



Technical report:
**Oracle 9i/RAC installation in
a Linux environment with
an IBM System Storage N series with
SnapRestore**

Best practices for data protection

• • • • • • • • •

Document NS3189-0

July 30, 2007



Table of contents

Abstract	1
Introduction	1
Assumptions	1
The server/filer environment	2
Requirements	3
Minimum hardware requirements	3
Software requirements.....	3
Preinstallation setup tasks (for IBM N series Storage System)	4
Preinstallation setup tasks (for cluster nodes)	5
Installation procedure	7
Known issues and workarounds	13
Best practices and recommendations	14
Appendix	15
Trademarks and special notices	16



Abstract

The certified configuration of Oracle 9i Real Application Cluster on a Red Hat Advanced Server with Linux operating system benefits from the high value and performance of an IBM System Storage N series storage system. This white paper outlines installation and best practices.

Introduction

This white paper covers the installation of an IBM® System Storage™ N series system in an Network Attached Storage (NAS) environment and Intel® architecture-based enterprise-class servers using Red Hat Advanced Server 2.1, Linux® operating system running Oracle 9i™ Real Application Cluster, Release 2 (hereafter referred to as Oracle 9i RAC). SnapRestore® recovery software is included as a feature option of the IBM N series system.

This is a certified configuration, and, as such, the components presented in this paper have to be used in the same combination to gain support from all parties involved. The only exception to this is the application of certain patches (as defined and required by all the vendors in this configuration). This document will also cover the best practices and recommendations, once installed, for running Oracle 9i RAC on IBM N series filers in an NAS environment.

Assumptions

This white paper assumes the reader is familiar with Oracle 9i Release 2 and the operation of IBM N series filers. It also assumes familiarity with the operation of the Linux OS and installation of Oracle® patches and any relevant Linux patches. It is also important to be familiar with all networking terminology and implementations.

The server/filer environment

The configuration presented in this document is based on an Oracle 9i RAC certification environment.

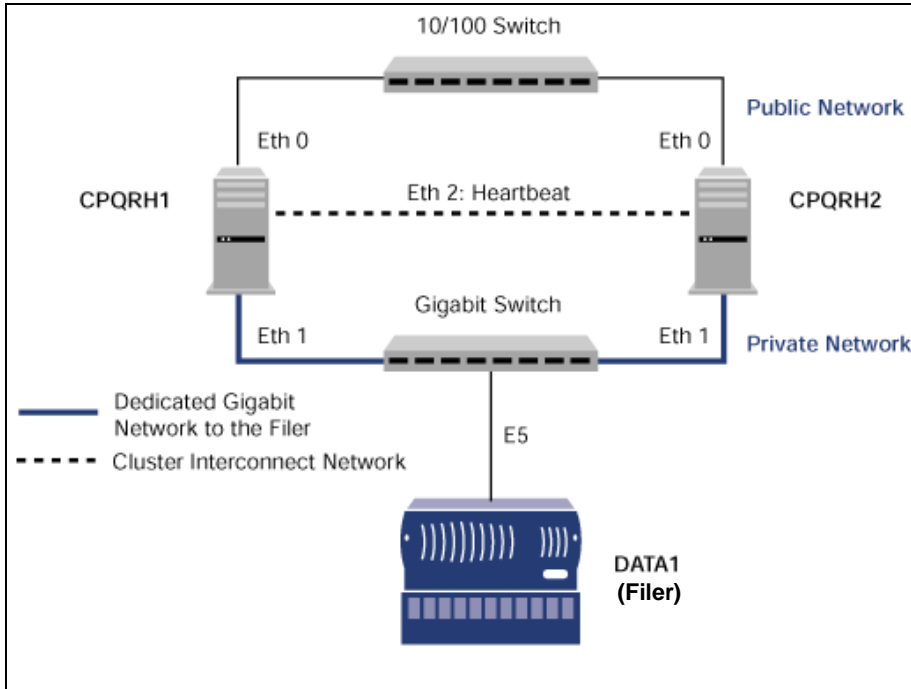


Figure 1. Oracle 9i RAC on Red Hat Advanced Server with an IBM N series filer.

Figure 1 illustrates a typical configuration of Oracle 9i RAC with an IBM N series filer and Intel architecture-based enterprise-class servers running Red Hat Advanced Server 2.1. This is a scalable configuration and allows the user to scale horizontally and internally in terms of processor, memory, and storage.

