

Installing Linux and Oracle9i RAC on IBM @server BladeCenter™

January 22, 2004

By:

Stephen R. Poon



**IBM / Oracle International Competency Center
San Mateo, California**



Contents

Notices	3
Trademarks	3
Special Notice	3
Acknowledgements	3
Introduction.....	4
About IBM ^ BladeCenter	4
About IBM TotalStorage FAStT900 Storage Server	5
Configuration	5
Hardware	5
Firmware Versions	7
Software.....	7
Linux Operating System	7
Fibre Channel Expansion Card Driver	7
IBM FAStT Storage Manager for x86 Linux	8
IBM FAStT Management Suite Java (MSJ) Diagnostic and Configuration Utility for Linux	8
Oracle9i.....	8
Setting Up the Hardware	8
BladeCenter	8
Fibre Channel Cabling	13
SAN Availability Considerations	16
FAStT900 Link Rate Switches	18
FAStT EXP700 Enclosure ID Switches	19
BladeCenter HS20.....	19
Setting Up Operating Environment.....	21
Installing Red Hat Enterprise Linux AS 2.1.....	21
Network.....	27
Fibre Channel Adapter Driver	28
Shared Storage.....	30
Data Files for Oracle9i RAC.....	31
Installing Oracle9i RAC	32
Pre-installation Tasks	32
Install Oracle Cluster Manager	33
Install Oracle9i RAC	42
Install Patch Set 9204.....	49
Obtain Patch Set.....	49
Oracle Universal Installer Patch.....	49
Cluster Manager Patch	53
Oracle9i RAC Patch	57
Post Installation	59
Configure Listener.....	59
Creating a Test Database	64
Using srvctl to manage RAC Instances.....	74
Adding a Node to Our Existing RAC Cluster	75
Summary	81
Appendix A	82
Appendix B.....	83
Adapter Port Name in Fast!UTIL	83
References	85

NOTICES

This document reflects the IBM / Oracle International Competency Center's understanding about Oracle products running on IBM hardware servers and storage. It was produced and reviewed by the members of the IBM organization called the IBM / Oracle International Competency Center and others.

This document is presented "As-Is" and IBM does not assume responsibility for the statements expressed herein. It reflects the experiences of the IBM / Oracle International Competency Center. If you have questions about the contents of this document, please direct them to the IBM / Oracle International Competency Center (ibmoracl@us.ibm.com).

TRADEMARKS

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

IBM®
IBM eServer™
BladeCenter™
xSeries™

Oracle is a registered trademark, and Oracle9i, Oracle*MetaLink*, and SQL*Plus are trademarks or registered trademarks of Oracle Corporation.

Linux is a registered trademark of Linus Torvalds.

Red Hat, the Red Hat "Shadow Man" logo, RPM and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc. in the United States and other countries.

SuSE is a registered trademark of SuSE Linux AG.

All other trademarks mentioned herein are the property of their respective owners.

SPECIAL NOTICE

The material in this document was produced by the IBM / Oracle International Competency Center based on Oracle, Red Hat and SuSE products and documentation. The material herein is copyrighted by IBM, Oracle, Red Hat, and SuSE.

ACKNOWLEDGEMENTS

Special thanks to the following people for their support and assistance with this project:

Robert F. Nelson	BladeCenter Product Marketing Manager
John A. McAbel	WW xSeries Product Marketing Manager

Thanks to Kathryn Arrell, Tim Spencer and Pat Blaney of the IBM/Oracle International Competency Center for their assistance and contribution to this project.

INTRODUCTION

This document is intended to help those who are installing Oracle9i Real Application Clusters (RAC) on an IBM eServer BladeCenter. It is based on experiences gained by installing a test environment at the IBM / Oracle International Competency Center and available documentation from IBM, Oracle, and Red Hat.

In the sections that follow, we describe activities needed to prepare for and successfully install Red Hat Enterprise Linux AS 2.1 and Oracle9i RAC on BladeCenter servers. Our shared storage subsystem consisted of an IBM FASTT900 Storage Server and two IBM FASTT EXP700 Storage Expansion Units. First, we describe the installation of a 2-node RAC cluster, and then the procedure for adding a third node.

ABOUT IBM eSERVER BLADECENTER

The unique BladeCenter design addresses today's customers' most serious issues: space constraints, efficient manageability, resiliency, and the physical environment, which includes cooling and power. The BladeCenter servers take less time to install, fewer people to manage and maintain, and cost less than traditional server solutions. These blade servers are so compact -- just about an inch thick -- and so easy to use that customers can increase the performance of their systems by simply sliding an additional BladeCenter server into the integrated chassis. It's just like putting books on a shelf. Since BladeCenter servers share a common, integrated infrastructure, basic components such as power, system ports and fans, they also lower power consumption and reduce system complexity.

Here are key features of the BladeCenter:

- IBM eServer BladeCenter provides a shared infrastructure designed to minimize single points of failure. Resiliency is maximized with a high availability midplane, and hot-swappable power and cooling modules.
- IBM eServer BladeCenter offers one of the densest form factors in the industry. Fourteen BladeCenter servers can occupy a chassis only 7U in height. That means a standard full height rack can hold 84 HS20 2-way blades.
- Optional integrated Gigabit Ethernet and Fibre Channel switches offer a broad range of networking options integrated into the chassis to simplify infrastructure complexity and manageability while lowering total cost of ownership.
- The optional Optical Pass-thru Module delivers alternative interconnect for clients needing to incorporate BladeCenter into existing Storage Area Networks and optical Ethernet environments.
- The BladeCenter chassis accepts both Intel and PowerPC processor-based blade servers:
 - Each IBM eServer BladeCenter HS20 server can have up to two Intel® Xeon processors, up to 80GB of IDE storage and 8GB of Chipkill memory.
 - Each IBM eServer BladeCenter JS20 server can have two PowerPC® 970 processors, up to 80GB of IDE storage and 4GB of Chipkill memory.
- Included with IBM eServer BladeCenter is IBM Director 4.1, which provides systems management tools and the ability to manage the entire BladeCenter from a single remote console.
- An optional SCSI storage expansion kit allows up to 144GB of storage in one BladeCenter server slot.

For further information on the IBM eServer BladeCenter, refer to:

http://www.ibm.com/pc/us/eserver/xseries/bladecenter_family

ABOUT IBM TOTALSTORAGE FAST900 STORAGE SERVER

The IBM® TotalStorage® FASt900 Storage Server is a member of the IBM FASt family disk storage products. The FASt900 is an enterprise-class storage server that is designed to provide performance and flexibility for today's demanding data-intensive computing environments.

Highlights of the FASt900 Storage Server include:

- Provides data protection with dual redundant components, multiple RAID levels, LUN masking and enhanced management options
- Supports 2Gb Fibre Channel and Serial ATA disk drive modules
- Supports 2 to 224 hard disk drives and provides more than 32TB of FC scalability or 56TB of Serial ATA scalability
- Includes FASt Storage Manager to help manage the storage subsystem
- Uses 2Gb Fibre Channel connectivity (up to 795MB/sec throughput from disk) for fast, responsive access to data
- Facilitates storage consolidation for Storage Area Network, Network Attached Storage and direct-attach storage environments
- Offers optional, advanced replication functions such as FlashCopy®, VolumeCopy and Remote Mirroring

The FASt900 builds on the capabilities of existing FASt systems, using a common storage management interface and similar disk expansion tools, providing investment protection throughout the FASt family of storage systems.

The FASt900 supports IBM ^ pSeries®, IBM ^ xSeries® and BladeCenter servers, as well as other popular Intel® processor-based, AIX® and UNIX®- based servers and operating systems.

Other servers in the FASt family include:

- The FASt700 delivers superior performance with 2 Gbps Fibre Channel technology. The FASt700 scales from 36GB to over 32TB to support growing storage requirements created by e-business applications.
- The FASt600 is a mid-level storage server that can scale to over eight terabytes of Fibre Channel disk using 3 EXP700s and over sixteen terabytes of Fibre Channel disk with the Turbo feature using 7 EXP100s. The host interface on the base FASt600 is 2Gb, with the Turbo option auto senses to connect to 1Gb or 2Gb.
- The FASt200 is designed for workgroup and departmental servers that require an external storage solution. The single controller model provides a cost-effective solution, while the FASt200 High Availability (HA) model features a fully redundant configuration with dual-active controllers. The FASt200 scales from 18GB to 1.47TB in a compact 3U size with a maximum system capacity of 9.6TB.

For further information on the IBM TotalStorage products, refer to <http://www.storage.ibm.com/index.htm>.

CONFIGURATION

HARDWARE

Our hardware configuration is shown in Figure 1. We installed Oracle9i RAC on two BladeCenter HS20s. On each HS20, one of the two on-board Gigabit Ethernet adapters was used for the production network and the other for the cluster interconnect. Each HS20 had a Fibre Channel Expansion Card for attachment to the FASt900 Storage Server. The FASt900 was attached to two FASt EXP700s, each with fourteen 73.4GB hard drives.

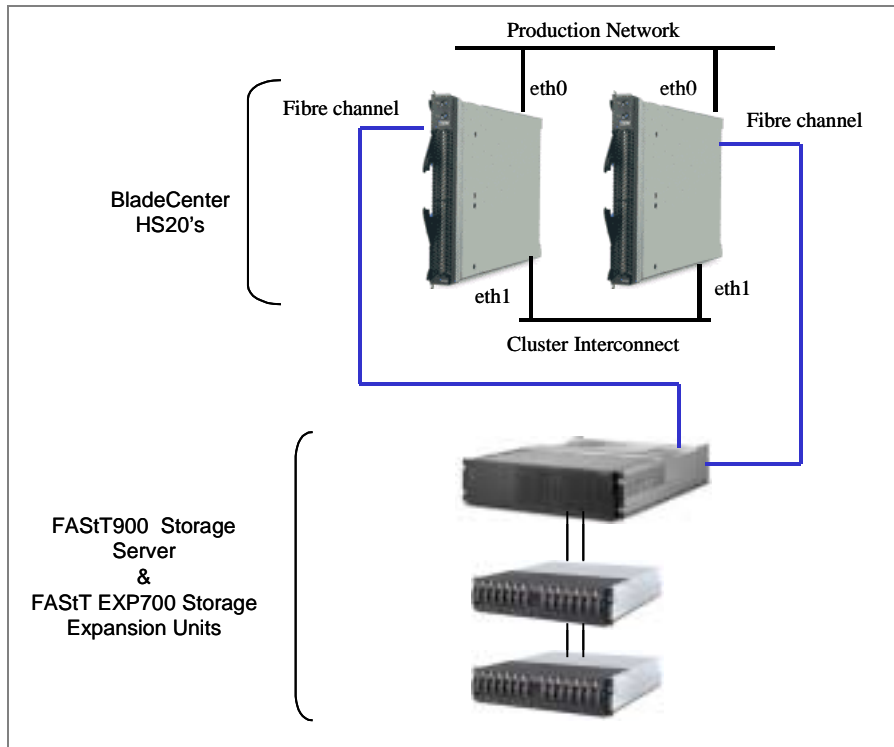


Figure 1. Hardware Configuration

Table 1 lists the parts in our configuration.

Part	Description	Qty
86771XX	IBM eServer BladeCenter™ Chasis	1
48P7062	IBM BladeCenter™ 2-port Fibre Channel Switch Module	1
48P7054	IBM BladeCenter™ 4-port GB Ethernet Switch Module	2
02R9080	IBM BladeCenter™ Optical Passthru Module	1
73P6033	IBM BladeCenter™ Optical Passthru Module LC Cable	1
	Blade 1	
867821X	IBM BladeCenter™ HS20 Xeon 2.0GHz/400MHz, 512KB, 512MB, O/Bay, Rack	1
48P7060	2.0GHz/400Mhz, 512KB L2 Cache Xeon Processor	1
33L5039	1GB PC2100 CL2.5 ECC DDR SDRAM RDIMM	2
48P7061	IBM HS20 Fibre Channel Expansion Card	1
48P7063	IBM HS20 40GB 5400-rpm ATA-100 HDD	1
	Blade 2	
867821X	IBM BladeCenter™ HS20 Xeon 2.0GHz/400MHz, 512KB, 512MB, O/Bay, Rack	1
48P7060	2.0GHz/400Mhz, 512KB L2 Cache Xeon Processor	1
33L5039	1GB PC2100 CL2.5 ECC DDR SDRAM RDIMM	2
48P7061	IBM HS20 Fibre Channel Expansion Card	1
48P7063	IBM HS20 40GB 5400-rpm ATA-100 HDD	1
	Blade 3	
867821X	IBM BladeCenter™ HS20 Xeon 2.0GHz/400MHz, 512KB, 512MB, O/Bay, Rack	1
48P7060	2.0GHz/400Mhz, 512KB L2 Cache Xeon Processor	1
33L5039	1GB PC2100 CL2.5 ECC DDR SDRAM RDIMM	2

48P7061	IBM HS20 Fibre Channel Expansion Card	1
48P7063	IBM HS20 40GB 5400-rpm ATA-100 HDD	1
	Blade 4	
867821X	IBM BladeCenter™ HS20 Xeon 2.0GHz/400MHz, 512KB, 512MB, O/Bay, Rack	1
48P7060	2.0GHz/400Mhz, 512KB L2 Cache Xeon Processor	1
33L5039	1GB PC2100 CL2.5 ECC DDR SDRAM RDIMM	2
48P7061	IBM HS20 Fibre Channel Expansion Card	1
48P7063	IBM HS20 40GB 5400-rpm ATA-100 HDD	1
	FAStT	
174290U	IBM TotalStorage FAStT900 Storage Server	1
17401RU	FAStT700 Storage Expansion Unit	2
32P0768	2Gbps FC 73.4 GB 15K Hot-Swap HDD	28
19K1271	Short Wave SFP Module	12
	LC-LC Fibre Channel Cables	6

Table 1. Configuration parts list

Firmware Versions

The following are the firmware versions in our installation. Check <http://www.ibm.com/support> for the latest versions.

- IBM eServer BladeCenter HS20 - Flash BIOS Update version 1.05
- IBM eServer BladeCenter HS20 - Fibre Channel Expansion Card BIOS version 1.29
- FAStT900 Firmware – 5.30.09.00
- FAStT900 NVSRAM – M1742F900R830V04

SOFTWARE

Linux Operating System

We installed Red Hat Enterprise Linux Advanced Server 2.1 (new name for Red Hat Advanced Server 2.1) for our test. The BladeCenter is also certified for Red Hat Enterprise Linux 3. Besides Red Hat, the BladeCenter is also certified for SuSE Linux Enterprise Server 8 (SLES8).

For the latest information on supported operating systems for the BladeCenter, refer to:

<http://www.pc.ibm.com/us/compat/nos/matrix.shtml>

For the latest information regarding Red Hat and SuSE certifications, refer to:

Red Hat's web site:

<http://hardware.redhat.com/hcl/?pagename=hcl&view=certified&vendor=4&class=8#list>

SuSE's web site:

http://www.suse.com/de/business/certifications/certified_hardware/ibm/bladecenter_8677/

Fibre Channel Expansion Card Driver

We used IBM eServer BladeCenter HS20 Type 8832, 8678 - Fibre Channel Expansion Card drivers version 6.05.00 for Linux. This package is available for download from <http://www.ibm.com/support>.

IBM FAStT Storage Manager for x86 Linux

The FAStT Storage Manager provides a GUI utility for configuring, monitoring, and managing one or more FAStT Storage Servers. FAStT Storage Manager V8.3 is included with the FAStT900.

IBM FAStT Management Suite Java (MSJ) Diagnostic and Configuration Utility for Linux

FAStT MSJ assists you in configuring a SAN environment. This application is specifically designed for IBM FAStT Fibre Channel Host Bus Adapters and the BladeCenter HS20 Fibre Channel Expansion Card (HBA) in a SAN environment.

FAStT MSJ can be used to configure remote systems. This networking capability enables centralized management and configuration of the entire SAN.

IBM FAStT MSJ can be used to manage and control one or more HBAs that are installed on local or remote computers. IBM FAStT MSJ can also be used in diagnosing Fibre Channel (FC) subsystem problems in adapters, controllers, and devices that are attached to the FC fabric or loop. FAStT MSJ also allows you to perform the following operations on the devices in your system:

- Disable (unconfigure) a device on a host bus adapter
- Enable a device
- Designate a path as an alternate or preferred path

Oracle9i

Oracle9i Release 2 is the current release of Oracle's database product on the 32-bit Linux/Intel platform. Version 9.2.0.1 is the installable version, available on CD's or by download from Oracle Technology Network (OTN). 9.2.0.4 is the latest version, available as a patch set, and can be downloaded from Oracle *MetaLink*.

For the latest information on Oracle product certification, refer to:

<http://otn.oracle.com/support/metalink/content.html>

SETTING UP THE HARDWARE**BLADECENTER**

The first step was to provide 220V power to the BladeCenter. Once plugged into the power source, the BladeCenter powers up automatically.

Figure 2 shows a front view of the BladeCenter with four BladeCenter HS20s (blade servers). The control panel on each blade server contains buttons used to switch the keyboard, monitor and mouse to that blade. The control panel also features a cover, which can be flipped down to access the power button for that blade server.




Figure 2. BladeCenter front view

Figure 3 shows a rear view of the BladeCenter. The BladeCenter comes with one hot-swap management module in the management bay, located on the upper right hand corner in Figure 3.




Figure 3. BladeCenter rear view

The management module functions as a system-management processor and a keyboard/video/mouse-multiplexing switch for the blade servers. It provides the following external connections:

- Keyboard, video, and mouse ports for a local console
- One RJ-45 port for a 10/100 Ethernet connection, which provides access to the system management processor

The system management processor communicates with other BladeCenter components, providing functions such as:

- Status monitoring of the blade servers, switch modules, power modules, blower modules
- Blade server management and control, e.g. power/restart, upgrading firmware, switching the keyboard/video/mouse, etc., in conjunction with the blade server service processors
- Switch modules configuration, such as enabling/disabling external ports
- Remote console

Communication with other BladeCenter components occurs over an internal management network. The system management process and managed components are assigned default IP addresses, which can be changed for deployment in a larger network. Though not required for deployment, it is recommended that the management network be on a separate network segment, with access limited only to those responsible for administering the BladeCenters. For our test, we did not change the default IP addresses on the internal management network.

The BladeCenter also supports four I/O modules that enable it to communicate with external resources. These I/O modules are housed in bays 1 through 4. Several I/O modules are available, designed specifically to communicate with a specific environment. The left hand side of Figure 3 shows two optional 4-port Gigabit Ethernet switch modules in bays 1 and 2. These provide connections from the blade servers' integrated Gigabit Ethernet ports to external networks. An optional Fibre Channel Switch module is shown in the bay 3, for connecting the optional Fibre Channel Expansion Cards on the blade servers to external storage, in our case the FAStT900 Storage Server. A second Fibre Channel switch module may be installed in bay 4 to provide connection for the second port on the Fibre Channel Expansion Cards.

The Gigabit Ethernet controller located nearer the bottom of the blade server, routed through switch module 2, is associated with the first Ethernet interface in Linux (eth0). The Ethernet controller located near the top of the blade, routed through switch module 1, is associated with the second Ethernet interface (eth1). This association is summarized in the table below.

Ethernet controller location on blade server	Routed through switch module	Linux interface
Top of blade server	1	eth1
Bottom of blade server	2	eth0

By default, the external ports on the switch modules are disabled. To connect to external networks and shared storage, the external ports on the Ethernet Switch Modules and Fibre Channel Switch Module need to be enabled. However, since we used eth1 on the blade servers for the (private) cluster interconnect, the external ports on Ethernet switch module 1 did not need to be enabled.

Enabling of the switch module external ports was done through the system-management processor as follows:

1. Connect another system or server to the management module Ethernet port. We used a crossover cable to connect a laptop running Windows 2000 to the management module (upper right hand corner shown in Figure 3). Change the laptop's IP address to 192.168.70.124, subnet mask 255.255.255.0, in order to be on the same subnet as the management module's default IP address and subnet. Using Netscape, type the address 192.168.70.125 in the Location field and press Enter to access the BladeCenter Management Module.
2. A popup window is displayed, prompting for a userid and password. Enter:

Userid: USERID
 Password: PASSWORD (zero in place of the letter O)

Be sure that you enter the information in capital letters.

3. The BladeCenter Management Module welcome screen, shown in Figure 4, is displayed. Change the Inactive session timeout value to no timeout, and click **Continue**.

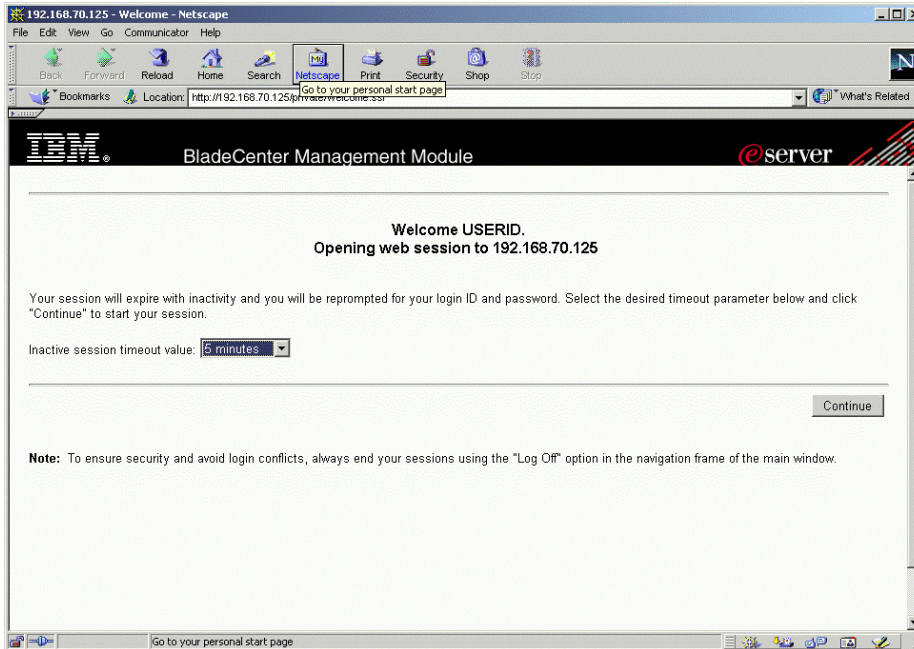


Figure 4. BladeCenter Management Module welcome screen

4. The System Status Summary screen is displayed, as shown in Figure 5. Click **Management** on the left hand panel, under **Switch Tasks**.

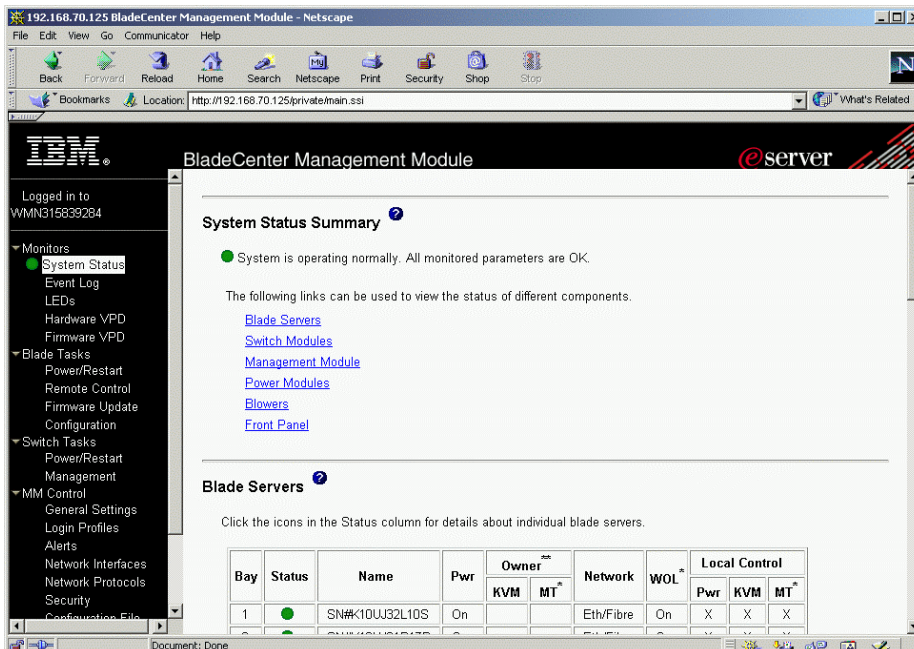


Figure 5. System Status Summary

- The Switch Management screen, shown in Figure 6, is displayed. To enable the external ports on the switch module associated with eth0, click **Switch Module 2** to access the section for switch module 2.

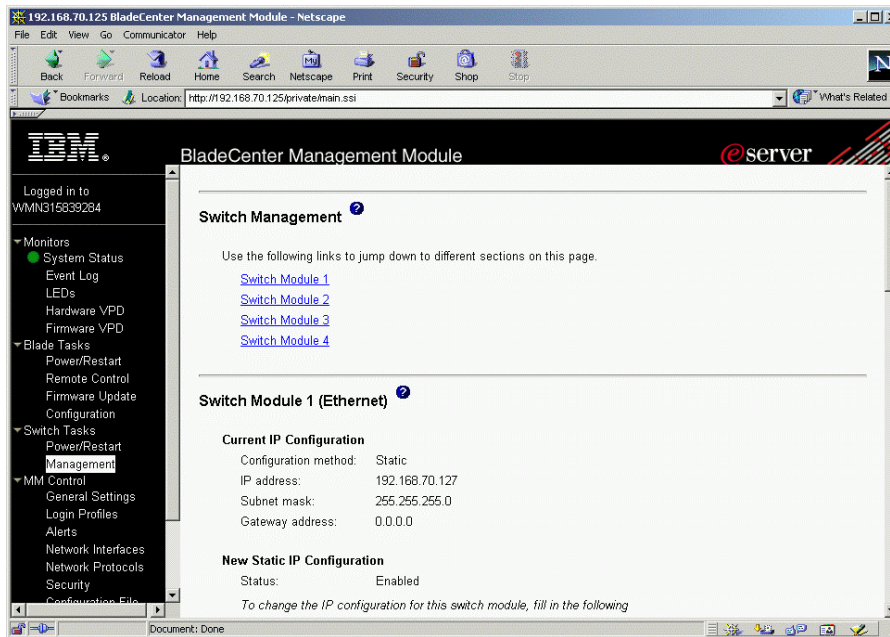


Figure 6. Switch Management

- On the Switch Module 2 section (Figure 7), click **Advanced Switch Management** or scroll down to the Advanced Switch Management section.

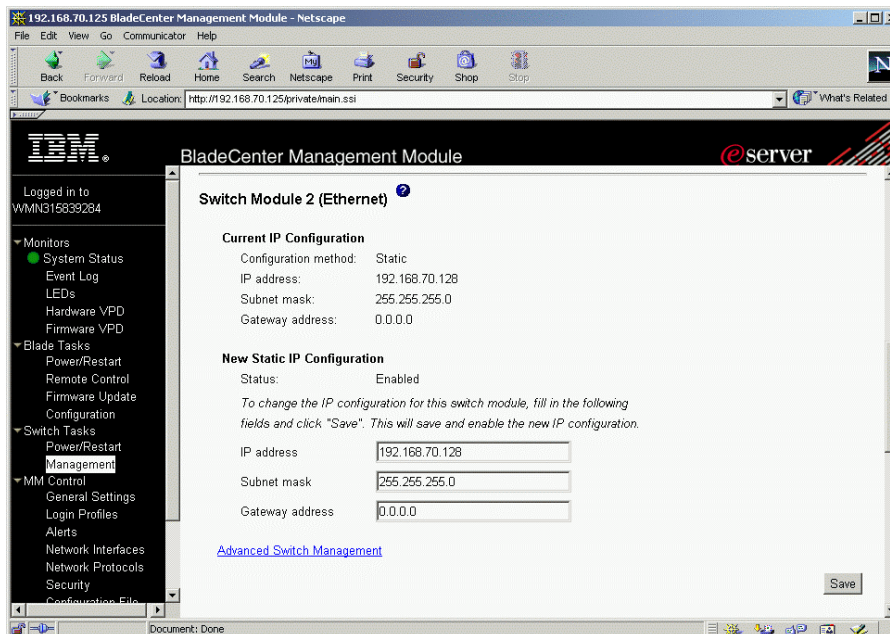


Figure 7. Switch Module 2 (Ethernet)

- On the Advanced Setup section (Figure 8), click the pull-down list for External ports and select **Enabled**, as shown. Click **Save** to activate the change.

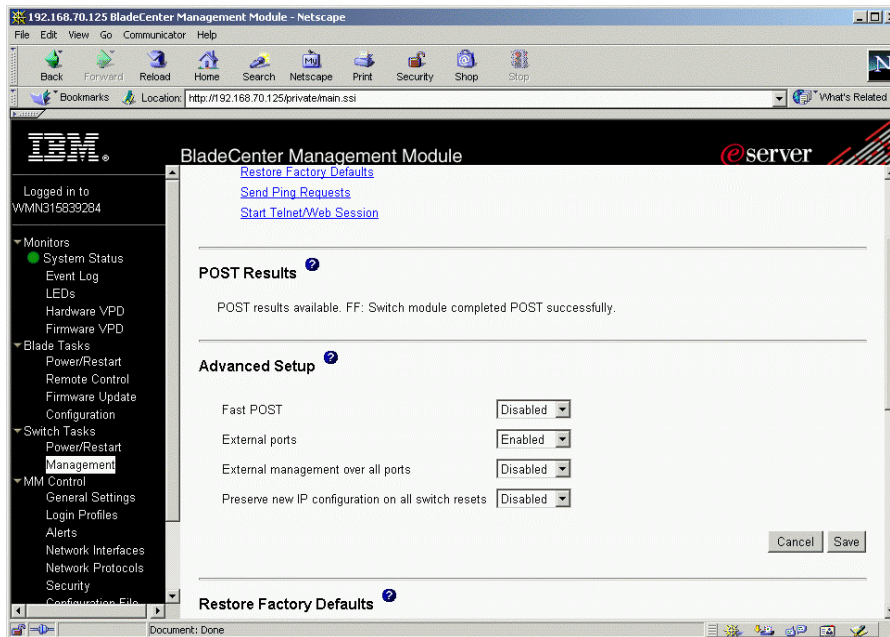


Figure 8. Advanced Setup

- Scroll up to the top of the screen and repeat the procedure to enable external ports on Switch Module 3 (Fibre Channel Switch Module).

