HP Operations Manager

High Availability Manager

Software Version: 9.10

for the UNIX and Linux operating systems

Legal Notices

Warranty.

Hewlett-Packard makes no warranty of any kind with regard to this document, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be held liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

A copy of the specific warranty terms applicable to your Hewlett-Packard product can be obtained from your local Sales and Service Office.

Restricted Rights Legend.

Use, duplication or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c) (1) (ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013.

Hewlett-Packard Company
United States of America

Rights for non-DOD U.S. Government Departments and Agencies are as set forth in FAR 52.227-19(c)(1,2).

Copyright Notices.

Trademark Notices.

Adobe® is a trademark of Adobe Systems Incorporated.

Intel®, Itanium®, and Pentium® are trademarks of Intel Corporation in the U.S. and other countries.

Java is a registered trademark of Oracle and/or its affiliates.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

Conventions

The following typographical conventions are used in this manual:

Table 1   **Typographical Conventions**

Font	Meaning	Example
Italic	Book titles and manual page names	For more information, see the *HPOM Administrator's Reference* and the *opc(1m)* manual page.
	Emphasis	You *must* follow these steps.
	Variable that you must supply when entering a command (in angle brackets)	At the prompt, enter **rlogin *\<username\>***.
	Parameters to a function	The *oper_name* parameter returns an integer response.
Computer	Text and other items on the computer screen	The following system message displays: `Are you sure you want to remove current group?`
	Command names	Use the `grep` command ...
	Function names	Use the `opc_connect()` function to connect...
	File and directory names	Edit the `itooprc` file... `/opt/OV/bin/OpC/`
	Process names	Check to see if `opcmona` is running.
Computer Bold	Text that you enter	At the prompt, enter **ls -l**.

Table 1   **Typographical Conventions (Continued)**

Font	Meaning	Example
Keycap	Keyboard keys	Press **Return**.
	Menu name followed by a colon (:) means that you select the menu, and then the item. When the item is followed by an arrow (->), a cascading menu follows.	From the menu bar, select **Actions: Filtering -> All Active Messages**.
	Buttons in the user interface	Click **OK**.

In This Document

This document describes the High Availability Manager (HA Manager). The HA Manager enables you to switch a high availability resource group (HARG) between the nodes that make up a virtual cluster.

The following high availability and cluster terms are used in this document:

Virtual Cluster

Virtual clusters consist of independent nodes grouped together by the HA Manager, which runs on each node. Virtual clusters also include a HARG (that is, a high availability application in a virtual cluster). Applications that are integrated into a virtual cluster become high availability applications when they are active on one node in the cluster in a particular period of time. If there is a problem with the node on which the application runs, the application activity is moved to another node, which results in the application running smoothly at all times.

HARG

A high availability resource group that represents a resource defined in the "cluster world" that can be linked to an application instance. The HARG runs on a cluster and can be switched from one cluster node to another. A cluster package is usually also linked to an element from the "networking world" known as a virtual node.

Virtual Node

A virtual node is the network representation of an application package running on an HA cluster. A virtual node typically has a hostname and an IP address, it is known to the name resolution, and it can be addressed like an ordinary system.

Virtual IP Address

A virtual IP address is typically part of a HARG and can be switched from one cluster node to another cluster node.

Reference Node

A reference node, which does not belong to the HA Manager environment, is used primarily for checking the connection between the nodes. If a local node fails to make a connection to another cluster node or if the connection is lost, the local node tries to access the reference node (for example, a local DNS server) to pinpoint on which node a problem occurred. The reference node check is used to avoid the "split-brain syndrome" (that is, to find out if the network is down or if the network is up and running, but the other cluster node is down). A "split brain" means that both cluster nodes believe that the other one is down and each of them takes over the control, but the real problem lies in the network connection between the nodes and therefore they cannot communicate to each other.

The information in this document covers the following topics:

1 High Availability Manager

HA Manager and HARG Concepts

High availability is a general term used to characterize environments that represent business-critical systems protected against downtime through different redundant resources. The HA Manager is a light-weight solution that allows the configuration of an automatic failover of the virtual IP address in a server pooling setup in a similar way as in a regular failover cluster. The HA Manager is not an additional cluster software such as the HP Serviceguard or the Red Hat Cluster Suite, but it is an alternative. It represents a cluster without special hardware with redundancy and a shared disk. Therefore, the HA Manager feature enables you to do the following:

❑ Switch an IP address from one node to another node within a server pooling environment. In this case, no hardware cluster is needed and HP Operations agents and HPOM GUIs can communicate using that high availability virtual IP address.

For details, see the *HP Operations High Availability Through Server Pooling* document.

❑ Control other resources besides virtual IP addresses and make them high available.

For more information about high availability, see the *HPOM Concepts Guide*.

