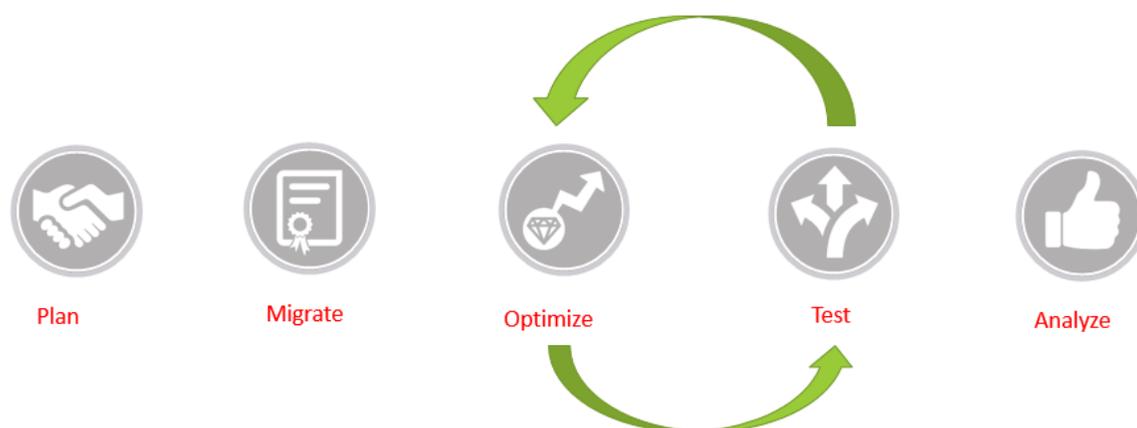


Introduction

Exadata Machine was first introduced by Oracle in 2008 and now it has become one of the most popular database platforms to host Oracle databases. Exadata machine is like a mini data center in itself, comprising of Database Servers, Storage Servers, InfiniBand Switches and Ethernet switches working together to deliver outstanding database performance. Exadata is a big investment for customers, realizing that Oracle sometimes offers on-site/off-site loaner Exadata machine for the proof of concept purposes. One should understand that there are only few week window period to complete POC and if there is no proper plan for POC, one may end up losing some valuable time.

Overview



The purpose of this article is to help readers to do a successful POC and get maximum performance results out of Exadata machine. Let's start with a brief description of Exadata POC process and later we will discuss them in detail. Basically each Exadata POC should start with detail planning and then move into migrating data using a suitable migration methods. Once data migration is completed, start testing different Exadata features like Compression, Storage indexes, Offloading and Smart Scan. It is important to capture all the statistics during the whole POC process, so that it can be analyzed and compared later on.

Phase 1: Plan

Start planning early. This is very important. In most cases you will only get 2-3 weeks to do an Exadata POC. If you are not ready, you will end up losing very valuable time. I have seen many Exadata POCs, where a customer ran out of time and was not able to test all the valuable Exadata features. Testing Exadata features is important because even as a customer you are asking your business users to commit a larger sum of money for an Exadata machine, and it be in everyone's best interest to achieve the optimum result from POC. It is important to note that dedicated resources also play a key role in a successful Exadata POC. Also have your migration, testing and monitoring scripts ready. If possible, test

Published in RMOUG SQL > UPDATE Magazine

your scripts on NON-Exadata system before the POC. It will save you some time that you might better spend on troubleshooting those scripts during POC.

During this phase, you should identify all the processes and sql statements, which will be part of your POC. How you go about picking targets for your POC will depend on your target application. For example, your application can be OLTP, running a lot of DML statements randomly. In this case DML statements are key to your Exadata POC. Similarly your application can be a data warehouse, running a lot of complex select statements and overnight batch processes. In this case select statements and batch process are good targets for your POC. It is important to understand that you might not be able to test all statements and processes running on target system during your POC, so it is best to pick your targets based on criticality, elapse time and resource utilization.

