

IBM PowerHA SystemMirror Rapid Deploy Cluster Worksheets

IBM Redbooks Solution Guide

The IBM® PowerHA® SystemMirror® V7 Cluster software is the next evolution of IBM AIX® clustering. To provide tighter integration with the AIX operating system, a new kernel level layer was developed that is called Cluster Aware AIX (CAA). The Cluster software uses this new foundation for its heartbeat and message communication. Running at this kernel level ensures that the cluster communication receives top priority and is not affected in the event of a memory leak or rogue application consuming systems resources.

This redesign enables health monitoring across all of the network interfaces along with the ability to react to the loss of `rootvg` when booting from external SAN storage. In addition, new target mode capabilities in the Fibre Channel adapters allow for Storage Framework communication for health monitoring over the storage area network (SAN).

This IBM Redbooks® Solution Guide is intended to assist with the rapid deployment of a PowerHA SystemMirror V7 cluster by using the new `clmgr` command-line interface. It also provides examples of common administrative tasks, sample configuration files, and useful logs.

Figure 1 shows the IBM PowerHA SystemMirror V7 topology to show the components that are mentioned in this document.

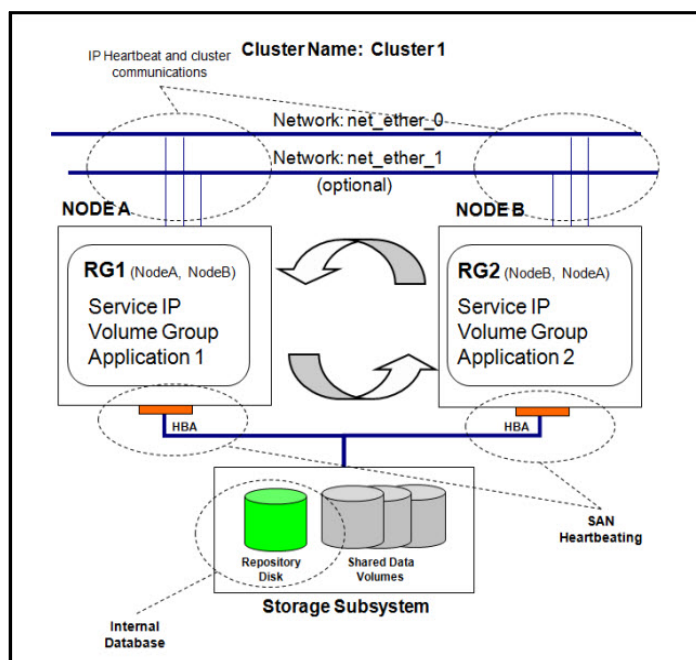


Figure 1. IBM PowerHA SystemMirror V7 topology

Did you know?

The new `clmgr` utility in PowerHA SystemMirror V7 helps expedite the deployment of the cluster configuration. PowerHA SystemMirror V7 clusters may be created entirely from the new command-line interface (CLI).

Business value

The objective behind a high availability solution is to provide near-continuous application availability for both planned and unplanned outages. Business-critical applications are configured into a cluster, which typically has at least two systems (or nodes), and the cluster monitors the critical resources for changes that might indicate a failure, a pending failure, or a possible configuration change. The cluster is monitored for health and for configuration changes that must be made consistent across the cluster. A cluster can be configured for disaster recovery (DR) by providing clustering capabilities across geographically dispersed locations. As a preferred practice, data centers conduct periodic disaster recovery tests to demonstrate compliance with corporate policies. Compliance tests can be both operationally expensive (tying up critical resources during the test) and impact the business. Many companies simply cannot afford to have their IT operations unavailable for an extended DR test and therefore implement a cluster to simplify and shorten the test.

All high availability products provide the same basic functions for monitoring and recovery of mission-critical applications. The strength of PowerHA for AIX lies in its tight integration with AIX and IBM Power Systems™ hardware. General-purpose solutions that run on many different hardware platforms and operating systems can offer only the least common denominator set of features. No other product can provide the features, performance, and reliability of PowerHA on IBM Power Systems servers and AIX.

