

# Top Exadata Features You Need to Know

**25<sup>th</sup> April , 2012**



**Anuj Mohan**

**Sr. Oracle DBA**

**Dbadirect Inc.**

**Oracle certified Exadata Implementer**

**Oracle Certified RAC Expert**

**Oracle 11g Certified Professional (OCP)**

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

[www.collaborate12.org](http://www.collaborate12.org)

[www.collaborate12.ioug.org](http://www.collaborate12.ioug.org)


# About dbaDIRECT

**Our Oracle Exadata consulting service can completely revolutionize the value your enterprise is able to extract from its data environment by providing you with:**

- Database administrators certified by Oracle to provide Exadata services
- Expertise gained from a broad customer experience base
- Documented best practices specific to your installation
- 24x7 monitoring of your Exadata solution
- Reduce IT costs through consolidation of your database systems and hardware

**dbaDIRECT has provided over a decade of remote database administration and DBA consulting services. We offer 24x7 support for all major platforms, including Oracle, MS SQL Server, DB2, Sybase, and MySQL. Our customers rely on us for:**

- Remote Database Administration and ONdemand Services
- Staffing
- Consulting
- Upgrades
- Performance Tuning
- System Administration



This presentation is intended for informational purposes only. The contents provided are based on my experiences and has no association with Dbadirect or Oracle

## What is Exadata ?

### Definition...

Oracle Exadata Database Machine provides an optimal solution for all database workloads, ranging from scan-intensive data warehouse applications to highly concurrent online transaction processing (OLTP) applications.

Combination of smart Oracle Exadata Storage Server Software, complete and intelligent Oracle Database software, and the latest industry-standard hardware components, Oracle Exadata Database Machine delivers extreme performance in a highly-available, highly-secure environment.

## Types of Exadata Database Machines

X 2-8 Exadata Database Machine Full Rack

X 2-2 Exadata Database Machine Full Rack

X 2-2 Exadata Database Machine Half Rack

X 2-2 Exadata Database Machine Quarter Rack

## X2-2 Full Rack

8 x64 dual processor Sun Fire X4170 M2 database servers

- Each database server has 2 cpu with 6 cores

- 8 servers \* 2 cpu \* 6 cores/cpu =96 cores (12 per database server)

- 8 servers \* 96G RAM = 768GB Memory

  - \*\*\*Expandable to 144GB/Server

- PCI Flash - 5.3TB

14 Exadata storage servers each with 12 disks

- All 600GB SAS disks for high performance ~15000RPM

- All 3TB SAS Disks for High capacity ~7400RPM

  - \*\*\* You cannot mix and match either all high performance or all high capacity

3 Infiniband switches with 36 ports

1 Cisco Ethernet switch –Admin network

redundant power distribution units (PDUs)

## X2-2 Half Rack

4 x64 dual processor Sun Fire X4170 M2 database servers

- Each database server has 2 CPU with 6 cores

- 4 servers \* 2 CPU \* 6 cores/CPU =48 cores (12 per database server)

- 4 servers \* 96G RAM = 384GB Memory

  - \*\*\*Expandable to 144GB/Server

- PCI Flash – 2.6TB

7 Exadata storage servers each with 12 disks

- All 600GB SAS disks for high performance ~15000RPM

- All 3TB SAS Disks for High capacity ~7400RPM

  - \*\*\* You cannot mix and match either all high performance or all high capacity

3 Infiniband switches with 36 ports

1 Cisco Ethernet switch –Admin network

redundant power distribution units (PDUs)



## X2-2 Quarter Rack

2 x64 dual processor Sun Fire X4170 M2 database servers

-Each database server has 2 CPU with 6 cores

-2 servers \* 2 cpu \* 6 cores/cpu =24 cores (12 per database server)

-2 servers \* 96G RAM = 192GB Memory

\*\*\*Expandable to 144GB/Server

-PCI Flash – 1.1TB

3 Exadata storage servers each with 12 disks

-All 600GB SAS disks for high performance ~15000RPM

-All 3TB SAS Disks for High capacity ~7400RPM

\*\*\* You cannot mix and match either all high performance or all high capacity

2 Infiniband switches with 36 ports

1 Cisco Ethernet switch –Admin network

redundant power distribution units (PDUs)

References: <http://www.oracle.com/technetwork/database/exadata/dbmachine-x2-2-datasheet-175280.pdf>

## X2-8 Full Rack

2 Servers with 8 - Ten-Core Intel® Xeon® E7-8870 Processors (2.40 GHz)

- Each database server has 8 CPU with 10 cores

- 2 servers \* 8 cpu \* 10 cores =160 cores (80 per database server)

- 2 servers \* 2TB RAM = 4TB Memory

- PCI Flash - 5.3TB

14 Exadata storage servers each with 12 disks

- All 600GB SAS disks for high performance ~15000RPM

- All 3TB SAS Disks for High capacity ~7400RPM

\*\*\* You cannot mix and match either all high performance or all high capacity

3 Infiniband switches with 36 ports

1 Cisco Ethernet switch –Admin network

redundant power distribution units (PDUs)

References: <http://www.oracle.com/technetwork/database/exadata/dbmachine-x2-8-datasheet-173705.pdf>

## Exadata Implementation process

### STEP 1

Complete the DBM Configuration spreadsheet .Oracle needs it to understand customer environment and later on generate configuration scripts. This step involves Oracle and Customer both.

**Next**



## Exadata Implementation process

### STEP 2

Based on configuration sheet and Inputs from Oracle, customer needs to do site preparation. This includes allocating IP addresses .Exadata uses a lot of IP addresses for example X2-2 Full Rack needs 70+ IP addresses.

**Next**



## Exadata Implementation process

### STEP 3

Oracle Engineer will come to onsite to review site preparation/survey.

**Next**



## Exadata Implementation process

### STEP 4

Exadata machine will be delivered to your facility. Unpacking and transferring to its final location will be done by oracle.

**Next**



## Exadata Implementation process

### STEP 5

Oracle and customer will plan installation date and final last checks. They will not come to this step if any of network preparation is pending.

