



Dell Compellent Storage Center

Oracle Best Practices

Best Practices Guide

Dell Compellent Technical Solutions Group
July 2013

THIS BEST PRACTICES GUIDE IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2013 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Dell, the DELL logo, and the DELL badge are trademarks of Dell Inc. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims any proprietary interest in the marks and names of others.

Table of Contents

1	Preface	1
1.1	Audience	1
1.2	Purpose.....	1
1.3	Customer Support.....	1
2	Introduction.....	2
3	Features of Dell Compellent Storage Center.....	3
3.1	Dynamic Block Architecture.....	3
3.2	Storage Virtualization.....	3
3.3	Dynamic Capacity (Thin Provisioning).....	4
3.4	Data Instant Replay™ (DIR).....	4
3.5	Data Progression™	5
3.6	Fast Track.....	5
3.7	Consistency Group.....	6
3.8	Benefits of Oracle on Compellent Storage Center.....	7
4	Storage Setup and Configuration.....	8
4.1	Disk Drives & RAID Recommendations	8
4.2	RAID levels and data Progression™ for Oracle databases.....	9
4.2.1	Solid State Disks (SSD)	9
4.2.2	SAS of Fibre Channel Disks	9
5	Recommended Oracle Disk Configuration.....	10
5.1	Oracle with Cooked File Systems.....	10
5.2	Oracle with SSD	11
5.3	Oracle with Automatic Storage Management (ASM)	12
5.4	ASM Disk Group.....	12
5.5	ASM Disk Group Devices.....	13
5.6	Filesystem.....	13
5.7	Database Setup and Configuration – Direct I/O and Async I/O	13
5.8	Putting it all Together	14
6	Using Compellent Data Progression™ and Data Instant Replay Features	16
6.1	SSD Configuration.....	16

6.2	Non-SSD Configuration	17
7	Oracle RAC Tested Configuration.....	19
8	Conclusion	20
9	Additional Resources.....	21
9.1	Oracle Performance Testing with ORION	21

Document Revisions

Date	Revision	Author	Comments
01/6/2012	1.0		Initial release
04/19/2012	1.1	Mark Tomczik	Content and format change
07/23/2013	1.2	Mark Tomczik	Added Orion information provided by Copilot, converted doc to new format

1 Preface

1.1 Audience

The audience for this document is intended for database administrators, system administrators and storage administrators that need to understand how to configure Oracle databases on Dell Compellent Storage Center. Readers should be familiar with Dell Compellent Storage Center and have prior experience in configuring and operating the following:

- Oracle Architecture
- Oracle 10g and 11g
- Real Application Clusters (RAC) or Single Instance
- Oracle's Automated Storage Management (ASM)
- General understanding of SAN technologies

1.2 Purpose

This document explains how to combine Oracle with Dell Compellent storage technology, and will describe the features and terminology from Dell Compellent and Oracle. It also gives an example of installing and configuring Oracle with Compellent Storage Center.

1.3 Customer Support

Dell Compellent provides live support 1-866-EZSTORE (866.397.8673), 24 hours a day, 7 days a week, 365 days a year. For additional support, email Dell Compellent at support@compellent.com. Dell Compellent responds to emails during normal business hours.

2 Introduction


When designing the physical layer of a database, DBAs must consider many storage configuration options. The storage solution must facilitate high performance I/O. It must protect against failure of storage hardware, such as disks, host bus adapters (HBAs), and fabric switches. The strains of a growing and ever changing workload imposed upon the storage solution require the SAN configuration to be dynamic. These same strains also create additional work for the storage administrators and require them to be much more responsive to the needs of the business. Having storage related tasks automated reduces the risk of human error, and allows DBAs to concentrate on other business critical database administration functions and services.

3 Features of Dell Compellent Storage Center

Dell Compellent provides a number of features that can be taken advantage of in an Oracle environment. Several features are discussed below.


3.1 Dynamic Block Architecture

Dynamic Block Architecture records and tracks specific information, called metadata, about every block of data, and provides the system intelligence on how every block is being used. Metadata about blocks is gathered without system overhead and can be extensive, including time written, the type of disk drive used, the type of data stored, RAID level, and more. All of this metadata or “data about the data” enables Storage Center to take a more sophisticated and intelligent approach to storing, recovering and managing data. For example, metadata that describes that specific blocks of data have been accessed the most can be used by Storage Center to move the blocks to a faster tier of storage. Metadata describing inactive blocks can be used by Storage Center to automatically place the blocks on lower cost, lower performance drives.



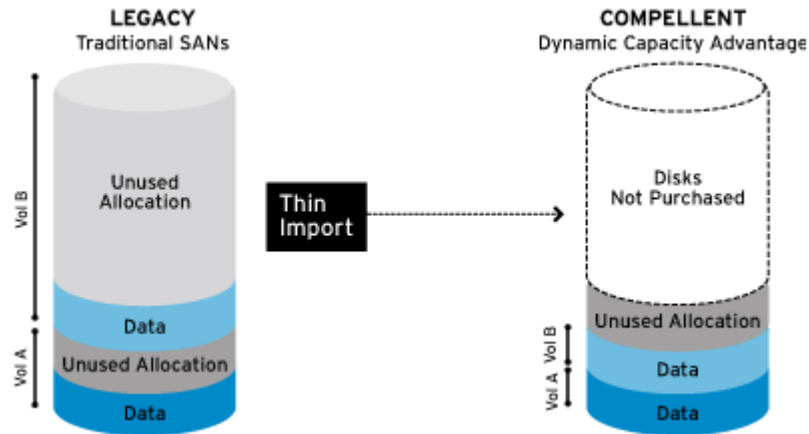
3.2 Storage Virtualization

Compellent Storage virtualizes enterprise storage at the disk level, creating a dynamic pool of storage resources shared by all servers. Because read/write operations are spread across all available drives within the same tier, multiple requests are processed in parallel, boosting system performance. With Compellent Storage Virtualization, users can create hundreds of volumes in seconds to support any server platform and optimize the placement of applications. Storage Center automatically restripes all data across all drives in the storage pool when adding disks, and can dynamically scale the storage pool without disruptions. Virtualization also allows Storage Center to change RAID levels on the fly.