PowerHA SystemMirror Standard Edition and PowerHA SystemMirror Enterprise Edition with Cluster Aware AIX (CAA), kernel-based health management, Graphical Management, IBM HyperSwap®, and other integrated features, provide a robust high availability disaster recovery (HADR) environment that is focused on ease of implementation and ease of use.

Solution overview

Although the new `clmgr` utility in PowerHA SystemMirror V7 helps expedite the deployment of the cluster configuration, here is a cluster resource checklist that you might find useful when you plan to implement an IBM PowerHA SystemMirror high availability clustering solution:

- IP address planning:
 - Request IPs (number of boot/base, persistent, and service IPs).
 - Register DNS names.
 - Update configuration files: `/etc/hosts` and `/etc/cluster/rhosts`.
 - Hard set IPs on the network interfaces.
- Shared storage planning:
 - Determine space requirements (number of data LUNs, and the size of the cluster repository disk).
 - Identify driver and multipath requirements.
 - Determine the LUN mappings / SAN zoning.

- o Create / import shared volume group, logical volume, and file system information. Use unique names for resources that are imported across cluster members.
- o (optional) Identify and implement requirements for SANCOMM.
- Highly available applications planning:
 - o Identify installation location and space requirements.
 - o Identify user and permission settings.
 - o Test and deploy application start / stop scripts.
 - o (optional) Test and deploy application monitoring scripts.
- PowerHA SystemMirror cluster deployment:
 - o Identify and install AIX level requirements (including CAA and RSCT packages) on all nodes.
 - o Identify and install PowerHA SystemMirror code level on all nodes.
 - o Reboot LPARs to pick up kernel bos updates.
- From node 1:
 - o Define a cluster name.
 - o Define a cluster Repository Disk.
 - o Define a multicast address (automatic or manual).
 - o Define node names.
 - o Define networks.
 - o Define interfaces.
 - o Define application controllers.
 - o Define service IPs.
 - o Define resource groups.
 - o Define resources to the resource group (RG)
 - o Verify / Synchronize the cluster.
 - o Start cluster services on all nodes.
- Fallover testing: Graceful stop with Takeover and RG moves (soft) versus `reboot -q` (hard)
- Monitor the environment through different available tools.

Prepare the following configuration files:

- `/etc/hosts`

The contents of this file should include all of the cluster IP addresses and their corresponding IP labels, as it is preferable to have the cluster resolve locally and then revert to DNS if necessary.

- `/etc/cluster/rhosts`

Populate the file on both nodes and refresh the cluster communication daemon (by running `refresh -s clcomd`). Explicitly defined cluster IPs in each line help avoid name resolution issues. Ensure that only valid, accessible cluster IPs are defined in this file.

- `/usr/es/sbin/cluster/netmon.cf`

This file is used by the cluster in single adapter networks to attempt to determine adapter status in the event of a failure. Virtualized environments should deploy this file to point to default gateways or IPs outside of the physical frame to validate outside connectivity.

Prepare and check the following IP addresses:

- Multicast address (automatically or manually assigned)

The cluster heartbeat on PowerHA SystemMirror V7 clusters uses IP multicasting and by default assigns a multicast address during the cluster creation process. It attempts to avoid duplication across clusters by defining an address that is based on the first IP that it detects on your network interfaces. For example, en0 - 9.10.10.1 base IP results in a 228.10.10.1 multicast address. If you want to define your own multicast address, you may do so during that portion of the cluster configuration. PowerHA SystemMirror V7.1.3 makes this multicast requirement optional.

- Base IP addresses

Every adapter in AIX typically has an IP address on it stored in the object data manager (ODM) and set to come online during the system boot sequence. These adapters can be defined within the cluster definitions as base / boot adapters if they are to be within a PowerHA network. CAA attempts to use all interfaces within the LPAR unless the administrator has explicitly defined them in a PowerHA private network. VLANs that have interfaces that host a potential service IP must have IP multicasting enabled, otherwise CAA considers these interfaces down and never attempts to acquire the service IP alias.

