White Paper

Dell Reference Configuration

Deploying an Oracle® Optimized Warehouse With Oracle Database 10*g* and Enterprise Linux 4 on Dell™ PowerEdge™ Servers and Dell/EMC Storage

Abstract

This white paper provides an architectural overview and configuration guidelines for deploying Oracle Optimized Warehouses with Red Hat Enterprise Linux 4 and Oracle Enterprise Linux 4 on Dell PowerEdge servers with Dell/EMC storage. Using the knowledge gained through joint development, testing and support with Oracle and EMC, this Dell Reference Configuration documents "best practices" that can help speed data warehouse solution implementation and help simplify operations, improve performance and availability.

October 2007

October 2007 Rev. A01

Table of Contents

Introduction

The Oracle Optimized Warehouse Initiative provides a validated, balanced approach that combines software, hardware, storage, I/O and networking into a configuration optimized for pre-determined data warehouse requirements. Using their extensive joint field experience and technical knowledge, Dell, EMC and Oracle have developed configurations for data warehouses with varying raw data, average concurrent user and workload complexity. By offering configurations suited for different profiles, operating systems and hardware configurations, customers can select the one that best suits their specific price, performance and compatibility requirements.

The Oracle Optimized Warehouse Initiative on Dell and EMC is an Oracle data warehouse reference configuration that delivers a prescribed database, server and storage mix optimized for data warehouses. This reference configuration is based on a building block approach with the goal of speeding and simplifying the implementation of data warehouses. Each building block is capable of handling 1-2TB of data. Balancing the servers, storage and interconnect aspects of the infrastructure using this approach can help to minimize bottlenecks.

For more information on the Oracle Optimized Warehouse Initiative please visit http://www.oracle.com/solutions/business_intelligence/optimized-warehouse-initiative.html

Dell™ PowerEdge™ servers and Dell/EMC storage systems are ideal choices to deploy highly reliable and sustainable Oracle Optimized Warehouses. This Reference Configuration white paper is intended to help IT professionals design and configure Oracle Optimized Warehouses for data warehousing solutions using Dell servers and storage that apply "best practices" derived from laboratory and real-world experiences. This white paper documents Dell's recommended approach for implementing a tested and validated solution for Oracle 10*g* data warehousing on Dell PowerEdge 9th generation servers, Dell/EMC storage systems and Enterprise Linux 4.

Dell Solutions for Oracle Database 10*g*

Dell Solutions for Oracle Database 10*g* are designed to simplify operations, improve utilization and cost-effectively scale as your needs grow over time. In addition to providing price/performance leading server and storage hardware, Dell Solutions for Oracle Database 10*g* include:

- **Dell Tested and Validated Configurations for Oracle** – in-depth testing of Oracle 10*g* configurations for the most in-demand solutions; documentation and tools that help simplify deployment
- **Integrated Solution Management** – standards-based management of Dell Solutions for Oracle 10*g* that lower operational costs through integrated hardware and software deployment, monitoring and update
- **Oracle Licensing** multiple licensing options that can simplify customer purchase
- **Dell Enterprise Support and Professional Services for Oracle** – offerings for the planning, deployment and maintenance of Dell Solutions for Oracle Database 10*g*

For more information concerning Dell Solutions for Oracle Database 10*g*, please visit www.dell.com/oracle.

Overview of this White Paper

The balance of this white paper will provide the reader with a detailed view of the Dell Reference Configuration for an Oracle Optimized Warehouse based upon Oracle Database 10*g* and Enterprise Linux 4, best practices for configuring the hardware and software components and pointers for obtaining more information.

Architecture Overview - Dell Solution for Oracle Optimized Warehouse

The Dell Reference Configuration for Oracle Optimized Warehouse is intended to validate the following solution components:

- Single-node and multi-node cluster of Dell PowerEdge 2950 dual-core systems.
- Dell/EMC CX3-10 Fibre Channel storage system
- Brocade fibre-channel switches
- Red Hat Enterprise Linux (RHEL) 4 Update 4 x86_64 or Oracle Enterprise Linux (OEL) 4 Update 5 x86_64 (referred to in this whitepaper as "Enterprise Linux 4".)
- Oracle Database 10*g* R2 Standard Edition (10.2.0.3) x86_64.

Single Building Block
A single building block for the Oracle Optimized Warehouse for Dell-EMC consists of one PowerEdge 2950 server, two Brocade fibre-channel (FC) switches, and one Dell-EMC CX3-10 storage system with 37 disks. The two FC switches can support up to nine additional building blocks.

In this configuration, the Oracle database is installed in "Real Application Clusters" mode so that adding new servers or storage systems is seamless. In other words, Oracle Cluster Ready Services (CRS) is installed on the single-instance. The architecture is made of the following components:

- Client systems that will access data stored within the Oracle single-instance RAC database
- Client-server network made up of network controllers, cables and switches
- Single node Dell PowerEdge 2950 server running on Enterprise Linux 4
- Gigabit Ethernet switches for cluster interconnect network
- Brocade Fibre Channel switches for a SAN environment
- Dell/EMC CX3-10 Fibre channel storage system with 37 disks