Before you start using the HA Manager, you should be familiar with the following concepts:

❑ **HA Manager communication**

The HA Managers that run on different nodes use a file transfer for communication because it represents a simple, reliable, and standardized way of exchanging data between the systems or the nodes.

The following two directories are used for the file transfer on each node:

- `/var/opt/OV/hamanager/input`

    This directory contains the files that the local HA Manager receives from other nodes. The HA Manager automatically reads the received files and deletes them from this directory after reading them.

- `/var/opt/OV/hamanager/output`

    This directory contains the files that the local HA Manager sends to other nodes. The HA Manager deletes the files from this directory after sending them.

❏ **HARG status and status synchronization or update**

The statuses of all nodes are stored in the `/var/opt/OV/hamanager/status` directory. For each node, there is a subdirectory that contains the statuses of all HARGs:

`/var/opt/OV/hamanager/status/<nodename>`

The status of each HARG is stored in the `harg_status.<HARG_name>` file.

The local HA Manager updates the statuses of HARGs for a local node, whereas the remote HA Manager sends the statuses of HARGs from the remote nodes.

❏ **HA Manager alive check**

For each node, an alive check is performed by checking the time of the last status update from a selected node. The alive check ensures that a virtual cluster runs smoothly by taking an immediate and adequate action in case a problem occurs on a certain node. A local node or a remote node is detected as FAULTED when the following happens:

- A local node is detected as FAULTED when a reference node is down and there is no connection to other nodes (that is, there is no status update from other nodes for some time).

- A remote node is detected as FAULTED when there is no
  connection to a selected node (that is, there is no status update
  from the selected node for some time) and when a local node can
  access a reference node, or when there is no ping reply from the
  selected node.

❑ **HARG online, offline, and monitor procedures**

When a HARG online, offline, or monitor procedure is performed, the
HA Manager runs one of the following commands on a selected node:

/opt/OV/lbin/ovharg -start *<HARG_name>*

/opt/OV/lbin/ovharg -stop *<HARG_name>*

/opt/OV/lbin/ovharg -monitor *<HARG_name>*

When starting, stopping, or monitoring the HARG, the ovharg tool
performs all the start (S), stop (K), or monitor (M) operations as
specified by a sequence number.

The return code is 0 when the action is performed successfully,
otherwise the return code is 1.

The general return code is a collection of individual return codes.
When all return codes are 0, the general return code is also 0.
However, if there is at least one return code that is 1, the general
return code is also 1.

NOTE   HARG monitoring is performed only on the node on which the HARG
is ONLINE.

❑ **HARG autostart**

A HARG is started automatically on a local node that is ONLINE in
the following cases:

- The HARG is OFFLINE and the local node is the primary node.

- The HARG is OFFLINE and the primary node is FAULTED.

- The HARG is FAULTED on some other node, but OFFLINE on the
  local node.

- The HARG is ONLINE on a FAULTED node and either the local
  node is the primary node or the primary node is FAULTED.

Keep in mind the following specifics:

- When sending the local status to the remote HA Manager fails, the remote node is not marked as FAULTED after one failure, but after a specified number of failures (the default value is 3). To specify after how many failures the remote node should be marked as FAULTED, set the MAX_COMM_PROBLEMS variable in the HA Manager configuration file to a desired value.

- The node alive timeout (that is, time during which the remote node status must be updated) can be set by using the NODE_ALIVE_TIMEOUT variable in the HA Manager configuration file (the default value is 60 seconds). If the node status is not updated in specified time, the node becomes FAULTED.

- The node sends its local status every 15 seconds (the default value). You can set another value by using the MAX_SEND_LOCAL_STATUS_TIME variable in the HA Manager configuration file.

- The HARG that is FAULTED on a local node can be automatically cleared only if it is ONLINE on some other node. This behavior is disabled by default. To enable it, set the HARG_AUTOCLEAN_FAULTED_TIME variable in the HA Manager configuration file to a number that is greater than zero. This number represents the number of seconds that elapses from the moment the HARG becomes FAULTED until the moment the autoclean is performed.

- The HA Manager configuration is automatically reloaded when it is modified.

❏ **HARG autostop**

A HARG is stopped automatically on a local node if the HARG is ONLINE on both the local node that is not the primary node and the primary node. This is because the HARG cannot run on more than one node at the same time.

❏ **Disabled HARG**

When a HARG is disabled, the HA Manager cannot control the HARG. Therefore, HARG online, offline, and monitor procedures are disabled.