- Persistent IPs

This is a cluster node-specific alias that is available on system boot whether PowerHA services are running or not. These IPs can be used as administrative IPs for each node or as IPs to hold the route for the routable subnet in the event of a cluster failover. For some time now, PowerHA has allowed single adapter networks to define the base/boot IP and service IPs on the same routable subnet, so the need for a persistent IP is not as prevalent as in earlier releases, so they are not typically required.

- Service IPs

Any defined service IP address aliases are managed by the cluster if they are defined within a resource group. Where the resource group and its corresponding resources are being hosted determines the location of the service IP alias or aliases.

Prepare and check the following shared disks configurations:

- CAA repository disk (size requirement: minimum 512 MB to a maximum of 460 GB)

This is a new CAA requirement that must be visible to all cluster members. The common practice is for this LUN to be defined as a standard LUN size in the environment if it is within the minimum and maximum size requirements. At the time of the first verify and synchronize operation, the cluster creates a private volume group on the device.

- Shared data volumes

All cluster managed shared volume groups must be created or converted to Enhanced Concurrent Mode and mapped and imported onto all cluster nodes. The corresponding LUNs should be defined to have no reservations set in their back-end multipath drivers. During cluster processing, the cluster manages the devices with its own disk fence registers, and allows only file systems to mount on the node hosting the resource group.

Also, check the cluster resource group policies. A resource group in the cluster configuration is a container for the different highly available resources. The different resource group start, fallover, and fallback policies should be established during the planning phase, and should be fully understood, as shown in Table 1.

Table 1. Resource group startup, fallover & fallback policies

Resource group policy	Available options
Start policy	Online on home node only Online on first available node Online using distribution policy Online on all available nodes
Fallover policy	Fall over to the next priority node in the list Fall over using dynamic node priority Bring offline
Fallback policy	Fall back to higher priority node in the list Bring offline

Prepare and check the following clustered application controller definitions:

- Start and stop scripts

The application controller scripts must be in a common path in all participating cluster members. They must also be executable by the root user. The content of the scripts does not need to match between all cluster members. However, if the content needs to match based on application requirements, then the PowerHA file collection function may be used to ensure that changes are synchronized automatically between the cluster members every 10 minutes.

- (Optional) Application monitoring scripts

The Cluster software delivers an optional application monitoring framework that can be used in any deployment. The cluster runs a `clappmon` process for every monitor that is defined on the node hosting its resource group and corresponding application controller. Any monitor scripts should be executable by root, be thoroughly tested, have proper script termination, and be in a common location on all cluster members.

Here is a useful checklist for Cluster Aware AIX (CAA) heartbeat communication:

- Repository disk

PowerHA SystemMirror V7 cluster communication requires the use of a shared LUN (repository disk) for heartbeating and to store cluster configuration information. The size requirement for the PowerHA SystemMirror V7.1.1 and V7.1.2 releases are a minimum size of 512 MB, with up to 460 GB. It is common for clients to use their standard LUN size rather than designating a volume smaller than their current data LUNs.

- IP interfaces

CAA based releases before PowerHA SystemMirror V7.1.3 require the enablement of IP multicasting on the layer 2 devices backing the network interfaces. CAA uses all the interfaces on the system by default unless they are defined to an HA private network. IP network definitions are required for the cluster to perform IP address takeover between the cluster members. The cluster does not bring a service IP alias online on an interface if multicast communication is not working (in a release relying on multicast communication) because the interface is considered unavailable.

- (Optional) Storage framework communication (SANCOMM)

SAN-based communication is an additional heartbeating option in PowerHA SystemMirror V7 clusters. If properly enabled, the storage framework communication passes heartbeats between the Fibre Channel adapters within the shared SAN environment to provide an additional heartbeat communication path. This configuration is supported only over SAS or 4 GB and 8 GB Fibre Channel adapters, and works in dedicated host-based adapters (HBAs) or virtualized adapters using VSCSI or NPIV.