FIBRE CHANNEL CABLING

Figure 9 shows the cabling for our RAC implementation.

The IBM HS20 Fibre Channel Expansion Card is a 2Gb Fibre Channel device that has two configurable adapter ports. Communication signals are routed from the blade server to switch module bays 3 and 4 in the BladeCenter chassis unit. Fibre Channel connections are supported by Fibre Channel Switch Modules in bays 3 and 4. Since our BladeCenter had one Fibre Channel Switch Module, we only had a connection from one of the Fibre Channel adapter ports.

On the host side, we connected the BladeCenter Fibre Channel Switch Module in bay 3 directly to the FAST900 host side mini-hubs. We connected two FC cables from the Switch Module to the FAST900, one to Controller A mini-hub and the other to Controller B mini-hub. This provides visibility from the one Fibre Channel adapter to both FAST900 controllers, allowing the blade server to access logical drives in the event of controller failover.

Note: If you are connecting to an existing SAN fabric, refer to IBM eServer BladeCenter - Fibre Channel Switch Interoperability Guide version 2.00. This document may be downloaded from IBM at <http://www.ibm.com/support>.




Figure 9. Cabling (with FC Switch Module)

Our BladeCenter was originally configured with one Fibre Channel Switch Module. Shortly before the end of this project, we received an Optical Passthru Module (OPM), which we were able to test in our RAC installation. The OPM provides an unswitched/unblocked optical connection between each blade server and a SAN fabric. The OPM is compatible with the TotalStorage family of products including FAST family, Enterprise Storage Server and IBM SAN Switches.

The OPM provides four ports external bi-directional interface ports. The optical pass-thru module cable that connects to each port on the Optical Pass-thru Module provides four channel pairs of optical connectors for connecting to the SAN infrastructure. Figure 10 shows the OPM and OPM cable with SC couplers. We used an OPM cable with LC couplers for our 2Gb SAN.

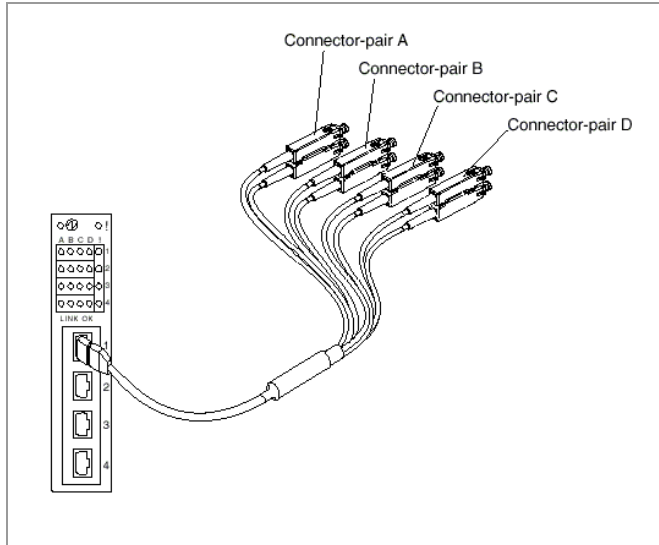


Figure 10. OPM Ports and OPM cable (with SC couplers)

Each blade server bay (BladeCenter slot) is associated with an optical pass-thru module cable connector-pair. These associations are shown in Table 2. For example, blade server in BladeCenter slot 1 (e.g. blade1) is associated with OPM port 1, connector-pair A.

Blade server	OPM Port	Cable pair
1	1	A
2	1	B
3	1	C
4	1	D
5	2	A
6	2	B
7	2	C
8	2	D
9	3	A
10	3	B
11	3	C
12	3	D
13	4	A
14	4	B

Table 2. Blade server and associated OPM cable pairs

The OPM is compatible with a variety of I/O expansion cards, including Gigabit Ethernet, Fibre Channel, and Myrinet. Due to this compatibility, installing two pass-thru modules would enable use of several different types of compatible I/O expansion cards within the same BladeCenter unit.

We installed the OPM in the BladeCenter bay 4 and connected the OPM cables to an IBM TotalStorage SAN Fibre Channel Switch Model 2109 F16, as shown in Figure 11. We connected two LC-LC cables from the SAN switch to the FAST900 host side mini-hubs, one to each controller. As with the earlier configuration using the Fibre Channel Switch Module, this provides visibility from the Fibre Channel Adapter port to both FAST900 controllers in the event of controller failover.

Note that mixing of I/O modules in paired bays (e.g., 1 and 2, 3 and 4) is not currently supported by IBM. One example of mixing I/O modules is installing the FC Switch Module in bay 3 and the OPM in bay 4. If your environment requires such a configuration, please contact your IBM Representative.

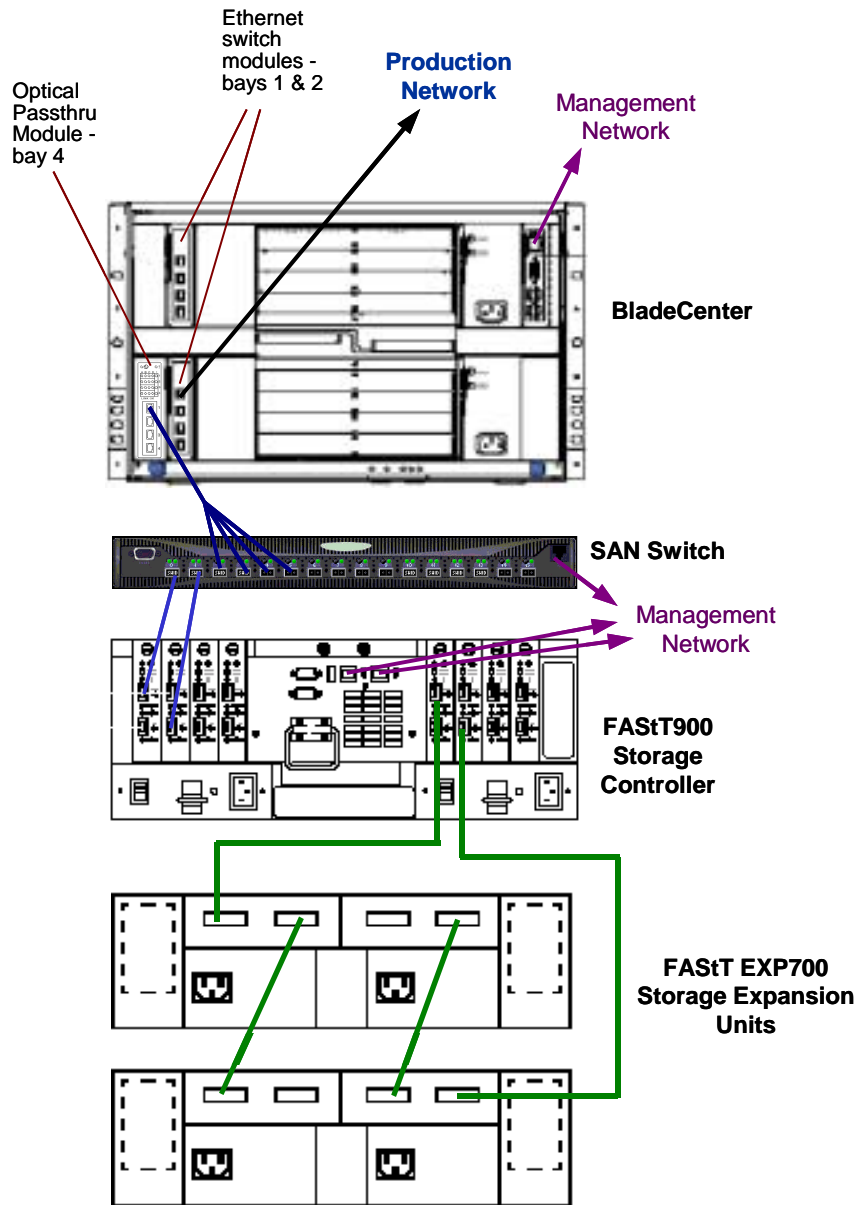


Figure 11. Cabling (with OPM)

SAN Availability Considerations

In our configuration, one Fibre Channel Switch Module was adequate for our purposes. In a production environment, the configuration should include two I/O modules. For direct connections to a supported SAN, two Fibre Channel Switch Modules are recommended. If integration into an existing SAN fabric is required, the configuration may include two Optical Passthru Modules or the two FC Switch Modules may be configured for interoperability mode.

Having two I/O modules for the Fibre Channel allows you to use both ports on the Fibre Channel Expansion Card and provides redundancy at the I/O module level. Redundancy at the I/O module is important because all the blade servers use the same I/O modules for external connections. With FAST MSJ and the

BladeCenter HS20 Fibre Channel Expansion Card driver version 6.05.00, you can configure multi-path failover and load-balancing over two adapter ports.

Figure 12 shows the cabling for a RAC implementation with two Fibre Channel Switch Modules and Figure 13 shows the cabling for a RAC implementation using two Optical Passthru Modules.

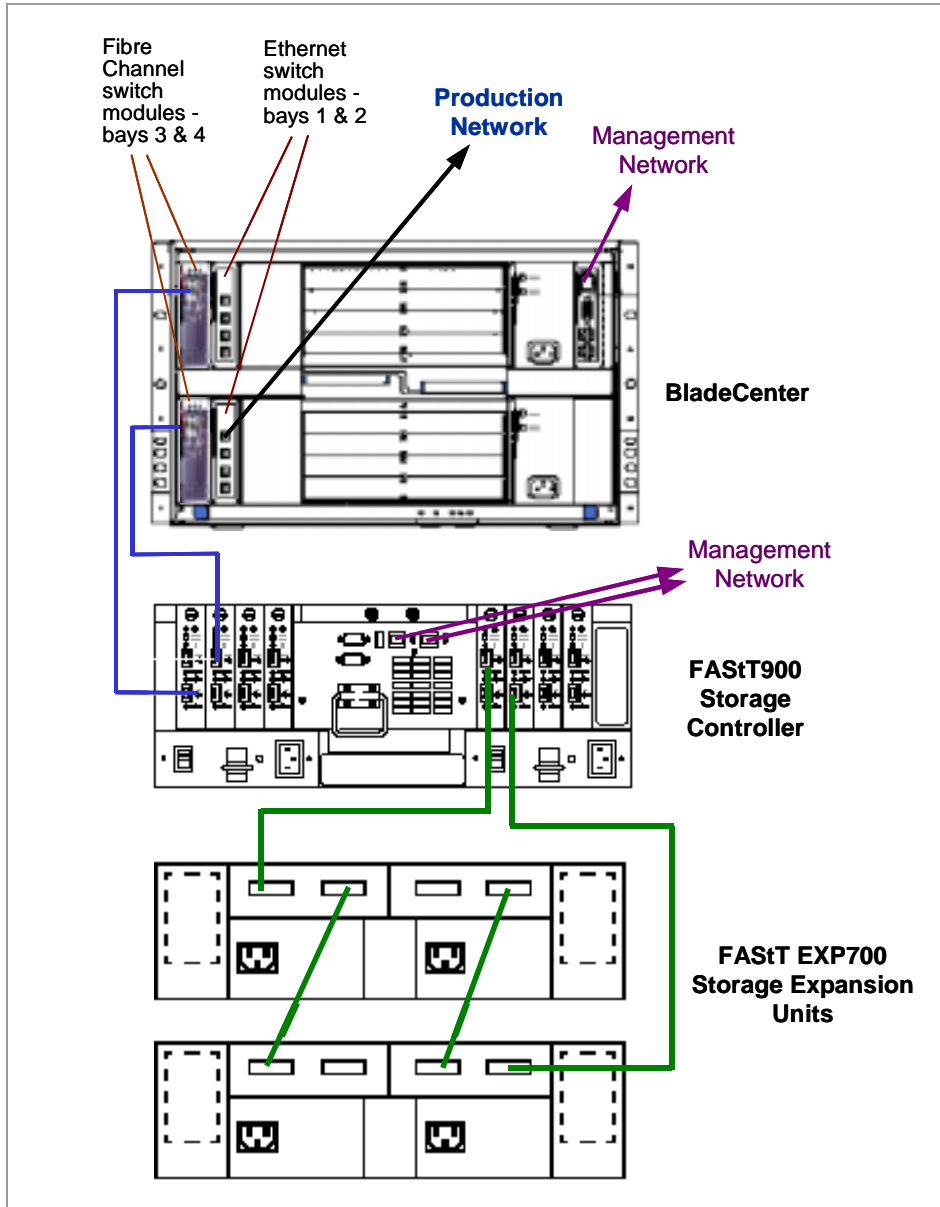


Figure 12. Cabling with 2 FC Switch Modules

Note from Figure 13 that two SAN switches have been included to eliminate the switch as a single point of failure.

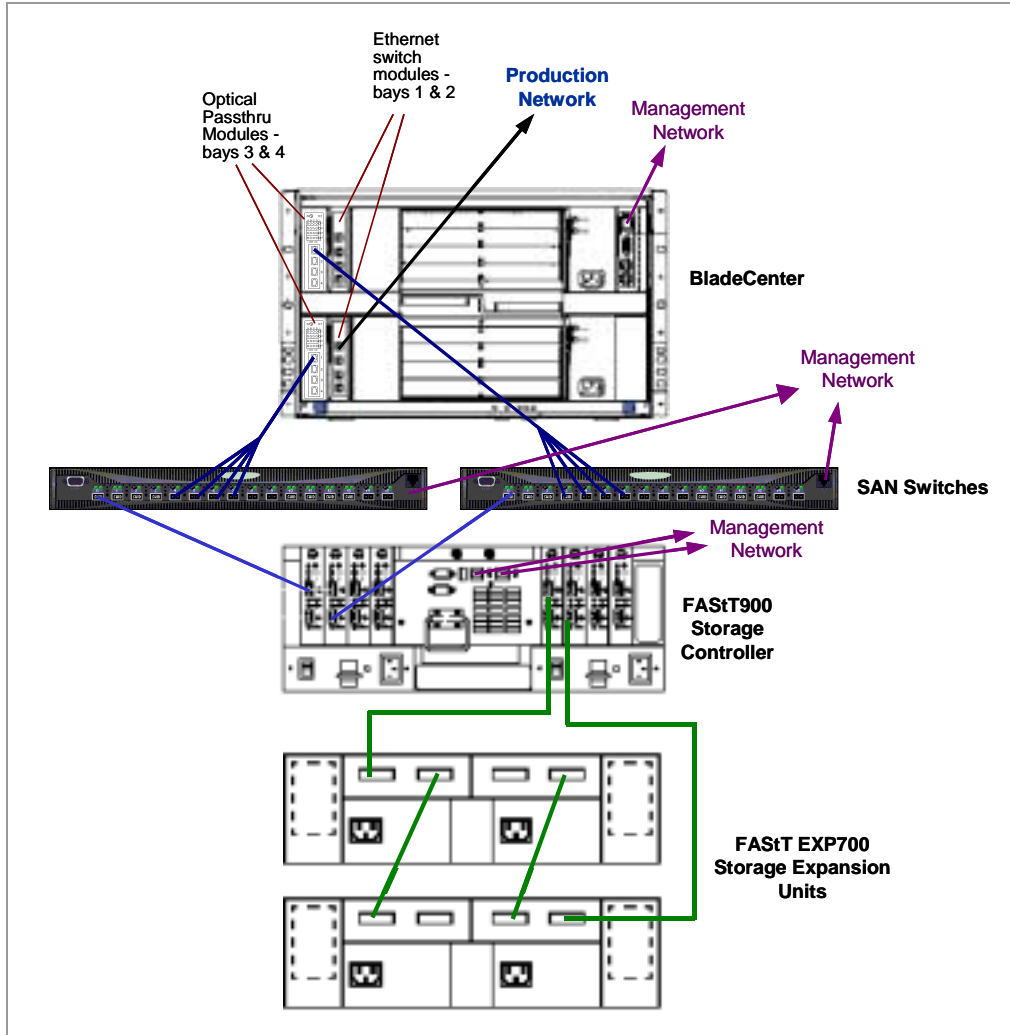


Figure 13. Cabling with 2 OPMs

FAST900 LINK RATE SWITCHES

On the FAST900, set the link rate interface switches to 2Gb/s. These switches, shown in Figure 14, are located on the host side and disk side mini-hubs.




Figure 14. Link rate interface switch on the FAST900

FASTT EXP700 ENCLOSURE ID SWITCHES

The FASTT EXP700 has two enclosure ID switches that are used to identify the FASTT EXP700 on a Fibre Channel loop, and to assign physical addresses to the drives. The enclosure ID switches are located on the back of the expansion unit, as shown in Figure 15. Our configuration consisted of two EXP700s, and we set the Enclosure Id switches to 00 and 01 respectively.




Figure 15. FASTT EXP700 Enclosure ID switches

BLADECENTER HS20

When you have determined which blade server to install, perform the following:

1. Connect a keyboard, video monitor and mouse to the ports on the management module, shown in Figure 3, upper right hand corner.
2. On the blade server being installed, use the blade server's control panel buttons (see Figure 2) to obtain ownership of the keyboard/video/mouse and CD-ROM/diskette drive.
3. To power on the blade server, use the power button behind the flip-down door on the blade server control panel. This can also be done using the BladeCenter Management Module.
4. This can also be done using the BladeCenter Management Module. Do this on the original browser window showing the BladeCenter Management Module screen (e.g. Figure 8).

- a. Click **Remote Control** on the left hand panel, under Blade Tasks. Then click Redirect Server Console, shown in Figure 16. This will launch a new browser window for the remote console.

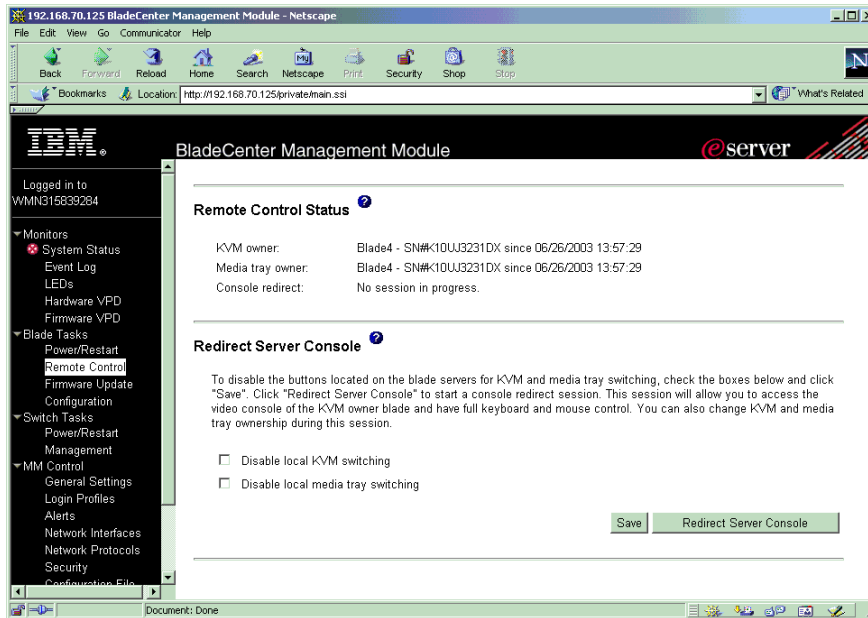


Figure 16. Remote Control screen

- b. On the field labeled **Change KVM owner**, click the pull-down list and select the blade server you want to access. The local console screen will be displayed on your browser. To switch the CD-ROM and diskette drive, select the blade server in the field labeled **Change media tray owner**. Refer to the top of Figure 17.

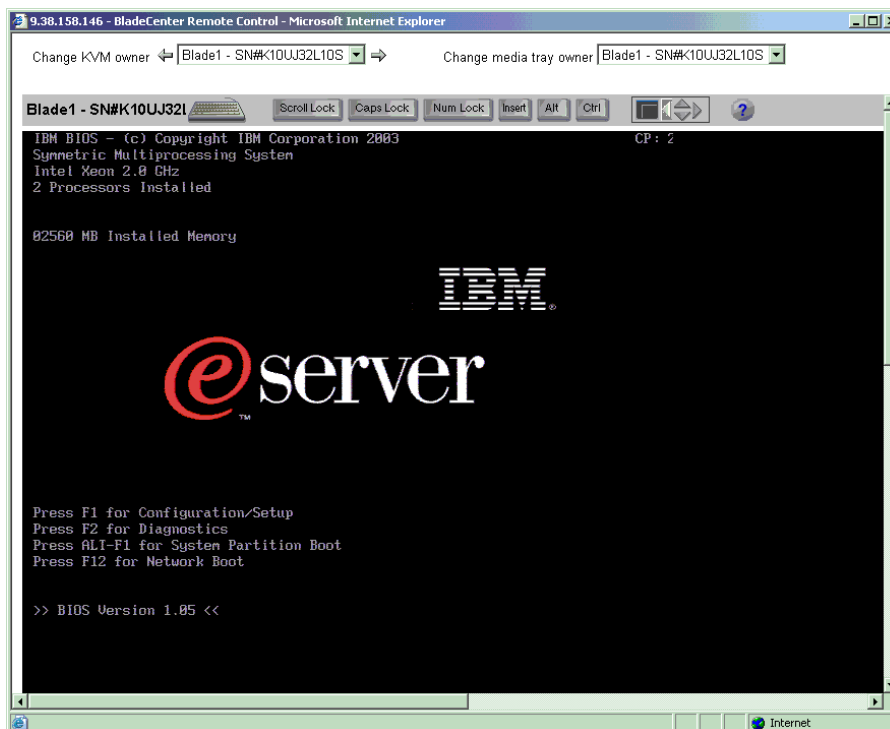


Figure 17. Remote console

If you do not see the console screen, make sure that your browser is enabled for Java Plugin. For example, in Netscape 4.8 Communicator, this is found by clicking **Edit → Preferences → Advanced**. Check **Enable Java**, then **Enable Java Plugin**. In Microsoft Internet Explorer 5.5, click **Tools → Internet Options → Advanced** and uncheck **Use Java 2 1.4.1_01** under **Java (Sun) for <applet>**.

- c. To power up the blade server, go to the original browser window showing the BladeCenter Management Module screen (e.g. Figure 8). On the left hand panel, under Blade Tasks, click **Power/Restart**. Bootup messages will be displayed on the remote console window (Figure 17).

SETTING UP OPERATING ENVIRONMENT

INSTALLING RED HAT ENTERPRISE LINUX AS 2.1

This section describes our installation of Linux on the blade server.

Install Red Hat Enterprise Linux AS 2.1 as follows:

1. Refer to the document entitled IBM eServer BladeCenter – Installing Red Hat Enterprise Linux AS Version 2.1 (References, 1) for instructions on setting up the BladeCenter hardware, updating the firmware, and configuring the disk drives using the IBM ServeRAID and LSI SCSI controllers as applicable.

If you are using the original CDs for Red Hat Advanced Server 2.1, you will need to obtain boot disk and driver disk images as described in the current version of this document, dated April 9, 2003. We did not need to do this since we were using Red Hat Enterprise Linux AS (ELAS) 2.1 Update 2. The remaining steps describe our installation procedure using Red Hat ELAS Update 2.

2. Power on the BladeCenter HS20 and insert the first Red Hat ELAS CD into the CDROM drive.
3. The Red Hat text-mode welcome screen is displayed. Press Enter to install Red Hat Linux in graphical mode.
4. We will not describe all the steps to install Red Hat ELAS here, but rather refer you to The Red Hat Enterprise Linux AS Installation Guide (References, 27). However, a number of steps were important to our installation, and these were as follows:

- a. In the Mouse Configuration window (Figure 18), verify selection of a USB mouse, either Generic – 3 Button Mouse (USB) or Generic – 2 Button Mouse (USB).

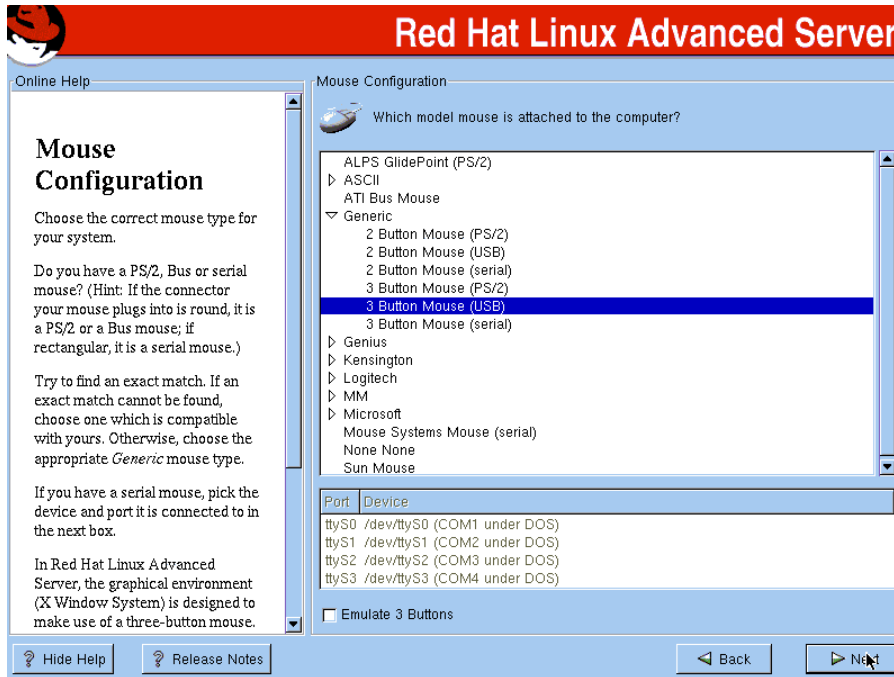


Figure 18. Mouse Configuration

- b. In the Disk Partitioning Setup screen, we chose to have the installer automatically partition the disk, as shown in Figure 19. Click **Next** to continue.

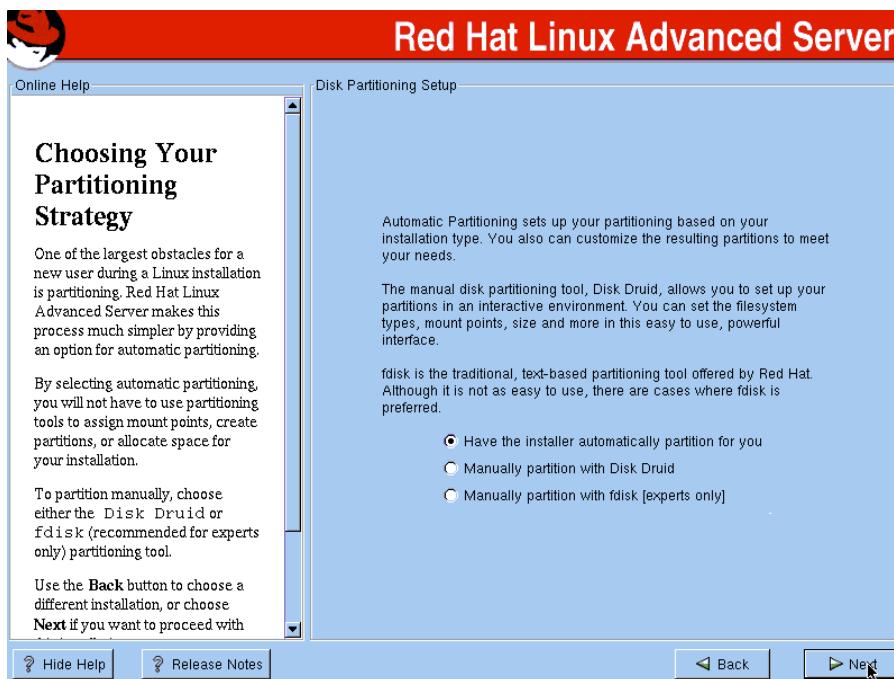


Figure 19. Disk partitioning Setup

- c. On the Automatic Partitioning screen (Figure 20), check **Review** and click **Next**.

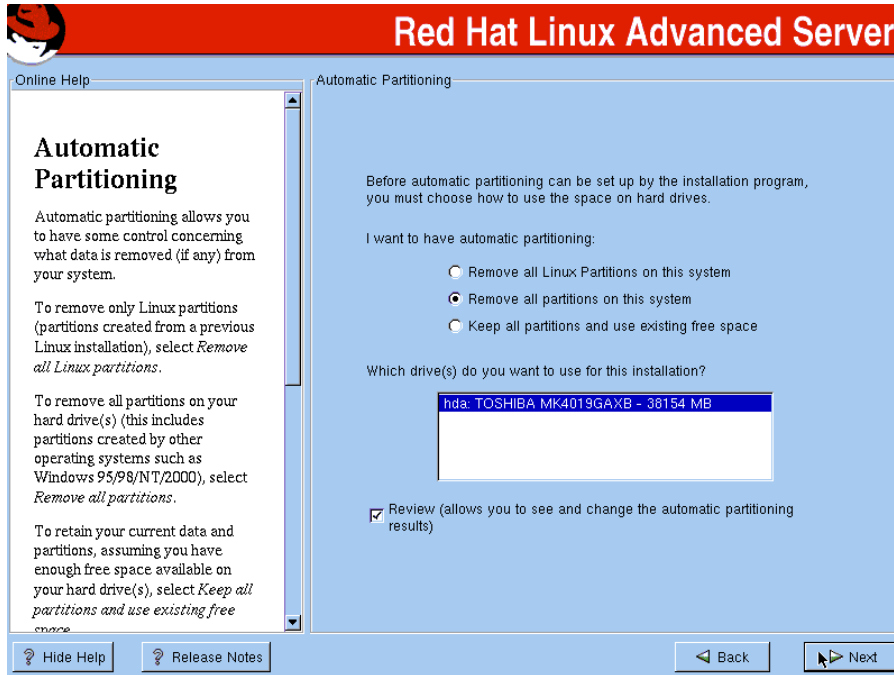


Figure 20. Automatic partitioning

- d. On the Disk Setup screen, shown in Figure 21, notice the sizes of / (root filesystem), /user, and /home. The Oracle installer requires up to 400MB of temporary space, so we increased the root filesystem to provide a large enough /tmp. We also reduced both /usr and /home to create a new partition for our ORACLE_HOME.




Figure 21. Disk Setup

Figure 22 shows the Disk Setup with our changes.

