



Installation Guide for Solaris™ Cluster 3.2 Software and Oracle® 10g Release 2 Real Application Clusters

*Fernando Castano
June 2007
Sun Microsystems, Inc.*

Abstract: *This document is a detailed step-by-step guide for installing the Solaris 10 11/06 Operating System, Solaris Cluster (formerly Sun Cluster) 3.2 software, the QFS 4.5 cluster file system, and Oracle 10g Release 2 Real Application Clusters (Oracle 10gR2 RAC).*

Copyright © 2007 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

U.S. Government Rights - Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements. Use is subject to license terms. This distribution may include materials developed by third parties.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd. X/Open is a registered trademark of X/Open Company, Ltd.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon architecture developed by Sun Microsystems, Inc.

ORACLE is a registered trademark of Oracle Corporation.

Sun, Sun Microsystems, the Sun logo, Java, Solaris, StorEdge, and Sun Fire are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

This product is covered and controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Table of Contents

1 Introduction.....	4
2 Setup Description.....	4
3 Hardware Stack.....	5
4 Software Stack.....	5
5 What Is Provided With This Installation?.....	6
5.1 Redundancy for Public Interfaces.....	6
5.2 Number of Oracle RAC Nodes.....	6
5.3 Clusterware.....	6
5.4 Redundancy for Private Interfaces.....	7
5.5 Redundancy for I/O Paths.....	7
5.6 Raw Devices.....	7
5.7 Volume Manager.....	7
5.8 Cluster File System (QFS).....	7
5.9 Automatic Storage Manager (ASM).....	7
5.10 Latest Software Versions.....	8
6 Pre-installation Requirements.....	8
6.1 Firmware Update.....	8
6.2 Setup Information.....	8
6.3 Configuring the Storage.....	8
6.4 Documentation and Installation Software.....	8
6.4.1 Documents.....	9
6.4.2 Installation Software.....	9
7 Sun Software Installation.....	10
7.1 Installing the Solaris OS.....	10
7.2 Installing Cluster Control Panel in the Administrative Console.....	11
7.3 Installing Packages for Sun Cluster Framework and Data Service for Oracle RAC.....	13
7.4 Installing QFS Packages.....	13
7.5 Creating a Cluster.....	14
7.6 Preparing for Oracle UDLM Package Installation.....	16
7.7 Installing the UDLM Package.....	17
7.8 Configuring RAC Framework and Solaris Volume Manager Resource.....	17
7.9 Configuring a Solaris Volume Manager Metaset and Metadevices.....	19
7.10 Configuring QFS.....	20
8 Preparing for Oracle RAC Installation.....	23
9 Installing Oracle Clusterware (CRS).....	23
10 Installing the Oracle RDBMS Server.....	25
11 Creating a Test Database.....	27
12 About the Author.....	28
13 For More Information.....	28
14 Appendix A: Server Information Table.....	29
15 Appendix B: Storage Information Table.....	30

1 Introduction

This document is a detailed step-by-step guide for installing the Solaris™ 10 11/06 Operating System (also called Update 3), Solaris Cluster (formerly Sun™ Cluster) 3.2 software, the QFS 4.5 cluster file system, and Oracle 10g Release 2 Real Application Clusters (Oracle 10gR2 RAC, hereafter referred to as Oracle RAC). This article also provides detailed instructions on how to configure QFS and Solaris Volume Manager so they can be used with Oracle RAC.

The procedures described here can be used for SPARC® or x64 installations with up to eight nodes using IP over InfiniBand (IPoIB) or Gigabit Ethernet (GbE) for the private interface and with any level of I/O redundancy. After you complete the procedures, you can use any combination of raw devices, Solaris Volume Manager, Automatic Storage Manager (ASM), or QFS cluster file systems to store data files. This document does not replace the official documentation for Sun Cluster software, QFS, Solaris Volume Manager, or Oracle RAC.

It is assumed that readers are familiar with Oracle RAC, its installation, and its underlying concepts, such as voting device, private interface, optical character recognition (OCR), and Virtual IP (VIP). It also is assumed that readers have a fair understanding of all the software components used during the installation as well as UNIX®.

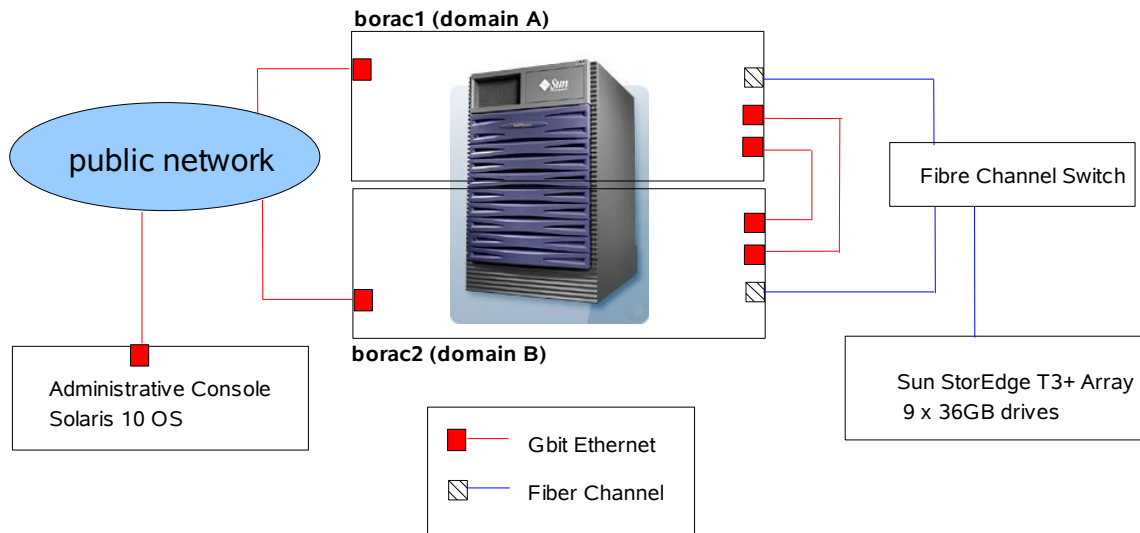
This document uses a two-node setup to show the installation process, but the same procedure can be used for setups with a different number of nodes and different hardware components.

2 Setup Description

The setup used in this installation was very simple: A Sun Fire™ 4800 Server was sub-domained in two domains. Each domain was used as a node for the two-node Sun Cluster software and Oracle RAC installation. A single Sun StorEdge™ T3+ array was connected to a fibre channel (FC) switch and used for shared storage. This setup is referred to as the *borac setup* in the rest of this document.

Extrapolating this procedure to setups with more storage, more servers, or both, or extrapolating to different types of Sun servers or architectures is trivial in most cases. Simple, mature hardware was selected on purpose, based on the assumption that this would be the lowest common denominator for many variants. Considerations for the most common variants are addressed through the document.

