

## Dell Reference Configuration

# Deploying Oracle® Database 10g R2 Standard Edition Real Application Clusters with Red Hat® Enterprise Linux® 4 Advanced Server x86\_64 on Dell™ PowerEdge™ Servers and Dell PowerVault™ MD3000i iSCSI Storage

### Abstract

This white paper provides an architectural overview and configuration guidelines for deploying a two-node Oracle 10g R2 Standard Edition Real Application Clusters (RAC) database with Red Hat Enterprise Linux 4 Advanced Server x86\_64 on Dell PowerEdge servers with Dell PowerVault MD3000i iSCSI storage. Using the knowledge gained through joint development, testing and support with Oracle, this Dell Reference Configuration documents “best practices” that can help speed Oracle solution implementation and help simplify operations, improve performance and availability.

December 2007

THIS WHITE PAPER 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.

© 2007 Dell Inc. All rights reserved.

Reproduction in any manner whatsoever without the written permission of Dell Inc. is strictly forbidden.

Trademarks used in this text: *Dell*, the *DELL* logo, *PowerEdge* and *PowerVault* are trademarks of Dell Inc.; *Intel* is a registered trademark of Intel Corporation; *EMC* is a registered trademark of EMC Corporation; *Oracle* is a registered trademark of Oracle Corporation and/or its affiliates.

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 Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

December 2007 Rev. A01

# Table of Contents

|   |           |
|---|-----------|
| <b>ABSTRACT .....</b>   | <b>1</b>  |
| <b>INTRODUCTION .....</b>   | <b>4</b>  |
| <b>DELL SOLUTIONS FOR ORACLE DATABASE 10G .....</b>   | <b>4</b>  |
| <b>OVERVIEW OF THIS WHITE PAPER .....</b>   | <b>4</b>  |
| <b>ARCHITECTURE OVERVIEW - DELL SOLUTION FOR ORACLE 10G ON DELL<br/>POWERSHARE MD3000I ISCSI STORAGE .....</b>            | <b>5</b>  |
| <b>HARDWARE CONFIGURATION - DELL SOLUTION FOR ORACLE 10G ON DELL<br/>POWERSHARE MD3000I ISCSI STORAGE .....</b>           | <b>7</b>  |
| <b>STORAGE CONFIGURATION.....</b>   | <b>7</b>  |
| <i>Configuring PowerVault MD3000i iSCSI Storage Connections with Dual Gigabit Ethernet NIC for<br/>iSCSI Network.....</i> | <i>7</i>  |
| <i>Configuring Virtual Disks.....</i>   | <i>9</i>  |
| <i>Configuring Challenge Handshake Authentication Protocol (CHAP) .....</i>   | <i>12</i> |
| <b>GIGABIT ETHERNET SWITCH CONFIGURATION .....</b>  | <b>12</b> |
| <i>Configuring Virtual Local Area Network (VLAN).....</i>   | <i>12</i> |
| <b>SERVER CONFIGURATION .....</b>   | <b>12</b> |
| <i>Configuring Fully Redundant Ethernet Interconnects.....</i>  | <i>12</i> |
| <i>Configuring Dual NIC Ports for MD3000i iSCSI Storage Network.....</i>  | <i>13</i> |
| <i>Configuring Dual NIC Ports for Private Network.....</i>  | <i>13</i> |
| <b>SOFTWARE CONFIGURATION - DELL SOLUTION FOR ORACLE 10G ON DELL<br/>POWERSHARE MD3000I ISCSI STORAGE .....</b>           | <b>14</b> |
| <b>OPERATING SYSTEM CONFIGURATION .....</b>   | <b>14</b> |
| <i>Configuring the Private NIC Teaming.....</i>   | <i>14</i> |
| <i>Configuring the Same Public Network Interface Name on All Nodes.....</i>   | <i>14</i> |
| <i>Configuring SSH and RSH.....</i>   | <i>14</i> |
| <i>Configuring Shared Storage for Oracle Clusterware using the RAW Devices Interface.....</i>                             | <i>14</i> |
| <i>Configuring Shared Storage for the Database using the ASM Library Driver.....</i>                                      | <i>15</i> |
| <i>Enabling Full Network Startup.....</i>   | <i>15</i> |
| <b>ORACLE 10G R2 CONFIGURATION .....</b>  | <b>15</b> |
| <i>Reconfiguring Clusterware Parameter Disktimeout to 300.....</i>  | <i>15</i> |
| <b>CONFIGURATION DELIVERABLES LIST - DELL SOLUTION FOR ORACLE 10G ON DELL<br/>POWERSHARE MD3000I ISCSI STORAGE .....</b>  | <b>16</b> |
| <b>CONCLUSION.....</b>  | <b>18</b> |
| <b>TABLES AND FIGURES INDEX .....</b>   | <b>19</b> |
| <b>REFERENCES .....</b>   | <b>19</b> |

# Introduction

Dell™ PowerEdge™ servers and Dell PowerVault™ MD3000i iSCSI SAN storage systems are ideal choices for deploying highly reliable and sustainable Oracle 10g databases. This Reference Configuration white paper is intended to help IT professionals design and configure Oracle 10g database 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 10g database on Dell PowerEdge 9th generation servers, Dell PowerVault MD3000i iSCSI SAN storage systems and Red Hat® Enterprise Linux® 4 (RHEL 4).