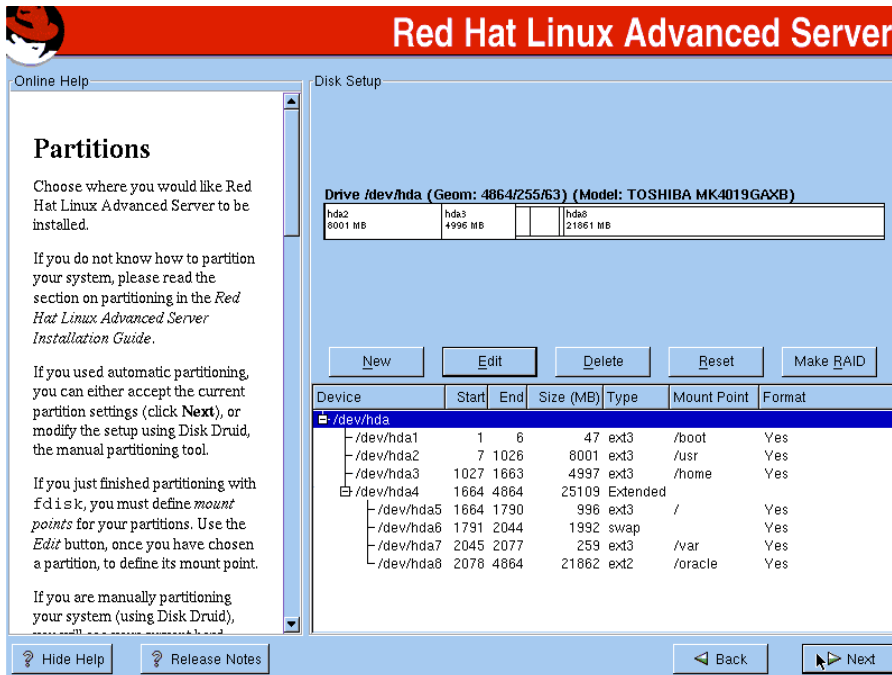


Figure 22. Disk setup with modified partitions

- e. On the Firewall Configuration screen, take note of the default selections. Leaving the default settings will interfere with setup activities later on that require network communications. Either configure the firewall at this time, or defer it for later by choosing No firewall. To simplify our installation, we chose No firewall (Figure 23).




Figure 23. Firewall configuration

- f. When the Package Group Selection screen is displayed (Figure 24), select Software Development, which you will need for the Oracle9i installation. Check Select individual packages to review and modify selections as needed.

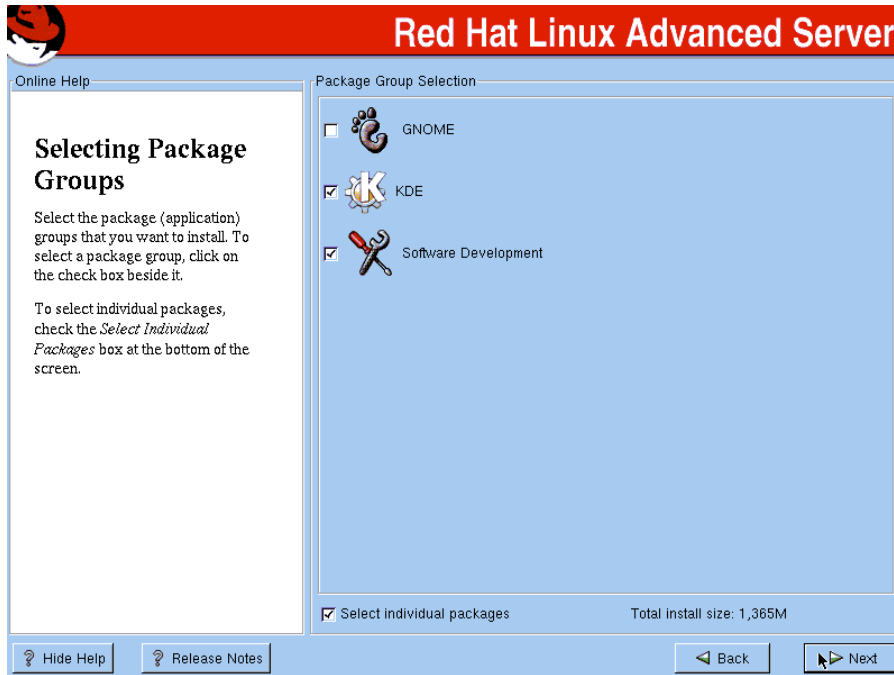


Figure 24. Package Group Selection

- g. The Individual Package Selection screen is displayed. Review the selections and select additional packages as required. We added the ftp server, which is not one of the default selections. To do this, click **System Environment** → **Daemons** on the left hand panel. Scroll down on the right hand panel and select the wu-ftp daemon, wu-ftpd.

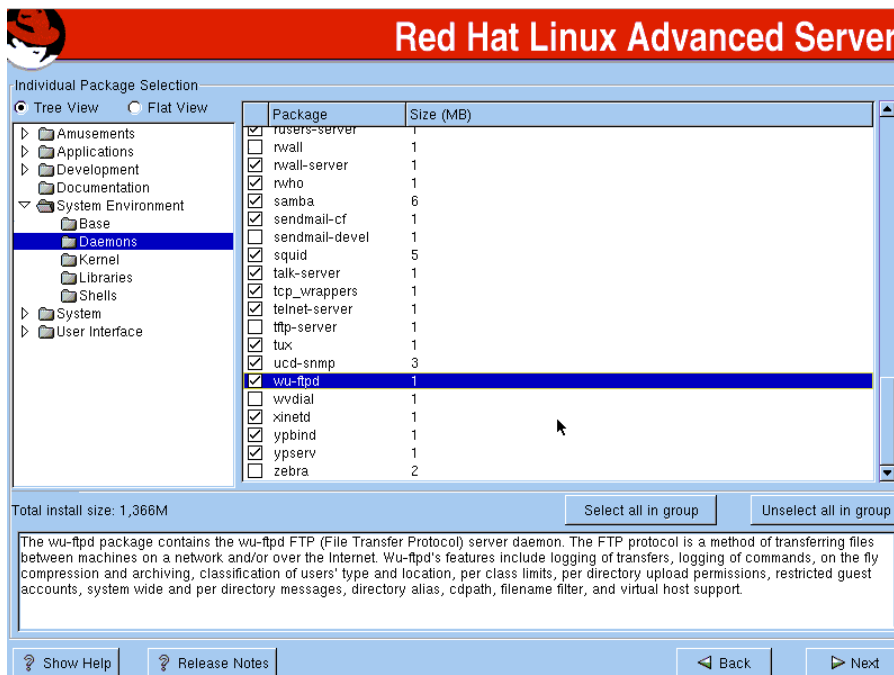


Figure 25. Select Daemons

- h. The kernel source package is required to compile the driver for the Fibre Channel Expansion Card. Click **Development** → **System** on the left hand panel. Select kernel-source on the right hand panel as shown

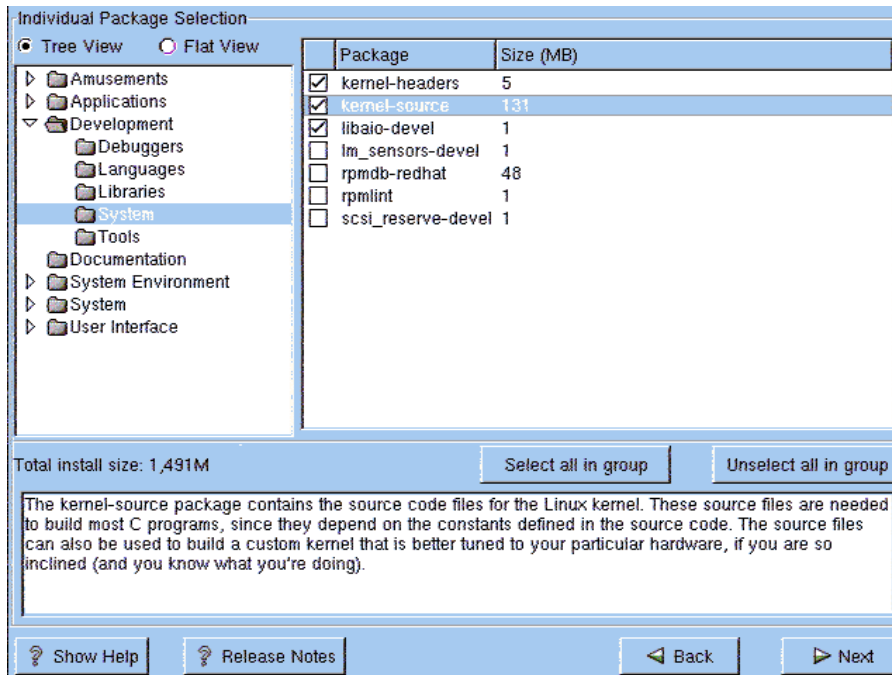


Figure 26. Select kernel source package

- i. When the Red Hat installation is complete, a Congratulations window is displayed. Click **Exit** and remove the CD and diskette from their respective drives. Red Hat restarts the server.

Note: During our installation, the system did not shut down completely. The last message on the console was "Unmounting file systems:". We used Ctrl-Alt-Del to reboot the server manually.

5. Post-installation

a. CD-ROM

After installing Red Hat Enterprise Linux AS 2.1, the CD-ROM fails to mount if the drive has been previously switched to another blade server. To resolve this, do the following as root user:

- i) Create the mount point for the CD-ROM under /mnt. Type the command:

```
mkdir /mnt/cdrom
```

- ii) Edit /etc/fstab. Change the device in the entry for the CD-ROM from /dev/cdrom to /dev/scd0.

b. Diskette drive

Perform the following steps as root to enable mounting of the diskette drive:

- i) Create the mount point for the diskette drive under /mnt. Type the command:

```
mkdir /mnt/floppy
```

- ii) Edit /etc/fstab and add the following statement:

```
/dev/sdx /mnt/floppy auto noauto,user 0 0
```

where x is the letter that will be the last SCSI drive on the system. If you create 5 logical drives on the FAStT, then the floppy will be /dev/sdf.

- c. Check <http://rhn.redhat.com> for new and updated packages that available on Red Hat Network (RHN).

We installed the kernel errata e.25 on our system. We also installed the matching kernel source and kernel header packages.

- d. Create the operating system group for the Oracle Universal Installer Inventory and the user ID that will be used to install, configure and run Oracle. In our test, we created the oinstall and dba groups, and the oracle user ID. We also made oracle user a member of the disk group, to enable access to raw devices.

You will need to be root user to do this. From a shell prompt, type the following commands:

```
groupadd oinstall
groupadd dba
useradd -g oinstall -G dba,disk -m oracle
```

Note that we have made the `oracle` user a member of the `disk` group. This allows the `oracle` user to access the raw devices.

Change oracle's password to enable login. Type the following command:

```
passwd oracle
< type the desired oracle password when prompted>
```

NETWORK

As mentioned earlier, we used `eth0` for the production (public) network and `eth1` for the private cluster interconnect. The interfaces, IP addresses and hostnames are shown in Table 3.

Node	Public LAN			Private LAN		
	Hostname	Interface	IP address	Hostname	Interface	IP address
1	blade1	eth0	192.168.10.10	iblade1	eth1	1.1.1.1
2	blade2	eth0	192.168.10.20	iblade2	eth1	1.1.1.2

Table 3. Network addresses

We set up networking as follows:

1. As root user, edit `/etc/sysconfig/network-scripts/ifcfg-eth0` as shown below for blade1:

```
DEVICE=eth0
BOOTPROTO=static
IPADDR=192.168.10.10
NETMASK=255.255.255.0
ONBOOT=yes
```

Example 1. `/etc/sysconfig/network-scripts/ifcfg-eth0`

2. Edit `/etc/sysconfig/network-scripts/if-cfg-eth1`, shown below for blade1.

```
DEVICE=eth0
BOOTPROTO=static
IPADDR=1.1.1.1
NETMASK=255.255.255.0
ONBOOT=yes
```

Example 2. `/etc/sysconfig/network-scripts/ifcfg-eth1`

3. Restart the network with the commands:

```
service network stop
```

```
service network start
```

4. Repeat the steps 1 - 3 for blade2.
5. Add the host names to /etc/hosts on both nodes as shown in Example 3:

```
192.168.10.10 blade1.sanmateo.ibm.com blade1
192.168.10.20 blade2.sanmateo.ibm.com blade2
1.1.1.1 iblade1.localdomain iblade1
1.1.1.2 iblade2.localdomain iblade2
```

Example 3. /etc/hosts

6. Make sure that the network interfaces have started. Use the ping command to check that both interfaces are working on both nodes.
7. Set up host equivalence. Do this on both nodes by editing /etc/hosts.equiv and adding the production and private hostnames as shown in Example 4.

```
blade1
blade2
iblade1
iblade2
```

Example 4. /etc/hosts.equiv

8. Make sure that the rlogin and rsh services are started automatically. Use the chkconfig command to check as follows:

```
chkconfig --list rlogin
chkconfig --list rsh
```

If off, turn the service on using the chkconfig command, for example, for rlogin,

```
chkconfig rlogin on
```

9. Log in as oracle user and test host equivalence from both nodes using the rsh, rcp and rlogin commands. For example, from blade1,

```
[oracle@blade1 oracle]$ rsh blade1 date
Tue Dec 2 17:00:11 PST 2003
[oracle@blade1 oracle]$ rcp .bash_profile blade2:/tmp/test1
[oracle@blade1 oracle]$ rlogin blade2
Last login: Tue Dec 2 17:00:21 from blade1
[oracle@blade2 oracle]$ ls /tmp/test1
/tmp/test1
[oracle@blade2 oracle]$
```

FIBRE CHANNEL ADAPTER DRIVER

The BladeCenter HS20 Fibre Channel Expansion Card (Fibre Channel adapter) provides two Fibre Channel ports. The first port operates at I/O base address 0x2400 and is supported by an I/O module in bay 3. The second operates at I/O base address 0x2600 and is supported by an I/O module in bay 4.

We downloaded IBM eServer BladeCenter HS20 (Type 8832, 8678) - Fibre Channel Expansion Card drivers version 6.05.00 for Linux from the IBM support page. This driver has failover support enabled by default.

To find the download URL, access <http://www.ibm.com/support> and search for "bladecenter fibre channel". The driver package consists of two files, a compressed tar file and a README text file.

Review the README file prior to starting the installation. Follow the procedure described in the README text file to install and configure the Fibre Channel adapter driver.

We followed these steps to build and install the device driver:

1. Move the downloaded files to blade1.
2. As root user, rebuild kernel dependencies –

```
cd /usr/src/linux-2.4/configs
cp kernel-2.4.9-i686-smp.config ../.config
make dep
cd <directory containing compressed tar file>
```

3. Unpack the compressed tar file.

```
tar xzvf 25k9265.tgz
cd i2x00-v6.05.00-fo
./libinstall
./drvsetup
```

This creates another directory named `i2x00-v6.05.00-fo` containing the driver source files.

4. Compile the driver as follows:

```
cd i2x00-v6.05.00-fo
make all SMP=1
```

5. Copy the compiled driver to the kernel driver directory:

```
cd /lib/modules/2.4.9-e.25smp/kernel/drivers/addon/qla2xxx
mv qla2300.o qla2300.o.save
cd -
cp qla2300.o /lib/modules/2.4.9-e.25smp/kernel/drivers/addon/qla2xxx
```

6. Edit `/etc/modules.conf`. Add the lines shown in bold type below (Example 5).

```
Alias parport_lowlevel parport_pc
alias eth0 tg3
alias eth1 tg3
alias scsi_hostadapter qla2300
alias usb-controller usb-ohci
options scsi_mod max_scsi_luns=128
```

Example 5. `/etc/modules.conf`

7. Recreate the initial ramdisk image:

```
cd /boot
mkinitrd -f <ramdisk image file name> <kernel version>
```

In our implementation, we did not rename the default ramdisk image file. We recreated the ramdisk image with the following command:

```
mkinitrd -f initrd-2.4.9-e.25smp.img 2.4.9-e.25smp
```

If you create a new ramdisk image file, be sure to update `/etc/grub.conf` to point to your new ramdisk file. Update the line beginning with `initrd=` for all the appropriate kernel entries.

Note: If you are using LILO for your boot manager, do the same in `/etc/lilo.conf`. Be sure to run the `lilo` command before rebooting.

8. Load the device driver as follows:

```
insmod qla2300 displayConfig=1
```

9. Obtain the Fibre Channel adapter host port identifiers or port names with the following command:

```
grep adapter-port /var/log/messages
```

Record the hexadecimal values for use during FASt900 configuration.

You can also use the Fast!UTIL utility during bootup to find the Fibre Channel adapter's host port identifiers. See Appendix B for more details on Fast!UTIL.

10. Reboot the server.
11. After the server reboots, log in as oracle and switch to root user. Verify that the Fibre Channel adapter driver has been loaded. Use the `lsmod` command to display the loaded kernel modules. You should see `qla2300` in the displayed list of modules.
12. Transfer the driver module compiled in step 4 to blade2, and copy it to `/lib/modules/2.4.9-e.25smp/kernel/drivers/addon/qla2xxx` on blade2, following step 5.
13. Repeat steps 6 - 11 on blade2.

SHARED STORAGE

Configure the FASt900 storage with the IBM Storage Manager client (SMclient). We will not go into detail in this paper, but refer you to IBM Redbook, Implementing Linux with IBM Disk Storage, and IBM Redpaper, Implementing IBM [^] xSeries SANs (References, 13 and 14).

IBM Storage Manager Version 8.3 has three components:

- Storage Manager 8.3 Client (SMclient),
- Storage Manager 8.3 Runtime (SMruntime)
- Storage Manager 8.3 Utility software (SMutil)

We installed IBM FASt Storage Manager Version 8.3 on blade1 and configured storage from blade1. We recommend that you install SMutil and SMruntime (required for SMutil) on the other nodes for the `SMdevices` command, which we found to be very useful for displaying the logical drives (or LUNs).

To use Storage Manager, run as root:

```
export DISPLAY=0.0
SMclient
```

Install IBM FASt MSJ Diagnostic and Configuration Utility version 2.0 for Linux (we used release 40). This will allow you to configure failover and to monitor the Fibre Channel adapters. You will only need to install the GUI on a server, which could be one of the nodes. Install the agent on all the nodes.

The MSJ installer sets up the agent, called `qlremote`, as a service started during bootup. If you do not want to start `qlremote` automatically, remove the symbolic links with the name `S99qlremote` in `/etc/rc.d/rc3.d` and `/etc/rc.d/rc5.d`. (`chkconfig` does not work properly)

The agent can be started manually. Start it as root with the command:

```
qlremote
```

To use the GUI program, run as root:

```
cd /usr
./FASTT_MSJ
```

Once you have configured the FASTT, reload the Fibre Channel Expansion Card driver to force detection of the logical drives. Do this as root user on all server nodes, using the following commands:

```
rmmmod qla2300
modprobe qla2300
```

The logical drives (LUNs) will appear as SCSI disk drives, e.g. /dev/sda, /dev/sdb, etc. Use the `SMdevices` command to verify that the logical drives are visible to the operating system. If you chose not to install Storage Manager SMutil on the blade server, you can also use the `fdisk -l` command to do this.

After creating the logical drives, remember to edit the diskette drive entry (/mnt/floppy) in /etc/fstab, as noted in step 1.a.ii) on page 26.

Data Files for Oracle9i RAC

Oracle9i RAC stores the database files in raw devices that are shared by all cluster nodes. You can also use Oracle Clustered File System (OCFS), which is available for download from Oracle Technology Network (OTN). We used raw devices for our test installation.

Note that the number of partitions that may be created on a SCSI hard disk (or logical drive) is limited to 15 in Linux/i386. This limit includes the extended partition, so the actual number of usable partitions (thus, raw devices) for each SCSI disk is 14. This limit on the number of partitions is documented in /usr/src/linux-2.4/Documentation/devices.txt.

Installation of the Oracle9i RAC product requires two raw devices. These devices are for the quorum file (used by Oracle Cluster Manager), and the Shared Configuration File, (used for RAC configuration). When the test database is created later in this paper, additional raw partitions will be needed for the database control files, logs and tablespaces.

1. Create two partitions on the FASTT900/EXP700 SAN using `fdisk` as root user. After creating the partitions, use the `raw` command to bind the disk partitions to the raw devices. To make the bindings persistent, /etc/sysconfig/rawdevices needs to be set up so that they will be recreated during bootup. As root, we edited /etc/sysconfig/rawdevices and added the lines for raw1 and raw2 below (Example 6):

```
## raw device bindings
## format: <rawdev> <major> <minor>
##          <rawdev> <blockdev>
## example: /dev/raw/raw1 /dev/sda1
##          /dev/raw/raw2 8 5
/dev/raw/raw1 /dev/sda1
/dev/raw/raw2 /dev/sda2
```

Example 6. /etc/sysconfig/rawdevices

2. Bind the disk partitions to the raw devices - issue the `service` command shown below as root user:

```
[root@bladel1 root]# service rawdevices start
Assigning devices:
    /dev/raw/raw1 --> /dev/sda1
/dev/raw/raw1: bound to major 8, minor 1
    /dev/raw/raw2 --> /dev/sda2
/dev/raw/raw2: bound to major 8, minor 2
done
[root@bladel1 root]#
```

3. Check the bindings with the `raw -q` command as shown:

```
[root@bladel1 root]# raw -qa
/dev/raw/raw1: bound to major 8, minor 1
/dev/raw/raw2: bound to major 8, minor 2
```

4. We created symbolic links to the raw devices for ease of management. This is not necessary, since you can also specify raw device names during installation. As `oracle` user,

```
mkdir $ORACLE_HOME/oradata
ln -s /dev/raw/raw1 $ORACLE_HOME/oradata/quorum
ln -s /dev/raw/raw2 $ORACLE_HOME/oradata/srvm
```

Table 4 summarizes the raw devices needed for RAC. Additional raw devices will be created later for our test database. See Appendix A for consolidated table of raw devices in our example.

Database file	Size (MB)	Disk Partition	Raw partition	Symbolic link
quorum	110	/dev/sda1	/dev/raw/raw1	\$ORACLE_HOME/oradata/quorum
srvm	110	/dev/sda2	/dev/raw/raw2	\$ORACLE_HOME/oradata/srvm

Table 4. Raw devices for RAC

INSTALLING ORACLE9i RAC

PRE-INSTALLATION TASKS

To prepare the system for Oracle9i installation, perform steps 1 through 6 on all nodes:

1. Login as `root` user.
2. Edit `root` user's `.bash_profile`. Add the environmental variable `ORACLE_HOME`:

```
export ORACLE_HOME=/oracle
```

3. Modify the kernel parameters that affect resources used by Oracle9i: This task must be performed as `root` user.
 - a. Review kernel parameter recommendations in Reference 18. Use the `sysctl` command to do look at the current settings. For example, the default values for `kernel.sem` are shown below:

```
bladel1:~ # sysctl kernel.sem
kernel.sem = 250          32000    32         128    ← semmsl semmns semopm semmni
bladel1:~ #
```

- b. In our test, we modified the kernel parameters shown in Example 7, which may be different for your environment. Add the statements shown to `/etc/sysctl.conf`, which will cause the parameters to be set upon reboot.

```
kernel.sem = 250 32000 100 128
kernel.shmmax = 1318072320
```

Example 7. Kernel parameters

4. Set the parameters now using the `sysctl -p` command, as shown below:

```
[root@bladel1 root]# sysctl -p
net.ipv4.ip_forward = 0
net.ipv4.conf.default.rp_filter = 1
kernel.sysrq = 0
kernel.sem = 250 3200 100 128
fs.file-max = 81920
kernel.shmmax = 1287180000
```

```
[root@blade1 root]#
```

5. Create ORACLE_HOME directory and make `oracle` user the owner. Issue the commands shown below as `root` user:

```
blade1:~ # mkdir /oracle  
blade1:~ # chown -R oracle:oinstall /oracle
```

6. Set up environmental variables for the `oracle` user. Do this as `oracle` user.
 - a. Modify `.bash_profile` in `oracle`'s home directory, adding the statements shown in Example 8.

```
export ORACLE_BASE=/oracle  
export ORACLE_HOME=$ORACLE_BASE  
export ORA_NLS33=$ORACLE_HOME/ocommon/nls/admin/data  
  
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$ORACLE_HOME/lib  
  
PATH=.:$ORACLE_HOME/bin:$PATH:$HOME/bin  
export PATH
```

Example 8. `oracle` user environmental variables

- b. Verify that the environmental variables are set correctly by logging in as `oracle` user and echoing an environmental variable to the console. For example:

```
oracle@blade1:~> echo $ORACLE_HOME  
/oracle  
oracle@blade1:~>
```

7. Oracle 9i Release 2 for Linux is supplied on three CDs, so it is necessary to switch product CDs during installation. Oracle Universal Installer (OUI) will manage the switching between CDs and ask you for the CD location when necessary.

To install Oracle9i RAC, you need to run the OUI from one node. For the most part, OUI handles installation to the other cluster nodes. There are a number of steps that need to be done on the other nodes and these are called out by OUI at various points during the process.

For our installation, we ran OUI from `blade1`. We copied the Oracle9i product CDs to the hard disk on `blade1`, into three subdirectories named `/u/stage9i/Disk1`, `/u/stage9i/Disk2` and `/u/stage9i/Disk3`, respectively.

If you chose to install from the CDs, do not run OUI with the current directory under the CD's mount point, or you will not be able to switch CDs when prompted.

INSTALL ORACLE CLUSTER MANAGER

1. On `blade1`, log in as `oracle` and invoke the Oracle9i Universal Installer (OUI) as follows:

```
cd /u/stage9i/Disk1  
./runInstaller
```

- The Welcome screen in Figure 27 is displayed. Click **Next** to continue.

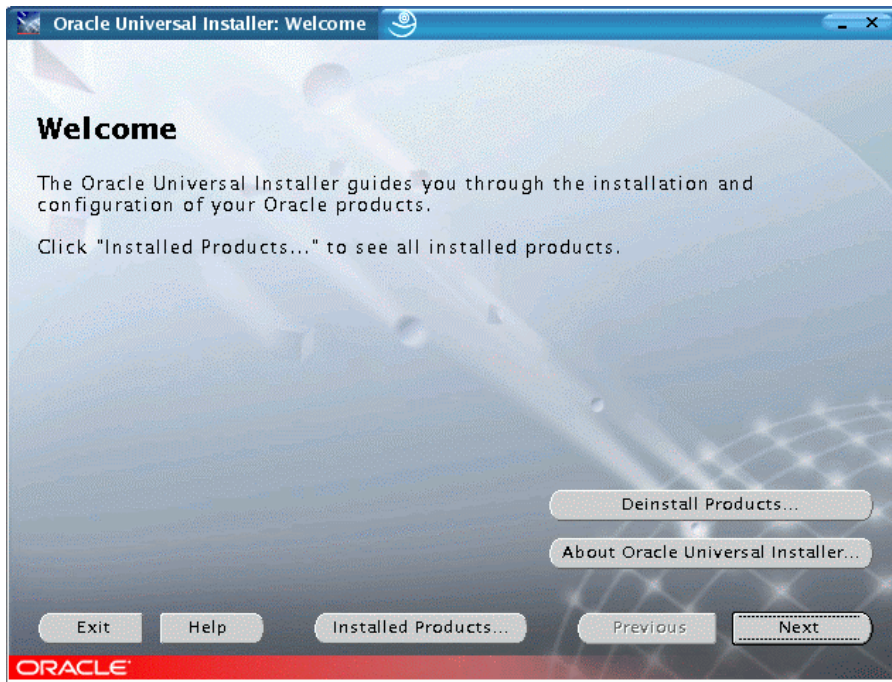


Figure 27. Welcome Screen

- Since this is the first installation, the Inventory Location screen is displayed, as shown in Figure 28. The default directory for installation files is a subdirectory under ORACLE_HOME. Leave the default setting and click **OK** to continue.



Figure 28. Inventory Location

Note: There are several files that are used to keep track of choices you make when using OUI. If you have to restart the installation process, information in these files will direct the OUI to default values. The best way to restart is to use OUI to de-install, then restart. However, if you have to manually restart the installation, look at these files and remove or modify the entries.

- oraInst.loc, located in either /etc or /var/opt/oracle, contains the following:
inventory_loc=/oracle/9i/oraInventory
inst_group=oinstall
- /etc/oratab has entries for **ORACLE_HOME** and **ORACLE_SID**

4. The screen entitled UNIX Group Name, Figure 29, is displayed. Enter oinstall as shown. Click **Next** to continue.

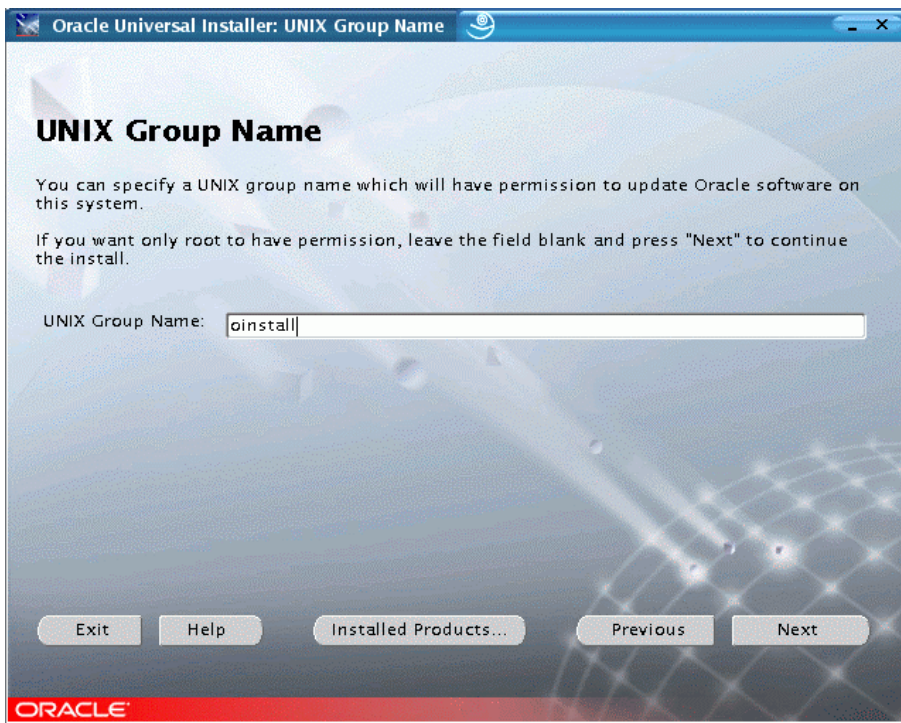


Figure 29. UNIX Group Name

5. A popup window appears, asking to execute a shell script as `root`. Go to another window, switch to `root` and execute the specified shell script. When completed, return to the popup window and click **Continue**.

