Cost and Performance benefits of Dell™ Compellent™ Automated Tiered Storage for Oracle® OLAP Workloads

This Dell technical white paper discusses performance and cost benefits achieved with Dell™ Compellent™ Automated Tiering technology for Oracle® OLAP workloads.

Balamurugan B and Ramamohan Reddy K

Database Solutions Engineering

Dell Product Group

April 2012

April 2012| Rev 1.0

Contents

Tables

Figures

Executive summary

Data growth in the Enterprise world is enormous. In the traditional system, to achieve high performance needs, all the data must reside in the high performing disks irrespective of the access pattern. However, all the data residing in these disks may not be accessible all the time as a result high performing disks are used inefficiently. To avoid this, many storage vendors have adopted a concept called *tiering* . Tiered storage consists of different disks that differ in cost, performance, and capacity. Tiered storage enables placing the data in the right tier based on usage pattern, performance, and cost requirements but it requires a manual intervention to move and reformat data.

In Dell™ Compellent™ storage systems, this process of tiering is automatic. You get the benefit of automatic tiering by enabling data progression, which automatically moves data to the right place at the right time based on usage pattern and performance needs. This paper emphasizes Dell Compellent Automatic Tiering and provides the cost and performance benefits that are achieved using the automatic storage tiering features like Fast Track, Data Instant Replay, and Data Progression for Oracle® database OLAP workload.

Introduction

Storage tiering is an important strategy for optimizing storage expenditures, which helps organizations achieve lower storage costs by storing less important data on less costly disks. However, the hidden costs of data classification and the manual movement of data between tiers, along with the purchase and integration of multiple point products, quickly eliminate any potential savings. Without an intelligent and automated solution, the benefits of tiered storage are immediately outweighed by the resulting cost and management burden. Because Compellent Storage Center operates at the block level, Dell Compellent storage has the intelligence to move infrequently accessed data to lower tier disks automatically. With Compellent Data Progression enabled, you can take advantage of a complete automated tier storage solution and eliminate the need for manual data classification and transfer.

Oracle Database workloads are classified as online transaction processing (OLTP) or online application processing (OLAP) or a mix of both workloads. For a typical OLTP workload, only a small set of data is active, which must be placed in a high performing storage tier. In an OLAP database, most of the data is accessed only *once in a while*. This OLAP data is placed in high-capacity and low-cost tiers. In addition, in an OLAP work load, the amount of data that gets written is small, so you can get a cost benefit by placing data in cost efficient RAID5 storage.

This paper outlines how each of the Compellent automatic storage tiering features helps the Oracle Database OLAP workload to achieve better usage:

Fast Track	Fast Track places frequently accessed data on the outermost sector of the drives to improve the performance.
Data Instant Replay	Data Instant Replay moves read-only data from RAID 10 to RAID 5 within the same tier, which maintains the same read performance and frees up space on the RAID level.
Data Progression	Data Progression deals with moving infrequently accessed blocks of data to a lower storage tier and RAID level, or to a different RAID level within the same tier. Whether data is moved to a new tier or RAID level, the migration frees up space on the higher tier or RAID level.

Technology Overview

Dell Compellent Fluid Data Architecture – feature description

Dell Compellent created an enterprise storage solution that automates the movement and management of data throughout its life cycle with Fluid Data Architecture. Fluid Data Architecture enables true storage virtualization, providing a single pool of storage resources that spans across all disks and RAID levels. It also enables a dynamic block architecture that actively and intelligently manages data at more granular level. With the automated storage tiering, data is placed intelligently and efficiently among the multiple tiers and RAID levels.

The following features describe Fluid Data Architecture from Dell Compellent.

Built-in virtualization

A virtualized storage environment provides an ideal foundation for tiered storage. With a truly virtualized environment, there are no constraints on where data can reside, as data is not confined to particular disk group. You can move data without limitation to a tier with a particular drive type or performance level, or to a particular RAID level within a storage tier, depending on usage patterns.

Figure 1. Fluid Data Architecture

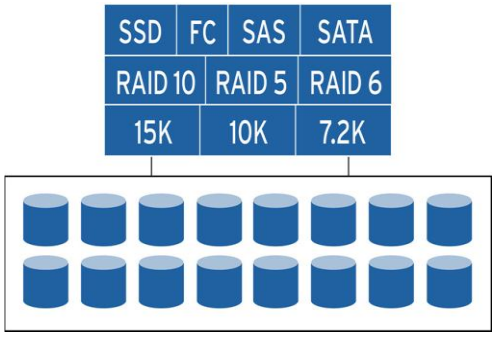

Block-level data management

Fluid Data Architecture enables true storage virtualization and facilitates a highly granular approach to storage tiering. Dell Compellent divides data into pages, which are preformatted collections of allocated and unallocated disk blocks. The pages are 2 MB by default, though administrators have the option to manage storage as 512 KB or 4 MB pages, and details about each block are captured in action. This granular approach optimizes tiering. Using real-time system intelligence, the Dell Compellent SAN identifies very small data blocks that are eligible for movement based on access frequency and then moves just the small pages that include those blocks to the appropriate tier or RAID level, creating a finely tuned tiered storage environment. Using small pages also increases the efficiency of data movement since it is more efficient to read, write, and migrates smaller pages rather than larger ones. Static storage architectures are limited to much larger blocks of data, ranging from 16 MB to 1 GB.

