

4

Oracle Exadata Storage Server Migration Overview

ORACLE®

Copyright © 2009, Oracle. All rights reserved.


Objectives

After completing this lesson, you should be able to:

- Describe the general steps to migrate your database to Exadata storage
- List the main approaches for migrating your database to Exadata storage
- Identify the most appropriate approach for different circumstances

ORACLE

Migration Best Practices Overview



Legacy storage to Oracle Exadata Storage Server




ORACLE

Migration Best Practices Overview

The slide depicts the four important phases to migrate your database to Exadata.

Performing Capacity Planning

1



4 - 4

Copyright © 2009, Oracle. All rights reserved.

ORACLE

Performing Capacity Planning

Before migrating to Oracle Exadata Storage Server, you must understand the I/O characteristics of your current environment.

Collect the size and throughput of your current system. The key measures of I/O throughput are I/Os per second (IOPS) and megabytes per second (MBPS). Use the system statistic physical IO disk bytes to derive the current MBPS for the system. Use the system statistics physical reads and physical writes to determine the current IOPS for the system. These statistics are available in an Automatic Workload Repository (AWR) report.

After you understand the capacity of your current system, you can configure the correct number of Exadata cells. It is important to size both for performance and capacity. One Sun Oracle Exadata Storage Server configured with 12 SAS drives provides approximately 2 TB of mirrored uncompressed usable storage, 1,500 MBPS throughput for large I/Os, and 3,600 IOPS for small I/Os, however much greater effective throughput can be achieved by making efficient use of Exadata Smart Flash Cache and Exadata Hybrid Columnar Compression. Verify that the cell provides the expected throughput by running the CALIBRATE command.

It is also important to consider failures when planning capacity. Exadata cell and disk failures are transparently tolerated using Automatic Storage Management (ASM) redundancy. However, it is a best practice to ensure that post-failure I/O capacity is sufficient to meet the redundancy requirements and performance service levels.

ASM and Database Configuration

- ASM software must run release 11.2.0.1.
- ASM disk group attributes:
 - COMPATIBLE.ASM=11.2.0.0.0
 - COMPATIBLE.RDBMS=11.2.0.0.0
 - CELL.SMART_SCAN_CAPABLE=TRUE
 - AU_SIZE=4M
- Database initialization parameters:
 - COMPATIBLE=11.2.0.0.0
 - DB_BLOCK_CHECKSUM=TYPICAL|FULL
- Database extent sizes should be a multiple of 4 MB.

ORACLE

4 - 5

Copyright © 2009, Oracle. All rights reserved.

ASM and Database Configuration


Here is a list of configuration requirements and recommendations for Exadata V2:

- The Oracle ASM software must be running release 11.2.0.1 or later.
- For the ASM disk group containing Exadata grid disks, specify the attribute settings shown in the slide. Although these are not mandatory, they are highly recommended in order to enable all of the Exadata features and to optimize performance.
- To enable all Exadata features in your database, your Oracle Database software must be release 11.2.0.1 or later, and you must set the COMPATIBLE initialization parameter as shown in the slide. Because Oracle Exadata Storage Server is HARD (Hardware Assisted Resilient Data) compliant, you should set DB_BLOCK_CHECKSUM to TYPICAL or FULL to ensure that database checksums are computed and stored in the blocks. Note that this initialization parameter is set to TYPICAL by default.
- Oracle Exadata Storage Server performs best when scanning at least 4 MB contiguous chunks. To ensure that this occurs within ASM, the allocation unit size (AU_SIZE) should be set to 4 MB. To ensure the same at the database level, the database extent sizes should be a multiple of 4 MB.

Note: The parameters in red on the slide are required to enable smart storage capabilities.

Choosing the Right Migration Path

2



ORACLE

4 - 6

Copyright © 2009, Oracle. All rights reserved.

Choosing the Right Migration Path

Migration to Oracle Exadata Storage Server can be done logically or physically.

Using a logical approach, you can change the database extent size and alter other physical characteristics of your database which is not possible using a physical migration approach.

If the source database does not contain extent sizes that are a multiple of 4 MB, you should probably choose a suitable logical migration approach.

If the source database already contains extent sizes that are a multiple of 4 MB, you can make effective use of a physical migration approach.

The following slides briefly explain each of the options.

Logical Migration Approaches

Migration approach	Source database must be 11.2	Can adjust ASM AU size to 4MB	Supports heterogeneous database server platforms	Time and effort requirement	Migration outage time
Create new tablespaces <i>Best in conjunction with rolling partitions</i>	Yes	Yes	No	Lowest	Zero
Logical standby database and Data Pump <i>Minimal downtime</i>	No	Yes	No	Medium	Low
Streams and Data Pump <i>Minimal downtime</i>	No	Yes	Yes	High	Low
Oracle Data Pump <i>For migration during a planned maintenance period</i>	No	Yes	Yes	Low	High

ORACLE

4 - 7

Copyright © 2009, Oracle. All rights reserved.