Figure 1: Diagram of the Borac Setup



3 Hardware Stack

Table 1 identifies the hardware stack used in this project.

Table 1: Hardware Stack

Attribute	Node 1	Node 2
Server	Sun Fire 4800 Server (domain A)	Sun Fire 4800 Server (domain B)
CPU	6x USIII Cu 1.2GHz	6x USIII Cu 1.2GHz
RAM	16GB	16GB
Host channel adaptor (HCA)	1 dual-port HCA 2Gbps	1 dual-port HCA 2Gbps
Ethernet ports	ce0, c31, ce2, and ce3	ce0, ce1, ce2, and ce3
Storage	Sun StorEdge T3+ array with 9 x 36GB drives	

4 Software Stack

Table 2 describes the software stack used in this project.

Table 2: Software Stack

Role	Vendor	Product Version
Operating system	Sun Microsystems	Solaris 10 11/06 s10s_u3wos_10 SPARC
Database server	Oracle	Oracle RAC RDBMS server 10.2.0.2 SPARC 64-bit
Clusterware	Sun Microsystems	Sun Cluster 3.2 software
Cluster file system	Sun Microsystems	QFS 4.5
Volume manager	Sun Microsystems	Solaris Volume Manager, which is part of the Solaris 10 OS

5 What Is Provided With This Installation?

The installation procedures described here achieve a complete Oracle RAC installation. This installation addresses the needs of most Oracle RAC installations on the Solaris OS. In addition, it leverages the Solaris OS stack using Sun Cluster 3.2 software, Solaris Volume Manager, and QFS, eliminating the need for any third-party products. This section discusses the different aspects of the installation, what is provided, and how to modify the installation to obtain Remote Access Services (RAS) features.

5.1 Redundancy for Public Interfaces

The redundancy for public interfaces is something desirable in many setups since it avoids single points of failure in the network path for the server's public access. Redundancy for public interfaces was not implemented in this setup but it is totally independent of the procedures described here and can be achieved using IP network multipathing (IPMP). See the IPMP section of *System Administration Guide: IP Services* (<http://docs.sun.com/app/docs/doc/816-4554?q=816-4554-13>) for details. If IPMP is desired, it should be implemented before installing the Sun Cluster software. To configure Oracle RAC so that it uses the redundant public interfaces provided by IPMP, see document ID 283107.1 at the Oracle MetaLink web site (<https://metalink.oracle.com>), requires registration).

5.2 Number of Oracle RAC Nodes

The current setup has only two nodes, but up to eight nodes can be used without any modification.

5.3 Clusterware

Sun Cluster 3.2 software is the cluster solution and it is integrated well with the Solaris OS. Sun Cluster software can bring more robustness to Oracle RAC by providing many advantages over other cluster solutions. For detailed information about those advantages, read the white paper *Sun™ Cluster 3.2 Software: Making Oracle Database 10g R2 RAC Even More "Unbreakable"* (http://www.sun.com/software/whitepapers/solaris10/solaris_cluster.pdf).

5.4 Redundancy for Private Interfaces

This setup provides redundancy for private interfaces because it uses Sun Cluster software. Sun Cluster software requires at least two separate paths for the private interface, and it automatically manages failover and load balancing across the different paths. If more than two nodes are used, then two Ethernet switches are required to avoid a single point of failure.

The current setup uses GbE, but the procedure does not change if IPoIB is used. With IPoIB, two HCA ports per server and InfiniBand (IB) switches are required instead of GbE cards and switches.

5.5 Redundancy for I/O Paths

The setup configuration presented here does not provide redundant paths to the storage. To provide fully redundant paths, each node needs two HCA cards with each connected to a different FC switch, and there needs to be connections from each storage array to each FC switch through different RAID controllers. The procedures described here would not change if any level of I/O redundancy is introduced in the setup because I/O multipathing (MPxIO) is enabled through these procedures and the Solaris OS hides any complexity introduced by redundancy in the I/O paths. Regardless of the level of redundancy, the Solaris OS and the Sun Cluster software always present one device for each shared device, and the failover mechanism is handled automatically by the Solaris OS.

5.6 Raw Devices

Raw devices are the simplest but most cumbersome way of creating a database with Oracle RAC. This setup allows the creation of data files using raw devices in shared storage, if that is desired.

5.7 Volume Manager

In this setup, Solaris Volume Manager is configured so that raw devices in the shared storage can be used to create metadevices, which can be used for data files. Even though Solaris Volume Manager is part of the Solaris OS, it can be used for Oracle RAC only if Sun Cluster software is managing the cluster.

5.8 Cluster File System (QFS)

These procedures also configure and provide a cluster file system that can be used to store data files or any other files. QFS is a generic file system in which any kind of file can be stored. QFS cannot be used for Oracle RAC unless Sun Cluster software is installed and is managing the cluster.

5.9 Automatic Storage Manager (ASM)

ASM can be used to hold database files after this installation. After all the software is installed, ASM can be configured with any raw devices not already used in the shared storage. If the raw devices provided to ASM are already mirrored in hardware by taking advantage of the RAID capabilities of the storage controller, then ASM can be configured for external redundancy. Otherwise, ASM can be configured to mirror the data by creating failing groups and configuring for normal redundancy.

5.10 Latest Software Versions

This setup delivers the latest software stack currently certified and publicly available: the Solaris 10 11/06 OS for SPARC platforms, Sun Cluster 3.2 software, Oracle 10g Release 2 Real Application Clusters, and QFS 4.5. The same stack is available, certified, and supported on the Solaris 10 OS for x64 platforms, so the procedures change very little (and changes are described throughout the document) for x64 platforms.

6 Pre-installation Requirements

Before starting the installation procedures, ensure you complete the steps outlined in this section.

6.1 Firmware Update

Update the firmware version of all your hardware components for a Solaris 10 11/06 OS installation, including storage arrays, FC switches, IB switches, PCI cards, and system controllers. If your hardware is different from the hardware used in this setup, check the Sun Cluster 3.2 hardware collection to make sure your hardware is supported by the Sun Cluster 3.2 software and to verify that there are no special procedures required for your hardware.

6.2 Setup Information

Create tables for your setup with similar information to the server information table and storage information table presented in Appendix A and Appendix B. You will need all this information during the installation and it is better to plan it all before starting your installation.

6.3 Configuring the Storage

Check to make sure that your storage arrays are supported and to determine if there are any special procedures to use: see *Sun Cluster 3.1-3.2 Hardware Collection for Solaris OS (x86 Platform Edition)* at <http://docs.sun.com/app/docs/coll/1142.2>. Read the manuals for your storage arrays to make sure they support MPxIO.

