Consolidating Oracle OLTP with Dell Infrastructure

This white paper discusses OLTP database performance improvements using the Dell PowerEdge R730.

Phani MV, Dell Global Solutions Engineering

Dell Technical Solutions
September 2014

A Dell Reference Architecture

Revisions

Date	Description
September 2014	Initial release

Table of contents

Executive summary

Business need addressed

With every leap forward, technology offers a rich set of features and capabilities to customers. Improved IT infrastructure helps customers grow their businesses, and achieve business results faster with reduced TCO and higher margins.

The Dell™ enterprise portfolio is evolving to incorporate better performing, more energy efficient, and more highly available products. With the introduction of Dell's latest PowerEdge™ 13th generation server product line, customers have an opportunity to improve their businesses by consolidating distributed legacy environments.

Dell strives to simplify IT infrastructure by consolidating legacy production environments and reducing data center complexity while still meeting customers' needs. The tools and procedures described in this white paper can help administrators test, compare, validate, and implement the latest hardware and database solution bundles. Dell established these procedures and guidelines based on lab experiments and database workload simulations performed by the Dell Database Solutions Engineering team. The tools and procedures described in this document help customers to optimize their database workloads.

Solution overview

The solution discussed in this paper describes an OLTP workload scenario, and provides guidance in building the reference architecture. This paper calls out the benefits of consolidating Oracle OLTP workload onto the PowerEdge R730 server.

The key results of this white paper are included in the snapshot below:

PowerEdge R730

- Delivers up to 2 times the TPS compared to PowerEdge R710
- Delivers up to 2 times the database userload support
- Query response time is 85% less compared to the PowerEdge R710 configuration
- 2:1 server consolidation
- 50% rack space savings

1 Introduction

1.1 Objective

The purpose of this white paper is to introduce and demonstrate some of the features and capabilities of the PowerEdge R730 and how the system can be leveraged to consolidate and accelerate Oracle OLTP infrastructures.

1.2 Scope

- Document the architecture, design, and components used in the solution
- Discuss design and benefits of the latest Dell PowerEdge 13th generation R730 server
- Characterize the performance of the Dell PowerEdge R730 for OLTP workloads

1.3 Assumptions

This document assumes familiarity with server, storage, and database concepts. The reader should have a high level of familiarity with concepts such as RAID, Fibre Channel, Serial Attached SCSI (SAS), LUN, and general IP networking concepts. Familiarity with the Dell Compellent Storage Platform and Oracle Real Application Clusters, although not necessary, can be beneficial in understanding the solution presented in this document.

1.4 Audience

This document is intended for Oracle DBA's, database architects, storage administrators, architects, and customers seeking to deploy cost effective Oracle OLTP infrastructures. The solution presented in this paper provides the audience a more cost effective solution to address the demanding performance requirements of Oracle OLTP infrastructures.

2 Developing an Oracle OLTP reference architecture

The data access pattern of a typical OnLine Transaction Processing (OLTP) system is mostly transactional with short read and write IO operations of small block sizes. OLTP applications are characterized by high select, insert, or update intensive operations. Typical examples of OLTP system operations are banking system ATM transactions, a reservation system, a stock market, and retail sales systems.

The power, scalability, and effectiveness of an OLTP system can be determined based on the following characteristics:

- Server computing power
- Server memory foot print
- Concurrent database user load support
- Query response time

2.1 Design goals

This paper discusses the design and development of OLTP reference architecture with PowerEdge R730 and Dell Compellent storage array. The primary goal is to characterize the benefits of using the PowerEdge R730 when compared to the PowerEdge R710.