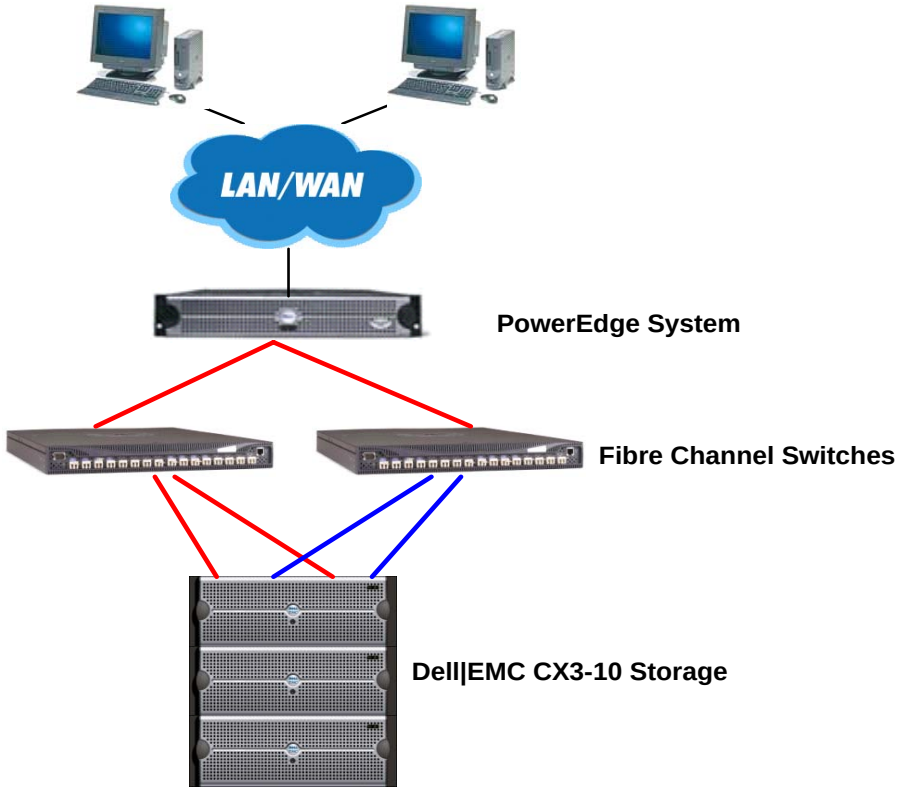

Figure 1 – Architectural Overview of Single Block Oracle Optimized Warehouse

Multiple Building Blocks

Depending on the performance and capacity requirements of a specific implementation, multiple building blocks may be needed. Each additional building block will require an additional PowerEdge 2950 server and an additional Dell/EMC CX3-10 with 37 disks. Additional FC switches are not required unless the number of building blocks exceeds 10.

This architecture illustrates two building blocks with the following components:

- Client systems that will access data stored within the Oracle two node RAC database
- Client-server network made up of network controllers, cables and switches
- Two node Dell PowerEdge 2950 cluster running on Enterprise Linux 4
- Gigabit Ethernet switches for cluster interconnect network
- Brocade Fibre Channel switches for a SAN environment
- Two Dell/EMC CX3-10 Fibre channel storage systems with 74 disks

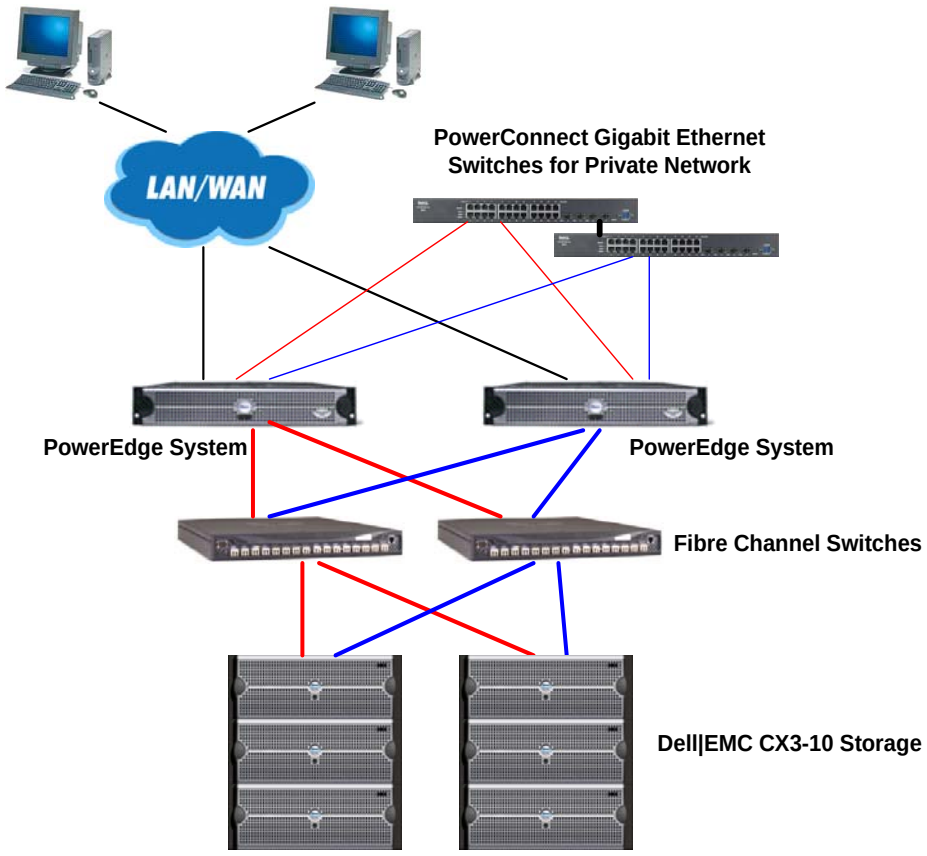

Figure 2 - Architectural Overview of Multi-block Oracle Optimized Warehouse

Dell PowerEdge servers are designed to deliver the highest performance for critical enterprise applications such as data warehousing, messaging, web services and infrastructure applications. As proprietary systems are increasingly replaced by industry-standard systems, applications like databases, high performance computing clusters and messaging systems can take advantage of the performance and scalability of the PowerEdge servers. Combined with Dell storage systems, customers can easily deploy these PowerEdge servers as building blocks of a scalable enterprise, consolidating and virtualizing both the computing resources as well as the storage resources.