HA Manager Tool

When switching a HARG between the nodes that make up a virtual cluster, use the opchamgr tool that can be found at the following location:

```
/opt/OV/bin/OpC
```

The syntax of the opchamgr tool is as follows:

```
opchamgr -daemon
        -kill
        -monitor <harg>
        -init
        -status
        -node list
        -node join <active_cluster_node>
        -node unregister <nodename>
        -harg list|add|delete|disable|enable <harg>
        -harg primary <harg> <node>
        -online <harg> [<nodename>]
        -offline <harg>
        -clear <harg> [<nodename>]
        -update
        -ping [<nodename>]
        -trace [enable|disable]
        -h|-\?|-help
```

You can use the following options with the opchamgr tool:

-daemon	Monitors the HA Manager daemon and starts it if it does not run.
-kill	Stops the HA Manager on a local node.
-monitor <harg>	Performs HARG status monitoring.
-init	Initializes a cluster on a local node.
-status	Shows the HA Manager status, nodes, and HARGs.
-node list	Lists virtual cluster nodes.
-node join <active_cluster_node>	Joins a local node to a virtual cluster that already runs on an active cluster node.

`-node unregister` `<`*`nodename`*`>`	Unregisters a node from a cluster.

`-harg`
`list|add|delete|`
`disable|enable <harg>`
	Lists, adds, deletes, disables, or enables a HARG.

`-harg primary <`*`harg`*`>` `<`*`node`*`>`	Sets a primary node for a HARG.
`-online <`*`harg`*`>` `[<`*`nodename`*`>]`	Makes a HARG online on a selected node.
`-offline <`*`harg`*`>`	Makes a HARG offline in a virtual cluster.
`-clear <`*`harg`*`>` `[<`*`nodename`*`>]`	Clears the FAULTED status of a HARG in general or on a selected node.
`-update`	Distributes local configuration to all nodes.
`-ping [<`*`nodename`*`>]`	Checks communication with a cluster node.
`-trace` `[enable|disable]`	Enables or disables tracing.
`-h|-\?|-help`	Shows the usage.

Configuring an HA Manager Environment

To configure an HA Manager environment, follow these steps:

1. On each node, create the following directories:

   mkdir -p /etc/opt/OV/hamanager

   mkdir -p /var/opt/OV/hamanager/input

   mkdir -p /var/opt/OV/hamanager/output

2. After you create the directories, create the following HA Manager configuration file on each node:

   /etc/opt/OV/hamanager/hamanager.conf

3. Set the following variables in the hamanager.conf file:

   HAMGR_COMMUNICATION_TYPE=<*selected_communication_type*>
   REFERENCE_NODE=<*reference_node*>

   For example:

   HAMGR_COMMUNICATION_TYPE=SSH
   REFERENCE_NODE=dnsnode

   In this instance, HAMGR_COMMUNICATION_TYPE defines a communication type and REFERENCE_NODE is used for testing communication.

   When selecting the communication type, you can choose between the following:

   - SSH

     Basic SSH communication channel for which you must use the following command:

     scp <*filename*> <*target_node*>:<*target_dir*>

   - REMSH

     Basic REMSH communication channel for which you must use the following command:

     rcp <*filename*> <*target_node*>:<*target_dir*>

- CUSTOM_TOOL

  When HAMGR_COMMUNICATION_TYPE is set to CUSTOM_TOOL, an additional variable must be set, CUSTOM_TRANSFER_TOOL. This variable must contain the full path to the tool performing a file transfer from a local node to a target node (that is, from /var/opt/OV/hamanager/output to /var/opt/OV/hamanager/input).

  The tool usage is as follows:

  <tool> <target_node> <file_to_send>

  For example, if the CUSTOM_TRANSFER_TOOL is set to /opt/OV/bin/OpC/utils/hamgr_transfer, the tool usage is as follows:

  /opt/OV/bin/OpC/utils/hamgr_transfer nodeA /var/opt/OV/hamanager/output/test.file

- CUSTOM_COMMAND

  When HAMGR_COMMUNICATION_TYPE is set to CUSTOM_COMMAND, an additional variable must be set, CUSTOM_TRANSFER_COMMAND. This variable must contain the full command for transferring a selected file from the /var/opt/OV/hamanager/output directory on a local node to the /var/opt/OV/hamanager/input directory on a target node.

  For example:

  **scp /var/opt/OV/hamanager/output/${FILE} \
  ${NODE}:/var/opt/OV/hamanager/input**

  In this instance, ${FILE} and ${NODE} are replaced with a filename and a nodename.

NOTE   If a non-root user is used for a file transfer, set permissions of the /var/opt/OV/hamanager/input directory to 777.

To check if the communication channel works, on each node, perform the following steps (assuming that the file is commCheck.file and the target node is nodeB):

a. Create the following file:

    /var/opt/OV/hamanager/output/commCheck.file

b. Replace $FILE with commCheck.file and $NODE with nodeB.

    For example, to do this, run the following command:

    **scp /var/opt/OV/hamanager/output/commCheck.file \
    nodeB:/var/opt/OV/hamanager/input**

c. On nodeB, the following file should exist:

    /var/opt/OV/hamanager/input/commCheck.file

d. Delete commCheck.file on both nodes.