**Next**



## Exadata Implementation process

### STEP 6

Oracle hardware engineer will come onsite typically 2-3 days before Oracle ACS engineer is scheduled for the actual installation. The engineer will inspect the hardware and bring up the system. They will also configure KVMs Infiniband and PDUs

**Next**






## Exadata Implementation process

### STEP 7

Oracle ACS engineer will come onsite typically for 3-5 business days. They will complete the rest of the installation. Once this process is completed you will have an Exadata machine with one database up and running.

**Next**



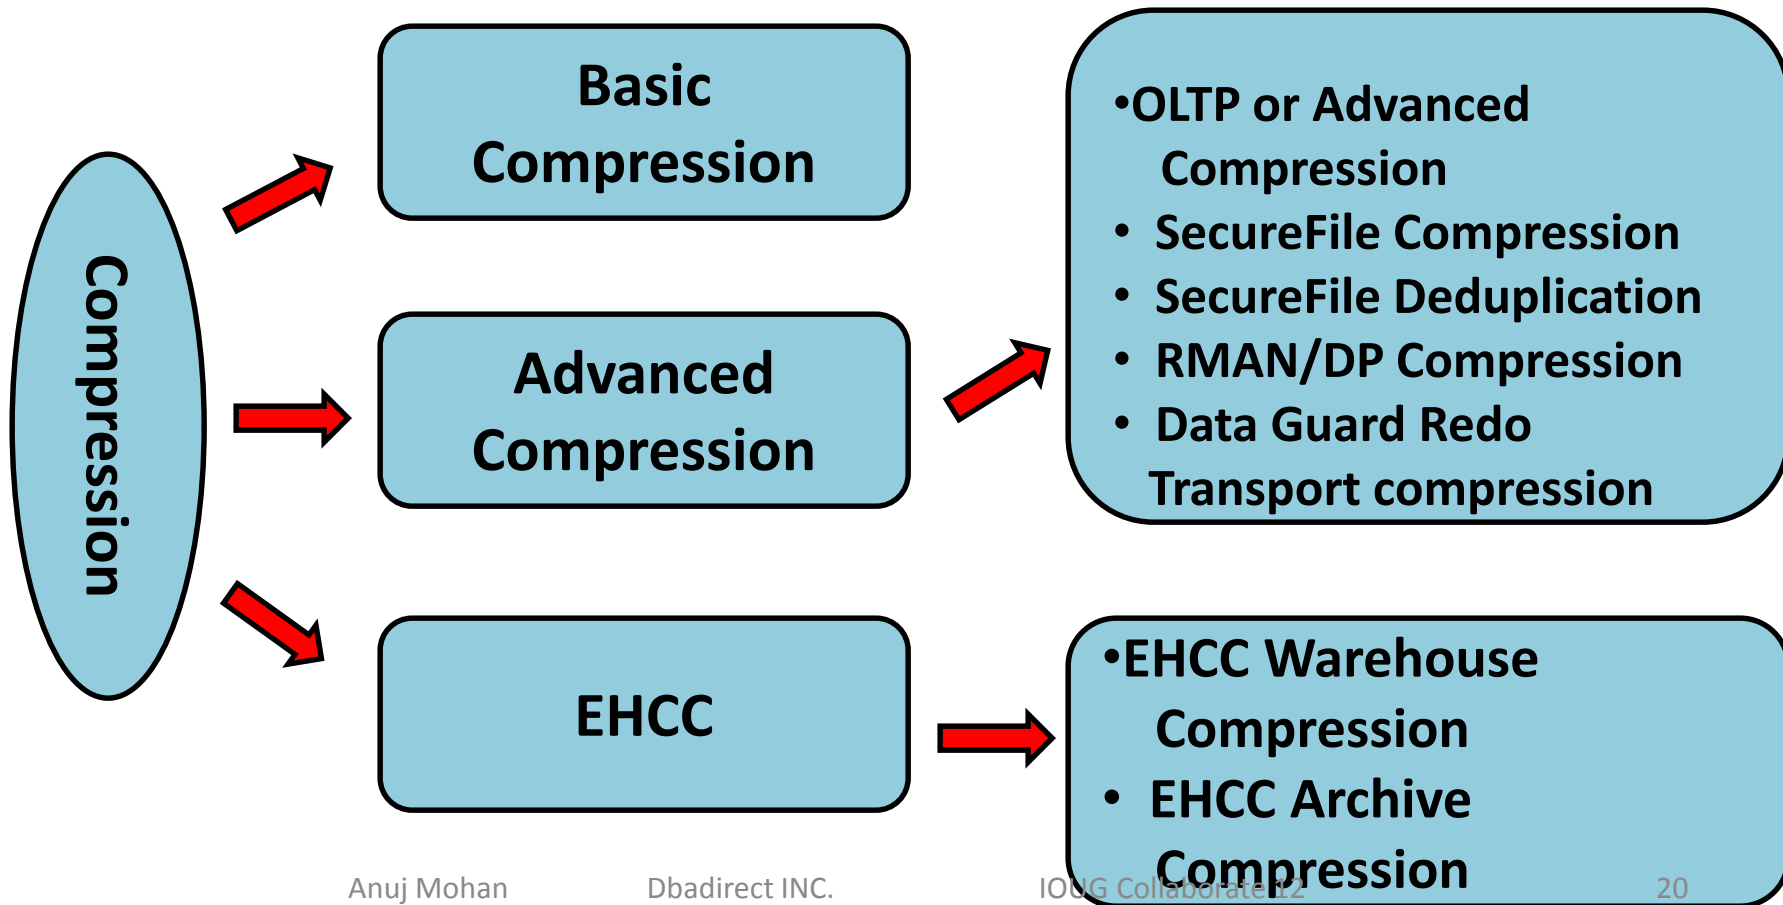


## Exadata Hybrid Columnar Compression (EHCC)

Exadata Hybrid Columnar Compression (EHCC) is a new compression technology exclusive to Exadata. Oracle explains that you can achieve 10x to 50x compression based on what type of compression you choose and your data.

Exadata Hybrid Columnar Compression technology (EHCC) is a new method for storing data within a db block. Data within the data block is neither stored row wise sequentially nor column wise. As the name implies, EHCC utilizes a combination of both row and Columnar methods for storing data. Because of the hybrid approach to achieve this compression we get the benefits of columnar storage while avoiding the performance drawbacks of a pure columnar format.

