

# **Deploying an Oracle 10g RAC Database with ASM on Windows Server 2003 with an LSI Logic Storage Array**

**By Joshua Weiser**

**Installation and Configuration Guide**

25103-00, Rev. A  
February 2007



# Contents

## DEPLOYING AN ORACLE 10G DATABASE WITH ASM ON WINDOWS SERVER 2003 WITH AN LSI LOGIC STORAGE ARRAY

Introduction .....	1
Benefits.....	2
Disclaimers .....	3
Test Bed.....	3
Test Bed Overview.....	3
Test Bed Details .....	4
Software .....	4
Three Servers .....	4
Two LSI Logic Storage Arrays .....	4
HBAs .....	5
Configuring the Host and the Storage Array.....	5
Configuring the Host.....	5
Configuring the Storage Array.....	6
Verifying the Host and Storage Array Setup .....	9
Oracle Setup and Installation .....	9
Preparing for Installation .....	10
Configuring the Storage .....	12
Disabling Write Caching .....	12
Enabling Auto Mounting .....	13
Formatting All Disk Drives .....	13
Installing the Oracle Clusterware.....	14
Installing the Oracle Database Software.....	17
Creating the Oracle Database with DBCA .....	17
Conclusion.....	19
Contact Information .....	20
North American Headquarters .....	20
LSI Logic KK Headquarters .....	20

LSI Logic Europe Ltd. European Headquarters .....	20
Appendix A: Oracle RAC Troubleshooting .....	21
Multiple problems .....	21
Reinstalling the Oracle Clusterware.....	21
Oracle Clusterware complains that “nodes are not clusterable” .....	21
Services fail to start at the end of Oracle Clusterware installation .....	22
Appendix B: References.....	23

# *Deploying an Oracle 10g Database with ASM on Windows Server 2003 with an LSI Logic Storage Array*

---

## **Introduction**

IT departments that already run an Oracle 10g database with ASM on Windows Server 2003 might want to upgrade their storage array to an LSI Logic storage array. Any proposed upgrade raises the question of interoperability with the existing system.

Using a test setup, this document shows that it is possible to successfully install Oracle 10g Clusterware and the Oracle 10g database application with an LSI Logic Fibre Channel storage array.

This document shows how to overcome the following key obstacles to a successful installation:

- **Incorrect file permissions for the various components.** When the Oracle Cluster Register OCR file permission is set improperly, installation of Oracle Cluster might fail, or, even if the installation completes, CRS does not run properly. The same rules apply to the voting disks.
- **Not getting disk access to the storage array from all shared devices.** If you do not map all shared volumes to the host group, not all hosts see the correct LUNs. Also, if the Host Bus Adapter (HBA) driver you install is not compatible with the BIOS and firmware or does not use the recommended settings, the host might have problems discovering the LUNS or might experience unexpected errors. Furthermore, if you forget to disable auto failback when you install the RDAC failover driver, LUN thrashing could result in the event of a path failure on a single host.
- **Incorrect partition of an LSI Logic storage array.** Map all shared LUNs through SANshare storage partitioning to the host group to which all hosts in the cluster belong. If a shared LUN is mapped only to a single host, the other LUNs can not access this LUN, causing problems during installation.
- **Problems installing the Oracle Clusterware and the Oracle Database software.** Errors during configuration of a host setup and SANshare storage partitioning cause various problems (see “[Appendix A: Oracle RAC Troubleshooting](#)” on page 21) during installation of Oracle Clusterware and Oracle Database software. As mentioned in

previous bullets, disk access to the correct hosts is crucial. File permissions and ownership are just as important. Kernel parameters must meet the minimum levels that Oracle recommends. An experienced database administrator can identify the correct kernel parameter settings for your environment.

## Benefits

This method of deploying Oracle 10g database with ASM on Windows Server 2003 allows you to use an LSI Logic storage array and thus reap its many benefits. The LSI Logic storage array combined with Oracle RAC provide a reliable solution for the high availability needs of the small-to-medium business installation.

Performance benefits of the LSI Logic storage array:

- Performs among the fastest storage arrays in its class
- Includes front-to-back 4-gigabit storage array controller that is capable of sustained throughput of up to 1600MB/s
- Delivers up to 42,254 sustained IOPS (per SPC-1 benchmarking)

Benefits of the SANshare Storage Partitioning premium feature:

- Provides the ability to map individual LUNs (volumes) to a specific host or to a group of hosts, also known as Host Groups
- Specifies each host with an operating system allowing a single LSI Logic storage array to support volumes mapped to multiple hosts safely running a myriad of operating systems. The Oracle Database application is supported across all operating systems for a given controller firmware. Please refer to the compatibility matrix for the supported operating system and the controller firmware you are running. Applications do not manage or manipulate the data path. Oracle RAC is not considered an application and therefore is supported on only the operating systems that have been tested. Tested operating systems are as follows:
  - Windows 2003
  - Redhat Linux Advanced Server 3 and 4
  - SUSE Linux Enterprise Server 8 and 9
  - Solaris Sparc 8,9, and 10

Benefits of the Remote Volume Mirroring (RVM) premium feature:

- Allows data mirroring from one storage array to another.
- Relieves the host of creating a mirror copy and allows the system administrator to keep an exact copy of the production data.