The Dell/EMC storage subsystem delivers advanced storage capabilities including simple management tools, continuous data availability and integrity, industry-leading price/performance, data mobility, and scalability between multiple storage tiers. The Dell/EMC storage subsystem is offered in various models, ranging from affordable entry-level solutions to high-performance, maximum-capacity configurations for your most demanding requirements. All Dell/EMC CX3 series arrays support advanced software including local replication for backup/restore, remote replication for disaster recovery and data mobility. The Dell/EMC is architected with two storage processors to guard against a single point of failure.

Hardware Configuration – Dell Solution for Oracle Optimized Warehouse

Storage Configuration

Configuring Dell/EMC CX3 Fibre Channel Storage Connections with Dual HBAs and Dual Fibre Channel Switches

Figure 3 illustrates the fiber cabling of the two-building-block PowerEdge cluster hosting Oracle Database 10*g* and the CX3 storage array where the data resides. As mentioned in the Architectural Overview section above, each CX3 storage array has two storage processors (SP), called SPA and SPB, which can access all of the disks in the system and provide the physical storage capacity for the Oracle RAC database. Before data can be stored, the CX3 physical disks must be configured into components, known as RAID groups and LUNs. A RAID group is a set of physical disks that are logically grouped together. Each RAID group can be divided into one or more LUNs, which are logical entities that the server uses to store data. The RAID level of a RAID group is determined when binding the first LUN within the RAID group. It is recommended binding one LUN per RAID group for database workloads to avoid disk spindle contention.[1] For details on LUN configuration, please refer to the "Configuring LUNs" section below.

In the CX3 array, the LUNs are assigned to and accessed by the Oracle Database 10*g* cluster nodes directly through one storage processor. In the event of a storage processor port failure, traffic will be routed to another port on the same SP if the host is connected to more than one SP port and the EMC PowerPath multi path software is used. In the event of a storage processor failure, LUNs on the failed processor will trespass to the remaining storage processor. Both events could result in interrupted service unless multiple I/O paths are configured between the Oracle 10*g* RAC database hosts and the CX3 array. Therefore, it is crucial to eliminate any single point of failures within the I/O path.

At the interconnect level, it is recommended that each node of the Oracle 10*g* RAC database cluster have two HBAs with independent paths to both storage processors. With the EMC PowerPath software installed on the cluster node, I/O can be balanced across HBAs as well. It is a best practice to hard set both the HBA ports and the switch ports to operate at 4 GB/s throughput to allow proper I/O path auto recovery. It is also recommended that two Fibre Channel switches are used because in the event of a switch failure in a single Fibre Channel switch fabric environment, all hosts will lose access to the storage until the switch is physically replaced and the configuration restored. In addition, since I/O blocks are now split across multiple switches, the I/O throughput pushed to the FC backend will not be bottlenecked by the switches.

[1] "Designing and Optimizing Dell/EMC SAN Configurations Part 1", Arrian Mehis and Scott Stanford, Dell Power Solutions, June 2004. http://www.dell.com/downloads/global/power/ps2q04-022.pdf

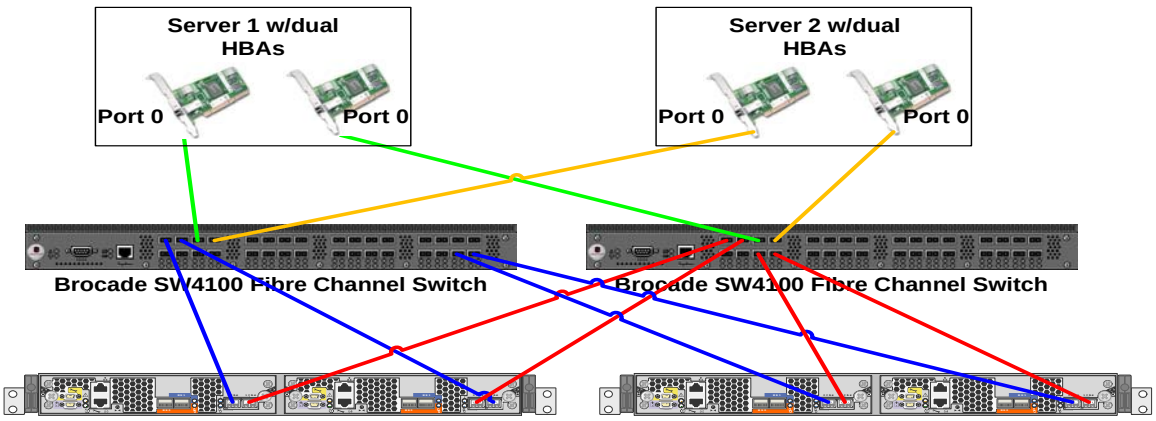

Figure 3 - Cabling a Two-Block Oracle Optimized Warehouse

Configuring LUNs

A LUN is a logical unit of physical disks presented to the host connected to the CX3 storage. The storage for an Oracle Database 10*g* data warehouse can be divided into the following three areas:

- The first area is for the Oracle Cluster Registry (OCR), the Clusterware Cluster Synchronization Services (CSS) Voting Disk, and the Server Parameter File (SPFILE) for the Oracle Automatic Storage Management (ASM) instances. The OCR stores the details of the cluster configuration, including the names and current status of the database, associated instances, services, and node applications, such as the listener process. The CSS Voting Disk is used to determine the nodes that are currently available within the cluster. The SPFILE for ASM instances is a binary file which stores the ASM instance parameter settings.
- The second area is for database data that are stored in the Oracle database physical files including datafiles, online redo log files, control files, SPFILE for the database instances, and temp files for the temporary tablespaces.
- The third area is for the Oracle Flash Recovery Area which is a storage location for all recovery-related files. The disk based database backup files are stored in the Flash Recovery Area. The Flash Recovery Area is also the default location for all archived redo log files.

