Considerations in Upgrading HACMP/6000 to 2.1

Special Notices

Information in this document is correct to the best of our knowledge at the time of this writing. Please send feedback by fax to "AIXServ Information" at (512) 823-4009.

Please use this information with care. IBM will not be responsible for damages of any kind resulting from its use. The use of this information is the sole responsibility of the customer and depends on the customer's ability to evaluate and integrate this information into the customer's operational environment.

> **Note:** The information in this document has **not** been verified for AIX 4.1.

About this Document

This document describes upgrade considerations for migrating from HACMP/6000 1.1 or 1.2 to HACMP/6000 2.1. Please note that the process described here applies to upgrades from either 1.1 or 1.2. For the rest of this document, these two releases will be generically referred to as 1.n.

Important: HACMP/6000 2.1 and the conversion utility described below REQUIRE a base operating system level of at least AIX 3.2.4. If your operating system level is less than 3.2.4, you must upgrade AIX before installing/upgrading to HACMP/6000 2.1.

Tasks Before Upgrading

Before upgrading your 1.n installation to HACMP/6000 2.1, be sure to do the following:

1. For each node, archive /usr/sbin/cluster to a readily accessible place on disk. This allows you to easily retrieve and compare localized script and configuration files. Also, capture and store the output from /usr/sbin/cluster/clgetenv and /usr/sbin/cluster/cllsif, as a snapshot of the node's view of its HACMP world.
2. If your 1.n installation is applied but not committed, you must commit it now. This makes it possible for 2.1 to be installed over it.
3. Do a mksysb of each node.

Configuration Information Moved to the ODM

Under 1.n, cluster configuration information was stored in /usr/sbin/cluster/cluster.cf, a stanza-based ASCII file. The easiest way to access it was through the /usr/sbin/cluster/cllsif command. Changes were made via "smit hacmp".

Node information was accessible from a link in /etc/objrepos that pointed to /usr/lib/obj/hacmp6000, through /usr/sbin/cluster/clgetenv.

HACMP/6000 2.1 moves all of this configuration information into the ODM, creating seven new object classes in /etc/objrepos. These are:

HACMPadapter
HACMPcluster
HACMPcommand
HACMPevent
HACMPnetwork
HACMPresource
HACMPserver

HACMP/6000 2.1 also provides you with /usr/lpp/cluster/tools/clconvert, a utility to convert 1.n configuration information into the appropriate 2.1 ODM object classes.

Automatic Conversion During Upgrade

Important: Installing 2.1 over a committed 1.n installation causes an AUTOMATIC UPDATE of your 1.n configuration information to 2.1 format. cluster.cf and node info (/etc/objrepos/hacmp6000) are used to populate the appropriate 2.1 object classes. You can back out an automatic conversion with the standard SMIT interface used to modify/delete configuration information.

The conversion fails if either /usr/sbin/cluster/cluster.cf.OLD or /etc/objrepos/hacmp6000 is not available. If there is already configuration information on one of the 2.1 ODM object classes, it is overwritten during the installation.

If the original cluster.cf is correct and /etc/objrepos/hacmp6000 exists, the new ODM classes are updated. Verify this by running "clgetres -A" for resource information and "cllsif" for cluster topology information.

Also, the following output should appear both in your smit.log and on the screen:

```
Adding new HACMPcommand ODM entries.
Adding new HACMPevent ODM entries.
clconvert successful.
```

Note: If you do not want to automatically run the conversion tool, rename /usr/sbin/cluster/cluster.cf and /etc/objrepos/hacmp6000 BEFORE you install 2.1.

Note: /etc/objrepos/hacmp6000 is a symbolic link that points to /usr/lib/objrepos/hacmp6000.

Files Saved

Now you can install your 2.1 release as described in the 2.1 documentation. The following files are saved with a .OLD extension:

```
clinfo.rc.OLD
cluster.cf.OLD
netchng.rc.OLD
node.servers.OLD
node.vars.OLD
topchnc.rc.OLD
```

Any customized subdirectories (/usr/sbin/cluster/local, for example) are left untouched by the installation procedure.