## Disclaimers

This document serves as an example only, *not* as a best practices guide or as a configuration guide. The steps in this document were performed on a test setup. There was no live data involved. This document implies no warranties. These instructions are not guaranteed to work in every situation. For the latest information and instructions on installing Oracle products, refer to Oracle installation guides. Go to: <http://otn.oracle.com>

This document does not cover the following situations:

- An existing non-RAC Oracle installation with an LSI Logic storage array  
(This situation requires support from Oracle to convert the database to a RAC database.)
- An existing Oracle RAC and/or Oracle install with an LSI Logic storage array

The complexity of the installation described in this document requires that the installer have considerable prior expertise. Do not attempt to follow this procedure unless you are properly qualified. Installing this type of configuration for production requires special expertise for tuning an LSI Logic storage array, the Oracle Database parameters, and the host server to best fit the needs of your configuration. This type of installation requires expertise on both data migration and on Oracle. The installer must have worked with Oracle Databases, Oracle Clusterware, Windows Server 2003, and an LSI Logic Fibre Channel storage array.

## Test Bed

The operating system used in this test setup is Windows Server 2003 R2 SP1. The installer must be familiar with all aspects of the hardware and software configuration before attempting to follow the instructions in this document.

### Test Bed Overview

The test setup consists of Windows 2003 R2 SP1. For the latest patches and updates, go to Oracle's web site: <http://otn.oracle.com>.

Oracle Cluster File System (OCFS) and Oracle Automatic Storage Management (ASM) were used to configure the shared devices. Both were developed by Oracle for Oracle database products.

- OCFS – This general purpose cluster file system stores the ASM parameter file. However, this storage method is optional. It is entirely possible to have Oracle databases that do not use OCFS. For available storage options, check the *Oracle Clusterware and Oracle Real Application Clusters Installation Guide, 10 g, Release 2 (10.2) for Microsoft Windows*. For documentation on OCFS, go to: <http://www.oracle.com>

- ASM – A feature in Oracle Database 10g that provides the database administrator with a simple storage management interface that is consistent across all server and storage platforms. As a vertically integrated file system and a volume manager, purpose-built for Oracle database files, ASM provides the performance of asynchronous I/O with the easy management of a file system. ASM provides capability that saves the database administrator time and provides flexibility to manage a dynamic database environment with increased efficiency. File systems and volume managers utilize the disk that is served up from the storage array. The disk is then filtered through the multi-path driver. As long as the operating system, the HBA, and the multi-path driver are supported by LSI Logic, Oracle ASM will work.
  - For supported HBAs, go to:  
[http://www.engenio.com/html/partners/compatible\\_matrix/interop.asp](http://www.engenio.com/html/partners/compatible_matrix/interop.asp)
  - For documentation on ASM go to:  
<http://www.oracle.com/technology/products/database/asm/index.html>
- The test setup uses the cluster software supplied with Oracle 10g R2. In previous releases of Oracle, this cluster software was called Cluster Ready Services (CRS). However, the cluster software has been renamed to Oracle Clusterware. Therefore, the terms *Clusterware* and *CRS* are used interchangeably in this document.

## Test Bed Details

### Software

Oracle 10g R2 DB – 10.2.0.1.0

Oracle 10g R2 Clusterware – 10.2.0.1.0

### Three Servers

Server 1 – 2 Intel Xeon 1.4 GHz CPUs, 1 GB RAM, 18 GB SCSI hard disk drive

Server 2 – 2 Intel Xeon 3.06 GHz CPUs, 3.75 GB RAM, 18 GB SCSI hard disk drive

Server 3 – 2 Intel Xeon 3 GHz CPUs, 2 GB RAM, 36 GB SCSI hard disk drive

All servers have Windows Server 2003 R2 SP1

### Two LSI Logic Storage Arrays

Firmware – 06.16.88.00

NVSRAM – N6091-616800-014

SANtricity Storage Manager – 09.16.G5.70

Multipathing (RDAC/MPP) – 09.01.30.30

## **HBAs**

One dual-port HBA in each host

2 Qlogic 2462 cards

Driver – 9.1.1.11 (SCSIport driver)

Firmware – 4.00.12

BIOS – 1.04

1 Emulex 9802DC

Driver – 5.5.20.10 (SCSI Miniport driver)

Firmware – 1.91a1

BIOS – 5.01a7

# **Configuring the Host and the Storage Array**

## **Configuring the Host**

To completely configure the host, you must install and configure the operating system, the HBA driver, and the SANtricity Storage Manager software.

- 1** Install the operating system. Refer to the operating system documentation. Verify with Oracle that the operating system is certified.
- 2** Install the HBA driver. For supported HBAs and drivers, refer to the Engenio Interoperability Matrix. Go to:  
[http://www.engenio.com/html/partners/compatible\\_matrix/interop.asp](http://www.engenio.com/html/partners/compatible_matrix/interop.asp)
- 3** Install the SANtricity Storage Manager software. Also known as SMinstaller, it consists of the following individual packages:
  - SMclient – Manages the storage array
  - SMutils – Collection of utilities to verify what the host sees from the storage array. Installing SMutils is highly recommended, and SMutils can be installed independently of the other packages.
  - SMruntime – Java runtime environment
  - SMagent – Required for inband management of the storage array
  - RDAC – Host multipath failover driver
- 4** Configure the HBA and the registry parameters as specified by the HBA Settings document. You might want to modify some of these settings to match your environment.