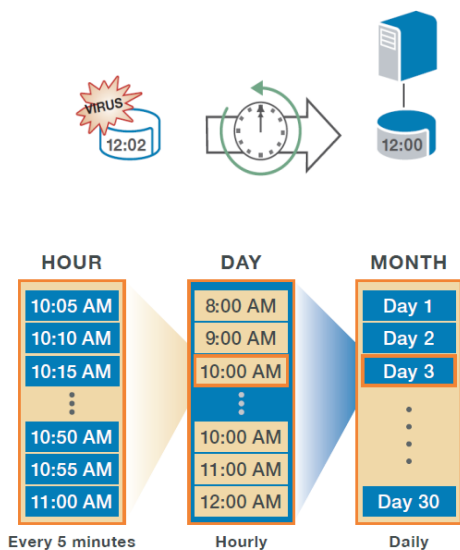
3.3 Dynamic Capacity (Thin Provisioning)

With traditional storage systems, Administrators must purchase, allocate and manage capacity upfront, speculating where to place storage resources and creating large, underutilized volumes with long term growth built in. This practice leaves the majority of disk space



allocated yet unused, and only available to specific applications. Dell Compellent's Thin Provisioning, called Dynamic Capacity, delivers the highest storage utilization possible by eliminating allocated but unused capacity. Dynamic Capacity completely separates storage allocation from utilization, enabling users to create any size virtual volume upfront, yet only consume actual physical capacity when data is written by the application.

3.4 Data Instant Replay™ (DIR)



Data Instant Replay™ is most often compared to "snapshot" technology and provides continuous space-efficient data protection using a feature called Replays. Replays create point-in-time copies of your volumes where further changes to a volume are journaled in a way that allows the volume to be rolled back to its original state when the replay was created. Replays can be mounted as volumes, called View Volumes, for the sake of partial or full volume data restore. There are various uses of Replays, but with respect to an Oracle database, database backups and cloning are easily performed. Policy based schedules, with varying intervals and expiration, can be created to manage Replays and provide greater recovery capabilities to a previous known state, and there is no limit on the number


or replays taken. Replays, when used in conjunction with another feature called Consistency Groups, provide data integrity of data spanning multiple volumes. Once a replay is created, an unlimited number of new volumes, called view volumes, can be created from the Replay.

For more information on view volumes, see section “Consistency Group”.

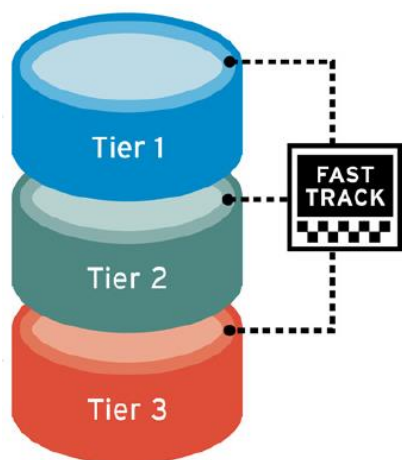
3.5 Data Progression™

Fluid data storage or automated tiered storage is another feature, called Data Progression™, offered by Dell Compellent. Data Progression™ automatically migrates enterprise data to the most optimal storage tier based on a set of predefined or custom policies or storage profiles. Over time,

infrequently accessed data migrates to lower-cost drives (SATA), while the most active, mission-critical data migrates to high-performance (expensive FC/SAS) drives. The number of days used to alert Data Progression™ to migrate the data can easily be changed. By default, data that has been accessed within the last four progression cycles is stored on the highest tier, composed of the fastest disks. Data that has not been accessed for the last 12 progression cycles is gradually migrated down to the lowest tier, composed of slower, cheaper, larger disks. Data progression™ eliminates the need to manually classify and migrate data to different storage tiers while reducing drive number and cost, and cooling and power costs.