Map the logical unit numbers (LUNs) to the controllers according to the storage information table you created. Make sure you can see all the LUNs in all the nodes. If you see a LUN more than once in the same system, it is because you have redundant paths and MPxIO is not yet enabled. This will be resolved later.

Try to leverage the RAID controllers by creating RAID0+1 or RAID5 LUNs. If you want to eliminate single points of failure, map your LUNs to both controllers in the storage array and connect each controller to a different FC switch. Also, connect to each FC switch from a different host bus adapter (HBA) in each server, so that the HBA does not become a single point of failure.

6.4 Documentation and Installation Software

This section lists the documentation that is referenced through this document. It also lists all the software that you need to use during the installation process. Obtain all the installation software listed here before starting the installation.

6.4.1 Documents

The following documents are referenced throughout the installation procedures:

- System Administration Guide: IP Services:
<http://docs.sun.com/app/docs/doc/816-4554?q=816-4554-13>
- *Sun Cluster Software Installation Guide for Solaris OS*:
<http://docs.sun.com/app/docs?q=819-2970-10>
- *Sun Java™ Enterprise System 5 Installation Guide for UNIX*:
<http://docs.sun.com/app/docs/doc/819-4891?q=Sun+Java+Enterprise+System+5+Installation+Guide>
- *Sun Cluster Data Service for Oracle RAC Guide for Solaris OS*:
<http://docs.sun.com/app/docs?q=819-2981-10>
- *Sun StorEdge QFS Installation and Upgrade Guide*:
<http://docs.sun.com/app/docs?q=819-2757-10>
- *Sun Cluster System Administration Guide for Solaris OS*:
<http://docs.sun.com/app/docs/doc/819-2971?q=Sun+Cluster+System+Administration+Guide+for+Solaris+OS&s=t>

6.4.2 Installation Software

The following software is installed:

- Solaris 10 11/06 OS for SPARC or x64 platforms, both available for download at sun.com (<http://www.sun.com/>)
- Sun Java Availability suite for SPARC or x86 platforms (Solaris Cluster) (suncluster-3_2-ga-solaris-sparc.zip or suncluster-3_2-ga-solaris=x86.zip), both available for download at sun.com (<http://www.sun.com/download/products.xml?id=4581ab9e>)
- Oracle patch 5389391 with the 3.3.4.9 UDLM for SPARC platforms (p5389391_10202_SOLARIS64.zip), available for download at the Oracle Metalink web site (<https://metalink.oracle.com/>) and provided in the installation kit
- QFS packages for SPARC or x86 (StorEdge_QFS_4.5_sparc.iso or StorEdge_QFS_4.5_x64.iso), available for download at sun.com (<http://www.sun.com/download/products.xml?id=4429b168>)
- The following patches, available from the SunSolve web site (<http://sunsolve.sun.com/pub-cgi/show.pl?target=tous>), which are needed for official support of QFS 4.5 on Sun Cluster 3.2 software:
 - 122807-05, if you are installing on the Solaris OS for SPARC platforms
 - 122808-05, if you are installing on the Solaris OS for x64 platforms
- Oracle installation software, available at oracle.com: (<http://www.oracle.com/index.html>)

7 Sun Software Installation

7.1 Installing the Solaris OS

Note: This section builds on procedures in the *Sun Cluster Software Installation Guide for Solaris OS*.

1. Install Solaris 10 11/06 s10s_u3wos_10 Generic_118833-33 (entire software group) on all the servers.

The Solaris 10 OS can be downloaded from sun.com for free. Make sure that you create three slices in the boot disk, as described by the server information table in Appendix A of this document. These slices will be needed for Sun Cluster software and other software components. Do not install any Solaris patches at this point.

2. Enable MPxIO:

- a) After the Solaris OS is installed on all nodes, execute as root the following command on all nodes:

```
/usr/sbin/stmboot -e
```

The nodes reboot.

- b) On all nodes (one at the time) reboot using the following command:

```
boot -- -r
```

- c) Verify that on each node you see one and only one path to each of the LUNs presented by the storage arrays. If that is not the case, solve the problem before going any further.

3. Configure `/.rhosts` so that all nodes and the administrative console (admin console) can use `rsh` as root among themselves:

```
# hostname  
borac1  
# more /.rhosts  
+ root  
#
```

4. Configure shared memory.

On each node, add the following line to `/etc/system` and then reboot each node:

```
set shmsys:shminfo_shmmax=desired_SGA_size_in_bytes
```

For example:

```
# tail /etc/system  
* set ce:ce_taskq_disable=1  
  
* End of lines added by SUNWscr  
* BEGIN SUNWsamfs file system addition  
* DO NOT EDIT above line or the next 2 lines below...  
forceload: fs/samfs  
* END SUNWsamfs file system addition  
  
set shmsys:shminfo_shmmax=10737418240
```

The `shmmax` setting provides a system-wide limit of 10GB shm segment size. The Database Configuration Assistant (DBCA) reads this value to calculate the maximum possible System Global Area (SGA) size.

5. Make the installation software available to all the nodes.

Select one node on which to place the software installation files in `/stage/`. Place all the installation software in `/stage` and share `/stage` through NFS. Mount the installation software in `/stage` on all the other nodes.

```
# hostname
borac1
# share
- /stage rw ""

# hostname
borac2
# grep stage /etc/vfstab
borac1:/stage - /stage nfs - yes soft,bg
```

6. Install the following patch, available from the SunSolve web site, for official support of QFS 4.5 on Sun Cluster 3.2 software:

- 122807-05, if you are installing on the Solaris OS for SPARC platform
- 122808-05, if you are installing on the Solaris OS for x64 platform

7.2 Installing Cluster Control Panel in the Administrative Console

Note: This procedure follows information in the *Sun Cluster Software Installation Guide for Solaris OS*.

This procedure describes how to install the Cluster Control Panel (CCP) software on an admin console. The CCP provides a single interface from which to start the `cconsole(1M)`, `ctelnet(1M)`, and `crlogin(1M)` tools. Each tool provides a multiple-window connection to a set of nodes, as well as a common window. You can use the common window to send input to all nodes at one time.

1. Become superuser on the admin console.

2. Unzip the Sun Java Availability Suite in `/stage`:

Unzip the file `suncluster-3_2-ga-solaris-sparc.zip` or `suncluster-3_2-ga-solaris-x86.zip`, depending on your installation, inside `/stage/sc_rac_install/sw/`.

3. Change to the following directory:

- If installing on a SPARC console:

`/stage/sc_rac_install/sw/Solaris_sparc/Product/sun_cluster/Solaris_10/Packages/`

- If installing on an x86 console:

`Solaris_x86/Product/sun_cluster/Solaris_10/Packages/`

4. Install the SUNWcccon package:

```
adminconsole# pkgadd -d . SUNWcccon
```

5. (Optional) Install the SUNWscman package:

```
adminconsole# pkgadd -d . SUNWscman
```

When you install the SUNWscman package on the admin console, you can view Sun Cluster man pages from the admin console before you install Sun Cluster software on the cluster nodes.

