



Oracle10g™ Real Application Clusters Release 1 with HP-UX 11iV2 on NetApp Storage Systems

Network Appliance Inc., Tushar Patel, TR 3432

Contents

1. Introduction	3
2. Assumptions	3
3. The Server/Storage system Environment	3
4. Requirements	4
4.1. Hardware Used for Tests	4
Cluster Nodes	4
Storage Infrastructure	4
4.2. Software Used	4
5. Setup for Netapp Storage	5
6. Preinstallation Setup Tasks (Cluster Nodes)	5
7. Known Issues (Patches)	7
8. Installation Procedure	7
9. Appendix.....	8
10. Acknowledgements.....	8
11. Disclaimer	8

1. Introduction

This technical report will cover the installation of a NetApp storage system in a NAS environment for HP-UX IA64 or PA-RISC enterprise-class servers. The servers are running the HP-UX 11iv2 (11.23) operating system with Oracle10g Real Application Clusters (Oracle 10g RAC) release 1 with Oracle Cluster Resource Services (OCRS). This is now a certified configuration, and therefore 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 patches and recommendations for running Oracle10g RAC on NetApp storage systems in a NAS environment.

2. Assumptions

We assume that readers are familiar with Oracle10g release 1 and the operation of Network Appliance™ storage systems. We also assume that readers are familiar with the operation of the HP-UX environment and installation of Oracle® patches and any relevant HP patches. It is also important to be familiar with all networking terminology and implementations.

3. The Server/Storage System Environment

The configuration presented in this document is based on the Oracle10g RAC certification environment specified by Oracle and Network Appliance.

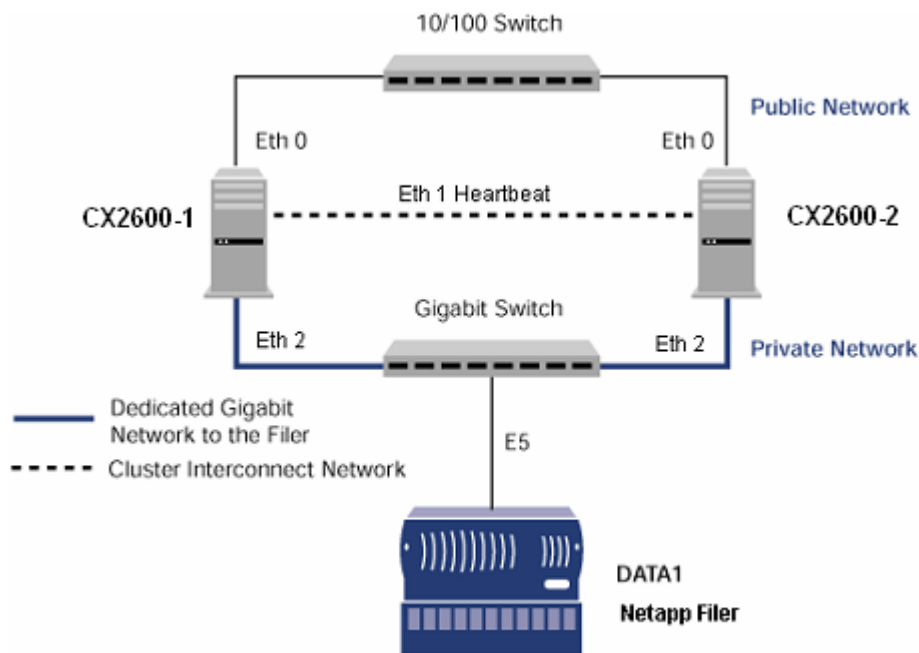


Figure 1) Oracle10g RAC on HP-UX server with a Network Appliance storage system.

Figure 1 illustrates a typical configuration of Oracle10g RAC with a NetApp storage system and HP-UX servers running HP-UX 11iV2. 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 Oracle10g RAC servers and the storage system. This is accomplished using a dedicated Gigabit network (with a Gigabit switch) connected to the storage system. A dedicated network connection is beneficial for the following reasons:

- In an Oracle10g 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.