It is important to capture all the statistics during POC, you will need them later to analyze and compare POC results. You can use many native Oracle tools like AWR Report, ASH report and SQL Tuning set. The AWR can be used to capture statistics on instance level and can provide you a list of top queries by elapsed time, by CPU utilization and by IO waits. The ASH report can be used to capture similar statistics for any particular query or a session. Oracle SQL Tuning Set can be used to collect all SQL run during particular time window or from a particular session. SQL Performance Analyzer can be used with SQL tuning set to analyze and compare results between non-Exadata and Exadata systems. You have to start with capturing a baseline from target system. You can capture a baseline from the current production system but only if you can simulate the same production load on the Exadata machine. From my personal experience, a lot of my clients were not able to produce a production load during their POC's. So I recommend using a test or UAT environment with the same hardware profile to capture your baseline. Also understand that you might be only testing part of your production load during your POC. For example your production system might have 100 batch processes running, but you should only pick the top critical process for your POC. You can pick all 100 batch process for your POC, but then you will need more time and resources to analyze the result of your POC.

Plan



After planning, you will need to migrate data to the Exadata machine for your POC. It will be best if you choose same migration method that you are planning to use for the actual migration. I have seen customers run into issues during the data migration process and end up wasting valuable time during the POC. So make sure you have a migration plan and migration scripts ready before the start of POC.

Similarly, you should also have clearly define success criteria before the start of your POC. Work with developers and business partners to come up with well-defined success criteria and share it between all the parties involving in POC including Oracle. Elapse time is very important success criteria but you can achieve better elapse time just changing few parameters inside the database, so it is important to pay attention to resource utilization in addition to elapse times.

Phase 2: Migration

There are several methods to migrate data to a POC Exadata machine. Data migration methods can be categorized as physical migration or logical migration. Each migration method has its own pros and cons, so analyze them carefully based on your requirements. If you want to keep the physical database structure the same, you should choose migration methods like Data Guard or RMAN. If you are want to migrate from Big Endian to Little Endian, you should choose Transportable Tablespaces or Data Pump. The migration methods you choose during this phase can be critical to your POC down the road. For example if you choose to migrate data using methods like RMAN backup and recovery or physical data guard, then you will be bringing all the physical characteristics of the target database, including database parameters which might not be optimal settings for databases running on an Exadata machine. Please note that an Exadata machine comes with its own set of best practices with regards to running

databases. In case you decide to migrate data using a physical migration method, you should run the exachk utility to review and implement Exadata best practices. Similarly, if you decide to use a physical migration method to migrate databases, you also have to go through additional steps to test Exadata features like EHCC compression, storage indexes and offloading.

Migrate

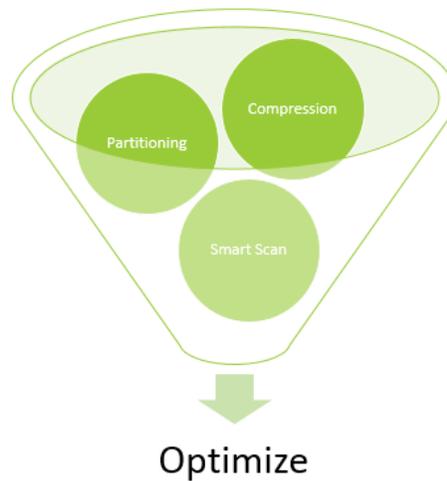


There are few things you need to keep in mind about moving data into an Exadata machine. Your POC Exadata machine might not be sitting in your own Data Center and if you are a bank or healthcare provider with customer data, moving data out of your data center and transferring it to a POC Exadata machine might not be straight forward. You might have to get proper approvals, which may take few days or weeks. That is why it is important to complete this task before you commit to any dates for an Exadata POC. It is also important to mask or encrypt data before you move into an Exadata POC machine. If the target database size is in terabytes, you should look into migrating data using a secure storage device. I would also like to point out that the dataset size you pick for your POC is very important, don't pick something too small that you are not able to run long queries during your POC and don't pick too big that it will take days to transfer data to the POC machine.

