

Exadata Implementation Strategy

BY UMAIR MANSOOB

Who Am I

- Work as Senior Principle Engineer for an Oracle Partner
- Oracle Certified Administrator from Oracle 7 – 12c
- Exadata Certified Implementation Specialist since 2011
- Oracle Database Performance Tuning Certified Expert
- Oracle Business Intelligence Foundation Suite 11g Certified Implementation Specialist
- Oracle Database Data Warehousing Certified Implementation Specialist
- Multiple Exadata Implementations / POC's for large financial organizations
- Migrate / Upgrade databases between various versions of Oracle
- Database Consolidation to Exadata / 12c Platform
- Architect Databases for OLTP and OLAP applications
- **Not an Oracle Employee or Nor I represent Oracle in any way**

Migration Vs Implementation

- Are you only looking to migrate database to Exadata Machine ?
- There is difference between Database migration and Exadata Implementation
- Lift and shift databases to Exadata Machine without making any significant change to databases
- Sometime you don't have a choice , if you are using 3rd party vender application

Why do you need a Strategy

- Understand that you need to use Exadata native features to achieve extreme performance
- You might not be able to take full advantages to Exadata machine if you perform only Data migration
- Even it's a third party application like SAP , you should have strategy in place to complete a successful migration
- Mostly like to meet deadlines and avoid any loss of revenue because of delays