As shown in the network diagram, we recommend that you dedicate a private network connection between the Oracle 9i RAC servers and the filer. This is accomplished using a dedicated Gigabit network (with a Gigabit switch) to the filer. A dedicated network connection is beneficial for the following reasons:

- In an Oracle 9i RAC environment, it is important to eliminate any contentions and latencies.
- Providing a separate network ensures security.

The cluster interconnect is used to monitor the heartbeat of the two servers in the cluster. This is a typical configuration that can be deployed in a customer's environment.



Requirements

Minimum hardware requirements

Cluster Nodes

- Two 2-way 900 MHz, 2GB RAM Intel architecture–based enterprise-class servers
- Two Intel Pro1000 Gigabit Ethernet NICs (for cluster interconnect)
- Two Intel Pro1000 Gigabit Ethernet NICs (for the IBM N series filer I/O)

Storage Infrastructure

- One IBM N series filer and IBM System Storage N series with Data ONTAP™ 7.1 or later with SnapRestore feature
- One Gigabit switch with at least four ports
- One Gigabit NIC in the filer
- One or more disk shelves, based on the disk space requirements

Software requirements

(For both nodes in the participating cluster unless specified otherwise)

- Red Hat Advanced Server 2.1, kernel 2.4.9-e.34 or later
- Oracle 9i Release 2 (9.2.0.1), with Real Application Cluster license
- Oracle 9i RAC 9.2.0.4 patch (patch: 3095277)



Preinstallation setup tasks (for IBM N series Storage System)

1. Install the N series and connect the disk shelves to it.
2. Make sure each shelf has a unique ID. This should be on the back of the shelf.
3. Connect the shelves to the Fibre Channel adapter in a daisy chain. The output of one shelf should be connected to the input of the other. The last shelf's output should be terminated. The input to the first shelf should come from the Fibre Channel adapter in the filer.
4. Install Data ONTAP 7.1 or later and configure the filer by giving it host names such as `Data1`.
5. Install the Network File System (NFS) and IBM System Storage N series with SnapRestore license keys.
6. Configure the Gigabit network interface card (NIC) in the filer.
 - a. Ensure that the Gigabit NIC is installed in the filer. This document assumes the NIC is inserted in slot 5 in the filer `Data1`.
 - b. Configure the Gigabit interfaces: e.g., `e5` on `Data1`.
 - c. Specify the following IP address for the interface (Run the "setup" command on the filer console to configure the network interfaces.):

```
Data1>e5 - ip: 10.1.1.103, netmask: 255.255.255.0
```
7. Create and export volumes for storing Oracle Database files on the filers:
 - d. Create a volume on the `Data1` filer for storing Oracle Database files. This volume will be used to store all the data files, control files, and log files for the Oracle Database. Use the following commands at the filer console:

```
Data1> vol create oradata 14
```
 - e. Edit the `/etc/exports` file on `Data1` and add the following entries to that file:

```
/vol/oradata -anon=0
```
 - f. Execute the following command at the filer console:

```
Data1> exportfs -a
```



Preinstallation setup tasks (for cluster nodes)

This information is intended for both nodes in the participating cluster unless specified otherwise.

1. Install the Pro1000 Gigabit network interface cards (NICs) in the cluster nodes (two per node).
 - a. Connect one Gigabit NIC back-to-back to the other node for cluster interconnects.
 - b. Connect one Gigabit NIC to the Gigabit switch, which will connect to the Gigabit NIC on the filer.
2. Install Red Hat Advanced Server 2.1 on the cluster nodes. (The cluster nodes in this document will be referred to as `cpqrh1` and `cpqrh2`.) Refer to Red Hat Advanced Server installation manuals for more details. Install the latest errata kernel, e34 or a later kernel from Red Hat.
3. Configure the two Gigabit network interfaces on `cpqrh1` and `cpqrh2` using Red Hat Advanced Server 2.1 network management tools.

- a. Configure the two network interfaces and specify the following IP addresses:

```
cpqrh1
eth1 - ip: 10.1.1.101, netmask: 255.255.255.0
eth2 - ip: 20.1.1.101, netmask: 255.255.255.0
cpqrh2
eth1 - ip: 10.1.1.102, netmask: 255.255.255.0
eth2 - ip: 20.1.1.102, netmask: 255.255.255.0
```

Where:

Interface `eth1` on both cluster nodes is connected to the Gigabit switch for filer I/O. Interface `eth2` on both cluster nodes is connected back-to-back for cluster interconnects.