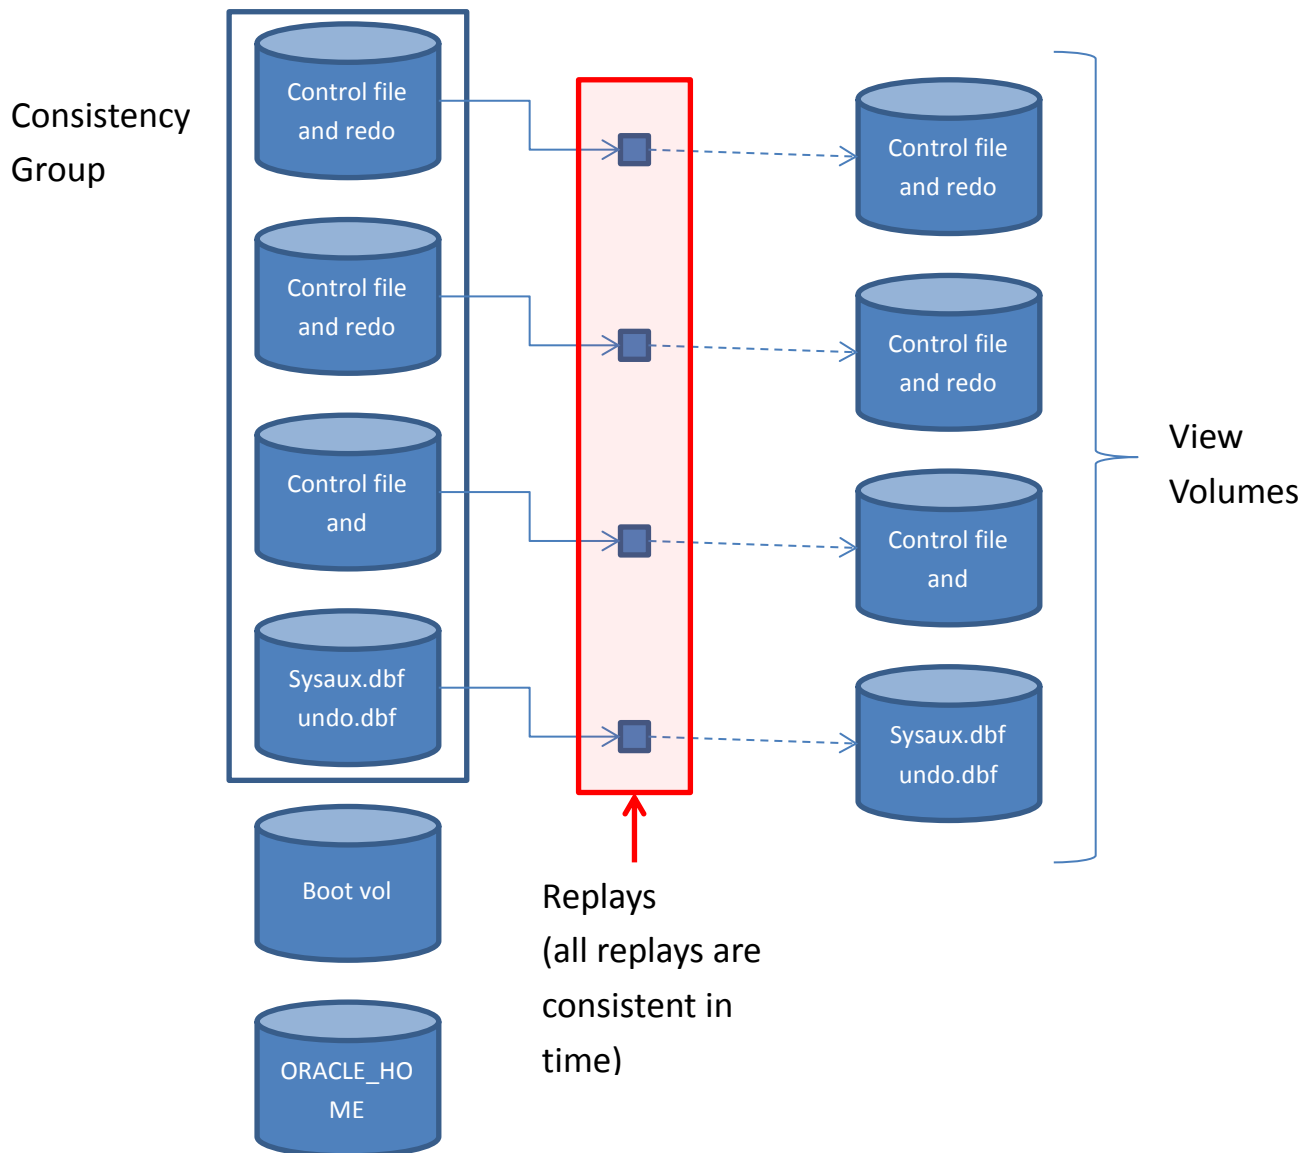
3.6 Fast Track



Dell Compellent Fast Track technology enhances Automated Tiered Storage by dynamically placing the most frequently accessed data on the fastest, or outer, tracks of each disk drive. Meanwhile, the least active blocks of data remain on the inner tracks of all the drives in each storage tier. The result is ready access to frequently used data with fewer drives than conventional storage systems.

3.7 Consistency Group

Dell's Compellent Consistency Group feature allows storage administrators to take a snapshot of an Oracle database atomically. When creating a snapshot of a running Oracle, make sure that all storage volumes (LUNs) that make up the database be atomically. Without a consistency group, a usable snapshot of a running database cannot be made. If Storage Center is on 4.x code, there is no consistency group available when using data Instant Replay™ for snapshot. In order to take a snapshot of a database while it is running, the Oracle database online redo log files, and control files must exist on the same volume. Multiplexing online redo logs and control files across volumes cannot be done. If the Storage Center is on 5.x code or higher, then there is no restriction. Also note that testing need to be done to determine the number of LUNs required for your database for optimal configuration in terms of performance since every operating system is different.



3.8 Benefits of Oracle on Compellent Storage Center

Some of the benefits of deploying Oracle in a Compellent SAN solution are listed in the below table.

Benefit	Details
Lower total cost of ownership (TCO)	Reduces acquisition, administration, and maintenance costs
Greater manageability	Ease of use, implementation, provisioning, and management
Simplified RAC Implementation	Provides shared storage (raw or filesystems)
High availability and scalability	Clustering provides higher levels of data availability and combined processing power of multiple server for greater throughput and scalability
Compellent Information Life Cycle (ILM) benefits	Provides tiered storage, dynamic capacity, data progression™, thin provisioning, instant replay (snapshot) and more

4 Storage Setup and Configuration

The below table is a guideline that can help you get started with placement of Oracle files on Compellent. One should understand that given information is only a recommendation, and actual placement of files will be dictated by your business needs.