Figure 2. Identifying very small data blocks that are eligible for movement

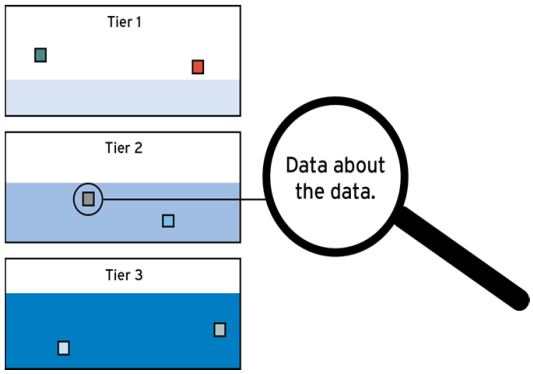

Data Progression

Data Progression automatically classifies and migrates data to the optimum tier of storage, retaining frequently accessed data on high performance storage and storing infrequently accessed data on lower cost storage. Data Progression reduces overall storage cost up to 74 percent by classifying and moving data at the block level between tiers of storage based on access frequency. This complete automated tiered storage solution eliminates time consuming data classification and the repetitive, manual transferring of data between tiers. Data Progression helps to provide a fully automated sub-LUN tiering. You can choose different storage profiles to drive the placement and movement of data and apply profiles to a single LUN, a group of LUNs, or multiple LUNs within a group, fine-tuned by application if desired. The profiles can specify not only the storage tiers to use, but also the various disk types, rotational speeds and/or RAID levels within each tier. You can change profile settings at any time without downtime or disruption to the production environment.

Migration occurs automatically at a set time you define, or on demand, while the system is still online. The migration process runs in the background and does not affect data availability or application performance. There is no need to bring down an application, pause I/O, or wait for a minimum I/O requirement.

Figure 3. A Built-in migration engine provides hands-free management through tiering profiles

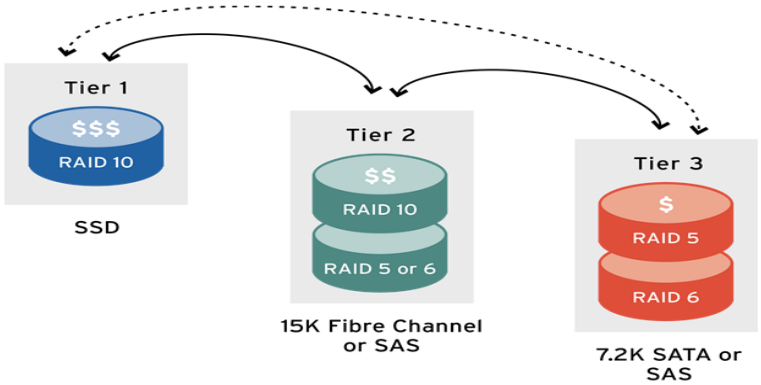

Data Instant Replays

Dell Compellent Data Instant Replay technology plays a key role in delivering outstanding performance in an automated tiered storage environment. Data Instant Replay protects data without wasting capacity. You create space-efficient snapshots (Replays) of changes in data for continuous protection from loss or corruption. Replays use minimal storage space because only data written to the volume since the previous Replay was taken is captured on the array.

In the Dell Compellent architecture, new data is written by default to tier 1, RAID 10 storage to provide the best write performance. Over time, according to the Storage profile, infrequently accessed blocks of data are moved to a lower storage tier and RAID level, or to a different RAID level within the same tier. Moving this read-only data from RAID 10 to RAID 5 within the same tier enables the solution to maintain the same read performance. Whether data is moved to a new tier or RAID level, the migration frees up space on the higher tier or RAID level. When new data is written to an existing block that was converted to read-only and migrated to a lower tier, those writes are redirected to the tier 1, RAID 10 storage. A new writable block is automatically allocated to provide the highest transaction performance.

Figure 4. Snapshot blocks are periodically moved to the lower tiers

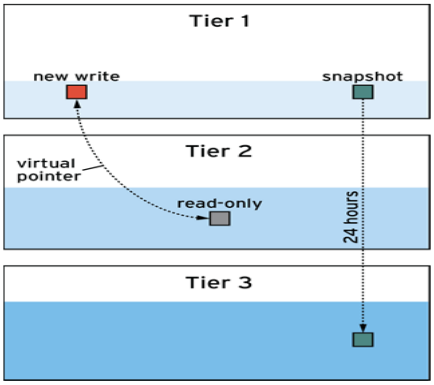

Fast Track

Dell Compellent Fast Track boosts automated tiered storage by delivering optimal placement of data on each individual disk. Fast Track uses the intelligence continuously collected by the Dell Compellent system to identify the most active, frequently accessed blocks of data on each spindle. It then places those data blocks together on the outer tracks of each drive. Keeping the head focused on that one area of the disk, where the active data resides, and delivers better performance than if the head was forced to move all over the disk.