4. Start the HA Manager daemon on both nodes by running the following command:

    /opt/OV/bin/OpC/opchamgr -daemon

5. Add the opchamgr -daemon command to crontab to enable regular HA Manager monitoring and an automatic restart if the HA Manager is stopped.

    For example, for the opchamgr -daemon command to be run every hour, add the following line to crontab:

    0 * * * * /bin/sh /opt/OV/bin/OpC/opchamgr -daemon

6. Initialize a cluster on one node by running the following command:

    /opt/OV/bin/OpC/opchamgr -init

7. After a few moments, check the status on the local node by typing the following:

    /opt/OV/bin/OpC/opchamgr -status

    The status of the local node should be ONLINE and no HARGs are configured at this point.

8. On the second node, run the following command to join this node to the cluster:

    **/opt/OV/bin/OpC/opchamgr -node \
    join <*first_active_cluster_node*>**

9. After a few moments, check the status on both nodes by typing the following:

/opt/OV/bin/OpC/opchamgr -status

Both nodes should be listed and their statuses should be ONLINE. No HARGs are configured at this point.

Configuring the HP Operations Management Server HARG in a Server Pooling Environment

In a server pooling environment, HP Operations management servers are configured identically and the role of the primary manager is assigned to a virtual interface. Managed nodes send their messages not to a physical server but to its virtual interface.

NOTE   The procedure described in this section is an example of configuring the HP Operations management server HARG in a server pooling environment.

For detailed information about server pooling, see the *HP Operations High Availability Through Server Pooling* document.

To configure the HP Operations management server HARG in a server pooling environment, first define a HARG name (for example, `hpom-server`), and then follow these steps:

1. On each node, create HARG directories and configuration files. To do this, follow these steps:

    a. Create the `/var/opt/OV/hacluster/<HARG_name>` directory by running the following command:

       `mkdir -p /var/opt/OV/hacluster/<HARG_name>`

    b. Change to this newly created directory by running the following command:

       `cd /var/opt/OV/hacluster/<HARG_name>`

IMPORTANT   Before you continue with the next step, make sure that you are familiar with the tools described in "Resources" on page 23.

c. In the /var/opt/OV/hacluster/*<HARG_name>* directory, create the following links:

   **ln -s /opt/OV/bin/OpC/utils/ha/ha_virtual_ip \
   S100_virtual_ip**

   **ln -s /opt/OV/bin/OpC/utils/ha/ha_virtual_ip \
   K200_virtual_ip**

   **ln -s /opt/OV/bin/OpC/utils/ha/ha_virtual_ip \
   M100_virtual_ip**

   **ln -s /opt/OV/bin/OpC/utils/ha/ha_mon_ovserver \
   M200_server**

d. *Outbound-only communication only:* In the /var/opt/OV/hacluster/*<HARG_name>* directory, create the following links:

   ln -s /opt/OV/bin/OpC/utils/ha/ha_start_cb S200_ovrg

   ln -s /opt/OV/bin/OpC/utils/ha/ha_stop_cb K100_ovrg

   ln -s /opt/OV/bin/OpC/utils/ha/ha_mon_cb M300_ovrg

e. Set the virtual IP resource configuration in the following file:

   /var/opt/OV/hacluster/*<HARG_name>*/
   resource_virtual_ip.conf

   The following variables must be set:

   ADDRESS=*<virtual_IP_address>*
   DEVICE=*<device_to_which_virtual_IP_is_attached>*
   NETMASK=*<netmask_address>*

   For example:

   ADDRESS=192.168.1.100
   DEVICE=e1000g0:1
   NETMASK=255.255.0.0

   The following variable is optional:

   OPTIONS=*<custom_options_to_ifconfig_command_when_activating
   _IP>*

   For details, see "Resources" on page 23.

f. *Outbound-only communication only:* Set the OVRG resource configuration in the following file:

   /var/opt/OV/hacluster/*<HARG_name>*/resource_ovrg.conf

The following variable must be set:

```
OVRG=<OV_resource_group_name_used_for_outbound-only_configu
ration>
```

For example:

```
OVRG=virt
```

For details, see "Resources" on page 23.

2. Check virtual IP activation or deactivation.

   At this point, the HARG is not under the control of the HA Manager, so you can only check if the start, stop, and monitor commands work. To do this, select a node for the check, make sure that the virtual IP and the OVRG are not active, and then follow these steps:

   a. Activate the HARG:

      /opt/OV/lbin/ovharg -start <HARG_name>

      The return code should be 0.

   b. Check the HARG status:

      /opt/OV/lbin/ovharg -monitor <HARG_name>

      The return code should be 0.