Overview




Plan



Migrate



Optimize

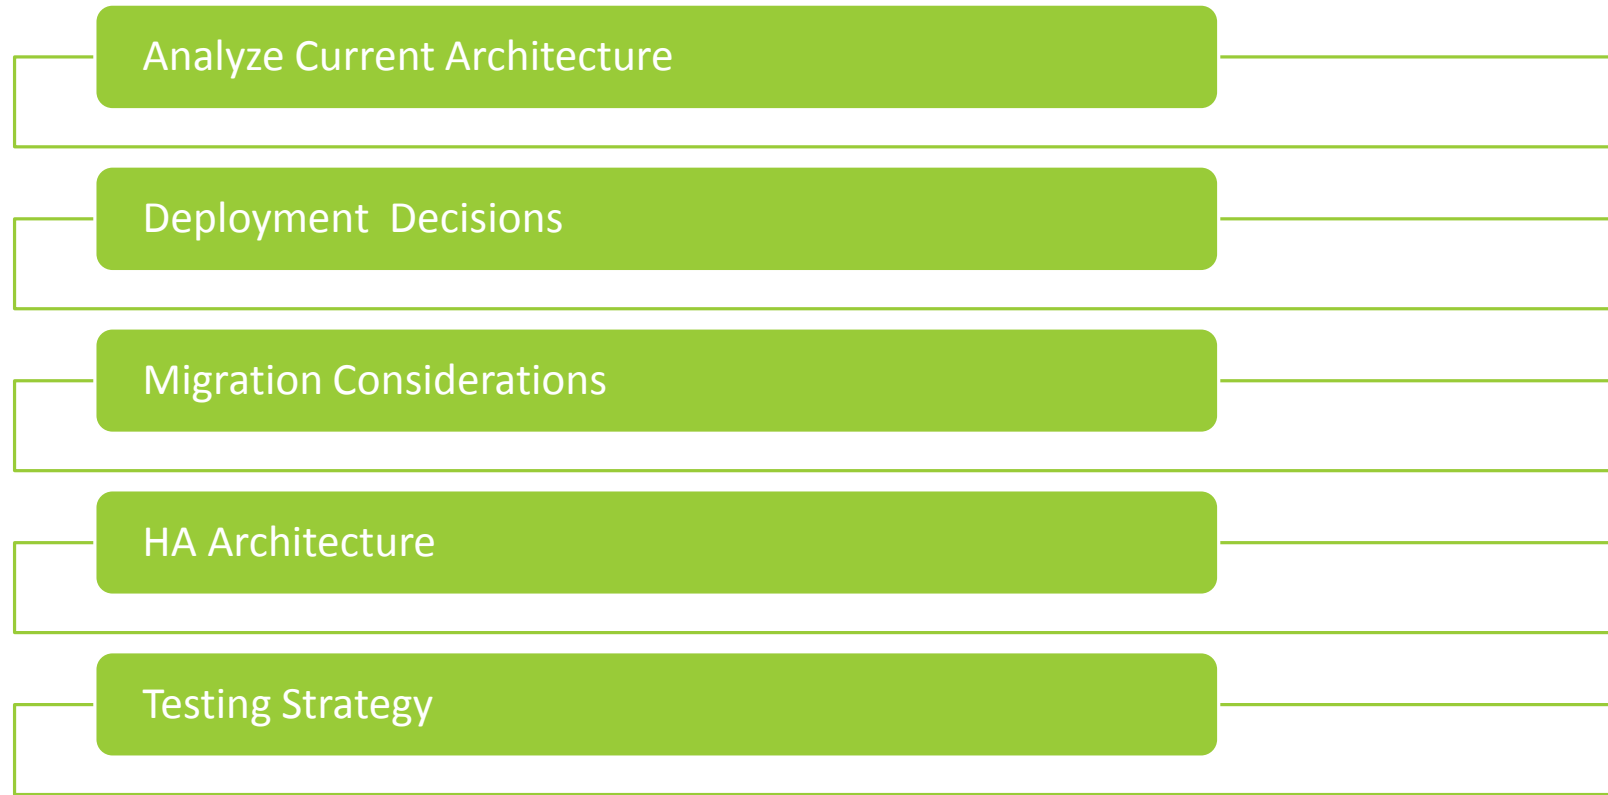


Test



Cutover

Plan



Analyze Current Architecture

- Analyze source database characteristic like database size and parameter settings, especially hidden parameters
- Analyze source host system characteristics like memory foot and IO through put and resource utilization.
- Is there a compliance requirement to encrypt or isolate database?
- Gather baseline from existing application system to compare result later.
- Gather AWR, ASH and other statistics from current environment

Deployment Decisions

- if you start your planning phase too late in the implementation, you might not be able to change any of the following
- Do you need to virtualize Exadata machine to isolate target database.
- What type of redundancy level do you need (High or Normal), External is not supported, depends on your application critically and DR
- Do you need to implement resource management using DBRM or IORM
- Do you need to isolate network using InfiniBand partitioning or Vlan tagging for compliance reasons.

Migration Considerations

- How big is your migration window ?
- Will you be needing Golden Gate to replicate database for **0 downtime** ?
- Do you have access to Database backups from Exadata Machine ?
- How is network bandwidth between source and target Exadata System, 1GB or 10G it matters
- Inform business user and other stake holders about upcoming migration, so they can plan for outage
- What is Recovery Time Objective (RTO)

