Technical Report

Setting Up IBM HACMP Cluster on AIX with NetApp Storage (NFS) for High Availability

Suresh Vundru, NetApp
April 2010 | TR-3577

TABLE OF CONTENTS

1 INTRODUCTION

More and more NetApp customers recognize the value of high availability in their mission- and business-critical applications. IBM provides host-clustering software known as High-Availability Cluster Multiprocessing (HACMP), which minimizes downtime by quickly restoring services when a system, component, or application fails.

This technical report documents the installation and configuration of HACMP version 5.4 and 5.5 software in the following environment:

- AIX version 5.3 and 6.1 operating system platform with NetApp® storage
- Data ONTAP® version 7.2.x or 7.3.x
- Network File System (NFS) protocol

2 ASSUMPTIONS

This technical report assumes readers are familiar with the following:

- HACMP software
- NetApp storage systems
- AIX operating system on IBM Systems p-series and i-series servers
- General networking concepts

3 HIGH-AVAILABILITY CLUSTER MULTIPROCESSING FOR AIX

The IBM HACMP software provides a computing environment that enables quick recovery of mission-critical applications from hardware and software failures. HACMP has two major components:

- High availability (HA)
- Cluster multiprocessing (CMP)

The high-availability system combines custom software with industry-standard hardware to minimize downtime by quickly restoring services when a system (or node), component, or application fails. In an HACMP cluster, to make sure of the availability of the applications, the applications are put under HACMP's control. HACMP takes measures to make sure that the applications remain available to client processes even if a component in a cluster fails. To make sure of availability in the event of a component failure, HACMP moves the application (along with resources that enable access to the application) to another node in the cluster.

The primary goal of HACMP is to minimize or eliminate application downtime. HACMP software optimizes availability by allowing for the dynamic reconfiguration of running clusters. Most routine cluster maintenance tasks, such as adding or removing a node or changing the priority of nodes participating in a resource group, can be applied to an active cluster without stopping and restarting cluster services.

In addition, you can keep an HACMP cluster online while making configuration changes by using the Cluster Single Point of Control (C-SPOC) facility. C-SPOC simplifies cluster management by allowing you to make changes to shared volume groups, users, and groups across the cluster from a single node. The changes are then propagated transparently to other cluster nodes.

The following cluster configurations are discussed in this technical report:

- **Standby configurations.** These are the traditional redundant hardware configurations where one or more standby nodes stand idle, waiting to replace the work of a failed server node in the cluster.
- **Takeover configurations.** All nodes in a takeover configuration participate in processing the cluster's workload. There are no standby nodes. Furthermore, each node has takeover capacity. If a node in the cluster fails, a surviving node takes over the resource groups owned by the failed node. Takeover configurations use hardware resources more efficiently than standby configurations since there is no idle processor. However, more load-balance planning needs to take place to enable the takeover hardware to manage the extra workload.

4 SERVICE IP LABEL AND IP ADDRESS TAKEOVER

HACMP uses the IP Address Takeover (IPAT) operation if the physical network interface adapter on one node fails and there are no other accessible physical network interface adapters on the same network on the same node.

IPAT is a mechanism for recovering a service IP label by moving it to another NIC on another node when the initial NIC fails. IPAT makes sure of the availability of a service IP label that provides services to the client nodes. A service IP label is a label that maps to the service IP address and is used to establish communication between client nodes and the server node.

HACMP supports two methods for performing IPAT:

- IPAT using IP Aliases (default)
- IPAT using IP Replacement (known in previous releases as IPAT, or Traditional IPAT)

With IPAT using IP Aliases, when the resource group containing the service IP label fails over from the primary node to the target node, the initial service IP labels that were used at boot time are added as alias addresses on the specified NIC on the target node and removed from the initial node. Unlike in IPAT using IP Replacement, this allows a single NIC to support more than one service IP label placed on it as an alias. Therefore, the same node can host more than one resource group at the same time.

IPAT using IP Aliases has the following advantages over the IPAT using IP Replacement:

- **Speed.** It takes considerably more time moving the IP address and the hardware address than it does to simply move the IP address.
- **Availability.** IP aliasing allows coexistence of multiple service labels on the same network interface. Upon failover, HACMP equally distributes aliases between available network interface adapters.

A limitation of HACMP is that a service IP cannot be on the same subnet as the bootable adapters, so you must create a service IP address that is on a different subnet from other bootable network adapters.

5 HEARTBEATING IN HACMP

A **heartbeat** is a type of communication packet that is sent between cluster nodes. Heartbeats are used to monitor the health of the nodes, networks, and network interfaces as well as to prevent cluster partitioning. In order for an HACMP cluster to recognize and respond to cluster failures, it must continually check the health of the cluster.

Heartbeats can be sent over:

- **TCP/IP networks.** RSCT topology services use the HACMP network topology to dynamically create a set of heartbeat paths that provide coverage for all TCP/IP networks. These paths form heartbeat rings so that all components can be monitored without an excessive number of heartbeat packets. Refer to “HACMP Configuration” on page 12 for configuration instructions.
- **Point-to-point networks.** The non-IP point-to-point network connections that directly link cluster nodes can be configured as alternate heartbeat paths in HACMP. The RS-232 serial point-to-point network can be configured as a heartbeat path to HACMP. The configuration steps are not discussed in this document. You can find the detailed information about this in the *HACMP v5.4 Installation Guide (SC23-5209-00)*.
- **Shared disks.** Heartbeat is supported on any shared disk that is part of an AIX-enhanced concurrent volume group. Heartbeat over disk provides another type of non-IP point-to-point network for failure detection. In a disk heartbeat network, two nodes connected to the disk periodically write heartbeat messages and read heartbeat messages (written by another node) on a small portion of the disk. In this document, the iSCSI protocol is used over which physical disks from NetApp storage are made available to both nodes of the cluster. Refer to “Configuring a Disk Heartbeat” on page15 for configuration instructions.

6 APPLICATION MONITORING IN HACMP

HACMP application monitoring monitors specified applications and attempts to restart them upon detecting process death or application failure. Application monitoring works in one of two ways:

- **Process application monitoring** detects the termination of one or more processes of an application using RSCT Resource Monitoring and Control (RMC).
- **Custom application monitoring** checks the health of an application with a custom monitor method at user-specified polling intervals.

Multiple application monitors can be configured and associated with one or more application servers.

7 THE SERVER ENVIRONMENT

In this report and testing environment, the servers are running AIX 5.3 operating system for HACMP v5.4 and AIX v6.1 operating system for HACMP v5.5. This is now a certified configuration, so the components presented in this document must be used in the same combination to gain support from all parties involved. The only exception is the application of certain patches as defined and required by all the vendors in this configuration. Two NetApp storage systems are configured in a cluster to operate in an NFS environment.