On the supported HBAs, you must enable target mode on the LPARs that own the cards and ensure that the SAN zoning provides visibility between all applicable adapters in all cluster members. To do so, run the following commands:

```
chdev -l fscsi# -a dyntrk=yes -a fc_err_recov=fast_fail -P
chdev -l fcs# -a tme=yes -P (a reboot is required)
```

Note: The `-P` option is used to update only the AIX ODM when there are existing child devices on the HBAs, which is why a reboot is required for the setting to take effect.

Virtualized environments require the use a reserved Ethernet VLAN (3358) between the client LPARs and the corresponding virtual I/O (VIO) servers. A virtual Ethernet adapter must be defined on the client LPAR and on the VIO server to create a bridge that allows SAN heartbeat communication to reach the physical HBAs on the VIO servers. The virtual Ethernet adapters are not required to have an IP address defined on them.

For the storage packets to pass between cluster members that are defined across physical server frames, the SAN zoning must include all corresponding HBA worldwide port names (WWPNs). In a virtualized environment, the physical WWPN for the HBAs in each VIO server (not the client virtual WWPNs) must be defined within the same SAN zone. Review the current online documentation or recent IBM Redbooks publications (refer to the "Related information" section) for examples using this feature.

Solution architecture

With each product release, some of the most common questions that arise include the following ones:

- Which release should you install/use?
- What are the key features that make it worthwhile to install the latest release?

The complete list of new features is outlined in the product announcement letters, documentation, and IBM Redbooks publications. However, some key features that might impact which release to deploy are shown in Table 2.

Table 2. Differences between PowerHA SystemMirror V7 releases

Version considerations	Available updates in the different Version 7 releases
Minimum operating system requirements	If the environment is not running the minimum code levels that are outlined in this table, you cannot install that release. The decision about which PowerHA SystemMirror release to implement often revolves around the certified and installed AIX levels for the environment in question.
Heartbeating protocol	Early PowerHA SystemMirror V7 clusters changed the cluster's communication protocol to use IP multicasting and made it a fixed requirement for communication between the network interfaces. PowerHA SystemMirror V7.1.3 relaxes that requirement by using Unicast by default and optionally allowing for IP Multicast communication.
Host name changes	CAA initially did not accept a system host name change during failover operations without having to re-create the cluster. PowerHA SystemMirror V7.1.3 ignores temporary host name changes by default to provide greater flexibility.
Site-specific IPs for IP address takeover (IPAT) in local clusters	Before PowerHA SystemMirror V7 releases, a client could use the Standard Edition of the product, define sites, and use site-specific IPs for IPAT in environments where the local cluster members were in separate IP network segments. PowerHA SystemMirror V7.1.2 and later can define site definitions in a stretched cluster topology, which brings back the capability for PowerHA SystemMirror V7 clusters that are using only the Standard Edition but require the use of site-specific IPs.
Application scripts (Smart Assistants)	These configuration wizards detect installed applications and generate a cluster configuration with start/stop/monitor scripts. The most current PowerHA SystemMirror release typically supports the most recent application versions.
Disaster recovery (DR) - Automation requirements	For Cluster Aware AIX (CAA) based PowerHA SystemMirror releases, the Enterprise Edition was not made available until Version 7.1.2. For integration and automated handling of an IP or storage level replication offering, an Enterprise Edition release is required.
Disaster recovery (DR) - Tie breaker support	The ability to define an external LUN in to the cluster with different Split Merge policies was introduced in PowerHA SystemMirror V7.1.2. This feature enables the support of SCSI reservations to enforce which site remains active in the event of a loss of communication.
Disaster recovery (DR) - Failover user confirmation	PowerHA SystemMirror V7.1.3 introduced the ability to define "MANUAL" Split Merge policies that enable PowerHA administrators to control the behavior of a failover in an environment using site definitions.