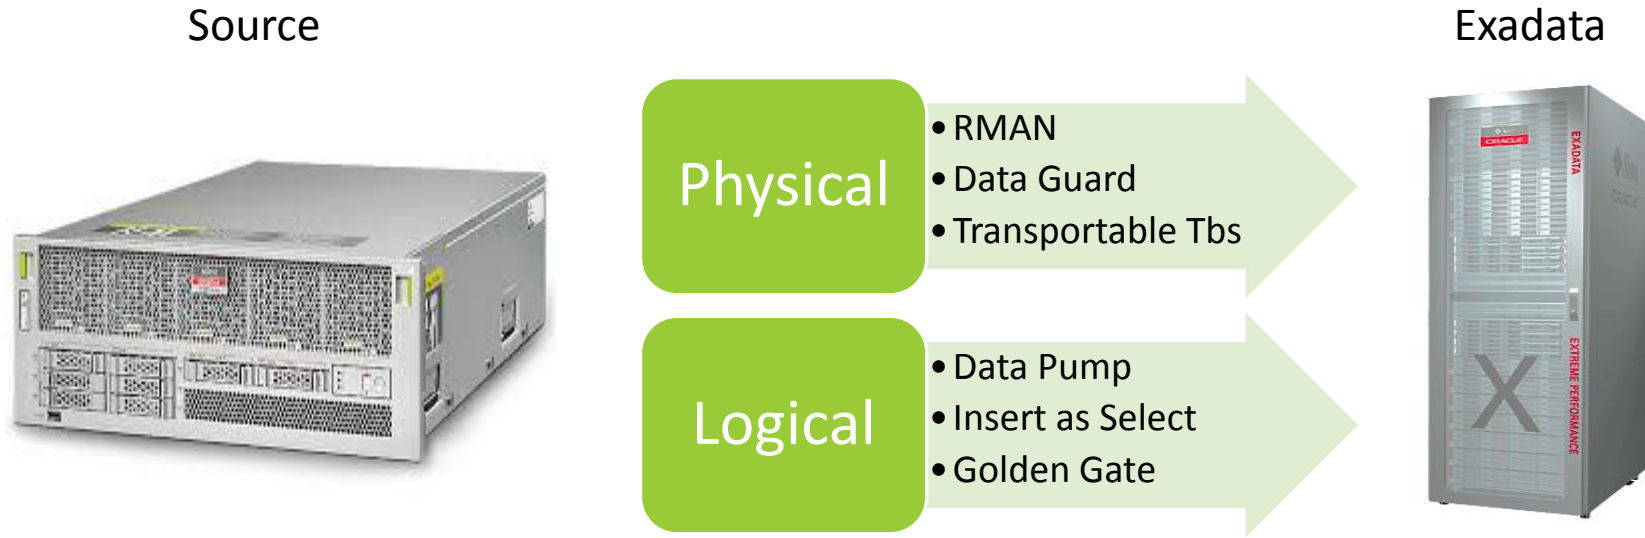
High Availability Architecture

- What will be your backup strategy , Tape or Exadata Storage , ZFS
- Depending on the database size , backing it up on Exadata storage can be very expensive
- Is Data Guard going to be part of your DR strategy ?
- Can you offload backups or reporting services to DR site ?
- Are migrating from Single instance to RAC database ?
- How are you planning to use cluster services to manage your work load
- Are you planning to enable flashback database, it can impact application performance if you are not already using it

Testing Strategy

- Moving to Exadata is a significant architectural change for you application, (OS , Storage , RAC)
- How are you planning to test your application on Exadata ?
- Do you need to develop additional test plans ?
- Will you be introducing code or architecture changes during the migration
- How are you planning to capture statistic during testing ?
- Most of application see immediate performance improvement but

Migrate



Migrate

- Each migration method has its own pros and cons, so analyze them carefully based on your requirements
- Data migration methods can be categorized as physical migration or logical migration.
- Physical migration is a block by block copy of data, popular methods are data guard and RMAN.
- Popular logical methods are data pump or insert as select.
- Pick migration method based to migration window and network bandwidth
- Network bandwidth can major road block , figure this out early

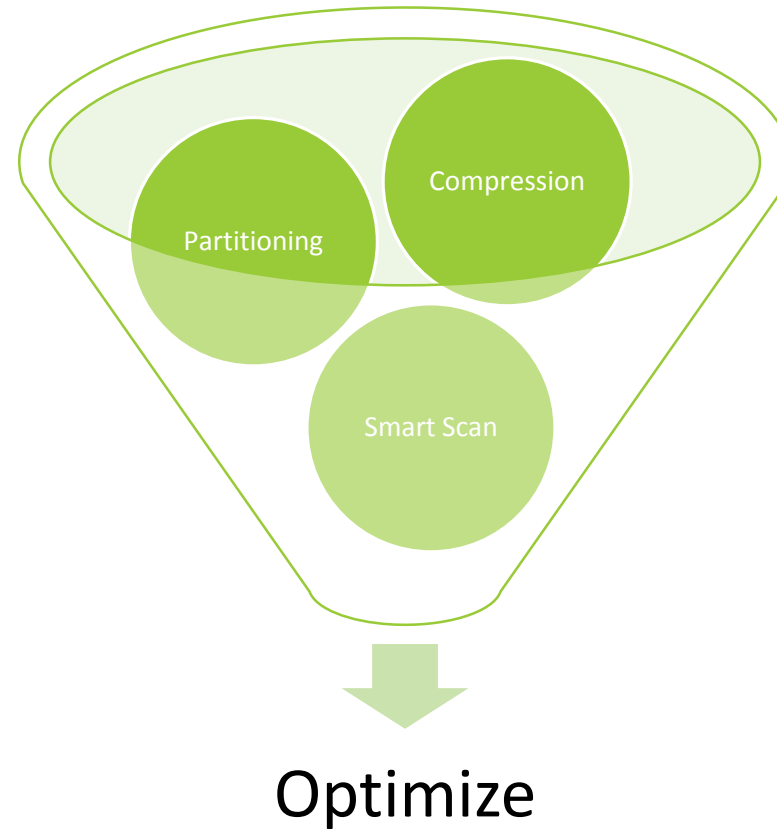
Physical Migration

- Pro : Migrating data using physical methods can be very simple and straight forward.
- Pro : Support cross platform Migrations using RMAN transportable database.
- Con : Bring all the characteristics of target database, which might not be optimal like hidden parameters
- Con : Will require extra effort to implement Exadata best practices and enable features like compression.

Logical Migration

- Pro : Exadata best practices will already be in place , if database was created using **DBCA utility**.
- Pro : Can optimize your workload for Exadata Machine using features like compression and partitioning during migration
- Con : Not simple , will require extra effort and time to bring data over to Exadata machines, **DB Size & Invalid Objects**
- Con : Since its not a block by block copy , **execution plans** can change.