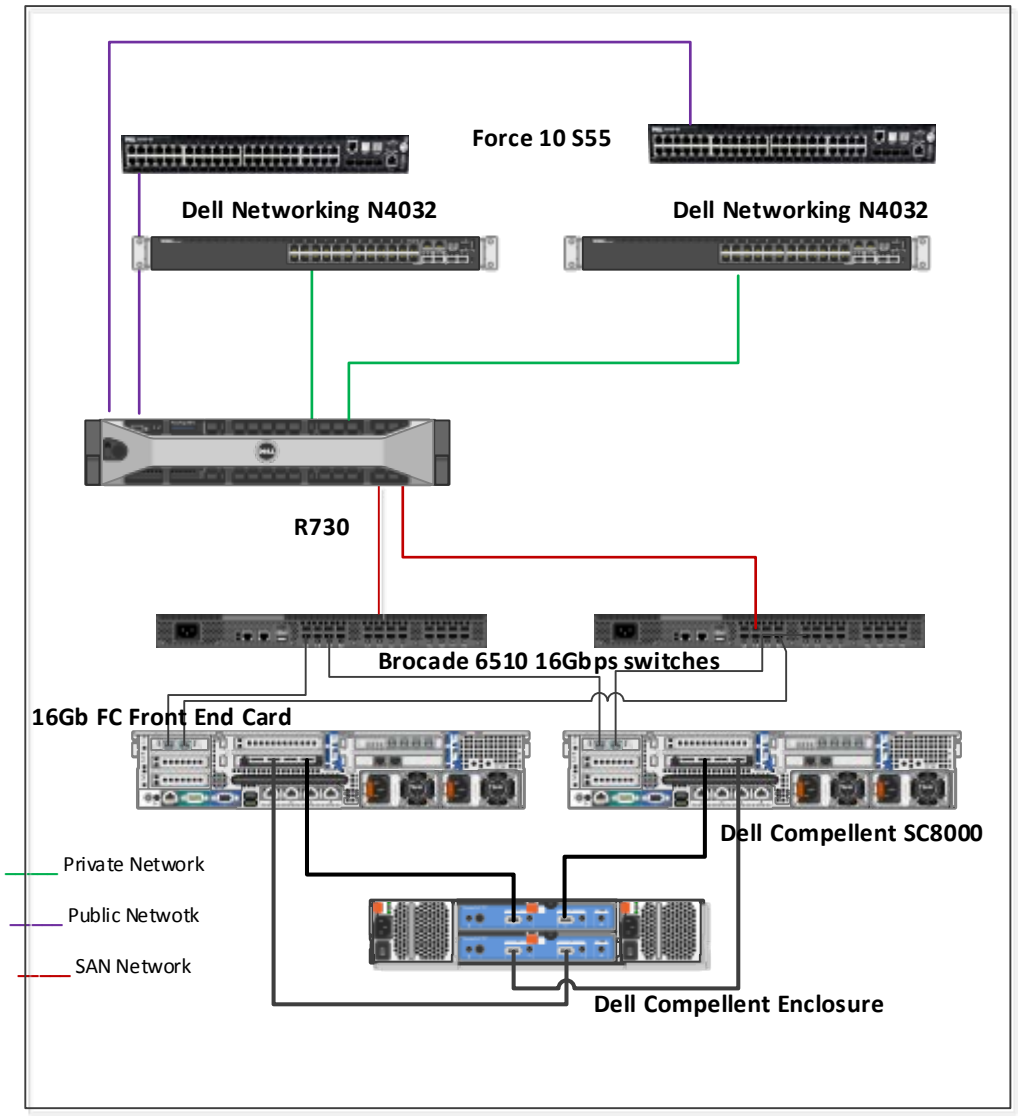

Figure 1 OLTP reference architecture schematic diagram

Detailed storage and database configuration information of the reference architecture is presented in Appendix A.

NOTE: Figure 1 is a schematic representation of the PowerEdge R730 solution. The PowerEdge R710 solution uses the same components apart from the server itself.

2.2 What is the Dell PowerEdge R730?

The Dell PowerEdge R730 is powered by the next generation of Intel® Xeon® processor E5 v3(Haswell-EP) product family. It is a 2-socket, 2U rack server ideally suited to complex and demanding workloads.

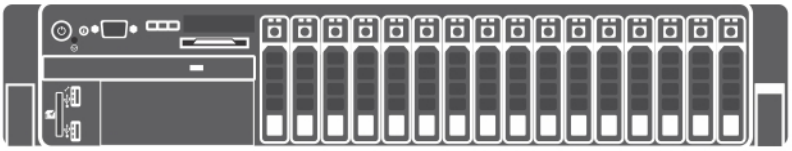

Figure 2 PowerEdge R730 front vew

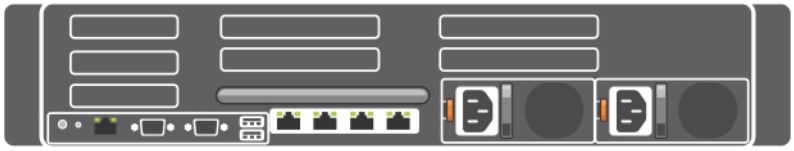

Figure 3 PowerEdge R730 back view

More memory options are available than ever before with Dell PowerEdge R730 including greater capacities, high frequencies, and more flexibility. At the time of market release, It supports up to 768 GB of memory with twenty-four, 32 GB DIMMs and speeds upto 2133 MT/s providing high performance for various applications.