## Dell Solutions for Oracle Database 10g

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

- **Dell Configurations for Oracle** – in-depth testing of Oracle 10g configurations for high-demand solutions; documentation and tools that help simplify deployment
- **Integrated Solution Management** – standards-based management of Dell Solutions for Oracle 10g that can 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 10g

For more information concerning Dell Solutions for Oracle 10g Database, please visit [www.dell.com/oracle](http://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 Oracle 10g with Dell PowerVault MD3000i iSCSI SAN storage, best practices for configuring the hardware and software components, and pointers for obtaining more information.

# Architecture Overview - Dell Solution for Oracle 10g on Dell PowerVault MD3000i iSCSI Storage

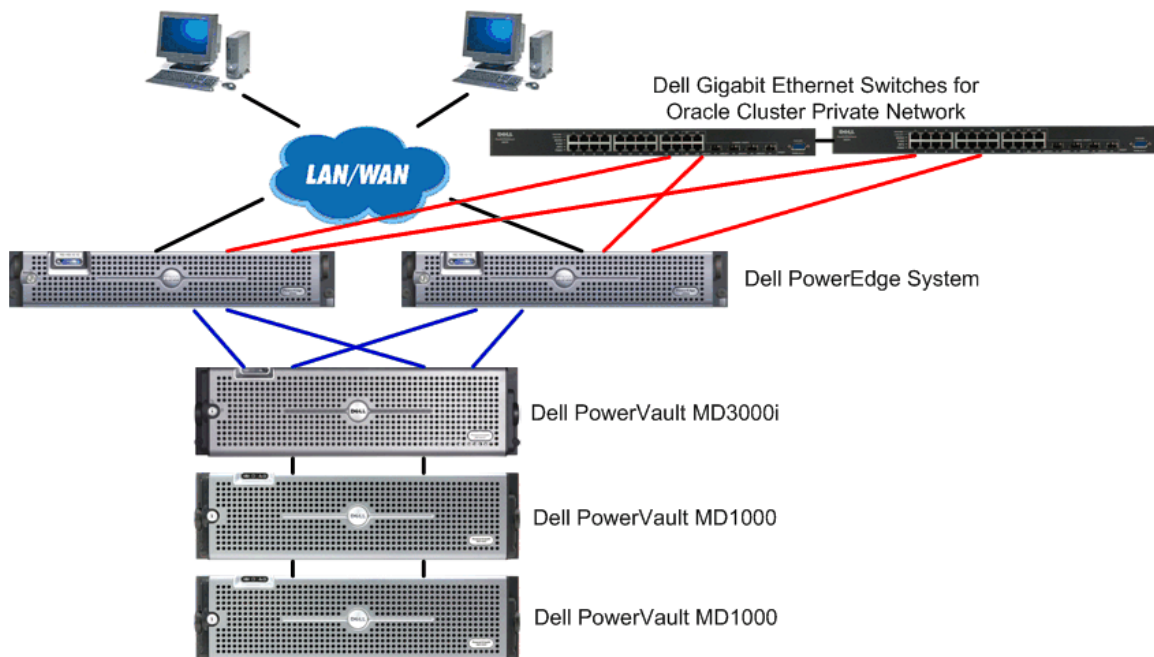
The Dell Reference Configuration for Oracle 10g on Dell PowerVault MD3000i iSCSI storage is intended to validate the following solution components:

- Two-node cluster of Dell PowerEdge 2950 quad-core systems.
- Dell PowerVault MD3000i iSCSI storage system with MD1000 expansion units.
- Red Hat Enterprise Linux 4 Update 4.
- Oracle Database 10g R2 Standard Edition (10.2.0.3) x86\_64.

The MD3000i storage can be directly attached to the host servers, or it can also be attached to host servers through an IP Storage Area Network (SAN) Gigabit Ethernet switch.