# Exadata Hybrid Columnar Compression (EHCC)



## Exadata Hybrid Columnar Compression (EHCC)

DBMS\_Compression package can calculate compression benefits in Advanced compression or EHCC

```
DBMS_COMPRESSION.GET_COMPRESSION_RATIO (
scratchtbsname => An_scratchtbsname,
ownname       => An_ownname,
tablename     => An_tabname,
partname      => An_partname,
comptype      => An_comptype,
blkcnt_cmp    => An_blkcnt_cmp,
blkcnt_uncmp  => An_blkcnt_uncmp,
row_cmp       => An_row_cmp,
row_uncmp     => An_row_uncmp,
cmp_ratio     => An_cmp_ratio,
comptype_str  => An_comptype_str );
```

## Exadata Hybrid Columnar Compression (EHCC)

Exadata Hybrid Columnar Compression (EHCC) is new compression technology exclusive to Exadata. Oracle explains that you can achieve 10x to 50x compression based on what kind of compression you choose and your data.

### How it works

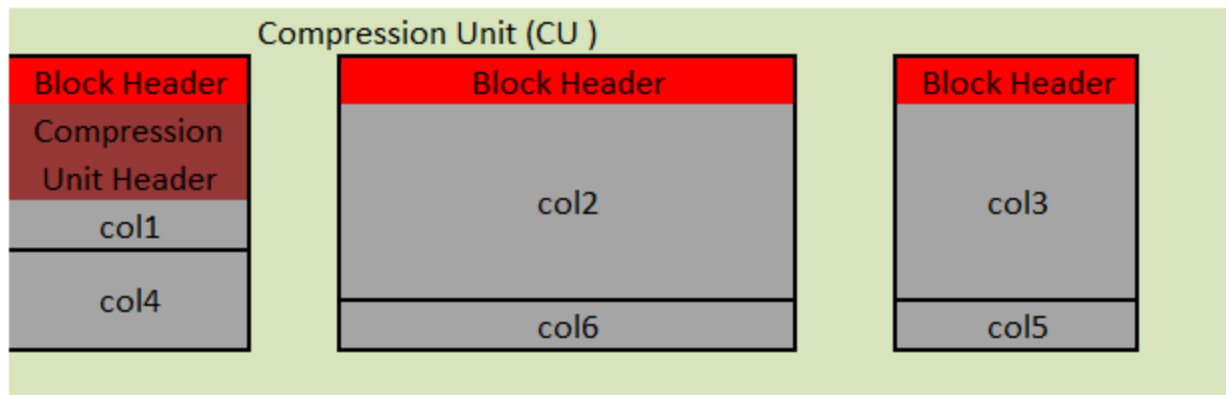
- Table blocks are organized into compression units ( CU )
- Within a compression unit, data is stored by columns instead of rows
- Column wise storage allows similar data to stay together which saves space.
- How many rows that will fit into a compression unit is determined at load time based on the size of the row and the estimated compression.
- Each column will compress separately

## Exadata Hybrid Columnar Compression (EHCC)

- To compress in EHCC you need to load data using direct load
- Parallel DML , Insert /\*+APPEND\*/ and direct path SQL\*LDR
- Individual row lookup need more CPU as it has to reconstruct row

### Negative points

- Updates migrate row to different location with lower compression
- All rows in CU will be locked for update
- Regular insert will give lower compression



## Exadata Hybrid Columnar Compression (EHCC)

### EHCC Supports

- B-Tree/Bitmap indexes
- MV
- Partitioning
- Parallel Query /DML/DDL
- Data Guard and physical standby support



## Exadata Hybrid Columnar Compression (EHCC)

### Two types of compression

#### Warehouse Compression

- 10X Average compression

- Good for warehouse

- Similar space gain in storage

- IO performance increase on same scale

- 50K IO/s will reduce to 5K IO/s

#### Archive compression

- 15-70X Compression

- Good for storing historical information

- Data which was supposed to go to tape will be online

- Runs your queries without recovering from tape

## EHCC –Syntax for Exadata

### Warehouse Compression syntax

```
Create table DP_Details (  
-----  
-----  
) Compress for QUERY [LOW/HIGH];
```

```
Create table DP_Details (  
-----  
-----  
) Compress for ARCHIVE [LOW/HIGH];
```

## Exadata Hybrid Columnar Compression (EHCC)

### Data Archiving Strategy

#### OLTP

- Partitioning
- Archive compression


#### Warehouse

- Partitioning
- Warehouse compression
- Archive compression

## Exadata Hybrid Columnar Compression (EHCC)

### Quick Tips

- Compression Advisor may have issues with Uniform TS
- If you want to use EHCC in OLTP, partition it and use EHCC on old partitions
- Alter table move can fix EHCC
- `Dbms_compression.get_compression_type` package can be used to know compression
- To increase load speed –use query low or load uncompressed and compress later
- Make sure you use direct path load
- Staging data on DBFS can give better performance on load
- You will see best benefit on IO bound queries



## Database Resource Manager ( DBRM )

With DBRM you can set priority of resources to database users or group of database users. You not only prioritize resources to users or user groups rather you can guarantee certain resources to certain users/groups.

Useful for consolidating workloads on the Exadata db machine . It is foundation for IORM which is Exadata only feature.