Fast Track works on every spindle-based drive type within the system. Data movement happens automatically, on a daily basis, without any manual intervention. With Fast Track, organizations optimize performance of each platter based on actual usage patterns, reducing seek times for the most active blocks of data. Fast Track also helps reduce long-term storage costs by maximizing performance on cost-effective drives.

Figure 5. Delivers optimal placement of data on each individual disk

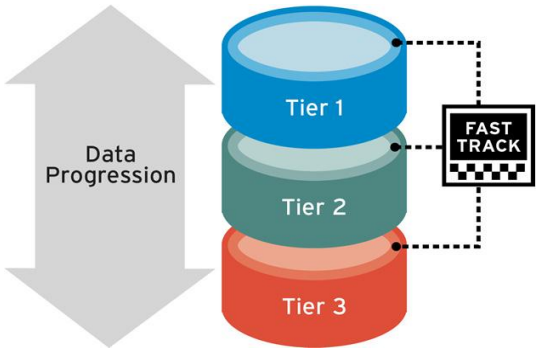

Test configuration

A two-node Oracle 11gR2 Real Application Cluster (RAC) was deployed for the database configuration in this paper. The following hardware and software components were used for the cluster setup.

Table 1. Hardware and software configurations

Component	Details
Servers	Two Dell PowerEdge R720 servers with: • Two Intel Xeon 8-core 2.20 GHz Processors • 64 GB of RAM
Storage	Dell Compellent SC40 with: • Two controllers • One 3.5” SAS enclosure with 12x450GB 15K RPM SAS 3.5” disks • One 2.5” SAS enclosure with 5x450GB 10K RPM SAS 2.5” disks and 6x1TB 7.2K RPM SAS 2.5” disks • Storage Center v5.5.3
HBA	8Gb Emulex LPE12002
Network switch	Brocade 5100 for Fiber channel storage connect
Operating system	RHEL 5.7
Oracle software	Oracle 11g R2 11.2.0.3 RAC
Workload	• Quest Benchmark Factory 6.1.1 TPC-H workload • User connections: One user with 15 Queries

Component	Details
Volume Configuration	3 X OCRVOTE - 2GB Control files - 3GB Online Redo log - 4GB DATA - 700GB FRA - 700GB