An architectural overview of the Dell Solution for Oracle 10g using direct attached Dell PowerVault MD3000i iSCSI storage is shown in Figure 1 below. This architecture is made of the following components:

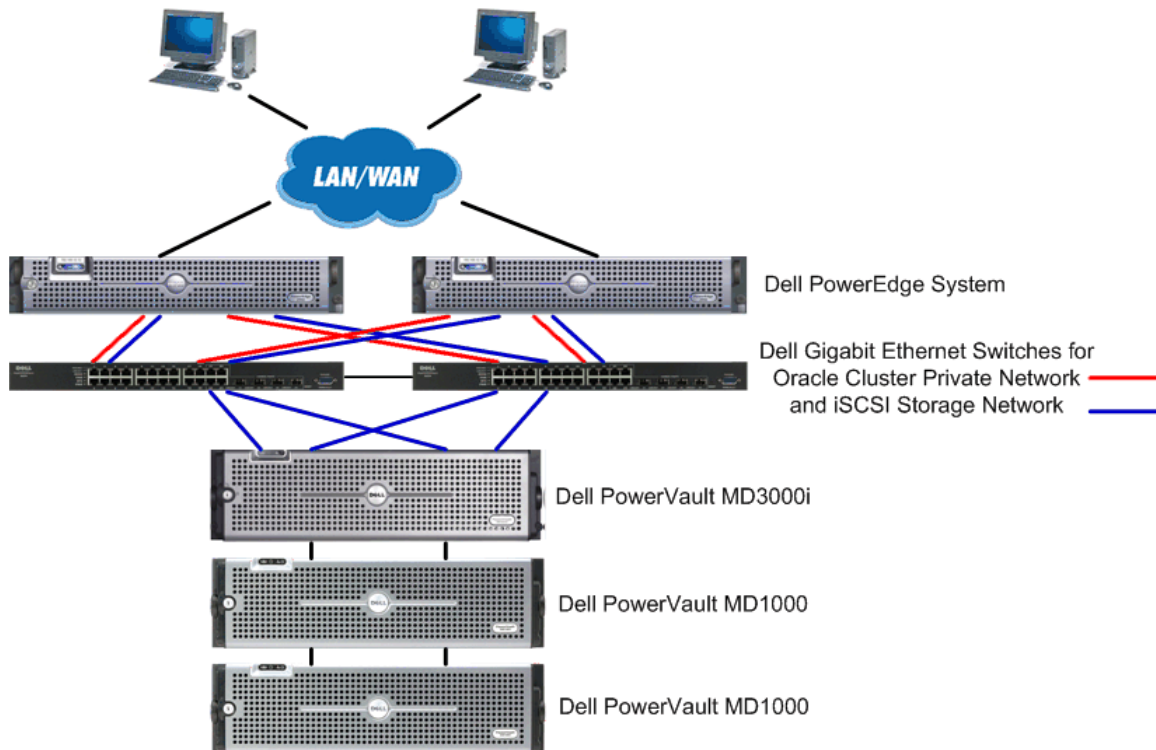
- Client systems that will access data stored within the Oracle 10g RAC two-node database
- Client-server network made up of network controllers, cables, and switches
- Two-node Dell PowerEdge 2950 quad-core servers running on RHEL 4 update 4
- Gigabit Ethernet switches for Oracle cluster interconnect network
- Server-storage interconnect using direct-attached Internet SCSI (iSCSI) storage
- Dell PowerVault MD3000i iSCSI storage array with 2 MD1000 expansion units, supporting up to 45 drives



**Figure 1 - Architectural Overview of Oracle 10g on Direct Attached Dell PowerVault MD3000i iSCSI Storage**

An architectural overview of the Dell Solution for Oracle 10g using network switch attached Dell PowerVault MD3000i iSCSI storage is shown in Figure 2 below. This architecture is made of the following components:

- Client systems that access data stored within the Oracle 10g RAC two-node database
- Client-server network made up of network controllers, cables, and switches
- Two-node Dell PowerEdge 2950 quad-core servers running on RHEL 4 update 4
- Gigabit Ethernet switches for Oracle cluster interconnect network and iSCSI SAN network
- Server-storage interconnect using switch-attached Internet SCSI (iSCSI) storage
- Dell PowerVault MD3000i iSCSI storage array with 2 MD1000 expansions, supporting up to 45 drives



**Figure 2 - Architectural Overview of Oracle 10g on Switch Attached Dell PowerVault MD3000i iSCSI Storage**

The Dell PowerVault MD3000i iSCSI storage is based on the Internet Small Computer Systems Interface (iSCSI) protocol which transports block-based storage over an Internet Protocol (IP) network. The MD3000i supports up to 16 standalone servers. The MD3000i is designed to allow multiple servers access to a common, shared pool of storage without the cost and complexities of dealing with a storage infrastructure network such as Fibre Channel. The MD3000i includes dual active/active RAID controllers with mirrored cache and host multi-path I/O management which helps ensure that storage processing continues without disruption. Other high-availability features include hot-pluggable, redundant power supplies, and active disk scrubbing.

Dell PowerEdge servers are designed to deliver the highest performance for critical enterprise applications like database, 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.