4.1 Disk Drives & RAID Recommendations

Description	RAID10 SSD	RAID10/FC/SAS 15K rpm	RAID10/FC/SAS 10K rpm	RAID5/FC/SAS 15K rpm	RAID5/FC/SAS 10K rpm	All RAID FC/SAS/SATA
Data files	OK (W)	Recommended (W)	OK (W)	DP	DP	DP
Control files	Recommended (W)	OK (W)	OK (W)	Not Required	Not Required	Not Required
Online Redo Logs	Recommended (W)	OK (W)	OK (W)	Avoid	Avoid	Avoid
Archived Redo Logs	Not Required	Recommended (W)	OK (W)	Not Required	Not Required	DP
Flashback Recovery	Not Required	Recommended (W)	OK (W)	Not Required	Not Required	DP
OCR files / Voting Disk	Not Required	Recommended (W)	OK (W)	Avoid	Avoid	Avoid

Abbreviation:

W → Writes

DP → Data Progression™

- If Fast Track is licensed, then it is enabled by default and will be utilized behind the scenes. No manual configuration is required.
- Drives with higher RPM provide higher overall random-access throughput and shorter response times than drives with lower RPM.
- Because of better performance, SAS or Fibre Channel drives with 15K rpm are always recommended for storing Oracle datafiles and online redo logs.
- Serial ATA and lower cost Fibre Channel drives have slower rotational speed and therefore recommended for Oracle archived redo logs and flashback recovery area with Data Progression™.

4.2 RAID levels and data Progression™ for Oracle databases

4.2.1 Solid State Disks (SSD)

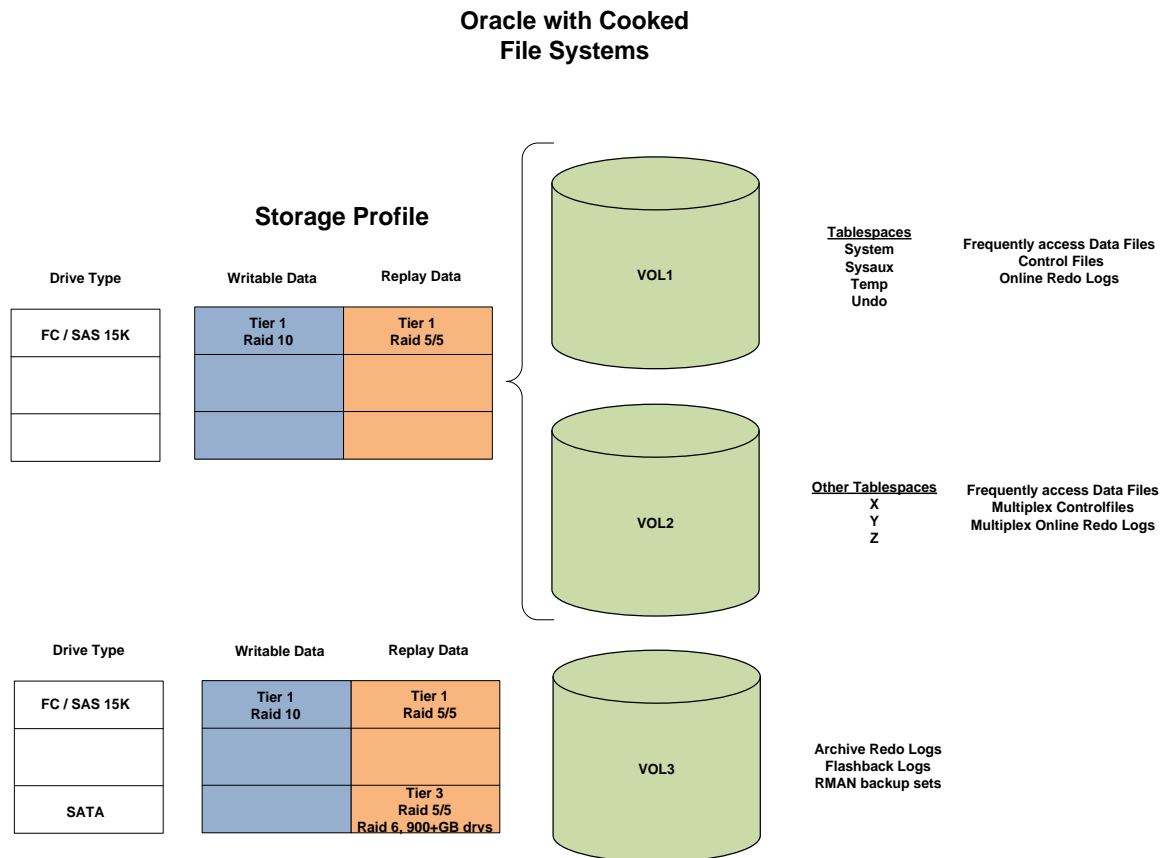
- Before implementing SSD on Compellent Storage Center for Oracle databases, determine if the database in question warrants the high performance that SSD provides. Since the price per gigabyte of SSD is more expensive than SAS or Fibre Channel, you need to carefully evaluate your database performance.
- Compellent recommends using SSD RAID10 for database online redo logs and frequently used indices if the database is highly transactional and is IO constrained.
- Data Warehouse databases should not use SSD for online redo logs as that will not have any performance gain unless the whole database resides on SSD which may be very costly depending on how large the database is.
- If using SSD for online redo logs, do not configure Data Progression™ on this volume as there will not be any gain or any space back due to the nature of online redo logs.