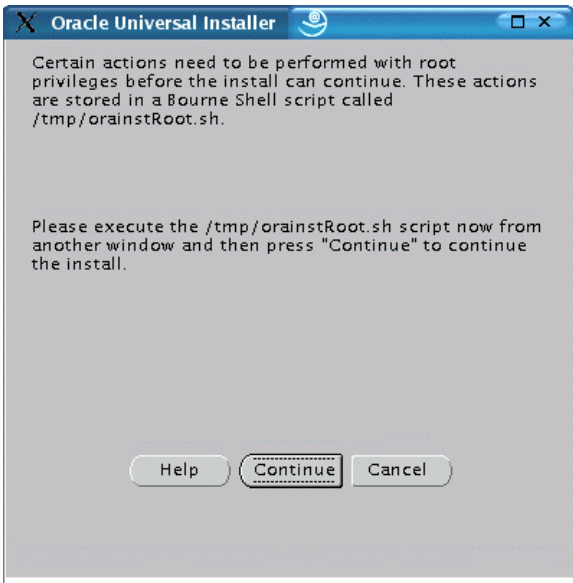


Figure 30. Popup window – root script execution

6. Verify information on the File Locations screen, shown in Figure 31. Click **Next** to continue.

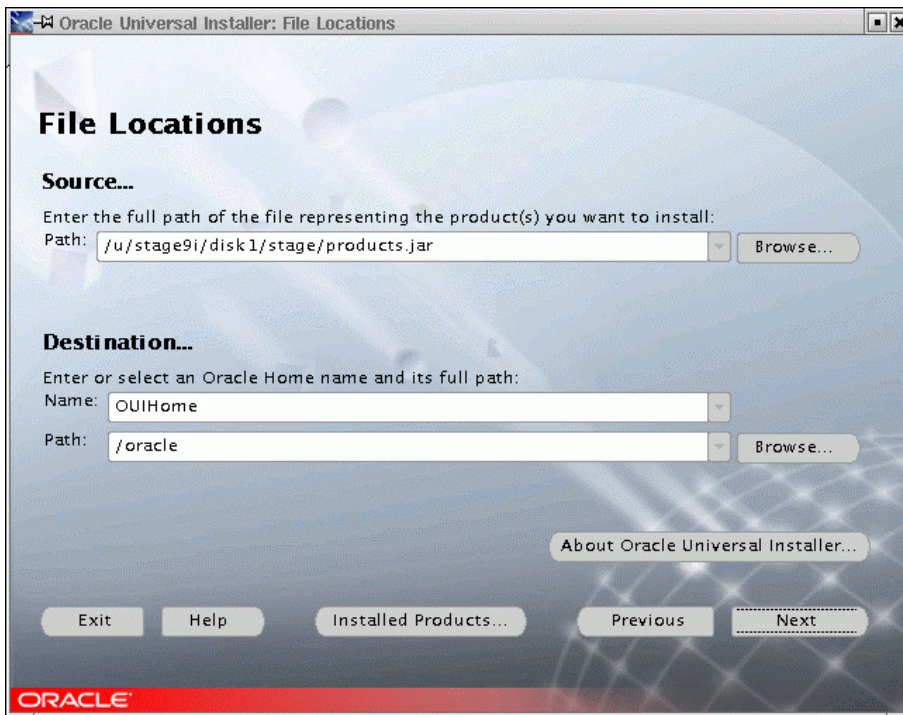


Figure 31. File Locations

7. Select Oracle Cluster Manager 9.2.0.1.0 on the Available Products screen shown in Figure 32. Click **Next** to continue.

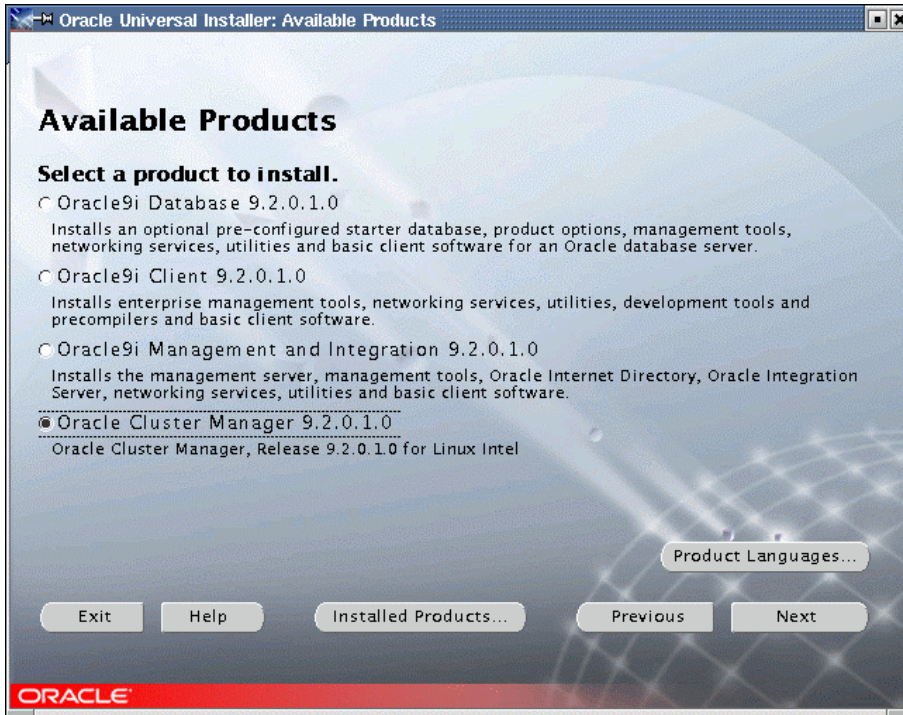


Figure 32. OUI Available Products

8. The Public Node Information window lets you specify the node names. In our case, the two HS20 nodes are named blade1 and blade2 (Figure 33). Enter the information as shown and click **Next**.

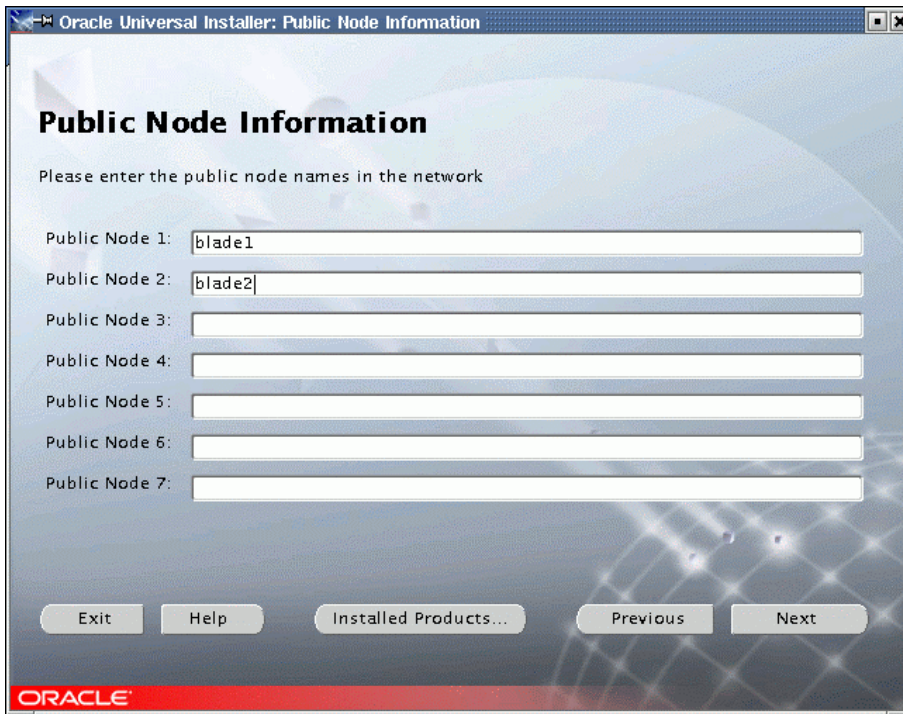


Figure 33. Public Node Information

- The Private Node Information window configures the interfaces to be used for the cluster interconnect. In our example, the hostnames for the cluster interconnect are `iblade1` and `iblade2`. Enter the information as shown and click **Next** to continue.

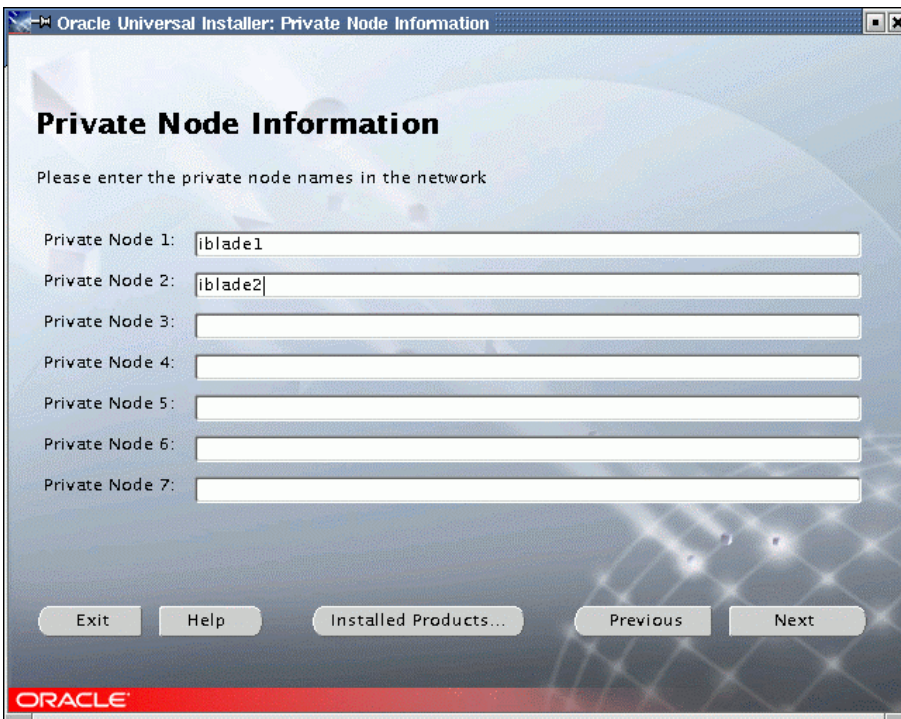


Figure 34. Private Node Information

- The Watchdog parameter window (Figure 35) asks you for the watchdog parameters. Simply accept the default value (60000) and click **Next** to continue.

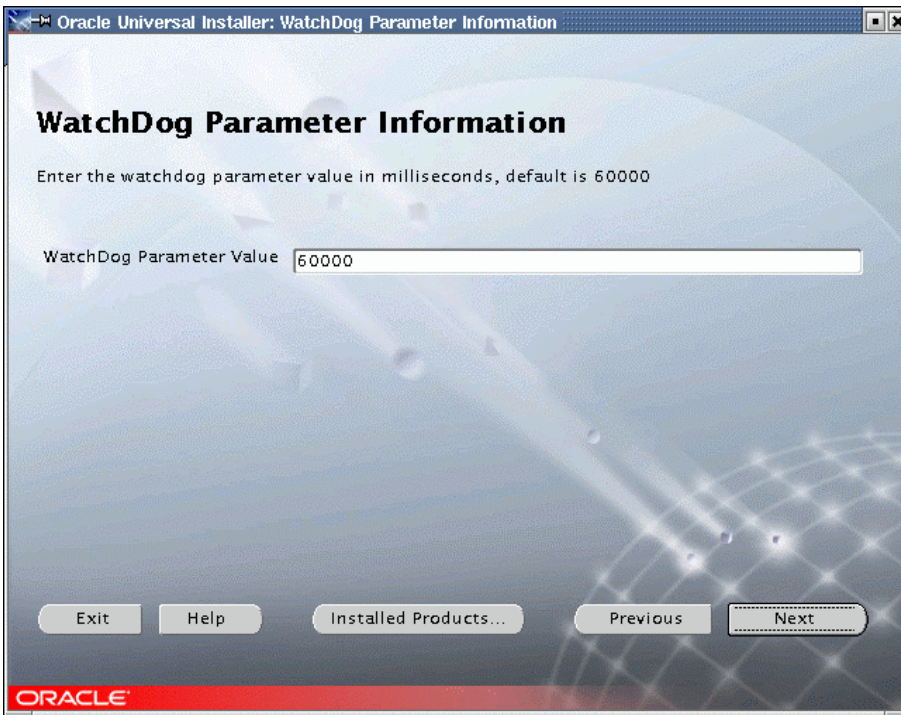


Figure 35. WatchDog Parameter

- The Quorum Disk Information window displays next. In this case, our quorum device will be `/oracle/oradata/quorum`. Enter this as shown and click **Next** to continue.

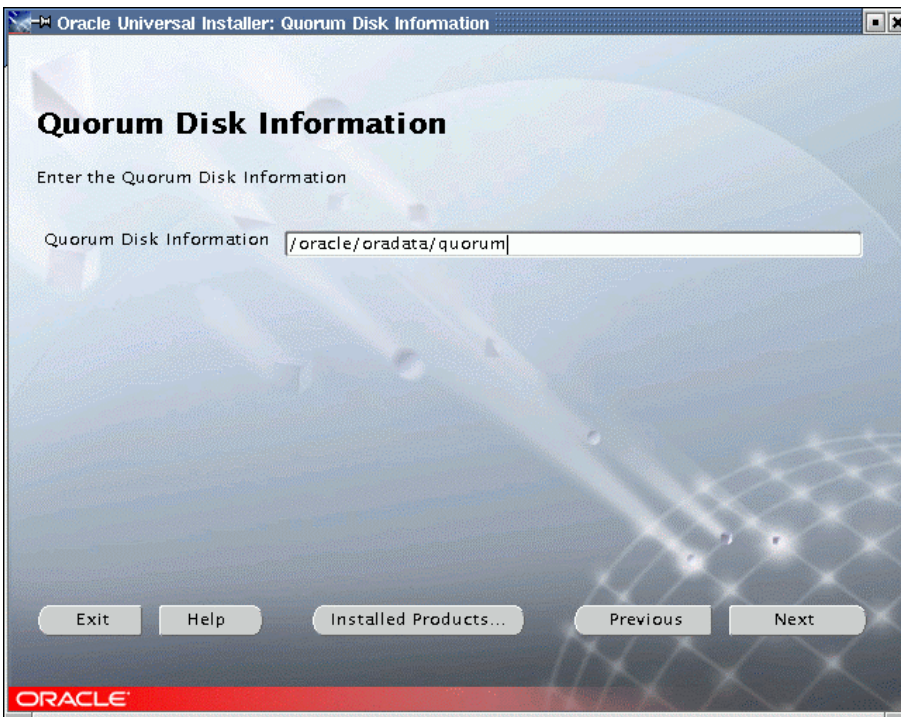


Figure 36. Quorum Disk

- The summary screen, shown in Figure 37, is displayed. Click **Install** to start the installation.

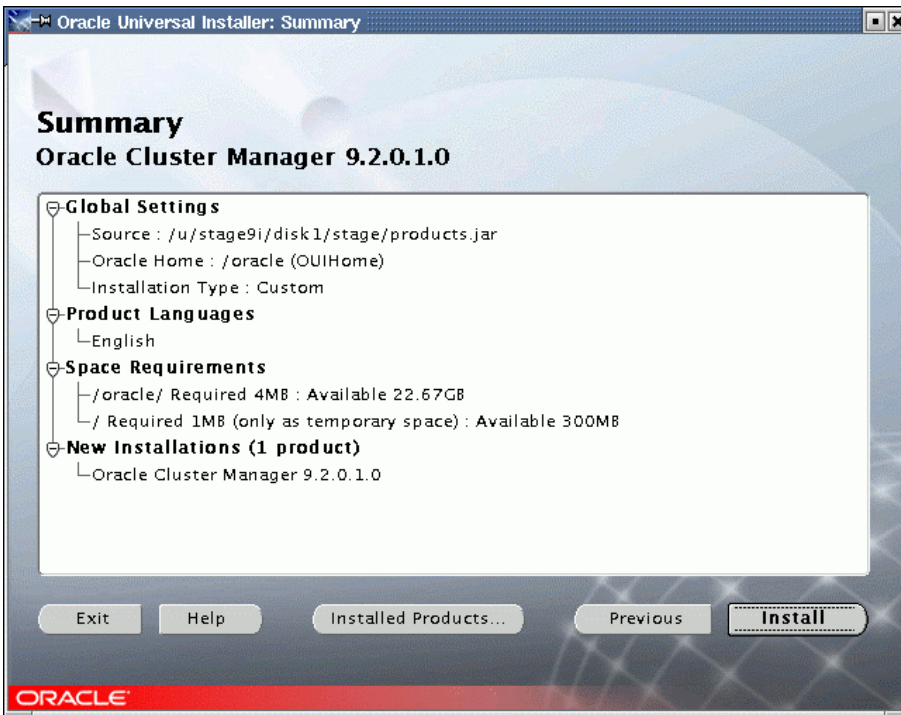


Figure 37. Summary

13. A progress screen will be briefly shown, then the End of Installation window shown in Figure 38. Click **Exit**, then **Yes** on the confirmation window.

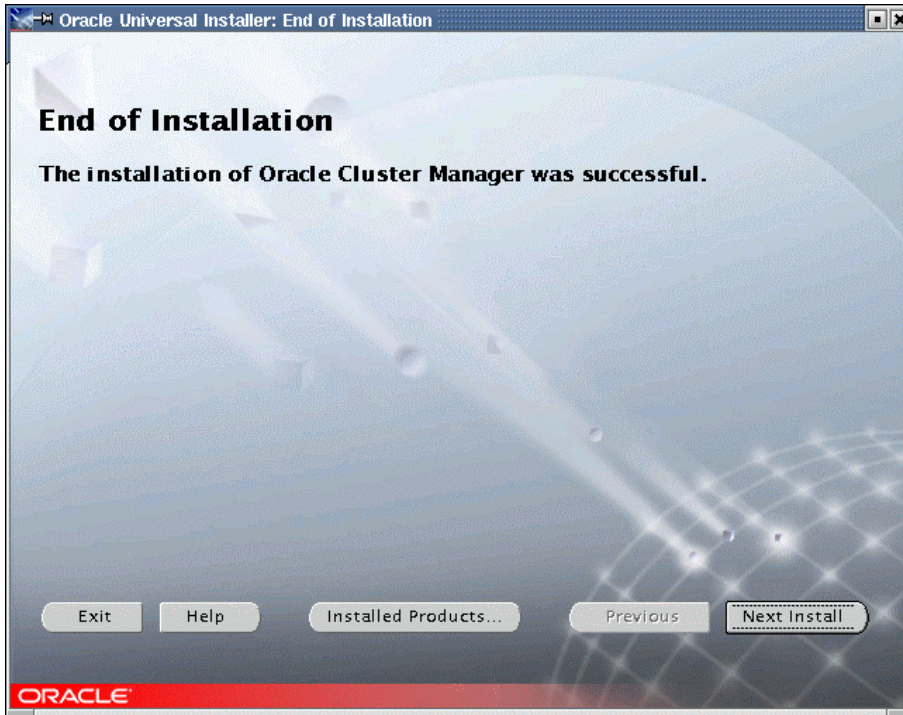


Figure 38. End of Installation

14. As `oracle` user, check to see if the `$ORACLE_HOME/oracm/log` directory exists on all nodes. If not, create `$ORACLE_HOME/oracm/log`.
15. Switch to root user and remove references to the watchdog daemon from the Cluster Manager configuration files on all nodes. As root user,
- Edit `$ORACLE_HOME/oracm/admin/cmcfgr.ora`. Comment the following lines:

```
#WatchdogSafetyMargin=5000
#WatchdogTimerMargin=60000
```
 - Edit `$ORACLE_HOME/oracm/admin/ocmargs.ora`. Comment the following line:

```
#watchdogd
```
 - Edit `$ORACLE_HOME/bin/ocmstart.sh`. Comment all the lines related to `watchdogd`, as shown in Example 9 below:


```

.
.
.
# watchdogd's default log file
#WATCHDOGD_LOG_FILE=$ORACLE_HOME/oracm/log/wdd.log

# watchdogd's default backup file
#WATCHDOGD_BAK_FILE=$ORACLE_HOME/oracm/log/wdd.log.bak

# Get arguments
#watchdogd_args=`grep '^watchdogd' $OCMARGS_FILE |\
# sed -e 's+^watchdogd *++'`

.
.
.
# Check watchdogd's existence
#if watchdogd status | grep 'Watchdog daemon active' >/dev/null
#then
# echo 'ocmstart.sh: Error: watchdogd is already running'
# exit 1
#fi

.
.
.
# Backup the old watchdogd log
#if test -r $WATCHDOGD_LOG_FILE
#then
# mv $WATCHDOGD_LOG_FILE $WATCHDOGD_BAK_FILE
#fi

# Startup watchdogd
#echo watchdogd $watchdogd_args
#watchdogd $watchdogd_args

.
.

```

Example 9. \$ORACLE_HOME/oracm/bin/ocmstart.sh

16. Start Cluster Manager by issuing the following command as root user:

```
/oracle/oracm/bin/ocmstart.sh
```

Do not start the watchdog daemon.

17. Check to make sure that Cluster Manager is running. On both nodes, use the command:

```

[root@bladel root]# ps -ef |grep oracm
root      2626      1  0 20:01 pts/3    00:00:00 oracm
root      2628    2626  0 20:01 pts/3    00:00:00 oracm
root      2629    2628  0 20:01 pts/3    00:00:00 oracm
root      2630    2628  0 20:01 pts/3    00:00:00 oracm
root      2631    2628  0 20:01 pts/3    00:00:00 oracm
root      2632    2628  0 20:01 pts/3    00:00:00 oracm
root      2633    2628  0 20:01 pts/3    00:00:00 oracm
root      2634    2628  0 20:01 pts/3    00:00:00 oracm
root      2635    2628  0 20:01 pts/3    00:00:00 oracm
root      2642    2628  0 20:02 pts/3    00:00:00 oracm
root      2692    2644  0 20:02 pts/4    00:00:00 grep oracm
[root@bladel root]#

```

If you do not see the `oracm` processes running, check `$ORACLE_HOME/oracm/log/cm.log` for error messages.

INSTALL ORACLE9i/RAC

1. As `oracle` user, start the Oracle Universal Installer on blade1 again. Click **Next** on the Welcome screen and you will be prompted for the cluster nodes on which to install. The Cluster Node Selection window is shown in Figure 39. Select the nodes in the cluster and click **Next** to continue. Note that blade1, where OUI is being run, does not need to be highlighted, since it is always selected.

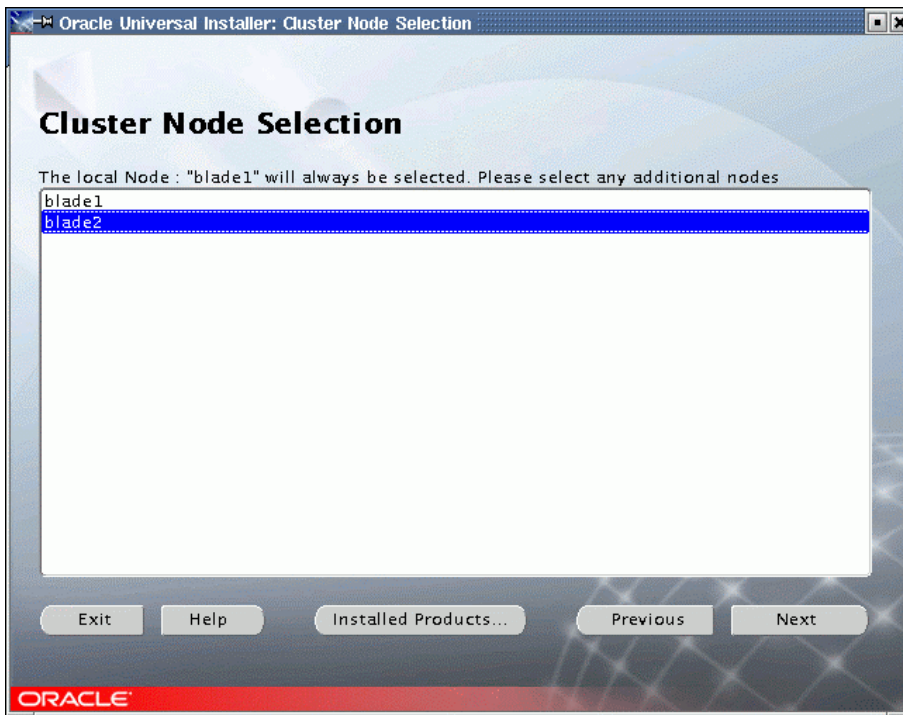


Figure 39. Cluster Node Selection

The next window will be the File Location Screen that we have already seen in Figure 31. Click **Next** to continue.

2. Select **Oracle9i Database 9.2.0.1.0** from the Available Products window (Figure 40) and click **Next** to continue.

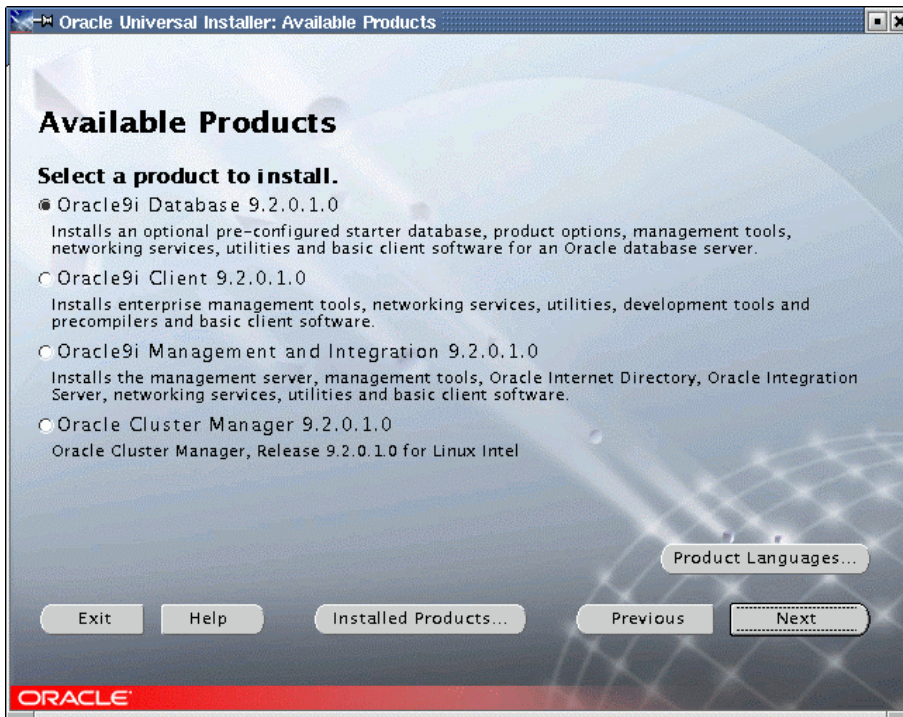


Figure 40. Available Products

3. The next window displayed is for Installation Types (Figure 41). Select **Custom** to have full control over the components being installed. Click **Next** to continue.

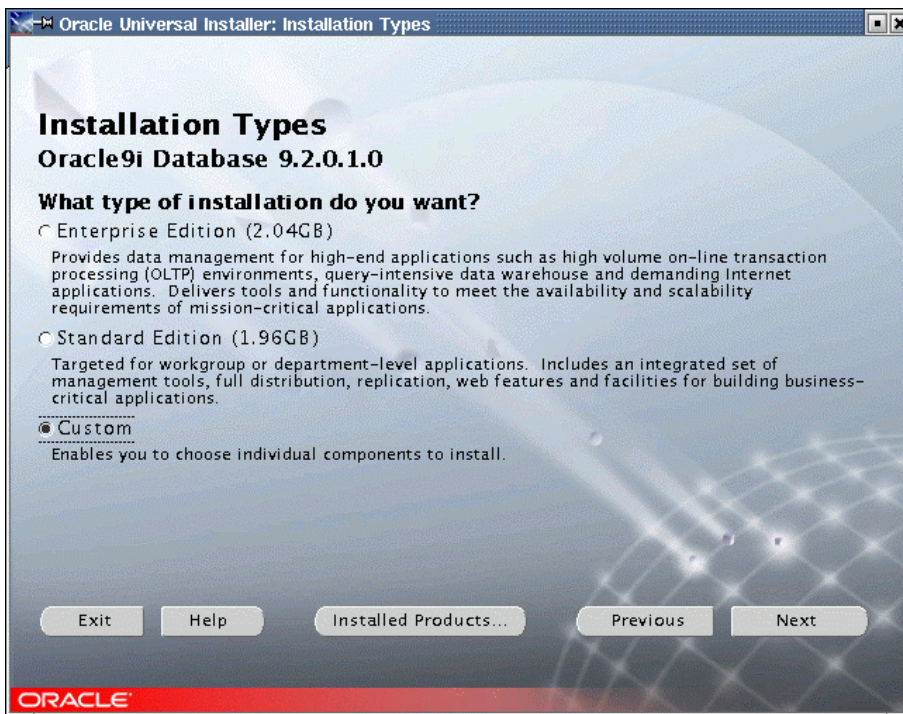


Figure 41. Installation Types

- The Available Product Components window (Figure 42) allows you to select components to be installed. For a Real Application Cluster configuration, be sure to select at least the Oracle Database, Real Application Clusters, Partitioning and Netlistener components. When all desired components have been selected, click **Next** to continue.