# Hardware Configuration - Dell Solution for Oracle 10g on Dell PowerVault MD3000i iSCSI Storage

## Storage Configuration

### Configuring PowerVault MD3000i iSCSI Storage Connections with Dual Gigabit Ethernet NIC for iSCSI Network

The MD3000i RAID array enclosure is connected to host via two hot-pluggable RAID controller modules in an active/active redundant configuration. As illustrated in Figure 3, the RAID controller modules are identified as RAID controller module 0 and RAID controller module 1. From the rear of the unit, RAID controller module 0 will be on the left side. Each RAID controller module has two iSCSI in port connectors, which provide either direct connection to the hosts or connection by way of Ethernet switches.




Figure 3 - PowerVault MD3000i iSCSI RAID Array Enclosure Back View

The default IP addresses of the Ethernet connectors on the RAID Controllers are as follows:


- RAID Controller Module 0
  - Management Port – 192.168.128.101
  - iSCSI Port 0 – 192.168.130.101
  - iSCSI Port 1 – 192.168.131.101
- RAID Controller Module 1
  - Management Port – 192.168.128.102
  - iSCSI Port 0 – 192.168.130.102
  - iSCSI Port 1 – 192.168.131.102

The IP addresses of the above Ethernet connectors could be set manually if users choose not to use the default settings. As a best practice, the user should isolate the iSCSI network from the public LAN, as well as the private heartbeat network used for the Oracle cluster interconnect. If a separate network cannot be dedicated to iSCSI traffic, it is recommended to assign iSCSI network to a separate virtual local area network (VLAN), which creates independent logical networks within a physical network. For details on VLAN configuration, please refer to the “Configuring Virtual Local Area Network (VLAN)” section below.



In order to have complete pathway redundancy, two separate, isolated networks need to be used. Therefore, it is a best practice to configure a separate subnet for each iSCSI controller port. For example, the iSCSI Port 0 of both RAID controller modules should be configured on the same subnet. Similarly, the iSCSI Port 1 of both RAID controller modules should be configured on the same subnet which is different from the subnet used by the iSCSI Port 0.


Figure 4 illustrates the direct attached interconnection of the two-node PowerEdge cluster hosting the Oracle 10g database to a PowerVault MD3000i storage system with two MD1000 expansion units where the data resides. Using dual Gigabit Ethernet NIC ports in each server with an active/active configuration, each server has a redundant pathway to the controllers and the virtual disks. If either an iSCSI NIC fails or an MD3000i RAID controller module fails, operations continue using the remaining iSCSI NIC – MD3000i storage controller pair. As mentioned above, it is a best practice to configure separate subnets for iSCSI Port 0 and iSCSI Port 1 on the RAID controller module. In order to configure a redundant iSCSI connection on each host server, one of the host iSCSI NIC port connects to iSCSI Port 0 of one MD3000i RAID controller module, and the other iSCSI NIC port connects to iSCSI Port 1 of the other MD3000i RAID controller module.



**Figure 4 - Cabling a Direct Attached PowerVault MD3000i Storage Array with Two PowerVault MD1000 Expansion Units**

Host servers can also be attached to the MD3000i through an IP storage area network (SAN) industry-standard 1GB Ethernet switch. Figure 5 illustrates the switch attached interconnection of the two-node PowerEdge cluster hosting the Oracle 10g database to a PowerVault MD3000i storage system with two MD1000 expansion units. Each host server has a separate connection to the redundant Gigabit Ethernet switches and to duplex RAID controllers. In this configuration, there is pathway redundancy at the host, at the two Ethernet switches, and at the duplex RAID controllers.





**Figure 5 - Cabling a Switch-Connect PowerVault MD3000i Storage Array with Two PowerVault MD1000 Expansion Units**

The MD3000i can be daisy-chained with up to two MD1000 expansion units, providing access to a maximum of 45 disks in the entire storage array. Figure 3 and Figure 4 illustrate the cabling method in which two MD1000 enclosures can be cascaded from an MD3000i enclosure.

The physical disks in the MD3000i storage array provide the physical storage capacity for the Oracle 10g RAC database. Before data can be stored, the MD3000i physical storage capacity must be configured into components, known as virtual disk groups and virtual disks. A disk group is a set of physical disks that are logically grouped and assigned a RAID level. Each disk group created provides the overall capacity needed to create one or more virtual disks, which are logical entities that the server uses to store data. For details on virtual disk configuration, please refer to the “Configuring Virtual Disks” section below.

## Configuring Virtual Disks

As discussed in the Section “Configuring PowerVault MD3000i iSCSI Storage Connections with Dual Gigabit Ethernet NIC for iSCSI Network” above, a Virtual Disk is a logical unit of physical disks presented to the host in the MD3000i storage.