It is a best practice to separate the above three storage areas onto their own LUNs on separate Disk Groups / RAID Groups. The separation can enable better I/O performance by ensuring these files do not share the same physical disks. Table 1 shows a sample LUN configuration with three LUNs for each of the three storage areas described above.

LUN	Minimum Size	RAID	Number of Partitions	Used For	OS Mapping
First LUN	1024 MB	10, or 1	Three of 300 MB each	Voting disk, Oracle Cluster Registry (OCR), and SPFILE for ASM instances	Three raw devices for Voting Disk, OCR, and SPFILE
Second LUN	Larger than the size of your database	10, or 5 for read-only	One	Data	ASM disk group DATABASEDG
Third LUN	Minimum twice the size of your second LUN/Virtual Disk	10, or 5 for read-only	One	Flash Recovery Area	ASM disk group FLASHBACKDG

Table 1 - LUNs for the Cluster Storage Groups / RAID Groups

Figure 4 illustrates a sample RAID group and LUN configuration on a Dell/EMC CX3-10 storage array with three Disk Array Enclosures (DAEs). There are separate partitions for the three storage areas described in Table 1. Spindles 0 through 4 in the DAE 0 of the CX3-10 contain the operating system for the storage. These spindles are also used during power outage to store the system cache data. It is not recommended to use the operating system spindles for as data, or Flash Recovery Area, or OCR / Voting Disk / SPFILE drives. DAE 1 and DAE 2 contain the ASM diskgroups for the database files and flash recovery area.

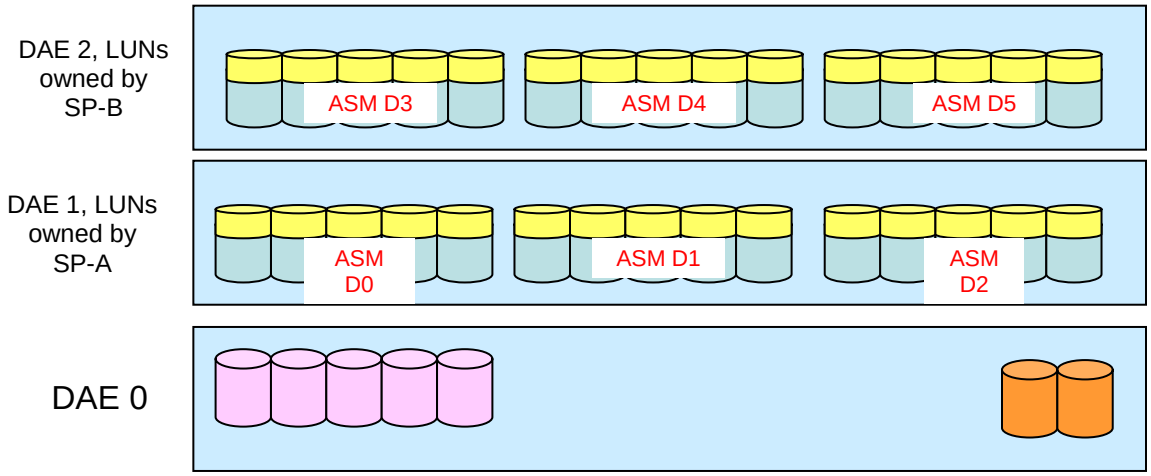

Figure 4 - Disk Array Layout

Figure 5 below provides the RAID configurations for the ASM diskgroups that make up the data disks and flash recovery area. The vault drives house to the operating system of the storage system.

Figure 5 RAID Configurations for each ASM Diskgroup

Each virtual disk / LUN created in storage will be presented to all the Oracle 10*g* RAC hosts and configured at the OS level. For details on the shared storage configuration at the OS level, please refer to the "Configuring Shared Storage for the Oracle Clusterware using the RAW Devices Interface" section and the "Configuring Shared Storage for the Database using the ASM Library Driver" section below.

Server Configuration

Configuring Fully Redundant Ethernet Interconnects

Each Oracle Database 10*g* RAC server needs at least two network interface cards (NICs), one for the external interface and one for the private interconnect network. The servers in an Oracle RAC are bound together using cluster management software called Oracle Clusterware, which enables the servers to appear as though they are a single server. Servers in the cluster communicate with each other using a dedicated private network also known as the cluster interconnect. One of the servers in the Oracle RAC cluster is assigned as the master node. In the event of a interconnect NIC failure in a single interconnect NIC environment, the server lost communication to the master node, and the master node will initiate recovery of the failed database instance on the server. In the event of a network switch failure in a single private network switch environment, a scenario will result equivalent to the failure of every single node in the cluster except for the designated master node. The master node will then proceed to recover all of the failed instances in the cluster before providing a service from a single node which will result in a significant reduction in the level of service available. Therefore, it is recommended the implementation of a fully redundant interconnect network configuration, with redundant private NICs on each server and redundant private network switches.[2]

Figure 5 illustrates the CAT 5E/6 Ethernet cabling of a fully redundant interconnect network configuration of a two-node PowerEdge Oracle RAC cluster, with two private NICs on each server, and two private network switches. For this type of redundancy to operate successfully, it requires the implementation of the Link Aggregation Group, where one or more links are provided between the switches themselves.