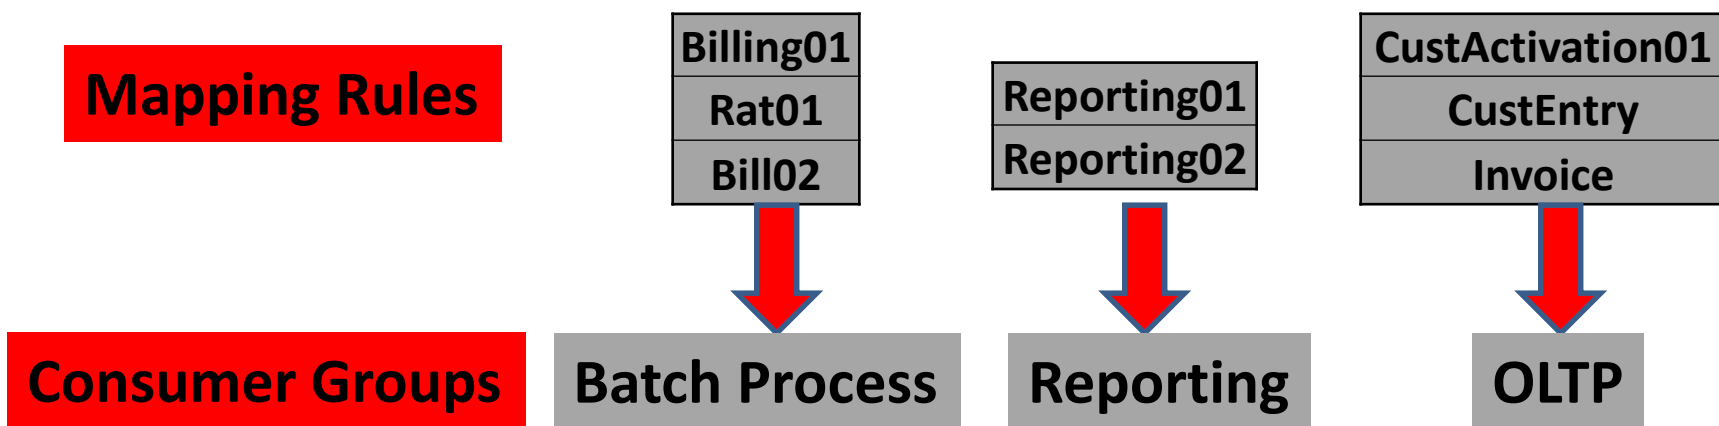
## Database Resource Manager ( DBRM )

### DBRM's Quick Scheduling

- 100ms Quantum
- Low priority session / High priority session switch in 1 quantum
- BG Processes are not managed by DBRM and they do not consume a lot of CPU
- If a high priority consumer group is not using CPU then it will be given to low priority consumer group based on resource plan
- Assign DOP for queries of particular consumer group
- Control percentage of parallel servers a consumer group can use

## DBRM – Basic Implementation Steps

- Group database sessions/users/service with common performance goals into one consumer group





## DBRM – Basic Implementation Steps

- Allocating resources to newly created consumer groups
- 3 Key methods

### Method 1: Ratio

25%



Reporting

15%



Batch Process

60%



OLTP

### Method 2: Priority

Priority 2



Priority 3



Priority 1



## DBRM – Basic Implementation Steps

- Allocating resources to newly created consumer groups
- 3 Key methods


		Method 3: Hybrid	
		P1	P2
Reporting	→	15%	50%
Batch Process	→	05%	40%
OLTP	→	80%	10%

## DBRM – Basic Implementation Steps

- Enable Resource Plan
- Two ways to Enable
- Manual – set parameter `resource_manager_plan`
- Automatic – Associate them with scheduler windows

### References

[http://docs.oracle.com/cd/B19306\\_01/server.102/b14231/dbrm.htm](http://docs.oracle.com/cd/B19306_01/server.102/b14231/dbrm.htm)



## I/O Resource Manager (IORM) – Basics

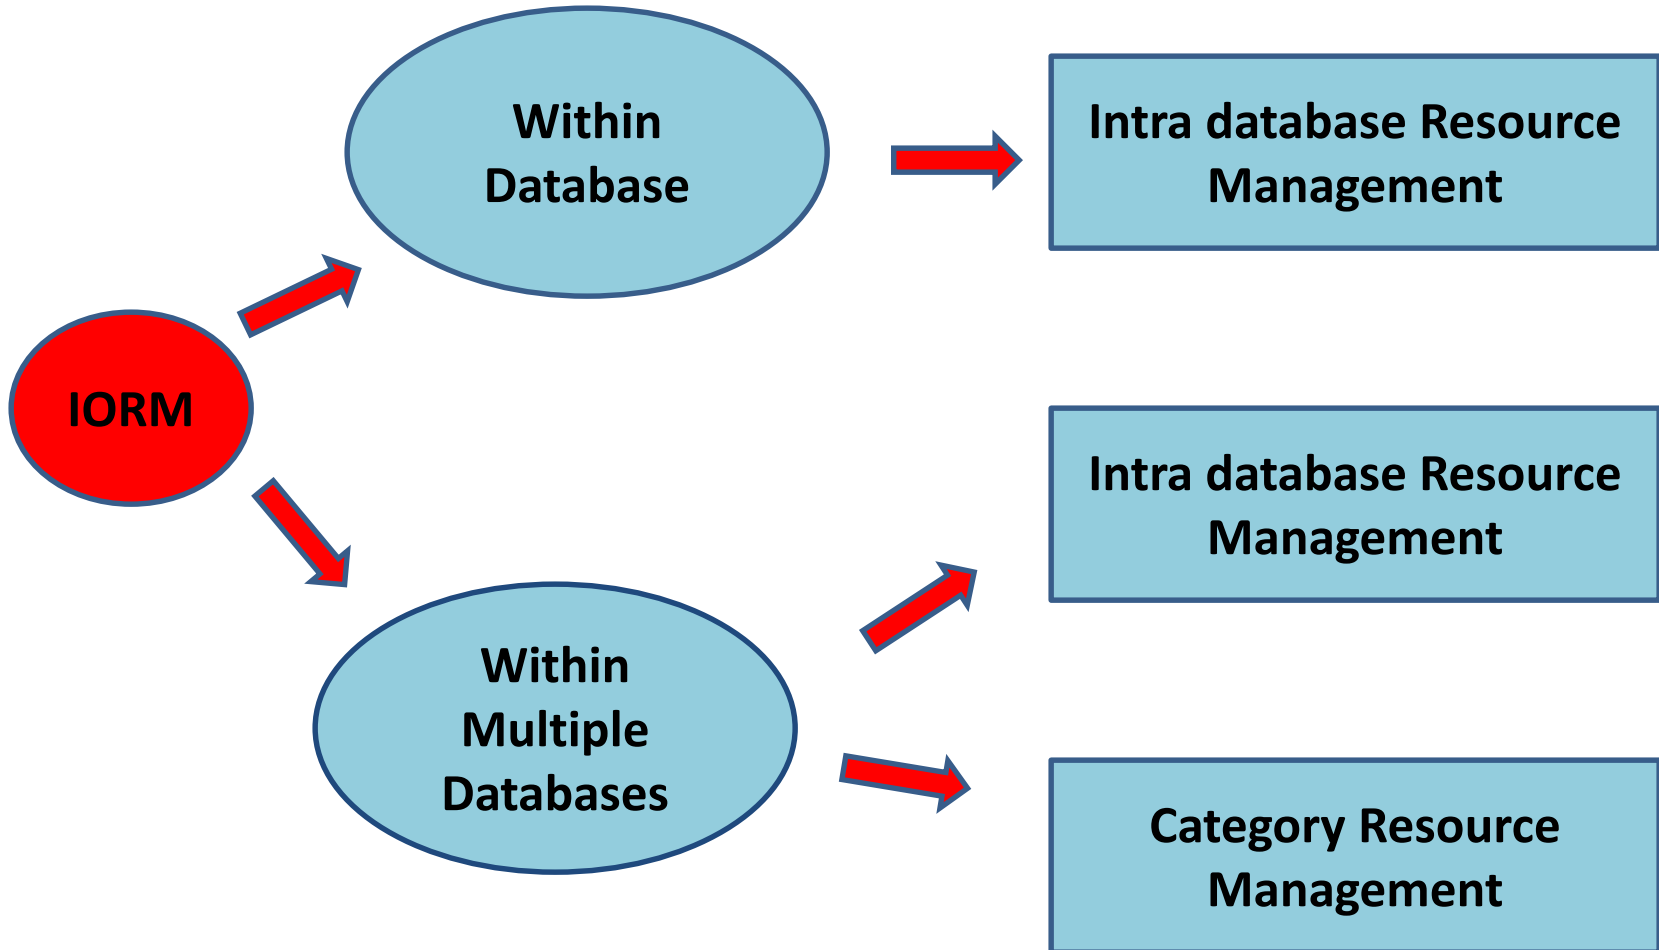
- You can manage resources using DBRM within database by creating consumer groups and resource plans.
- With Consolidations “a business need”, when you run multiple databases on single server you need to manage resources across databases (inter database). There are two Exadata specific features which can help on this goal. First is Instance Caging and second IORM. Let’s start with IORM