6. Create an `/etc/clusters` file on the admin console:

Add your cluster name and the physical node name of each cluster node to the file:

```
adminconsole# vi /etc/clusters
clustername node1 node2
```

See the `/opt/SUNWcluster/bin/clusters(4)` man page for details.

7. Create an `/etc/serialports` file:

Add an entry for each node in the cluster to the file. Specify the physical node name, the host name of the console-access device, and the port number. Examples of a console-access device are a terminal concentrator (TC), a System Service Processor (SSP), and a Sun Fire system controller. Make sure the consoles are configured for `telnet` and not for `ssh` access.

```
adminconsole# vi /etc/serialports
node1 ca-dev-hostname port
node2 ca-dev-hostname port
node1, node2                (Physical names of the cluster nodes)
ca-dev-hostname             (Host name of the console-access device)
port                        (Serial port number)
```

8. (Optional) For convenience, set the directory paths on the admin console:

a) Add the `/opt/SUNWcluster/bin/` directory to the `PATH`.

b) Add the `/opt/SUNWcluster/man/` directory to the `MANPATH`.

c) If you installed the SUNWscman package, also add the `/usr/cluster/man/` directory to the `MANPATH`.

9. Start the CCP utility:

```
adminconsole# /opt/SUNWcluster/bin/ccp &
```

Click the `cconsole`, `crlogin`, or `ctelnet` button in the CCP window to launch that tool. Alternatively, you can start any of these tools directly. For example, to start `ctelnet`, type the following command:

```
adminconsole# /opt/SUNWcluster/bin/ctelnet &
```

See the procedure "How to Log In to Sun Cluster Remotely" in the "Beginning to Administer the Cluster" section of the *Sun Cluster System Administration Guide for Solaris OS* for additional information about how to use the CCP utility. Also see the `ccp(1M)` man page.

7.3 Installing Packages for Sun Cluster Framework and Data Service for Oracle RAC

Repeat this procedure sequentially on each node. This procedure installs the packages for Sun Cluster framework and the data service for Oracle RAC.

Note: Steps 1 through 3 come from information in the *Sun Cluster Software Installation Guide for Solaris OS*. Steps 4 through 6 come from information in the *Sun Java Enterprise System 5 Installation Guide for UNIX*.

1. Ensure that the display environment of the cluster node is set to display the GUI on the admin console:

a) On the admin console, execute:

```
# xhost +
```

b) On the cluster node, execute:

```
# setenv DISPLAY adminconsole:0.0
```

Note: If you do not make these settings, the installer program runs in text-based mode.

2. Become superuser on the cluster node.

3. Start the installation wizard program:

```
# cd /stage/sc_rac_install_kit/sw/ Solaris_sparc    (OR Solaris_x86)
# ./installer
```

4. Select only Sun Cluster 3.2 and the agent for Sun Cluster support for Oracle RAC.

5. Select all shared resources and accept all the defaults in the next screens of the installer.

The installer cannot configure Sun Cluster software, so it displays a message about it this. You can safely ignore this message.

6. Make sure the directory `/usr/cluster` exists after the installation is completed. If it does, proceed with the next node. If it does not, the Sun Cluster software was not installed correctly. In such a case, see the *Sun Java Enterprise System 5 Installation Guide for UNIX*.

7.4 Installing QFS Packages

Note: This procedure follows information in the *Sun StorEdge QFS Installation and Upgrade Guide*.

1. Place the QFS 4.5 installation software in `/stage/sc_rac_install_kit/sw` and untar it.
2. Become superuser in the first (or next) cluster node.
3. Change to the `/stage/software_kit/qfs-bits/2.10/` directory.
4. Use the `pkgadd(1M)` command to add the `SUNWqfsr` and `SUNWqfsu` packages.

```
# pkgadd -d . SUNWqfsr SUNWqfsu
```
5. Enter `yes` or `y` as the answer to each of the questions.
6. Repeat steps 2 through 5 for the next cluster node.

7.5 Creating a Cluster

Note: This procedure follows information in the *Sun Cluster Software Installation Guide for Solaris OS* and the *Sun Cluster Data Service for Oracle RAC Guide for Solaris OS*.

1. Unplumb all the communication devices (NICs or IPoIB) that the Sun Cluster software will use for the private interface. For each device in each node use the command `ifconfig device-name down unplumb`, and ensure the following:
 - Make sure that this is a dedicated subnet and that no network traffic is present.
 - Make sure the IP address and netmask `172.16.0.0/255.255.248.0` do not conflict with the other networks in the lab. If they do, select another network that does not conflict.
 - Make sure that no `/etc/hostname.dev-name` files exist for the communication devices that will be used by the Sun Cluster software for the private interface. The Sun Cluster software owns these devices and takes care of plumbing and configuring them.
2. Run `/usr/cluster/bin/scinstall` from the node you want to be named last (`borac2`). By running `scinstall` in `borac2`, you guarantee that `borac2` will be recognized as node 2 in the cluster. This follows the same naming convention as the one used for the host names assigned to the nodes.
3. Select the option to create a new cluster.
4. Select a custom installation.
5. When asked for the other nodes, provide them in the reverse order to that which you want them named. The last one you provide become node 1 in the cluster.
6. Provide the name of the communication devices to use in each node for private interface, and provide the switch to which they are connected. (If there are only two nodes, a point-to-point connection can be used instead of a switch.) This information should be in the server information table you already filled out.

7. Allow for automatic quorum selection.

8. Provide the installer with the name of the raw device representing the slice in the boot disk of each node where the global devices file system (`/globaldevices`) will be created. This information should be in the server information table you already filled out.

After a successful installation, all nodes reboot and join the cluster after they boot up.

9. Determine which device was selected for quorum device using the following command:

```
# /usr/cluster/bin/clquorum show
```

10. (Optional) Change the quorum device.