Use the content of the old scripts to adjust and customize the new 2.1 scripts appropriately. The new scripts are discussed in the section “New Scripts” of this document. Especially important to review is the information about node.servers and topchng.rc.

New Scripts

With 2.1, the concept of cluster "events" is introduced. A cluster event is a change within the cluster that causes the clstrmgr to take an appropriate action, such as swap the IP address of a communication adapter. These "events" replace the scripts (topchng,rc, node.servers, etc.) that clstrmgr called in 1.n.

The new events, or scripts, are located in /usr/sbin/cluster/samples. Here is a listing of that directory:

```
acquire_service_addr       node_down_remote
acquire_takeover_addr      node_down_remote_complete
clinfo.rc                  node_up*
config_too_long*           node_up_complete*
event_error*               node_up_local
fail_standby*              node_up_local_complete
get_disk_vg_fs             node_up_remote
join_standby*              node_up_remote_complete
network_down*              release_service_addr
network_down_complete*     release_takeover_addr
network_up*                release_vg_fs
network_up_complete*       start_server
node_down*                 stop_server
node_down_complete*        swap_adapter*
node_down_local            swap_adapter_complete*
node_down_local_complete   unstable_too_long*
```

Those in the above list that are followed by an asterisk are "events", scripts that are called directly by the clstrmgr in response to a change in cluster status. The other scripts are "sub-events", called by event scripts.

For example, when clstrmgr is started on a node, the following sequence of events and sub-events occurs on that node:

```
clstrmgr-->
          node_up-->
                    node_up_local
                         |
                         clstrmgr-->
                             node_up_complete-->
                                        node_up_local_complete
```

These event scripts replace netchng.rc and topchng.rc as the clstrmgr's response to a change in cluster status. Sub-events can be compared to the functions within topchng.rc and netchng.rc.

Note: clinfo.rc remains virtually the same in function and intent as in previous versions. It is neither an event nor a sub-event but a script that runs on clients to flush a stale arp cache.

Application Servers Replace node.servers

In 1.n releases, local applications were typically started out of node.servers, which was called by topchng.rc. However, in 2.1 applications they are viewed as resources, in much the same way that file systems or IP addresses are.

For example, let's say that your 1.n release called /usr/sbin/cluster/local/start_db_server out of node.servers on a node up. And /usr/sbin/cluster/local/stop_db_server stops the same server on a node down. To achieve the same functionality with 2.1, do the following:

1. Using SMIT/hacmp, register an application server.
2. Give it the logical name "db_app_server".
3. On the local node, define it as an "owned" resource.
4. Cite the path as /usr/sbin/cluster/local/start_db_server for the "Start Script" field, /usr/sbin/cluster/local/stop_db_server for the "Stop Script" field.
5. On the takeover node, do the same, but cite it as a "takeover" resource.
6. Be sure that each node's version of this resource is entered and configured correctly for that node, i.e., paths, local variables, hostname considerations, etc.

In the example in the last section, the last thing that happened was the sub-event node_up_local_complete. This sub-event calls "clcallev start_server $APPLICATION", where $APPLICATION corresponds to the logical name "db_app_server".

clcallev goes into the HACMP ODM classes, finds that the logical "db_app_server" is mapped to /usr/sbin/cluster/local/start_db_server on this node, and runs that script.

Reader's Comments

Please fax this form to (512) 823-4009, attention "AIXServ Information". You may also e-mail comments to: elizabet@austin.ibm.com. These comments should include the same customer information requested below.

Use this form to tell us what you think about this document. If you have found errors in it, or if you want to express your opinion about it (such as organization, subject matter, appearance) or make suggestions for improvement, this is the form to use.

If you need technical assistance, contact your local branch office, point of sale, or 1-800-CALL-AIX (for information about support offerings). These services may be billable. Faxes on a variety of subjects may be ordered free of charge from 1-800-IBM-4FAX. Outside the U.S. call 415-855-4329 using a fax machine phone.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Note: If you have a problem report or item number, supplying that number may help us determine why a procedure did or did not work in your specific situation.

Problem Report or Item #: Branch Office or Customer #:

Be sure to print your name and fax number below if you would like a reply:

Name: Fax Number:

END OF DOCUMENT (hacmp21.upgrade.ha_, 4FAX# 2682)