- b. Update the `/etc/hosts` file on the cluster nodes and add the following entries:

```
cpqrh1
cpqrh1-eth1 10.1.1.101
cpqrh1-eth2 20.1.1.101
data1 10.1.1.103
cpqrh2-eth1 10.1.1.102
cpqrh2-eth2 20.1.1.102

cpqrh2
cpqrh1-eth1 10.1.1.101
cpqrh1-eth2 20.1.1.101
data1 10.1.1.103
cpqrh2-eth1 10.1.1.102
cpqrh2-eth2 20.1.1.102
```

4. Make sure the `eth1` interface is connected to the switch and can communicate with the filer Data1. Use the `/usr/sbin/ping` command to verify the connectivity to the filer over the Gigabit network from both the cluster nodes, i.e., `cpqrh1` and `cpqrh2`.

```
$/usr/sbin/ping data1
```



5. The other interface, `eth2` on `cpqrh1`, should be able to communicate with its counterpart, `eth2` on `cpqrh2`. Use the `/usr/sbin/ping` command to verify the cluster interconnect from both the cluster nodes, i.e., `cpqrh1` and `cpqrh2`.

```
$/usr/sbin/ping cpqrh1-eth2
```

```
$/usr/sbin/ping cpqrh2-eth2
```

`cpqrh1-eth2` and `cpqrh2-eth2` will be used as private node names for the cluster interconnects. These interface names will be used for the `HostName` and `PrivateNodeNames` parameters in the cluster configuration file, `cmcfg.ora`, for the Oracle 9i RAC.

6. Create mount points and mount volumes.
 - a. Update the `/etc/fstab` file on all the server nodes and add the following entries:

```
data1:/vol/oradata - /oradb nfs - yes  
rw,bg,hard,nointr,rsize=32768,wsiz=32768,tcp,noac,vers=3,timeo=600
```

Where:

`Data1` is the name of the filer.

`oradb` is the mount point on the cluster nodes.

The mount options that are required for the Oracle 9i RAC are:

`noac`: This mount option disables caching on the client side.

`tcp`: Mount the file system using the `tcp` option.

- b. Create the following mount point:

```
#mkdir /oradb
```
 - c. Add the following entry to the `/etc/modules.conf` file on all the cluster nodes. Reboot the cluster nodes after making the modification.

```
options nfs nfs_uncached_io=1
```

- d. Mount the exported volume on the mount point created above on the cluster nodes.

```
#mount /oradb
```
 - e. Check the `/var/log/messages` file to verify if unca



Installation procedure

For more information about installing Oracle 9i Release 2 on Linux, refer to Oracle 9i Installation Guide Release 2 (9.2.0.1) for UNIX[®] Systems [Part No:A96167-01] at <http://otn.oracle.com/docs/content.html>.

1. Preparing to install the Oracle 9i RAC on cluster nodes

- a. This document assumes the Oracle user account and the group to be "oracle" and "dba," respectively, on both cluster nodes, cpqrh1 and cpqrh2. The user ID and group name for the "oracle" account should be same on both cluster nodes. A sample "oracle" user .profile file is provided in the appendix.
- b. The \$ORACLE_HOME directory is not shared across the cluster nodes. The \$ORACLE_HOME directory is local to the cluster nodes. Only the \$ORACLE_HOME/dbs directory is shared on the cluster nodes.

- c. Grant appropriate permissions to the "oracle" user on the shared mounted volume, /oradb.

```
#chown -R oracle:dba /oradb
#chmod -R 755 /oradb
```

- d. Set up account equivalence between the cluster nodes for the "oracle" user account. Add the following entries to the /etc/hosts.equiv file on both cluster nodes.

```
cpqrh1 oracle
cpqrh2 oracle
cpqrh1-eth2 oracle
cpqrh2-eth2 oracle
```

- e. Test the "oracle" account equivalence using a remote shell utility such as rsh after logging in as "oracle" user from both cluster nodes.

```
cpqrh1:
#su - oracle
$rsh cpqrh1 pwd
$rsh cpqrh2 pwd
$rsh cpqrh1-eth2 pwd
$rsh cpqrh2-eth2 pwd
```

```
cpqrh2:
#su - oracle
$rsh cpqrh1 pwd
$rsh cpqrh2 pwd
$rsh cpqrh1-eth2 pwd
$rsh cpqrh2-eth2 pwd
```



- f. Create the following two shared files on `/oradb`. The "oracle" user on both cluster nodes should have read/write permissions on these files.

```
#su - oracle
$touch /oradb/SharedConfigFile
$dd if=/dev/zero of=/oradb/CmDiskFile bs=10k count=1000
```

Where:

`/oradb/SharedConfigFile` will be used by cluster management utilities such as `srvctl` and `oem` to store the cluster configuration and `/oradb/CmDiskFile` will be used by Oracle Cluster Manager to store the quorum disk information.