4.2.2 SAS of Fibre Channel Disks

- Disk types are determined by the performance requirement for a database. In a highly transactional database, Dell Compellent recommends using RAID10 with 15K rpm disks for datafile volumes, online redo log volumes (if SSD is not available), and archived log volumes for all production databases. If SSD drives are available, place online redo logs and frequently used indices on SSD.
- For non-production systems, the use of RAID10 with 10K rpm disks is sufficient.
- Whether 15K rpm disks or 10K rpm disks are being used, the same rules apply to Data Progression™ and Data Instant Replay™.
- Please refer to the table 4 below under the Data Progression™ and Data Instant Replay™ section for more information on how to configure the settings appropriately.

5 Recommended Oracle Disk Configuration

5.1 Oracle with Cooked File Systems

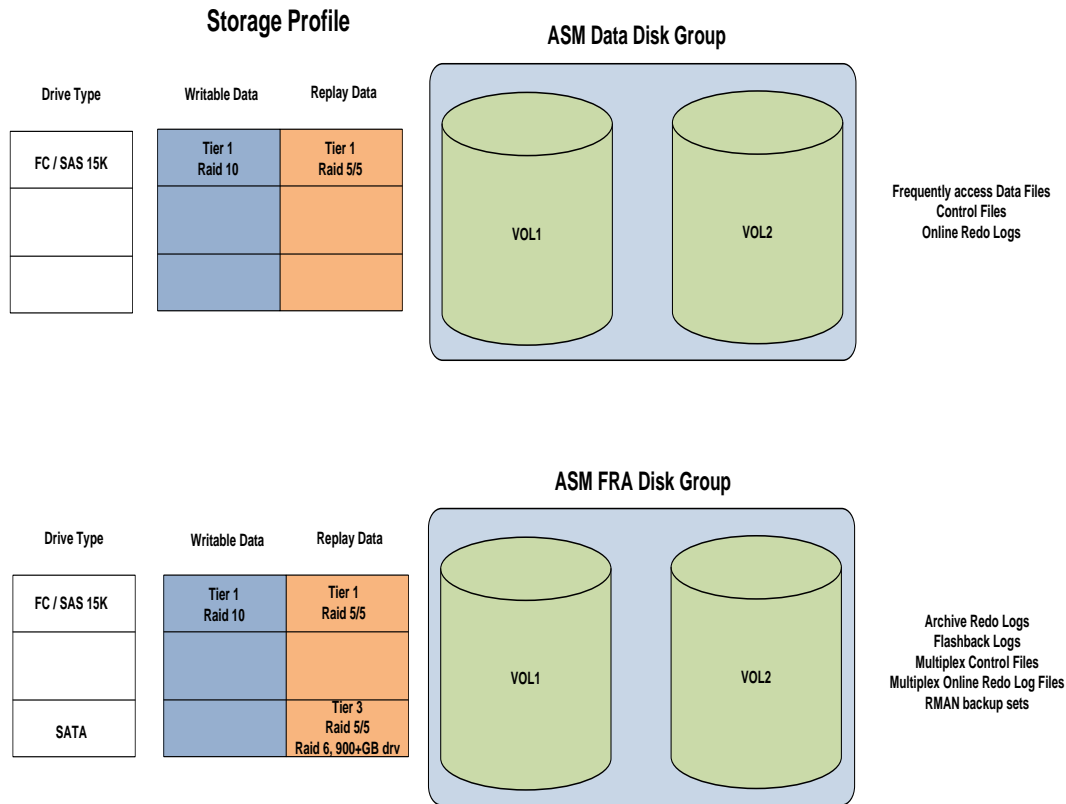


5.2 Oracle with SSD



5.3 Oracle with Automatic Storage Management (ASM)

Oracle with Automatic Storage Management (ASM)



5.4 ASM Disk Group

ASM Disk Group	Contents	RAW Device Support	ASMLib Support
+DATA	Oracle Datafiles	Yes	Yes
+FLASH	Flashback Recovery Area	Yes	Yes

5.5 ASM Disk Group Devices

- Dell Compellent Storage Center supports RAW devices with or without ASMLib service.
- Oracle recommends that generally no more than two disk groups be maintained and managed per RAC cluster or single ASM instance.
- Select External Redundancy when creating ASM disk group.
- When creating LUNs for ASM disk group, make sure the LUNs are on the same disk characteristics (ex. FC 15K rpm or 10K rpm and 146GB or 300GB). LUNs should not contain mixed speed drives.
- When creating multiple LUNs for an ASM disk group, make sure to create the same size LUN to avoid imbalance.
- Create larger LUNs to reduce LUN management overhead. This will reduce LUN management and will allow future growth. Since the Dell Compellent Storage Center provides thin provisioning, disk space only taken up when there is data written to it.

5.6 Filesystem

Dell Compellent Storage Center supports various filesystems and raw devices:

- JFS & JFS2
- UFS
- VxFS
- EXT2 & 3 & 4
- ReiserFS
- NTFS
- OCFS
- Etc.

5.7 Database Setup and Configuration – Direct I/O and Async I/O

- Oracle recommends using direct I/O and async I/O. The use of Direct I/O bypasses the filesystem buffer cache, hence reduces CPU overhead on your server.

- Direct I/O is very beneficial to Oracle's log writer, both in terms of throughput and latency. Async I/O is beneficial for datafiles I/O.
- When mounting filesystems for your Oracle database, make sure to understand how your Operating System supports Direct I/O. Some Operating Systems require certain mount options to enable direct I/O or async I/O.