The Dell PowerEdge R730 supports as many as 7 PCIe IO slots, and provides customers with the flexibility to get started with a minimal number of network cards. The PowerEdge R730 server can be the right platform for workload consolidation or transaction intensive applications, with features that are essential to building a successful OLTP system.

2.3 What is Dell Complellent storage?

Dell Compellent storage is based on fluid data architecture powered by dynamic block architecture and offers a true enterprise class virtualized storage solution with features that far exceed traditional storage capabilities. Dell Compellent products are designed to be intelligent enough to take automated decisions in optimal data placement across different class of hard drives.

Dell Compellent storage features

- True thin provisioning
- Dynamic capacity
- Fast track
- Tiered storage
- Data progression
- Data instant replay

2.3.1 Why Dell Compellent flash optimized storage?

Typical OLTP deployments require a query response time less than 1 second, and in some cases even in the order of milliseconds. In order to address the increasing needs of customers, OLTP system storage has to be properly designed and sized to match server computing power. Storage infrastructure has to be designed in such a way that it can deliver to highly scalable performance and capacity requirements.

Dell Compellent flash optimized storage built on core fluid data architecture has changed the economics of flash storage by making flash storage highly affordable. Dell-designed storage is an optimal and fitting choice for OLTP systems: Dell storage solutions offer flash performance and higher storage capcities (with flash drives in the order of TBs) at the price of a spinning HDD. Along with SSD drives, Dell Compellent also supports spinning drives of different types such as SAS and SATA. With a wide variety of drives supported, customers can either design storage with all flash drives, a mix of flash and spinning drives or spinning drives alone.

Flash drives supported by Dell Compellent are:

- Read-intensive or MLC (Multi Layered cell) SSD: These drives are the best fit for read operations.
- Write-Intensive or SLC (Single Layered Cell) SSD: These drives are optimized for write operations and are higher performing as compared to RI SSD. This performance comes at the cost of higher price as compared to RI SSDs.

Dell Compellent flash optimized solution

- Delivers high performance
- Eliminates overprovisioning
- Delivers cost effective performance compared to other flash offerings

2.4 Test methodology

In order to simulate an OLTP system, we chose a **Real World TPC-C style transactional workload.** Dell Quest Benchmark Factory[1] for databases was used to simulate a real TPC-C style workload. The workload was tested in an Oracle 12cR1 database environment.

The primary goal of our test methodology is to identify:

- maximum TPS delivered
- maximum database user load supported
- improvement in query response time

The test methodology used for this effort compares the performance of two Oracle OLTP database solutions. One solution deploys the latest Dell PowerEdge 13th generation R730, which is compared to

[1] http://www.quest.com/benchmark-factory/

second solution using a Dell PowerEdge 11th generation R710. In order to determine the performance benefits of using the PowerEdge R730 for OLTP systems, SAN layer switches and storage are configured identically for both solutions. For more information on the solution design and configuration refer Appendix A.

3 Results and analysis

Dell Quest Benchmark Factory[2] for databases was used to load data that reflects real TPC-C style workload. To determine the database servers performance (R730 vs R710 in this case), databases were stress tested with an incremental concurrent database userload test, starting with 100 users and increasing the user load count in increments of 100 (100,200,300,400 etc). The stress test was continued until the server gained the threshold mark of ~85% CPU utilization; at the same time we ensured that no performance bottlenecks from network and storage infrastructure occurred.

3.1 Concurrent database user load/CPU utilization graph

The graph in Figure 4 shows the CPU utilization of the two database servers nodes – R730 and R710 – against the concurrent database user load. The stress test was maintained with incremental userload until the server gained a threshold level of ~85% CPU utilization. We can conclude from the CPU utilization graph that:

- R710 scaled upto a maximum of 600 concurrent database userload, whereas
- R730 scaled upto a maximum of 1200 concurrent database userload