Optimize



Optimize

- You should look into testing following Exadata Features during your implementation or even after the successful migration
- Compression not only reduces your storage footprint but also improve performance.
- Even though offloading and smart scan are enabled by default but make sure they are turned on and performing as expected.
- Test Storage indexes by making but selected non primary key indexes invisible.

Compression

- Regardless of Exadata, Oracle has two native compression types, basic table compression and OLTP Compression
- You can get reasonable compression ratio with OLTP compression and it will also support DML operations.
- Please note that there will be some overhead and you will need advance compression license to use OLTP compression.
- You can get extremely good compression ratio with Hybrid Columnar compression but OLTP operations are not supported
- OLTP Compression Vs HCC

Smart Flash Cache

- Exadata smart flash cache has an ability to move data in and out from cache based on usage.
- It is enabled by default, you don't have to configure anything to enable it.
- If enable , Write back flash cache provides the ability to write I/Os directly to PCI flash
- Write-back cache can help reduce ""free buffer waits" waits for write intensive applications.

Offloading / Smart Scan

- Exadata extreme performance is archive through offloading and smart scan.
- Offloading means some of Oracle processes are offloaded to Exadata Storage node.
- Some of the Oracle processes that can be offloaded to storage nodes are incremental backups, Data File creation, decompression and decryption
- There are some pre-requisites for smart scan like direct path read and full table scan.
- Verify Offloading using OEM or v\$SQL

Storage Indexes

- Storage indexes are memory structure at storage level.
- Storage Indexes reduce Disk I/O by maintaining an entry for minimum and maximum value for data per 1 MB by default.
- You can drop most non-primary key indexes and force queries to use Storage indexes.
- Usually BITMAP indexes are ideal candidates for offloading queries to Storage Indexes.

Do you need Indexes on Exadata?

- Following balance approach when it comes to use of indexes on Exadata.
- Don't drop all the indexes
- Keep primary key / unique indexes
- You can drop bit map indexes
- Use invisible index options when possible
- Avoid indexes using SQL HINTS
- Drop and rebuild indexes during ETL load

Partitioning

- Partitioning pruning is the simplest and also the most substantial means to improve performance.
- Partitioning can also improve the performance of multi-table joins, by using a technique known as partition-wise joins.
- Partitioning Strategies Interval - Range , List , Hash , Composite, etc
- Partitioned database objects provide partition independence , an important part of a high-availability strategy.


Parallelism

- Utilized Exadata Machine power by executing your queries in parallel to speed up your work load
- You can enable parallel query execution at object level or you can use SQL hint
- AUTO DOP will automatically parallelize your queries based on a threshold, I like to do it manually
- The default of `parallel_min_time_threshold` parameter is 10 seconds

Resource Management

- Database consolidation might require some level of resource management
- Resource mgmt will help you get consistent performance across all workload and databases
- You can use Oracle native resource management tool called DBRM to manage CPU utilization, parallel queueing and long running queries
- You can use IORM Exadata native utility to manage I/O throughput and latency.

Test



Testing Considerations

- Remember you are moving your application to new hardware
- You might be changing underline hardware system for your database
- You might be moving from NON-RAC to RAC database system
- Performance testing provides most important statistics for your Migration
- Some customer like to perform failover and break test for their critical applications

Start Testing


- You can a test like for like, which means you are not planning to make any changes to your database architecture like compression, partition or encryption
- You can use your existing test plans for Exadata Migration testing.
- Testing should include should include all the databases you have introduced during the migration.
- There are different types of testing you can perform during Exadata Migration, like Performance test, load test, break test.
- Don't forget to capture statistics during the testing phase of Exadata Migration.

Capture Statistics

- Make sure to capture performance stats using tools like AWR, ASH, and SQL Performance analyzer
- AWR reports will provide you all the details you need to compare elapse time, IO wait and CPU utilization
- Also validate Exadata configuration through running Exachk
- ASH reports can provide you further details about execution plans and wait times
- Analyze Results and move forward cutover phase
- Remediate any performance issues discover during this phase , before cutover

Cutover

Source



Exadata



Cutover

- Make sure to backup both source and target databases
- Have a fallback plan, just in case if you encounter any issues after the cutover
- You will probably have to sync your target database just before the cutover ,if you are not using GoldenGate
- Based on the migration method, Database SCN number can play a key role during this phase
- You should be on guard for next 48 hours and ready to remediate any issues.

Thank You

Umair Mansoob

773-297-2061

umairmansoob@gmail.com

<http://blog.umairmansoob.com/>