The storage for an Oracle 10g RAC database can be divided into the following five areas:

- The first area is for two mirrored copies of Oracle Cluster Registry (OCR), three mirrored copies of 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. Oracle 10g R2 supports both Normal Redundancy level and External Redundancy level for the OCR and CSS Voting Disk. It is recommended to specify Normal Redundancy level to allow Oracle Clusterware to manage mirrored copies of the OCR and CSS Voting Disk on separate raw devices. The Oracle Universal Installer (OUI) allows two copies of OCS mirrored files and three copies of CSS Voting Disk mirrored files during the Clusterware installation.

- The second area is for database data that is 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 Virtual Disks on separate Disk Groups. The separation can enable better I/O performance by ensuring that these files do not share the same physical disks.

Table 1 shows a sample Virtual Disk configuration with three Virtual Disks for each of the three storage areas described above.

| Storage Area        | Minimum Size                                       | RAID                   | Number of Partitions     | Used For   | OS Mapping   |
|---------------------|--|------------------------|--------------------------|--|--|
| One Virtual Disk    | 2048 MB  | 10                     | Six of 300 MB each       | Three copies of Voting disk, Two copies of Oracle Cluster Registry (OCR), and SPFILE for ASM instances | Six raw devices for Voting Disk (three copies), OCR (two copies), and SPFILE |
| Three Virtual Disks | Larger than the size of your database              | 10, or 5 for read-only | One on each Virtual Disk | Data   | ASM disk group DATABASEDG  |
| Six Virtual Disks   | Minimum twice the size of your second Virtual Disk | 10, or 5 for read-only | One on each Virtual Disk | Flash Recovery Area  | ASM disk group FLASHBACKDG   |

**Table 1 - Virtual Disks for the Cluster Storage Groups**

Figure 6 illustrates a sample disk group and virtual disk configuration, which separates the three storage areas on distinct physical disks, on an MD3000i array with two MD1000 expansion units. As the need for storage increases, additional disks can be added to the database area. With the use of the Oracle Automatic Storage Management (ASM), expansion of the data storage area and the Flash Recovery Area can be simplified.



**Figure 6 - Separation of Disk Groups and Virtual Disks within a PowerVault MD3000i Storage Array and two MD1000 Expansion Units**

RAID 10 is often considered the optimal choice for Oracle 10g RAC virtual disk implementation because it offers fault tolerance, greater read performance, and greater write performance than other RAID configurations.<sup>1</sup> The disk group on which the data and Flash Recovery Area reside should be configured with RAID 10. The disk group on which the SFFILE resides could be configured with RAID 1, which provides protection from drive hardware failure. The disk groups on which the mirrored copies of OCR and Voting Disk reside could be configured with RAID 0 as Oracle has provided data redundancy by using multiple copies of OCR and Voting Disk files. As a best practice, the three virtual disks for OCR and Voting Disk should be distributed to different MD3000i RAID controller modules as preferred owners so that if one RAID controller module fails, the virtual disks can be reassigned to the remaining RAID controller module.

Because additional drives are required to implement RAID 10, it may not be the preferred choice for all applications. For the disk group of the virtual disk for the data storage area or Flash Recovery Area, RAID 5 provides a cost-effective alternative especially for predominantly read-only workloads such as a data warehouse database. However, RAID 5 is not suitable for heavy write workloads, such as in an OLTP database, as RAID 5 can have significantly lower write performance due to the reading and writing of parity blocks in addition to the reading and writing of database data.

Each virtual disk created in storage will be presented to all the Oracle 10g 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.

<sup>1</sup> “Pro Oracle Database 10g RAC on Linux”, Julian Dyke and Steve Shaw, Apress, 2006.

## Configuring Challenge Handshake Authentication Protocol (CHAP)

Challenge Handshake Authentication Protocol (CHAP) is an iSCSI authentication method which authenticates access between the storage array (target) and the iSCSI initiators on the host servers. CHAP is an optional feature and is not required to use iSCSI. CHAP validates the identity of remote clients. As a best practice, it is recommended to set up CHAP security to authenticate all access to the iSCSI ports on the storage array. This type of CHAP is called target CHAP where the storage array authenticates all requests for access issued by the iSCSI initiators on the host server via a CHAP secret.

## Gigabit Ethernet Switch Configuration

### Configuring Virtual Local Area Network (VLAN)

As discussed in the Section “Configuring PowerVault MD3000i iSCSI Storage Connections with Dual Gigabit Ethernet NIC for iSCSI Network” above, it is recommended to use a separate network for iSCSI storage needs. If a separate network cannot be dedicated, it is recommended to assign the storage function to a separate virtual local area network (VLAN), which creates independent logical networks within a physical network. As illustrated in the switch connection architecture overview in Figure 2 above, the two Gigabit Ethernet switches are shared by the iSCSI network and the Oracle RAC interconnect network. A VLAN can be set up on each of the two Ethernet switches and the ports used for iSCSI network can be assigned to the VLAN, as illustrated in Figure 5.