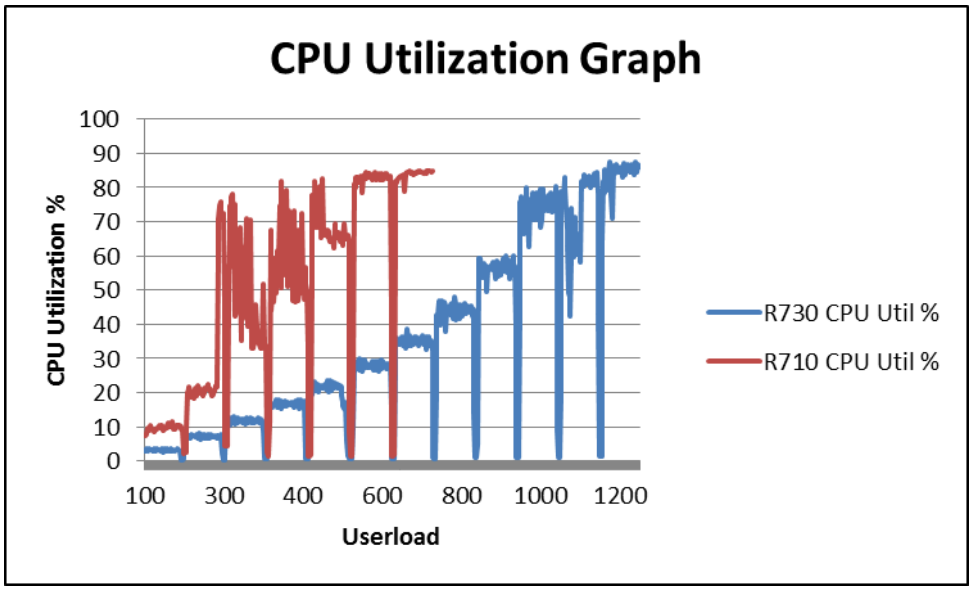

Figure 4 R730 vs. R710: Comparing CPU utilization against database user load

[2] http://www.quest.com/benchmark-factory/

3.2 TPS (Transactions Per Second)

This section talks about the maximum TPS delivered by the R710 and R730 at the maximum database user load supported—600 and 1200 respectively. The TPS delivered by the R710 scaled up until it reached 600 users, after which it declined. For R730, the TPS scaled untill it reached 1200 users.

- R710 delivered a maximum TPS of 2500
- R730 delivered a maximum TPS of 5240

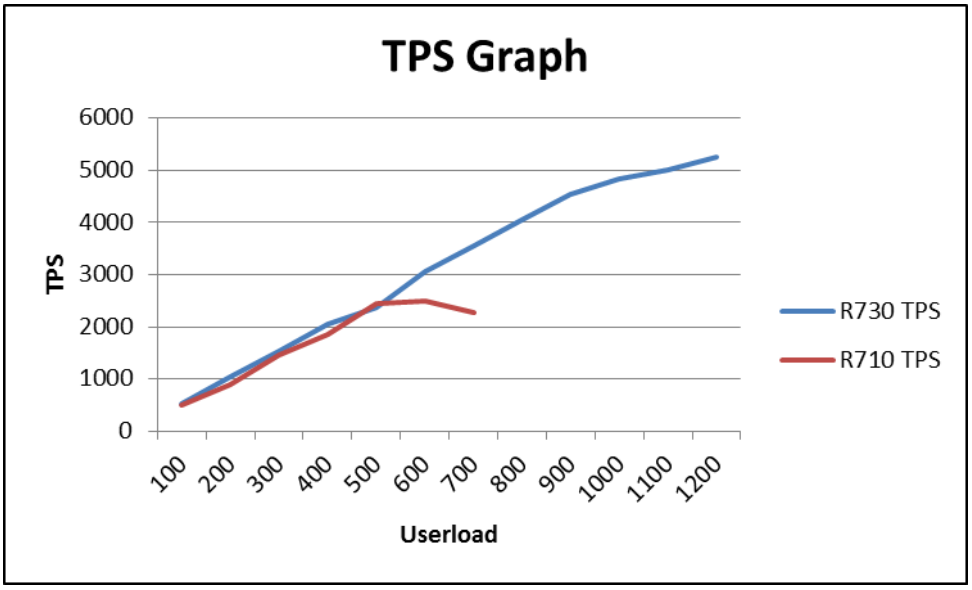

Figure 5 R730 vs. R710: comparing TPS

3.3 Query response time graph

The graph in Figure 6 shows the query response time data for the R730 vs the R710 against database user load. In the case of R710, until the saturated maximum userload of 600 was applied, response time was below 0.2 seconds, after which it started increasing steeply. This can be attributed to the CPU hitting its threshold levels and server was not able to scale up any more.