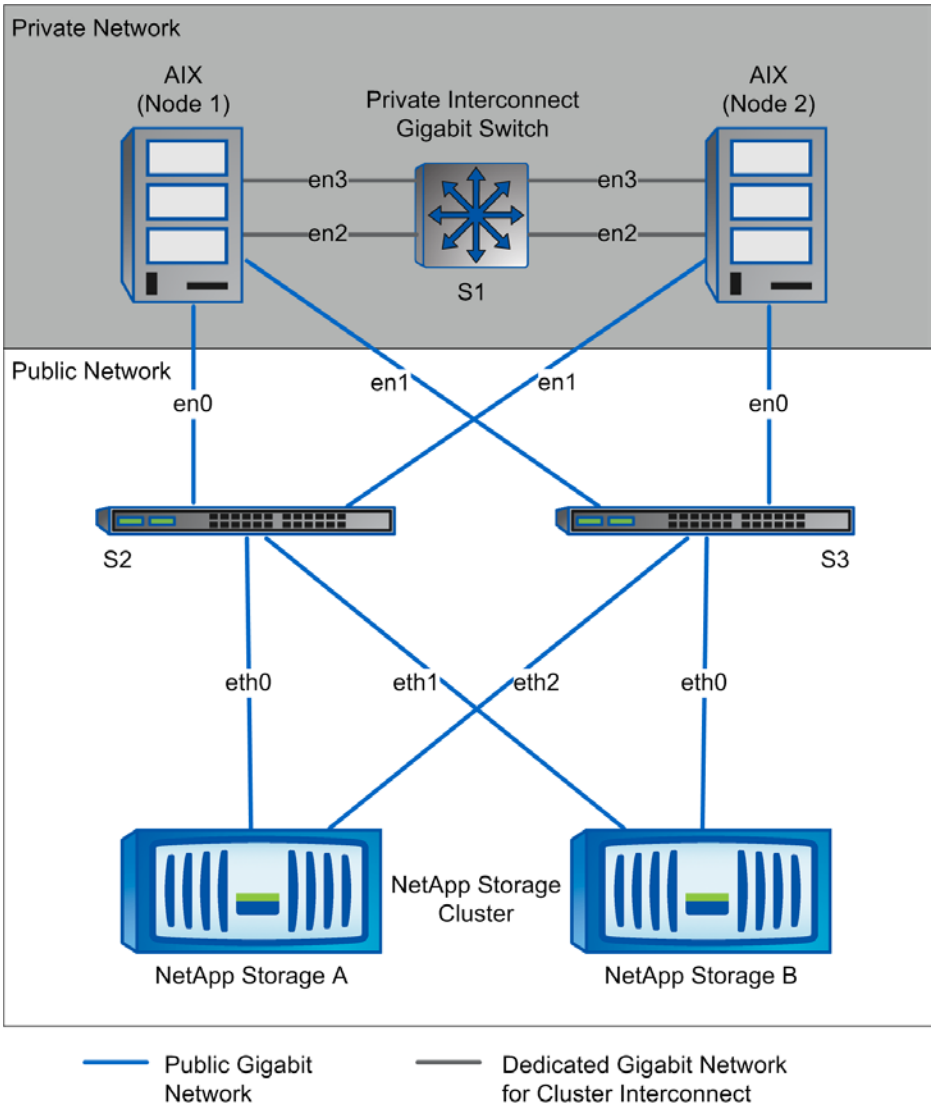

Figure 1) HACMP cluster of two nodes utilizing NetApp storage.

Figure 1 illustrates a typical configuration of a two-node HACMP cluster using the NetApp storage cluster in an NFS environment. This is a scalable configuration and allows users to scale horizontally and internally in terms of processor, memory, and storage.

Because this design couples NIC bonding on AIX hosts and NetApp storage clustering, it has no single point of failure between the host and the data on NetApp storage. Each AIX node has four Ethernet adapters, two of which (en2 and en3) are used as cluster interconnects and connected using a switch, S1. The other two Ethernet adapters (en0 and en1) are cabled to two different gigabit switches, S2 and S3. Ethernet adapter en0 is connected to switch S2, and Ethernet adapter en1 is connected to switch S3.

The AIX Etherchannel feature is used as a port aggregation method whereby two or more Ethernet adapters are defined as one Etherchannel. Remote systems view the Etherchannel as one IP and one MAC address. If any adapter fails, then traffic is automatically sent to the next available adapter in Etherchannel without disrupting user connections. For more information, refer to the Etherchannel configuration tech note document TD101785 at www.ibm.com.

This configuration, with each of the AIX nodes running Etherchannel, provides two hardware paths from each of the nodes to the NetApp storage cluster and greater reliability through a path failover mechanism. If one connection is lost between a node and NetApp storage, then the node continues to access the data over the other path until the failed path is repaired.

The feature of Data ONTAP called virtual interfaces (VIFS) implements link aggregation on the storage system and provides a mechanism to group multiple network interfaces into one logical interface. Refer to the *Network Management Guide* for configuring VIFS on NetApp storage at http://now.netapp.com.

8 REQUIREMENTS

8.1 HARDWARE

CLUSTER NODES

- Two IBM Systems p-Series servers running AIX v5.3 or v6.1 with latest maintenance level
- Two 10/100/1000 Base-TX Ethernet PCI adapters (service IP over which NetApp NFS mountpoints get connected)
- Two 10/100/1000 Base-T Ethernet PCI adapters (for private interconnect)

STORAGE INFRASTRUCTURE

- Two NetApp FAS2xx/F7xx/F8xx/FASF9x/FAS30xx/FAS60xx systems with Data ONTAP 7.2.x or 7.3.x
- Two gigabit NICs in each system
- One or more disk shelves based on the disk space requirements

8.2 SOFTWARE

The following requirements apply to each node in the participating cluster unless specified otherwise:

- AIX v5.3 (ML 4) or v6.1 (TL2-SP1) operating system
- High-Availability Cluster Multiprocessing v5.4 or v5.5 software (read the HACMP release notes for AIX version compatibility)
- Data ONTAP 7.2.x or 7.3.x

9 SETTING UP THE NETAPP STORAGE CLUSTER

9.1 STORAGE CLUSTER PREREQUISITES

Data ONTAP version 7.1.1 or later has a fine-grained NFS lock release mechanism where NFS locks are held and released per process IDs. But in Data ONTAP versions lower than 7.1.1, NFS locks are held or released per the hostname of the NFS client, so we recommend upgrading to Data ONTAP to 7.1.1 or later. In this report and testing environments, NetApp storage systems are running Data ONTAP version 7.2.2 and 7.3.1.

Refer to the support matrix for the latest supported configurations of HACMP and Data ONTAP versions at www.netapp.com/ftp/host-clustering-support-matrix.pdf.

NetApp recommends using flexible volumes in your database environment. NetApp FlexVol® technology pools storage resources automatically and enables you to create multiple flexible volumes on a large pool of disks. This flexibility means you can simplify operations, gain maximum spindle utilization and efficiency, and make changes quickly and seamlessly.

Refer to the NetApp storage system installation and setup guides at http://now.netapp.com.

Make sure following software is required to be installed on both the AIX nodes:

- The NetApp iSCSI AIX initiator Support Kit
- Dot Hill SANpath multipathing software

These software packages can be downloaded from the NetApp NOW™ Web site at http://now.netapp.com. If you do not have access to the above Web site, contact your NetApp sales representative.