      The virtual IP should be attached to the network interface. To check this, run the following command:

      ifconfig -a

      Outbound-only communication only: The OVRG should be active. To check this, run the following command:

      **/opt/OV/bin/ovbbccb -ovrg `hostname` | grep \
      <ovrg_name>**

   c. Deactivate the HARG:

      /opt/OV/lbin/ovharg -stop <HARG_name>

      The return code should be 0.

   d. Check the HARG status:

      /opt/OV/lbin/ovharg -monitor <HARG_name>

      The return code should be 1.

The virtual IP should not be attached to the network interface. To check this, run the following command:

`ifconfig -a`

Outbound-only communication only: The OVRG should not be active. To check this, run the following command:

`/opt/OV/bin/ovbbccb -ovrg `hostname` | grep \`
`<ovrg_name>`

After you configure the HP Operations management server HARG in a server pooling environment, put the HARG under the HA Manager control. For details, see "Putting a HARG under the HA Manager Control" on page 26.

Resources

You can use the following predefined tools for managing the HP Operations management server, the Oracle database, the virtual IP, and the OVRG that can be used within the HARGs:

❏ *HP Operations management server:*

- For monitoring the HP Operations management server:

  /opt/OV/bin/OpC/utils/ha/ha_mon_ovserver

- For checking the connection between the HP Operations management server and the database server:

  /opt/OV/bin/OpC/utils/ha/ha_mon_dbconn

- For starting the HP Operations management server:

  /opt/OV/bin/OpC/utils/ha/ha_start_ovserver

- For stopping the HP Operations management server:

  /opt/OV/bin/OpC/utils/ha/ha_stop_ovserver

❏ *Oracle database:*

- For monitoring the Oracle database:

  /opt/OV/bin/OpC/utils/ha/ha_mon_oracle

- For starting the Oracle database:

  /opt/OV/bin/OpC/utils/ha/ha_start_oracle

- For stopping the Oracle database:

  `/opt/OV/bin/OpC/utils/ha/ha_stop_oracle`

❏ *OVRG:*

- For monitoring the OVRG within the Communication Broker:

  `/opt/OV/bin/OpC/utils/ha/ha_mon_cb`

- For starting the OVRG within the Communication Broker:

  `/opt/OV/bin/OpC/utils/ha/ha_start_cb`

- For stopping the OVRG within the Communication Broker:

  `/opt/OV/bin/OpC/utils/ha/ha_stop_cb`

❏ *Virtual IP:*

For monitoring, starting, and stopping the virtual IP:

`/opt/OV/bin/OpC/utils/ha/ha_virtual_ip`

General Purpose Resource

Besides the predefined tools, you can also use a general purpose resource that represents a general resource with real tools and their usage specified in the resource configuration file.

To use the general resource, the following links must be created:

```
ln -s /opt/OV/bin/OpC/utils/ha/ha_resource \
K<sequence>_<name>
```

```
ln -s /opt/OV/bin/OpC/utils/ha/ha_resource \
S<sequence>_<name>
```

```
ln -s /opt/OV/bin/OpC/utils/ha/ha_resource \
M<sequence>_<name>
```

The tool naming convention in these links is predefined and is as follows:

`<operation><sequence>_<name>`

In these instances, `K`, `S`, or `M` (that is, `<operation>`) represents the type of action the link executes (stop (`K`), start (`S`), or monitor (`M`)), `<sequence>` is the number that indicates the position in the sequence of execution, and `<name>` is the name of the operation to start.

The general resource configuration is set in the following file:

`/var/opt/OV/hacluster/<HARG_name>/resource_<name>.conf`

In this instance, *<name>* is the same as *<name>* used in
<operation><sequence>_<name>.

The following variables must be set in this file:

```
START_COMMAND=<command_for_starting_this_resource>
STOP_COMMAND=<command_for_stopping_this_resource>
MONITOR_COMMAND=<command_for_monitoring_this_resource>
```

The commands defined in this configuration file return 0 when the action
is performed successfully and 1 when the action fails.

An example of a general purpose resource is a resource for managing a
web server. If you use this resource, create the appropriate links to the
`ha_resource` tool and set the following variables in the
`resource_web.conf` file:

```
MONITOR_COMMAND=<command_to_get_web_server_status>
START_COMMAND=<command_to_start_web_server>
STOP_COMMAND=<command_to_stop_web_server>
```

Putting a HARG under the HA Manager Control

To put a HARG under the HA Manager control, follow these steps:

1. On one of the nodes, add the HARG to the HA Manager by running the following command:

   /opt/OV/bin/OpC/opchamgr -harg add <*HARG_name*>

2. Check the HA Manager status on both nodes by typing the following:

   /opt/OV/bin/OpC/opchamgr -status

   In the status output, the HARG is listed and its status is OFFLINE on both nodes. There is no primary nor active node.