5.8 Putting it all Together

- When provisioning storage for the Oracle database server, take into consideration the following before deciding how to setup and configure the Compellent Storage Center:

- Operating System Types (UNIX, Linux, Windows, etc.)
- File System Types (NTFS, VxFS, JFS, Ext3, UFS, etc.)
- Number of Oracle databases per server
- Database Types (OLTP, Data Warehouse, Reporting, etc.)
- Database Usage (Heavy, Medium, Light)
- Archived Log Mode or No Archived Log Mode
- Database Block Size

- Based on a small workload OLTP database, create one volume for your Oracle data files, one volume for your Oracle online redo logs, and one volume for your Oracle archived redo logs in addition to other volumes such as Oracle binaries and any non-Oracle related data. Refer to the table 2 of this document to decide which RAID level you should be using for your volumes.

- When creating volumes for your Oracle database server, do not have to use any software striping at the Operating System level. The data in the Compellent volume is automatically striped depending on which RAID level you have selected for your volume. However, you should create multiple LUNs for your larger workload database and use Operating System striping mechanism (eg. LVM, VxVM) to get better performance because of multiple disk queues at the OS level if Oracle ASM is not configured.

- For a more complex database, create multiple volumes dedicated to your data files, assuming you separate your data tablespaces from your index tablespaces.

- Again, when creating these volumes, refer to the Operating System manual on how to mount these volumes with direct I/O and async I/O options. This will enhance your database performance. Also remember to set the filesystemio_options to setall or DirectIO in your database initialization parameter file.

- When creating volumes for your databases, it is recommended to initially create the volumes (datafiles, archived redo logs, flash recovery area) larger than needed for future database growth. Since the Compellent Storage Center has the ability to dynamically provision storage, disk space is not taken up until actual data has been written to it. This way you can create your tablespaces with the AUTOEXTEND parameter so you don't have to worry about running out of disk space in that volume.
- Dell Compellent Data Instant Replay™ is great for Oracle backup and recovery. With Data Instant Replay™ there is no limit on the number of replays (snapshots) taken on the Dell Compellent Storage Center. To perform an online backup (your database must run in archive log mode in order to take an online backup), the database must be placed in online backup mode. Take a replay of your data file volumes, redo log volumes, and archived log volumes, end your database online backup, and mount these replays on the backup server, and send the data to tape or you can just leave these replays as is and not expire them until they are copied to tape for offsite storage. If export and import are used frequently, create a separate volume dedicated for the export dump files.

*for more information on Oracle backup and recovery please refer to the document ***Dell Compellent - Oracle Backup & Recovery Best Practices***

6 Using Compellent Data Progression™ and Data Instant Replay Features

In order to utilize Data Progression™ effectively, you need to use Data Instant Replay™ whether or not backup and recovery is required. For Data Progression™ to work effectively, a replay (snapshot) should be taken at least once a week on all volumes. Below is the recommended Data Progression™ setting for your Oracle database volumes. For more information on Data Progression™ & Thin Provisioning please refer to the document [***Dell Compellent – Data Progression™ & Thin Provisioning with Oracle***](#)

6.1 SSD Configuration

In a SSD configuration, SSD will become Tier 1 and 15K & 10K rpm disks will become Tier 2, and SATA is Tier 3.

Tier 1 SSD	RAID	Writeable Data	Replay Data
	R10	Online Redo Logs /Control Files / Frequently Used Indices.	
	R5/5		
	R5/9		
Tier 2 15K/10K FC/SAS			
	R10	Datafiles / Archived Logs / FRA / OCR &	
	R5/5		Data Files
	R5/9		Data Files
Tier 3 SATA			
	R10		
	R5/5		Archived Logs / FRA
	R5/9		Archived Logs / FRA
	R6 (900GB+ Drives)		Archived Logs / FRA

When using SSD, for optimal database performance, create volumes with the following:

- Datafile volumes → Refer to the recommended disk configuration.
- Online Redo Log volumes → Refer to the recommended disk configuration.
- Archived Redo Log volumes → Refer to the recommended disk configuration.
- Flash Recovery Area → Refer to the recommended disk configuration.
- OCR & VOTE Disk → Create a new Storage Profile and select RAID10 only in the RAID Levels Used section and select Tier 2 in the Storage Tiers Used section and apply to volume.

6.2 Non-SSD Configuration

In a Non-SSD configuration, 15K rpm disks will be in Tier 1, and 10K rpm disks will become Tier 2, and SATA is still Tier 3.