9.2 ABOUT THE STORAGE CLUSTER

The storage configuration described in this document is a mirrored, active-active controller configuration of NetApp FAS6070 systems. The words “failover” and “takeover,” “failback” and “giveback” are used interchangeably throughout the document. The word “partner” refers to a storage controller in a clustered pair.

When one partner fails or becomes impaired, a takeover occurs, and the partner storage system continues to serve the failed storage system's data.

When the failed storage system is functioning again, the storage administrator initiates a giveback command that transfers resources (failed over resources) back to the original partner storage system to resume normal operation, serving its own data.

It is recommended not to configure both NetApp storage systems for automatic giveback. Giveback should be initiated manually by the administrator during planned downtime because the giveback process takes longer than the takeover process.

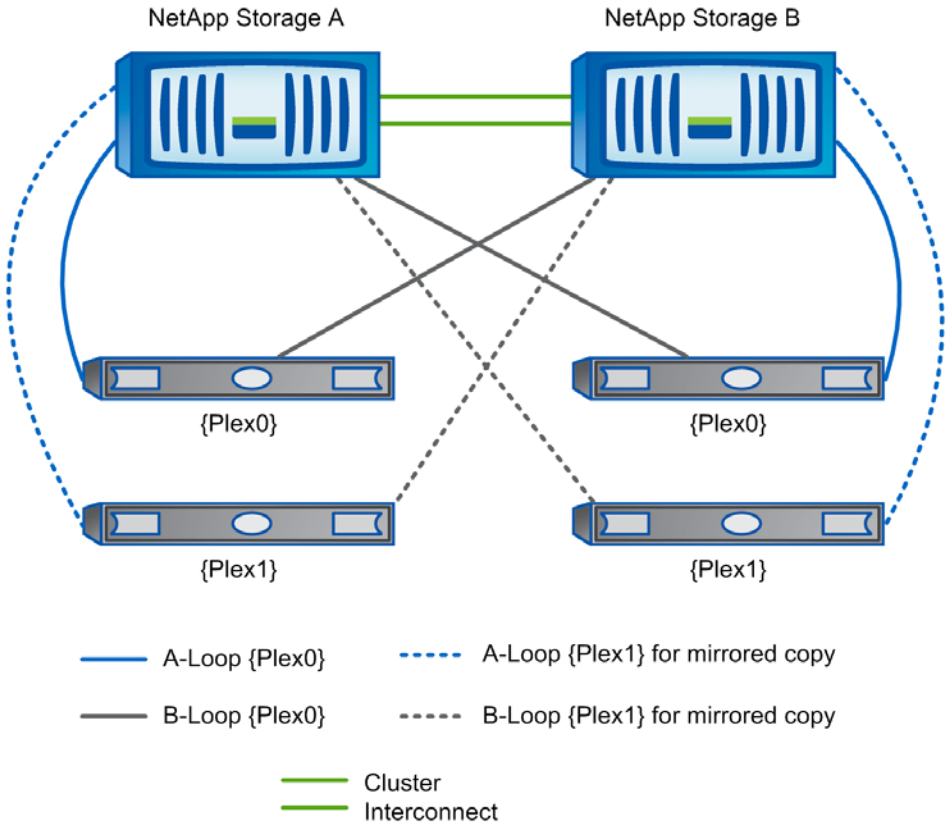

Figure 2) Hardware setup on mirrored active-active controllers.

9.3 CONFIGURING THE STORAGE CLUSTER

1. Configure a NetApp storage system running Data ONTAP 7.2.x or 7.3.x and with cluster, NFS, SnapMirror®, Snapmirror_sync, syncmirror_local, and SnapRestore® license keys.

   The cluster failover parameters on both NetApp storage systems should have following values:

   ```
   CF.GIVEBACK.AUTO.ENABLE                 OFF
   CF.GIVEBACK.CHECK.PARTNER               ON
   CF.TAKEOVER.DETECTION.SECONDS           15
   CF.TAKEOVER.ON_FAILURE                  ON
   CF.TAKEOVER.ON_NETWORK_INTERFACE_FAILURE  ON
   CF.TAKEOVER.ON_PANIC                    ON
   CF.TAKEOVER.ON_SHORT_UPTIME             ON
   ```

2. Create volumes for storing shared application files on the storage device.

   For example, create two volumes named “data” and “binary” using the following command at the NetApp storage console:

   ```
   Storage> vol create data 14
   Storage> vol create binary 8
   ```

   Note: We created volume “data” with 14 disks and volume “binary” with eight disks as an example. Create your volumes based on your workload and application needs.

3. Open the `/etc/exports` file on NetApp storage and add the following entries:

   ```
   /vol/data -anon=0
   /vol/binary -anon=0
   ```

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

   ```
   Storage> exportfs -a
   ```

5. Configure NetApp storage for disk heartbeating over iSCSI.

   HACMP needs an enhanced concurrent capable volume group to configure the disk heartbeat. The enhanced concurrent capable volume group requires a disk device on AIX hosts. Since a volume mounted over NFS protocol is considered to be a file system, it cannot be used for creating the volume group on AIX, so you must provide the AIX hosts with a physical disk over either FCP or iSCSI protocol. In this report and testing environment, iSCSI protocol is used for disk heartbeat.

   Note: A license for iSCSI is free for a NetApp storage system.

CONFIGURE THE AIX NODES TO ACCESS LUNS OVER THE ISCSI PROTOCOL

1. On one AIX node, log in as the root user and make the following command:

   ```
   btcppesrv5#> lsattr –El iscsi0 nodename.iqn.aixhost1.hostid.0a3cac46
   ```

   The AIX default initiator nodename does not fully comply with the iSCSI specifications and NetApp nodename requirements.

2. Change the default initiator nodename to include a date in the second field as in this example: `iqn.yyyy-mm.hostname.hostid.0a3cac46`.

   ```
   btcppesrv5#> chdev -l iscsi0 -a initiator_name=iqn.1996-
   04.aixhost1.hostid.0a3cac46
   ```

3. Repeat the previous steps for second AIX node.

   ```
   btcppesrv6#> chdev -l iscsi0 -a initiator_name=iqn.1996-
   04.aixhost2.hostid.0a3cac47
   ```

4. On the NetApp storage system, enable the iSCSI service:

   ```
   Storage> iscsi start
   Storage> iscsi nodename
   ```

   This will provide the iSCSI target nodename. In our example, it is:

   ```
   iqn.1992-08.com.netapp:sn.118042259
   ```

5. Create AIX igroups using the AIX host nodenames:

   ```
   Storage> igroup create -i -t aix aixhost1-iscsi iqn.1996-
   04.aixhost1.hostid.0a3cac46
   Storage> igroup create -i -t aix aixhost2-iscsi iqn.1996-
   04.aixhost2.hostid.0a3cac47
   ```

6. Create the LUNs:

   ```
   Storage> lun create -s 10g -t aix /vol/aix/one
   ```

7. Map the LUNs to both igroups created above:

   ```
   Storage> lun map /vol/aix/one aixhost1-iscsi
   Storage> lun map /vol/aix/one aixhost2-iscsi
   ```

8. Access the NetApp iSCSI LUNs from AIX hosts and edit the `/etc/iscsi/targets` file to add the NetApp target in the form of a destination IP address, port, and destination nodename:

   ```
   10.73.68.112 3260 iqn.1992-08.com.netapp:sn.118042259
   ```

9. Use the `cfgmgr` command on AIX to pick up the newly created iSCSI disk devices:

   ```
   btcppesrv5#> cfgmgr –l iscsi0
   btcppesrv6#> cfgmgr –l iscsi0
   ```

10. Check for the newly added iSCSI disk:

    ```
    btcppesrv5#> lsdev -Cc disk
    hdisk0 Available 09-08-00-5,0 16 Bit LVD SCSI Disk Drive hdisk1 Available
    09-08-00-8,0 16 Bit LVD SCSI Disk Drive hdisk2 Available Other iSCSI Disk
    Drive
    btcppesrv6#> lsdev –Cc disk
    ```

```
hdisk0 Available 09-08-00-5,0 16 Bit LVD SCSI Disk Drive hdisk1 Available
09-08-00-8,0 16 Bit LVD SCSI Disk Drive hdisk2 Available Other iSCSI Disk
Drive
```

The disk device hdisk2 is available for use on both AIX nodes. The enhanced concurrent capable volume group needs to be created for the HACMP disk heartbeat (see “Configuring a Disk Heartbeat” on page 15).

Note: The iSCSI protocol is used only for configuring the disk heartbeat. The application data resides on file systems mounted over NFS protocol.

10 PREINSTALLATION OS ACTIVITIES FOR HACMP

1. Make sure the following AIX base operating system (BOS) components are present for HACMP.

Table 1) AIX BOS components.

AIX BOS Component	AIX 5L v.5.3 for HACMP v5.4	AIX 6L v.6.1 for HACMP v5.5
bos.adt.libm	5.3.0.10	6.1.2.0
bos.adt.syscalls	5.3.0.10	6.1.2.0
bos.net.tcp.client	5.3.0.10	6.1.2.0
bos.net.tcp.server	5.3.0.10	6.1.2.0
bos.rte.SRC	5.3.0.10	6.1.2.0
bos.rte.libc	5.3.0.10	6.1.2.0
bos.rte.libcfg	5.3.0.10	6.1.2.0
bos.rte.libcur	5.3.0.10	6.1.2.0
bos.rte.libpthreads	5.3.0.10	6.1.2.0
bos.rte.odm	5.3.0.10	6.1.2.0
bos.clvm.enh	5.3.0.10	6.1.2.0

2. Install the Reliable Scalable Cluster Technology (RSCT) images before installing HACMP. Make sure that each node has the same version of RSCT.

Table 2) RSCT file sets.

AIX RSCT File Sets	AIX 5L v.5.3 for HACMP v5.4	AIX 6L v.6.1 for HACMP v5.5
rsct.basic.rte	2.4.5.0	2.5.2.0
rsct.basic.hacmp	2.4.5.0	2.5.2.0
rsct.basic.sp	2.4.5.0	2.5.2.0
rsct.compat.basic.rte	2.4.5.0	2.5.2.0
rsct.compat.basic.hacmp	2.4.5.0	2.5.2.0
rsct.compat.basic.sp	2.4.5.0	2.5.2.0
rsct.compat.clients.hac	2.4.5.0	2.5.2.0
rsct.compat.clients.rte	2.4.5.0	2.5.2.0
rsct.compat.clients.sp	2.4.5.0	2.5.2.0
rsct.core.sec	2.4.5.0	2.5.2.0
rsct.core.rmc	2.4.5.0	2.5.2.0

3. Use the following commands to determine whether the appropriate file sets are installed and their level:

```
/usr/bin/lslpp -l | grep rsct
/usr/bin/lslpp -l | grep bos
```

11 HACMP SOFTWARE INSTALLATION

The HACMP software installation medium contains the HACMP enhanced scalability subsystem images. This provides the services for cluster membership, system management, configuration integrity and control, failover, and recovery. It also includes cluster status and monitoring facilities for programmers and system administrators.

INSTALL THE HACMP SOFTWARE ON A SERVER NODE

1. Insert the CD into the CD-ROM drive and enter:

   ```
   smit install_all
   ```

   SMIT displays the first Install and Update from ALL Available Software panel.
2. Enter field values as follows.

   The fields not mentioned in Table 3 should be kept at their default values.

Table 3) Installation parameters.

Field	Values
INPUT device/directory for software	Enter the device name of the installation medium or install directory.
SOFTWARE to install	Select an option from the pick list or enter `all` to install all server and client images. Make sure to install the level of RSCT required for AIX. Refer to “Preinstallation OS Activities for HACMP” on page 10.
PREVIEW only?	Set to `yes` to use the preview option, which just checks to make sure that installation prerequisites are met. Set to `no` to perform the actual installation.
AUTOMATICALLY install requisite software	Use the default. Set this field to `no` if the prerequisite software for HACMP is already installed or if the OVERWRITE same or newer versions? field is set to `yes`. Otherwise, set this field to `yes` to install the required software.
ACCEPT new license agreements?	Select `yes` to proceed with installation.

3. When you are satisfied with the entries, press Enter.

   SMIT prompts you to confirm the selections.
4. Press Enter again to install.
5. Run the following commands to verify the software installation:

   ```
   lppchk -v
   lppchk -c “cluster.*”
   ```

   The `lppchk` command verifies that files for an installable software product (file set) match the Software Vital Product Data (SWVPD) database information for file sizes, checksum values, or symbolic links. If the installation is OK, both commands should return nothing.
6. Reboot each HACMP cluster node.

12 HACMP CONFIGURATION

12.1 CONFIGURING NETWORK INTERFACE ADAPTERS

In our example, we have two NICs, one used as cluster interconnects and other as a bootable adapter. The service IP address will be activated on the bootable adapter after cluster services are started on the nodes. The following IP addresses are used in the setup:

```
NODE1: hostname – btcppesrv5
Boot IP address                          10.73.70.155        btcppesrv5
Netmask                                  255.255.254.0
Interconnect IP address                  192.168.73.100      btcppesrv5i
Netmask                                  255.255.255.0       btcppesrv5sv
Service IP address                       10.73.68.222
Netmask                                  255.255.254.0
NODE2: hostname – btcppesrv6
Boot IP address                          10.73.70.156        btcppesrv6
Netmask                                  255.255.254.0
Interconnect IP address                  192.168.73.101      btcppesrv6i
Netmask                                  255.255.255.0
btcppesrv6sv
Service IP addresss                      10.73.68.223
Netmask                                  255.255.254.0
```

The hostname may not contain -, _, * or other special characters.

12.2 EDITING CONFIGURATION FILES FOR HACMP

All the files mentioned in this procedure need to be configured on both nodes of cluster. Refer to “Appendix A: Sample configuration files for HACMP” on page 22 for a sample file.

1. Enter all the IP addresses present in the network in these files:
   - `/usr/sbin/cluster/netmon.cf`
   - `/usr/sbin/cluster/etc/clhosts`
   - `/usr/sbin/cluster/etc/rhosts`
2. Enter all the IP addresses present in the network with username (that is, `root`) in this file: `/.rhosts.`
3. Enter all the IP addresses with their IP labels present in network in this file: `/etc/hosts.`

12.3 CREATING A CLUSTER USING SMIT

This section presents a sample HACMP configuration that might require customization for your environment. This section demonstrates how to configure two AIX nodes, btcppesrv5 and btcppesrv6, into an HACMP cluster.