3. Set a primary node for the HARG (that is, a node on which the HARG is ONLINE by default) by running the following command:

   /opt/OV/bin/OpC/opchamgr -harg primary <*HARG_name*> <*node*>

NOTE   When the primary node is set, the HARG is automatically started on the selected node. The autostart is performed within one minute.

4. Check the status of the HA Manager by running the following command:

   /opt/OV/bin/OpC/opchamgr -status

   The status output shows that the primary node is set. After a few moments, the HARG becomes ONLINE on the primary node.

When the HARG is under the control of the HA Manager, you can use the opchamgr tool to perform a HARG switchover. For details, see "Performing a HARG Switchover or Failover" on page 27.

Performing a HARG Switchover or Failover

Once a HARG is under the control of the HA Manager, use the opchamgr tool to perform a HARG switchover or failover. The difference between the switchover and the failover is that the switchover represents a controlled switch of a cluster package from one cluster node to another (for example, due to load balancing), whereas the failover represents an unplanned switch of a cluster package from one cluster node to another (for example, due to an application error).

To perform the HARG switchover, run the following command:

/opt/OV/bin/OpC/opchamgr -online *<HARG_name> <second_node>*

After you run this command, the HARG is switched between the nodes (that is, the HARG becomes ONLINE on the selected node).

Example 1-1   **Failover Scenario**

This example shows the failover scenario in which the failover occurs because the management server on the node on which the HARG is active stops running. After a few moments, the HA Manager detects that the HP Operations management server does not run and switches the HARG to a second node. The HARG status on the node on which the HP Operations management server is down is marked as FAULTED. This means that there is a problem with the HARG on this node and it is not possible to switch the HARG to this node. Make sure that you fix the problem with the HARG. For example, you can start the HP Operations management server and clear the FAULTED flag for the HARG by running the following command:

**/opt/OV/bin/OpC/opchamgr -clear *<HARG_name>* \
[*<node_where_HARG_is_FAULTED>*]**

If the HARG is FAULTED on both nodes and the FAULTED flag is cleared, the HARG is automatically started on the primary node within one minute.

To monitor the HA Manager status, you can use the following while loop:

```
# while true
# do
#    /opt/OV/bin/OpC/opchamgr -status > /HAMGR_status.txt
#    clear
#    cat /HAMGR_status.txt
#    sleep 2
# done
```

HA Manager Status

To check the HA Manager status, run the following command:

/opt/OV/bin/OpC/opchamgr -status

The following is an example output of the HA Manager status:

```
Nodes :
=======
nodeA : ONLINE
nodeB : ONLINE

HARGS :
=======
hpom-server :
 General status : ONLINE
 Active node    : nodeA
 Primary node   : nodeA
 Failover node  : nodeB

 ---------------
 nodeA : ONLINE
 nodeB : OFFLINE
```

The status values used for the node and HARG statuses are ONLINE, OFFLINE, FAULTED, and DISABLED.

Log Files and Trace Files

When troubleshooting, you can use a log file analysis that represents a useful methodology for understanding all the aspects of the HA Manager feature. To help you investigate the cause of problems, you can also use problem tracing. Trace and log files can help you pinpoint when and where problems occurred.

HARG Log and Trace Files

All errors performed during the HARG online, offline, or monitor procedure are written into the error.log file that you can find at the following location:

/var/opt/OV/hacluster/<*HARG_name*>

To enable HARG tracing on a local node, run the following command:

/opt/OV/lbin/ovharg -tracing <*HARG_name*> enable

Trace information is stored in the trace.log file that you can find at the following location:

/var/opt/OV/hacluster/<*HARG_name*>

If you want to disable HARG tracing on a local node, run the following command:

/opt/OV/lbin/ovharg -tracing <*HARG_name*> disable

HA Manager Trace File

To enable HA Manager tracing on a local node, run the following command:

/opt/OV/bin/OpC/opchamgr -trace enable

Trace information is stored in the hamgr-trace.log file that you can find at the following location:

/var/opt/OV/hamanager

To disable HA Manager tracing on a local node, run the following command:

/opt/OV/bin/OpC/opchamgr -trace disable

Data Flow

This section contains several examples of the data flow related to the HA Manager and the HARG, so that you can see how to address the issues shown by these examples.

Starting the HA Manager

When the HA Manager is started, it forwards the local status to all nodes. After some time, the HA Manager starts the HARG on a local node if the local node is set as a primary node.

Performing a HARG Switchover

This use case shows which steps are performed during a HARG switchover.