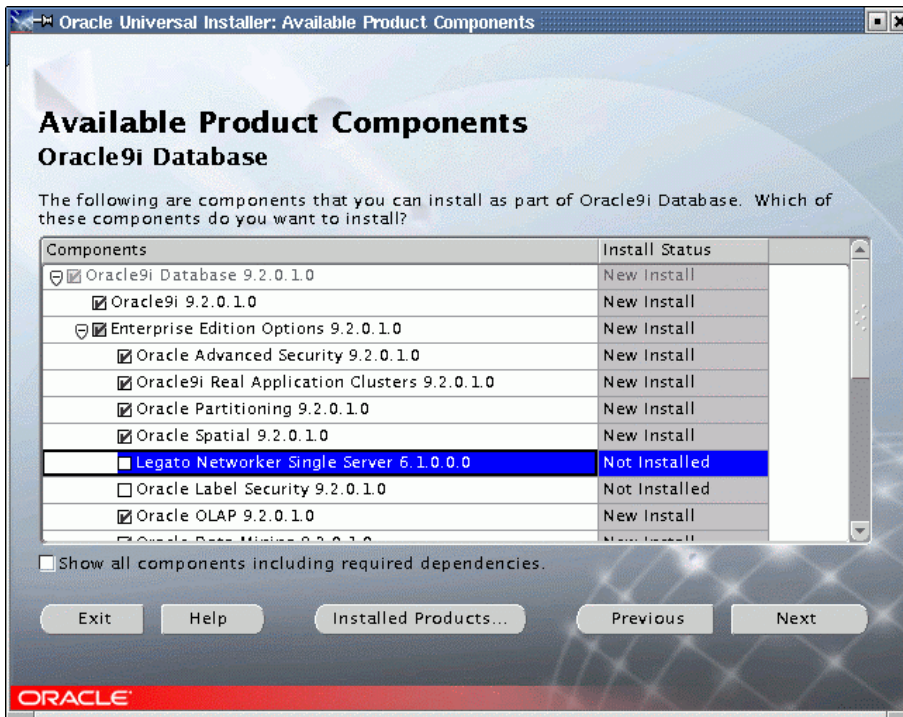


Figure 42. Available Components

- The Components Location window is displayed (Figure 43). Do not modify the values, click **Next** to continue.

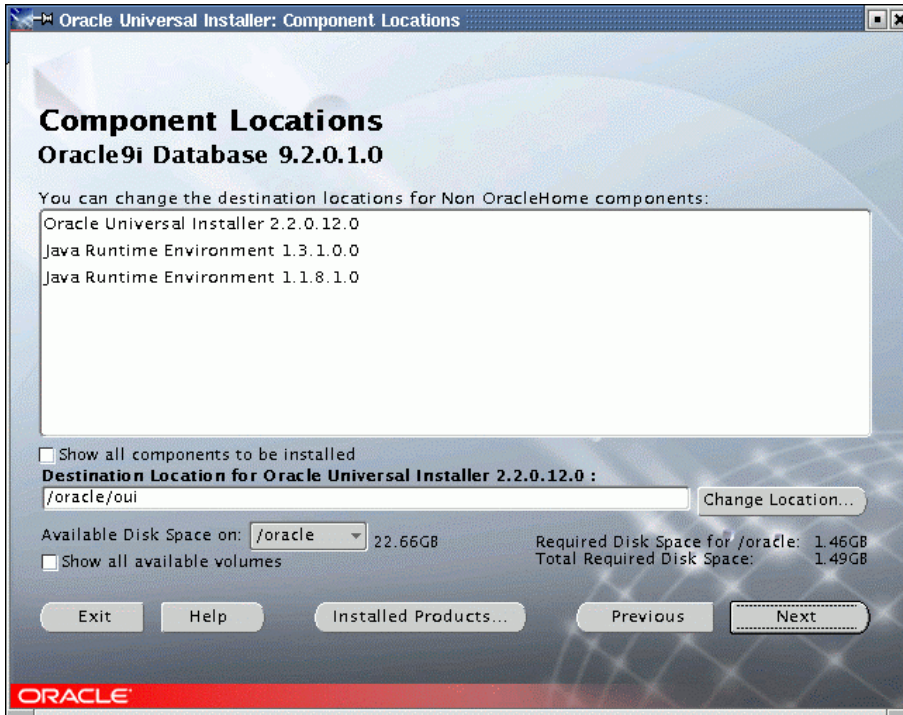


Figure 43. Component Locations

- The next screen (Figure 44) prompts for the name of the file or raw device for the shared configuration file. Enter the file name as shown and click **Next** to continue.



Figure 44. Shared Configuration File Name

- On the Privileged Operating System Groups window (Figure 40) select the operating system group(s) for your database operator and administrator. Normally, these values are set to dba. Enter the appropriate information and click **Next** to continue.

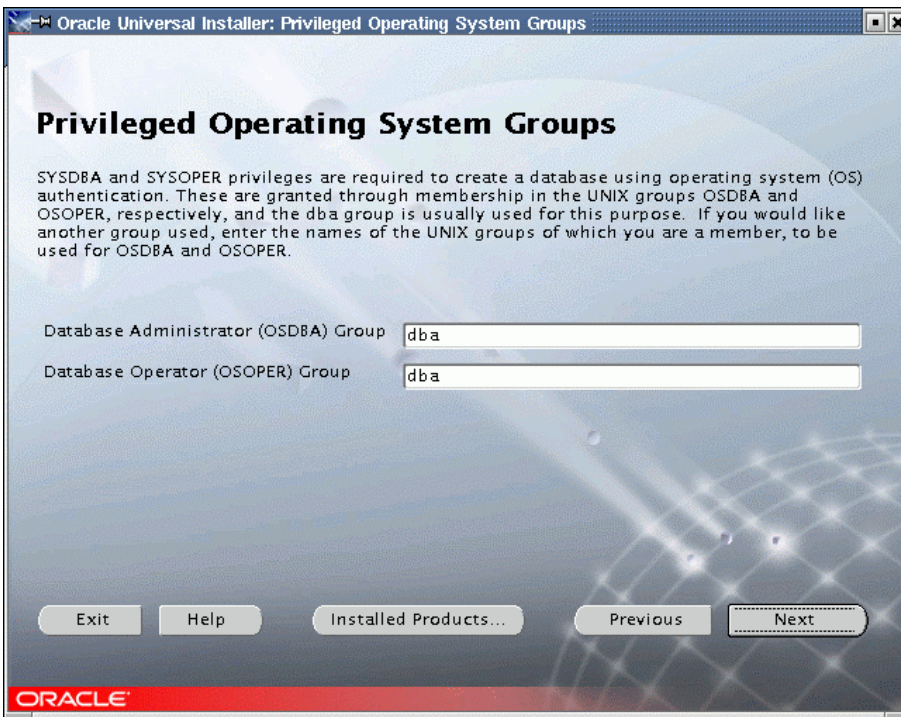


Figure 45. Privileged OS Groups

- The next screen asks whether to run the Database Configuration Assistant (Figure 46). We chose not to at this point in time. Click **Next** to continue.

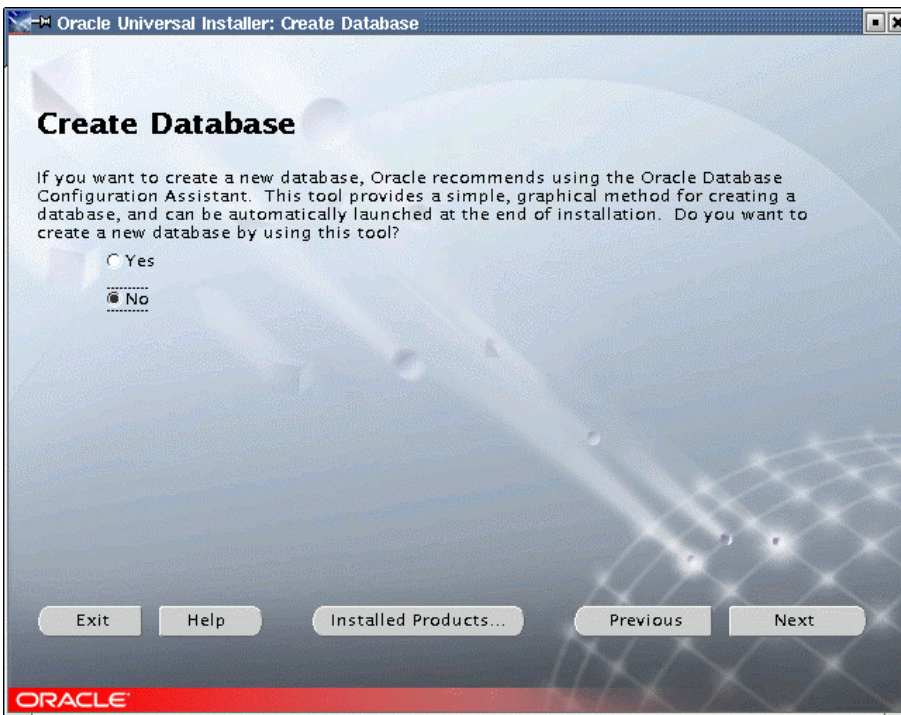


Figure 46. Create Database

- Finally, a Summary screen shows the selections made to this point. Review the settings, and if you need to add something else, go back by clicking on Back. Otherwise, if everything is OK, click **Install** to start the installation.

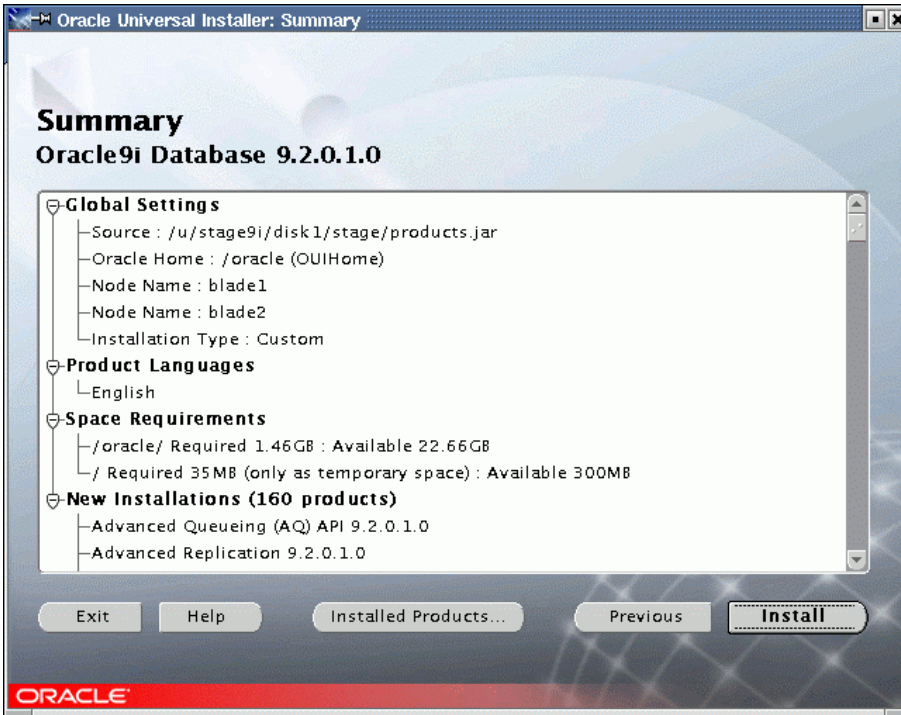


Figure 47. Summary

- During the installation process, you will see an installation progress window. After a few minutes, a pop-up window entitled Disk Location displays (Figure 48), asking for CD 2. Since we are installing from a staging directory, simply change `disk1` to `disk2` in the Path and click **OK** to continue.

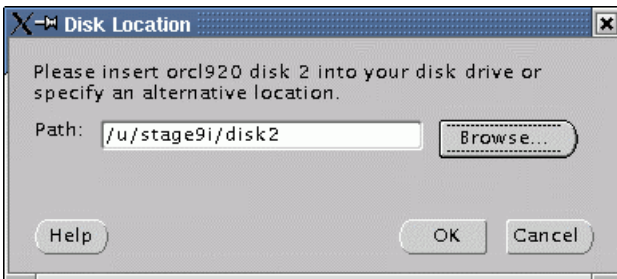


Figure 48. Disk Location

- After another few minutes, another Disk Location pop-up window is displayed for the location of CD #3's contents. Change the path as necessary and click **OK** to continue the installation.

- Once all the CD's are installed on the first node, the OUI will install the software on each node in your cluster. A pop-up window entitled Setup Privileges will ask you to execute `root.sh` (Figure 49).

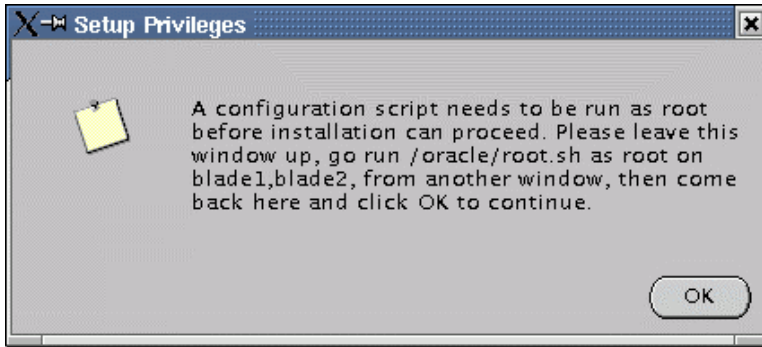


Figure 49. Setup Privileges

- Before executing `root.sh`, create the directory `/var/opt/oracle` and make `oracle` its owner on all nodes. Execute the following commands from the `root` shell prompt:

```
mkdir /var/opt/oracle
chown oracle:oinstall /var/opt/oracle
```

- Execute `root.sh` as `root` user on all nodes. When finished, click **OK** to continue.
- OUI will automatically configure the cluster and run the Network and Agent Configuration Assistants, shown in Figure 50.

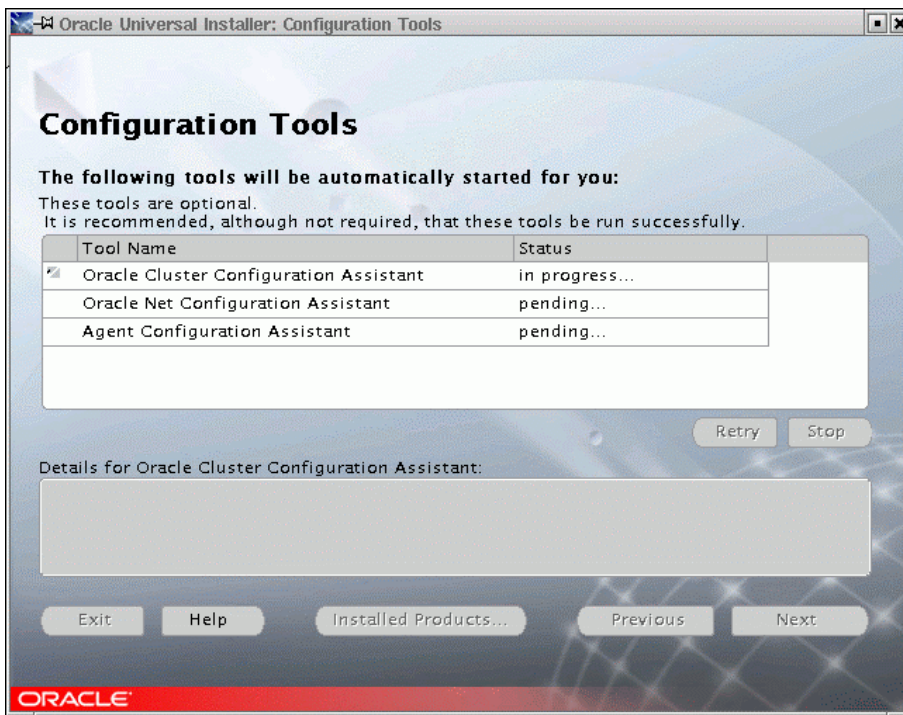


Figure 50. Configuration Tools

16. After cluster configuration, the attempt to run Network and Agent Configuration Assistants will fail. We ignored this error message, since we were going to run Network Configuration Assistant later on. Click **OK** to continue, then **Exit** to finish this installation.

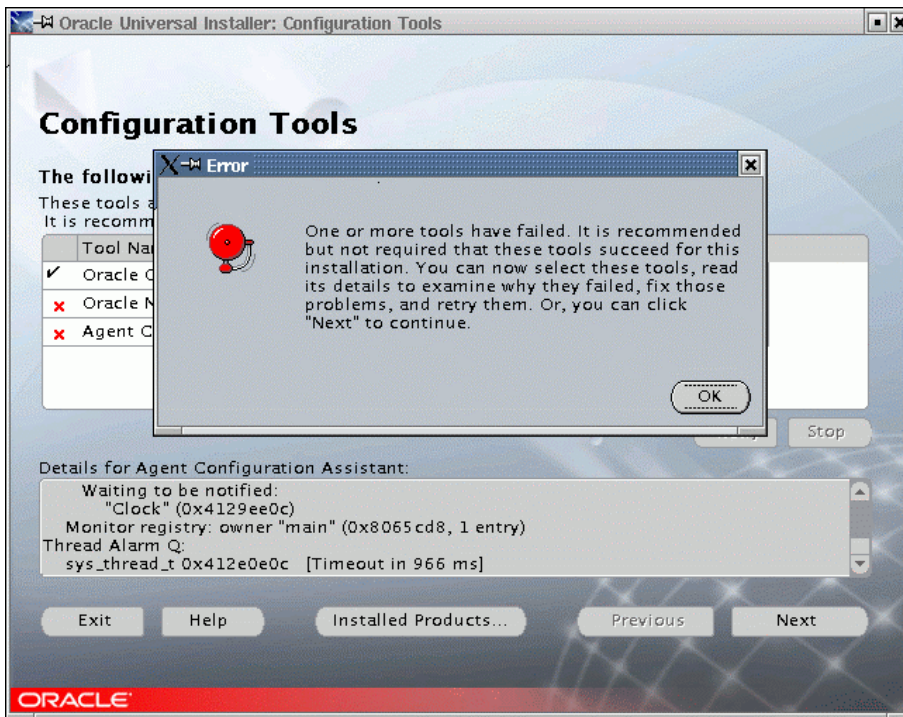


Figure 51. OUI Error

17. Some of the log directories may not have been created on all nodes during the installation. Make sure then that the following directories exist on each node:

- \$ORACLE_HOME/rdbms/audit
- \$ORACLE_HOME/rdbms/log
- \$ORACLE_HOME/network/log
- \$ORACLE_HOME/network/trace
- \$ORACLE_HOME/Apache/Apache/logs
- \$ORACLE_HOME/Apache/Jserv/logs

If they do not exist, create them, as oracle user, in order to have full logs on your system.

INSTALL PATCH SET 9204

Obtain Patch Set

We downloaded Patch Set 9204 from Oracle MetaLink. The file containing this patch set is named `p3095277_9204_LINUX.zip`. Transfer the file to one of the nodes, typically where the installation was done (blade1 in our example). As `oracle` user, uncompress and unpack the download file:

```
unzip p3095277_9204_LINUX.zip
cpio -idmv < 9204_lnx32_release.cpio
```

A subdirectory named `p3095277_9204_LINUX` is created.

Oracle Universal Installer Patch

1. Run OUI from the patch set.

```
cd /u/p3095277_9204_LINUX/Disk1
```

```
./runInstaller
```

The Welcome screen shown in Figure 52 is displayed. Click **Next** to continue.

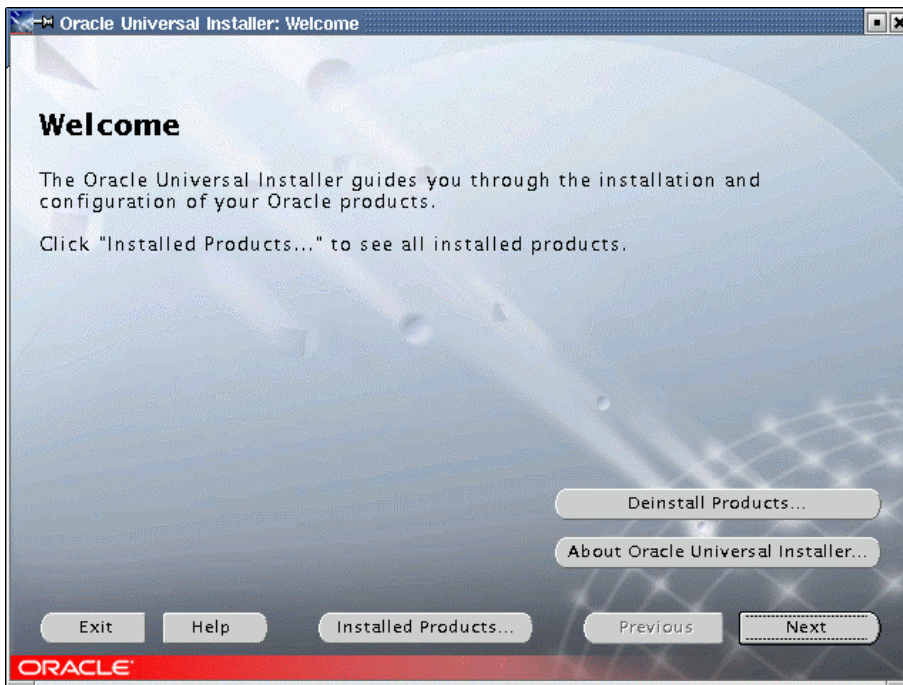


Figure 52. Welcome

2. The Cluster Node Selection screen is displayed (Figure 53). Our local node, blade1 is always selected, so select blade2 as shown. Then click **Next** to continue.

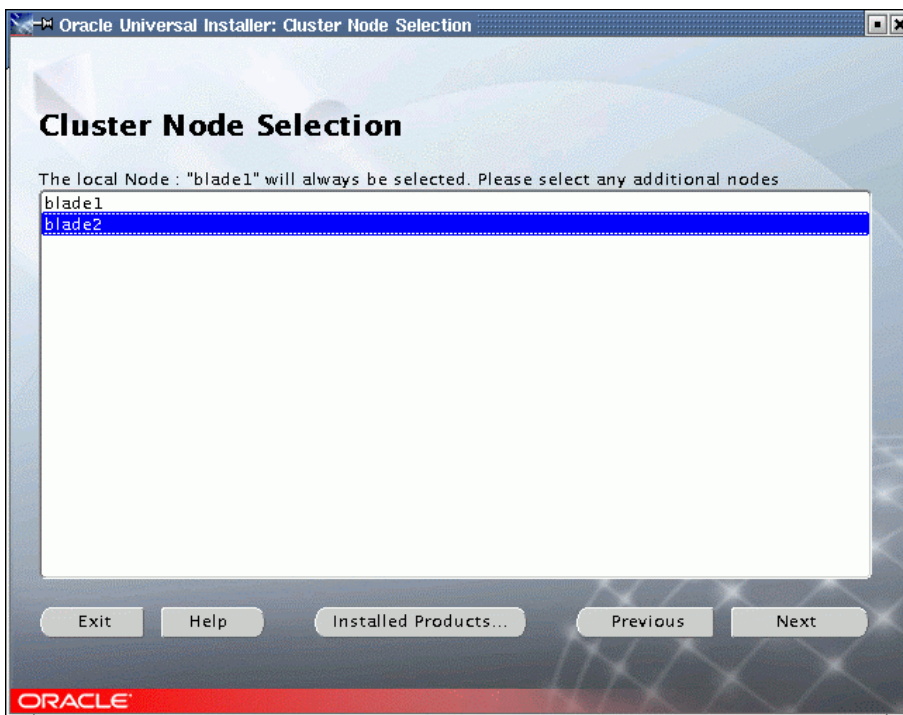


Figure 53. Cluster Node Selection

- The File Locations screen asks for the path of the install file. The unzip'd patch file is in /u, so the full path name is /u/p3095277_9204_LINUX/Disk1/stage/products.jar. Modify the screen (Figure 54) as needed and click **Next** to continue.

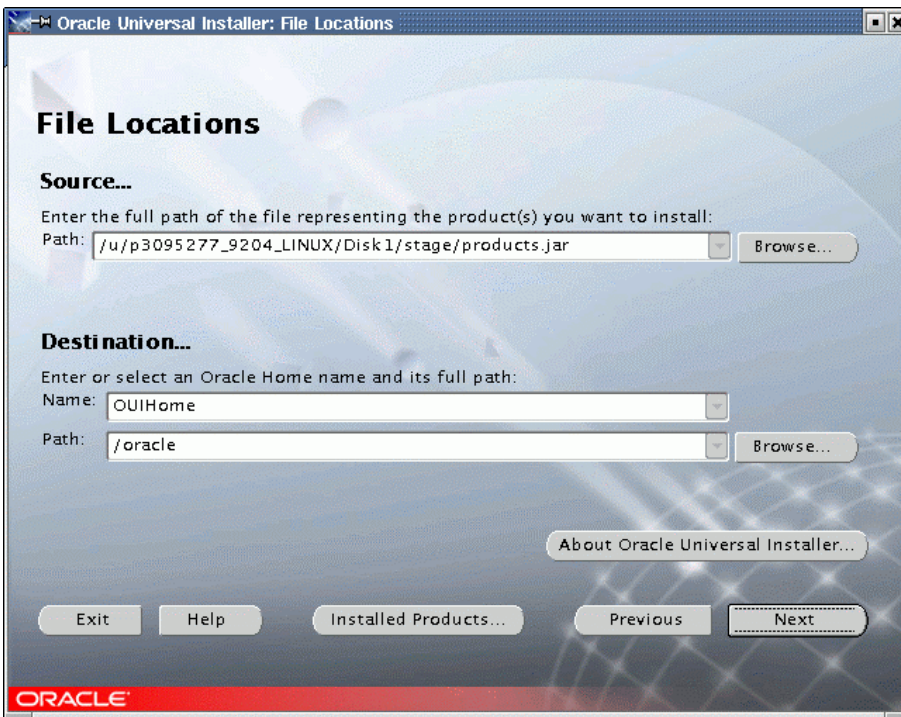


Figure 54. File Locations

- The Available Products screen is displayed next (Figure 55). Select Oracle Universal Installer 2.2.0.18.0 and click **Next** to continue.

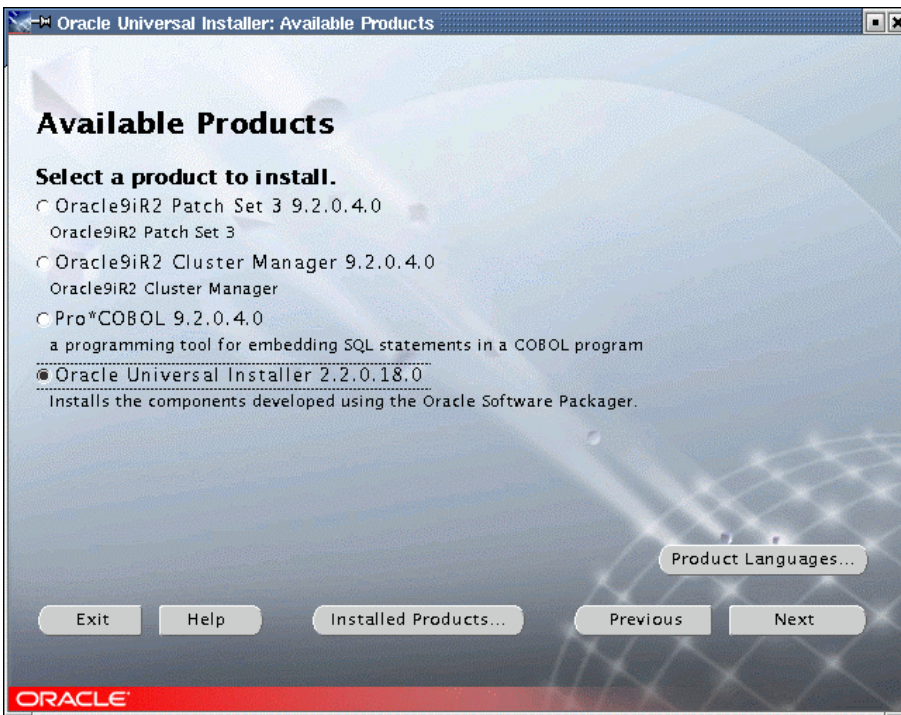


Figure 55. Available Products

- The Component Locations screen is displayed. Click **Next** to continue.

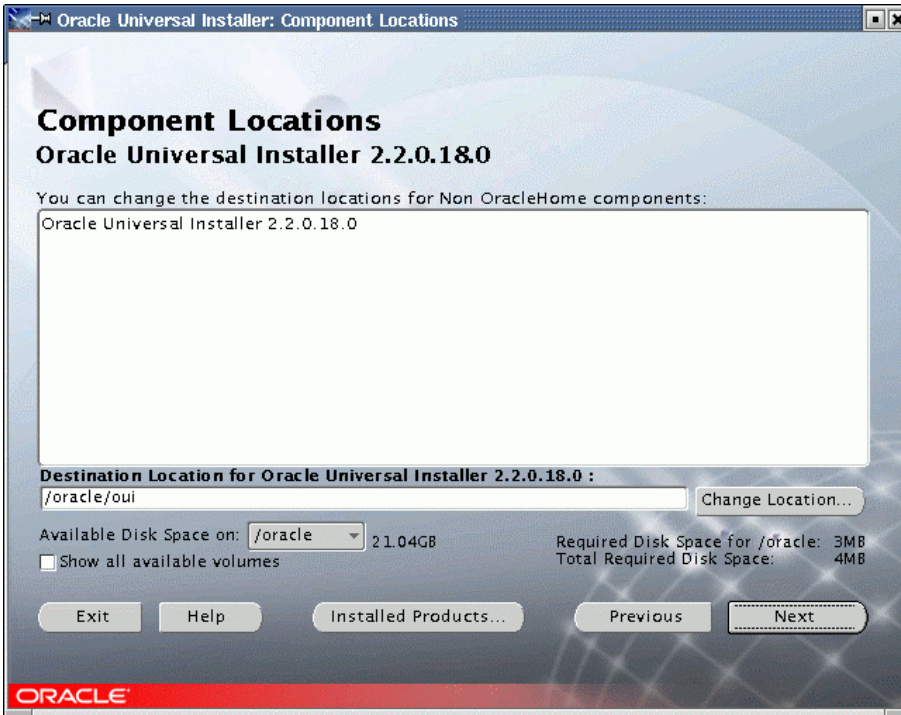


Figure 56. Component Locations

- A Summary screen is displayed (Figure 57). Click **Install** to begin the installation.