4. Requirements

4.1. Hardware Used for Tests

Cluster Nodes

- For IA64, two two-way 1.3GHz, 4GB RAM Intel® IA64 enterprise-class servers

- For PA-RISC, two four-way (PA 8500 CPU Modules 2.4 @ 440MHz) HP Server 9000/800/L2000-44 Series servers

- Two Intel Pro1000 Gigabit Ethernet network interface cards (NICs) for cluster interconnect

- Additional NIC on each machine for public network (Eth0)

- Two Intel Pro1000 Gigabit Ethernet NICs (for Network Appliance storage system I/O)

Storage Infrastructure

- One Network Appliance FAS200, FAS700, FAS800, or FAS900 series storage system with Data ONTAP® 7.0 or later

- One gigabit switch with at least four ports.

- One Gigabit Ethernet NIC in the storage system

- One or more disk shelves, based on the disk space requirements

4.2. Software Used

For both nodes in the participating cluster unless specified otherwise:

- HP-UX 11i Version 2 (11.23) with September 2004 patch bundle

- Oracle10g release 1 (10.1.0.2), with Real Application Clusters license

5. Setup for NetApp Storage

For more information, refer to the Network Appliance installation and setup guides at <http://now.netapp.com>.

1. Configure a NetApp storage system running Data ONTAP 7.0 and with NFS and SnapRestore® license keys.
2. Create and export volumes for storing Oracle Database files on the storage system:
 1. Create three volumes on the `Data1` storage system as listed below and mount them on all the cluster nodes:

```
volhome  Shared Oracle and CRS Home
voldb    Shared Oracle datafiles and control files
vollog   db logs, a copy of control files and archive log
files.
```

To create volumes, use the following command at the storage system console:

```
Data1> vol create voldb 14
```

We created a volume `voldb` with 14 disks and volumes `volog` and `volhome` with three disks each. You can create your volumes based on your workload needs.

Edit the `/etc/exports` file on `Data1` and add the following entries to that file:

```
/vol/volhome -anon=0
/vol/voldb -anon=0
/vol/vollog -anon=0
```

2. Execute the following command at the storage system console:

```
Data1> exportfs -a
```

6. Preinstall Setup Tasks (Cluster Nodes)

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

1. Have two HP-UX 11.23 (11iV2) servers ready with latest recommended patches. For this certification, we used the September 2004 patch bundle.
2. Install the Pro1000 Gigabit Ethernet NICs in the cluster nodes (two per node).
 1. Connect one Gigabit Ethernet NIC back-to-back to the other node for cluster interconnects.
 2. Connect one Gigabit Ethernet NIC to the gigabit switch, which will connect to the Gigabit Ethernet NIC on the storage system.

3. Configure the two Gigabit Ethernet network interfaces on each node.

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

```
#cx2600-1
eth2 - ip: 10.1.1.101, netmask: 255.255.255.0
eth1 - ip: 172.68.30.1, netmask: 255.255.255.0
#cx2600-2
eth2 - ip: 10.1.1.102, netmask: 255.255.255.0
eth1 - ip: 172.68.30.2, netmask: 255.255.255.0
```

Where:

Interface `eth2` on both cluster nodes is connected to the gigabit switch for storage system I/O.

Interface `eth1` on both cluster nodes is connected back-to-back for cluster interconnects.

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

```
#cx2600-1
cx2600-1_i 172.68.30.1
cx2600-1_eth2 10.1.1.101
data1 10.1.1.103
cx2600-2_i 172.68.30.2
cx2600-2_eth2 10.1.1.102

#cx2600-2
cx2600-1_i 172.68.30.1
cx2600-1_eth2 10.1.1.101
data1 10.1.1.103
cx2600-2_i 172.68.30.2
cx2600-2_eth2 10.1.1.102
```

4. Make sure the `eth2` interface is connected to the switch and can communicate with the storage system `Data1`. Use the `ping` command to verify the connectivity to the storage system over the Gigabit Ethernet network from both the cluster nodes, i.e., `cx2600-1` and `cx2600-2`.
5. The other interface, `eth1` on `cx2660-1`, should be able to communicate with its counterpart, `eth1` on `cx2600-2`. Use the `ping` command to verify communication among private interconnects.
6. Create mount points and mount the volumes with the following mount options on all the cluster nodes. Update the `/etc/fstab` file on all server nodes and add following entries:

```
Data1:/vol/voldb /voldb nfs
rsize=32768,wsiz=32768,hard,nointr,rw,bg,vers=3,proto=tcp,noac,fo
rccdirectio,timeo=600 0 0
Data1:/vol/vollog /vollog nfs
rsize=32768,wsiz=32768,hard,nointr,rw,bg,vers=3,proto=tcp,noac,fo
rccdirectio,timeo=600 0 0
```

```
Data1:/vol/volhome      /volhome      nfs
rsize=32768,wsiz=32768,hard,nointr,rw,bg,vers=3,proto=tcp,noac,t
imeo=600 0 0
```

Where:

`Data1` is the name of the storage system.