For the HBA Settings document, contact Support. See “[Contact Information](#)” on page 20.

---

**IMPORTANT** If you are using LSI RDAC, set the DisableLunRebalance registry parameter to 3 to prevent volume thrashing and resulting I/O failures.

---

## Configuring the Storage Array

- 1 Using the SANtricity Storage Manager Client, create volumes on the LSI Logic storage array.
  - a Oracle home – Install the Oracle Clusterware and the Oracle Database software on volumes from the storage array. Each host must have two volumes, one for the Oracle Clusterware and one for the Oracle Database software. You also can choose to install these softwares on the host drive itself.
  - b Oracle system and data files – Create separate volumes for voting, registry, and ASM disks. Configurations for production vary depending on the following factors:
    - Percentage of reads and writes
    - Size of the database
    - Expected growth
    - Purpose of the database, that is, data warehousing, OLTP and so forth
    - Number of simultaneous connections
    - Server configuration (memory, number of processors, speed of the processors)
    - Network infrastructure
    - Whether queries are optimized to utilize indexes

---

**IMPORTANT** Consult with an experienced Oracle database administrator to determine the best configuration.

---

- 2 In the SANtricity Storage Manager main window, right-click the Unconfigured Space icon, and choose **Create Volume** to open the **Create Volumes** wizard. Follow the prompts to create the volumes for this installation of Oracle RAC. Table 1 shows the options you must choose on each page of the wizard.

Table 1 Create Volumes Wizard

Page	Action
<b>Introduction</b>	Click <b>Next</b> .
<b>Specify Array</b>	Choose the RAID Level and the drive selection on this screen. The recommended drive selection is <b>Automatic</b> . <b>NOTE</b> Interpret the term <i>Array</i> to mean a storage array.
<b>Specify Capacity Name</b>	Specify the capacity and the name for the volume.
<b>Specify Volume-to-LUN Mapping</b>	Specify the host type and choose to map the volumes to LUNs at a later time. In the test setup, “Windows Non-clustered” is the host type.

- 3 Create host port partitions on the storage array to map the volumes to a single host or to a group of hosts for Oracle RAC. Map the appropriate volumes to the hosts and the host groups.

Figure 1 shows the storage layout used in the test setup.