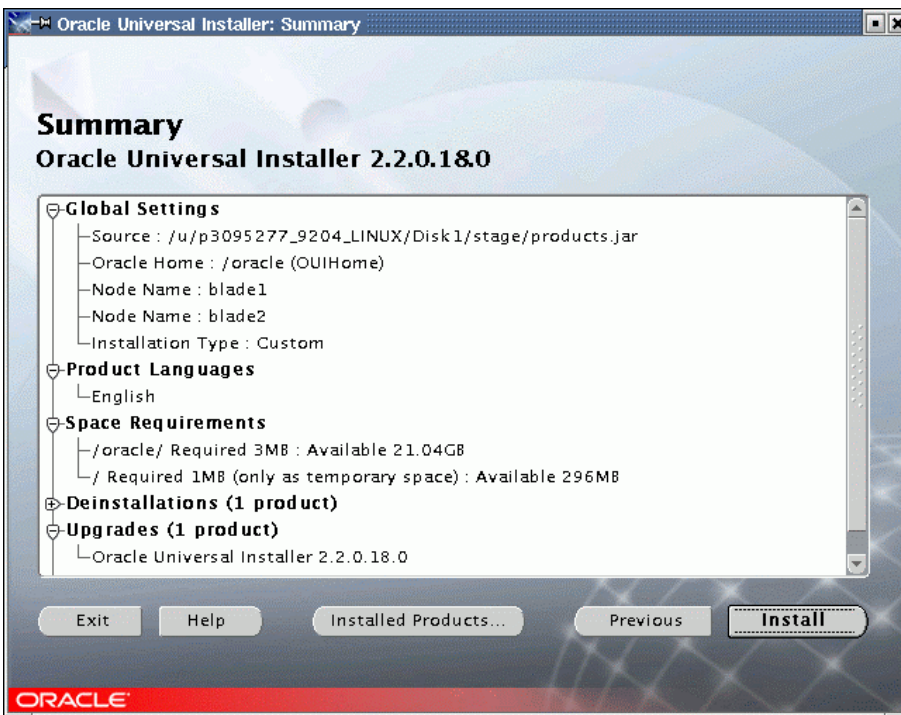


Figure 57. Summary

7. A progress screen will be displayed briefly, then the End of Installation screen (Figure 58). Click **Exit**, then **Yes** on the confirmation window.



Figure 58. End of Installation

8. Create the symbolic link as follows,

```
cd $ORACLE_HOME/oui/bin/linux
ln -s libclntsh.so.9.0 libclntsh.so
```

Cluster Manager Patch

1. Stop Cluster Manager on all nodes. As root user, use the `killall` command to stop all `oracm` processes.

```
killall oracm
```

Use the `ps -ef |grep oracm` command to make sure that all the `oracm` processes have been stopped.

2. Start OUI from `$ORACLE_HOME/bin`. Do this as `oracle` user.

```
$ORACLE_HOME/bin/runInstaller
```

The procedure is similar to that in the previous section. It starts with the same screens as from step 1, page 49 (Figure 52, Welcome screen) through step 4, page 51 (Figure 55, Available Products screen).

- On the Available Products screen, select Oracle9iR2 Cluster Manager 9.2.0.4.0, as shown in Figure 59. Click **Next** to continue.

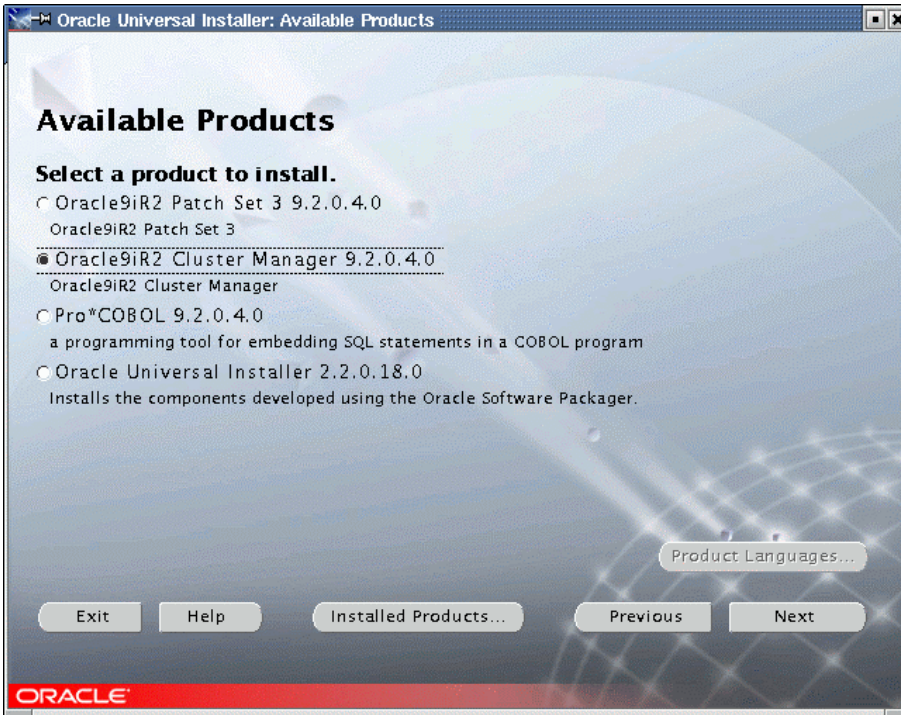


Figure 59. Available Products

- The next screens will prompt for public node information and private node information. Enter the information shown in Figure 60 and Figure 61.




Figure 60. Public Node Information




Figure 61. Private Node Information

5. A Summary screen displays next, shown in Figure 62. Click **Install** to begin the installation. A progress screen is displayed, followed by the End of Installation screen. Click **Exit**, then **Yes** on the confirmation window.

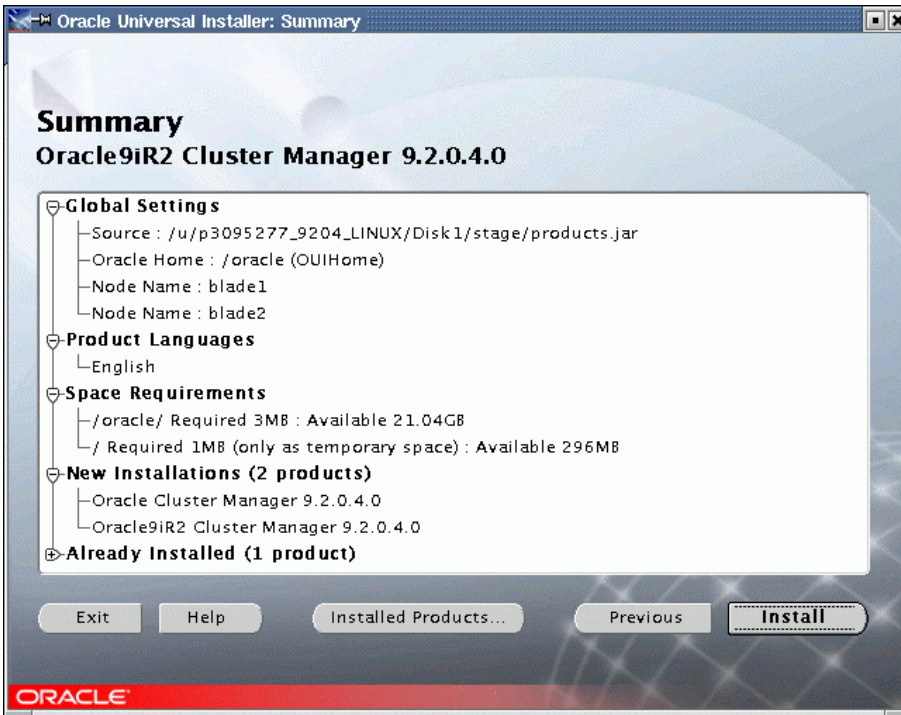


Figure 62. Summary

Note: If you encounter the popup window shown in Figure 63, this means that you have not stopped Cluster Manager. Go to another window and as root, issue the command `killall oracm`. Verify

using the `ps -ef | grep oracm` command. Be sure to stop Cluster Manager on all the nodes. When you have stopped Cluster Manager on all nodes, return to this popup window and click **Retry**.

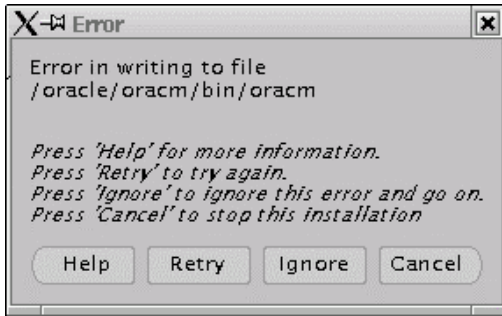


Figure 63. Error - oracm

6. Edit the Cluster Manager configuration file `$ORACLE_HOME/oracm/admin/cmcfg.ora` on all nodes as root. Add the following lines:

```
KernelModuleName=hangcheck-timer
MissCount=210
```

7. Start Cluster Manager on all nodes. Do this as root user,

```
cd /oracle/oracm/log
rm ocmstart.ts
cd ../bin
./ocmstart.sh
```

8. Verify that the Cluster Manager starts and that the `hangcheck-timer` module is loaded. As root user use the commands as shown below for blade1:

```
[root@blade1 log]# ps -ef |grep oracm
root      15374      1  0 11:53 pts/0    00:00:00 oracm
root      15376 15374  0 11:53 pts/0    00:00:00 oracm
root      15377 15376  0 11:53 pts/0    00:00:00 oracm
root      15378 15376  0 11:53 pts/0    00:00:00 oracm
root      15379 15376  0 11:53 pts/0    00:00:00 oracm
root      15380 15376  0 11:53 pts/0    00:00:00 oracm
root      15381 15376  0 11:53 pts/0    00:00:00 oracm
root      15382 15376  0 11:53 pts/0    00:00:00 oracm
root      15383 15376  0 11:53 pts/0    00:00:00 oracm
root      15396 15029  0 11:53 pts/0    00:00:00 grep oracm
[root@blade1 log]# lsmod
Module                Size  Used by    Not tainted
hangcheck-timer       2432   0 (unused)
sg                    35076  0 (autoclean)
pcmcia_core           57440  0
autofs                13796  0 (autoclean) (unused)
tg3                   50528  2
usb-storage           61996  0
mousedev              5824   1
keybdev               2848   0 (unused)
hid                   22272  0 (unused)
input                 6560   0 [mousedev keybdev hid]
usb-ohci              23392  0 (unused)
usbcore               68864  1 [usb-storage hid usb-ohci]
ext3                  74176  5
jbd                   55304  5 [ext3]
qla2300               250656  2
```

```
sd_mod          13888    2
scsi_mod        126252   5 [sg usb-storage qla2300 sd_mod]
[root@blade1 log]#
```

9. If Cluster Manager does not start, check the log file `$ORACLE_HOME/oracm/log/cm.log` for more information.
10. To start Cluster Manager during bootup, edit `/etc/rc.local`. Add the lines below Example 10:

```
#!/bin/sh
#
# This script will be executed *after* all the other init scripts.
# You can put your own initialization stuff in here if you don't
# want to do the full Sys V style init stuff.

touch /var/lock/subsys/local

/sbin/inssmod hangcheck-timer hangcheck_tick=30 hangcheck_margin=180

. /root/.bash_profile
if [ -f /oracle/oracm/log/ocmstart.ts ]; then
    echo "removing ocmstart.ts"
    rm /oracle/oracm/log/ocmstart.ts
fi
#
/oracle/oracm/bin/ocmstart.sh
#
```

Example 10. `/etc/rc.local`

Oracle9i RAC Patch

1. Start OUI from `$ORACLE_HOME/bin`. Do this as oracle user.

```
$ORACLE_HOME/bin/runInstaller
```

The procedure that follows is similar to that in the previous section. It starts with the same screens as from step 1, page 49 (Figure 52, Welcome screen) through step 4, page 51 (Figure 55, Available Products screen).

- On the Available Products screen, select Oracle9iR2 Patch Set 3 9.2.0.4.0, as shown in Figure 64. Click **Next** to continue.

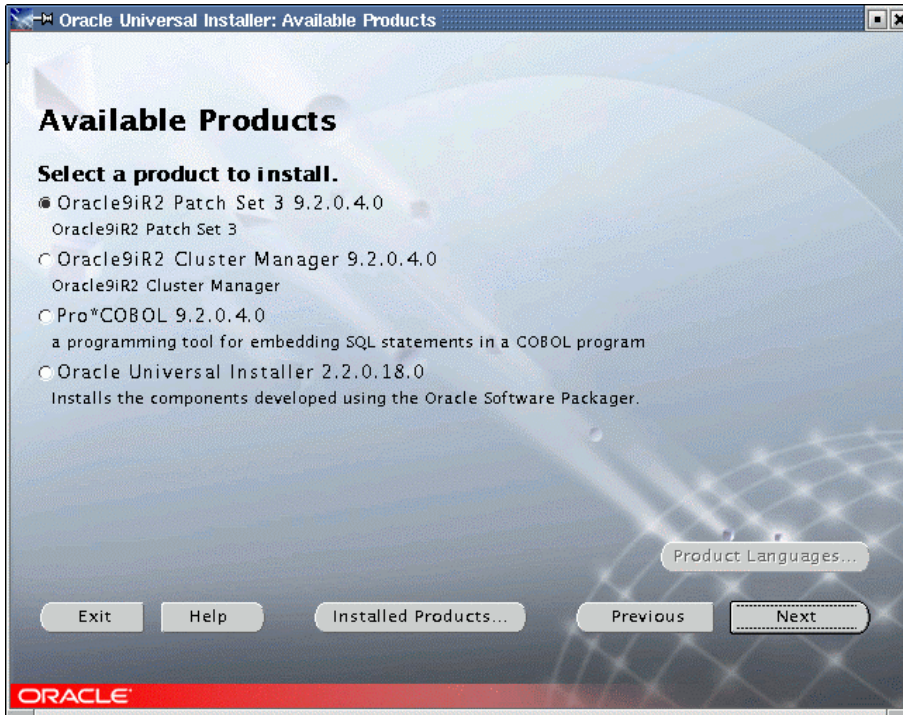


Figure 64. Available Products

- The Summary screen is displayed (Figure 65). Click **Install** to begin installation.

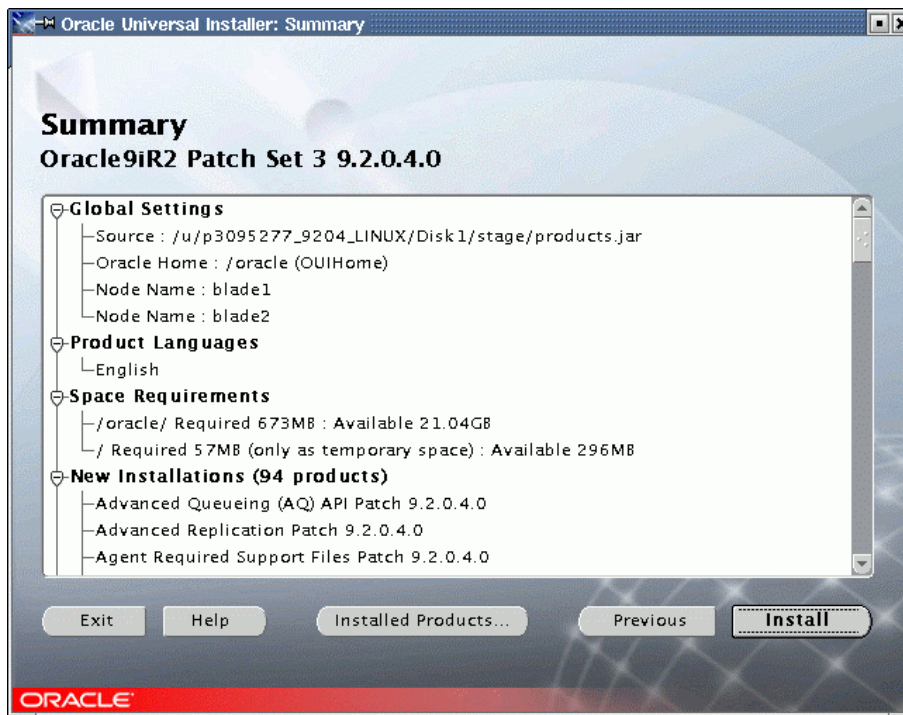


Figure 65. Summary

4. A progress screen is displayed. After a few minutes, a popup window is displayed, asking to execute a shell script as `root` user on all nodes (Figure 66). Execute the requested shell script as `root` user, then click **OK** on the popup window.

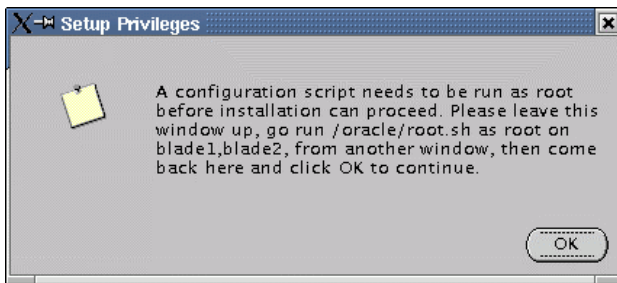


Figure 66. Setup Privileges

5. The End of Installation screen is displayed. Click **Exit** and **Yes** to confirm.
6. Start the Global Services Daemon. As `oracle` user, type “`gsdctl start`” on each node of your cluster to start the Global Services Daemon.
7. In order to have Cluster Manager and Global Services Daemon running each time you start your machine, edit `/etc/rc.local`. As `root` user,

```
#!/bin/sh
#
# This script will be executed *after* all the other init scripts.
# You can put your own initialization stuff in here if you don't
# want to do the full Sys V style init stuff.

touch /var/lock/subsys/local

/sbin/inssmod hangcheck-timer hangcheck_tick=30 hangcheck_margin=180

. /root/.bash_profile
if [ -f /oracle/oracm/log/ocmstart.ts ]; then
    echo "removing ocmstart.ts"
    rm /oracle/oracm/log/ocmstart.ts
fi
#
/oracle/oracm/bin/ocmstart.sh
#
su - oracle -c "gsdctl start"
```

Example 11. Start gsd in `/etc/rc.local`

POST INSTALLATION

Configure Listener

Before the database creation we will create the listeners that will be needed to create our database. Make sure the “LANG” variable is not set by typing “`unset LANG`”. Then you can start the Network Configuration Assistant. This tool is the one that failed during the Oracle9i installation.

1. As oracle user, type `netca` from a shell prompt. The first `netca` screen, shown in Figure 67, asks for the type of configuration. Select **Cluster configuration** and click **Next** to continue.

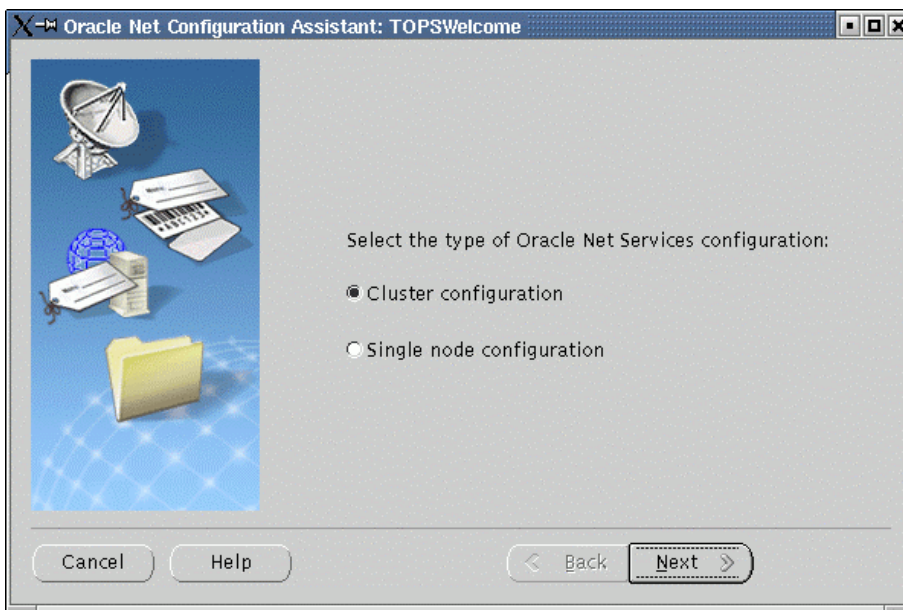


Figure 67. netca – configuration type

2. The screen shown Figure 68 displays, asking for node selection. Click **Select all nodes**, then **Next** to continue.

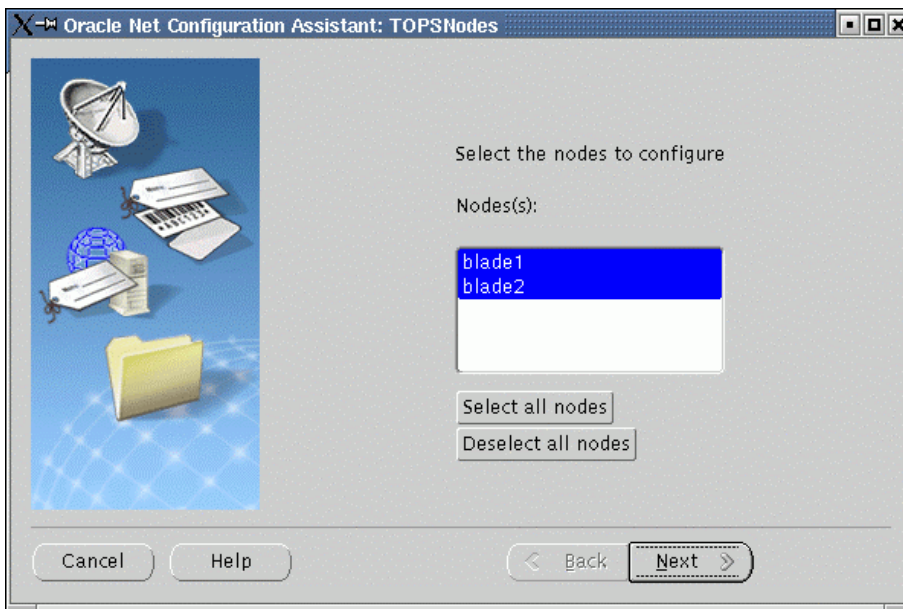


Figure 68. netca – node selection

- The netca welcome screen is displayed (Figure 69). Select **Listener configuration** and click **Next** to continue.

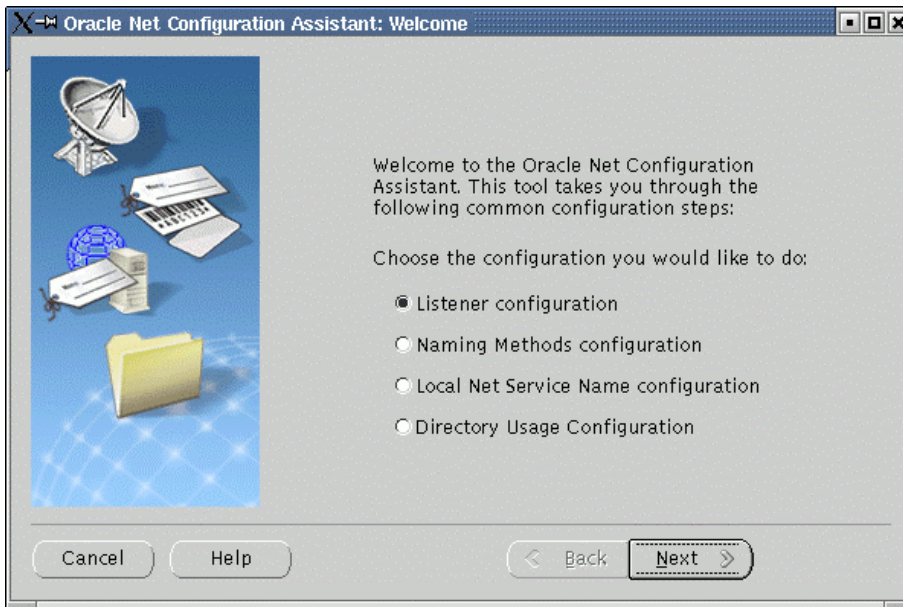


Figure 69. netca - Welcome screen

- The Listener Configuration window is displayed (Figure 70). Select **Add** and click **Next** to continue.

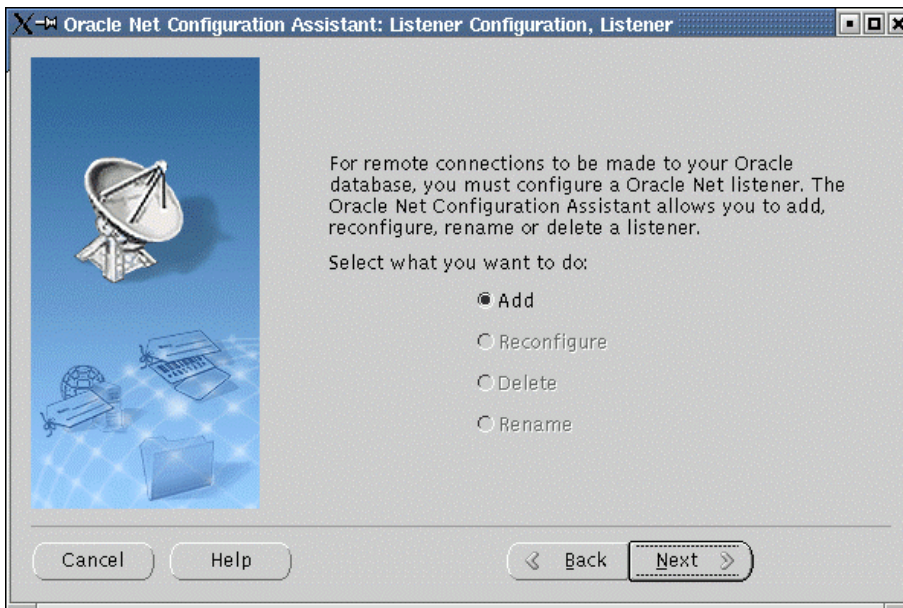


Figure 70. netca - listener configuration

- The next screen (Figure 71) asks for the listener name. Type in LISTENER as shown and click **Next** to continue.

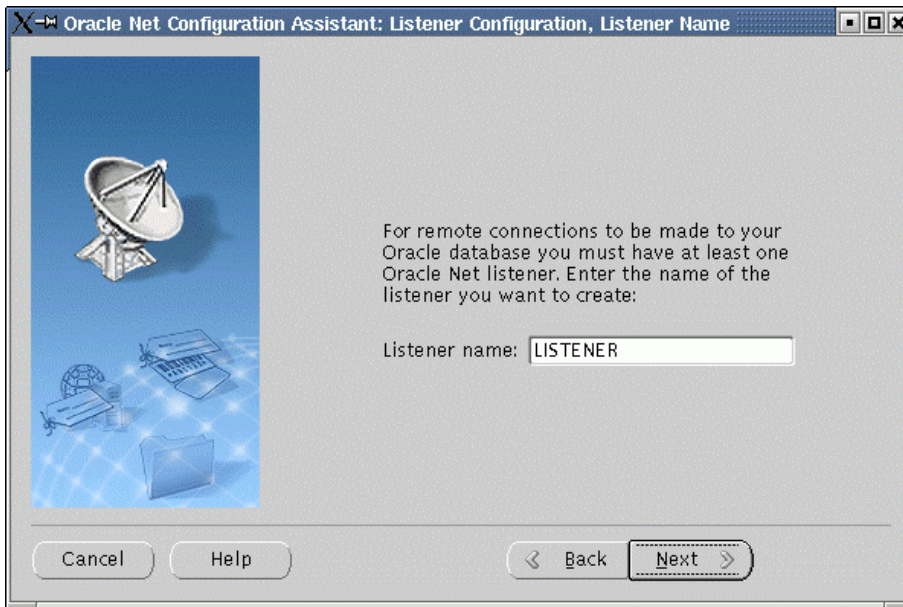


Figure 71. netca - listener name

- The Protocol Selection screen, shown in Figure 72, is displayed, with tcp/ip selected as default. Leave the default as shown and click **Next** to continue.

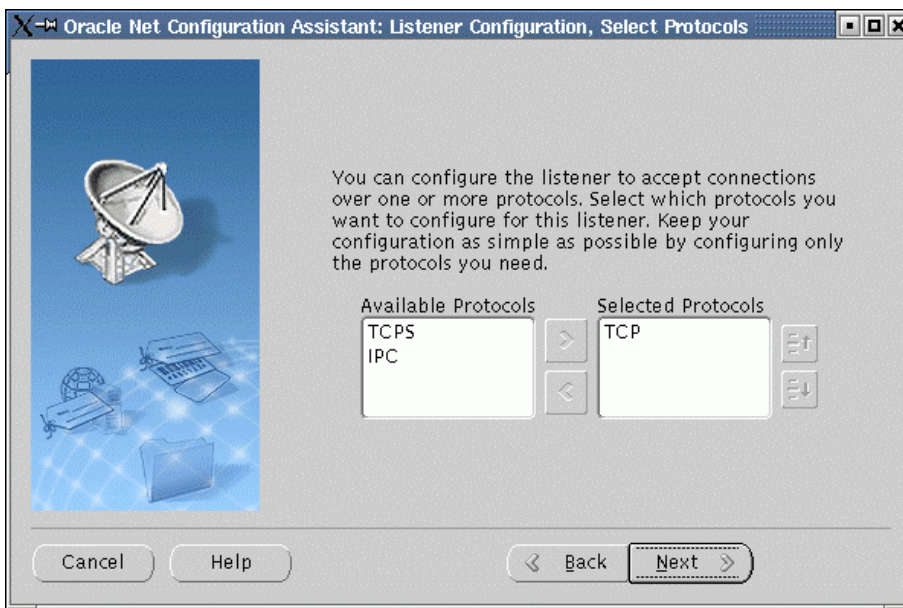


Figure 72. netca - protocol selection

- The next screen, Figure 73, asks for the tcp/ip port. Leave the default and click **Next** to continue.

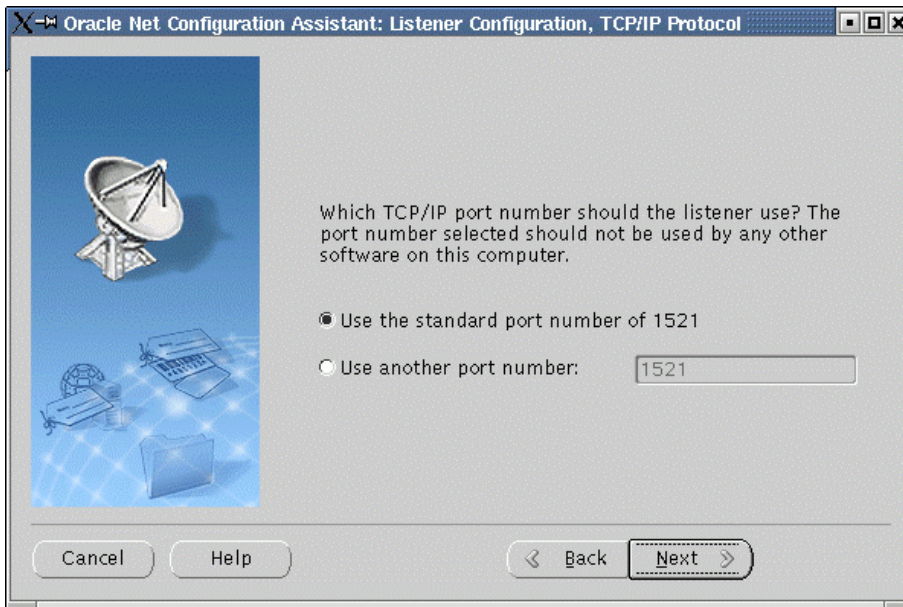


Figure 73. netca - tcp/ip protocol

- This should complete listener configuration. netca will then ask if you want to create another listener (Figure 74). Select **No** and click **Next**.