The architecture overview of the test configuration is shown in Figure 6

Figure 6. Test configuration architecture

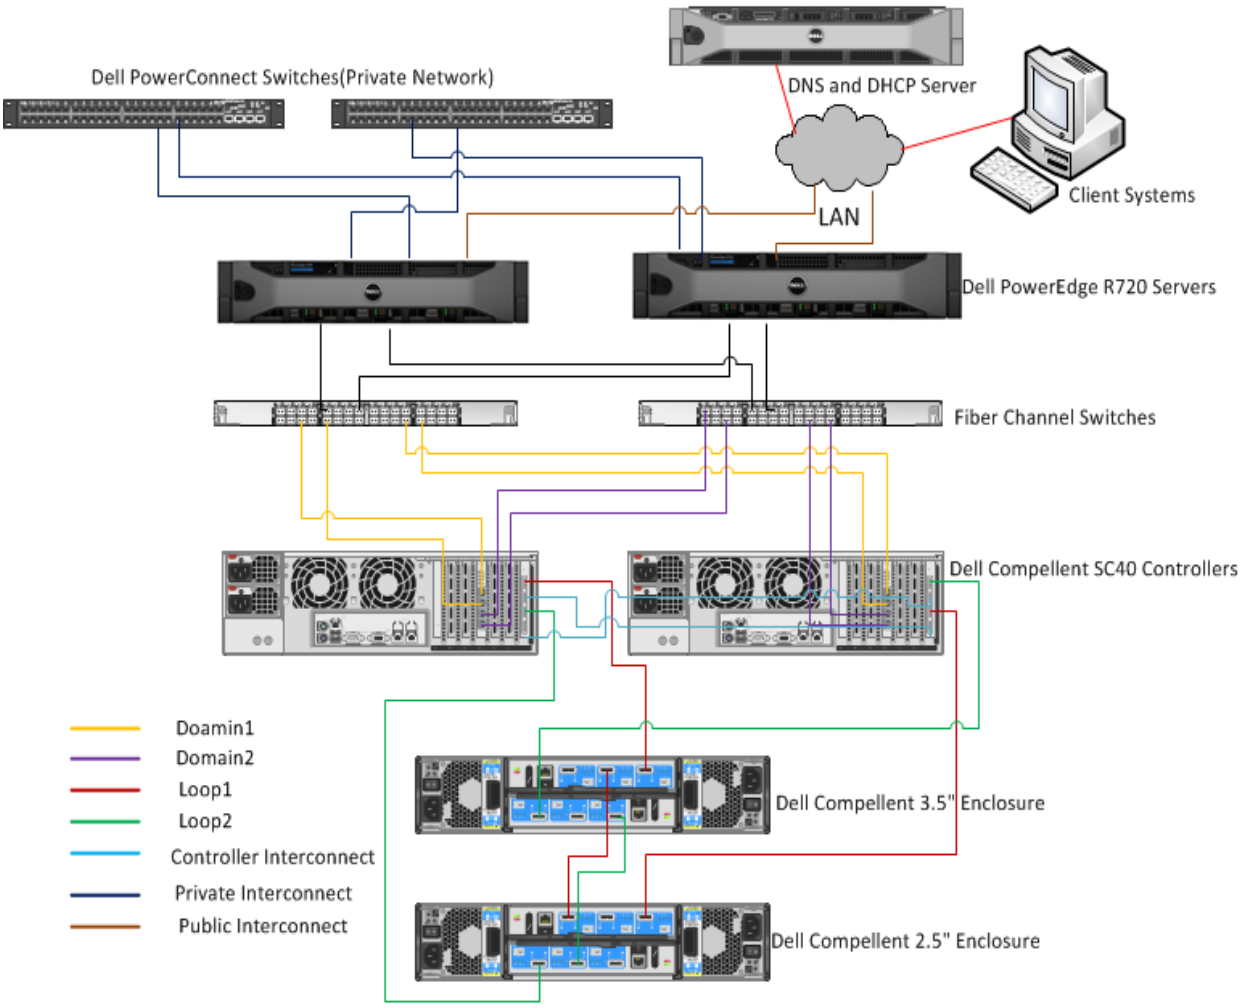

Test methodology

To observe the performance and cost benefits of Dell Compellent Automated Tiering Storage for an Oracle OLAP workload, various tests were performed by database solution engineers using Quest Software Benchmark Factory TPC-H, a load-generating utility that simulates OLAP workload on a database. Initially the tests were conducted on single storage tier to get the baseline numbers. Then the engineers enabled each of the Dell Compellent storage features like Fast Track, Data Instant

Replay, and Data Progression to find how the Compellent Fluid Architecture helps in auto-tiering resulting in performance improvements and cost benefits.

The test methodology used is as follows:

1. To simulate the traditional SAN storage, we configured a single tier (12 X 15K SAS drives) on the Dell Compellent storage and deployed Oracle database volumes such as OCRVOTE, DATA, FRA , Control file, and Redo logs . We used Benchmark Factory to load 500GB of data and stressed the database using TPC-H to gather the query response time and the disk space usage on the storage.
2. To obtain the benefits of the Compellent Fast Track feature, we enabled Fast Track by applying the license on the Storage Center. We performed multiple Quest Benchmark Factory TPC-H runs to move the frequently accessed blocks to the outer sector of the disk to get the optimum performance. We then gathered the query response times and the disk space usage on the storage.
3. We applied the Data Instant Reply profile (Daily once expires in 3 days) for the database volumes and performed multiple Quest Benchmark factory TPC-H runs to find out the cost savings and performance benefits by moving the read only data to the RAID 5-9, and gathered the query response time, and the disk space usage on the storage.
4. We introduced additional two tiers, Tier 2 Storage (5 X SAS 10K disks) and Tier 3 storage (6 X SAS 7.2 K Disks) to the existing system and applied the below mentioned storage profiles to the database volumes.

   Storage Profile1 (Tier1 RAID10, Tier1 RAID5-9): For OCR, Redo logs and control files

   Storage Profile2 (Tier1 RAID10, Tier1 RAID5-5 and Tier2 RAID5-5): For DATA

   Storage Profile3 (Tier1 RAID10, Tier3 RAID6-6): For FRA

Table 2. Volume distribution in multiple tiers

		Writeable Data	Replay Data
Tier1	RAID 10	OCR, Redo logs, Control files, DATA and FRA	
	RAID 5-9		OCR, Redo logs, Control files and DATA
Tier2	RAID 10		
	RIAD 5-5		DATA
Tier3	RAID 10		
	RIAD 6-6		FRA

We performed multiple Quest Benchmark factory TPC-H runs to find out how Compellent Auto-tiering helps to distribute the database volumes across different tiers based on the IO pattern and observed how it helps to reduce cost and improve performance.

Test results and analysis

As explained in the test methodology, we used the Quest benchmark TPC-H tool with 500 GB of data loaded. After loading the data, the total size of all database volumes was around 950GB including FRA. Initially, we started testing the OLAP database on a single tier (SAS15K RPM) with Data Progression, Data Instant Replay, and Fast Track disabled. This was to get baseline results like traditional storage without any Dell Compellent auto-tiering features. Then we conducted Quest benchmark TPC-H tests on the database by enabling each of those tiering features one-by-one as mentioned above and gathered performance numbers in terms of average query response time. For each test scenario, we captured the overall disk space consumed by all database volumes in the Compellent storage.

Figure 7 shows the response time comparison between baseline configurations without any Compellent auto tiering features enabled (to simulate a traditional storage system) and with Fast track enabled.