## (IORM) – Limiting I/O Bandwidth

- IORM Resolves Oversubscription in same manner as DBRM



# IORM – Broad Picture



## IORM –Intra or Within database


- IORM can manage multiple consumer groups in single database
- Works at database level implemented by DBRM and Resource Plan
- Consumer groups –made up of sessions of similar performance Objective
- Resource plan with I\O prioritization for consumer groups




## IORM – Consumer Groups

- Group database sessions/users/service with common performance goals into one consumer group
- Consumer groups – Batch Process / Reporting / OLTP
- You need to create rules to map sessions to consumer groups based on session attributes


### Mapping Rules



**Batch Process**



**Reporting**



**OLTP**

### Consumer Groups

## IORM – Resource Plan

- Allocating resources to newly created consumer groups
- 3 Key methods


### Method 1: Ration

25%



Reporting

15%



Batch Process

60%



OLTP

### Method 2: Priority

Priority 2

Priority 3

Priority 1

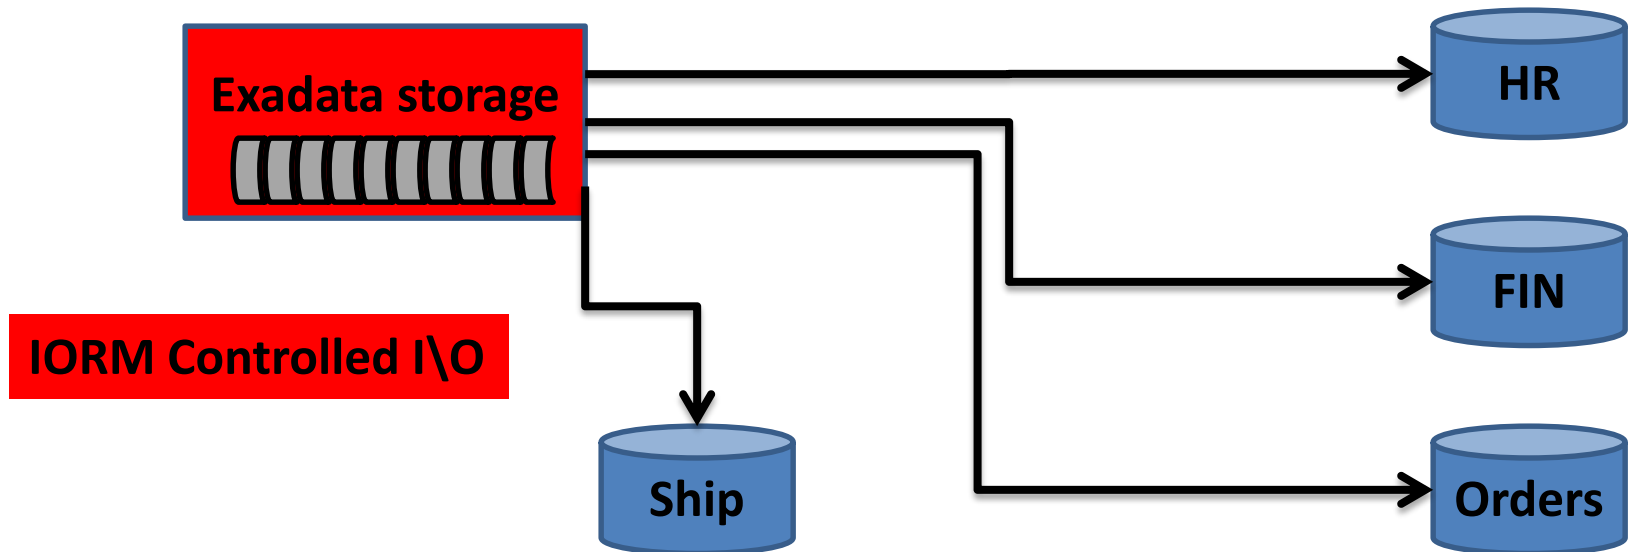
## IORM – Hybrid Plan

- Allocating resources to newly created consumer groups
- 3 Key methods

		Method 3: Hybrid	
		P1	P2
Reporting	→	15%	50%
Batch Process	→	05%	40%
OLTP	→	80%	0%