Figure 74. netca - more listeners

- Click **Next** on the next screen (Figure 75) to complete Listener Configuration.

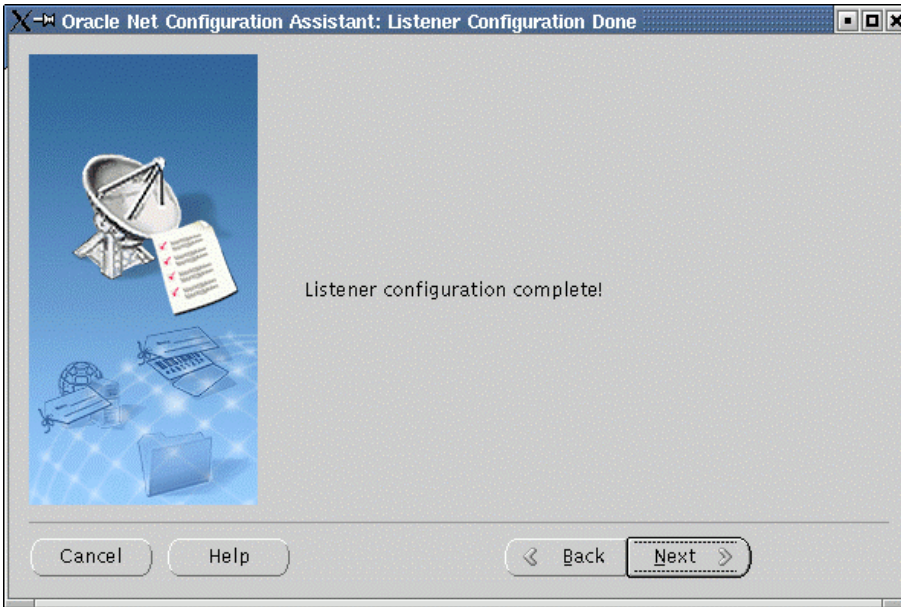


Figure 75. netca - completion

CREATING A TEST DATABASE

To test our installation of Oracle9i RAC, we created a database using Oracle’s Database Configuration Assistant, dbca. Oracle recommends the use of DBCA to create a database. This is because the DBCA preconfigured databases optimize the environment to take advantage of Oracle9i features such as the server parameter file and automatic undo management. DBCA also enables you to define your own tablespaces as part of the database creation process. Even if the datafile requirements differ from those offered in one of the DBCA templates, Oracle recommends to use DBCA. You can also execute user-specified scripts as part of the database creation process.

In order to run DBCA, you must have the GSD daemon started. If not, be sure you run it before starting database creation.

To help configure and create the raw devices, we created the table shown in Table 5. This will help us determine how to configure the shared storage device, how many partitions to create, what size to make them, and to document the purpose of each raw device.

The first column in Table 5 contains descriptions of the database files that will be on the raw devices, the second contains the sizes of the database files. The third and fourth columns contain the disk partitions and the raw devices. The fifth column contains the symbolic links to the raw devices. (Note the table contains the two partitions, quorum and srvm, created earlier, prior to software installation.)

We chose to name our test databases RACDBx, where x = the instance number. Therefore, our symbolic links were placed in a subdirectory name RACDB as shown in Table 5 below.

Database file	Size (MB)	Disk Partition	Raw device	Symbolic Link
quorum	110	/dev/sda1	/dev/raw/raw1	\$ORACLE_HOME/oradata/quorum
srvm	110	/dev/sda2	/dev/raw/raw2	\$ORACLE_HOME/oradata/srvm
spfile	5	/dev/sda3	/dev/raw/raw3	\$ORACLE_HOME/oradata/RACDB / spfile
control file 1	110	/dev/sda5	/dev/raw/raw4	\$ORACLE_HOME/oradata/RACDB / control01
control file 2	110	/dev/sda6	/dev/raw/raw5	\$ORACLE_HOME/oradata/RACDB / control02

control file 3	110	/dev/sda7	/dev/raw/raw6	\$ORACLE_HOME/oradata/RACDB/control03
system	250	/dev/sda8	/dev/raw/raw7	\$ORACLE_HOME/oradata/RACDB/system
temp	40	/dev/sda9	/dev/raw/raw8	\$ORACLE_HOME/oradata/RACDB/temp
tools	10	/dev/sda10	/dev/raw/raw9	\$ORACLE_HOME/oradata/RACDB/tools
users	25	/dev/sdb1	/dev/raw/raw10	\$ORACLE_HOME/oradata/RACDB/users
indx	25	/dev/sdb2	/dev/raw/raw11	\$ORACLE_HOME/oradata/RACDB/index
redo log thread 1 group 1	120	/dev/sdb3	/dev/raw/raw12	\$ORACLE_HOME/oradata/RACDB/redo1_1
redo log thread 1 group 2	120	/dev/sdb5	/dev/raw/raw13	\$ORACLE_HOME/oradata/RACDB/redo1_2
redo log thread 2 group 1	120	/dev/sdb6	/dev/raw/raw14	\$ORACLE_HOME/oradata/RACDB/redo2_1
redo log thread 2 group 2	120	/dev/sdb7	/dev/raw/raw15	\$ORACLE_HOME/oradata/RACDB/redo2_2
undo tablespace thread 1	200	/dev/sdb8	/dev/raw/raw16	\$ORACLE_HOME/oradata/RACDB/undo1
undo tablespace thread 2	200	/dev/sdb9	/dev/raw/raw17	\$ORACLE_HOME/oradata/RACDB/undo2

Table 5. Database files and raw devices

1. As `root` user, create the disk partitions using `fdisk`.
2. Bind the raw devices to the disk partitions. Do this by editing `/etc/sysconfig/rawdevices` (Example 12) as `root` user.

```

## raw device bindings
## format: <rawdev> <major> <minor>
##           <rawdev> <blockdev>
## example: /dev/raw/raw1 /dev/sda1
##           /dev/raw/raw2 8 5
/dev/raw/raw1 /dev/sda1
/dev/raw/raw2 /dev/sda2
/dev/raw/raw3 /dev/sda3
/dev/raw/raw4 /dev/sda5
/dev/raw/raw5 /dev/sda6
/dev/raw/raw6 /dev/sda7
/dev/raw/raw7 /dev/sda8
/dev/raw/raw8 /dev/sda9
/dev/raw/raw9 /dev/sda10
/dev/raw/raw10 /dev/sdb1
/dev/raw/raw11 /dev/sdb2
/dev/raw/raw12 /dev/sdb3
/dev/raw/raw13 /dev/sdb5
/dev/raw/raw14 /dev/sdb6
/dev/raw/raw15 /dev/sdb7
/dev/raw/raw16 /dev/sdb8
/dev/raw/raw17 /dev/sdb9
    
```

Example 12. `/etc/sysconfig/rawdevices` with devices for database

3. Bind the raw devices with the following command, as `root` user,

```
service rawdevices start
```

If Cluster Manager is running at this time, an error message “Error setting raw device (Device or resource busy)”. Ignore this message.

Check the bindings with the `raw -qa` command.

4. Create symbolic links to the raw devices, as `oracle` user. Create a shell script with the commands shown in Example 13 and run it on all cluster nodes.

```
ln -s /dev/raw/raw3 $ORACLE_HOME/oradata/RACDB/spfile
ln -s /dev/raw/raw4 $ORACLE_HOME/oradata/RACDB/control01
ln -s /dev/raw/raw5 $ORACLE_HOME/oradata/RACDB/control02
ln -s /dev/raw/raw6 $ORACLE_HOME/oradata/RACDB/control03
ln -s /dev/raw/raw7 $ORACLE_HOME/oradata/RACDB/system
ln -s /dev/raw/raw8 $ORACLE_HOME/oradata/RACDB/temp
ln -s /dev/raw/raw9 $ORACLE_HOME/oradata/RACDB/tools
ln -s /dev/raw/raw10 $ORACLE_HOME/oradata/RACDB/users
ln -s /dev/raw/raw11 $ORACLE_HOME/oradata/RACDB/index
ln -s /dev/raw/raw12 $ORACLE_HOME/oradata/RACDB/redol_1
ln -s /dev/raw/raw13 $ORACLE_HOME/oradata/RACDB/redol_2
ln -s /dev/raw/raw14 $ORACLE_HOME/oradata/RACDB/redo2_1
ln -s /dev/raw/raw15 $ORACLE_HOME/oradata/RACDB/redo2_2
ln -s /dev/raw/raw16 $ORACLE_HOME/oradata/RACDB/undol
ln -s /dev/raw/raw17 $ORACLE_HOME/oradata/RACDB/undo2
```

Example 13. Create symbolic links to raw devices

5. The Database Configuration Assistant (DBCA) makes use of a file that maps database tablespaces to datafiles. Create this file to suit your own tablespaces and make sure it is owned by the oracle user. See Example 14 below for a listing of the file we used in our example, which we named `dbca_mapping.cfg`.

```
spfile=/oracle/oradata/RACDB/spfile
control1=/oracle/oradata/RACDB/control01
control2=/oracle/oradata/RACDB/control02
control3=/oracle/oradata/RACDB/control03
system=/oracle/oradata/RACDB/system
temp=/oracle/oradata/RACDB/temp
tools=/oracle/oradata/RACDB/tools
users=/oracle/oradata/RACDB/users
indx=/oracle/oradata/RACDB/indx01
redol_1=/oracle/oradata/RACDB/redol_1
redol_2=/oracle/oradata/RACDB/redol_2
redo2_1=/oracle/oradata/RACDB/redo2_1
redo2_2=/oracle/oradata/RACDB/redo2_2
undotbs1=/oracle/oradata/RACDB/undotbs01
undotbs2=/oracle/oradata/RACDB/undotbs02
```

Example 14. DBCA mapping file – `dbca_mapping.cfg`

6. The name of this mapping file is passed to DBCA by the environmental variable `DBCA_RAW_CONFIG`. As `oracle` user, set this variable by entering:

```
export DBCA_RAW_CONFIG=~/dbca_mapping.cfg
```

- As `oracle` user, type the `dbca` command at the shell prompt. The dbca welcome screen is displayed, as shown in Figure 76. Select **Oracle Cluster Database** and click **Next** to continue.

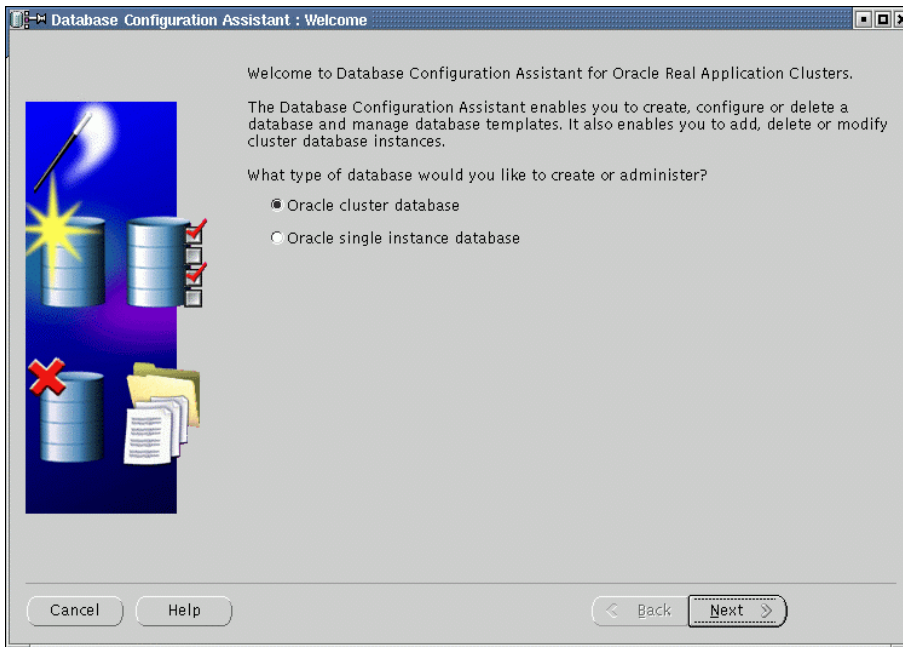


Figure 76. dbca welcome

- A screen entitled Operations is displayed (Figure 77). Select **Create a database** and click **Next** to continue.

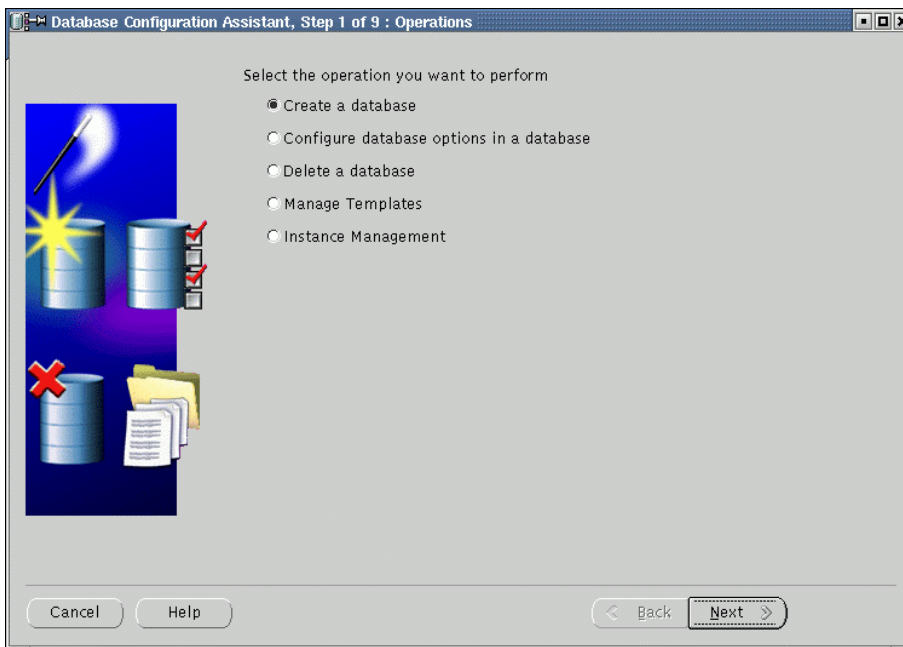


Figure 77. dbca operations

- The next screen, Figure 78, asks for the nodes on which to create the database. In our case, blade1 and blade2 are listed, as they should be. Click **Select All**, then **Next** to continue. If your nodes are not listed, use the `lsnodes` command to determine which nodes are active. Cancel dbca, resolve the problem, and restart dbca.

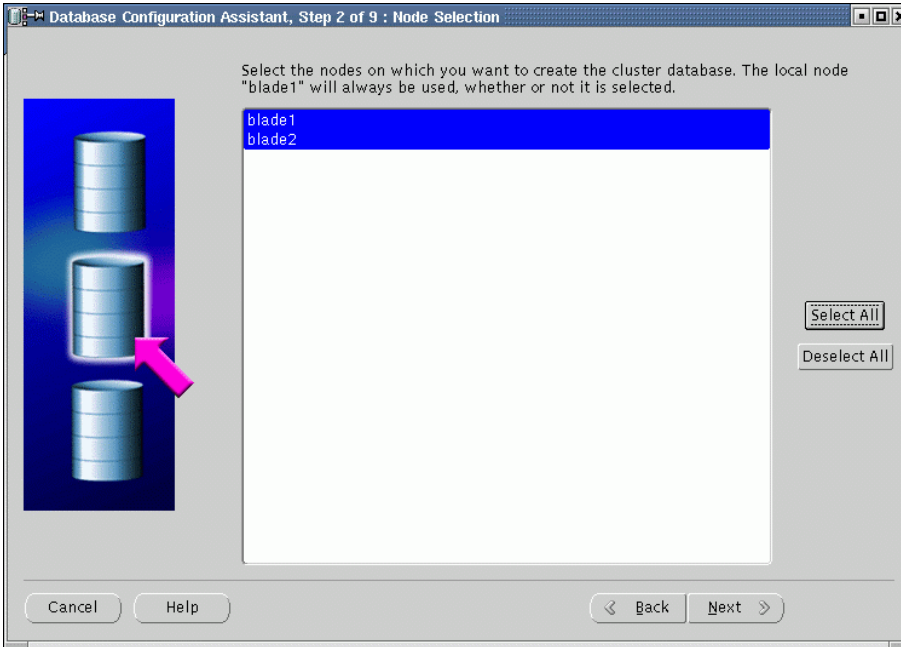


Figure 78. dbca node selection

- The Database Templates screen is displayed. Select **New Database** and click **Next** to continue.

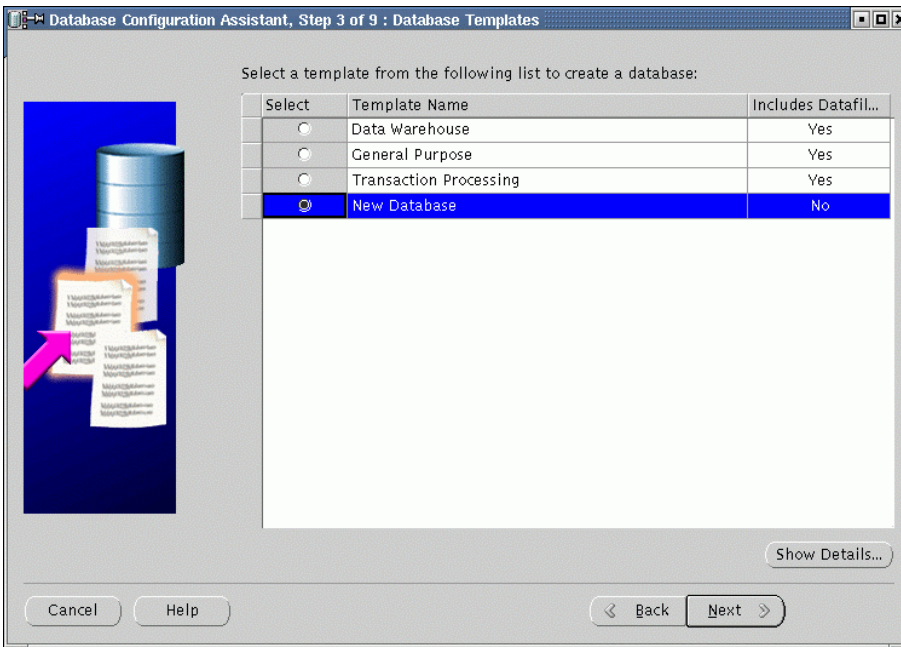


Figure 79. dbca database templates

- The next screen (Figure 80) asks for the names of our database. We entered **RACDB** for Global Database Name and left SID Prefix as suggested by dbca.




Figure 80. dbca database identification

- The Database Features screen (Figure 81) is displayed. In our example, we deselected all the features shown. Click **Standard database features** to deselect standard features.

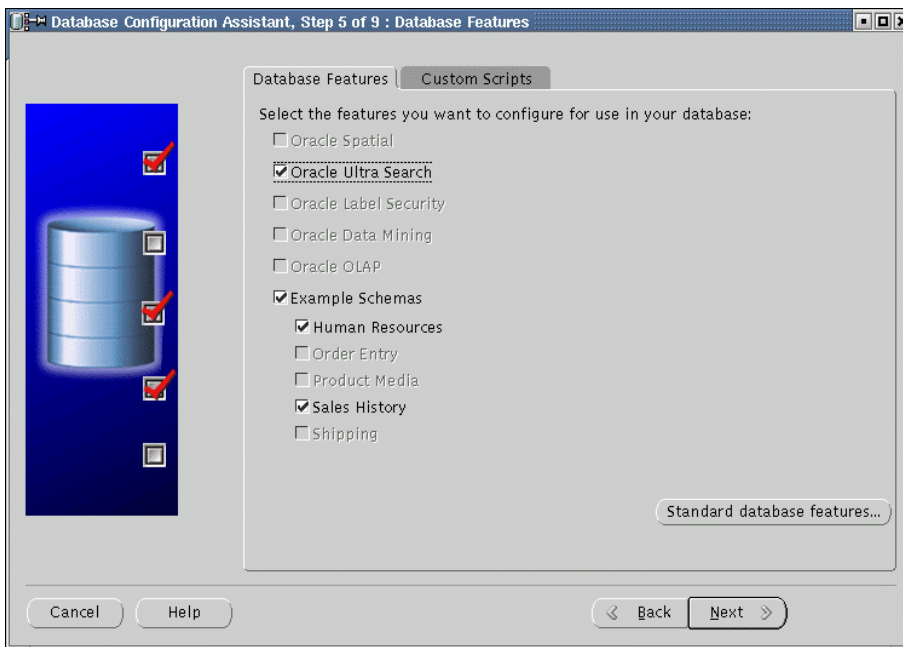


Figure 81. dbca database features

13. A popup window for Standard database features is displayed (Figure 82). Again in our example, we deselected all standard features. Click **OK** to continue.

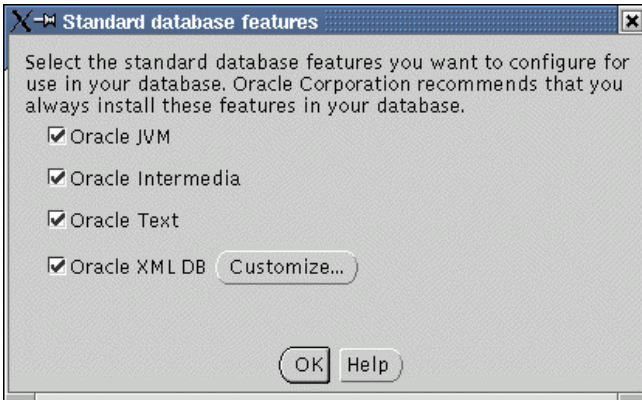


Figure 82. dbca standard database features

14. The next screen is for Database Connection Options (Figure 83). Select **Dedicated Server Mode** and click **Next** to continue.

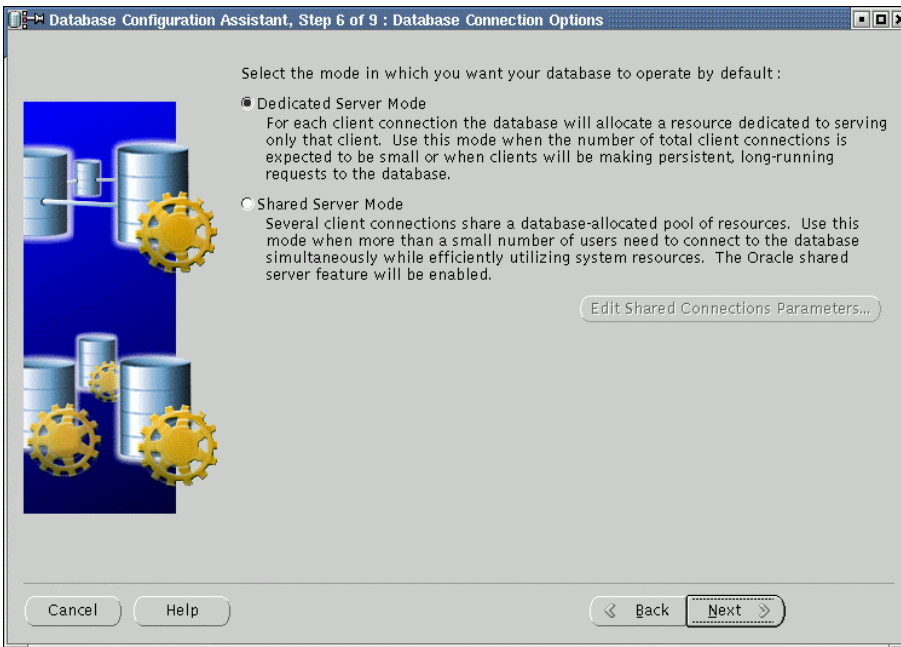


Figure 83. dbca database connection options

- The next screen displays initialization parameters (Figure 84). Click the File Locations tab. Note that **Create server parameters file (spfile)** is checked by default. **Server Parameters Filename** should point to the correct datafile (symbolic link for spfile). If not, either DBCA_RAW_CONFIG is not set correctly or the mapping file contains errors. Click on the other tabs and verify all the settings are according to your requirements. When this is done, click **Next** to continue.

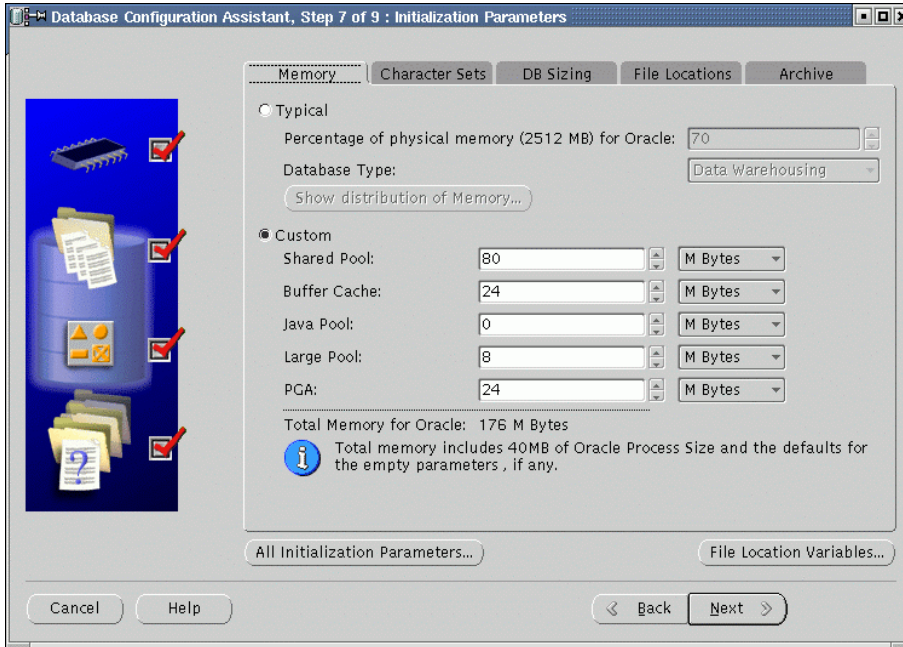


Figure 84. dbca initialization parameters

- DBCA now shows the Database Storage Window (Figure 85). The tablespaces and datafiles should be as specified in the mapping file, `dbca_mapping.cfg`. Select from the left hand panel and review the information shown on the right hand panel. When finished, click **Next** to continue.

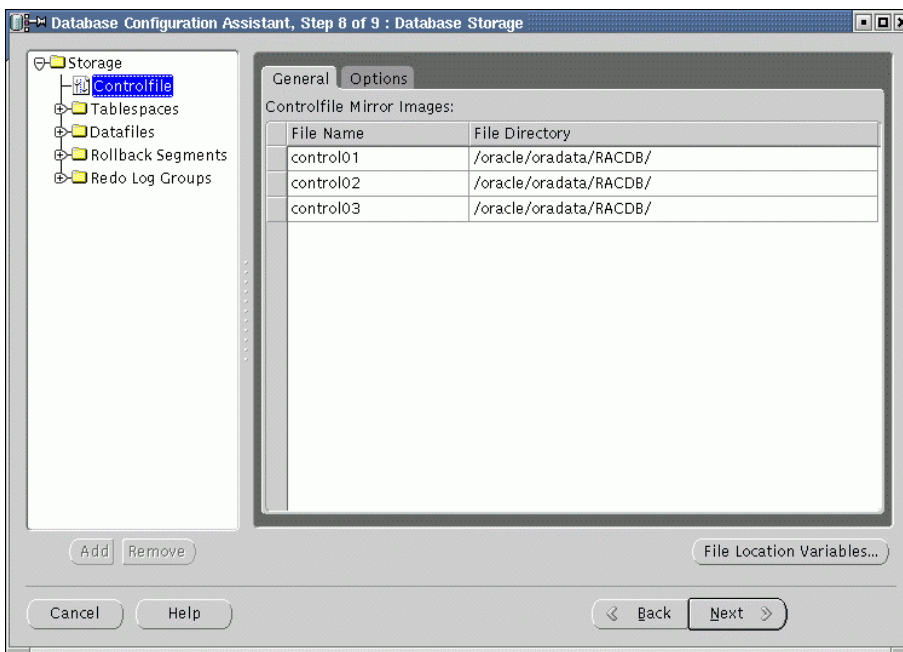


Figure 85. dbca database storage – control files

17. The next screen (Figure 86) displays Creation Options. Ensure that **Create Database** is checked and click **Finish**.

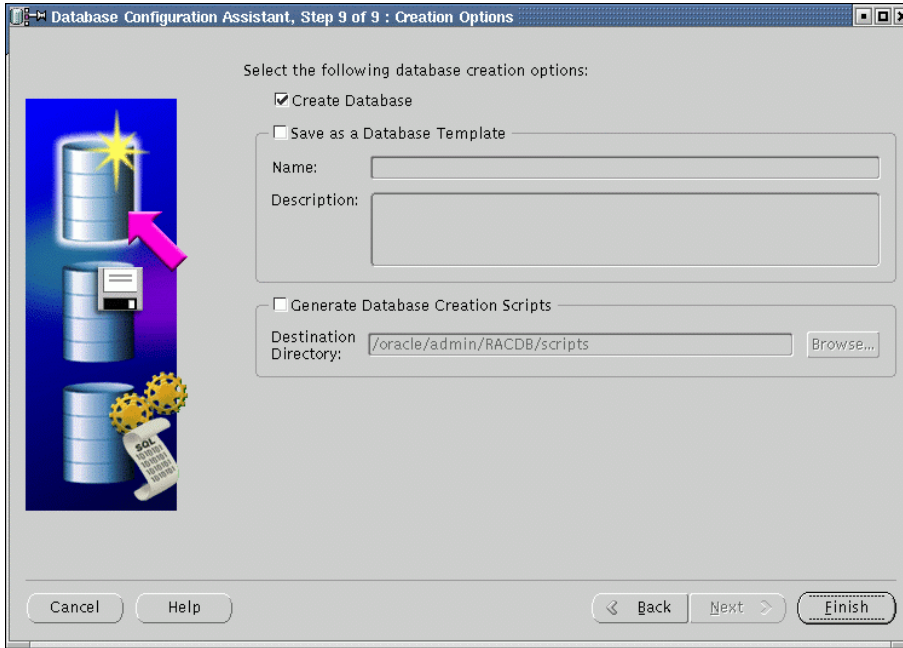


Figure 86. dbca creation options

18. A summary screen is displayed (Figure 87). Use the scroll bar on the right hand side to view all information, then click **OK** to continue.

