenhofer, Database Administration Manager, University of Minnesota

### Server Consolidation

40 servers → 2 Exadata



235 databases → 29

### Reduced Administration



5 DBAs → 3

### Standardized Configurations

Security, Backup, HA/DR

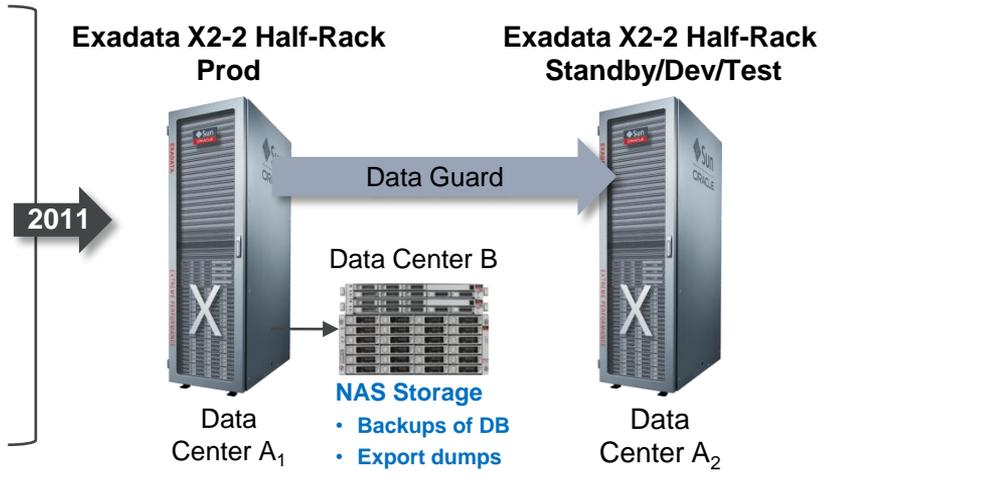


## Business Objectives

- Eliminate server proliferation
- Provide consistent service levels
- Free up redundant administrators
- Standardize roles and permissions

## Solution

- Consolidate 40 DB servers onto 2 Exadata X2-2 Half-Racks
- Create a standard Schema-as-a-Service with 150+ schemas in 8 databases, serving all departments



# CBA: DBaaS

## Benefits

*“As the business seeks more services on demand, we needed to provide that resilient platform and Exadata provides us a consolidated platform for database services that can meet all those needs”*

- Nicolas Tan, Head, Infrastructure and Platform Services

OPEX

50~60%  
reduction



ROI

150% in 5  
years



Rapid  
Provisioning

Months to  
Minutes



Simplified  
Management

Two Oracle Homes

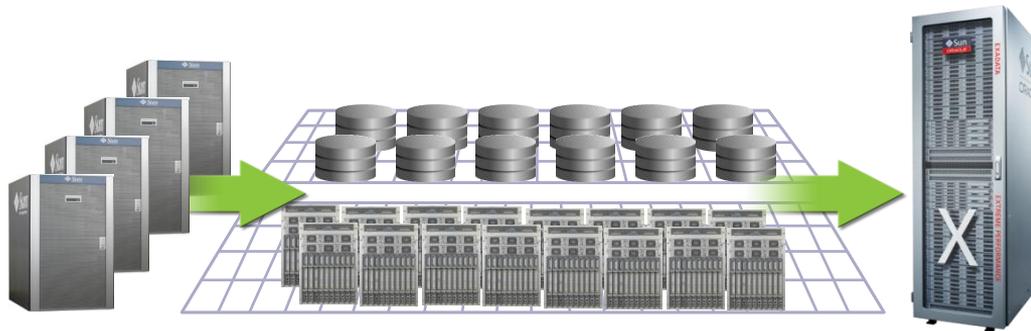


## Business Objectives

- Cost Savings
- Rapid Provisioning
- Extreme Standardization

## Solution

- Exadata for
  - Rapid, reliable deployments
  - Standardized environments



Enterprise servers  
configured by CBA

Commodity clusters  
configured by CBA

Commodity clusters  
configured by Oracle

# Major Semiconductor Company: Manufacturing Data Warehouse

## Benefits

*“Exadata transformed production data analysis. Exadata improved performance better than 5x and reduced costs by more than 50%.”*

### Faster Applications

- 10-15 x Faster Data Load
- 5-50x Faster Queries



### Data Center Cost Savings

- 5x Less Hardware
- 80% Less Power



### Storage Savings

- ~ 9 x Data Compression



### Reduced Admin Serviceability

- 40% Savings Single Patch Platinum Services



## Business Objectives

- Improve yield management by correlating data across fabs
- Reliability, availability, scalability
- Reduce support costs and improve serviceability
- Lower data center costs