The device selected by the Sun Cluster software for quorum device can be used for any other purpose, since Sun Cluster software does not use any of its cylinders. If you want to move the quorum device to another shared device (maybe for RAS or performance reasons), use the following commands:

```
# /usr/cluster/bin/clquorum add dx (where dx is the global ID of new the device)
# /usr/cluster/bin/clquorum remove dx (where dx is the global device reported by Sun
Cluster software)
```

To find how shared storage maps to Sun Cluster global devices use the following command:

```
# /usr/cluster/bin/scdidadm -l
```

11. Verify the quorum list and status:

a) Run the following command, which should return a list with all the nodes and the quorum device in it:

```
# /usr/cluster/bin/clquorum list
d4
borac2
borac1
```

b) Verify the status of the quorum device to make sure that the cluster has a working quorum setup, where each node has a vote and the quorum has a vote too.

```
# /usr/cluster/bin/clquorum status

=== Cluster Quorum ===

--- Quorum Votes Summary ---

          Needed   Present   Possible
          -
          2         3         3

--- Quorum Votes by Node ---

Node Name      Present   Possible   Status
-----
borac2         1         1         Online
```

```

borac1          1          1          Online

--- Quorum Votes by Device ---

Device Name      Present      Possible      Status
-----
d4                1            1            Online

```

12. To configure Network Time Protocol (NTP) to synchronize time among all cluster nodes, do the following on all nodes and then reboot them:

```

# cp /etc/inet/ntp.conf /etc/inet/ntp.conf.orig
# cp /etc/inet/ntp.cluster /etc/inet/ntp.conf

```

13. Bypass Network Information Service (NIS) name service to allow proper operation of the data service for Oracle RAC.

On each node, modify the following entries in the `/etc/nsswitch.conf` file:

```

passwd:      files nis [TRYAGAIN=0]
group:       files nis [TRYAGAIN=0]
publickey:   files nis [TRYAGAIN=0]
project:     files nis [TRYAGAIN=0]

```

7.6 Preparing for Oracle UDLM Package Installation

Note: If you are installing on an x86 based system, you can skip this section because there is no need for the UNIX Distributed Lock Manager (Oracle UDLM) package on the Solaris OS for x86 based systems.

Now that the cluster has been created and the data service for Oracle RAC is in place, the Oracle UDLM package needs to be installed. This package is the "connector" between Oracle RAC and the Sun Cluster software. There is one version of this package for Oracle 10.2.0.1 and another version for Oracle 10.2.0.2. Since the package for Oracle 10.2.0.2 is backwards compatible with the one for 10.2.0.1, you will install only the UDLM package for 10.2.0.2. The UDLM package is usually included in the Oracle tarball for Oracle Clusterware (formerly called CRS for "Cluster Ready Services"). Unfortunately, Oracle forgot to place it in the 10.2.0.2 version of Oracle Clusterware for SPARC, so they released a patch that contains it. The patch number is 5389391 and it contains the 3.3.4.9 UDLM (included in the installation kit).

1. Create the DBA group on all nodes.

```

# groupadd -g 120 dba

```

2. (Optional) Create `/export/home` on all nodes.

```

# mkdir /export/home

```

3. Create the Oracle user on all nodes.

```

# useradd -u 5000 -g dba -d /export/home/oracle -m -s /bin/csh oracle

```


4. Define a password for the Oracle account on all nodes.

```
# passwd oracle
```

5. Create `/var/opt/oracle` on all nodes.

```
# mkdir /var/opt/oracle
# chown oracle:dba /var/opt/oracle
```

6. Obtain the 3.3.4.9 UDLM package.

7.7 Installing the UDLM Package

Note: If you are installing on an x86 based system, you can skip this section because there is no need for the UDLM package on the Solaris OS for an x86 based system.

Perform the following procedure on each node, one at the time.

Note: This procedure follows information in the UDLM readme file.

1. Unpack the file `p57117016_10202_SOLARIS64.zip` into the `/stage/sc_rac_install_kit/sw` directory.

2. Install the patch by adding the package as root:

```
# cd /stage/sc_rac_install_kit/sw/
# pkgadd -d . ORCLudlm
```

7.8 Configuring RAC Framework and Solaris Volume Manager Resource

The RAC framework and Solaris Volume Manager can be configured with command line interface (CLI) commands or with an interactive text menu using `clsetup`. This document uses CLI commands, but if you want to use the menu driven process, you can follow the directions in the *Sun Cluster Data Service for Oracle RAC Guide for Solaris OS*.

Note: This procedure follows information in the *Sun Cluster Data Service for Oracle RAC Guide for Solaris OS*.

1. Become superuser on any of the nodes.

2. Create a scalable resource group.

```
# clresourcegroup create -s my-rac-fmwk-rg
```

3. Register the `SUNW.rac_framework` resource group.

```
# clresourcetype register SUNW.rac_framework
```

4. Add an instance of the `SUNW.rac_framework` resource type to the resource group just created.

```
# clresource create -g my-rac-fmwk-rg -t SUNW.rac_framework my-rac-fmwk-rs
```

5. Register the `SUNW.rac_udlm` resource type.

```
# clresourcetype register SUNW.rac_udlm
```

6. Add an instance of the `SUNW.rac_udlm` resource type to the resource group just created.

```
# clresource create -g my-resource-group \  
-t SUNW.rac_udlm \  
-p resource_dependencies=my-rac-fmwk-rs my-rac-udlm-rs
```

7. Register and add an instance of the resource type that represents the Solaris Volume Manager.

```
# clresourcetype register SUNW.rac_svm
```

8. Add an instance of the `SUNW.rac_svm` resource type to the resource group just created.

```
# clresource create -g my-rac-fmwk-rg \  
-t SUNW.rac_svm \  
-p resource_dependencies=my-rac-fmwk-rs my-rac-svm-rs
```

9. Bring online and in a managed state the RAC framework resource group and its resources.

```
# clresourcegroup online emM my-rac-fmwk-rg
```

10. Verify that the resource group and the resources are online.

```
# /usr/cluster/bin/scstat -g  
  
-- Resource Groups and Resources --  
  
      Group Name      Resources  
-----  
Resources: rac-framework-rg rac-framework-rs rac-udlm-rs rac-svm-rs  
  
-- Resource Groups --  
  
      Group Name      Node Name      State      Suspended  
-----  
Group: rac-framework-rg borac2      Online      No  
Group: rac-framework-rg borac1      Online      No  
  
-- Resources --  
  
      Resource Name    Node Name      State      Status Message  
-----  
Resource: rac-framework-rs borac2      Online      Online  
Resource: rac-framework-rs borac1      Online      Online  
  
Resource: rac-udlm-rs      borac2      Online      Online  
Resource: rac-udlm-rs      borac1      Online      Online  
  
Resource: rac-svm-rs      borac2      Online      Online  
Resource: rac-svm-rs      borac1      Online      Online
```

7.9 Configuring a Solaris Volume Manager Metaset and Metadevices

Note: This procedure follows information in the *Sun Cluster Data Service for Oracle RAC Guide for Solaris OS*.