## IORM – Inter database

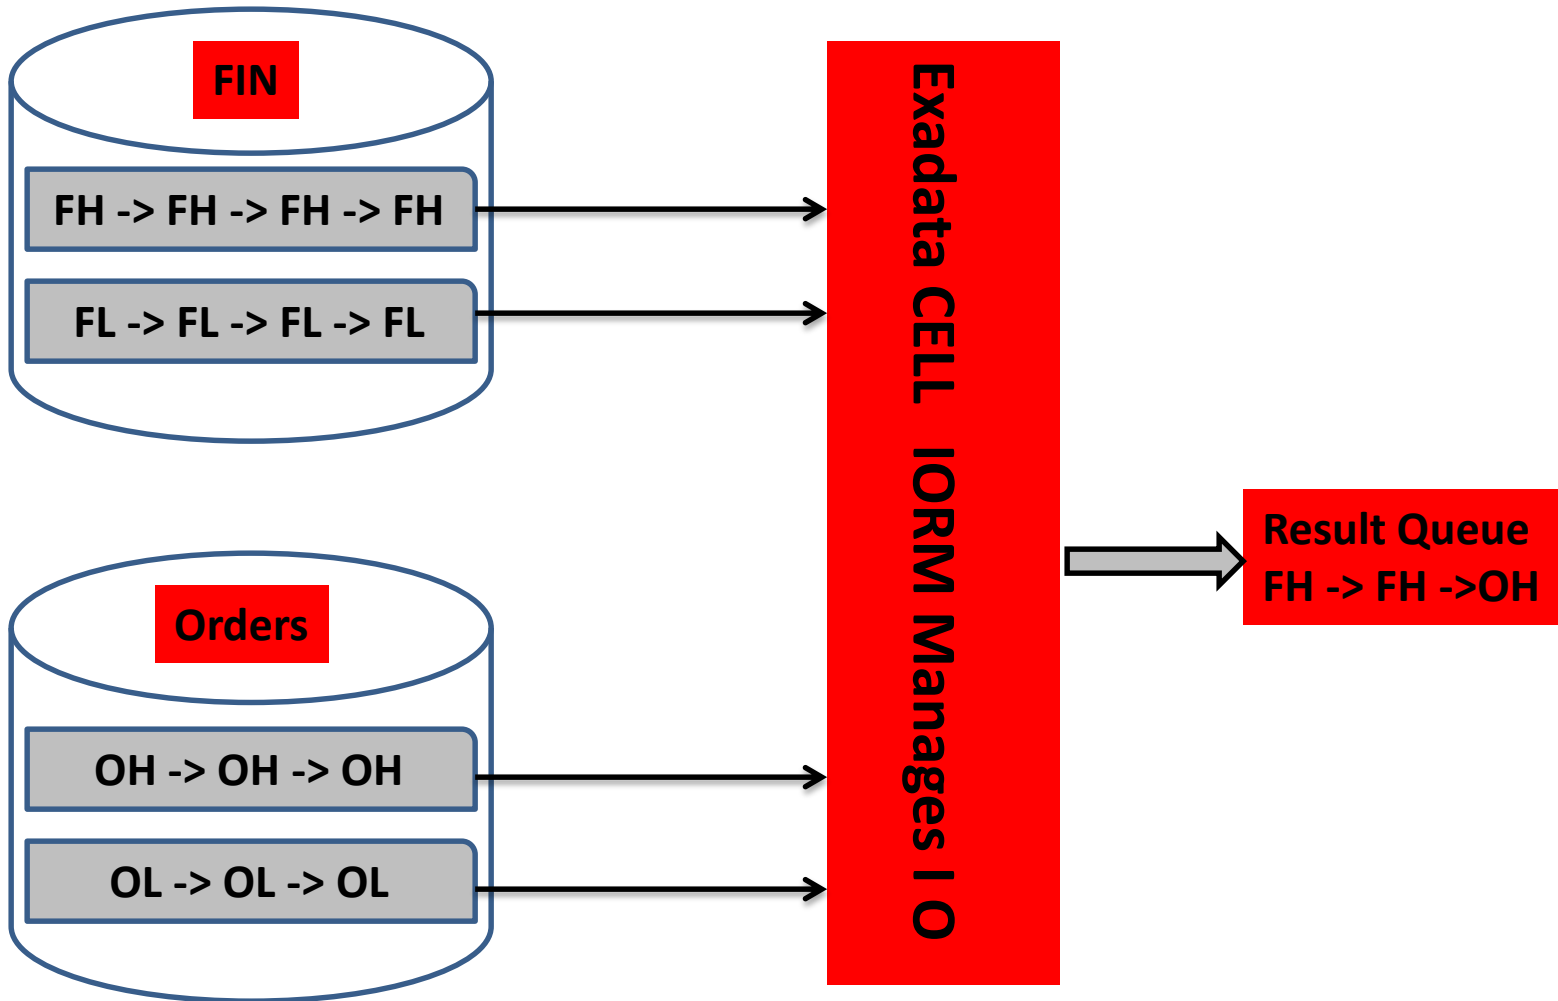
- IORM can manage multiple database on common Exadata storage
- Can partition I\O resources among multiple production instances
- Non production instances will be restricted not to take production resources



## IORM –Inter database

- Configure Inter database plans using CellCLI on each Exadata cell
- Enable plans using CellCLI on each Exadata cell
- Have multi-level plans where lower level plans only get resources which are not consumed by the higher priority plans
- If the inter DB plan is all that is configured , I\O resources will be distributed among databases as per the inter DB plan specifies.
- If Inter and Intra plans are both configured.
  - I\O resources will be distributed among databases as defined by the specified plans.
  - Resources are distributed in databases among consumer groups as per resource plan defined in the intra database plan.