Phase 3: Optimize

If you really want to take full advantage of your Exadata POC, you should test Exadata / Database features like Compression, Partitioning and Auto DOP during your POC. For example: compression not only reduces your storage footprint but also improves performance. Parallel execution will also help you with performance and the same goes for properly caching tables to the Exadata Flash Cache. Even though offloading and smart scan are enabled by default, make sure they are turned on and performing as expected. Some of these features are Oracle database native features, not specific to Exadata but you might not be using it in your existing database. Here is a brief overview of some of the features you should test during your POC:

Optimize



Compression: Regardless of Exadata, Oracle has two native compression types, Basic Table Compression and OLTP Compression. You will not get a good compression ratio with Basic table compression and it will not support DML operations. 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 an advance compression license to use OLTP compression. Exadata also comes with its own compression called Hybrid Columnar Compression. You can get an extremely good compression ratio with Hybrid Columnar compression but OLTP operations are not supported.

Partitioning: Oracle supports many types of partitioning techniques including range, list, composite and hash. Partitioning can provide you many benefits like better performance, high availability and ease of maintenance. Partitioning can also help you achieve better performance through partition pruning. You can also perform certain maintenance tasks like truncate a partition. Partitioning can also improve the performance of multi-table joins, by using a technique known as partition-wise joins. Gather stats and rebuilding a local index just for a particular partition, this will hence provide you ease of maintenance and high availability.

Auto DOP: You can execute your queries in parallel to speed up your load. If you are not already using the parallel query feature, you should look into enabling this feature. You can enable parallel query execution at the object level or you can use a SQL hint. You can also let Oracle determine the degree of parallelism based on a set of criteria and some initialization parameter settings. This feature is called AUTO DOP and it will automatically parallelize your queries based on a threshold. The threshold that is prominently mentioned above is set by parallel min time threshold. The default of this parameter is 10 seconds. If you want to run more statements in parallel, make sure to reduce that number so that more plans qualify for parallel evaluation.

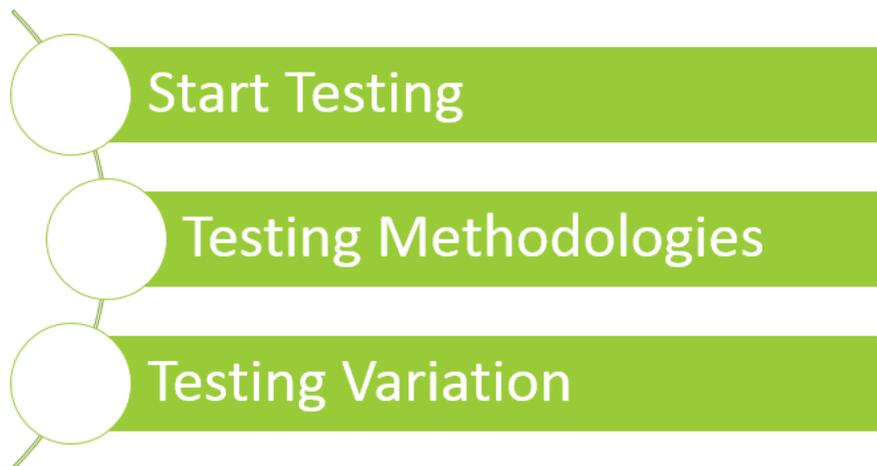
Storage Indexes: Storage indexes are memory structures at the storage level. They reduce Disk I/O by maintaining an entry for minimum and maximum values for data per 1 MB by default. It's kind of like a database indexes, that is why they are called Storage Indexes. There are many ways you can force Oracle

queries to use storage indexes and in most cases they will be faster than using a traditional Oracle index. You can drop most indexes or make them invisible and force queries to use Storage indexes. Usually bitmap indexes are ideal candidates for offloading queries to Storage Indexes. Dropping indexes is a physical change to a database structure and many times I find customers hesitant to test this feature during a POC. Keep in mind, you might need primary or unique key indexes for data integrity purposes.