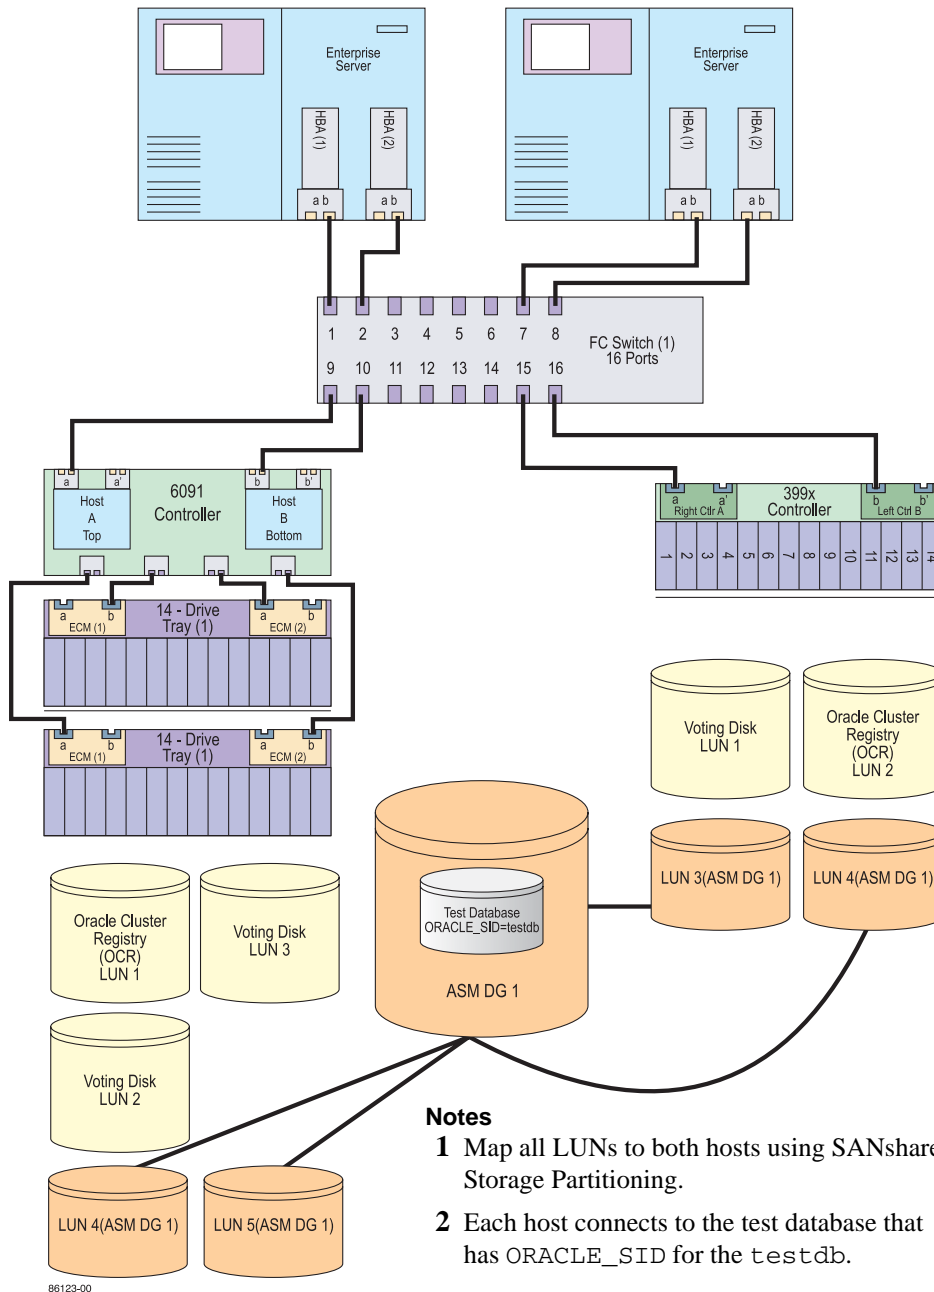


Figure 1 Mapping for the Test Setup

The voting disks, the registry disks, and the ASM disks are split between the two storage arrays to provide better tolerance against failure of a single storage array. No software home volumes are necessary because the software files are installed on the hard drive.

Figure 1 also shows the host group Windows\_Oracle mappings on both storage arrays. These volumes are shared between the three hosts in this RAC. No volumes are mapped to the individual hosts because the host's hard drive is used for all host-specific files, namely, the Oracle Clusterware and the Oracle Database software.

The results of the mapping in this step are as follows:

- The host group Windows\_Oracle contains three hosts: win2k3\_node1, win2k3\_node2, and win2k3\_node3.
- Each host has two host adapter ports associated for mapping LUNs.
- Each host can see the volumes mapped to the host group, as shown in the **Accessible By** column. There are no volumes mapped to a single host. All volumes are shared.

---

**IMPORTANT** The LUN numbering for the volumes that each host sees should always start with 0 (zero) on at least one storage array. In addition, LUNs must be sequential. Therefore, do not skip LUN numbers when mapping. Problems can occur with hosts not seeing LUNs if these conditions are not met.

---

## Verifying the Host and Storage Array Setup

Verify the storage array access by verifying that the hosts have access to the mapped volumes.

The `<root_drive>\Program Files\Storage Manager\util\SMdevices.bat` command shows what the host sees. After the initial mapping of the volumes, you can either reboot the server or run the `hot_add` command. Rebooting is preferred.

## Oracle Setup and Installation

To install Oracle RAC remotely, enable either Telnet or Terminal Service/Remote Desktop services. For details, see *Oracle Clusterware and Oracle Real Application Clusters Installation Guide, 10g Release 2 (10.2) for Microsoft Windows*.

Clusterware Verification Utility (CVU) is a command line utility provided by Oracle to verify that the requirements are being met at certain defined stages of the Oracle Clusterware installation process. CVU is located at `<media>\clusterware\cluvfy`, where `<media>` is the location of the Oracle Clusterware installation files. For the CVU commands for each stage, see *Oracle Clusterware and Oracle Real Application Clusters Installation Guide, 10g Release 2 (10.2) for Microsoft Windows*. Some of the CVU commands are used in the following steps.

## Preparing for Installation

- 1 Check the Oracle installation manual for hardware and software requirements. For the latest certified list of hardware and software, go to:

<http://metalink.oracle.com>

- 2 Synchronize the clocks.

Make sure that the clocks on all of the hosts are synchronized. You can do this either manually or by using Windows time servers. If the clocks do not match, the Oracle Clusterware installation gives you errors.

- 3 Set up the host network.

Each host must have at least two network adapters, one for public interconnect and one for private interconnect.

- a Configure three static IP addresses for each host: one public, one private, and one virtual.

- Configure the public and virtual IP addresses on the same subnet. To configure the virtual IP addresses, simply specify the virtual IP addresses in the `%SystemRoot%\system32\drivers\etc\hosts` file. The remainder of the virtual IP address setup is completed during the software installation steps.
- Configure the private IP address on a separate subnet.

- b Select the same interfaces on all hosts. That is, on each host, `eth0` must be public and `eth1` must be private. Interface names must be identical on all hosts, with public and private interfaces having separate names. Interface names are case sensitive. In the test setup, all the public interfaces were named “Public” and all the private interfaces were named “Private.”

- c List the public interface first in the bind order—the order in which network services access the host. Use the following steps to verify the bind order.

- 1 Select **Start >> Control Panel >> Network Connections >> Advanced >> Advanced Settings**.

- 2 Click the **Adapters and Bindings** tab.

- 3 In the **Connections** area, observe whether the public interface appears before all others and whether the private interface is second. If the public and the private interface do not appear in the proper order, modify the order.

- d Disable Windows Media Sensing, as required by Oracle, by performing the following steps on each host.

- 1 Open `Regedit.exe`.