Logical Migration Approaches

The above table summarizes the logical migration approaches:

- **Create new tablespaces on Exadata storage:** This is usable mainly for data warehouses using roll-in and roll-out partitions stored on different tablespaces.
- **Logical standby database and Data Pump:** If your application service-level agreements permit little or no downtime, then use an Oracle Data Guard logical standby database to move the data, and track and merge the changes while the applications continues to run. The procedure documented in My Oracle Support bulletin 737460.1 can be used to change physical storage attributes of the database, such as segment extent sizes.
- **Oracle Streams and Data Pump:** If your application service-level agreements permit little or no downtime, you can also use Oracle Streams with Oracle Data Pump to move the data, and to track and merge the changes while the applications continues to run. For more information about this approach, see Appendix D in *Oracle Streams Concepts and Administration 11g Release 2 (11.2)*.
- **Oracle Data Pump:** If your application service-level agreements permit enough downtime to move all of the data during a planned maintenance window, then use Oracle Data Pump to move the data in bulk from the legacy storage to Oracle Exadata Storage Server.

Physical Migration Approaches

Migration approach	Source database must be 11.2	Can adjust ASM AU size to 4MB	Supports heterogeneous database server platforms	Time and effort requirement	Migration outage time
ASM Online Migration <i>Best when already using ASM and 4 MB AU's</i>	Yes	No	No	Lowest	Zero
Physical standby database <i>Migration of DB server and storage</i>	Yes	Yes	Limited	Low	Low
Recovery Manager <i>Tablespace level migration</i>	Yes	Yes	No	Low	Medium
Transportable Database <i>Migration to a different platform with the same endian format</i>	Yes	Yes	Yes (same endian)	Low	Medium
Transportable Tablespace <i>Migration to a different platform with different endian format</i>	No	Yes	Yes	Medium	Medium

ORACLE

4 - 8

Copyright © 2009, Oracle. All rights reserved.

Physical Migration Approaches

The above table summarizes the physical migration approaches:


- **ASM Online Migration:** This method is only applicable if your database is already using ASM and you do not wish to adjust the ASM allocation unit (AU) size. Simply add new grid disks and drop existing ones in your ASM disk groups to migrate the storage.
- **Physical standby database:** Create a physical standby database on Exadata storage and perform a Data Guard switchover to migrate. See My Oracle Support bulletin 413484.1 for information about heterogeneous platform support.
- **Recovery Manager:** Create RMAN backups and use the RMAN SWITCH...TO COPY command to perform the migration. See Chapter 8 of the *Oracle Database Storage Administrator's Guide 11g Release 2 (11.2)* further information on using RMAN for ASM migrations.
- **Transportable Database:** Use the Transportable Database feature to migrate the entire database to a different platform with the same endian format. See the *Platform Migration using Transportable Database* white paper at <http://www.otn.oracle.com/goto/maa> for more information on Transportable Database.
- **Transportable Tablespaces:** Use the Transportable Tablespace feature to migrate tablespaces from your current system to a new database based on Exadata storage.

Creating New Tablespaces on Exadata Server

3

Before starting migration:

- COMPATIBLE must be set to 11.2.0.0.0



4 - 9

Copyright © 2009, Oracle. All rights reserved.

ORACLE

Creating New Tablespaces on Exadata Server

If your application inherently supports the logical roll-in of new data and age-out of old data at the table level, a strategy commonly used with the Oracle Partitioning option, then migrate to Oracle Exadata Storage Server by creating all your new partitions in a new tablespace based on Exadata storage. This is illustrated in the slide. As you create more new partitions and delete old one you slowly migrate across to Exadata storage. In this scenario, the migration process can take a long time during which your database is stored partly on original storage and partly on Exadata storage.


This approach has the following characteristics:

- No application downtime is required.
- Exadata is gradually phased-in to the environment so risk is minimised.
- The database must:
 - Run Oracle Database release 11.2.0.1 or later software.
 - Have the COMPATIBLE initialization parameter set to 11.2.0.0.0 before the migration process begins.

ASM Online Migration

Before starting migration:

- `COMPATIBLE` must be set to `11.2.0.0.0`
- `CELL.SMART_SCAN_CAPABLE` attribute must be `FALSE`



ORACLE

4 - 10

Copyright © 2009, Oracle. All rights reserved.

ASM Online Migration


Using ASM rebalancing, it is possible to migrate entire databases from their current storage platform to Exadata without suffering any down time.

Before migration, your database and ASM environments should be running version 11.2.0.1 or greater. The ASM disk group attribute `CELL.SMART_SCAN_CAPABLE` attribute must be set to `FALSE` for your existing disk groups. This is because during the migration process the disk group will contain original storage and Exadata storage for some period of time.

