




Allasso Consulting, LLC

An IBM Business Partner

Configuring HACMP



*A overview of the steps to build an
HACMP Cluster of two or more nodes*

Author: Chuck Shoecraft

Configuring HACMP V5

Introduction

HACMP is organized by three concepts:

1. **Topology**, which is the cluster from a network view. Topology consists of the cluster itself, the nodes that make up the cluster and the networks that connect them. The networks are of two types, IP networks and non-IP networks. IP networks are typically multi-drop networks and use the standard TCP/IP protocols. Non-IP networks are point to point networks which use a protocol which is proprietary to HACMP. The purpose of the non-IP networks is to eliminate the TCP/IP software in a node from becoming a single point of failure.
2. **Resources**, which are the things being made highly available and are normally used to provide services to the client population of systems. Some examples of resources include, the Service IP Label/Address which clients will use to access a service, File systems and the Volume Groups that contain them, Application Servers which are the mechanism HACMP uses to identify and locate the start and stop scripts which control the application.
3. **Resource Groups**, which are the method of managing all the resources associated with an application, and identifying the nodes where the application can run and the rules used to identify the priority for selecting the node for a given situation.

These three concepts and the preparations necessary to support them will provide the outline for the construction of an HACMP Cluster in this paper.

Prior to configuration of the HACMP cluster there is some preparatory work required such as, installing the hardware, connecting and configuring the TCP/IP networks, and connecting the storage devices. Please refer to the paper on cluster design for details on these subjects.

This paper will address cluster configuration in four steps:

1. Preparation and configuration of the hardware, AIX and the applications..
2. Topology configuration.
3. Resource definition.
4. Resource Group configuration.

Configuring HACMP V5

Hardware and Software Preparation

Prior to ordering hardware for an HACMP Cluster you should determine which applications are going to be made highly available, what if any relationships they have to each other, and what resource requirements each has. The hardware should then be selected which will provide the needed resources and connectivity.

Each system must be installed individually including the installation of and basic configuration of AIX. The IP network connections should be completed, as should any serial network connections you may have planned. And the connections to the shared storage media completed.

You should have planned for two subnets to be used for heartbeats. Each node should have two connections to each IP network and each connection on the node should be assigned to a different subnet. This will provide you with the IP addresses to set into your IP adapters at startup. Use **smit tcpip → Minimum Configuration & Startup** menu to set the IP address, netmask and hostname for the first interface on each cluster node. Use **smit tcpip → Further Configuration → Network Interfaces → Network Interface Selection → Change / Show Characteristics of a Network Interface** to set the IP address and netmask for the remaining interfaces. Test connectivity between the addresses in each subnet for all machines using a ping command.

Now create an `/etc/hosts` file on one of the systems that contains all of the IP addresses used in the cluster. An example hosts file for a 2 node cluster with 1 IP network is shown on the next page. Subnets 192.168.1 and 192.168.2 are heartbeat subnets. Subnet 192.168.50 is used for persistent aliases and service addresses.

When the host file is complete transfer it to the other nodes using ftp or rcp. Test name resolution by issuing ping commands using names instead of IP addresses. Do this on each node.

If your application does not require DNS name resolution remove the `resolve.conf` file from each node. If the application does require DNS name resolution you will need to use one of the following methods to assure that local name resolution occurs before other methods. Either create a `/etc/netsvc.conf` file with the line `hosts=local,bind` or create the system variable `NSORDER=local,bind`. If multiple resolution order methods are used the `NSORDER` method will be the last checked and therefore the deciding method.

Install the HACMP code, from the distribution CD, on all nodes. Check the README files and other prerequisite information for system requirements. Reboot each system after installing HACMP. After HACMP 5.4 has been installed and rebooted this will no longer be required.

Configuring HACMP V5

Sample /etc/hosts file

```
# Hosts file used for local name resolution in a two node HACMP Cluster
# with one IP network.
127.0.0.1 loopback localhost
# <Cluster Name>
# Netmask = 255.255.255.0 (use a mask that will match your environment.)
# Node 1 addresses
192.168.1.1 node1-if1
192.168.2.1 node1-if2
192.168.50.1 node1-per
# Node 2 addresses
192.168.1.2 node2-if1
192.168.2.2 node2-if2
192.168.50.2 node2-per
# Service Labels/Addresses
192.168.50.15 Application-1-service
```

Configuring HACMP V5

Make the connections for the shared volume groups to all systems that will need access to them and run `cfgmgr` on one system. Use the `lspv` command to determine if the shared disks have pvid's. If they do not have them to each disk with the command `chdev -l hdiskX -a pv=es`. Run `cfgmgr` on each of the other nodes the disks are attached to.

It is desirable to install and test the applications to be made highly available prior to configuring the HACMP Cluster. Install the applications on the nodes that will be primary to them, create and populate the Volume Groups and file systems required for the applications using the shared disks and test the running of the application. Determine what is required to start and stop the applications and write scripts for each operation that will completely start or stop the application from a single command (script call). The start and stop scripts should be in directories on local disks. Once tested copy the scripts to the same exact location on each node where the application will run.

If the application was installed on a local volume group install it on the other node(s) in the same location. Import the shared volume group(s) and test the application on the other node(s).

Topology Configuration

Select a node for your cluster console. Use this node for all cluster configuration actions.