Note: These instructions assume that you are using the graphical user interface to SMIT (`smit –M`). If you are using the ASCII interface to SMIT (smit –C), then modify these instructions accordingly.

1. Configure two AIX nodes to allow the user root to use the `rcp` and `remsh` commands between themselves without having to specify a password.
2. Log in as the user root on AIX node btcppesrv5.
3. Enter the following command:

   ```
   # smit hacmp
   ```
4. Click Initialization and Standard Configuration.
5. Click Configure an HACMP Cluster and Nodes.

6. In the Cluster Name field, enter `netapp`.
7. In the New Nodes (using selected communication paths) field, enter `btcppesrv5` and `btcppesrv6`.
8. Click OK.
9. Click Done.
10. Click Cancel.
11. Select Exit > Exit SMIT Menu.

CONFIGURE THE HEARTBEAT NETWORKS AS PRIVATE NETWORKS

1. Enter the following command:
   ```
   # smit hacmp
   ```
2. Click Extended Configuration.
3. Click Extended Topology Configuration.
4. Click Configure HACMP Networks.
5. Click Change/Show a Network in the HACMP cluster.
6. Select net_ether_01 (192.168.73.0/24).
7. In the Network Attribute field, select private.
8. Click OK.
9. Click Done.
10. Click Cancel.
11. Select Exit > Exit SMIT Menu.

CONFIGURE THE SERVICE IP LABELS/ADDRESSES

1. Enter the following command:
   ```
   # smit hacmp
   ```
2. Click Initialization and Standard Configuration.
3. Click Configure Resources to Make Highly Available.
4. Click Configure Service IP Labels/Addresses.
5. Click Add a Service IP Label/Address.
6. In the IP Label/Address field, enter `btcppsrv5sv`.
7. In the Network Name field, select net_ether_02 (10.73.70.0/23).
   The Service IP label will be activated on network interface 10.73.70.0/23 after the cluster service starts.
8. Click OK.
9. Click Done.
10. Click Cancel.
11. Repeat these steps to add a second service IP label, `btcppsrv6sv`**.**
12. Select Exit > Exit SMIT Menu.

CREATE EMPTY RESOURCE GROUPS WITH NODE PRIORITIES

1. Enter the following command:
   ```
   # smit hacmp
   ```
2. Click Initialization and Standard Configuration.
3. Click Configure HACMP Resource Groups.
4. Click Add a Resource Group.
5. In the Resource Group Name field, enter `RG1`.
6. In the Participating Nodes (Default Node Priority) field, enter `btcppesrv5` and `btcppesrv6`.

The Resource Group RG1 will be online on btcppsrv5 first when the cluster service starts. In the event of failure, RG1 will be taken over by btcppesrv6 as the node priority for RG1 is assigned to btcppesrv5 first.

7. Click OK.
8. Click Done.
9. Click Cancel.
10. Repeat these steps to add a second Resource group, RG2 with node priority first assigned to `btcppesrv6`.
11. Select Exit > Exit SMIT Menu.

MAKE SERVICE IP LABELS PART OF THE RESOURCE GROUPS

1. Enter the following command:
   ```
   # smit hacmp
   ```
2. Click Initialization and Standard Configuration.
3. Click Configure HACMP Resource Groups.
4. Click Change/Show Resources for a Resource Group (standard).
5. Select a resource Group from pick list as RG1.
6. In the Service IP Labels/Addresses field, enter `btcppesrv5sv`. As btcppesrv5sv service IP label must be activated on first node btcppesrv5.
7. Click OK.
8. Click Done.
9. Click Cancel.
10. Repeat these steps to add a second Service IP Label `btcppesrv6sv` in Resource Group RG2. A btcppesrv6sv service IP label must be activated on second node btcppesrv6.
11. Select Exit > Exit SMIT Menu.

12.4 VERIFYING AND SYNCHRONIZING A CLUSTER

This section demonstrates how to verify and synchronize the nodes in an HACMP cluster. This process verifies that the HACMP configuration is done from one node and then synchronizes to other node in the cluster. Whenever any changes are necessary in the HACMP cluster, they must be done from a single node and then synchronized with other nodes.

1. Log in as user root on AIX node btcppesrv5.
2. Enter following command:
   ```
   # smit hacmp
   ```
3. Click Initialization and Standard Configuration.
4. Click Verify and Synchronize HACMP Configuration.
5. Click Done.
6. Select Exit > Exit SMIT Menu.

12.5 STARTING CLUSTER SERVICES

This section demonstrates how to start an HACMP cluster on both the participating nodes.

1. Log in as user root on AIX node btcppesrv5.
2. Enter following command:
   ```
   # smit cl_admin
   ```
3. Click Manage HACMP services.
4. Click Start Cluster Services.
5. In the Start Now, on System Restart or Both fields, select now.

6. In the Start Cluster Services on these nodes field, enter `btcppesrv5` and `btcppesrv6`.

   The cluster services can be started on both nodes simultaneously.
7. In the Startup Cluster Information Daemon field, select true.
8. Click OK.
9. Click Done.
10. Click Cancel.
11. Select Exit > Exit SMIT Menu.

12.6 STOPPING CLUSTER SERVICES

This section demonstrates how to stop an HACMP cluster on both participating nodes.

1. Log in as user root on AIX node btcppesrv5.
2. Enter following command:

   ```
   # smit cl_admin
   ```
3. Click Manage HACMP services.
4. Click Stop Cluster Services.
5. In the Stop Now, on System Restart or Both fields, select now.
6. In the Stop Cluster Services on these nodes field, enter `btcppesrv5` and `btcppesrv6`. The cluster services can be stopped on both the nodes simultaneously.
7. Click OK.
8. In the Are You Sure? dialog box, click OK.
9. Click Done.
10. Click Cancel.
11. Select Exit > Exit SMIT Menu.

12.7 CONFIGURING A DISK HEARTBEAT

For configuring disk heartbeating, you must create the Enhanced Concurrent Capable Volume group on both AIX nodes. As explained in “Setting up the NetApp Storage Cluster” on page 7, the physical disks are available for both the AIX nodes.

To use HACMP C-SPOC successfully, a basic IP-based topology must already exist and the storage devices must have their PVIDs on the ODMs of both systems. This can be verified by running the `lspv` command on each AIX node. If a PVID does not exist on any AIX node, then run the `chdev –l <devicename> -a pv=yes` command on each AIX node.

```
btcppesrv5#> chdev –l hdisk3 –a pv=yes btcppesrv6#> chdev –l hdisk3 –a
pv=yes
```

This will allow C-SPOC to match up the devices as known shared storage devices.