Here are some additional highlights:

- CAA interface monitoring enhancements with the `clras` command (PowerHA SystemMirror V7.1.2)
- CLI enhancements to `clmgr` (PowerHA V7.1.2 | PowerHA SystemMirror V7.1.3)
- Flexibility to suspend CAA daemons to release the disks during multipath driver updates (PowerHA SystemMirror V7.1.3),
- DS8800 HyperSwap single node, ACTIVE/ACTIVE, and HyperSwap support in an LPM capable environment (PowerHA SystemMirror V7.1.3)

Usage scenarios

PowerHA SystemMirror V7 clusters may be created entirely from the new CLI. In this example, IPs already are appended to the `/etc/hosts` file, the volume group already is imported in to all cluster members, and the application scripts already are written and propagated to the common `/usr/local/hascrpts` directory in each of the cluster nodes. Resource group policies in this example are set to the most commonly used policies of Online on Home Node Only (this is the default in the command, so its input is not required), Fallover to the Next Available Node, and Never Fallback. The syntax options in the example may be modified to include additional cluster features.

The following CLI rapid deployment instructions create a basic two node cluster. Table 3 shows the network parameters that used for the two-node cluster.

Table 3. Network parameters for the two-node cluster

Network	Label	Function	Interface	Node
net_ether_01	nodea_base1	boot	en0	nodea
	nodeb_base1	boot	en0	nodeb
	sharedIP	service	alias	shared

Table 4 shows the resource group names that are used for the two-node cluster configuration.

Table 4. Resource group names

Resource group names	DB_app1_rg
Startup policy	Online on Home Node Only
Fallover policy	Fallover to the Next Priority Node
Fallback policy	Never Fallback
Participating nodes	nodea nodeb
Service IP labels	sharedIP
Volume groups	sharedvg
Application controllers	DB_app1

To create the two-node cluster, complete the following steps:

1. Create a cluster by running the following command: `clmgr add cluster SampleCluster repository=hdisk10 nodes=nodea.dfw.ibm.com, nodeb.dfw.ibm.com`
2. Add a service IP by running the following command: `clmgr add service_ip sharedIP network=net_ether_01`
3. Define an application controller by running the following command: `clmgr add application_controller DB_app1 startscript="/usr/local/hascripts/DB_app_start.sh" stopscript="/usr/local/hascripts/DB_app_stop.sh"`
4. Create a resource group by running the following command: `clmgr add rg DB_app1_rg nodes=nodea.dfw.ibm.com, nodeb.dfw.ibm.com startup=ohn fallback=nfb service_label=sharedIP volume_group=sharedvg application=DB_app1`
5. Verify and synchronize cluster by running the following command: `clmgr sync cluster`

Note: The CAA private volume group that is created on the Repository Disk shows up only after the first time that the cluster definitions are synchronized. This is a "hands-off" volume group and should not be modified, mirrored, or extended through AIX LVM.

Common administrative tasks, sample configurations, and logs

This section provides commands and menu commands that are used during common administrative tasks while implementing and managing the PowerHA SystemMirror cluster. The command syntax often may allow for different options to achieve the same behavior. In particular, the `clmgr` command is loaded with useful tips.

- PowerHA SystemMirror SMIT menus:
 - `smitty sysmirror`
 - `smitty cl_admin`
- Start Cluster Services (different choices):
 - `clmgr start cluster`
 - `clmgr online node nodea`
 - `clmgr start node nodea`
 - `smitty clstart`
 - `clmgr start cluster START_CAA=yes`
- Stop Cluster Services:
 - `clmgr stop cluster`
 - `clmgr offline node nodea`
 - `clmgr stop node nodea`
 - `smitty clstop`
 - `clmgr stop cluster STOP_CAA=yes`
- Verify / Synchronize Cluster:
 - `clmgr verify cluster`
 - `clmgr sync cluster`

- Move Resource Group: (different choices)
 - `clmgr move rg rgA, rgB node=nodeA` With multiple resource groups (RGs), the move is done serially.
 - `clRGmove -g RGname -n nodeA -m`