To implement a fully redundant interconnect configuration requires the implementation of NIC teaming software at the operating system level. This software operates at the network driver level to provide two physical network interfaces to operate underneath a single IP address.[3] For details on configuring NIC teaming, please refer to the "Configuring the Private NIC teaming" section below.

[2] Dyke and Shaw, op. cit.
[3] Dyke and Shaw, op. cit.

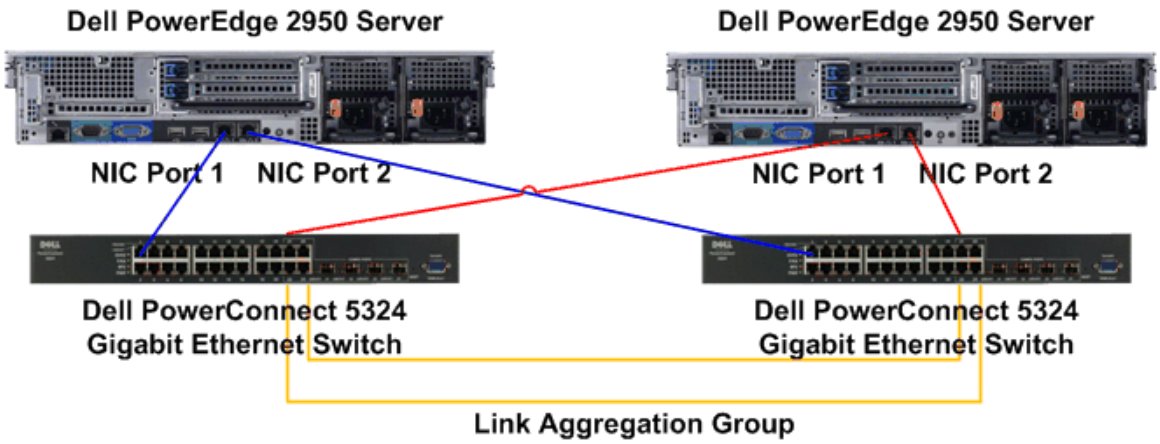

Figure 5 - Ethernet Cabling a Fully Redundant Private Interconnect Network

Configuring Dual HBAs for Dell/EMC CX3 storage

As illustrated in Figure 3, it is recommended that two HBAs be installed on each of the PowerEdge servers hosting the Oracle 10*g* RAC database because in the event of an HBA failure in a single HBA fabric environment, the host will lose access to the storage until the failed HBA is physically replaced. Using dual HBAs provides redundant links to the CX3 storage array.

Configuring Dual NICs for Private Network

As illustrated in Figure 5, it is recommended that two private NICs be installed on each of the PowerEdge servers hosting the Oracle 10*g* RAC database to provide redundant private network links. In the event of a NIC failure in a single private NIC environment, Oracle Clusterware will remove the node from the cluster.

Software Configuration – Dell Solution for Oracle Optimized Warehouse

Operating System Configuration

Configuring the Private NIC Teaming

As mentioned in the Section "Configuring Fully Redundant Ethernet Interconnects" above, it is recommended to install two physical private NICs on each of the Oracle 10*g* RAC cluster servers, to help guard against private network communication failures. In addition to installing the two NICs, it is required to use NIC teaming software to bond the two private network interfaces together to operate under a single IP address, providing failover functionality. If a failure occurs, affecting one of the NIC interfaces – examples include switch port failure, cable disconnection, or failures of the NIC itself – network traffic is routed to the remaining operable NIC interface. Failover occurs transparently to the Oracle RAC database with no network communication interruption or changes to the private IP address.

Configuring the Same Public Network Interface Name on All Nodes

It is important to ensure that all nodes within an Oracle 10*g* RAC cluster have the same network interface name for the public interface. For example, if "eth0" is configured as the public interface on the first node, then "eth0" should also be selected as the public interface on all of the other nodes. This is required for the correct operation of the Virtual IP (VIP) addresses configured during the Oracle Clusterware software installation.[4]

Configuring SSH and RSH

During the installation of Oracle 10*g* RAC software, the Oracle Universal Installer (OUI) is initiated on one of the nodes of the Oracle RAC cluster. OUI operates by copying files to and running commands on the other servers in the cluster. In order to allow OUI to do that, the secure shell (SSH) and remote shell (RSH) must be configured, so no prompts or warnings are received when connecting between hosts via SSH or RSH as the oracle user. To prevent unauthorized users from accessing the systems, it is recommended that RSH be disabled after the Oracle software installation.