`voldb`, `vollog`, `volhome` and `ora10g` are the mount points on the cluster nodes.

The mount options that are required for the Oracle10g RAC are `noac` (do not cache file attributes), `timeo=600` and `tcp` (mount file system using the `tcp` option).

Note: HP-UX 11iV2 has a patch that fixes several NFS issues, including `forcedirectio` support. That patch must be applied on both HP-UX nodes for the `forcedirectio` mount option to work. Please refer to the “Known Issues (Patches)” section below for details.

7. Create the following mount points on all cluster nodes:

```
#mkdir /voldb
#mkdir /vollog
#mkdir /volhome
```

Mount exported volumes on the mount points created above on the cluster nodes and verify mount options by using the `mount` command.

7. Known Issues (Patches)

Before the Oracle install, we applied the following patches on the HP-UX 11iV2 servers with the September 2004 patch bundle.

For HP-UX IA64, we applied the following patches:

```
PHSS_32502
PHSS_31087
PHNE_32055 (Required for NFS certification)
PHCO_32426 (A kernel panic patch for bug JAGaf61327)
```

For HP-UX PA-RISC, we applied the following patches:

```
PHNE_32055 (Required for NFS certification)
PHSS_32765
```

To download these HP-UX patches, visit the patch/firmware database at <http://www1.itrc.hp.com>.

8. Installation Procedure

For PA-RISC platforms, Oracle must be installed with `ignoreSysPreReqs`. For example,

```
./runInstaller -ignoreSysPreReqs
```

For step-by-step instructions for the Oracle RAC install and the recommended kernel parameter settings for the Oracle environment, go to

http://download-west.oracle.com/docs/html/B10766_08/toc.htm.

9. Appendix

Sample `.profile` file for the Oracle10g user:

```
ORACLE_BASE=/volhome/ora10g; export ORACLE_BASE
ORACLE_PRODUCT=$ORACLE_BASE/product; export ORACLE_PRODUCT
ORACLE_HOME=$ORACLE_PRODUCT/10.1.0/db_1; export ORACLE_HOME
ORACLE_CRS=$ORACLE_PRODUCT/10.1.0/crs_1; export ORACLE_CRS
ORACLE_SID=rac1; export ORACLE_SID
LD_LIBRARY_PATH=$ORACLE_HOME/lib; export LD_LIBRARY_PATH
PATH=$PATH:$ORACLE_HOME/bin:$ORACLE_HOME:$ORACLE_CRS/bin; export PATH
```

10. Acknowledgments

The author thanks the following individuals for their contribution in this technical report:

Network Appliance Inc.

Daniel Morgan, Ricky Stout

11. Disclaimer

Each environment has its own specific set of requirements, and no guarantees can be given that the results presented in this report will work as expected on other platforms. This paper should assist in the research and troubleshooting that may be required in a particular case and serve as a checklist of items to be aware of. Please forward any errors, omissions, differences, new discoveries, and comments about this paper to tusharp@netapp.com.



www.netapp.com

© 2005 Network Appliance, Inc. All rights reserved. Specifications subject to change without notice. NetApp, the Network Appliance logo, DataFabric, FAServer, FilerView, NetCache, NearStore, SecureShare, SnapManager, SnapMirror, SnapRestore, SpinCluster, SpinFS, SpinHA, SpinMove, SpinServer, and WAFL are registered trademarks and Network Appliance, ApplianceWatch, BareMetal, Camera-to-Viewer, Center-to-Edge, ContentDirector, ContentFabric, Data ONTAP, EdgeFiler, HyperSAN, InfoFabric, MultiStore, NetApp Availability Assurance, NetApp ProTech Expert, NOW, NOW NetApp on the Web, RoboCache, RoboFiler, SecureAdmin, Serving Data by Design, Smart SAN, SnapCache, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapMigrator, Snapshot, SnapSuite, SnapVault, SohoCache, SohoFiler, SpinMirror, SpinShot, SpinStor, The evolution of storage, Vfiler, VFM, Virtual File Manager, and Web Filer are trademarks of Network Appliance, Inc. in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.