- 2 Navigate to  
HKEY\_LOCAL\_MACHINE\System\CurrentControlSet\Services\Tcpip\Parameters.
- 3 Add a new DWORD key: DisableDHCPMediaSense=1
- 4 Reboot the host to enable changes to the registry.
- e On each host, enter all three IP addresses of each host in the %SystemRoot%\system32\drivers\etc\hosts file.
- f Make sure that the loopback line in \etc\hosts does not contain the host name.

Table 2 Sample \etc\hosts File

IP Address	Host Name	Abbreviated Host Name
127.0.0.1	localhost.localdomain	localhost
172.22.215.186	win2k3_node1.ks.lsil.com	win2k3_node1
172.22.215.2	win2k3_node3.ks.lsil.com	win2k3_node3
172.22.215.3	win2k3_node2.ks.lsil.com	win2k3_node2
172.22.215.199		win2k3_node1vip
172.22.215.200		win2k3_node3vip
172.22.215.201		win2k3_node2vip
192.168.1.186		win2k3_node1prv
192.168.1.2		win2k3_node3prv
192.168.1.3		win2k3_node2prv

- 4 To download the Oracle Clusterware and Oracle Database software packages, go to: <http://otn.oracle.com> and then click **Downloads**.
- 5 Establish user equivalency. *User equivalency* means that a user can log on to any host, and from that host access all other hosts without using a password. Oracle installation requires user equivalency for remote copying of files.
  - a Set the same administrative user account and password on all RAC hosts.
  - b To verify that you have administrative privileges on all other hosts from the host where you will be installing the Oracle Clusterware software, use the \\<nodename>\c\$ command.
- 6 Use CVU commands to verify that all requirements up to this point are met.
  - a To check hardware requirements, use this command.
 

```
<media>\clusterware\cluvfy\runcluvfy.bat stage -post hws -n win2k3_node1,win2k3_node2,win2k3_node3 [-verbose]
```

- b** To check network setup, use this command.

```
<media>\clusterware\cluvfy\runcluvfy.bat comp nodecon -n  
win2k3_node1,win2k3_node2,win2k3_node3 [-verbose]
```

- 7** Disable the Distributed Transaction Coordinator service for the duration of the setup, because it might interfere with Oracle Clusterware installation. Do these steps on each host.

- a** Right-click **My Computer** and select **Manage >> Services and Applications >> Services**.

- b** Find the **Distributed Transaction Coordinator** service. Change the startup type to manual, and stop the service.

You can ignore any error or warning about virtual IP addresses.

## Configuring the Storage

To see how storage was laid out for the RAC installation of the test setup, see “[Configuring the Storage Array](#)” on page 6. To decide what layout best fits your environment, review all of the storage options available to you by referring to the Oracle Clusterware installation manual. After you have decided how much space you need and what volumes must be mapped in SANtricity Storage Manager, you are ready to configure the storage.

Configure the storage in three steps.

- 1** Disable write caching.
- 2** If you are using Windows 2003, enable auto mounting.
- 3** Format all disk drives.

These steps are described in detail in the following sections.

### Disabling Write Caching

Disable write caching for all disk drives that share data between hosts.

- 1** Right-click **My Computer** and select **Manage >> Device Manager >> Disk Drives**. If you are using RDAC, the mapped disk drives appear as RDAC *Virtual Disks*.
- 2** Under **Disk Properties** for each RDAC virtual disk, go to the **Policies** tab and remove the check from **Enable write caching**.

---

**NOTE** It is not necessary to disable write caching for software home volumes that you configured on the storage array.

---

## Enabling Auto Mounting

If you are using Windows 2003, enable auto mounting.

- 1 Run the following commands.

```
c:\> diskpart
DISKPART> automount enable
Automatic mounting of new volumes enabled
```

- 2 Type `exit` to log off the session.
- 3 Repeat step 1 and step 2 on each host in the cluster.
- 4 After all hosts have had automount enabled, reboot all hosts.

---

**IMPORTANT** You *must* reboot all hosts at this point.

---

## Formatting All Disk Drives

Format all disk drives, using one or both of the following methods.

- Format volumes used for the Oracle software installation as NTFS if these volumes reside on the storage array. Map two volumes to each host, one for the Oracle Clusterware and one for the Oracle Database software. These disk drives (or directories if you choose to install on the hard drive) must use the same drive letter and path across all hosts.
  - Format all other disk drives as raw devices initially, no matter what storage option (ASM, OCFS, or raw devices) you choose.
- 1 Create `crs` and `db1` directories in the `C:\` drive of all hosts. These will be the test setup homes for the Oracle Clusterware and the Oracle Database software, respectively.
  - 2 Create raw devices.
    - a Right-click **My Computer** and select **Manage >> Disk Management**. All of the disk drives from the storage array should appear as unformatted volumes.
    - b Format all disk drives as basic disks.
    - c Right-click a disk drive and select **Create disk >> Extended Partition**. Use the entire disk drive for the extended partition.
    - d Right-click the partition that you just created and choose **New Volume**. Choose no drive letter and do not format the partition. After this, if you check the disk drive properties, the File System column should show RAW.
    - e Repeat step d for all shared disk drives in your setup.

- f Reboot all hosts after creating raw devices.
- g Check on all hosts to see that your disk drives are discovered and recognized as raw devices after reboot.