Configuring Shared Storage for the Database using the ASM Library Driver

Oracle Automatic Storage Management (ASM) is a feature of Oracle Database 10*g* that provides the database administrator (DBA) with a simple storage management interface that is consistent across all server and storage platforms. ASM virtualizes the database storage into disk groups. ASM distributes data evenly across all disks within a disk group to optimize performance and utilization. ASM enables the DBA to change the storage configuration without having to take the database offline. ASM automatically rebalances files across the disk group after disks have been added or dropped.[5]

As discussed in the Section "Configuring LUNs" above, two LUNs are created for the data storage area, and the Flash Recovery Area, respectively. It is recommended that these two LUNs be configured as ASM disks to benefit from the capabilities of ASM. During ASM Instance creation, it is necessary to set the following ASM parameters:

[4] Dyke and Shaw, op. cit.
[5] "Oracle Database 10*g* – Automatic Storage Management Overview", Oracle TechNet. http://www.oracle.com/technology/products/manageability/database/pdf/asmov.pdf

Name	Value
_asm_ausize	8388608
_asm_stripesize	524288

Table 2 - ASM Configuration Parameters

Configuring Database Parameters for Data Warehousing Applications

During database creation using DBCA, the following parameters need to be set:

Database Parameters	
Name	**Value**
control_files	+DATA/emcdw/controlfile/cureent.299.628107111
db_block_size	8192
db_cache_size	1024M
db_file_multiblock_read_count	128
db_files	5000
db_name	EMCDW
log_archive_dest	+FRA
log_buffer	18321408
Max_dump_file_size	Unlimited
parallel_adaptive_multi_user	FALSE
parallel_execution_message_size	16384
parallel_max_servers	256
parallel_min_servers	64
Pga_aggregate_target	2560M
Processes	500
sga_target	0
shared_pool_size	500M

Table 3 - Database Configuration Parameters

Database tablespaces sizes	
Tablespace name	**Size**
System	32700M
Sysaux	32700M
Undo1	32700M
Undo 2	32700M
Undo 3	32700M
Undo 4	32700M
TEMP	800 GB
TS_DATA	1.5 TB
Redo logs	32700M

Table 4 - Tablespace Configuration Parameters

Adding Building Blocks to the Dell Solution for Oracle Optimized Warehouse

Hardware Configuration of Shared Storage

A single Dell building block for an Oracle Optimized Warehouse has two fibre channel switches that can support up to nine additional building blocks. In order to provide host side multi-pathing, balance the I/O load placed on the CX3 storage processors and greater reliability in case of HBA, switch or storage processor faults, it is recommended that the initial building block include two fibre channel switches. Figure 6 shows the cabling of a single block with shared storage in a high-availability configuration:

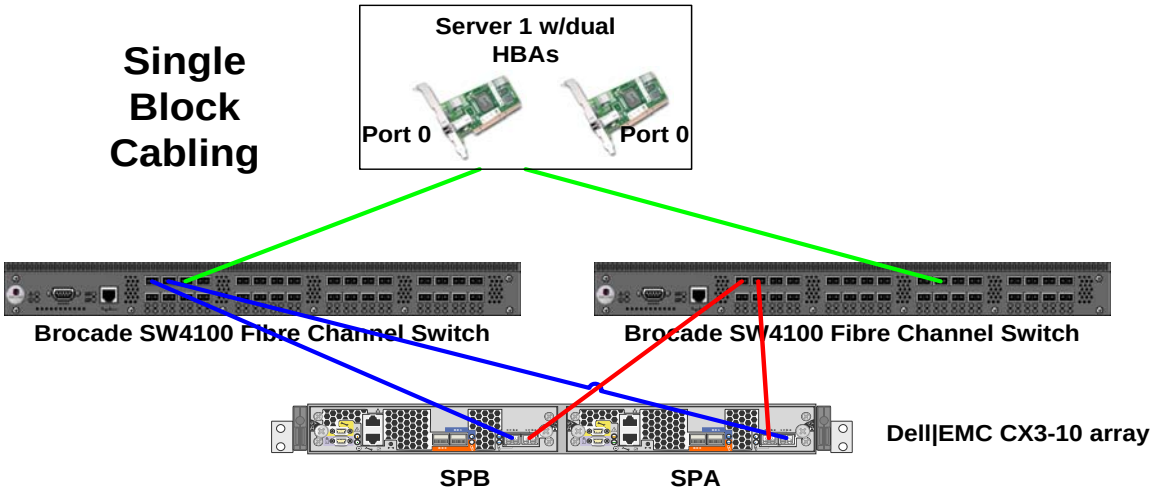

Figure 6 - Single Block SAN Configuration

Figure 7 below highlights a two-block cabling configuration. In this configuration, all the HBAs have complete access to all the shared storage.

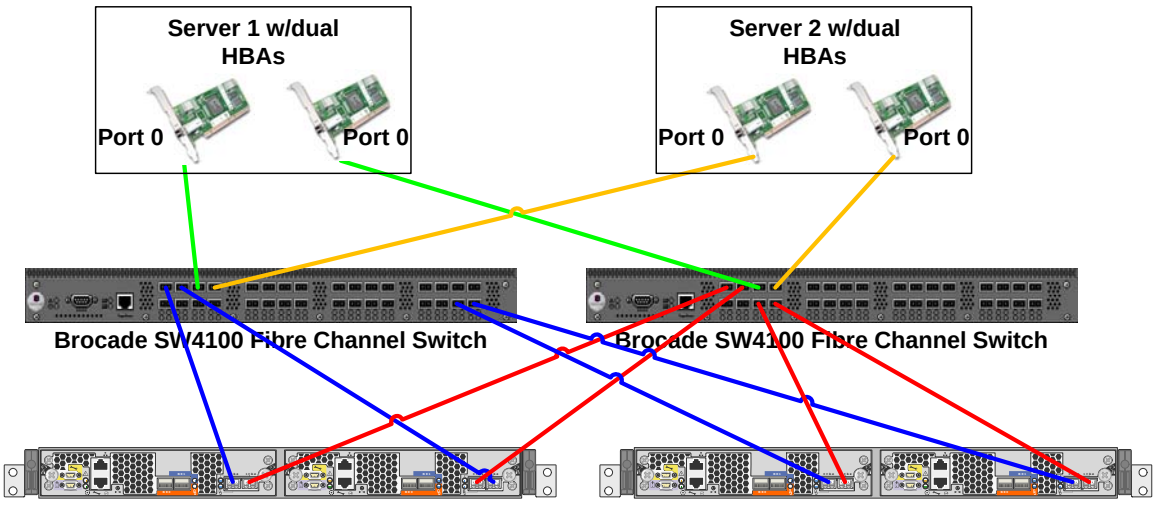

Figure 7 - Two block SAN Configuration

Software Configuration of Shared Storage

After the new storage is added to the cluster, it needs to be configured so that all the host servers have access to the storage. In order to add shared storage to the existing host servers in the cluster, steps 1 and 2 need to be taken on each host server. After that, complete steps 3 and 4 on the primary host server.