## IORM – Inter database



## IORM –Category

- Consumer groups can be grouped into categories based on the nature of their workload.
- The Ideal way to implement this is to define the priority of the workload based on its type and not on the basis of where it is running.

Database	Consumer Group	Workload Type
HR	Payroll	V. Critical
	Benefits	Important
Orders	NewCust	V. Critical
	RetCust	V. Critical
	Reporting	Critical
Fin	DSS	Critical
	Maintenanace	Important

## IORM –Category

- Category plans are created and activated using CellCLI
- Category is an attribute of each consumer group
- Category plan defines resources for every category

### Priority based category plan

Reporting



Priority 2 ( Critical )

Batch Process



Priority 3 ( Important )

OLTP




Priority 1 ( V. Critical )



## IORM – Category /Inter database /Intra database

### How do they all work together?

- **First level** -If category is present, it will do an I/O request from the cell
- **Second level** –If inter database plan is present, it will govern I/O Distribution between databases.
- **Third level** –The remaining I/O which can be allocated after both the category allocation and inter database plan will be allocated as per the intra database plan within the cell



## Instance Caging

- Instance caging is a technique with which you can run multiple Oracle instances on a multi-CPU Oracle server and map instances with a desired number CPUs. This is actually an 11g R2 Feature but became more popular with Exadata.
- This is an important feature as Exadata promotes consolidation of oracle servers. The easiest way to limit CPU consumption of a database instance is by implementing Instance caging.  
CPU\_COUNT – An instance initialization parameter which defines how many CPUs that particular instance can use.
- If you are running production and non-production databases on the same server, this feature can safeguard your production databases by limiting the non-production to utilize no more CPU's than allocated.

## Instance Caging

### Combination of Instance caging with IORM and DBRM

If you run into a situation where you cannot simply distribute CPUs between instances then you can use Resource Manager with instance caging. Resource manager allocates CPU among the various database sessions according to the resource plan that you set for the instance. Instance caging and the Resource Manager together are an effective way to manage multiple instances on a single server.

## Instance Caging

There are two ways you can implement instance caging.

**Partitioning:** You can distribute the number of CPUs the database server has to multiple instances. For example, you have an 8 CPU server and give 5 CPU to the production instance, 2 CPU to preprod and 1 CPU to development. With this configuration, any amount of load on preprod and development will never consume more than 3 CPU. The production database instance will still have 5 CPUs reserved for its use.


## Instance Caging

**Over-provisioning:** This method is used for non critical systems or low load production systems. In this method you allocate more CPUs to the instances than you actually have available on the server.

For example: You can allocate 3 to prod, 3 to dev, 3 to test and 2 to UAT for a total of 11 CPU's. In reality the server only physically has a total of 8 CPUs available. Using this method, instances can impact each others performance including your production instance. Utilizing instance caging can minimize the performance impact to production but not eliminate it under the over provisioning scenario.

### Enabling Caging

- Cpu\_count initialization parameter in init.ora
- ALTER SYSTEM SET CPU\_COUNT = 2;



## Exadata Smart Flash Cache

- Exadata Smart Flash Cache provides a caching mechanism for frequently-accessed data on each Exadata Cell.
- The Exadata Smart Flash Cache area on flash disks is automatically created on Exadata Cells during start up.
- Oracle Exadata Storage Servers are equipped with high performance flash disks in addition to traditional rotational hard disks.
- Most frequently-accessed data on the spinning disks are automatically cached in the Exadata Smart Flash Cache area on high-performance flash disks.
- Regular disk gives a few hundred IO/Sec but Flash can offer tens of thousands of I/Os per second.
- Exadata software transparently moves hot data to flash cache.



## Exadata Smart Flash Cache

- There are 4 - 24GB Flash memory modules per card and a total of 4 cards are installed per storage server.
- 6GB is reserved for failures with features like advanced wear leveling and bad block mapping
- PCIe card to avoid disk controller limitations
- Super capacitor backup battery
- Exadata full rack system has 14 storage servers. Total 5.3 TB
- 50GB/S throughput
- Can manage up to 1 Million IOPs

## Exadata Smart Flash Cache

- Cell disks are made up of flash cache and can be used for:
  - A) Smart flash cache which uses all space by default and is managed automatically
  - B) Grid disks made up of flash is an expensive option because you lose space for RAID and it requires planning & maintenance
- Exadata storage server software helps to make sure:
  - Only frequently accessed data is cached.
  - Skips data which is not frequently accessed
  - Does not cache very large files
  - DBA can fine tune caching
  - DB awareness algorithm promotes caching of data which is likely to be used again.
- If using hardware flash instead of Exadata you will not get the performance gains provided by Exadata and therefore not get the full benefit of the investment.

## Exadata Smart Flash Cache

- Flash cache is administered by a command line interface called CellCLI
- The “Create Cell” command creates flash cache by default
  - CellCLI > CREATE CELL CELL\_NAME
  - CellCLI > Drop flashcache
  - CellCLI > Create flashcache all
  - CellCLI > LIST CELLDISK DETAIL
  - CellCLI > CREATE GRIDDISK ALL FLASHDISK –  
PREFIX = 'FD', SIZE = 15G

## Exadata Smart Flash Cache

- You can set the following priorities  
Default/Keep/none for Table/Index/Lobs
- You can modify all of the above with an alter statement

Create table Tran\_record

```
-----  
-----
```

Storage (CELL\_FLASH\_CACHE KEEP)

Alter table Tran\_record Storage (CELL\_FLASH\_CACHE KEEP)

Alter table Tran\_record Storage (CELL\_FLASH\_CACHE Default)

Alter table Tran\_record Storage (CELL\_FLASH\_CACHE None)

## Exadata Smart Flash Cache

### How it works

- 11gR2 optimizer is not aware of flash cache or Exadata storage enhancements
- A SQL Statement is parsed and then sent to the storage server with all of the referenced objects flash prioritization
- If the storage server performs a smart scan then that implies that the query is doing a full table scan.
- If the query is a smart scan candidate, then blocks will not go to the flash cache
- If it is not a smart scan candidate then data will be checked in the flash cache.
- Some queries get data from the flash cache, as well as, from disk.
- Flash cache has no impact on write operations

## Exadata Smart Flash Cache

### How it works

- Exadata software continuously checks if an object needs to be in cache or not
- If an object is not in flash cache and has the attribute “keep” it will be stored in flash cache
- If the attribute is “none” it will not be stored in flash. If attribute is “default” and output is small I/O, then it will be stored in flash cache
- For small I/O it increases the response time
- For large I/O it will not increase response time but it may increase bandwidth. If some part of I/O is eligible for Flash then CellSRV will increase the bandwidth of the overall operation.
- Log writes or control file writes will never be cached in flash cache.