- g. Verify if the `/var/opt/oracle` directory is present on the cluster nodes. If not, create the `/var/opt/oracle` directory on the cluster nodes and grant read access to the "oracle" user on this directory.

2. Installing the Oracle 9i RAC Release 2, 9.2.0.1 Cluster Manager

In Oracle 9i Release 2, the Oracle Cluster Manager has to be installed, configured, and started before you install the Oracle 9i RAC. The Oracle 9i Release 2 Cluster Manager is packaged along with the Oracle 9i Release 2 software distribution. Refer to the Oracle 9i Installation Guide Release 2 (9.2.0.1) for UNIX Systems [Part No:A96167-01] at <http://otn.oracle.com/docs/content.html> for more information on installing Oracle Cluster Manager on Linux. This section highlights the specifics when integrating with an IBM N series filer.

- a. Start the Oracle 9i Release 2 installer from either of the nodes. This document assumes the Oracle installer is started from `cpqrh1`.
- b. From the available products window, select Oracle Cluster Manager 9.2.0.1.0.
- c. When prompted to enter the public node names, enter the names of nodes accessible from the public network, in this case, `cpqrh1` and `cpqrh2`.
- d. For the private node names enter the names of the nodes accessible over the private network between the cluster nodes or the cluster interconnect, in this case, `cpqrh1-eth2` and `cpqrh2-eth2`.
- e. Accept the default value for the watchdog parameter value.
- f. When prompted to enter a quorum disk filename, enter the full path to the quorum disk file created above, i.e., `/oradb/CmDiskFile`.
- g. The Oracle installer will install all the required files for Oracle Cluster Manager on both the nodes.

3. Configuring and starting Oracle Cluster Manager

After installing Oracle Cluster Manager, you have to configure and start it before you can install the Oracle 9i RAC. Refer to the Oracle 9i Administrator's Reference Release 2 (9.2.0.1) for UNIX Systems [Part No:A97297-01], Appendix F, for more information on configuring Oracle Cluster Manager on Linux.

- a. Make the following changes to the `$ORACLE_HOME/oracm/admin/ocmargs.ora` on `cpqrh1` and `cpqrh2`.

```
watchdogd -g dba -d /dev/null -l 0
oracm /a:0
```

- b. Rename the `$ORACLE_HOME/oracm/admin/cmcfg.ora.tmp` file on `cpqrh1` to

```
$ORACLE_HOME/oracm/admin/cmcfg.ora
cpqrh1
$mv $ORACLE_HOME/oracm/admin/cmcfg.ora.tmp
$ORACLE_HOME/oracm/admin/cmcfg.ora
```

- c. Copy the `cmcfg.ora` file over to the other cluster node. In this case copy it over to `cpqrh2`.

```
$rcp $ORACLE_HOME/oracm/admin/cmcfg.ora cpqrh2:$ORACLE_HOME/oracm/admin
```

- d. Edit the `cmcfg.ora` file on the cluster nodes and make appropriate changes.

```
cpqrh1:
$ORACLE_HOME/oracm/admin/cmcfg.ora
HeartBeat=15000
ClusterName=Oracle Cluster Manager, version 9i
PollInterval=1000
MissCount=20
PrivateNodeNames=cpqrh1-eth2 cpqrh2-eth2
PublicNodeNames=cpqrh1 cpqrh2
ServicePort=9998
WatchdogSafetyMargin=5000
WatchdogTimerMargin=60000
CmDiskFile=/oradb/CmDiskFile
HostName=cpqrh1-eth2
```

```
cpqrh2:
$ORACLE_HOME/oracm/admin/cmcfg.ora
HeartBeat=15000
ClusterName=Oracle Cluster Manager, version 9i
PollInterval=1000
MissCount=20
PrivateNodeNames=cpqrh1-eth2 cpqrh2-eth2
PublicNodeNames=cpqrh1 cpqrh2
ServicePort=9998
```

```
WatchdogSafetyMargin=5000
WatchdogTimerMargin=60000
CmDiskFile=/oradb/CmDiskFile
HostName=cpqrh2-eth2
```

Where:

`/oradb/CmDiskFile` is the quorum disk file accessible from both cluster nodes.