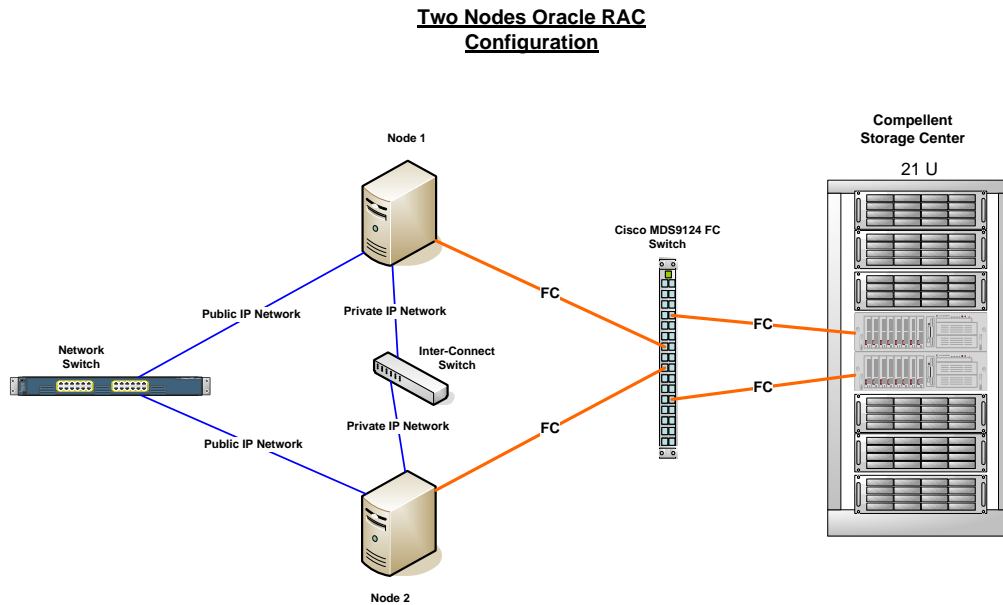
Tier 1 15K FC/SAS	RAID	Writeable Data	Replay Data
	R10	Datafiles / Archived Logs / FRA / OCR & VOTE	
	R5/5		Data Files
	R5/9		Data Files
Tier 2 10K FC/SAS			
	R10		
	R5/5		Data Files
	R5/9		Data Files
Tier 3 SATA			
	R10		
	R5/5		Archived Logs / FRA
	R5/9		Archived Logs / FRA
	R6 (900GB+ Drives)		Archived Logs / FRA

When not using SSD, for optimal database performance you should create volumes with the following:

- Data file volumes → Refer to the recommended disk configuration.
- Online Redo Log volumes → Refer to the Recommended Disk Configuration Layout.
- Archived Redo Log volumes → Refer to the recommended disk configuration.
- Flash Recovery Area → Refer to the recommended disk configuration layout.
- OCR & VOTE Disk → Create a new Storage Profile and select RAID10 only in the RAID Levels Used section and select Tier 1 in the Storage Tiers Used section and apply to volume.

7 Oracle RAC Tested Configuration

One configuration was tested consisting of a two node cluster running Oracle RAC 11g Release 2 with ASM running Oracle Enterprise Linux.



VLAN ID or Separate Switches	Description	CRS Setting
1 or switch A	Client Network	Public
2 or switch B	RAC Interconnect	Private

- If possible, configure jumbo frames for private network. Note: when configuring jumbo frames, you need to configure on all legs of the RAC interconnect networks (i.e. servers, switches, etc.).
- If configuring jumbo frames is not possible, you need to configure the interconnect network with at least 1Gbps link.

8 Conclusion

Coupling Compellent Storage Virtualization and Dynamic Block Architecture with Oracle best practices, where the data files should be spread across as many spindles as possible makes Storage Center a very robust and scalable SAN solution for Oracle Databases. Having the ability and granularity of looking at each individual block gives Compellent the advantage in ensuring that data is always spread evenly across all spindles at all times without having the DBA or storage administrators continually revisit or dedicate certain spindles of the storage subsystem to maintain performance of the database. Benefits include dramatically reducing the time and effort required for capacity planning and volume management. Traditional systems are prone to "hot spots" because there are very few drives allocated to a specific volume for that server. Spreading the access across multiple drives eliminates these hot spots and increases performance. With Dell Compellent, volumes can be created in seconds and performance is optimized with all the drives acting together.

9 Additional Resources

9.1 Oracle Performance Testing with ORION

ORION, an acronym for ORacle I/O Numbers calibration tool, is similar to the Iometer tool developed by the Intel Corporation or dd for Linux. Like Iometer and dd, ORION runs I/O performance tests, however it is tailored to measure Oracle RDBMS workloads using the same software stack that Oracle uses. It is not required that Oracle or even a database be installed in order to use ORION.

ORION is available in Windows, Linux and Solaris platforms for use in predicting the performance of an Oracle database, typically prior to installing Oracle. ORION is primarily made available for use in a pre-production or test environment, as the use of ORION is inherently dangerous in a live environment. This is due to the fact that the testing of writes is potentially destructive to existing data.

ORION was made available as a separate download with Oracle 10g, but has now been included in Oracle 11g. ORION can be used to test Storage Area Networks, Network Attached Storage or Direct Attached Storage. Although one may use the preconfigured test profiles, more accurate predictive results require configuring the test to reflect the particulars of the hardware and software to be used, as well as a determination of what type of database is needed. Typically databases fall into two main categories, depending upon how they will be used. It is very important to determine which type is needed so the ORION test can be tailored accordingly and produce meaningful results.

First, the Online Transaction Processing (OLTP) database, characterized by many small transactions (often 8k), reads and writes, in rapid succession and allowing multi-user access. An OLTP database generally consists of highly normalized data ordered with a non-clustered index. Think of a bank's ATM database. These databases are optimized to maximize the speed of myriad small transactions, thus the throughput is generally scrutinized according to the number of I/Os Per Second (IOPS).

Secondly, we have the Online Analytical Processing (OLAP) database, which typically processes fewer much larger (1MB) transactions, primarily reads. With an OLAP database, data is generally denormalized, often using a star schema and organized with a clustered index. An example would be a Data Warehouse, used by Business Intelligence for Data Mining or Reporting. These databases are designed to run complex queries on large amounts of data as efficiently as possible. Throughput considerations here focus on Megabytes Per Second (MBPS). Often it is the case that a single database, to some degree, will be providing both functionalities. ORION's workload options can be configured to test the performance one can expect from a database which have to bear such a mixed workload.