Figure 7. Average response time: Baseline vs. Fast Track enabled

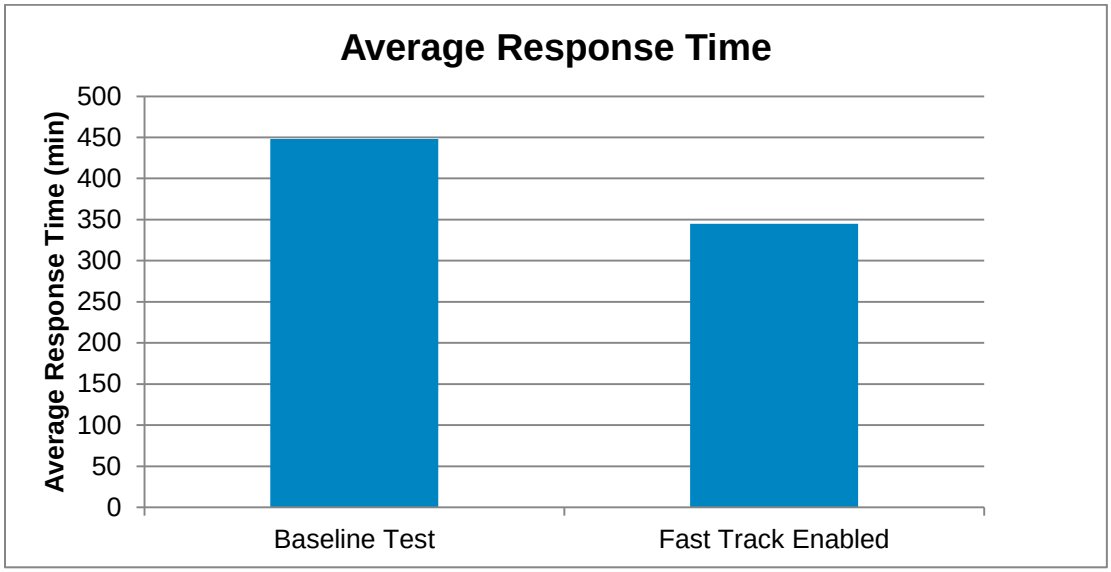

The above test result reveals that the baseline test carried out without Fast Track on all volumes residing on the SAS 15K RPM drives took an average response time of 448 minutes whereas on the same storage configuration when the Compellent Fast Track feature is enabled, after couple of test iterations, we observed that the average response time was reduced to 345 min. This result is because Dell Compellent Fast Track moves frequently accessed data blocks to the outer sectors of the drives. In the above test, there was clearly a 23 percent improvement in the performance by enabling Dell Compellent Fast Track. In addition, we observed that Dell Compellent Fast Track does not bring a cost advantage by means of disk space savings.

Figure 8 shows the response time comparison between baseline configurations without any Compellent auto-tiering features enabled (to simulate a traditional storage system) and Dell Compellent Fast Track and Replay enabled.

Figure 8. Average response time: Baseline vs. Fast Track enabled vs. Replay enabled

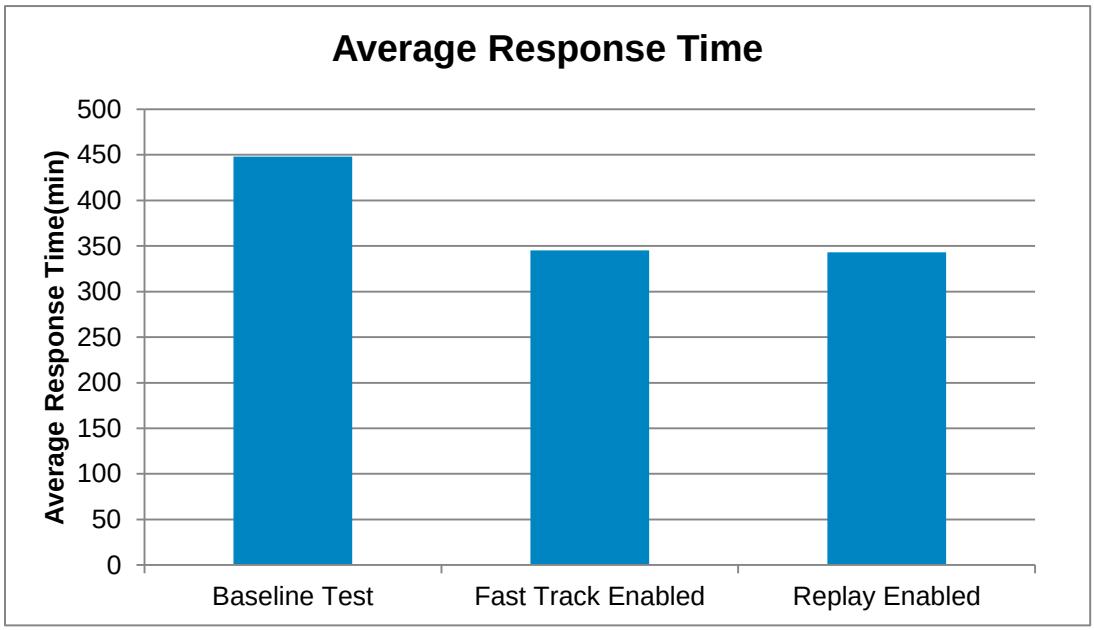

Based on the above test results, when the volumes were enabled with the Replay profile, we observed that the query response time was 343 minutes, which is well below the baseline test and also the same as the response time with Dell Compellent Fast Track enabled. With the above results, we conclude that by enabling Replay, there is no impact on the performance.

Figure 9 shows the comparison of disk space consumed (in TB) for the baseline test and with Data Instant Replay enabled.