## Solution

- 2013: 2 x X3-8 Full Rack; X3-2 Half Rack; 2 x ZFS Storage

### Pre-Exadata Data Warehouse (2012)



### Exadata X3-8 Data Warehouse



- 7 identical DWs (1 per fab)
- 63 servers, 14 storage racks

### Exadata X3-8 Active Data Guard



2013

- Production analysis
- 7 production DBs

- DR and additional reporting
- Stdby 1<sup>st</sup> patches & rolling upgrades

### Exadata X3-2 Half Rack Dev/Test



- Oracle Platinum Services

ORACLE



# Loblaw: JDA on Exadata

## Benefits

"Exadata is fast. We achieved 6x performance improvements over our existing platform in our forecasting and replenishment batch run, enabling our store level forecasting to complete consistently within SLA and scaling up seamlessly to full volume without performance degradation".

- David Markwell, VP of IT, Technology Services Management

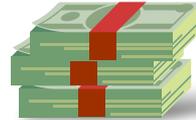
Faster Applications

2.5X -6X



Operation Cost Savings

Reduced Admin by 40%



Simplified Support

1 Week Deployment



Data Center Cost Savings

80% Less Power  
30 m<sup>2</sup> Less Space



## Objectives

- Eliminate Bottlenecks in Supply Chain Planning
- Reduce the long running batch window
- Reduce support time and cost

## Solution: 2013

- 2013: Migrated from IBM-P7 to X2-8 for production.
- X2-2+ for DR (using Active Data Guard.) and Dev/Test

Pre-Exadata



- IBM P7
- Hitachi USP V w/ SSDs

2012

Exadata X2-8 Production



- JDA (Supply Chain Planning)
- Performance boost with no application change using Exadata Smart Flash Cache and Smart Scan

Active Data Guard

Exadata X2-2+X3-2 Disaster Recovery



- ½ rack: X2-2 ¼ rack upgraded with X3-2 ¼ rack

Exadata X2-2+ X3-2 Dev/Test



- ½ rack: X2-2 ¼ rack upgraded with X3-2 ¼ rack

Turning 2<sup>nd</sup> Day Delivery into Next Day Delivery

# Starwood: DW and Mixed Workload

## Benefits

*“14x performance improvements enable our field managers to take proactive actions not possible before. Reporting with ‘real time’ data is rapidly becoming the expectation at Starwood .”*

Gordon Light, Sr. Dir., Datawarehouse, Starwood

Faster  
Reports

14X



Faster  
ETL

16 hours to 3



Simplified  
Support

Single Vendor  
> 25% cost reduction



## Business Objectives

- Improve application response time
- More current reports
- More responsive marketing

## Solution

- 2011: X2-2 full rack for prod and DR
- 2013: full Storage Expansion Rack; 2 x ZFS Backup Appliance

Pre-Exadata  
(2011)



Exadata X2-2 Full Rack  
Production  
Dev/Test



Exadata X2-2 Full Rack  
Disaster Recovery



2011

Storage  
Expansion

2013

Active Data Guard  
Golden Gate

WAN @  
2,700 miles

Remote Backup

2013

- 2 servers @128 cpu each
- DR:RMAN to disk, NetBackup to tape
- 45 TB on 2 frames
- DB 10.2
- HA: RAC; DR SRDF

- Apps: Loyalty program; marketing & revenue analytics
- DB 11.2

- Offsite reporting
- RMAN to ZFS; ZFS-ZFS

# WestJet: Siebel on Exadata



## Benefits

“WestJet consolidated 16 databases while delivering exceptional and consistent performance to our online ticketing, customer rewards and loyalty programs. Performance is at least 19 times faster.”

-Kris Trzesicki, DBA, WestJet

Faster Response

19x – 260x



Data Center Cost Savings

18 servers → 2 Exadata



Reduced Admin Serviceability

70% Savings  
Single Patch  
Platinum Services



Zero Unplanned Downtime



## Business Objectives

- Consolidate DBs
- Consistency of performance
- Availability, especially during rolling upgrades
- Time to market
- Lower data center costs

## Solution

- 2013: 3 x X2-2s ¼ Rack

Pre-Exadata (2012)



- HP Unix (2) + 23+ x64 servers
- DB 11g R1

Exadata X2-2 DB Consolidation



2012 /13

- Siebel, OBIEE, Booking & Notification, + other
- DB 11gR2
- Consolidated 16 DBs
- 1/2 rack: ¼ staging and ¼ prod'n