CREATE AN ENHANCED CONCURRENT VOLUME GROUP

1. Log in as user root on AIX node btcppesrv5.
2. Enter following command to create Enhanced concurrent VG.

   ```
   # smit vg
   ```
3. Click Add Volume Group.
4. Click Add an Original Group.
5. In the Volume group name field, enter `heartbeat`.
6. In the Physical Volume Names field, enter `hdisk3`.
7. In the Volume Group Major number field, enter `59`.

This number is the number available for a particular AIX node; it can be found on the available list in the field.

8. In the Create VG concurrent capable field, enter YES.
9. Click OK.
10. Click Done.
11. Click Cancel.
12. Select Exit > Exit SMIT Menu.
13. On btcppesrv5 AIX node, check the newly created volume group using command `lsvg`.
14. On second AIX node, enter the `importvg –V <major number> -y <volume group> <device name>` command to import the volume group:

    ```
    btcppesrv6#> importvg -V 59 -y heartbeat hdisk3
    ```

Since the enhanced concurrent volume groups are available for both the AIX nodes, we will use the discovery method of HACMP to find the disks available for heartbeat.

CONFIGURE THE DISK HEARTBEAT IN HACMP

1. Log in as user root on AIX node btcppesrv5.
2. Enter following command to configure Disk heartbeat:

    ```
    # smit hacmp
    ```
3. Click Extended Configuration.
4. Click Discover HACMP-related information from configured Nodes.

    This will run automatically and create `/usr/es/sbin/cluster/etc/config/clvg_config` file that contains the information it has discovered.
5. Click Done.
6. Click Cancel.
7. Click Extended Configuration.
8. Click Extended Topology Configuration.
9. Click Configure HACMP communication Interfaces/Devices.
10. Click Add Communication Interfaces/Devices.
11. Click Add Discovered Communication Interfaces and Devices.
12. Click Communication Devices.
13. Select both the Devices listed in the list.
14. Click Done.
15. Click Cancel.
16. Select Exit > Exit SMIT Menu.

ADD THE VOLUME GROUP INTO HACMP RESOURCE GROUP AND SYNCHRONIZE THE CLUSTER

1. Enter the following command to create Empty Resource Groups with different policies than what we created earlier.

    ```
    # smit hacmp
    ```
2. Click Initialization and Standard Configuration.
3. Click Configure HACMP Resource Groups.
4. Click Add a Resource Group.
5. In the Resource Group Name field, enter `RG3`.
6. In the Participating Nodes (Default Node Priority) field, enter `btcppesrv5` and `btcppesrv6`.
7. In the Startup policy field, enter `Online On All Available Nodes`.
8. In the Fallover Policy field, enter `Bring Offline (On Error Node Only)`.

9. In the Fallback Policy field, enter `Never Fallback`.
10. Click OK.
11. Click Done.
12. Click Cancel.
13. Click Change/Show Resources for a Resource Group (Standard).
14. Select RG3 from the list.
15. In the Volume Groups field, enter `heartbeat`, which is the concurrent capable volume group which was created earlier.
16. Click OK.
17. Click Done.
18. Click Cancel.
19. Select Exit > Exit SMIT Menu.

Follow the steps described in “Verifying and Synchronizing a Cluster” on page 14 to propagate modifications in HACMP cluster to other nodes.

TESTING DISK HEARTBEAT CONNECTIVITY

After the device and network definitions have been created, you should test it to make sure communication is working properly.

The `/usr/sbin/rsct/bin/dhb_read` command is used to test the validity of a diskh connection as follows:

```
dhb_read -p devicename -r #receives data over diskhb network
dhb_read -p devicename -t #transmits data over diskhb network
```

1. On btcppesrv5, run the following:

   ```
   dhb_read -p hdisk3 -r
   ```

2. On btcppesrv6, run the following:

   ```
   dhb_read -p hdisk3 -t
   ```

   If the link from btcppesrv5 to btcppesrv6 is operational, both nodes will display `Link operating normally`.

3. Run this again and swap which node transmits and which one receives.

After the cluster is up and running, one can monitor the activity of the disk heartbeats using the `lssrc –ls topsvcs` command.

An example of the command output is as follows:

```
btcppesrv5#> lssrc –ls topsvcs
Subsystem Group PID Status topsvcs topsvcs 32108 active
Network Name Indx Defd Mbrs St Adapter ID Group ID disknet1 [ 3] 2 2 S
255.255.10.0 255.255.10.1 disknet1 [ 3] rvpath3 0x86cd1b02 0x86cd1b4f
HB Interval = 2 secs. Sensitivity = 4 missed beats
Missed HBs: Total: 0 Current group: 0
Packets sent : 229 ICMP 0 Errors: 0 No mbuf: 0
Packets received: 217 ICMP 0 Dropped: 0
NIM's PID: 28724
```

12.8 CREATING APPLICATION SERVER FOR NFS MOUNTPOINTS

This section demonstrates how to mount the NFS file systems onto HACMP cluster nodes. As described in “Setting up the NetApp Storage Cluster” on page 7, there are two volumes, `/vol/data` and `/vol/binary`, residing on the NetApp storage system. Both these volumes are NFS exported at the NetApp storage. The requirement is to mount these volumes on HACMP cluster nodes when Cluster services get started. In this situation, NetApp storage acts as NFS server and HACMP cluster nodes act as NFS clients.

When one node mounts the NFS file system from NetApp storage and that node fails, then the other node will detect the failure and will mount the NFS file system onto it and start the application processes, thus maintaining the application availability.

HACMP version 5.4 has the feature of Application Servers as highly available resources, so you must create a startup and stop script on both nodes that will be part of an Application server. The startup script contains the AIX mount command with the required mount options to mount the NFS file system. The stop script contains the AIX umount command to unmount the NFS file system. The sample startup and stop scripts are shown in “Appendix A: Sample configuration files for HACMP” on page 22.

CONFIGURE THE APPLICATION SERVER

1. Log in as user root on AIX node btcppesrv5.
2. Enter following command:

   ```
   # smit hacmp
   ```
3. Click Initialization and Standard Configuration.
4. Click Configure Resources to Make Highly Available.
5. Click Configure Application Servers.
6. Click Add Application Server.
7. In the Server Name field, enter `App1`.
8. In the Start Script field, enter `/scripts/nfs_start`.

   The path of the startup script must be changed. This script must be copied to both the nodes in the same location. This is required in case of node failure.
9. In the Stop Script field, enter `/scripts/nfs_stop`.

   The path of the stop script is to be provided. This script must be copied to both the nodes in the same location. This is required in case of node failure.
10. Click Done.
11. Click Cancel.
12. Repeat these steps to add other application servers.
13. Select Exit > Exit SMIT Menu.

ADD AN APPLICATION SERVER RESOURCE AS PART OF AN EXISTING RESOURCE GROUP

1. Log in as the user root on AIX node btcppesrv5.
2. Enter following command:

   ```
   # smit hacmp
   ```
3. Click Initialization and Standard Configuration.
4. Click Configure HACMP Resource Groups.
5. Click Configure Application Servers.
6. Click Change/Show Resources for a Resource Group (standard).
7. Select a resource Group from pick list as RG1.
8. In the Application Servers field, enter `App1`.

   By selecting App1 in Resource Group RG1, the HACMP mounts the NFS file system on node btcppesrv5 since the App1 startup script contains the NFS mount command.
9. Click OK.
10. Click Done.
11. Click Cancel.
12. Select Exit > Exit SMIT Menu.