- e. Start the Oracle Cluster Manager by running the following script on `cpqrh1` and `cpqrh2`:

```
#su - oracle
$su
#.profile
#ocmstart.sh
```

- f. Verify that the Oracle Cluster Manager is started by checking the `$ORACLE_HOME/oracm/log/cm.log` file on cluster nodes.

4. Installing the Oracle 9i RAC

- a. After making sure that Oracle Cluster Manager has started on the cluster nodes, restart the Oracle 9i Release 2 installer from `cpqrh1`.
- b. In the cluster node selection you should see all the cluster nodes, in this case, `cpqrh1` and `cpqrh2`. If the nodes are not displayed in the cluster node selection, then Oracle Cluster Manager is not configured or started on those cluster nodes. Configure and restart the Oracle Cluster Manager and rerun the Oracle installer.
- c. Select all the available nodes in the cluster node selection window, i.e., `cpqrh1` and `cpqrh2`.
- d. Select "Oracle 9i Database 9.2.0.1.0 -> Enterprise Edition -> Customized" in the next three Oracle installer windows.
- e. When prompted to enter the shared configuration filename, enter the full path to the shared configuration file created above, i.e., `/oradb/SharedConfigFile`.
- f. After installing the required Oracle 9i Release 2 executables and libraries on the cluster nodes, the Oracle installer will start executing the tools and utilities in the following order:
 - i. Cluster Configuration Assistant (CLUSTCA)
 - ii. Network Configuration Assistant (NETCA)
 - iii. Database Configuration Assistant (DBCA)

Let the Oracle installer complete running and configuring CLUSTCA and NETCA. But when running DBCA, cancel out of the installation. The DBCA will have to be run later.



5. Upgrading to the Oracle 9i RAC 9.2.0.4

At the time of writing this document, the latest Oracle patch set for the Oracle 9i RAC 9.2.0.1 is 9.2.0.4.

- a. Stop all the Oracle processes on the cluster nodes. Stop the Oracle 9i RAC 9.2.0.1 cluster manager, watchdog, and any other Oracle process on the nodes.
- b. Download the Oracle 9.2.0.4 patch set from <http://metalink.oracle.com>. You will need a userid and password to access the metalink Web site. The patch number for the Oracle 9.2.0.4 patch is 3095277.
- c. Follow the README file for instructions on how to apply the Oracle 9.2.0.4 patch. The README file is packaged with the patch set.
- d. After upgrading your current environment to Oracle 9.2.0.4, restart the Oracle Cluster Manager on all nodes. The watchdog module is replaced with the hangcheck timer module in the Oracle 9i RAC 9.2.0.4. This module is included with the Red Hat Advanced Server Linux kernel 2.4.9-e12 and higher.

```
#su - oracle
$su
#.profile
#ocmstart.sh
```

- e. Verify that the Oracle Cluster Manager is started by checking the `$ORACLE_HOME/oracm/log/cm.log` file on cluster nodes.

6. Running DBCA to create an Oracle 9i RAC database

The Oracle 9i Release 2 Database Configuration Assistant (DBCA) can be used to create a cluster-aware database. Refer to Oracle 9i Real Application Clusters Setup and Configuration Guide Release 2 [Part No: A96600-01] for more information. This document will highlight the specifics of storing all the Oracle Database files in a shared directory on the IBM N series filer.

- a. Create a symbolic link to `/oradb` on `cpqrh1` and `cpqrh2`. Execute the following commands:

```
#su - oracle
$mv $ORACLE_HOME/dbs $ORACLE_HOME/dbs.org
$ln -s /oradb $ORACLE_HOME/dbs
```



- b. Before you run the DBCA utility to create the Oracle 9i RAC database, check whether the global services daemon (gsd) and the listener have started on the cluster nodes. Execute the following command to check the status of the gsd on `cpqrh1` and `cpqrh2`:

```
$gsdctl stat
```

If the `gsd` has not started on the cluster nodes, start it before running the DBCA utility by executing the following command on `cpqrh1` and `cpqrh2`:

```
$gsdctl start
```

To check the status of the listener, execute the following command on the cluster nodes:

```
$lsnrctl status LISTENER
```

If the listener has not started, start it by executing the following command on `cpqrh1` and `cpqrh2`:

```
$lsnrctl start LISTENER
```

- c. Run the DBCA utility to create an Oracle 9i RAC database from `cpqrh1`:

```
$dbca -datafileDestination $ORACLE_HOME/dbs/
```

The `datafileDestination` argument specifies the destination directory where all Oracle Database files, such as data files, log files, control files, server parameter files, initialization parameter files, and password files, will be stored. This directory has to be shared by all the cluster nodes. Refer to the Oracle 9i Real Application Clusters Setup and Configuration Guide Release 2 [Part No: A96600-01] for more details on using DBCA or creating an Oracle 9i RAC database manually using scripts.