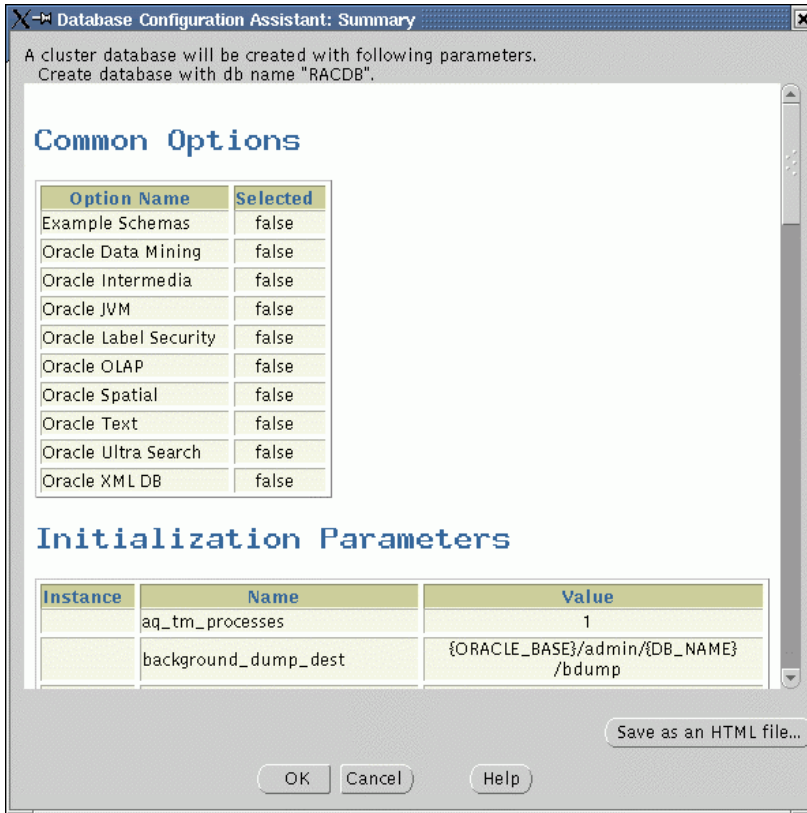


Figure 87. dbca summary

19. Database creation begins. A progress screen (Figure 88) is displayed, showing the various steps being performed.




Figure 88. dbca progress display

20. We encountered the popup window, shown in Figure 89, warning of error message ORA-29807. According to Metalink docid 237486.1, this message may be ignored. Click **Ignore** to continue.




Figure 89. dbca error message

21. When the database creation is completed, `dbca` displays the screen shown in Figure 90, asking for new passwords. Enter the passwords and click **Exit**.

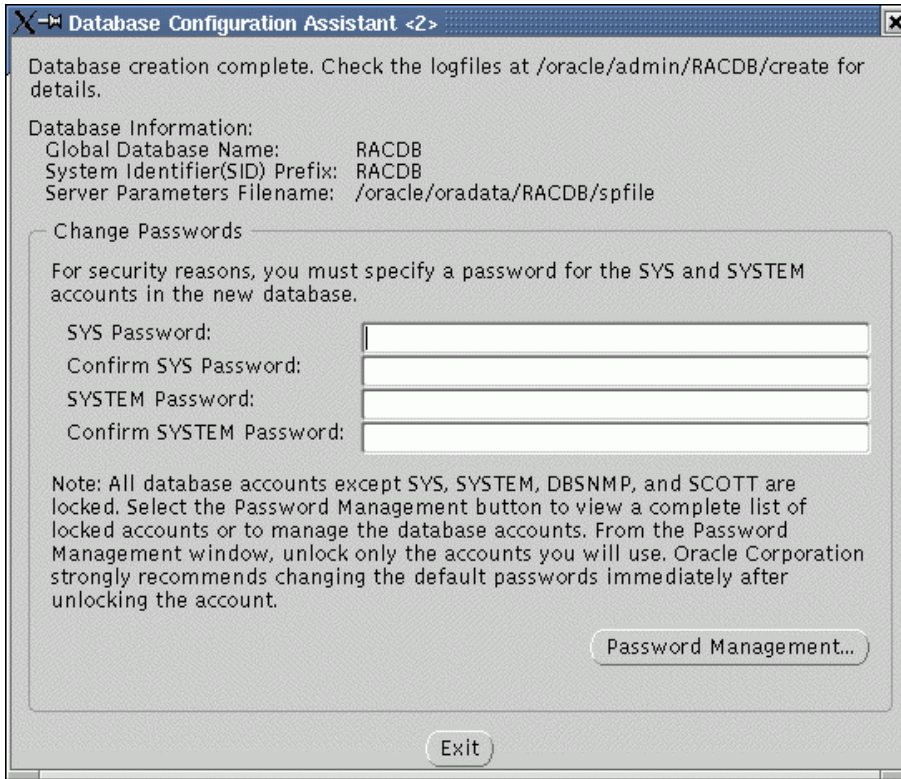


Figure 90. dbca password management

Once this step finished, you have successfully created the clustered database. It can now be accessed by Oracle SQL*PLUS or other applications designed to work with an Oracle RAC database. You should make sure to execute start-up scripts during system startup as root and other scripts or commands as oracle user as described earlier in this document.

22. The database has been created and each node has its own instance. Update oracle user's `.bash_profile` on each node's with the correct `ORACLE_SID` value. The instance name for a node is the SID prefix concatenated with the instance number. In our example, the instance on blade1 will be `RACDB1` and the instance on blade2 will be `RACDB2`.

For example, in blade1, we added the line

```
export ORACLE_SID=RACDB1
```

to oracle user's `.bash_profile`. Source the new `.bash_profile` in your current shell by issuing the command:

```
. ~/.bash_profile
```

USING SRVCTL TO MANAGE RAC INSTANCES

Oracle Corporation recommends that the `srvctl` command be used to manage a Real Application Clusters database environment. `srvctl` manages configuration information in the Server Management (SRVM) configuration repository, also used by Oracle tools such as Oracle Enterprise Manager and Oracle Intelligent Agent.

Before using `srvctl`, ensure that the Global Services Daemon (GSD) is running. To use `srvctl`, you must have already created the configuration information for the database. This would have been done either by using the Oracle Database Configuration Assistant (DBCA) or the `srvctl add` command. If you followed the instructions in this document, `dca` will have added your database and instances.

To display information in the SRVM, use the `srvctl config` command as shown:

```
[oracle@blade1 oracle]$ srvctl config database -d RACDB
blade1 RACDB1 /oracle
blade2 RACDB2 /oracle
[oracle@blade1 oracle]$
```

Examples of starting and stopping RAC follow:

```
[oracle@blade1 oracle]$ srvctl status database -d RACDB
Instance RACDB1 is running on node blade1
Instance RACDB2 is running on node blade2
[oracle@blade1 oracle]$ srvctl stop database -d RACDB
[oracle@blade1 oracle]$ srvctl status database -d RACDB
Instance RACDB1 is not running on node blade1
Instance RACDB2 is not running on node blade2
[oracle@blade1 oracle]$ srvctl start database -d RACDB
[oracle@blade1 oracle]$ srvctl status database -d RACDB
Instance RACDB1 is running on node blade1
Instance RACDB2 is running on node blade2
[oracle@blade1 oracle]$ srvctl stop instance -d RACDB -i RACDB2
[oracle@blade1 oracle]$ srvctl status database -d RACDB
Instance RACDB1 is running on node blade1
Instance RACDB2 is not running on node blade2
[oracle@blade1 oracle]$
```

For further information on `srvctl` see the Oracle9i Real Application Clusters Administration Release 2 (9.2) manual (References, 20).

ADDING A NODE TO OUR EXISTING RAC CLUSTER

One of the key features of Oracle9i RAC is the capability to increase processing capacity by adding another node, or “scale out”. This is easily done in the BladeCenter - simply add (assuming there are available slots in the BladeCenter chassis) another blade server, or redeploy an existing blade server.

In our example, we redeployed an existing blade server, `blade3`, to our two-node RAC implementation. Our approach was to make the third node identical to the first two, as much as possible, to simplify system management. For example, we kept many of the system files identical, including `/etc/sysctl.conf`, `/etc/sysconfig/rawdevices`, `/etc/rc.local`, `/etc/hosts`, etc.

Because ours is a lab environment, we are able to stop or start our systems without worrying about disrupting production work and impacting service level agreements. Following is the procedure we used. However, if you are changing a production environment, additional work is recommended to minimize the potential risk of disrupting production.

1. As a precaution, back up all database files, if this has not already been done. The `dd` command may be used to do this. Be sure to stop the database and Cluster Manager prior to using the `dd` command. For example, as `oracle` user, on `blade1`:

```
srvctl stop database -d RACDB                                ← first stop the database
su -
```

```

<password>
killall oracm
exit
dd if=/oracle/oradata/<DB file> of=<backup directory>/<DB file>
.
dd if=/oracle/oradata/RACDB/<DB file> of=<backup directory>/RACDB/<DB file>
.
.
    
```

2. Set up the blade3 as described in **Setting Up Operating Environment** beginning on page 21. The key areas are:
 - Operating system and operating system patches
 - Network interfaces and equivalence
 - Fibre channel adapter driver

Be sure to verify that:

- Both eth0 and eth1 are functional. For our example, use the `ping` command to test blade3 and iblade3 host names.
- `rsh` and `rlogin` can be run between all the nodes, as `oracle` user, without a password. Do this for both the public and private interconnect networks
- The logical drives on the FASiT are visible to the new node. Use `fdisk -l` as root user to verify.

3. Install the Fibre Channel Expansion Card driver as described in steps 12 and 13 on page 30.
4. Add the Host and Host Port definitions for blade3 to the FASiT Storage Manager configuration.
5. On an existing node, e.g., blade1, create the following partitions shown in Table 6. Use `fdisk` as root user.

Database file	Size (MB)	Disk Partition	Raw device	Symbolic Link
redo log thread 3 group 1	120	/dev/sdb10	/dev/raw/raw18	\$ORACLE_HOME/oradata/RACDB/redo3_1
redo log thread 3 group 2	120	/dev/sdb11	/dev/raw/raw19	\$ORACLE_HOME/oradata/RACDB/redo3_2
undo tablespace thread 3	200	/dev/sdb12	/dev/raw/raw20	\$ORACLE_HOME/oradata/RACDB/undo3

Table 6. Raw devices for blade3

6. Update `/etc/sysconfig/rawdevices`. Add the raw devices in Table 6. Do this on all existing nodes (e.g. blade1 and blade2).
7. Create symbolic links in Table 6. Do this on all the existing nodes with the commands in Example 15:

```

ln -s /dev/raw/raw18 $ORACLE_HOME/oradata/RACDB/redo3_1
ln -s /dev/raw/raw19 $ORACLE_HOME/oradata/RACDB/redo3_2
ln -s /dev/raw/raw20 $ORACLE_HOME/oradata/RACDB/undo3
    
```

Example 15. Additional symbolic links for third node

8. Edit `$ORACLE_HOME/oracm/admin/cmcfg.ora` on all existing nodes. Add blade3 and iblade3 to the cluster definitions as shown in Example 16:

```

HeartBeat=15000
ClusterName=Oracle Cluster Manager, version 9i
KernelModuleName=hangcheck-timer
MissCount=210
PollInterval=1000
MissCount=20
PrivateNodeNames=iblade1 iblade2 iblade3
PublicNodeNames=blade1 blade2 blade3
ServicePort=9998
    
```

```
CmDiskFile=/oracle/oradata/quorum
HostName=ibladel
```

Example 16. Add blade3 in \$ORACLE_HOME/oracm/admin/cmcfgr.ora

9. Edit \$ORACLE_HOME/network/admin/tnsnames.ora on all existing nodes. Add the service name entries shown in bold type in Example 17 below, using the entries for blade1 or blade2 as templates.

```
# TNSNAMES.ORA.BLADE1 Network Configuration File:
/oracle/network/admin/tnsnames.ora.bladel
# Generated by Oracle configuration tools.

LISTENER_RACDB3 =
  (ADDRESS = (PROTOCOL = TCP)(HOST = blade3)(PORT = 1521))

LISTENER_RACDB2 =
  (ADDRESS = (PROTOCOL = TCP)(HOST = blade2)(PORT = 1521))

LISTENER_RACDB1 =
  (ADDRESS = (PROTOCOL = TCP)(HOST = blade1)(PORT = 1521))

RACDB =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP)(HOST = blade1)(PORT = 1521))
      (ADDRESS = (PROTOCOL = TCP)(HOST = blade2)(PORT = 1521))
      (ADDRESS = (PROTOCOL = TCP)(HOST = blade3)(PORT = 1521))
      (LOAD_BALANCE = yes)
    )
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = RACDB)
    )
  )

INST1_HTTP =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP)(HOST = blade1)(PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVER = SHARED)
      (SERVICE_NAME = MODOSE)
      (PRESENTATION = http://HRService)
    )
  )

LISTENERS_RACDB =
  (ADDRESS_LIST =
    (ADDRESS = (PROTOCOL = TCP)(HOST = blade1)(PORT = 1521))
    (ADDRESS = (PROTOCOL = TCP)(HOST = blade2)(PORT = 1521))
    (ADDRESS = (PROTOCOL = TCP)(HOST = blade3)(PORT = 1521))
  )

RACDB3 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = blade3)(PORT = 1521))
    (CONNECT_DATA =
      (SERVICE_NAME = RACDB)
      (INSTANCE_NAME = RACDB3)
    )
  )

RACDB2 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = blade2)(PORT = 1521))
    (CONNECT_DATA =
```

```

        (SERVICE_NAME = RACDB)
        (INSTANCE_NAME = RACDB2)
    )
)
RACDB1 =
(DESCRIPTION =
  (ADDRESS = (PROTOCOL = TCP)(HOST = blade1)(PORT = 1521))
  (CONNECT_DATA =
    (SERVICE_NAME = RACDB)
    (INSTANCE_NAME = RACDB1)
  )
)
)

```

Example 17. tnsnames.ora

10. Back up the following directories and files on an existing node:

- \$ORACLE_HOME,
- /var/opt/oracle,
- /opt/ORCLfmap
- /etc/oratab
- /etc/sysconfig/rawdevices
- /etc/sysctl.conf
- /etc/rc.local
- /home/oracle/.bash_profile

As root user, use the tar command as follows:

```

tar -czvf oracle_rac.tarz $ORACLE_HOME /var/opt/oracle /opt/ORCLfmap
/etc/oratab /etc/sysconfig/rawdevices /etc/sysctl.conf /etc/rc.local
/home/oracle/.bash_profile

```

In our example, we backed up the files from blade1, so the subsequent steps in this procedure will refer to files and parameters for blade1.

11. Restore the backed up files to blade3. Transfer the tar file to blade3 and, as root user, unpack it as follows:

```

cd /
tar -tzvf <directory/>oracle_rac.tarz

```

12. Modify the restored files for blade3. Perform the following steps as oracle user.

- a. Edit oracle user's .bash_profile and change the value of ORACLE_SID to RACDB3
- b. In the \$ORACLE_HOME/dbs directory, rename orapwRACDB1 to orapwRACDB3, initRACDB1.ora to initRACDB3.ora -

```

cd /oracle/dbs
mv orapwRACDB1 orapwRACDB3
mv initRACDB1.ora initRACDB3.ora

```

- c. Modify Cluster Manager configuration. Edit \$ORACLE_HOME/oracm/admin/cmcfg.ora - change the value of HostName to iblade3.
- d. Edit \$ORACLE_HOME/network/admin/listener.ora. Change all instances of blade1 to blade3
- e. Reboot blade3.

- f. After blade3 reboots, log in as oracle user and verify the following:
 - i) Raw devices bindings have been created. Switch to root user and issue the `raw -qa` command to display the bindings.
 - ii) Cluster Manager has started automatically. As oracle user, use the `ps -ef |grep oracm` to display the Cluster Manager processes. If no oracm processes are started, check `$ORACLE_HOME/oracm/cm.log` for any error messages.

13. Check the value of MAXINSTANCES in the control file.

- a. Start the database instance only on blade1. Do this as oracle user, with the `srvctl` command, shown below:

```
[oracle@blade1 oracle]$ srvctl status database -d RACDB
Instance RACDB1 is not running on node blade1
Instance RACDB2 is not running on node blade2
[oracle@blade1 oracle]$ srvctl start instance -d RACDB -i RACDB1
[oracle@blade1 oracle]$ srvctl status database -d RACDB
Instance RACDB1 is running on node blade1
Instance RACDB2 is not running on node blade2
[oracle@blade1 oracle]$
```

- b. Connect to the database as a user with DBA privileges and back up the control file. As oracle user, issue the commands shown below:

```
[oracle@blade1 oracle]$ sqlplus /nolog

SQL*Plus: Release 9.2.0.4.0 - Production on Thu Jan 1 15:48:32 2004
Copyright (c) 1982, 2002, Oracle Corporation. All rights reserved.

SQL> connect / as sysdba
Connected.
SQL> alter database backup controlfile to trace;

Database altered.

SQL>
```

The trace file will be placed in the `$ORACLE_HOME/admin/<database name>/udump` directory. Use the `ls -ltr` command to find the most recently created file.

- c. Edit the trace file. Check the MAXINSTANCES value. We found that our dbca-created test database set MAXINSTANCES=32, more than enough for three instances. If MAXINSTANCES had been less than 3, the control file would have to be recreated. Refer to MetaLink Docid 118931.1 for details on how to do this.

14. Modify the Oracle initialization file.

- a. Since the dbca-created database used an spfile, we will create a pfile, which we will edit. Do this as follows:

```
SQL> create pfile='/oracle/dbs/pfileRACDB.ora' from
  2  spfile='/oracle/oradata/RACDB/spfile';

File created.
```

```
SQL> exit
Disconnected from Oracle9i Enterprise Edition Release 9.2.0.4.0 -
Production
With the Partitioning and Real Application Clusters options
JServer Release 9.2.0.4.0 - Production
[oracle@bladel1 oracle]$
```

- b. Edit the pfile created in the previous step, /oracle/dbs/pfileRACDB.ora. Change the line,

```
*.cluster_database_instances = 2

to

*.cluster_database_instances = 3
```

Add the following lines:

```
RACDB3.instance_name='RACDB3'
RACDB3.instance_number=3
RACDB3.local_listener='LISTENER_RACDB3'
RACDB3.thread=3
RACDB3.undo_tablespace='UNDOTBS3'
```

Example 18. RACDB3 statements added to initialization file

- c. Create the spfile,

```
[oracle@bladel1 oracle]$ sqlplus /nolog

SQL*Plus: Release 9.2.0.4.0 - Production on Thu Jan 1 16:04:47 2004

Copyright (c) 1982, 2002, Oracle Corporation. All rights reserved.

SQL> connect / as sysdba
Connected.
SQL> shutdown;
Database closed.
Database dismounted.
ORACLE instance shut down.
SQL> create spfile='/oracle/oradata/RACDB/spfile' from
pfile='/oracle/dbs/pfileRACDB.ora';

File created.

SQL> startup
ORACLE instance started.

Total System Global Area 139531744 bytes
Fixed Size 452064 bytes
Variable Size 113246208 bytes
Database Buffers 25165824 bytes
Redo Buffers 667648 bytes
Database mounted.
Database opened.
SQL>
```

15. Add new redo logs, thread, undo log for blade3.

```
SQL> alter database
2 add logfile thread 3
3 group 5 '/oracle/oradata/RACDB/redo3_1' size 100M,
```

```

4 group 6 '/oracle/oradata/RACDB/redo3_2' size 100M
5

SQL> alter database
2 add logfile thread 3
3 group 5 '/oracle/oradata/RACDB/redo3_1' size 100M,
4 group 6 '/oracle/oradata/RACDB/redo3_2' size 100M;

Database altered.

SQL> alter database enable public thread 3;

Database altered.

SQL> create undo tablespace undotbs3 datafile
2 '/oracle/oradata/RACDB/undotbs03' size 200M;

Tablespace created.

SQL> exit
Disconnected from Oracle9i Enterprise Edition Release 9.2.0.4.0 - Production
With the Partitioning and Real Application Clusters options
JServer Release 9.2.0.4.0 - Production
[oracle@bladel1 oracle]$

```

16. Add the new instance to the Server Management repository (srvm):

```

[oracle@bladel1 oracle]$ srvctl status database -d RACDB
Instance RACDB1 is running on node bladel1
Instance RACDB2 is not running on node blade2
[oracle@bladel1 oracle]$ srvctl add instance -d RACDB -i RACDB3 -n blade3
[oracle@bladel1 oracle]$ srvctl start instance -d RACDB -i RACDB3
[oracle@bladel1 oracle]$ srvctl status database -d RACDB
Instance RACDB1 is running on node bladel1
Instance RACDB2 is not running on node blade2
Instance RACDB3 is running on node blade3
[oracle@bladel1 oracle]$ srvctl start instance -d RACDB -I RACDB2
[oracle@bladel1 oracle]$ srvctl status database -d RACDB
Instance RACDB1 is running on node bladel1
Instance RACDB2 is running on node blade2
Instance RACDB3 is running on node blade3

```

17. Examine \$ORACLE_HOME/admin/RACDB/bdump/alert_RACDBn.log on all nodes to see that the threads are successfully started.

SUMMARY

This concludes our installation test of Oracle9i RAC and Red Hat Enterprise Linux AS 2.1 on the IBM BladeCenter.

For questions and comments, please send an email note to the IBM/Oracle International Competency Center at ibmoracl@us.ibm.com.

APPENDIX A

Database file	Size (MB)	Disk Partition	Raw device	Symbolic Link
quorum	110	/dev/sda1	/dev/raw/raw1	\$ORACLE_HOME/oradata/quorum
srvm	110	/dev/sda2	/dev/raw/raw2	\$ORACLE_HOME/oradata/srvm
spfile	5	/dev/sda3	/dev/raw/raw3	\$ORACLE_HOME/oradata/RACDB/spfile
control file 1	110	/dev/sda5	/dev/raw/raw4	\$ORACLE_HOME/oradata/RACDB/control01
control file 2	110	/dev/sda6	/dev/raw/raw5	\$ORACLE_HOME/oradata/RACDB/control02
control file 3	110	/dev/sda7	/dev/raw/raw6	\$ORACLE_HOME/oradata/RACDB/control03
system	250	/dev/sda8	/dev/raw/raw7	\$ORACLE_HOME/oradata/RACDB/system
temp	40	/dev/sda9	/dev/raw/raw8	\$ORACLE_HOME/oradata/RACDB/temp
tools	10	/dev/sda10	/dev/raw/raw9	\$ORACLE_HOME/oradata/RACDB/tools
users	25	/dev/sdb1	/dev/raw/raw10	\$ORACLE_HOME/oradata/RACDB/users
indx	25	/dev/sdb2	/dev/raw/raw11	\$ORACLE_HOME/oradata/RACDB/index
redo log thread 1 group 1	120	/dev/sdb3	/dev/raw/raw12	\$ORACLE_HOME/oradata/RACDB/redo1_1
redo log thread 1 group 2	120	/dev/sdb5	/dev/raw/raw13	\$ORACLE_HOME/oradata/RACDB/redo1_2
redo log thread 2 group 1	120	/dev/sdb6	/dev/raw/raw14	\$ORACLE_HOME/oradata/RACDB/redo2_1
redo log thread 2 group 2	120	/dev/sdb7	/dev/raw/raw15	\$ORACLE_HOME/oradata/RACDB/redo2_2
undo tablespace thread 1	200	/dev/sdb8	/dev/raw/raw16	\$ORACLE_HOME/oradata/RACDB/undo1
undo tablespace thread 2	200	/dev/sdb9	/dev/raw/raw17	\$ORACLE_HOME/oradata/RACDB/undo2
redo log thread 3 group 1	120	/dev/sdb10	/dev/raw/raw18	\$ORACLE_HOME/oradata/RACDB/redo3_1
redo log thread 3 group 2	120	/dev/sdb11	/dev/raw/raw19	\$ORACLE_HOME/oradata/RACDB/redo3_2
undo tablespace thread 3	200	/dev/sdb12	/dev/raw/raw20	\$ORACLE_HOME/oradata/RACDB/undo3

Figure 91. Consolidated table - raw devices for our three-node RAC cluster

APPENDIX B

ADAPTER PORT NAME IN FAST!UTIL

1. Reboot the BladeCenter HS20. When you see the bootup message "Press <CTRL-Q> for Fast!UTIL", press Ctrl-q key sequence. The Fast!UTIL utility will be loaded, shown in Figure 92. Select the adapter at the appropriate I/O Address and press Enter.




Figure 92. Fast!UTIL Host Adapter Selection

2. The screen shown in Figure 93 is displayed. Select Configuration Settings and press Enter.




Figure 93. Fast!UTIL Options

- The Configuration Settings screen is displayed, shown in Figure 94. Select Host Adapter Settings and press Enter.




Figure 94. Fast!UTIL Configuration Settings

- The Host Adapter Settings screen, shown in Figure 95, is displayed. The 16-digit hexadecimal number labeled Adapter Port Name is the Host Port Identifier in Storage Manager. Obtain this number for each BladeCenter HS20 that will access the logical drives on the FASTT.




Figure 95. Fast!UTIL Host Adapter Settings

REFERENCES

IBM

1. IBM eServer BladeCenter HS20 (Type 8678), Installing Red Hat Enterprise Linux AS version 2.1, available at <http://www-1.ibm.com/support/docview.wss?uid=psg1MIGR-53347>

BladeCenter Documentation

2. IBM ^ BladeCenter Type 8677 Planning and Installation Guide, GA27-4327-01. Available at <http://www.ibm.com/support>. Click **Search technical support** and search for "8677 planning and installation".
3. IBM ^ BladeCenter HS20 Type 8678 Installation and User's Guide. Provided with the BladeCenter. Also available at <http://www.ibm.com/support>. Click **Search technical support** and search for "8678 installation".
4. IBM ^ BladeCenter Products FAQ, Hints and Tips, 02R3058, December 2002. Available at: <http://www-1.ibm.com/support/docview.wss?uid=psg1MIGR-45277>
5. IBM ^ BladeCenter 2-Port Fibre Channel Switch Module, HS20 Expansion Card Hardware Maintenance Manual and Troubleshooting Guide, Part Number 73P8584, February, 2003.
6. IBM ^ BladeCenter Optical Pass-thru Module Installation Guide, Part Number 59P6567, August 2003.

TotalStorage Documentation

7. IBM TotalStorage FASSt900 Fibre Channel Storage Server Installation Guide, GC26-7530-00, March 2003. Available at: <http://www.ibm.com/support>, enter MIGR-50214 in the search field.
8. IBM TotalStorage FASSt900 Fibre Channel Storage Server User's Guide, GC26-7534-00, March 2003. Available at: <http://www.ibm.com/support>, enter MIGR-50216 in the search field.
9. IBM BladeCenter ^ Fibre Channel Switch Interoperability Guide Version 2.0, September 2003. Available at: <http://www.ibm.com/support>, enter MIGR-46241 in the search field.
10. IBM TotalStorage FASSt EXP700 Storage Expansion Unit Installation and User's Guide, Part Number: 32P0178, July 2002. Available at: <http://www.ibm.com/support>, enter MIGR-44119 in the search field.
11. IBM TotalStorage FASSt Storage Manager Version 8.3, Installation and Support Guide for Linux, Part Number GC26-7519-01, April 2003.
12. IBM FASSt MSJ User's Guide, Part Number: 32P0081, August 2001.

Redpieces at www.redbooks.ibm.com

13. IBM Redbook, Implementing Linux with IBM Disk Storage, Dufrasne Bertrand, Ronald Annuss, James Goodwin, Paul McWatt, Arwed Tschoeke, June, 2003, SG24-6261-01.
14. IBM Redpaper, Implementing IBM ^ xSeries SANs, Steve Russell, Derek Cowan, Andrew Le Crerar, January 2002, REDP-0416-00.
15. IBM Redpaper, The Cutting Edge: IBM ^ BladeCenter, Rufus Credle, David Brown, Lowell Davis, David Robertson, Thomas Ternau, November 2002, REDP3581.
16. IBM Redpaper, IBM ^ BladeCenter Systems Management, Rufus Credle, David Brown, Lowell Davis, David Robertson, Thomas Ternau, November 2002, REDP3582.

Oracle

Available from Oracle at <http://otn.oracle.com>

17. Oracle9i Release Notes Release 2 (9.2.0.1.0) for Linux Intel (32-bit), June 2002, Part No. A97349-02.
18. Oracle9i Installation Guide, Release 2 (9.2.0.1.0) for UNIX Systems: AIX-Based Systems, Compaq Tru64 UNIX, HP 9000 Series HP-UX, Linux Intel and Sun Solaris, May 2002, Part No. A96167-01.
19. Oracle9i Real Application Clusters Setup and Configuration Release 2 (9.2), March 2002, Part No. A96600-01.

20. Oracle9i Real Application Clusters Administration Release 2 (9.2), March 2002, Part No. A96596-01. Available from Oracle MetaLink at <http://metalink.oracle.com>
21. Oracle *MetaLink* Note: 201370.1, Quick Start Guide, Oracle9i Release 2 (9.2.0) RDBMS Installation, Linux Operating System.
22. Oracle *MetaLink* Note: 184821.1, Step-By-Step Installation of 9.2.0.4 RAC on Linux.
23. Oracle *MetaLink* Note: 203226.1, RAC Survival Kit: Real Application Clusters Troubleshooting and Information.
24. Oracle *MetaLink* Note: 210022.1, How to add a new instance to the existing two nodes RAC database manually.
25. Oracle *MetaLink* Note: 136327.1, The server parameter file in a Real Application Clusters environment.
26. Oracle *MetaLink* Note: 118931.1, Recreating the Controlfile in RAC and OPS.

Red Hat

Available from Red Hat at <http://www.redhat.com/docs/manuals/enterprise/>

27. Red Hat Enterprise Linux AS 2.1, The Official Red Hat Enterprise Linux AS Installation Guide, Red Hat, Incorporated.