- Add an application monitor:

```
clmgr add mon appA_mon TYPE=Custom APPLICATION=appA MONITORINTERVAL=60
FAILUREACTION=fallover STABILIZATION=300 RESTARTINTERVAL=1200
CLEANUPMETHOD=/usr/local/hascripts/appA_cleanup.sh
RESTARTMETHOD=/usr/local/hascripts/appA_restart.sh RESTARTCOUNT=3
MONITORMETHOD=/usr/local/hascripts/appA_monitor.sh
```

- Suspend / Resume Application Monitoring:

- `clmgr manage application_controller suspend test_app1`
- `clmgr resume application_controller resume test_app1`

Note: The syntax that is used in this document is intended to show some of the options within the new CLI. There are many additional granular options that are available for the commands. For an extensive list of syntax options and examples, run `clmgr -help` or `man clmgr`.

For many of the commands, there are available aliases that allow for a shorter version of the command, such as `file_system` (alias `fs`), `move` (alias `mv`), and `sync` (alias `sy`).

- Add disks in to a Volume Group:

```
clmgr modify volume_group dataVG add=hdisk3
```

- Add a Logical Volume:

```
clmgr add lv new_lvname type=jfs2 volume_group=dataVG logical_partitions=10
```

- Add a file system (previously defined lv):

```
clmgr add fs /fsname type=enhanced logical_volume=prev_created_lv
lv_for_log=INLINE inline_log_size=4
```

Note: The `clmgr` operation automatically mounts the file system and update the ODM and `/etc/filesystems` file in the other cluster nodes. If the Volume Group is already defined to a resource group, the cluster automatically manages the file system.

- Validate IP multicast traffic: (must be run on each node):

- `mping -v -r -a 228.10.10.1` (nodeA - receive flag)
- `mping -v -s -a 228.10.10.1` (nodeB - send flag)

Note: In AIX V6.1 TL9 and AIX V7.1 TL3, usability is enhanced to allow users to specify the IP address of the interface they want to `mping` on (`-b <addr>`). In previous versions, the commands report back as successful if the software could `mping` down one of the interfaces on the server.

- Display / Modify tunables:

- `clctrl -tune -L` Displays default and sets tunable values.
- `clcmd netstat -in` Displays all interfaces and IPs from all cluster nodes.
- `clcmd lspv` Displays all PVIDs and VG information from all cluster nodes.

Note: For the Cluster Aware AIX (CAA) enhanced usability, the `bos.cluster.rte` CAA package introduces the `clcmd` command. This command allows administrators to add it before their commands and collection information from all cluster nodes from a single window.

- Replace a repository disk:

```
clmgr replace repository new_disk
```

Here are some commands that provide useful information to verify, monitor, and check the PowerHA SystemMirror cluster status:

- To find the product version:

```
o halevel -s
o lslpp -l cluster.es.server.rte
o lssrc -ls clstrmgrES | grep fix
o clmgr query version
```

- To query cluster settings and status:

```
o clmgr query cluster
o clmgr -v -a name,state,raw_state query node
o lssrc -ls clstrmgrES | grep state
o clshowsrv -v
```

- To display a cluster configuration:

```
o cltopinfo
o clmgr view report basic
o cllsif | clshowres Cluster Topology | Resource Group configuration legacy commands
```

- To find the location of resources:

```
clRGinfo -p
```

- Cluster Aware AIX (CAA) commands:

```
o lscluster -c      Cluster configuration (multicast address)
o lscluster -i      Status of cluster interfaces
o lscluster -d      Cluster storage Interfaces
o lscluster -m      Cluster node configuration information
o lscluster -s      Cluster statistics
```

Note: SANCOMM is working only if `sfwcom` is visible in the `lscluster -m` output as follows:

Interface	State	Protocol	Status
dpcom	DOWN	none	RESTRICTED
en0	UP	IPv4	none
sfwcom	UP	none	none

```
# clras sancomm_status
NAME                UUID                                STATUS
nodeA.dfw.ibm.com | e9b4d6a4-5e71-11-e2-af42-00145ee726e1 | UP |
```