Figure 9. Baseline disk space consumed vs. Replay enabled

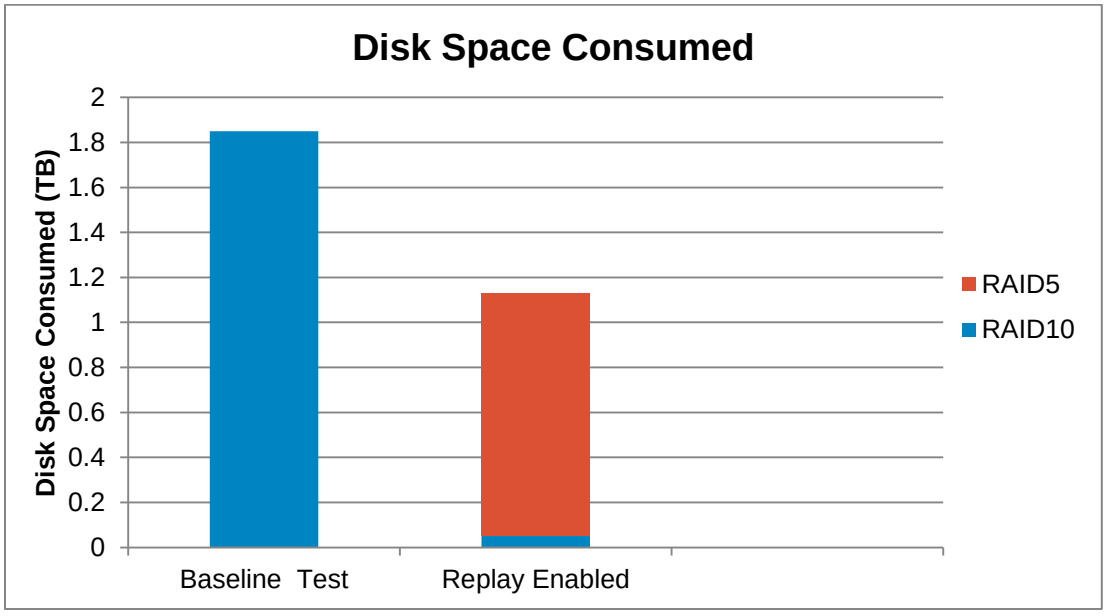

In the above graph we observed that after enabling Data Instant Replay, the data migrated to RAID 5-9 from RAID 10. Moving this read-only data from RAID 10 to RAID 5-9 within the same tier enables to maintain the same read performance and frees up space on the RAID level.

In the Baseline test, without enabling Replay, the disk space consumed for all volumes is 1.85TB with all the data residing on a RAID 10. After enabling Data Instant Replay, all the read-only data was identified and automatically moved to RAID 5-9. As a result, the total disk space consumed was reduced to 1.13 TB. Enabling Replays helped to shrink disk space by 39 percent and delivered the same performance.

Figure 10 shows the response time comparison between the baseline configuration and with all the Compellent auto-tiering features enabled.

Figure 10. Average response time: Baseline vs. Auto-tiering features enabled

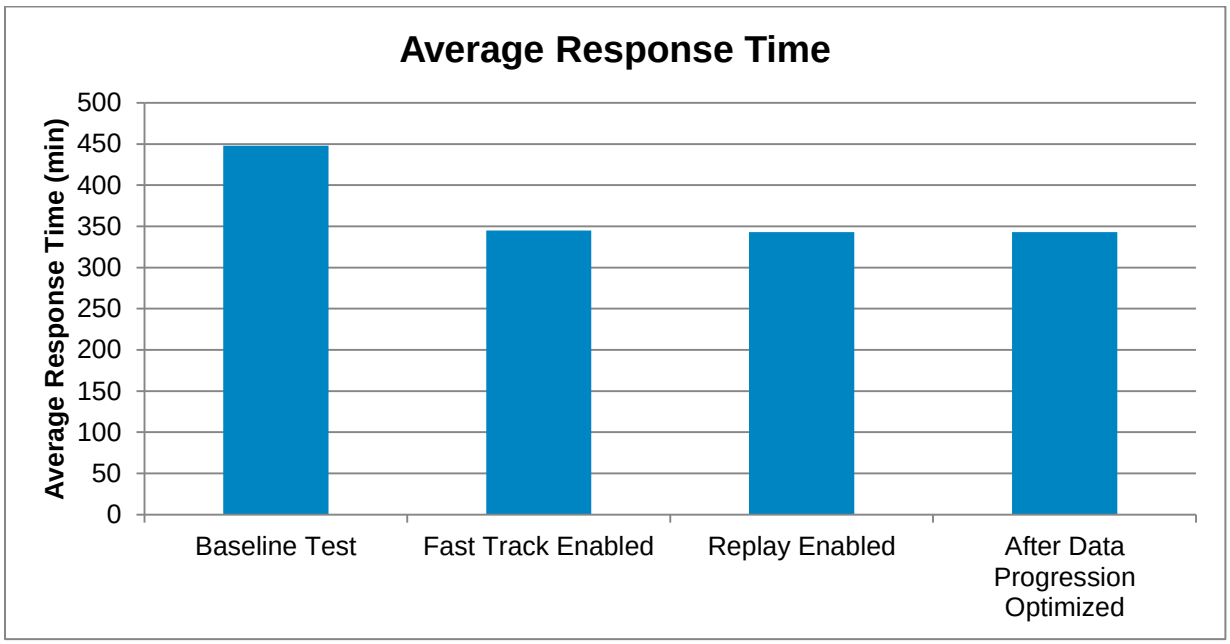

Based on the above test results, when we enable the Compellent auto-tiering features (Fast Track, Data Instant Replay, and Data Progression), we observed that there is no degradation in terms of query response time. The test results are well below the baseline test query response time and helped to save disk space.