## Server Configuration

### Configuring Fully Redundant Ethernet Interconnects

Each Oracle 10g RAC database server that is configured using the MD3000i iSCSI storage needs at least five NIC ports, one for the external interface, two for the private interconnect network, and two for the iSCSI storage 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 RAC cluster is assigned as the master node. In the event of a interconnect NIC failure in a single interconnect NIC environment, the server loses 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 to implement a fully redundant interconnect network configuration, with redundant private NICs on each server and redundant private network switches.<sup>2</sup>

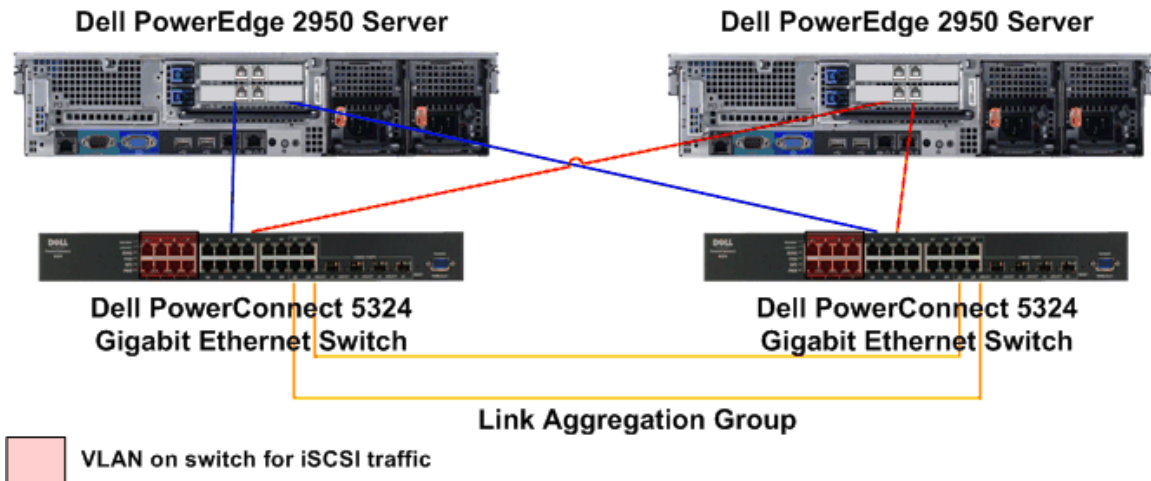
Figure 7 illustrates the CAT 5E/6 Ethernet cabling of a fully redundant interconnect network configuration of a two-node PowerEdge RAC cluster, with two private NIC ports on each server, and two private network switches. For this type of redundancy to operate successfully, it requires the implementation 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.<sup>3</sup> For details on configuring NIC teaming, please refer to the “Configuring the Private NIC teaming” section below.

---

<sup>2</sup> Dyke and Shaw, op. cit.

<sup>3</sup> Dyke and Shaw, op. cit.



**Figure 7 - Ethernet Cabling a Fully Redundant Private Interconnect Network**

### **Configuring Dual NIC Ports for MD3000i iSCSI Storage Network**

As illustrated in Figure 4 and Figure 5, it is recommended that two NIC ports be configured on each of the PowerEdge servers hosting the Oracle 10g RAC database to provide redundant links to the PowerVault MD3000i storage array. Using dual NIC ports helps protect against potential data loss or corruption if either a host iSCSI NIC port or an MD3000i RAID controller module fails.

### **Configuring Dual NIC Ports for Private Network**

As illustrated in Figure 7, it is recommended that two private NICs be installed on each of the PowerEdge servers hosting the Oracle 10g RAC database to provide redundant private network links. In the event of a NIC failure in a single private NIC environment, Oracle Clusterware can remove the node from the cluster.

# Software Configuration - Dell Solution for Oracle 10g on Dell PowerVault MD3000i iSCSI Storage

## 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 10g 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 affecting one of the NIC interfaces occurs – examples include switch port failure, cable disconnection, or a failure of the NIC itself – network traffic should be routed to the remaining operable NIC interface. Failover should occur transparently to the Oracle 10g 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 10g 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.<sup>4</sup>

### Configuring SSH and RSH

During the installation of Oracle 10g RAC software, the Oracle Universal Installer (OUI) is initiated on one of the nodes of the RAC cluster. OUI operates by copying files to and running commands on the other servers in the cluster. In order to allow OUI to perform, 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 Oracle Clusterware using the RAW Devices Interface

Before installing Oracle 10g RAC Clusterware software, it is necessary, at a minimum, for shared storage to be available on all cluster nodes for use by the Oracle Cluster Registry (OCR) and the Clusterware Cluster Synchronization Services (CSS) Voting Disk. The OCR file and the CSS Voting disk file can be placed on a shared raw device file. As discussed in the Section “Configuring Virtual Disks” above, one virtual disk is created for the OCR, Voting Disk, along with a SPFILE for the ASM instances; one virtual disk is created for the second mirrored copy of OCR and Voting Disk; and the other virtual disk is created for the third mirrored copy of Voting Disk. The partitions on these three virtual disks should be configured as raw disk devices.