1. **Device/Disk visibility on the servers**
   In order for the disks of the new storage system to be visible on the already running servers, the SCSI bus needs to be dynamically rescanned to pick up the new devices. For Qlogic HBAs, this can be achieved by running the Qlogic FC HBA LUN Scan Utility, which can be downloaded from the Linux Utilities section of the EMC Approved Software: Symmetrix, CLARiiON and Celerra section. For Emulex HBAs, this can be achieved by running the Emulex LUN Scan Script utility, which can be downloaded from the Tools for the Emulex Driver for Linux page. Upon running the utility, the LUNs from the new storage system will now be visible.

2. **EMC Multi-pathing configuration**
   To add new emcpower devices once the LUNs are visible, run the following powermt command:

   powermt config # this will automatically scan and multi-path all new added LUNs

3. **Adding new EMC devices to existing Oracle ASM diskgroups**
   In order to add the new EMC power devices, invoke sqlplus from the command line, and run the following command:

   Sqlplus> ALTER DISKGROUP DATABASEDG ADD Disk '/dev/emcpowerc1', '/dev/emcpowerd1';
   where '/dev/emcpowerc1' and '/dev/emcpowerd1' are the newly added EMC power devices

4. **Re-balancing ASM data within a diskgroup**
   Once new LUNs are added to a diskgroup, it is necessary to rebalance the data within the diskgroup. This can be done by executing the following sqlplus command:

   Sqlplus> alter diskgroup databasedg rebalance power 11;

Configuring host servers

In order to add new nodes to the current configuration, follow the instructions detailed in the "Adding and Removing Nodes" section of the appropriate Deployment Guide for Oracle Database 10*g* R2 on Enterprise Linux 4 as noted above. Performing the steps documented in the following sections is recommended:

- Adding a New Node to the Network Layer
- Configure Shared Storage on the New Node Using ASM
- Adding a New Node to the Clusterware Layer
- Adding a New Node to the Database Layer
- Reconfiguring the Listener
- Adding a New Node to the Database Instance Layer

Test Results - Dell Solution for Oracle Optimized Warehouse

As part of the joint effort among Dell, Oracle and EMC, the joint engineering teams conducted benchmark tests, on one, two, three and four building block configurations using I/O intensive stress tools that simulated Oracle database activity. Each building block in the tests housed 1 TB of raw data.

Table 2 illustrates the I/O throughput in terms of megabytes per second, as well as the hardware and software configurations of each of the building blocks. The results of these tests demonstrate the scalability of Dell PowerEdge servers running Oracle 10*g* RAC and ASM with Enterprise Linux 4.

Number of Building Blocks	1-Block Configuration	2-Block Configuration	3-Block Configuration	4-Block Configuration
I/O Throughput (MBytes / Sec)	620	1,174	1,695	2,213
Raw (atomic) user data	1TB	2TB	3TB	4TB
Architecture	x86-64			
# Servers / Model	1x PowerEdge 2950	2x PowerEdge 2950	3x PowerEdge 2950	4x PowerEdge 2950
CPU Type	Intel Xeon 5160 2.66/3GHz, 4M Cache, Dual-Core			
# CPUs / Node	2			
Total CPU Cores	4	8	12	16
Memory / Node	8GB			
Total Memory	16GB	32GB	48GB	96GB
FC Cards / Node	2	2	2	2
Node Interconnect	N/A	2 Dell Gigabit Ethernet Switches	2 Dell Gigabit Ethernet Switches	2 Dell Gigabit Ethernet Switches
# SAN Switches	2 x Brocade SW4100/DS4100			
# Ports / Switch	32 (Sufficient for up to 10 building blocks)			
# Storage Disk Array / Model	1 x Dell/EMC CX3-10c	2 x Dell/EMC CX3-10c	3 x Dell/EMC CX3-10c	4 x Dell/EMC CX3-10c
# Disk Array Enclosures	3	6	9	12
# Disk Controllers	2	4	6	8
# of Disks	37	74	111	148
Disk Speed	15K RPM			
Disk Capacity (Raw / Usable)	4.4 / 3 TB	8.8 / 7 TB	13.1 / 10 TB	17.5 / 13 TB
HBA Model	Emulex LP1150e			
# of HBA	2	4	6	8
OS Version	OEL 4 Update 5 or RHEL 4 Update 4			
Database Version	Oracle 10*g* R2 Enterprise Edition	Oracle 10*g* R2 Enterprise Edition with RAC		
Database Version	10.2.0.3			

Table 5 - Building Block Capacities and Throughput

Configuration Deliverables List - Dell Solution for Oracle Optimized Warehouse

This section contains the Solution Deliverables List (SDL) for the Dell Solution for Oracle Optimized Warehouse Initiative. It contains a detailed listing of server and storage hardware configurations, firmware, driver, OS and database versions.