Follow the steps described in “Verifying and Synchronizing a Cluster” on page 14 to propagate modifications in HACMP cluster to other nodes.

12.9 CONFIGURING APPLICATION MONITORS FOR NFS MOUNTPOINTS

The custom application monitoring is used to monitor the NFS mountpoints. This requires a shell script that will check the health of the NFS mountpoints periodically and provide a 0 exit code to HACMP if NFS mountpoints are healthy; otherwise it will provide any nonzero exit code. Depending on the exit code provided, HACMP will run the restart script to remount the NFS mountpoints. Refer to “Appendix A: Sample configuration files for HACMP” on page 22 for sample shell script (nfs_monitor) to monitor the health of the NFS mountpoints.

This demonstrates how to create a custom Application monitor for monitoring the application server which mounts the volume groups over NFS from NetApp storage system.

1. Log in as user root on AIX node btcppesrv5.
2. Enter following command.

   ```
   # smit hacmp
   ```
3. Click Extended Configuration.
4. Click Extended Resource Configuration.
5. Click HACMP Extended Resources Configuration.
6. Click Configure HACMP application servers.
7. Click Configure HACMP application monitoring.
8. Click Configure Custom Application Monitors.
9. Click Add Custom Application monitor.
10. In the Monitor Name field, enter `nfs_monitor`.
11. In the Application Servers to Monitor field, enter `App1`. Provide the name of the application server that is mounting the NFS volume groups.
12. In the Monitor mode field, enter `long-running monitoring`.

    In this mode, the application monitor periodically checks that the application server is running. The checking starts after the specified stabilization interval has passed.
13. In the Monitor Method field, enter `/scripts/nfs_monitor`. Provide the path of the shell script which monitors the health of the NFS mountpoints.
14. In the Monitor Interval field, enter `5`.

    This is the time in seconds after which every specified value in the shell script will be executed. The minimum value is 1 second.
15. In the Stabilization Interval, enter `200`.

    This is the period during which HACMP waits for an application to stabilize before beginning to monitor whether the application is running.
16. In the Restart Count field, enter `5`.

    This is the number of times to try restarting an application before taking any other action.
17. In the Action on Application failure field, enter `FAILOVER`.
18. In the Restart method field, enter `/scripts/nfs_start`.

    This script contains the NFS mount command that remounts the NFS mountpoints. This is the same script used by the application server App1.
19. Click OK.
20. Click Done.
21. Click Cancel.
22. Select Exit > Exit SMIT Menu.

Follow the steps described in “Verifying and Synchronizing a Cluster” on page 14 to propagate modifications in HACMP cluster to other nodes.

12.10 CONFIGURING PROCESS APPLICATION MONITORS

Process application monitoring uses the RSCT subsystem functions to detect the termination of a process and to generate an event. This section describes how to configure process application monitoring, in which one can specify one or more processes of a single application to be monitored. The Oracle® Database is considered as an example while configuring the Process application monitor.

For process monitoring, make sure to list the correct process names in the SMIT Add Process Application Monitor panel. Use processes that are listed in response to the `ps –el` command, and not the `ps –ef` command**.**

1. Log in as the user root on AIX node btcppesrv5.
2. Enter following command:

   ```
   # smit hacmp
   ```
3. Click Extended Configuration.
4. Click Extended Resource Configuration.
5. Click HACMP Extended Resources Configuration.
6. Click Configure HACMP application servers.
7. Click Configure HACMP application monitoring.
8. Click Configure Process Application Monitors.
9. Click Add Process Application monitor.
10. In the Monitor Name field, enter `oracle_monitor`.
11. In the Application Servers to Monitor field, enter the application server name that is starting the Oracle Database.
12. In the Monitor mode field, enter `long-running monitoring`.

    In this mode, the application monitor periodically checks that the application server is running. The checking starts after the specified stabilization interval has passed.
13. In the Processes to monitor field, enter `oracle`.

    Use correct process names as they appear in `ps-el` command.
14. In the process owner field, enter `oracle`.

    This is the owner of Oracle Database processes.
15. In the Instance Count field, enter `13`.

    This is the number of Oracle processes running on the AIX node.
16. In the Stabilization Interval, enter `200`.

    This is the period that HACMP waits for an application to stabilize before beginning to monitor that the application is running.
17. In the Restart Count field, enter `5`.

    This specified number of times to try restarting the application before taking any other action.
18. In the Action on Application failure field, enter `FALLOVER`.
19. In the Cleanup method field, enter `/scripts/oracle_clean`.

    This script will kill defunct Oracle processes if running. Refer to Appendix for sample script.
20. In the Restart method field, enter `/scripts/oracle_restart`.

    This script will start Oracle Database and listener. This is the same script used by the application server that starts the Oracle Database. Refer to “Appendix A: Sample configuration files for HACMP” on page 22 for a sample script.
21. Click OK.
22. Click Done.
23. Click Cancel.
24. Select Exit > Exit SMIT Menu.

Follow the steps described in “Verifying and Synchronizing a Cluster” on page 14 to propagate modifications in HACMP cluster to other nodes.

12.11 SETTING UP MULTIPLE ORACLE INSTANCES ON A TWO-NODE HACMP CLUSTER

If the setup consists of multiple Oracle instances running on same node but belonging to different resource groups, then use dedicated service IP addresses for each Oracle instance. This will help in the event of failure of any one Oracle instance to successfully getting started on another node without interrupting other Oracle instance resource groups.

To make a successful failover of a single resource group to another node, it is recommended to clean up the failed Oracle Database processes on the original node. Refer to “Appendix A: Sample configuration files for HACMP” on page 22 for a sample cleanup script (/scripts/oracle_clean). This script can be customized for a specific Oracle Database ID (SID).

13 HACMP CLUSTER BEHAVIOR DURING NETAPP CLUSTER FAILOVER (CFO)

The storage configuration used for testing environment is a mirrored, active-active controller configuration of NetApp FAS6070 systems. For more information about how this works, refer to “About the Storage Cluster” on page 7.

In both the processes of takeover and giveback, the NetApp storage systems are inaccessible to AIX hosts for a short duration since HACMP cluster heartbeat is configured to use disk heartbeat where AIX hosts access a shared disk over iSCSI protocol. The heartbeat will flow through TCP/IP network, and thus HACMP cluster will continue to work smoothly in takeover and giveback process.

Also, since volumes on NetApp storage are accessed over the NFS protocol by the application, they will not be served for the short duration, and once the takeover or giveback processes are completed, the applications will continue to work again.