1. Create a multi-owner disk set (metaset) from one node.

Solaris Volume Manager allows the creation of metadevices. To use these metadevices for Oracle RAC, a metaset needs to be created first. You need to include all the nodes in the cluster in this command and the name you want to give to the metaset (`racvm` in this example). For the borac setup, the command is:

```
#metaset -s racvm -M -a -h borac1.acme.com borac2.acme.com
```

2. Add raw devices to the metaset.

You can now add raw devices in the shared storage to this metaset, and then create metadevices with them. Look at the server information table to identify the devices you want to use with Solaris Volume Manager. Use the global device ID (DID) instead of the raw device name since DIDs are invariant across the cluster. For example, `borac1` presents the following global devices:

```
# /usr/cluster/bin/scdidadm -l
3      borac1:/dev/rdisk/c1t50020F2300002B39d1 /dev/did/rdsk/d3
4      borac1:/dev/rdisk/c1t50020F2300002B39d0 /dev/did/rdsk/d4
5      borac1:/dev/rdisk/c0t1d0                /dev/did/rdsk/d5
6      borac1:/dev/rdisk/c0t6d0                /dev/did/rdsk/d6
```

According to the storage information table in Appendix B, slice 4 of LUN2 can be used for Solaris Volume Manager. To add it to the metaset, you would do the following:

```
# metaset -s my_rac_vm -a /dev/did/dsk/d4s4
```

You can add as many raw devices in shared storage as needed, and using the `metainit` command, you can create metadevices with the raw devices in the metaset.

3. Verify the metaset.

To confirm that the metaset was correctly created and that it has all the raw devices you gave it, you can issue the `metaset` command. Here is the output for the borac setup:

```
# metaset

Multi-owner Set name = racvm, Set number = 1, Master = borac2

Host          Owner          Member
  borac2      multi-owner    Yes
  borac1      multi-owner    Yes

Driv Dbase

d4    Yes
```

The Solaris Volume Manager metaset is automatically registered with the Sun Cluster software. See *Sun Cluster Data Service for Oracle RAC Guide for Solaris OS* for details.

```

# scstat -D
-- Device Group Servers --

                Device Group          Primary          Secondary
                -----          -
-- Device Group Status --

                Device Group          Status
                -----          -

-- Multi-owner Device Groups --

                Device Group          Online Status
                -----          -
Multi-owner device group:  racvm          borac1,borac2

```

7.10 Configuring QFS

QFS is a generic cluster file system that can hold any kind of files. You can install the Oracle home in it if you like, use it only for database files, or store any other kind of information in it. To configure QFS, follow these steps:

1. Create on all nodes, a directory where you can mount the cluster file system.

```
# mkdir -p /db_qfs/Data
```

2. Ensure the directory is mounted after a reboot by adding these lines to `/etc/vfstab` on all nodes.

```

# RAC on perf shared QFS
#
Data    -      /db_qfs/Data    samfs    -      no      shared,notrace

```

3. Create `/etc/opt/SUNWsamfs/mcf` on all nodes.

Here is a brief explanation of what the entries in this file mean:

-ma: Represents the cluster file system name. You need only one line with `ma` and the name (`Data` in this example).

-mm: Represents the device or devices for storing QFS metadata. You need at least one but you should have more for redundancy. Use the DID provided by the Sun Cluster software.

-mr: Represents devices where you want QFS to store data. You should give QFS all the devices you want it to use for data storage (look at the storage information table to identify the devices for QFS). QFS puts all these devices in a pool and stripes them to create the cluster file system on top. Here is an example of the `mcf` file for the borac setup:

```

# hostname
borac1
# more /etc/opt/SUNWsamfs/mcf
Data      2      ma      Data    on      shared
/dev/did/dsk/d3s0    20     mm      Data    on
/dev/did/dsk/d3s3    21     mr      Data    on

```

```
# hostname
borac2
# more /etc/opt/SUNWsamfs/mcf
Data                2      ma      Data    on      shared
/dev/did/dsk/d3s0   20      mm      Data    on
/dev/did/dsk/d3s3   21      mr      Data    on
```

4. Create /etc/opt/SUNWsamfs/hosts.Data on all nodes.

In this file, you define which node is the QFS server and which nodes are backups (all the rest). Usually Node1 is defined as the manager and the rest as backups.

```
# hostname
borac1
# more /etc/opt/SUNWsamfs/hosts.Data
borac1 clusternode2-priv 1 0 server
borac2 clusternode1-priv 1 0
```

```
# hostname
borac2
# more /etc/opt/SUNWsamfs/hosts.Data
borac1 clusternode2-priv 1 0 server
borac2 clusternode1-priv 1 0
```

Notice that the host name needs to be mapped to the Sun Cluster node name. In the borac setup, the nodes names did not map in the same order (they were not presented in the correct order during `scinstall`). It is for that reason that `borac1` maps to `clusternode2-priv` instead of to `clusternode1-priv`. If you want to be sure you are mapping them correctly, you can find the Sun Cluster ID for any node by looking at the file `/etc/cluster/nodeid`.

```
# hostname
borac2
# more /etc/cluster/nodeid
1
```

5. Create /etc/opt/SUNWsamfs/samfs.cmd on all nodes.

Here is the file used in the borac setup. Copy this file onto all your nodes.

```
# hostname
borac2
# more /etc/opt/SUNWsamfs/samfs.cmd
stripe=1
sync_meta=1
mh_write
qwrite
forcedirectio
nstreams=1024
notrace
rdlease=300
wrlease=300
aplease=300
```

6. Create the file system on the node you defined as server in Step 4 (borac1 in this setup).

```
# /opt/SUNWsamfs/sbin/sammkfs -S -a 64 Data
# mount /db_qfs/Data
# chown oracle:dba /db_qfs/Data
```

7. Mount the file system on all other nodes:

```
# mount /db_qfs/Data
```

8. In all nodes, check the file system:

```
# hostname
borac2
# df -k | grep Data
Data                212324352 2688192 209636160    2%    /db_qfs/Data
```

9. Create the QFS metadata server (MDS) resource group for high availability. From node 1 do:

```
# clresourcetype register SUNW.qfs
# clresourcegroup create -p nodelist=pwitch1,pwitch2 Data-rg
# clresource create -t SUNW.qfs -g Data-rg -p QFSFileSystem=/db_qfs/Data
# clresourcegroup online -emM Data-rg
```

At this point the cluster resource group configuration looks like this:

```
# /usr/cluster/bin/scstat -g
```

```
-- Resource Groups and Resources --
```

	Group Name	Resources
Resources:	rac-framework-rg	rac-framework-rs rac-udlm-rs rac-svm-rs
Resources:	Data-rg	Data-rs

```
-- Resource Groups --
```

	Group Name	Node Name	State	Suspended
Group:	rac-framework-rg	borac2	Online	No
Group:	rac-framework-rg	borac1	Online	No
Group:	Data-rg	borac1	Online	No
Group:	Data-rg	borac2	Offline	No

```
-- Resources --
```

	Resource Name	Node Name	State	Status Message
Resource:	rac-framework-rs	borac2	Online	Online
Resource:	rac-framework-rs	borac1	Online	Online
Resource:	rac-udlm-rs	borac2	Online	Online
Resource:	rac-udlm-rs	borac1	Online	Online
Resource:	rac-svm-rs	borac2	Online	Online
Resource:	rac-svm-rs	borac1	Online	Online
Resource:	Data-rs	borac1	Online	Online - Service is online.
Resource:	Data-rs	borac2	Offline	Offline

10. (Optional) Configure OCR and voting in QFS.