## Installing the Oracle Clusterware

Before you install the Oracle Clusterware, you must ensure that all your servers are seeing the same devices. Due to multiple factors, sometimes different servers do not recognize the same LUN as the same device. This is a problem because in order for Oracle RAC to operate properly the servers must recognize each device through the same discovery path. To prevent this problem you must create links to each device on each server. These links must be the same on all servers.

Table 3 shows the devices that Oracle uses to configure the voting disks and the registry disks for RAC, with the links you must create. LUN 1 and LUN 4 have different names for each server. However, because the same link is used by each server, the situation remains workable. Oracle RAC references the devices under the **Link** column.

Table 3 Oracle Devices

Storage Array LUN #	Devices Seen by Server A	Devices Seen by Server B	Link
0	\\.\PhysicalDrive1	\\.\PhysicalDrive1	/reg_master
1	\\.\PhysicalDrive2	\\.\PhysicalDrive3	/reg_bkup
2	\\.\PhysicalDrive4	\\.\PhysicalDrive4	/vote_master
3	\\.\PhysicalDrive5	\\.\PhysicalDrive5	/vote_bkup1
4	\\.\PhysicalDrive6	\\.\PhysicalDrive7	/ vote_bkup2

Device persistence is another issue. After you reboot a single host, devices sometimes do not remain the same. To make sure that the underlying devices do not change upon a reboot or when new devices are added, you must use bindings at the host adapter level. For information on persistent binding, consult the manual for the Fibre Channel HBA in use. In addition, use the appropriate one of the following methods:

- Use these links as OCFS2 mount points, if you are using OCFS2.
- Create links at the operating system level from the correct device to the link location.

After you make sure that the essential devices are linked and are persistent, install the Oracle Clusterware according to the following instructions.

- 1 Go to the Oracle Clusterware image's location, and run `runInstaller` to start the **Oracle Universal Installer** wizard. Table 4 shows the options you must choose on each page of the wizard.

Table 4 Oracle Universal Installer Wizard

Page	Action
initial screen	Click <b>Next</b> .
<b>Specify Inventory Directory and Credentials</b>	Accept the default inventory location. This location is shared between the Oracle Clusterware home and the Oracle Database home. The location shown should be \$ORACLE_BASE/oraInventory.
<b>Specify Home Details</b>	Specify the Oracle Clusterware mount point (/crs in this case).
<b>Product-Specific Prerequisite Checks</b>	All checks should pass if everything is configured correctly. If not, fix the issue and click <b>Retry</b> .
<b>Specify Cluster Configuration</b>	Specify the hosts. Click <b>Use Cluster Configuration File</b> if you have many hosts. A “nodes are not clusterable” error appears if SSH or RSH are not configured correctly.
<b>Specify Network Interface Usage</b>	Specify the Interface Type appropriately as Public or Private.
<b>Specify Oracle Cluster Registry (OCR) Location</b>	Choose <b>External Redundancy</b> . In the <b>Specify OCR Location</b> field, type the pathway and filename to the registry mount point. In the <b>Specify OCR Mirror Location</b> field, type the pathway and filename to the backup registry mount points.
<b>Specify Voting Disk Location</b>	Specify the location of the voting disk files that you created when you created the volume on this volume group.
<b>Execute Configuration Scripts</b>	Follow the on-screen directions to run scripts on both servers. Scripts might take up to ten minutes to run completely. After the scripts have run, <i>do not</i> click <b>OK</b> yet.

- 2 Run the `vipca` application manually. Running `vipca` manually is necessary because Oracle 10.2.0.1 has a known bug (4437727) at the end of running the scripts. The last host you run the scripts on shows the following message:

---

```
The given interface(s), "eth0" is not public. Public
interfaces should be used to configure virtual IPs.
```

---

- a As root user on the last host the scripts were run on (Host2 in this setup), run the following command:

`$ORA_CRS_HOME/bin/vipca`

- b** Click **Next** on the initial screen.
- c** Use the **VIP Configuration Assistant** wizard to configure virtual IP addresses through `vipca`. Table 5 shows the options that you must choose on each page of the wizard.

Table 5 VIP Configuration Assistant Wizard

Page	Action
<b>Network Interface</b>	Select both interfaces, and then click <b>Next</b> .
<b>Virtual IPs for Cluster Nodes</b>	Configure the IP Alias Name and the virtual IP address as you specified in <code>/etc/hosts</code> .
<b>Summary</b>	Confirm your actions.
<b>Progress Dialog</b>	Watch as <code>Vipca</code> creates the appropriate application resources.
final screen shows configuration results	Click <b>Finish</b> .
<b>Execute Configuration Scripts</b>	Click <b>OK</b> . The configuration assistants verify proper configuration.

- 3** Check the Oracle Clusterware status by running the `$ORA_CRS_HOME/bin/crs_stat -t` command on either host.
- 4** Go to “[Installing the Oracle Database Software](#)” on page 17.

## Installing the Oracle Database Software

- 1 Verify that you unset the environment variables as you did before you ran the Oracle Clusterware installer. The Oracle Database software image or folder should be owned by `oracle:oinstall`.
- 2 As the `oracle` user, run `runInstaller` to open the **Oracle Universal Installer** wizard. Table 6 shows the options you must choose on each page of the wizard.