Starting with `smitty hacmp` select the **Initialization and Standard Configuration** menu.

Select **Add Nodes to a cluster** and enter the cluster name in the data block, move the cursor down and use F4 to select an interface to each node you want to configure.. Press enter and wait for the OK prompt.

If you are configuring more than two nodes you may find it easier to start with a two node cluster and add additional nodes after you have a running two node cluster.

Using the F3 key move to the main HACMP menu, then select **Extended Configuration**.

Next select **Extended Topology Configuration → Configure HACMP Communications Interfaces/Devices → Add Communications Interfaces/Devices**

At this point you have to choose between interfaces and devices. Devices means non-IP networks. Select devices and then select pairs of interfaces to create your non-IP networks. You must do this one network at a time.

Configuring HACMP V5

If your /etc/hosts file has interfaces identified that are not included in your cluster design you may have IP networks that you did not want to appear in your cluster configuration. You can delete these networks using the **Extended Configuration → Extended Topology Configuration → Remove a Network from the HACMP Cluster** selection.

Synchronize the cluster using **smitty hacmp → Extended Configuration → Extended Verification and Synchronization** and select interactive at the correct errors found in verification line. Synchronize the cluster.

Run the `cltopinfo` command to display the topology of the cluster you have just created.

Resources

Next you will identify all of the resources you want HACMP to make highly available. For now you will just attach a label to each resource so that HACMP will be able to use it later in a Resource Group. There are three primary resources that most resource groups have; an **application server**, a **service label/address**, and a **volume group and file systems**. Take the Extended Resources Configuration path.

smitty hacmp → Extended Configuration → Extended Resource Configuration → HACMP Extended Resources Configuration → continue with 1 & 2 below.

1. Application Server - Select **Configure HACMP Applications → Configure HACMP Application Servers → Add an Application Server**
 - a. Enter a name for the application server.
 - b. Enter the full path name of the start script for the application
 - c. Enter the full path name of the stop script for the application
 - d. Press the enter key
 - e. Press F3 three times
2. Service Label/Address - Select **→ Configure HACMP Service IP Labels/Addresses → Add a Service IP Label/Address**
 - a. Select **Configure on Multiple Nodes** press enter

Configuring HACMP V5

- b. Select the network where the Service Address will be active – press enter
- c. Press **F4**
- d. Select the Service Label for the application - press enter - press enter
- e. Press F3 three times.

Resource Groups

The next step is to create the Resource Group(s). This will include identifying the nodes the individual resource group can become active on and the runtime policies that the resource group will follow. If a resource group is going to be highly available it must have at least two nodes identified. The sequence, from left to right, in the input area for the nodes will identify their order of priority.

Run time policies must be chosen in three areas; Startup, Fallover, and Fall Back.

There are four start up policies to choose from:

1. Online On Home Node Only
2. Online On First Available Node
3. Online Using Distribution Policy
4. Online On All available Nodes

There are three fallover policies to choose from

1. Fallover To Next Priority Node In The List
2. Fallover Using Dynamic Node Priority
3. Bring Offline (On Error Node Only)

There are two fallback policies to choose from

1. Fallback To Higher Priority Node In The List
2. Never Fallback

Configuring HACMP V5

The default policies are:

- Online On Home Node Only
- Fallover To Next Priority Node In The List
- Fallback To Higher Priority Node In The List

Select **HACMP Extended Resource Group Configuration → Add a Resource Group**

Enter the Resource Group Name – **Move the cursor to the Participating Nodes** line and press **F4**. If you want to use the default runtime policies press enter. Otherwise select any policy you want to change, press **F4** and select the policy you want to use.. When all policies have been chosen, press enter. After the OK prompt press **F3** twice, and select **Change/Show Resources and Attributes for a Resource Group**

Select the Resource Group you are going to add resources to.

Press **F4** and select the **Service Label** you defined earlier.

Move the cursor to the application server line, press **F4** and select the **Application Server** you defined earlier.

Move the cursor to the Volume Groups line, press **F4** and select the Volume Group(s) for the application.

Press **F3** repeatedly until you are at the Extended Configuration Menu and select **Extended Verification and Synchronization** and select interactive at the correct errors found in verification line. Synchronize the cluster.

Press enter and wait for the ok prompt.

If no errors occurred and you have an ok prompt, you can press **F10** to exit smit and start the cluster. **smitty clstart**. Select the nodes you want to start, change **Startup Cluster Information Daemon** to true and press enter.

Configuring HACMP V5

Conclusion

Using the instructions in this paper you should be able to configure a two node cluster supporting a single application. However you will not have taken advantage of many of the newer features of HACMP. Tools like application monitoring or fallover for reasons other than node, network adapter or network failures, and a number of administrative tools that will monitor the health of the cluster and report changes that could affect that health.

Each of those items can be found in the HACMP Documentation, and is discussed in IBM training classes for Systems Administrators who will be responsible for HACMP clusters. Both the manuals and the training classes can be accessed through the IBM web pages.

Manuals can be found at:

<http://publib.boulder.ibm.com/infocenter/clresctr/vrx/index.jsp?topic=/com.ibm.cluster.infocenter.doc/library.html>

Training Classes can be found at:

<http://www-304.ibm.com/jct03001c/services/learning/ites.wss/us/en?pageType=page&c=a0000045>

3/3/2007