Phase 4: Test

Once you understand all the database and Exadata features available to you, it is the time to start testing. Testing should include all the variations you want to test during your POC. For example you can test like for like, which means you are not planning to make any changes to database architecture like compression, partition or encryption. Or you can test some or all the database / Exadata features during your POC. Depending how much time you have for the POC, you should plan your testing carefully. Some of the features I have mentioned earlier can drastically change the performance of your queries or batch processes, so I would highly recommend testing these features.

Test

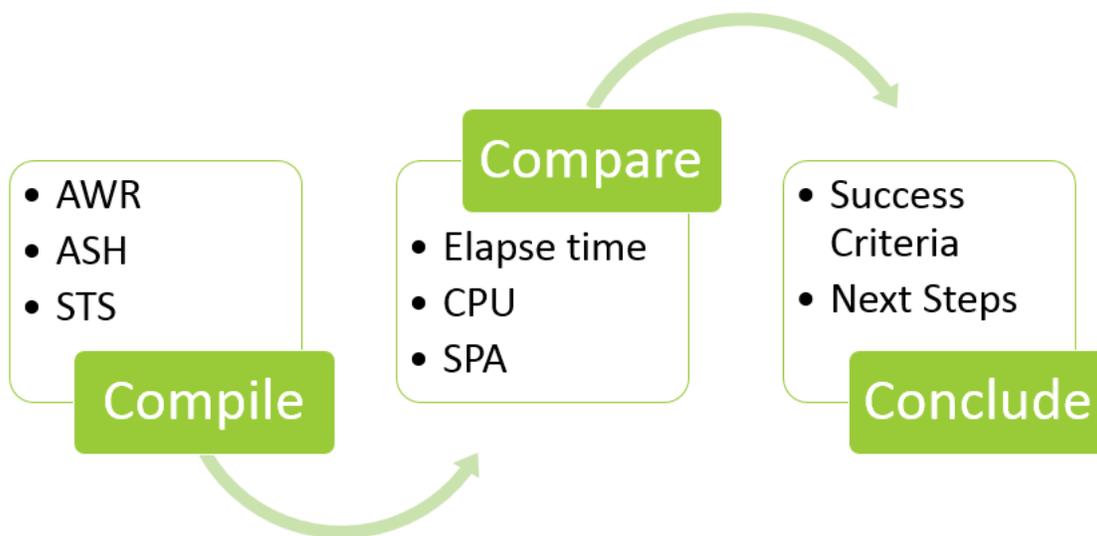


You might already be using one of the above database features like compression, partitioning and encryption. IF you are not, test them individually or in pairs to get the most out of your Exadata POC. During the POC, you might find some features that can improve your application performance many fold, and some can cause your processes to run slow. You will only know if you test them thoroughly.

Phase 5: Analyze

This is the final phase of your POC, it is time to compile your test results and analyze them for conclusion. You should start by comparing AWR reports between your baselines to Exadata POC. AWR reports will provide you all the details you need to compare elapsed time, IO waits and CPU utilization. Also compare critical processes and queries using ASH reports, it will provide you further details about execution plans and wait times. This phase should be very straight forward if you have clearly defined success criteria from the planning phase of this POC. It is possible that not all parts of your POC are successful, for example, you might not get the expected performance after implementing compression or write-back cache. It is also possible that 90% of your SQL queries or processes run better than expected and you are happy with the results. But of the other 10% of queries might not be as per your expectations.

Analyze



This is not unusual. Sometimes you don't get optimal execution plans after migrating databases to new hardware. You might have to spend some more time analyzing those queries or processes and tune them to the point where you are happy with their performances.

Conclusion

With recent Exadata capacity on-demand offerings, Exadata has become an affordable option for many small to mid-size organizations. Still, it's a considerable investment for many customers and they would really like to use a POC to make a good business case for buying an Exadata machine to run their Oracle databases. If you are one of those customers who get approved for Exadata POC's, you should treat your POC like a mini project and make sure to follow a comprehensive plan which can lead to a successful POC.