Table 6 Oracle Universal Installer Wizard

Page	Action
initial screen	Click <b>Installed Products</b> to make sure that the Oracle Clusterware is installed successfully.
<b>Select Installation Type</b>	Select <b>Enterprise Edition</b> .
<b>Specify Home Details</b>	Specify <code>\$ORACLE_BASE/&lt;any folder&gt;</code> as the location of the SW home directory for your database.
<b>Specify Hardware Cluster Installation Mode</b>	Select <b>Cluster Installation</b> and select all hosts.
<b>Product-Specific Prerequisite Checks</b>	All checks should pass if everything is configured correctly. If not, fix the issue and click <b>Retry</b> .
<b>Select Configuration Option</b>	Choose <b>Install database software only</b> . You will configure the database later with DBCA.
<b>Execute Configuration Scripts</b>	Follow the on-screen directions.

- 3 Go go [“Creating the Oracle Database with DBCA.”](#)

## Creating the Oracle Database with DBCA

For the test setup, the Database Configuration Assistant (DBCA) was used to create the clustered database. As the `oracle` user, run `dbca` (`$ORACLE_HOME/bin/dbca`) to open the Oracle Database Configuration Assistant, which guides you through the process of the creating the database. Table 7 on page 18 shows the options you must choose on each page of the wizard.

Table 7 Database Configuration Assistant Wizard (1 of 2)

Page	User Action
<b>Welcome</b>	Select <b>Oracle Real Application Clusters database</b> .
<b>Operations</b>	It is best to set up ASM <i>before</i> setting up the database. <ul style="list-style-type: none"> <li>• If you are using ASM, select <b>Configure Automatic Storage Management</b>.</li> <li>• If you are not using ASM, skip ahead in this table to the next instance of the Operations* page for instructions.</li> </ul>
<b>Node Selection</b>	Click <b>Select All</b> to select all hosts.
<b>Create ASM Instance</b>	Specify a password. The test setup uses an SPFILE (Oracle recommended) for the ASM instance. You must specify a shared OCFS2 location for the SPFILE. In the test setup, the location is /vote/dbfs, where /vote is the OCFS2 mount point for the voting disk.
<b>Create Disk Group</b>	Configure ASM disks into disk groups. DBCA recognizes these disks. Choose a <b>Redundancy</b> option:  If you choose <b>External</b> , create at least 2-3 disk groups, dividing the disks as you wish.  If you choose <b>High</b> or <b>Normal</b> redundancy, you must specify failover groups (see ASM documentation). Create at least one failover group per storage array used.
<b>ASM Disk Groups</b>	Click <b>Finish</b> to mount and use the ASM disk groups.
<b>Operations*</b>	Select <b>Create a Database</b> .
<b>Node Selection</b>	Click <b>Select All</b> to select all hosts.
<b>Database Templates</b>	Select the database type to create.
<b>Database Identification</b>	Name the database and SID. The SID prefix <i>must be exactly the same string</i> you specified in the \$ORACLE_SID environment variable without the numbers.
<b>Management Options</b>	Choose and configure the management options you want.
<b>Database Credentials</b>	Specify passwords.
<b>Storage Options</b>	Choose the storage option that is appropriate for your environment.
<b>ASM Disk Groups</b>	(This page only shows if you are using ASM.) Select one of the ASM disk groups you created to use as main storage for database files. You can create and modify new disks and groups at this point as well.
<b>Database File Locations</b>	Confirm that the storage option you choose is correct.
<b>Recovery Configuration</b>	Choose and configure a recovery option appropriate for the environment. If you intend to run hot backups of the database, you must check <b>Enable Archiving</b> .

Table 7 Database Configuration Assistant Wizard (2 of 2)

<b>Page</b>	<b>User Action</b>
<b>Database Content</b>	In the <b>Database Components</b> tab, you can select and configure components. However, these also can be added and configured later. If you are in a hurry or have not decided yet what you want, accept the defaults.  In the <b>Custom Scripts</b> tab, you can select any samples to include and any custom scripts you wish to run at the time of database creation. However, scripts also can be added or edited later. If you are in a hurry or have not decided yet what you want, accept the defaults.
<b>Database Services</b>	Create any database services that you want or just skip this page for now.
<b>Initialization Parameters</b>	Use the default values for initialization parameters unless you have good reasons to change them.
<b>Database Storage</b>	Click an item in the left panel to see its details in the right panel. You can reconfigure in the right panel if necessary.
<b>Summary</b>	View the summary information and save if you want a permanent record.
<b>Copying database files</b>	Watch as the database is created.
<b>Database creation complete</b>	Take notice when the database is successfully created.

## Conclusion

By following the instructions and the test setup provided in this paper, a system administrator can properly configure and install Oracle RAC using an LSI Logic storage array as the primary storage. Careful planning and execution of the prerequisites help to ensure a successful installation of this solution. The following items should be verified on the host setup:

- File permissions for all Oracle devices, such as the Oracle Cluster Register (OCR) disks and voting disks
- Storage array disk drive access for all shared devices. This includes the proper HBA driver and RDAC failover driver installation.

Provisioning the LSI Logic storage array is also an important step in the configuration. Verify that the SANshare storage partitioning of each storage array is configured to grant access to all shared LUNs by all hosts in the cluster. SANshare Storage Partitioning is a premium feature of an LSI Logic storage array that allows mapping of LUNs to a specific host or hosts.