Figure 11 shows the comparison of disk space consumed in a baseline test after Data Progression with multiple tiers. In addition, it shows the disk space tier distribution after introducing lower tiers (10K RPM SAS and 7.2K RPM SAS drives).

Figure 11. Disk Space Consumed Tier wise distribution

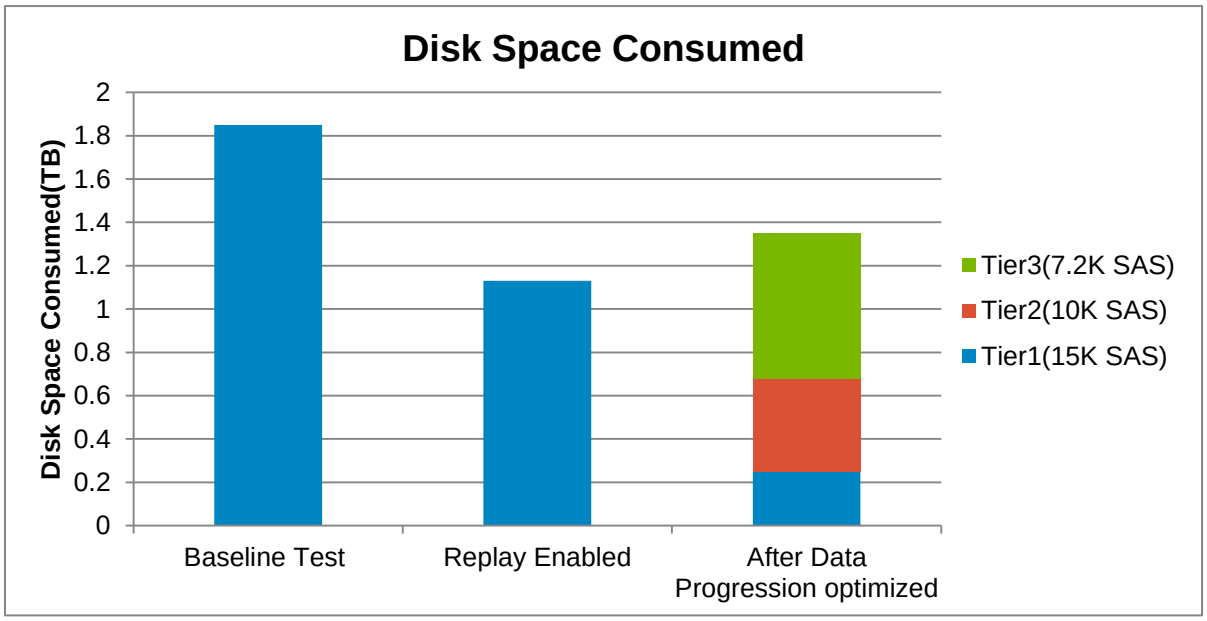

In the Figure 11, we observed that after introducing tiers and enabling Data Progression based on the parameters outlined in the storage profile that infrequently accessed blocks of data are moved to a lower storage tier and RAID level, or to a different RAID level within the same tier. Whether data is moved to a new tier or RAID level, the migration frees up space on the higher tier or RAID level.

After the Data Progression cycle on the tiered storage, the total disk space consumed was around 1.35 TB against 1.85 TB of the baseline configuration. The disk space consumed in the baseline test (1.85 TB) was on Tier 1 storage only where as in the tiered configuration it was observed that 1.35 TB of disk space is automatically distributed across different storage tiers (0.25 TB on Tier 1, 0.43 TB on Tier 2, 0.67 TB on Tier 3), which helped to save 70 percent disk space in the high cost tier1 storage.

This is due to the movement of infrequent accessed data to Tier 2 and Tier 3 storage space based on the storage profile defined by means of IO characteristics of the database volumes.

Summary

Table 3 shows the overall summary of all test cases in terms of performance, cost, and disk space usage across different tiers.

Table 3. Results summary: performance, cost, and disk space usage

	Overall volumes size(GB)	Volume size distribution in tiers and RAID types			Disk space consumed		Disk space saved in Tier1	Total cost($)	Cost per GB($)	Aver. response time (min)
Baseline configuration	950	T1	R10 Standard	950 GB	1.85 TB	All in T1	-NA-	2147	1.13	448
Fast Track Enabled	950	T1	R10 Fast R10 Standard	461 GB 491 GB	1.85 TB	All in T1	-NA-	2147	1.13	345
Replay Enabled	1013	T1	R10 Fast R5-9 Standard	21 GB 994 GB	1.13 TB	All in T1	737 GB	1312	1.13	344
After Data Progression has optimized	1018	T1	R10 Fast R5-9 Fast R5-9 Standard	21 GB 22 GB 167 GB	1.35 TB	T1 256GB T2 442GB T3 690GB	1.60 TB	770	0.55	343
		T2	R5-5 Fast R5-5 Standard	290 GB 63 GB						
		T3	R6-6 Fast R6-6 Standard	4 GB 451 GB						

Note: Refer Appendix for the sample cost used for the above cost calculations.