---

<sup>4</sup> Dyke and Shaw, op. cit.

## Configuring Shared Storage for the Database using the ASM Library Driver

Oracle Automatic Storage Management (ASM) is a feature of Oracle Database 10g 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.<sup>5</sup>

As discussed in the Section “Configuring Virtual Disks” above, two virtual disks are created for the data storage area, and the Flash Recovery Area, respectively. It is recommended that these two virtual disks be configured as ASM disks to benefit from the capabilities of ASM. For Oracle 10g R2 database running on Red Hat Enterprise Server 4, ASM requires the installation of a number of additional RPM packages including the following:

```
oracleasm-2.6.9-42.0.8.ELlargesmp-2.0.3-1
oracleasm-2.0.2-1
oracleasm-support-2.0.3-1
```

### Enabling Full Network Startup

The Dell Solutions for Oracle Database 10g includes a deployment CD which automates the deployment of server OS optimized for Oracle database configuration. The deployment CD installs a customized system service called networkwait on Linux platforms. The networkwait service allows the network to fully startup upon server reboot. Otherwise, the next sequential iSCSI service may not have network access and fail on mounting storage devices.

## Oracle 10g R2 Configuration

### Reconfiguring Clusterware Parameter Disktimeout to 300

As documented in the Oracle MetaLink note #294430.1, the Synchronization Services Component (CSS) of the Oracle Clusterware maintains two heartbeat mechanisms:

1. The disk heartbeat to the voting device.
2. The network heartbeat across the interconnect, which establishes and confirms valid node membership in the cluster

Both of these heartbeat mechanisms have an associated timeout value. In 10g R2, the network heartbeat timeout value is determined by the CSS parameter miscount. The disk heartbeat timeout value is normally determined by the setting of the CSS parameter disktimeout, except for the initial cluster formation or slightly before cluster reconfiguration. The default value of the CSS disktimeout is 200 seconds. For Oracle RAC 10g R2 with MD3000i iSCSI storage, it is recommended to set CSS disktimeout to 300. This is to allow ample time for the virtual disks to trespass to the remaining RAID Controller under the situation when one RAID Controller fails.

Please refer to Oracle MetaLink note #462616.1 for instructions on reconfiguring the CSS disktimeout of 10g R2 Clusterware for the Dell MD3000i iSCSI storage.

---

<sup>5</sup> “Oracle Database 10g – Automatic Storage Management Overview”, Oracle TechNet.  
<http://www.oracle.com/technology/products/manageability/database/pdf/asmov.pdf>



## Configuration Deliverables List - Dell Solution for Oracle 10g on Dell PowerVault MD3000i iSCSI Storage

This section contains the Solution Deliverables List (SDL) for the Dell solution for Oracle 10g on MD3000i iSCSI storage. It contains a detailed listing of server and storage hardware configurations, firmware, driver, OS, and database versions.

| <b>Recommended Hardware/Software Requirements</b><br>(For details, see below) |   |   |
|---|---|---|
|   | <b>Validated Component(s)</b>   | <b>Minimum Oracle RAC Configuration</b> |
| <b>PowerEdge Nodes</b>  | PowerEdge 2950  | 2                                       |
| <b>Memory</b>   | All valid PowerEdge 2950 memory configurations                            | 1Gig (per node)                         |
| <b>PowerVault Storage Array</b>   | PowerVault MD3000i with 2 PowerVault MD1000 expansion units <sup>1</sup>  | 1                                       |
| <b>Ethernet Ports</b>   | Intel or Broadcom Gigabit NICs  | 5                                       |
| <b>Ethernet Switches (For Private Interconnect and iSCSI Network)</b>         | Any Dell PowerConnect Gigabit-only Switches                               | 2                                       |
| <b>Raid Controllers (Used for internal storage only)</b>                      | PERC 5/i  | 1 (Per Node)                            |
| <b>Internal Drive</b>   | All valid PowerEdge 2950 internal storage configurations                  | 73 Gig/node                             |
| <b>Oracle Software &amp; Licenses</b>   | Oracle 10g R2 10.2.0.1 Standard Edition (Base) + Oracle Patchset 10.2.0.3 | RAC                                     |
| <b>Operating System</b>   | Red Hat Enterprise Linux Server 4 Update 4                                |   |
| <b>Recommended Support Contract</b>   | Dell Gold or Platinum Plus Enterprise Support                             |   |