If you plan to place the Oracle OCR and voting devices inside the QFS file system, follow this example on one of the nodes:

```
# mkdir /db_qfs/Data/crs_files
# mkdir /db_qfs/Data/db_files
# cd /db_qfs/Data/crs_files
# dd if=/dev/zero of=ocr bs=1024k count=120
# dd if=/dev/zero of=voting bs=1024k count=120
# chown -R oracle:dba voting
# chown -R root:dba ocr
# chmod -R 660 voting
# chmod -R 640 ocr
```

8 Preparing for Oracle RAC Installation

This section describes how to prepare a system for Oracle RAC installation with Oracle Clusterware (formerly called CRS for "Cluster Ready Services"). If other clusterware is being used, please follow the specific documentation for that product. Before proceeding with the installation steps make sure that the following hardware requirements are satisfied.

Unless otherwise stated, all the commands described here must be executed as root.

1. Clean up OCR and voting devices.

If you are not placing the OCR and voting devices in the QFS file system, do the following in the raw devices that will be used for this purpose:

```
# dd if=/dev/zero of=ocr_dev bs=1024k count=120
# dd if=/dev/zero of=voting-dev bs=1024k count=120
# chown -R oracle:dba voting-dev
# chown -R root:dba ocr-dev
# chmod -R 660 voting-dev
# chmod -R 640 ocr-dev
```

2. Allow Oracle users to use rsh on all systems.

Execute this command as Oracle user on all the servers that will be part of the RAC cluster.

```
# echo "+ oracle" > /export/home/oracle/.rhosts
```

3. Confirm that you can use rsh as Oracle.

For all nodes, execute the following command as Oracle user against all other nodes and confirm that there are no problems reported:

```
# rsh other-node hostname
```

9 Installing Oracle Clusterware (CRS)

1. Obtain the Oracle Clusterware tarball (for x86/x64 or SPARC systems) and place it in /stage/crs.

2. For x86/x64 systems, run the `rootpre.sh` script as root on each node:

```
# /stage/crs/rootpre/rootpre.sh
```

2. Run the following command.

```
# su - oracle
```

3. Change to the `/stage/crs/` directory.

4. Set the display to the system where you want to display the installer GUI.

5. Run the installer.

```
# ./runInstaller
```

For x86/64 systems, the installer prompts you for confirmation that the `rootpre` script has been executed. Answer `yes`.

6. Accept the path for the inventory directory and the group (`dba`), and click Next.

7. Set the name to `crs` and the path to `/export/home/oracle/crs`, and click Next.

8. See if there are any failures, and then click Next.

9. Verify that public, virtual, and private names resolve in `/etc/hosts`, NIS, or Sun Cluster database, and then click Next.

10. Set the public IP/device to public. Set the private IP/device to private. Select `clprivnet0` device for private and leave the ones used by the Sun Cluster software as "do not use."

11. Set OCR to the device selected in the storage table using the DID identifier (or the QFS file). Use external redundancy if possible.

12. Set voting to the device selected in the storage table using the DID identifier (or the QFS file). Use external redundancy if possible.

13. Click Install.

14. Run scripts as root in the order and on the nodes indicated.

15. If VIP fails while running scripts on the last node (known bug), do the following:

a) Set the `DISPLAY` variable appropriately.

b) Execute `# /export/home/oracle/crs/bin/vipca`.

c) Fill in again the VIP information, and `vipca` creates and starts the VIP, GSD (global services daemon), and ONS (Oracle name server) resources under Oracle Clusterware.

16. Deal with the GDS bug.

There is an Oracle bug related to the Sun Cluster software that manifests itself as the GDS resource not coming up. This resource is no longer needed in Oracle 10g, so it can be disabled. To avoid problems do the following:

a) With 10.2.0.1 running, unregister the GSD resources.

b) Change to the `$CRS_HOME/bin` directory.

c) Execute `# ./crs_stat | grep NAME`.

This command gives you the name of the resources. For the borac setup, this is the output:

```
# hostname
borac1
# ./crs_stat | grep NAME
NAME=ora.TST.TST1.inst
NAME=ora.TST.TST2.inst
NAME=ora.TST.db
NAME=ora.borac1.LISTENER_BORAC1.lsnr
NAME=ora.borac1.gsd
NAME=ora.borac1.ons
NAME=ora.borac1.vip
NAME=ora.borac2.LISTENER_BORAC2.lsnr
NAME=ora.borac2.gsd
NAME=ora.borac2.ons
NAME=ora.borac2.vip
```

d) As root on one node, execute `# crs_stop gsd-resource-name` for each of the GSD resources presented before. There should be one for each node.

e) As root on one node, execute `# crs_unregister gsd-resource-name` for each of the GSD resources presented before.

f) Reboot all the nodes. This step is necessary to make sure that the UDLM instance is running without the GSD processes impacting it.

10 Installing the Oracle RDBMS Server

1. Obtain the Oracle database 10.2.0.1.

2. Create the directory `/stage/10gdb` on the database server and place the database installation software there.

3. Change to the `/stage/10gdb` directory and unpackage the tarball.

4. Run the following command:

```
# su - oracle
```

5. Change to the `/stage/10gdb` directory.

6. Set the display to the system where you want to display the installer GUI.

7. Run the installer with the following command:

```
# ./runInstaller
```

8. Select Enterprise Edition, and click Next.

9. Define Oracle home as `/export/home/oracle/db`. Clear the check box for "create starter database," and click Next.

10. Select all nodes, and click Next.

11. Leave the inventory in `/export/home/oracle/oraInventory` and leave the operating system group name as `dba`. Click Next.

12. Disregard the failure for `nonexec_user_stack=1`. If there are no other warnings, click Next and click Yes in the warning popup dialog.

13. Select the option "install database software only" and click Next.

14. Click Install and wait until the installation finishes.

15. As root, execute the two commands presented by the installer, and click OK.

16. Exit the installer.

17. Now it is necessary to install the patchset 2. Obtain the tarball for this patchset. For x86/64 systems, get Oracle 10.2.0.2 Database Patchset. For SPARC systems, get Oracle 10.2.0.2 Patchset.