Minimum Hardware/Software Requirements (For details, see below)			
	Validated Component(s)	**Minimum Single Building Block Configuration**	**Minimum Multiple Building Block Configuration**
PowerEdge Nodes	2950	1	2
Memory	All valid PowerEdge 2950 memory configurations	16GB (per node)	16GB (per node)
Dell / EMC FC Storage Array	CX3-10c	1	2
Fibre Channel Switch	Brocade DS4100, SW4100	2	N/A (For Direct Attached) 2 (8 port switches for 2-4 nodes, and 16 port switches for 5 or above nodes) [1]
HBAs	LP 1150e	2 single-port HBAs (Per Node)[2]	2 single-port HBAs (Per Node)[2]
Ethernet Ports	Intel or Broadcom Gigabit NICs	3 (Per Node) [3]	3 (Per Node) [3]
Ethernet Switches (For Private Interconnect)	Gigabit-only Switches	N/A	2
Raid Controllers (Used for internal storage only)	PERC 5/e, PERC 5/i	1 (Per Node)	1 (Per Node)
Internal Drive	All valid PowerEdge 2950 internal storage configurations	73 Gig/node	73 Gig/node
Oracle Software & Licenses	Oracle 10*g* R2 10.2.0.1 Enterprise Edition (Base) + Oracle Patchset 10.2.0.3	Enterprise Edition	RAC
Operating System	Red Hat Enterprise Linux 4 Update 4 + errata OR Oracle Enterprise Linux 4.5		
Recommended Support Contract	Dell Gold or Platinum Plus Enterprise Support		

Notes:
1: Assumes there is no requirement for additional ports for other functions on the switch.
2: Two single port HBAs are recommended.
3: Assumes one NIC for public network, and two bonded NICs for private network.

Reference Configuration Details

Validated Servers					
PowerEdge Servers	**Model**	**BIOS**	**ESM/BMC Firmware**	**Internal SCSI RAID Firmware / Driver**	**Notes**
	PE2950	1.4.0 or higher	v1.33 or higher	PERC 5i –5.1.1-0040/00.00.02.03-RH1	Intel Dual Core Processors only
PERC 5/e, 5/i	Firmware = v5.1.1-0040; Driver = version 00.00.02.03-RH1				
Network Interconnect					
Ethernet Switches	All Dell Gigabit Ethernet Switches				
Intel NIC Drivers (100)	Driver = (e1000) 7.0.33-k2 NAPI				
Intel NIC Drivers (1000MT)	Driver = (e1000) 7.0.33-k2 NAPI				
Broadcom NIC Drivers (5708)	Driver = (bnx2)1.4.38				
NIC Bonding	NIC Kernel bonding for QU4 Version = 2.6.3				
Fibre Channel Host Bus Adapter (HBA)					
Emulex HBA LP1150e	Bios= 1.70.a3; Firmware=2.70a5; Driver = lpfc, version 8.0.16.27				
Fibre Channel Switches					
Brocade Fibre Channel Switch (SW4100, DS4100)	Firmware = v5.2.1 or higher				
Fibre Channel Storage					
Dell / EMC	**Validated Storage Arrays (with Software)**				
	Storage Subsystem: CX3-10c(Release 24 or later) **Storage Subsystem Disk Drives**: 15K RPM FC Drives **Drive Capacity** : 146GB **Data transport speed**: 4 GBit/s				
Database Software					
Oracle	10*g* R2 10.2.0.1 Enterprise Edition (Base) + 10.2.0.3 (PatchSet)				
ASMLib	oracleasm-2.6.9-42.0.8 (for RHEL4.4) oracleasm-2.6.9-55 (for OEL4.5)				
Operating system	RHEL4.4 (kernel-2.6.9-42.0.8.Elsmp) OEL4.5 (kernel-2.6.9-55.ELsmp)				
DKMS	Version 2.0.13-1				
EMC PowerPath	5.0.0-156 (available at www.emc.com)				

Conclusion

Dell Solutions for Oracle Database 10*g* are designed to help simplify operations, improve utilization and cost-effectively scale as your needs grow over time. This reference configuration white paper provides a blueprint for setting up an Oracle Optimized Warehouse for Dell/EMC with Oracle Database 10*g* R2 and Enterprise Linux 4 on Dell PowerEdge servers and Dell/EMC storage arrays.

The best practices described here are intended to help achieve optimal performance of Oracle Database 10g. To learn more about deploying Oracle Database 10*g* on PowerEdge servers and Dell storage, please visit www.dell.com/oracle or contact your Dell representative for up to date information on Dell servers, storage and services for Oracle 10*g* solutions.

Tables and Figures Index

References

1. "Designing and Optimizing Dell/EMC SAN Configurations Part 1", Arrian Mehis and Scott Stanford, Dell Power Solutions, June 2004.
   http://www.dell.com/downloads/global/power/ps2q04-022.pdf

2. "Pro Oracle Database 10*g* RAC on Linux", Julian Dyke and Steve Shaw, Apress, 2006.

3. "Oracle Database 10*g* – Automatic Storage Management Overview", Oracle TechNet.
   http://www.oracle.com/technology/products/manageability/database/pdf/asmov.pdf

4. "Benchmark Factory for Databases", Quest Software.
   http://www.quest.com/Quest_Site_Assets/PDF/Benchmark_Factory_5_TPCH.pdf

5. "Scaling out your Data Warehousing building blocks on Dell and EMC", Bob Ng and Maria Colgan.
   www.oracle.com/solutions/business_intelligence/docs/OOW_2006_WP_S282241_v2.pdf