Exadata X2-2 Active Data Guard



- DR: ¼ rack

Dev/Test



Oracle Platinum Services



# Ziraatbank: Core Banking System



## Benefits

*“Oracle Exadata enabled us to support more users with much better response time while at the same time improving uptime for our core banking OLTP system.”*

- Serdar Mutlu , Manager, Database Systems

Smaller Batch Window

60% less time



Consolidation of Databases

40% fewer Admin tasks



System Utilization

70% to 30%



Zero Data Loss HA Protection

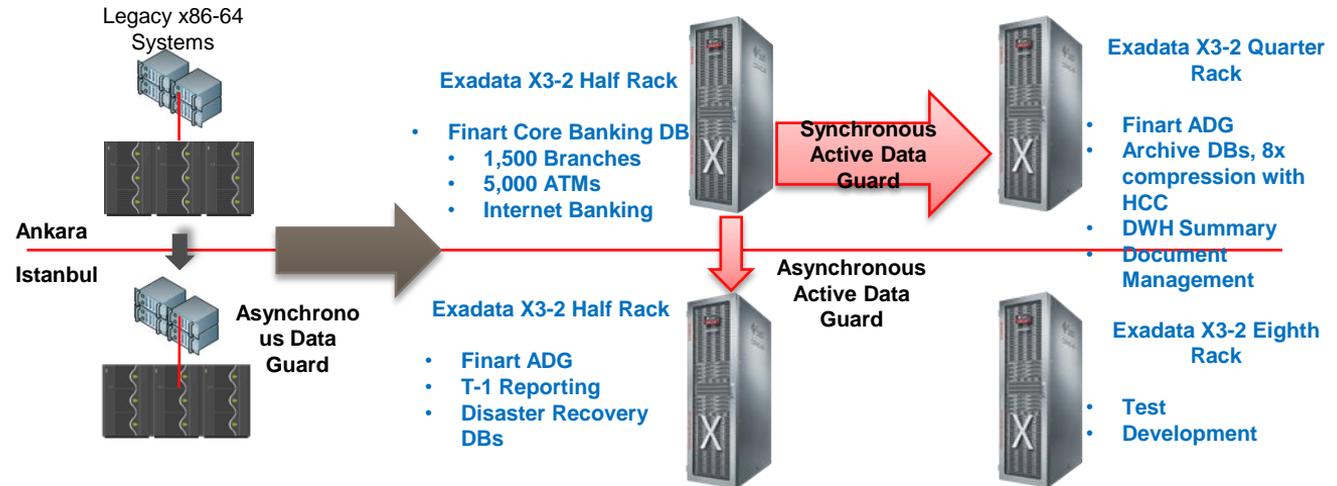


## Objectives

- Meet OLTP SLAs (99.95% availability, disk response time under 3ms, 25,000 users, 3,000 TPS)
- Zero Data Loss HA configuration
- Reduce overnight batch window from 8 hours to 3 hours
- Stability and availability with mixed workloads

## Solution

- 2013 Q2: 2 x Half Racks, and 1 x Quarter Rack X3-2 for production
- 2013 Q3: 1 x Eighth Rack for test and development



ORACLE

# Softbank Replaces 36 Teradata Racks

3 Exadata racks

150% more data capacity

12:1  
advantage

36 total racks

Twice the operational cost  
of Exadata



- Billions of CDRs processed in 7 hours (from 25 with Teradata)
- Power, cooling, space savings
- Maintenance charges slashed
- Up to 8x faster



# Exadata X4 vs. Traditional Storage Arrays

52  
GB/sec

100  
GB/sec

Largest Storage Array



1 Rack Exadata



# Exadata Flash Performance Scales Linearly

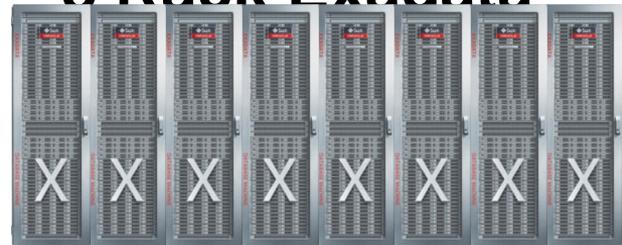
800  
GB/sec

52 GB/sec

Largest Storage Array



8 Rack Exadata



# Exadata Database Machine



- **The ultimate platform for all database workloads**
  - OLTP, Warehousing, Database as a Service
- **Most advanced hardware**
  - Fully scale-out servers and intelligent storage with unified InfiniBand connectivity and PCI flash
- **Most advanced software**
  - Database optimized compute, storage, and networking algorithms dramatically improve performance and cost
- **Standardized, optimized, hardened end-to-end**

**Hardware and Software**

**ORACLE®**

**Engineered to Work Together**

**ORACLE®**