18. Create the directory `/stage/ps2` and place the tarball in it.

19. Unpackage the tarball.

20. Run the following command:

```
# su - oracle
```

21. Set the display to the system where you want to display the installer GUI.

22. Change to the `/stage/ps2/Disk1` directory.

23. Install the patch on CRS home first, and then install it again on database home.

24. Execute the following command.

```
# ./runInstaller
```

25. Click Next.

26. Select CRS or db for the database (first install on Oracle Clusterware and then on database).

27. Click Install.

28. Repeat installation of patch 10.2.0.2 on database home.

29. As Oracle user, add the following entries to the file `/export/home/oracle/.cshrc`.

```
setenv ORACLE_BASE $HOME
setenv ORACLE_HOME $ORACLE_BASE/db
setenv CRS_HOME $ORACLE_BASE/crs          (if RAC present)
setenv PATH $ORACLE_HOME/bin:$PATH
setenv LD_LIBRARY_PATH $ORACLE_HOME/lib
setenv ORACLE_SID SID_for_your_database
```

30. (SPARC only) Install patch 5117016 after patch 10.2.0.2 on Oracle home only (mandatory).

31. Apply the workaround described in document ID 363840.1 at the Oracle MetaLink web site (<https://metalink.oracle.com>, requires registration), which involves the following:

a) Ensure all instances are shut down.

b) Manually copy library file `libknlopt.a` from the local node to the same directory on all remote nodes. On some ports, this library exists in `$ORACLE_HOME/rdbms/lib` and in other ports, this library exists in `$ORACLE_HOME/lib`. It is critical to copy the library file to the same location on all remote nodes as it exists on the installation node.

c) Relink Oracle on all remote nodes (it is not necessary to relink on the node where the installation was performed):

```
cd $ORACLE_HOME/rdbms/lib
make -f ins_rdbms.mk rac_on ioracle
```

32. Apply performance patches, which are downloadable from the Oracle MetaLink web site (<https://metalink.oracle.com>):

- For x86 based systems only, apply performance patch 4770693.
- For SPARC based systems only, apply patch 5559916 to improve the log writer daemon (LGRWR) performance.

33. Check the Oracle MetaLink web site for any other patches required for your installation.

11 Creating a Test Database

Using DBCA, create a test database. You can use raw devices, Solaris Volume Manager metadevices, ASM, or the QFS cluster file system to store the database files. If you decide to use ASM, configure it by providing the Sun Cluster DID instead of the `/dev/rdsk/*` path, since that path is not always constant across nodes. The Sun Cluster DID path is the same on all nodes.

12 About the Author

Fernando Castano is the Oracle technical project lead in the Sun ISV Engineering Organization. The mission of this group involves porting ISV applications to Sun platforms, certifying them on platforms, tuning them so they run better, benchmarking them and publishing the results, integrating them with Sun products (and certifying the integration), and creating collateral about using partner applications.

13 For More Information

- Solaris Cluster Hub on BigAdmin System Administration Portal
<http://www.sun.com/bigadmin/hubs/sc/>
- Blog on Sun Cluster Oasis – Oracle on Sun Cluster
http://blogs.sun.com/SC/entry/oracle_on_sun_cluster
- Training: Sun Cluster 3.2 Administration (ES-345)
<http://www.sun.com/training/catalog/courses/ES-345.xml>
- Training: Sun Cluster 3.2 Advanced Administration (ES-445)
<http://www.sun.com/training/catalog/courses/ES-445.xml>

14 Appendix A: Server Information Table

Attribute	Node 1	Node 2	Comments
Server type	Sun Fire 4800 Server domain A	Sun Fire 4800 Server domain B	
Host name	borac1	borac2	
CPUs	6x USIII Cu 1.2GHz	6x USIII Cu 1.2GHz	
Host IP address	10.6.222.11	10.6.222.67	
Root password	welcome	welcome	It helps if they are all the same
RAM	16GB	16GB	
Shared memory	10GB	10GB	At least 5GB less than total RAM
VIP alias	borac1-vip	borac2-vip	
VIP IP address	10.6.222.5	10.6.222.9	
OS version	Solaris 10 11/06	Solaris 10 11/06	Solaris 10 11/06 s10s_u3wos_10
Console IP address	10.6.222.84 (domain A)	10.6.222.84 (domain B)	Configure for telnet instead of ssh
Console password	welcome	welcome	It helps if they are all the same
HCA's	1 dual-port HCA	1 dual-port HCA	
NIC for public and VIP	ce0	ce0	Two NICs can be used with IPMP to provide redundancy
NICs for private interface	ce2, ce3	ce2, ce3	
QFS metadata server	yes	no	Choose one for server, the rest as backups
Boot disk name	c0t1d0	c0t0d0	Might not be the same, but that is OK
Boot disk, slice 0	27GB	27GB	Root file system
Boot disk, slice 1	4GB	4GB	Swap
Boot disk, slice 3	1GB	1GB	For Sun Cluster global DID
Boot disk, slice 4	1GB	1GB	Reserved for future use
Boot disk, slice 5	1GB	1GB	For Solaris Volume Manager metadb (add more if possible in extra local disks)
Boot disk, slice 6	0GB	0GB	Not used
Boot disk, slice 7	0GB	0GB	Not used

15 Appendix B: Storage Information Table

Attribute	Value	Comments
Array type	Sun StorEdge T3+	
Alias	t3-a	
Console IP address	18.1.1.1	
World Wide Name	50020f2000002b39	Needed to identify LUNs in the Solaris OS
Firmware version	3.2.5	
Enclosure firmware version	5.02	
Disk size	36GB	
Console password	welcome	
Disk 1 map	LUN1	
Disk 2 map	LUN1	
Disk 3 map	LUN1	
Disk 4 map	LUN1	
Disk 5 map	LUN2	
Disk 6 map	LUN2	
Disk 7 map	LUN2	
Disk 8 map	LUN2	
Disk 9 map	Spare for both LUNs	
LUN1 type	RAID 0+1	Sun Cluster quorum goes in LUN1. It requires no slice (maps to /dev/rdisk/c1t50020F2300002B39d0).
LUN2 type	RAID 0+1	Maps to dev/rdisk/c1t50020F2300002B39d1
LUN1, slice 0	1GB	Oracle voting device
LUN1, slice 1	1GB	To be used as raw devices or by Solaris Volume Manager
LUN1, slice 3	35GB	QFS data
LUN2, slice 4	35GB	To be used as raw devices or by Solaris Volume Manager
LUN2, slice 0	1GB	QFS metadata
LUN2, slice 1	1GB	Oracle OCR
LUN2, slice 3	35GB	QFS data
LUN2, slice 4	35GB	To be used as raw devices or by Solaris Volume Manager

LUN1 DID	d4	Global DID assigned by the Sun Cluster software after installation
LUN2 DID	d3	Global DID assigned by the Sun Cluster software after installation

Licensing Information

Unless otherwise specified, the use of this software is authorized pursuant to the terms of the license found at http://www.sun.com/bigadmin/common/berkeley_license.html