14 SUMMARY

When setting up an HACMP v5.4 or v5.5 on IBM System p-Series or i-Series using AIX v5.3 or v6.1, you should do the following:

- Try to get the latest maintenance level (ML) supported by IBM for AIX. Read the HACMP release notes for AIX version compatibility.
- Mount with NFS over TCP and NFS version 3 where possible, and use the largest rsize and wsize that still provide good performance. The NFS mount options should be specified per the requirement of the application one wants to run on the HACMP cluster. Refer to NetApp technical reports for different NFS mount options for different applications.
- Provide multiple heartbeat paths for HACMP to prevent node partitioning. Use a redundant non-IP network for cluster interconnects such as serial RS-232 cable. You can find information on configuring serial network for heartbeat paths in *HACMP v5.4 Installation Guide (SC23-5209-00)* and *HACMP v5.5 Installation Guide (SC23-5209).*
- Be sure the network between the HACMP cluster and NetApp storage drops as few packets as possible. This can be achieved by tuning the network parameters for your switches and NICs according to the network equipments best practice recommendation.
- Always look for the latest errata and bug fixes from the IBM support site (www.ibm.com/support) and watch for new NetApp technical reports at www.netapp.com/library/tr.

APPENDIX A: SAMPLE CONFIGURATION FILES FOR HACMP

1. /usr/sbin/cluster/netmon.cf

```
10.73.70.155
10.73.70.156
192.168.73.101
192.168.73.100
10.73.68.222
10.73.68.223
```

2. /usr/sbin/cluster/etc/clhosts

```
10.73.70.155
10.73.70.156
192.168.73.101
192.168.73.100
10.73.68.222
10.73.68.223
```

3. /usr/sbin/cluster/etc/rhosts

```
10.73.70.155
10.73.70.156
192.168.73.101
192.168.73.100
10.73.68.222
10.73.68.223
```

4. /.rhosts

```
10.73.70.155      root
10.73.70.156      root
192.168.73.101    root
192.168.73.100    root
10.73.68.222      root
10.73.68.223      root
```

5. /etc/hosts

```
10.73.70.156      btcppesrv6        #Boot IP
10.73.70.155      btcppesrv5        #Boot IP
192.168.73.100    btcppesrv5i       #Interconnect
192.168.73.101    btcppesrv6i       #Interconnect
10.73.68.222      btcppesrv5sv      #Service IP
10.73.68.223      btcppesrv6sv      #Service IP
```

6. /scripts/nfs_start

```
#!/bin/sh
#NetApp storage IP address is 10.73.68.105
mount –o “bg,hard,nointr,rw,proto=tcp,vers=3”\
10.73.68.105:/vol/data /data
```

7. /scripts/nfs_stop

```
#!/bin/sh
umount /data;
```

8. /scripts/nfs_monitor

```
#!/usr/bin/ksh
# NFS Mount point name is /oradata_9i
DF=`df | grep /oradata_9i | awk ‘{print $3}’`
if [${DF} = ””]
then
      exit 1
else
      exit 0
fi
```

9. /scripts/oracle_clean

```
#! /usr/bin/ksh
#Database SID is oradb.
#Storage is the hostname for NetApp storage.
#btcppesrv5 is hostname of the AIX node which holds locks on
storage.
#Enable rsh between storage and host. One can configure ssh also.
#The command sm_mon is used to clear stale NFS locks per process
IDs.

echo “Getting the oracle pid\n”;
for pid in `ps –ef | grep oradb | awk ‘{print $2}’`
do
echo $pid
rsh Storage “priv set advanced; sm_mon –l btcppesrv5 –p $pid;”
if [ $? –eq 0 ];
      then
            echo “pid - $pid cleared”
      else if
            echo “pid - $pid not cleared”
done
```

10. /scripts/oracle_restart

```
#!/usr/bin/ksh
#NODE19I is the name of the listener
su – oracle –c ‘export ORACLE_SID=node19i; sqlplus –S “/ as
sysdba”
@startup.sql;’
su – oracle –c ‘lsnrctl start NODE19I;’
```

Note: Modify these sample scripts as per your environment.

APPENDIX B: FREQUENTLY ASKED QUESTIONS

- **Q:** What characters should a hostname contain for HACMP configuration?
  A: The hostname cannot have following characters: -, _, * or other special characters.
- **Q:** Can the service IP and the boot IP be in the same subnet?
  A: No. The service IP address and boot IP address cannot be in the same subnet. This is the basic requirement for an HACMP cluster configuration. The verification process does not allow the IP addresses to be in same subnet and cluster will not start.
- **Q:** Can multiple service IP addresses be configured on single Ethernet cards?
  A: Yes. Using the SMIT menu, it can be configured to have multiple service IP addresses running on a single Ethernet card. It only requires selecting the same network name for specific service IP addresses in the SMIT menu.
- **Q:** What happens when a NIC having a service IP goes down?
  A: When a NIC card running the service IP address goes down, the HACMP detects the failure and fails over the service IP address to an available standby NIC on the same node or to another node in the cluster.
- **Q:** Can multiple Oracle Database instances be configured on a single node of an HACMP cluster?
  A: Yes. Multiple database instances can be configured on a single node of an HACMP cluster if you have separate service IP addresses over which the listeners for every Oracle Database will run. Hence, you can have separate resource groups which will own each Oracle instance. This configuration will be useful if there is a failure of a single Oracle Database instance on one node to be failed over to another node without disturbing other running Oracle instances.
- **Q:** Can HACMP be configured in an active-passive configuration?
  A: Yes. For an active-passive cluster configuration, do not configure any service IP on the passive node. Also, for all the resource groups on the active node, specify the passive node as the next node in the priority to take over in the event of failure of the active node.
- **Q:** Can a file system mounted over the NFS protocol be used for disk heartbeat?
  A: No. The volume mounted over the NFS protocol is a file system for AIX, and since a disk device is required for enhanced concurrent capable volume group for disk heartbeat, the NFS file system cannot be used for configuring the disk heartbeat. One needs to provide a disk device to AIX hosts over FCP or the iSCSI protocol.
- **Q:** Which are the HACMP log files available for troubleshooting?
  A: Following are log files which can be used for troubleshooting:
  - /var/hacmp/clverify/current/<nodename>/* contains logs from current execution of cluster verification.
  - /var/hacmp/clverify/pass/<nodename>/* contains logs from the last time verification passed.
  - /var/hacmp/clverify/fail/<nodename>/* contains logs from the last time verification failed.
  - /tmp/hacmp.out file records the output generated by the event scripts of HACMP as they execute.
  - /tmp/clstmgr.debug file contains time-stamped messages generated by HACMP clstrmgrES activity.
  - /tmp/cspoc.log file contains messages generated by HACMP C-SPOC commands.
  - /usr/es/adm/cluster.log file is the main HACMP log file. HACMP error messages and messages about HACMP related events are appended to this log.
  - /var/adm/clavan.log file keeps track of when each application that is managed by HACMP is started or stopped and when the node stops on which an application is running.
  - /var/hacmp/clcomd/clcomd.log file contains messages generated by HACMP cluster communication daemon.
  - /var/ha/log/grpsvcs.<filename> file tracks the execution of internal activities of the grpsvcs daemon.
  - /var/ha/log/topsvcs.<filename> file tracks the execution of internal activities of the topsvcs daemon.
  - /var/ha/log/grpglsm file tracks the execution of internal activities of grpglsm daemon.

ACKNOWLEDGEMENTS

The author would like to thank the following individuals for their contribution to the testing process and technical report: NetApp, Inc.

- Uday Shet
- Michelle Nguyen
- Sankar Bose

www.netapp.com