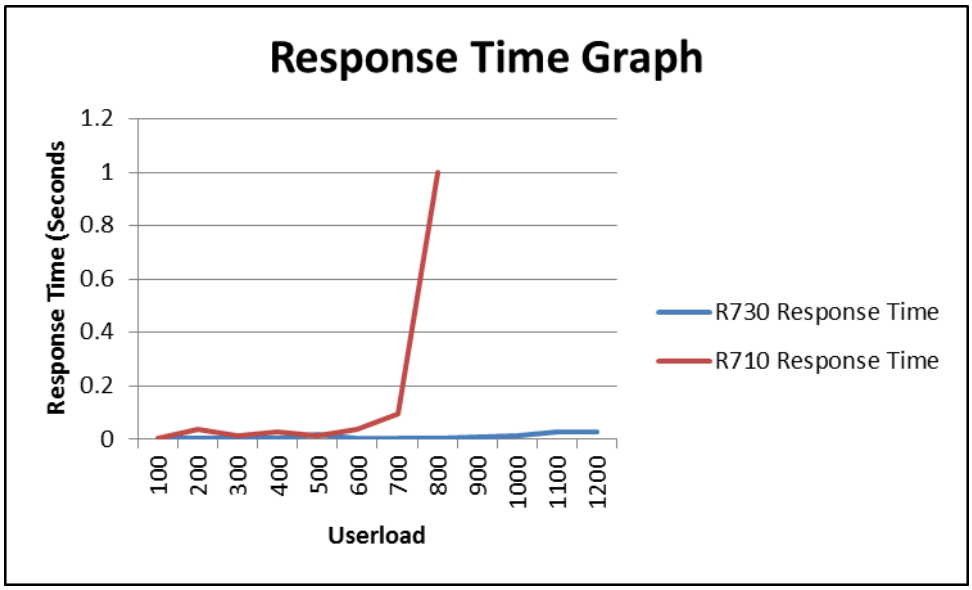

Figure 6 R730 vs. R710: Comparing average query response tIme

4 Conclusion

Below table presents a consolidated view of the results and analysis of the performance study done against R710 and R730.

Table 1 R710/R730 performance study analysis

Server	Max Userload	Max TPS	Response Time
R710	600	2503	0.039
R730	1200	5240	0.005

With Dell infrastructure technology advancing, each new generation offers a wide range of feature set to the customers. This detailed study and the results showcased in this whitepaper make it clear that the Dell PowerEdge R730 is an optimal choice for Oracle OLTP deployments.

Key takeaways

- PowerEdge R730 server is the optimal choice for OLTP deployments
- PowerEdge R730 delivers up to 2 times the TPS compared to the PowerEdge R710
- PowerEdge R730 supports up to 2 times the database userload compared to the PowerEdge R710
- PowerEdge R730 delivers 2:1 consolidation ratio
- PowerEdge R730 helps achieve 50% datacenter rack space savings

A Reference architecture for Oracle

A.1 Hardware and configuration

A.1.1 Physical server hardware

Table 2 Physical server hardware

Component	R730 Configuration	R710 Configuration
CPU	2 x Intel Xeon CPU E5-2680 v3 @ 2.50GHz, Cache: 30MB	2 x Intel Xeon X5675, @ 3.07 GHz Cache: 12M
Memory	256GB (16 x 16GB dimms @2133MHz)	96 GB (12 x 8GB DDR3 @ 1600 MHz)
Operating system	RHEL 6.5 (2.6.32-431.el6.x86_64)	RHEL 6.5 (2.6.32-431.el6.x86_64)
Fibre channel interfaces	Two dual-port QLogic® 16 Gb (QLE2662) Firmware: 02.00.84 Driver: 8.05.00.03.06.5-k2	Two dual-port QLogic 16 Gb (QLE2662) Firmware: 02.00.84 Driver: 8.05.00.03.06.5-k2
Cluster interconnect & client IP interfaces	Intel® 10 Gigabit X540-AT2	Intel® 10 Gigabit X540-AT2
OS disk	4x 146GB 15k SAS Disks in RAID10 Dell PERC H330 Mini RAID Controller (Embedded) Firmware: V4.240.01-3251, Driver: v06.700.06.00-rh1 (megaraid_sas)	4x 146GB 15k SAS Disks in RAID10 Dell PERC SAS6 IR Controller (Embedded) Firmware: 0.25.47.00-IR Driver: v06.700.06.00-rh1(megaraid_sas)

A.1.2 Storage hardware

Storage is provided by a dual-controller Dell Compellent array connected to a single SC220. Each controller is deployed with 2 dual port 16 Fibre Channel HBAs that serve front-end connectivity and IO. Back-end connectivity is provided by a quad-port 6 GB SAS controller per controller.