The final step is installing the Oracle Clusterware and the Oracle Database software. If the host and the storage array setup have been completed correctly, the installation should go smoothly without errors. Typically, errors encountered during this step of the process are a result of errors during host setup or storage array provisioning.

Use this document to configure the host and the LSI Logic storage array so that Oracle installation becomes simple. Plan to upgrade to an LSI Logic storage array so that you can take advantage of fast performance and premium features. If you have a complex system with multiple hosts and multiple operating systems, you can use the SANshare storage partitioning premium feature to map these to a single LSI Logic storage array. If you want to fully protect your production database, you can use the Remote Volume Mirroring premium feature to provide complete data replication on or off site.

## Contact Information

For more information and sales office locations, visit the LSI Logic web sites at:

<http://www.lsi.com/cm/ContactSearch.do>

### North American Headquarters

Milpitas, CA

Tel: 866 574 5741 (within U.S), 408 954 3108 (outside U.S.)

### LSI Logic KK Headquarters

Tokyo, Japan

Tel: 81 3 5463 7165

### LSI Logic Europe Ltd. European Headquarters

United Kingdom

Tel: 44 1344 413200

## Appendix A: Oracle RAC Troubleshooting

This appendix lists known issues and their resolutions.

### Multiple problems

More than one problem might occur at the same time, such as host reboot during I/O, vote access hang, and database creation failure. To possibly resolve these problems, modify the CSS miscount. The CSS miscount value indicates the number of heartbeats missed before CSS evicts a host.

The following formula calculates CSS miscount in Oracle versions 10.1.0.3 through 10.2.0.1:

$$\text{disktimeout\_in\_secs} = \text{MAX} ((3 * \text{CSS miscount})/4, \text{CSS miscount} - 15)$$

Adjust the CSS miscount value to a higher value using the following steps.

- 1 Shut down all hosts in the cluster except one.
- 2 As root user, observe the current miscount setting by running the following command:

```
<CRS_install_path>\bin\crsctl get css miscount
```

- 3 Set the miscount value. In the following example, the miscount value is set to 240.

```
<CRS_install_path>\bin\crsctl set css miscount 240
```

- 4 Reboot the host.
- 5 Bring up all other hosts in the cluster.

### Reinstalling the Oracle Clusterware

Make sure that before the new Oracle Clusterware installation you delete and recreate the vote, backup vote, registry, and backup registry files you are using. Otherwise, the CRS and CSS services fail to start and other unpredictable problems occur.

### Oracle Clusterware complains that “nodes are not clusterable”

This error message means that user equivalency between hosts is not configured properly. *User equivalency* means that a user can log on to any host, and from that host access all other hosts without using a password. To resolve the problem, make sure that all hosts have the same username and password. Also check that all network settings and connections are configured properly.

## **Services fail to start at the end of Oracle Clusterware installation**

At the end of Oracle Clusterware installation, the Configuration Assistant reports that services are in an invalid pending state and failed to start.

- 1** Make sure `DisableDHCPMediaSense=1` registry key is set properly and in the correct location.
- 2** Check to see that the bind order for the network interfaces is in the right order.
- 3** Check that the Distributed Transaction Coordinator service is disabled.

## Appendix B: References

Automatic Storage Management - Oracle Technical Library

<http://www.oracle.com/technology/products/database/asm/index.html>

HBA Settings

The HBA Settings document is not available on an external link. Customers who want it should contact Support. To contact Support, see “[Contact Information](#)” on page 20.

Engenio Interoperability Matrix

[http://www.engenio.com/html/partners/compatible\\_matrix/interop.asp](http://www.engenio.com/html/partners/compatible_matrix/interop.asp)

*OCFS2 User's Guide*

Go to: <http://oss.oracle.com>, and then click on OCFS2.

*Oracle Clusterware and Oracle Real Application Clusters Installation Guide, 10g Release 2 (10.2) for Microsoft Windows*

See Oracle Documentation Library.

*Oracle Database Installation Guide, 10g Release 2 (10.2) for Microsoft Windows (32-bit)*

See Oracle Documentation Library.

Oracle Documentation Library

<http://www.oracle.com/pls/db102/homepage>



## **Document Description**

The purpose of this document is to explain the setup and configuration of an LSI Logic storage array with an Oracle 10g RAC Database (“the Document”).

## **Ownership of Materials**

The Document is provided as a courtesy to customers and potential customers of LSI Logic Corporation (“LSI”). LSI assumes no obligation to correct any errors contained herein or to advise any user of liability for the accuracy or correctness of information provided herein to a user. LSI makes no commitment to update the Document. LSI reserves the right to change these legal terms and conditions from time to time at its sole discretion. In the case of any violation of these rules and regulations, LSI reserves the right to seek all remedies available by law and in equity for such violations. Except as expressly provided herein, LSI and its suppliers do not grant any express or implied right to you under any patents, copyrights, trademarks, or trade secret information. Other rights may be granted to you by LSI in writing or incorporated elsewhere in the Document.

## **Trademark Acknowledgments**

Engenio, the Engenio design, MegaRAID, HotScale, SANtricity, and SANshare are trademarks or registered trademarks of LSI Logic Corporation. All other brand and product names may be trademarks of their respective companies.

## **Performance Information**

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of LSI products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they want to purchase.