You can also check the sent and received storage packet counts by running the `lscluster -s` command:

```
# lscluster -s | grep storage
storage pkts sent: 168493709          storage pkts recv: 82575360
```

Here are sample configuration files that you can use while configuring and managing a PowerHA SystemMirror high availability cluster:

- `/etc/cluster/rhosts`

```
9.10.10.1
9.10.10.2
```
- `/etc/hosts`

```
127.0.0.1 loopback
# PowerHA      SystemMirror Cluster   IP Addresses
9.10.10.1     nodea.dfw.ibm.com             nodea # node A base address
9.10.10.2     nodeb.dfw.ibm.com             nodea # node B base address
9.10.10.10    shared_IP.dfw.ibm.com         shared_IP # Shared
                                                SAN Volume Controller IP address
```
- `/etc/netsvc.conf`

```
hosts= local, bind
```
- `/etc/resolv.conf`

```
nameserver 9.0.1.1
domain dfw.ibm.com
```
- `/usr/es/sbin/cluster/netmon.cf`

```
9.10.10.6
!REQD <owner> target          References /usr/sbin/rsct/samples/hats/netmon.cf
!IBQPORT <owner>              Documentation APARs: IZ01332 IZ01332
!IBQPORTONLY <owner>
```

Here are sample application controller scripts:

- `/usr/local/hascripts/appA_start.sh`

```
#!/bin/ksh
su - orastst -c "lsnrctl start"
su - tstadm -c "startsap"
exit 0
```

Basic SAP example
- `/usr/local/hascripts/appA_stop.sh`

```
#!/bin/ksh
su - tstadm -c "stopsap"
su - orastst -c "lsnrctl stop"
exit
```

Basic SAP example

- /usr/local/hascripts/appA_monitor.sh

```
#!/bin/ksh
...user provided logic ...
exit 0
```

Here are useful cluster log files for monitoring and troubleshooting the highly available cluster:

- /var/hacmp/log/hacmp.out Detailed EVENT processing


```
Aug 14 16:34:49 EVENT START: node_up nodea
:node_up [165] [[ high==high ]]
:node_up [165] version=1.10.11.32
:node_up [167] node_up_vg_fence_init
```
- /var/hacmp/adm/cluster.log High-level cluster EVENTS


```
Aug 14 16:34:49 nodea user:notice PowerHA SystemMirror for AIX: EVENT
START:
node_up nodea
Aug 14 16:34:51 nodea user:notice PowerHA SystemMirror for AIX: EVENT
COMPLETED:
node_up nodea
```
- /var/hacmp/log/clutils.log Generated by cluster utilities


```
CLMGR STARTED (9153:10254698:5177392) : Thu Aug 14 16:34:49 CET 2013
CLMGR USER (9153:10254698:5177392) : ::root:system
CLMGR COMMAND (9153:10254698:5177392) : clmgr online node nodea
CLMGR ACTUAL (9153:10254698:5177392) : start_node nodea
```
- /var/adm/ras/syslog.caa CAA logging and troubleshooting


```
Aug 14 16:34:28 nodea caa:info syslog: caa_query.c cl_get_capability
2594
There are two more capabilities that are defined at level 131072
Aug 14 16:34:49 nodea caa:info syslog: caa_query.c cl_get_capability 2594
There are two more capabilities that are defined at level 131072
```

Also, the following log files are useful to check while monitoring and troubleshooting the highly available cluster:

- /var/hacmp/clverify/clverify.log Detailed verification check output
- /var/hacmp/clcomd/clcomd.log Troubleshooting communication issues
- /var/hacmp/log/cspoc.log.long Detailed information from CSPOC
- /var/hacmp/log/clstrmgr.debug Generated by the clstrmgr daemon
- /var/hacmp/log/autoverify.log Generated by nightly verification

Supported platforms

PowerHA SystemMirror V7 is supported on the AIX V6.1 and AIX V7.1 operating systems. Here are the specific software requirements for PowerHA SystemMirror V7.1.3:

Operating system (one of the following):

- AIX V6.1 Technology Level 9 with Service Pack 1
- AIX V7.1 Technology Level 3 with Service Pack 1

The Cluster Aware AIX (CAA) kernel layer was introduced in AIX V6 TL6 and in the base AIX V7 code release. Each release has different AIX prerequisite levels to install the Cluster software, as shown in Table 5.