## Exadata Smart Flash Cache

### Monitoring commands

LIST FLASHCACHE DETAILS

LIST METRICCURRENT WHERE OBJECTTYPE='FLASHCACHE'

LIST METRICDEFINITION Metric\_Name

LIST FLASHCACHECONTENT WHERE OBJECTNUMBER =  
32323 DETAIL

### Retrieving Flash cache information from DB or DBM

AWR

Grid Control

V\$SYSSTAT

V\$SQL

## Exadata Smart Flash Cache

### V\$SYSSTAT


<b>cell flash cache read hits</b>	The number of read requests that were a cache hit on Exadata flash cache.
<b>cell IO uncompressed bytes</b>	The total size of uncompressed data that is processed on the cell. For scan on hybrid-columnar-compressed tables, this statistic is the size of data after decompression.
<b>cell physical IO interconnect bytes returned by smart scan</b>	The number of bytes that are returned by the cell for Smart Scan only, and does not include bytes for other database I/O.
<b>cell physical IO bytes saved by storage index</b>	The number of bytes saved by storage index.



## Exadata Smart Flash Cache

### Flash Cache Validation

- The Exadata software stores “check bytes” for every 4KB of data stored on the flash cache. Upon read, these “check bytes” are validated and if validation fails then the software directs the query to read from disk.
- You can also run `LIST METRICCURRENT FC_IO_ERRS` from CellCLI to find out if there are any current errors in the flash cache



## Storage Indexes

- Storage indexes are a very powerful feature of Oracle Exadata Storage Server Software. Storage indexes help to avoid I/O operations. Oracle Exadata Storage Server Software creates and maintains a storage index in Exadata memory. The storage index keeps track of minimum and maximum values of columns for tables stored on that Cell. This functionality is done transparently, and does not require any administration by the user.
- Data in Exadata storage is logically partitioned and each partition is 1MB. This 1 MB area is called a “Storage region”. Exadata Storage Software stores minimum and maximum values of column or columns in memory for every storage region.

## Storage Indexes

Table			Storage Index
Sr.#	Order#	Dt.	
	3011		Min=3011 Max=4119
	3013		
	4110		
	4119		
	4145		Min=4145 Max=5050
	4156		
	5050		

- Regular database Indexes give you a block address to a disk location. Where storage indexes do not have the block address, rather it locates the data by the process of eliminating storage regions.
- If the value in the where clause does not fall between Min-Max of storage region, that storage region will be eliminated

## Storage Indexes

- There will be separate Indexes for most of the columns. What column it will use to create the index is decided based on heuristics.
- As it eliminates storage regions this functionality is equivalent to partition pruning. Some queries get a very dramatic improvement while others may not be affected. Remember, they only work where there are full table scans occurring.

## Storage Indexes

We can evaluate I/O savings from v\$sysstat view and “Cell physical IO bytes saved by storage index” column.

<b>cell IO uncompressed bytes</b>	<b>The total size of uncompressed data that is processed on the cell. For scan on hybrid-columnar-compressed tables, this statistic is the size of data after decompression.</b>
<b>cell physical IO interconnect bytes returned by smart scan</b>	<b>The number of bytes that are returned by the cell for Smart Scan only, and does not include bytes for other database I/O.</b>
<b>cell physical IO bytes saved by storage index</b>	<b>The number of bytes saved by storage index.</b>

## Storage Indexes

### Storage Indexes – Clauses


- Works with smart scan queries –FTS.
- Works with >, >= , = , != , < , <= , Is Not Null
- Works with uncompressed table, EHCC and OLTP Compression
- NLS Columns and LOB data types will not be used in SI
- You cannot create, manage or update storage indexes because they are automatically managed by the Exadata Software.
- If you perform an insert or update on data which is compressed with EHCC or encrypted then the storage indexes will become invalidated. This is important to note because that would imply that using EHCC on an OLTP system is not necessarily advisable without consideration.

## Storage Indexes

### Storage Indexes – Clauses

- SI is lost if the storage cell is rebooted.
- After a reboot, if a smart scan query touches a storage region it will initiate SI creation automatically.
- Storage index with partitions: if you query a non partitioned column but data is sorted and both columns are similar, your query will not get the benefits of partition pruning but storage indexes will still likely improve performance.
- Loading data in sorted order provides better performance and takes full advantage of the storage index functionality.






## DBFS

- Oracle Database File System (DBFS) is an 11g database feature
- For users DBFS is a normal file system and can be mounted as NFS
- Runs on ASM
- Does not need to be a member of the cluster.
- DBFS uses FuSE (File system in UserSpaceE)
- FUSE is API and linux kernel module used to implement Linux file system.
- DBFS is the file system interface for SecureFiles
- Files are stored as SecureFiles LOBs
- Metadata is stored in database tables
- File Read/write/open/list are implemented using pl/sql
- You can mount file system by using DBFS and FUSE

## DBFS - Configuration

- Use of dedicated database for DBFS is recommended.
  - Use of ASMM is recommended .
  - 8GB SGA.
  - 2GB Redo logs.
  - Bigfile tablespace to hold DBFS
  - Create database user for DBFS with dbfs\_role,create session/table
  - Quota unlimited on bigfile tablespace
  - dbfs\_create\_filesystem\_advanced.sql in \$ORACLE\_HOME /rdbms/admin creates files system
  - \$ORACLE\_HOME/bin/dbfs\_client is used to mount this file system.
  - Below MOS Notes has all you need to create DBFS
1. **1191144.1** - Configuring a Database for DBFS on Oracle Database Machine
  2. **1054431.1** - Configuring DBFS on Oracle Database Machine



# Q & A

Anuj.Mohan@dbadirect.com