Getting started with ORION is relatively simple: download ORION, unzip/install the file, create a file listing the volumes to test, run ORION with or without workload options and view the output tables. Using ORION to best effect, however, is a bit more involved. After downloading ORION unzip or install it into a directory of your choice. This will create a folder with the ORION application and the associated Oracle libraries. Pick a unique name for this test run, say "iotest". Create a file in the same directory named iotest.lun. Edit the file to list the volumes you want to test, one volume per line, not delimited with commas, lines terminated with a carriage return. Thus, the contents of iotest.lun would be similar to the following:

```

/dev/raw/test1
/dev/raw/test2
/dev/raw/test3
/dev/raw/test4

```

IMPORTANT: Tests which include a write component can destroy data. Never include a volume that contains important data in this list, rather create new volumes for testing purposes. Check to be sure you are able to connect to the test volumes.

As the ORION test is dependent on asynchronous I/O, ensure that the platform you are using is capable of asynchronous I/O. With Windows this capability is built in, but in Linux or Solaris the library libaio needs to be in one of the standard lib directories, on the environment's library path. From the command line, navigate to the directory which holds the ORION application and run the test. There are 3 mandatory parameters: `run`, `testname`, and `num_disks`, where:

<p>-run</p>	<p>type of workload to run: simple, normal or advanced</p> <p>simple: tests random small (8k) I/Os, then tests random large (1MB) I/Os</p> <p>normal: tests random small (8k) and random large (1MB) I/Os simultaneously</p> <p>oltp: tests random small (8k) only</p> <p>dss: (decision support system) tests random large (1MB) only</p> <p>advanced: run workload with parameters specified by user</p>
-------------	--

-testname	the unique name of the test run, change each run or output files get overwritten
-num_disks	the number of actual disks (spindles). If not specified, it will default to number of volumes in .lun input file

So a minimalistic test run on volume(s) comprising 5 physical disks would look like this:

```
orion -run simple- testname iotest -num_disks 5
```

You'll see a response similar to this:

```
ORION: ORacle IO Numbers -- Version x.x.x.x.x
Test will take approximately xx minutes
Larger caches may take longer
```

When complete, the results will be outputted to files prefixed with your unique test run name. For example:

iotest_summary.txt	recap input parameters + latency, MBPS and IOPS results.
iotest_mbps.csv	large I/Os in MBPS , performance results
iotest_iops.csv	small I/Os in IOPS, performance results
iotest_lat.csv	latency of small I/Os
iotest_tradeoff.csv	large MBPS / small IOPS combinations possible with minimal latency
iotest_trace.txt	extended, unprocessed output

The .txt files are text files with summary data and unprocessed data, respectively. The summary text file is useful, providing the maximum value for each I/O type, allowing one to a best case scenario for performance expectations.

The .csv files are a list of comma-separated values that can be loaded into MS Excel, or another database application then converted into a graph for easier evaluation of results.

An ORION test can be tailored more closely to the environmental variables by using the 'advanced' designator with the -run command. Here are some of the options available:

-size_small	small I/O size in KB (default = 8)
-size_large	large I/O size in KB (default = 1024)
-type	type of large I/Os: rand: random (default) seq: sequential streams

-num_streamIO	concurrent I/Os per stream, i.e. threads (default = 1) number of cpus x number of parallel threads per cpu NOTE: only works if type is specified as seq
-simulate	how ORION forms the virtual test volumes concat: serial concatenation of volumes (default) if devices are already striped raid0: raid 0 mapping across all volumes if devices are not already striped. raid0 stripe size in KB (default = 1024)
-stripe	
-write	percentage of writes (default = 0). A typical percentage would be 'write 20' NOTE: Write tests destroy data on volumes.
-cache_size	size in MB of array's cache. 'cache_size 0' disables cache warming (recommended)
-duration	duration in seconds for data point (default = 60) NOTE: if -num_disks > 30, set -duration 120
-matrix	defines which data points to test: columns are fixed small I/O load rows are fixed large I/O load basic: test first row and first column detailed: entire matrix tested col: defined small I/O load with varying large I/O load (-num_small) row: defined large I/O load with varying small I/O load (-num_large) max: tests varying small and large I/O loads up to set limits (-num_small, -num_large)
-num_small	number of outstanding small I/Os (matrix must be point, col or max)
-num_large	number of outstanding large I/Os (matrix must be point, col or max)
-verbose	prints tracing information to standard output (default = not set)

Here is an example of an elaborate custom workload test setting in ORION:

```
orion -run advanced -testname iotest -num_disks 5 \  
-simulate raid0 -stripe 1024 -write 20 -type seq \  
-matrix row -num_large 0
```

NOTE: To get useful test results it is often advisable to disable the array's cache and set `'-cache_size 0'`, or at least set the cache size variable to a fraction of the actual cache size. This is due to the fact that in a production environment Oracle rarely monopolizes the entire cache, as there are other production demands. Testing with access the array's cache fully enabled and dedicated will tend to yield results that are too good to be meaningful outside a protected test environment.

Again, the ORION tool is meant to be used to test a pre-production environment, prior to installing Oracle and does not require that Oracle or a database be in place at the time of testing. Also, ORION tests writes as well as reads, so caution should be used as the data residing on the volumes will be destroyed.

A more detailed description of the use of ORION can be accessed from within ORION, as follows:

```
orion -help
```

Also, an excellent introduction to ORION can be found on Oracle's website:

http://docs.oracle.com/cd/E11882_01/server.112/e16638/iodesign.htm#BABFCFBC

The information in this section was compiled primarily with information from these two sources.