The following is the status before the HARG switchover:

```
Nodes :
=======
nodeA : ONLINE
nodeB : ONLINE

HARGS :
=======
hpom-server :
 General status : ONLINE
 Active node    : nodeA
 Primary node   : nodeA
 Failover node  : nodeB

 ---------------
 nodeA : ONLINE
 nodeB : OFFLINE
```

The HARG switchover procedure consists of the following steps:

1. To make the HARG online on a selected node, run the following command:

   /opt/OV/bin/Opc/opchamgr -online hpom-server nodeB

2. On `nodeA`, the HA Manager runs the following command:

   `/opt/OV/lbin/ovharg -stop hpom-server`

3. The `ovharg` command executes all `K*` tools in the `/var/opt/OV/hacluster/hpom-server` directory.

4. The HA Manager on `nodeA` informs the HA Manager on `nodeB` to start the HARG.

5. On `nodeB`, the HA Manager runs the following command:

   `/opt/OV/lbin/ovharg -start hpom-server`

6. The `ovharg` command executes all `S*` tools in the `/var/opt/OV/hacluster/hpom-server` directory.

7. When the `ovharg` command successfully executes all `S*` tools, the HARG becomes `ONLINE`.

The following is the status after the HARG switchover:

```
Nodes :
=======
nodeA : ONLINE
nodeB : ONLINE

HARGS :
=======
hpom-server :
 General status : ONLINE
 Active node    : nodeB
 Primary node   : nodeA
 Failover node  : nodeA

 ---------------
 nodeA : OFFLINE
 nodeB : ONLINE
```

Stopping a HARG Manually

This use case shows which steps are performed to stop a HARG manually.

The following is the status before the HARG is stopped:

```
Nodes :
=======
nodeA : ONLINE
nodeB : ONLINE

HARGS :
=======
hpom-server :
 General status : ONLINE
 Active node    : nodeA
 Primary node   : nodeA
 Failover node  : nodeB

 ---------------
 nodeA : ONLINE
 nodeB : OFFLINE
```

When stopping the HARG manually, the following steps are performed:

1. To make the HARG offline on a selected node, run the following command:

   `/opt/OV/bin/OpC/opchamgr -offline hpom-server`

2. On `nodeA`, the HA Manager runs the following command:

   `/opt/OV/lbin/ovharg -stop hpom-server`

3. The `ovharg` command executes all `K*` tools in the `/var/opt/OV/hacluster/hpom-server` directory.

The following is the status after the HARG is stopped:

```
Nodes :
=======
nodeA : ONLINE
nodeB : ONLINE

HARGS :
=======
hpom-server :
```

```
General status : OFFLINE
Active node    :
Primary node   : nodeA

----------------
nodeA : OFFLINE
nodeB : OFFLINE
```

Automatic HARG Failover

This use case shows what happens during an automatic HARG failover.

The following is the status before the HARG failover:

```
Nodes :
=======
nodeA : ONLINE
nodeB : ONLINE

HARGS :
=======
hpom-server :
 General status : ONLINE
 Active node    : nodeA
 Primary node   : nodeA
 Failover node  : nodeB

 ---------------
 nodeA : ONLINE
 nodeB : OFFLINE
```

The automatic HARG failover procedure consists of the following steps:

1. The HA Manager checks the HARG status by running the following command:

   /opt/OV/lbin/ovharg -monitor hpom-server

2. The ovharg command executes all M* tools in the /var/opt/OV/hacluster/hpom-server directory.

3. One of M* tools returns an error code 1.

4. The ovharg command returns an error code 1. This is the information for HA Manager that there is something wrong with the HARG.

5. The HA Manager switches the HARG to the failover node (that is, `nodeB`).

6. On `nodeA`, the HA Manager runs the following command:

   `/opt/OV/lbin/ovharg -stop hpom-server`

7. The `ovharg` command executes all `K*` tools in the `/var/opt/OV/hacluster/hpom-server` directory.

8. The HA Manager on `nodeA` informs the HA Manager on `nodeB` to start the HARG.

9. On `nodeB`, the HA Manager runs the following command:

   `/opt/OV/lbin/ovharg -start hpom-server`

10. The `ovharg` command executes all `S*` tools in the `/var/opt/OV/hacluster/hpom-server` directory.

11. When the `ovharg` command successfully executes all `S*` tools, the HARG becomes `ONLINE`.

The following is the status after the HARG failover:

```
Nodes :
=======
nodeA : ONLINE
nodeB : ONLINE

HARGS :
=======
hpom-server :
 General status : FAULTED
 Active node    : nodeB
 Primary node   : nodeA

 ---------------
 nodeA : FAULTED
 nodeB : ONLINE
```

FAULTED nodeB

This use case shows what happens when nodeB becomes FAULTED.

The following is the status before nodeB becomes FAULTED:

```
Nodes :
=======
nodeA : ONLINE
nodeB : ONLINE
HARGS :
=======
hpom-server :
 General status : ONLINE
 Active node    : nodeB
 Primary node   : nodeA
 Failover node  : nodeA

 ---------------
 nodeA : OFFLINE
 nodeB : ONLINE
```

When nodeb becomes FAULTED, the following steps occur:

1. The HA Manager on nodeB detects that the reference node is down
   and there is no status update from nodeA for some time. Therefore,
   the HA Manager assumes that there is something wrong with nodeB
   and marks nodeB as FAULTED.