Table 3 Storage hardware

Component	Description
Storage controller	2x SC8000
Storage firmware revision	Storage Center 6.4
Enclosure	1x SC220 (24x 2.5" Drives)
Write intensive drives	12 x 400 GB WI SSD (SLC)
Read intensive drives	12 x 1.6 TB RI SSD (MLC)
Back-end connectivity	One Quad port 6 Gb SAS per controller
Front-end connectivity	Two dual-port 16 Gb FC HBA per controller

A.1.3 Network switching hardware

Table 4 Switching components used for SAN and LAN connectivity.

Component	Description
Fibre Channel switch	Brocade 6510 Gen5 16 Gb FC switch
Cluster interconnect switch	Dell Networking ™ N4032 switch
Client IP switch	Dell Networking N4032 switch

A.2 Oracle RAC configuration

This architecture was designed for and tested with Oracle Database 12cR1 (12.1.0.1). A 350GB active dataset is used in the test. Dell Quest Benchmark Factory is used as load generator and performance testing tool to simulate real world TPC-C workload.

A.2.1 Oracle initialization parameters: processes, open cursors, transactions and sessions

Certain Oracle initialization parameters such as processes, sessions, open cursors, transactions and db writer processes etc. have to be carefully tuned. Failure to configure these paramaters according to the database environment requirements will result in a negative performance impact.

A.3 Oracle disk group storage layout

The logical storage layout used in this reference architecture is shown in Table 5. Oracle RAC was configured with ASM (external redundancy). A dedicated disk group was created for DATA, FRA, OCR and REDO logs.

Table 5 Oracle disk group storage layout

Disk Group	Storage Profile	Number of Luns	Total Disk Group Size
DATA, temp, index	Flash Optimized	2X250GB	500GB
FRA[3]	Flash Optimized	2X500GB	1TB
REDO Logs	Flash Optimized	2X50GB	100GB
OCR	Flash Optimized	1X15G	15GB

B Server Configuration Profile

Dell iDRAC8 with Lifecycle Controller provides the ability to generate a human readable snapshot of server configuration using the configuration XML feature. This single file contains all BIOS, iDRAC, LC, Network, and Storage settings. With some editing, this file can be captured from one server and applied to other servers, even across different server models.

Configuration XML operations are performed via both RACADM and WS-MAN commands and can be directed to or from network shares (NFS / CIFS).

This whitepaper also provides a server configuration profile (xml) to directly import the Dell Engineered configuration into Dell servers.

You can download the configuration profile here

C Additional resources

Support.dell.com is focused on meeting your needs with proven services and support.

DellTechCenter.com is an IT Community where you can connect with Dell Customers and Dell employees for the purpose of sharing knowledge, best practices, and information about Dell products and installations.

Referenced or recommended Dell publications:

- How to deploy Oracle 12cR1 on RHEL6/Oracle Linux 6: http://en.community.dell.com/techcenter/enterprise-solutions/w/oracle_solutions/4960.how-to-deploy-oracle-12c-release-1-on-rhel6oracle-linux-6.aspx
- Oracle Best Practices on Compellent Storage Center http://www.dellstorage.com/WorkArea/DownloadAsset.aspx?id=3055
- New economies of storage with the Compellent Flash-optimized solutions http://partnerdirect.dell.com/sites/channel/Documents/Dell-Compellent-New-Economies-of-Storage-Flash-Optimized-Solutions-Business-White-Paper.pdf

Referenced or recommended Oracle publications:

- Oracle 12c documentation http://docs.oracle.com/cd/E16655_01/server.121/e17906/chapter2.htm#NEWFT002
- Oracle Database Storage Administrators Guide http://download.oracle.com/docs/cd/B28359_01/server.111/b31107.pdf
- Oracle Database Performance Tuning Guide http://docs.oracle.com/cd/E11882_01/server.112/e16638/iodesign.htm#PFGRF95229