Known issues and workarounds

1. Cluster Manager installation

Issue:

After the Oracle Cluster Manager installation, the `$ORACLE_HOME/oracm/log` directory is missing from the other cluster nodes (`cpqrh2`). When you try to start the Oracle Cluster Manager on that node, it will fail to start.

Workaround:

Create the `$ORACLE_HOME/oracm/log` directory on the other cluster nodes (`cpqrh2`) before starting the Oracle Cluster Manager.

```
$mkdir $ORACLE_HOME/oracm/log
```

2. Oracle 9i RAC installation

Issue:

After the Oracle 9i RAC installation, the following directories are missing from the other cluster nodes (`cpqrh2`):

```
$ORACLE_HOME/network/log, $ORACLE_HOME/rdbms/log and  
$ORACLE_HOME/rdbms/audit
```

This will prevent you from starting the listener or connecting to any database from that node (`cpqrh2`).

Workaround:

Create the following directories on the other cluster nodes (`cpqrh2`):

```
$mkdir $ORACLE_HOME/network/log $mkdir $ORACLE_HOME/rdbms/audit $mkdir  
$ORACLE_HOME/rdbms/log
```

Best practices and recommendations

1. Using uncached I/O for an Oracle 9i RAC

The uncached I/O patch is introduced in Red Hat Advanced Server 2.1, e34 and up. It is mandatory to use uncached I/O for running an Oracle 9i RAC with IBM N series filers in an NAS environment. Uncached I/O does not cache data in the Linux file system buffer cache during read/write operations for volumes mounted with the “noac” mount option. To enable uncached I/O, you have to add the following entry to the `/etc/modules.conf` file and reboot the cluster nodes:

```
options nfs nfs_uncached_io=1
```

The volumes used for storing Oracle Database files should still be mounted with the “noac” mount option.

2. Using asynch I/O

Currently IBM N series does not support running an Oracle 9i RAC with asynch I/O when storing Oracle Database files on NFS mounted volumes. By default, an Oracle 9i RAC on a Linux installation disables asynch I/O. To use asynch I/O, the Oracle binaries have to be relinked, and the `initORA` parameter, “filesystemio_options,” has to be set to “asynch.” Make sure you DO NOT relink Oracle binaries to enable asynch I/O.

3. Using IBM N series clusters

IBM N series highly recommends running an Oracle 9i RAC on a highly available and scalable IBM N series cluster solution. If you are running IBM N series with Data ONTAP™, it is recommended to upgrade to the latest GA release.

4. Using a different socket buffer size

Enlarging the transport socket buffers your client uses for NFS traffic, helps reduce resource contention on the client, reduces performance variance, and improves maximum data and operation throughput. In Linux kernels after 2.4.20 the following procedure is not necessary, as the client will automatically choose an optimal socket buffer size.

- Become root on your client
- cd into `/proc/sys/net/core`
- `echo 262143 > rmem_max`
- `echo 262143 > wmem_max`
- `echo 262143 > rmem_default`
- `echo 262143 > wmem_default`
- Remount your NFS file systems on the client



Appendix

Sample .profile file for the “oracle” user:

```
ORACLE_BASE=/home/oracle; export ORACLE_BASE

ORACLE_PRODUCT=$ORACLE_BASE/product; export ORACLE_PRODUCT

ORACLE_HOME=$ORACLE_PRODUCT/9.2; export ORACLE_HOME

ORACLE_SID=rac1; export ORACLE_SID

LD_LIBRARY_PATH=$ORACLE_HOME/lib; export LD_LIBRARY_PATH

PATH=$ORACLE_HOME/bin:$ORACLE_HOME/oracm/bin:/usr/bin:/bin:/usr/bin/X11:/usr/local/bin:.
```



Trademarks and special notices

© International Business Machines 1994-2007. IBM, the IBM logo, System Storage, and other referenced IBM products and services are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. All rights reserved

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Network Appliance, the Network Appliance logo, Data ONTAP and SnapRestore are trademarks or registered trademarks of Network Appliance, Inc., in the U.S. and other countries.

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

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

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Any other trademarks, registered trademarks, company, product or service names may be trademarks, registered trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.