2. Because nodeB is FAULTED, the HA Manager stops all HARGs that
   are ONLINE on this node.

3. On nodeB, the HA Manager runs the following command:

   `/opt/OV/lbin/ovharg -stop hpom-server`

4. The ovharg command executes all K* tools in the
   /var/opt/OV/hacluster/hpom-server directory.

5. At this point, the status reported on nodeB is the following:

   ```
   Nodes :
   =======
   nodeA : OFFLINE
   nodeB : FAULTED

   HARGS :
   =======
   ```

```
hpom-server :
 General status : OFFLINE
 Active node    :
 Primary node   : nodeA

 ---------------
 nodeA : OFFLINE
 nodeB : OFFLINE
```

6. The HA Manager on nodeA detects that there is no status update from nodeB for some time. Therefore, the HA Manager assumes that there is something wrong with nodeB and marks nodeB as FAULTED.

7. Because nodeB is FAULTED and the HARG was ONLINE on nodeB, the HA Manager starts the HARG on nodeA, which is the primary node.

8. On nodeA, the HA Manager runs the following command:

   /opt/OV/lbin/ovharg -start hpom-server

9. The ovharg command executes all S* tools in the /var/opt/OV/hacluster/hpom-server directory.

10. When the ovharg command successfully executes all S* tools, the HARG becomes ONLINE.

The status reported on nodeA is the following:

```
Nodes :
=======
nodeA : ONLINE
nodeB : FAULTED

HARGS :
=======
hpom-server :
 General status : ONLINE
 Active node    : nodeA
 Primary node   : nodeA

 ---------------
 nodeA : ONLINE
 nodeB : OFFLINE
```

Troubleshooting

This section describes solutions to the specific problems you may encounter:

❑   HARG Status Is FAULTED

What to do if the status report shows that the HARG status is FAULTED.

❑   Node Status Is FAULTED

What to do if the status report shows that the node status is FAULTED.

HARG Status Is FAULTED

Problem

The status report shows the following:

```
Nodes :
=======
nodeA : ONLINE
nodeB : ONLINE

HARGS :
=======
hpom-server :
 General status : ONLINE
 Active node    : nodeA
 Primary node   : nodeA

 ---------------
 nodeA : ONLINE
 nodeB : FAULTED
```

Solution

To solve the problem, follow these steps:

1. On nodeB, check the HARG error file:

   /var/opt/OV/hacluster/<*HARG_name*>/error.log

The log file shows which monitor tool reported an error and it may also contain detailed information about the problem.

2. Depending on the monitor tool reporting a problem, a further investigation involving system administration must be performed.

3. When you find the cause of the problem and fix it, you can clear the FAULTED flag by using the opchamgr -clear *<HARG_name>* command.

If this is a repeated problem, follow these steps:

1. Make the HARG offline by running the following command:

   /opt/OV/bin/OpC/opchamgr -offline *<HARG_name>*

2. Disable the HARG by running the following command:

   /opt/OV/bin/OpC/opchamgr -harg -disable *<HARG_name>*

3. On the problematic node, manually start the HARG in the same way as the HA Manager starts it, that is, by running the following command:

   /opt/OV/lbin/ovharg -start *<HARG_name>*

   The return code should be 0.

4. Check the HARG status regularly by running the following command:

   /opt/OV/lbin/ovharg -monitor *<HARG_name>*

   Besides the status check, check also other system parameters that are useful for problem detection.

   When the HARG runs properly, the return code is 0. Otherwise, the return code is 1.

5. When the HARG status reports a problem (that is, when the return code is 1), analyze all provided data, find the cause of the problem, and then solve it.

6. Stop the HARG by running the following command:

   /opt/OV/lbin/ovharg -stop *<HARG_name>*

7. Enable the HARG by running the following command:

   /opt/OV/bin/OpC/opchamgr -harg -enable *<HARG_name>*

After a few moments, the HARG is automatically started on the primary node.

Node Status Is FAULTED

Problem

The status report shows the following:

```
Nodes :
=======
nodeA : ONLINE
nodeB : FAULTED

HARGS :
=======
hpom-server :
 General status : ONLINE
 Active node    : nodeA
 Primary node   : nodeA

 ---------------
 nodeA : ONLINE
 nodeB : OFFLINE
```

Solution

To solve the problem, follow these steps:

1. Check if all nodes are reachable by using the `ping` command, and then try to log on to problematic nodes.

2. Check if the HA Manager runs on all nodes by using the `opchamgr -daemon` command.

3. On each node, check if the remote shell or secure remote shell communication with the other nodes works (for example, try to copy several files between the nodes by using the remote shell copy).