After you initially configure your Exadata cells, you can add some Exadata grid disks to your existing ASM disk groups. Initially, the newly added Exadata grid disks will be empty as shown in the diagram on the slide.

ASM Online Migration

- After Exadata grid disks are added to your disk group ASM rebalances the data across your original storage and your new Exadata grid disks.




ORACLE

ASM Online Migration (continued)

The addition of the Exadata grid disks causes a rebalance operation to take place. After the rebalance operation, half the data will reside on the original storage and half the data will have moved to the Exadata grid disks.

ASM Online Migration

- The original disks can now be removed from your ASM disk group prompting the movement of all the data to the Exadata grid disks.



ORACLE

4 - 12


Copyright © 2009, Oracle. All rights reserved.

ASM Online Migration (continued)

To complete the migration, the original disks can now be removed from the ASM disk group. Removing the original storage prompts another rebalance operation to move the remaining data on to the new Exadata grid disks.

ASM Online Migration

- The original storage can now be physically removed.
- Now set `CELL.SMART_SCAN_CAPABLE` attribute to `TRUE` so smart scan is enabled.



ORACLE

4 - 13

Copyright © 2009, Oracle. All rights reserved.

ASM Online Migration (continued)

When all the data has been copied to the Exadata grid disks, the original disks are removed from the disk group and the original storage can be decommissioned and physically removed from the system. The entire process occurs while databases remain online. The rebalance operations can be configured so as to consume minimal system resources. On a system that is not being overly stressed, database end users will not be adversely impacted by the migration process. Alternatively, you can allocate a greater share of system resources to the rebalance operations. This will help to complete the migration more quickly; however, depending on the user workload, the more aggressive rebalance settings may cause users to experience degraded performance while the rebalance operation is being executed.

After the migration, the ASM disk group attribute `CELL.SMART_SCAN_CAPABLE` should be set to `TRUE` so that smart scan is enabled.

Note that while the preceding example describes the addition of the Exadata grid disks and the removal of the original disks as two separate steps, it is possible to add and drop disks using a single `ALTER DISKGROUP` statement. This approach is recommended in preference to separate operations wherever possible. Adding and dropping using the same command has the benefit of more efficient rebalancing.

Post-Migration Best Practices

4

- Check ASM disk groups for rebalance:
 - Manual rebalance may be needed.
 - Script available at My Oracle Support bulletin 367445.1
 - Enterprise Manager alert for disk group imbalance
- Assess index requirements:
 - With Exadata, full scans might be faster than using indexes.
 - Determine indexes that are not required and remove them.
- Configure I/O resource manager.

ORACLE

4 - 14

Copyright © 2009, Oracle. All rights reserved.

Post-Migration Best Practices

Perform the following tasks after completing the migration:

- One of ASM's functions is to ensure that data is evenly distributed across all disks in a disk group. This happens automatically, but occasionally a disk group may become imbalanced due to uncommon errors such as a failed rebalance. It is therefore an operational best practice to check disk groups on a regular basis and run a manual rebalance, if necessary. A script is available to check disk group balance in My Oracle Support bulletin 367445.1. Also, Enterprise Manager Grid Control displays an alert if a disk group becomes unbalanced beyond a customizable threshold.
- The superior scan rates available from Exadata make it possible that indexes, previously required for good performance, are no longer required. You should assess application execution plans that use indexes to see if they would run faster with smart scans. To determine if queries would be faster without an index, you can make the index invisible to the optimizer. An invisible index is maintained by DML operations, but it is not be used by the optimizer for queries. To make an index invisible, use the following command:

```
ALTER INDEX <index_name> INVISIBLE;
```
- After you perform the above tasks, you can set up and configure I/O resource manager.

Quiz

When migrating your storage to Exadata, you cannot have a mix of grid disks and nongrid disks in the same ASM disk group.

1. TRUE
2. FALSE

ORACLE

4 - 15

Copyright © 2009, Oracle. All rights reserved.

Answers: 2

Quiz

What are two important best practices when migrating to Exadata storage?

1. Double the size of your SGA.
2. Configure ASM allocation units to 4 MB.
3. Double the size of your disk space.
4. Configure database extents to multiples of 4 MB.

ORACLE

4 - 16

Copyright © 2009, Oracle. All rights reserved.

Answers: 2, 4

Summary

In this lesson, you should have learned how to:

- Describe the general steps to migrate your database to Exadata storage
- List the main approaches for migrating your database to Exadata storage
- Identify the most appropriate approach for different circumstances

ORACLE®

Demonstration Overview

The demonstrations associated with this lesson cover the following topics:

- A
- B
- C

ORACLE[®]