Table 5. PowerHA SystemMirror releases

PowerHA SystemMirror release	General announcement (GA) date	Minimum AIX level
PowerHA SystemMirror V7.1.0	September 2010	AIX V7.1 with RSCT 3.1.0.1 AIX V6.1 TL6 SP1 with RSCT 3.1.0.1
PowerHA SystemMirror V7.1.1	December 2011	AIX V7.1 TL1 SP3 with RSCT 3.1.2.0 AIX V6.1 TL7 SP3 with RSCT 3.1.2.0
PowerHA SystemMirror V7.1.2	November 2012	AIX V7.1 TL2 SP1 with RSCT 3.1.2.0 AIX V6.1 TL8 SP1 with RSCT 3.1.2.0
PowerHA SystemMirror V7.1.3	December 2013	AIX V7.1 TL3 SP1 with RSCT 3.1.5.0 AIX V6.1 TL9 SP1 with RSCT 3.1.5.0

The packages that are required for CAA functionality include the following ones:

- `bos.cluster.rte`
- `bos.ahafs`
- `bos.cluster.solid` (No longer required beyond Version 7.1.0)

Ordering information

For ordering, contact your IBM sales representative, an IBM Business Partner, or IBM Americas Call Centers at 800-IBM-CALL (Reference: RE001).

Ordering information is shown in Table 6.

Table 6. Ordering program name, program number, and version

Program name	PID number	Charge unit description
IBM PowerHA SystemMirror Enterprise Edition	5765-H37	7.1
IBM PowerHA SystemMirror Standard Edition	5765-H39	7.1

Related information

For more information, see the following documents:

- *IBM PowerHA SystemMirror 7.1.3 for AIX Best Practices and Reference Architectures*, SG24-8167
<http://www.redbooks.ibm.com/abstracts/sg248167.html>
- IBM PowerHA SystemMirror for AIX product page:
<http://www-03.ibm.com/systems/power/software/availability/aix/>
- IBM PowerHA SystemMirror for AIX V7.1.3 adds enhancement announcement letter:
http://www-01.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/6/897/ENUS213-416/index.html&lang=en&request_locale=en
- PowerHA Rapid Deployment Worksheets (online version with sample command outputs)
<https://www.ibm.com/developerworks/aix/tutorials/au-ibm-powerha-system-mirror/>
- IBM Offering Information page (announcement letters and sales manuals):
http://www.ibm.com/common/ssi/index.wss?request_locale=en

On this page, enter *IBM PowerHA SystemMirror for AIX offers solutions for high availability and disaster recovery for clusters*, select the information type, and then click **Search**. On the next page, narrow your search results by geography and language.

It is the goal of the IBM development team to continue to deliver product features that our clients demand and that provide the highest levels of resiliency. Feedback can be provided at any time by emailing hafeedbk@us.ibm.com.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service. IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you. This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

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. IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

© Copyright International Business Machines Corporation 2014. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

This document was created or updated on April 28, 2014.

Send us your comments in one of the following ways:

- Use the online **Contact us** review form found at:
ibm.com/redbooks
- Send your comments in an e-mail to:
redbook@us.ibm.com
- Mail your comments to:
IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400 U.S.A.

This document is available online at <http://www.ibm.com/redbooks/abstracts/tips1176.html> .

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the web at <http://www.ibm.com/legal/copytrade.shtml>.

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®
HyperSwap®
IBM®
Power Systems™
PowerHA®
Redbooks®
Redbooks (logo)®
SystemMirror®

The following terms are trademarks of other companies:

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

Other company, product, or service names may be trademarks or service marks of others.