**Table 2 – Solution Minimal Hardware/Software Requirements**

Notes:

1. PowerVault MD3000 with 2 PowerVault MD1000 expansion units (max 45 disks supported).

| <b>Validated Servers</b>                      |   |                |                            |              |
|---|---|----------------|----------------------------|--------------|
| <b>PowerEdge Servers</b>                      | <b>Model</b>  | <b>BIOS[*]</b> | <b>ESM/BMC Firmware[*]</b> | <b>Notes</b> |
|   |   | PE2950         | 1.5.1                      | 1.33         |
| <b>Internal Disks RAID</b>                    |   |                |                            |              |
| PERC 5/i                                      | Firmware version = 5.1.1-0040; Driver version = 00.00.02.03-RH1                             |                |                            |              |
| <b>Interconnect Network</b>                   |   |                |                            |              |
| Intel PRO/1000 Network Drivers                | Driver version = 7.0.33-k2-NAPI   |                |                            |              |
| Ethernet Channel Bonding Driver               | Driver version = v2.6.3   |                |                            |              |
| <b>iSCSI Network</b>                          |   |                |                            |              |
| Broadcom NetXtreme II Gigabit Ethernet Driver | Driver version = bnx2 1.4.38  |                |                            |              |
| <b>Direct Attached SAS Storage</b>            |   |                |                            |              |
| <b>PowerVault Storage</b>                     | MD3000i; Firmware = 06.70.10.60<br>MD1000; Firmware = A.03                                  |                |                            |              |
| <b>Software</b>                               |   |                |                            |              |
| Oracle  | 10g R2 10.2.0.1 Standard Edition (Base) + 10.2.0.3 (PatchSet)                               |                |                            |              |
| Oracle ASMLib                                 | oracleasm-2.6.9-42.0.8.ELlargesmp-2.0.3-1, oracleasm-lib-2.0.2-1, oracleasm-support-2.0.3-1 |                |                            |              |
| Operating system                              | RHEL 4u4 (2.6.9-42.0.8.ELlargesmp)  |                |                            |              |
| DKMS  | dkms-2.0.17-1   |                |                            |              |
| iSCSI initiator                               | iscsi_sfnet-4.0.1.11.3e-1dkms   |                |                            |              |
| iSCSI initiator utilities                     | iscsi-initiator-utils-4.0.3.0-4   |                |                            |              |
| RDAC MPP driver                               | linuxrdac-09.01.B6.70-1dkms   |                |                            |              |

**Table 3 – Solution Detailed Firmware, Driver and Software Versions**

NOTES:

\*: Minimum BIOS and ESM/BMC versions. For the latest BIOS updates go to <http://support.dell.com>

## Conclusion

The Dell PowerVault MD3000i iSCSI SAN array is designed to simplify the deployment and operations of solutions such as Oracle Real Application cluster (RAC) databases. The MD3000i enables storage to be shared by up to 16 servers using standard Gigabit Ethernet infrastructure unifying the hardware and management of storage and networking.

Dell Solutions for Oracle 10g are designed to 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 10g RAC database with Red Hat Enterprise Linux 4 on Dell PowerEdge servers and Dell PowerVault MD3000i iSCSI storage arrays.

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

## Tables and Figures Index

|   |    |
|---|----|
| Table 1 - Virtual Disks for the Cluster Storage Groups.....   | 10 |
| Table 2 – Solution Minimal Hardware/Software Requirements.....  | 16 |
| Table 3 – Solution Detailed Firmware, Driver and Software Versions .....  | 17 |
| Figure 1 - Architectural Overview of Oracle 10g on Direct Attached Dell PowerVault MD3000i iSCSI Storage.....                         | 5  |
| Figure 2 - Architectural Overview of Oracle 10g on Switch Attached Dell PowerVault MD3000i iSCSI Storage.....                         | 6  |
| Figure 3 - PowerVault MD3000i iSCSI RAID Array Enclosure Back View.....   | 7  |
| Figure 4 - Cabling a Direct Attached PowerVault MD3000i Storage Array with Two PowerVault MD1000 Expansion Units .....                | 8  |
| Figure 5 - Cabling a Switch-Connect PowerVault MD3000i Storage Array with Two PowerVault MD1000 Expansion Units .....                 | 9  |
| Figure 6 - Separation of Disk Groups and Virtual Disks within a PowerVault MD3000i Storage Array and two MD1000 Expansion Units ..... | 11 |

## 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 10g RAC on Linux”, Julian Dyke and Steve Shaw, Apress, 2006.
3. “Oracle Database 10g – Automatic Storage Management Overview”, Oracle TechNet.  
<http://www.oracle.com/technology/products/manageability/database/pdf/asmov.pdf>