T1 – Tier 1 Storage Space (SAS 15K RPM)

T2 – Tier 2 Storage Space (SAS 10K RPM)

T3 – Tier 3 Storage Space (SAS 7.2K RPM)

Figure 12 shows the overall summary of total storage cost.

Figure 12. Total storage cost comparison

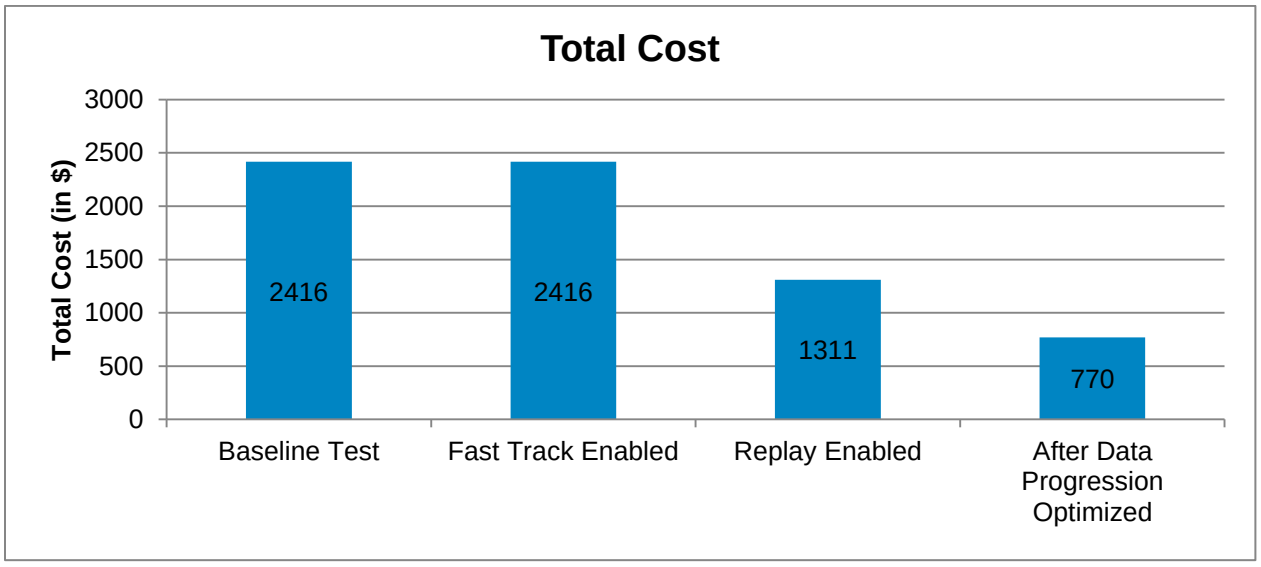

The table 3 shows the baseline configuration with only Tier 1 storage on a RAID 10 configuration. The response time was 448 minutes and the total cost per GB is 1.13$, whereas with the same configuration enabled with Compellent Fast Track and Replay features achieved an improved performance of 345

minutes with the same cost per GB. By enabling Compellent Data progression on the tiered storage, you can achieve the same performance with a reduced cost per GB of 0.55$. By enabling the Compellent auto-tiering, you can achieve lower cost per GB with improved performance.

From Figure 12, it is clear that by enabling Data Instant Replay, the total storage cost went down by 46 percent by moving the read-only data to the space efficient RAID 5 optimized for read operation. In addition, by applying Data Progression on the tiered storage helped to reduce the total cost by 64 percent with sub-LUN tiering.

The above study concluded that Configuring Oracle OLAP database, with Dell Compellent Automated Storage Tiering, the following benefits were achieved:

- With Fast Track enabled, there is a 23 percent improvement in performance.
- Compellent Auto Tiering helped to save disk space by 27 percent by moving read-only/inactive data to a cost efficient RAID5-9 and RAID6-6 from RAID10.
- Compellent Auto Tiering helped to reduce the Tier 1 storage space by 70 percent by moving infrequent accessed data to the lower tiers.
- Compellent Auto Tiering with data progression helped to reduced cost per GB by 51 percent without compromising performance.

References

1. The Architectural Advantages of Compellent Automated Tiered Storage

   http://www.compellent.com/~/media/www/Files/White_Papers/Dell_ATS_RR.ashx

2. Compellent Storage Center with Oracle Database

   http://kc.compellent.com/Pages/Download.aspx?DocID=790

3. Oracle Best Practices on Compellent Storage Center, a Dell Technical whitepaper

Appendix

Below are the approximate costs used for calculating costs.

Type	Size(GB)	Usable Capacity(GB)	Price($)
15K RPM SAS (3.5”)	450	419.19	475
10K RPM SAS (2.5”)	450	419.19	350
7.2K RPM SAS (2.5”)	1024	931.51	150