High Availability Through HPOM Server Pooling

for HP Operations Manager on UNIX® and Windows® Operating Systems

Server Pooling Concept, Scenarios and Configuration

Overview

This document introduces the virtualization of HP Operations Manager (HPOM) management servers for both the UNIX and Windows operating systems, using the backup server concept (server pooling).

This section describes the server pooling concept and the configuration details. It contains the following information:

- Requirements for setting up server pooling.
  See Requirements on page 2.
- Server pooling concept details.
  See About Server Pooling on page 3.
- Three server pooling scenarios.
  See Server Pooling Scenarios on page 4.
- Details about the basic server pooling setup which includes two physical servers and one virtual interface.
  See Setup Details on page 6.
- Procedures for setting up two physical servers with one virtual interface.
  For setting up HPOM for UNIX server pooling, see Configuring Server Pooling in HPOM for UNIX Environments on page 8.

  For setting up HPOM for Windows server pooling, see Configuring Server Pooling in HPOM for Windows Environments on page 17.

Requirements

The following requirements apply:

- Two or more physical servers, that *must* be located in the same subnet:
  - For HPOM for UNIX, HPOM for UNIX 8.24 or higher must be installed on the physical servers.
  - For HPOM for Windows, HPOM for Windows 8.10 or higher must be installed on the physical servers.
- One or more virtual interfaces.

About Server Pooling

Server pooling is an enhancement of the HPOM Backup Server concept, which is an option to implement high availability for HPOM.

Typically, in a backup server scenario, two or more HPOM management servers are configured identically. The main installation is referred to as the primary manager and the others as backup servers. If the primary manager is temporarily inaccessible for any reason, you can configure HPOM to transfer messages from the managed nodes to one or more designated backup management servers.

In a server pooling scenario, HPOM management servers are also configured identically, but the role of primary manager is assigned to a virtual interface. Messages from managed nodes are not sent to a physical server but to its virtual interface.

If a physical server with a virtual interface is temporarily inaccessible for any reason, you can switch the virtual interface to another physical server. Agents on managed nodes reconnect to the virtual interface automatically, where no manual interaction is required. This is one of the main benefits of server pooling.

If you have to perform some maintenance tasks on one physical server, simply transfer its virtual interface to another physical server and proceed with your maintenance work.

HTTPS-based buffered message forwarding (hot message synchronization) is set up between all physical servers. Every message delivered to one physical server is transferred to all other physical servers, even when they are not accessible at the moment.

When your maintenance work is finished and the HPOM management server is running again, this physical server will get all those missed messages from the other physical server.

All physical servers are also defined as responsible managers. This allow all physical servers to perform the configuration deployment and the action execution on all managed nodes.

Every physical server can have more than one virtual interface at once. This allows implementing various scenarios, not only failover or maintenance from one HPOM management server to another. Virtual interfaces for the HPOM management server can also be leveraged for implementing load balancing. If the load from incoming messages is too high, load balancing software can move some of the virtual interfaces to other, less utilized physical servers.

Server Pooling Scenarios

This section describes several server pooling scenarios, which vary from simple to more complex. You can adapt any of these scenarios to meet your specific needs.

Scenario 1

Figure 1 shows two physical servers (`server1`, `server2`) with one virtual interface (`virtual_server`). This is similar to the classic backup server scenario. If you need to restart `server1`, you can simply switch the `virtual_server` to `server2`. After restart, `server1` receives all missed messages.

Figure 1. Scenario 1

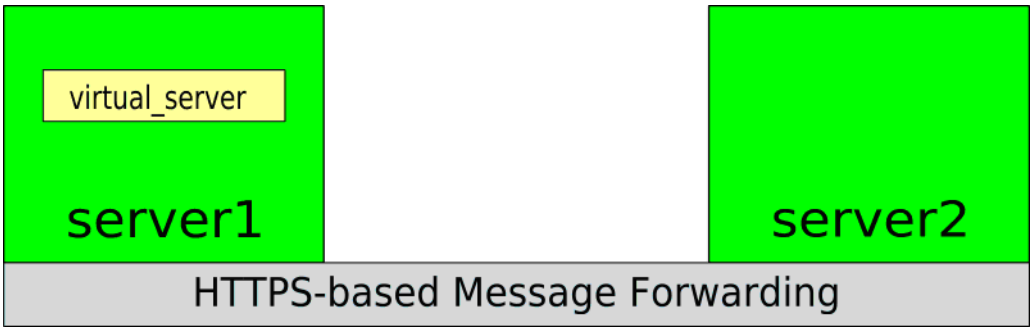

Scenario 2

Figure 2 shows two physical servers (`server1`, `server2`) with two virtual interfaces (`virtual_server1`, `virtual_server2`). In this case, load balancing is achieved with all managed nodes connected to the `virtual_server1` and all Java GUIs connected to the `virtual_server2`.

Figure 2. Scenario 2

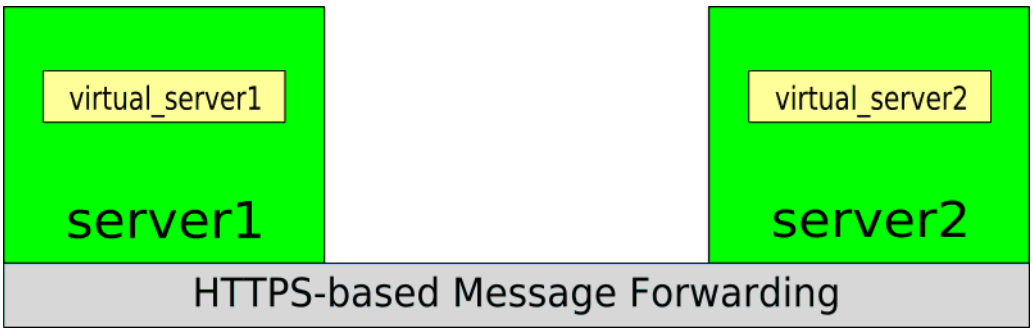

Scenario 3

Figure 3 shows a more complex scenario. It assumes three physical servers where each physical server has two virtual interfaces. Each managed node sends messages to one of the virtual interfaces. Physical servers are using event correlation to filter out related messages.

When load balancing software detects a high server load on the physical server `server1`, it switches one of the virtual interfaces from that server to another server. When the load on the physical server `server1` decreases, the load balancing software switches the virtual interface back to the `server1`.

Figure 3. Scenario 3

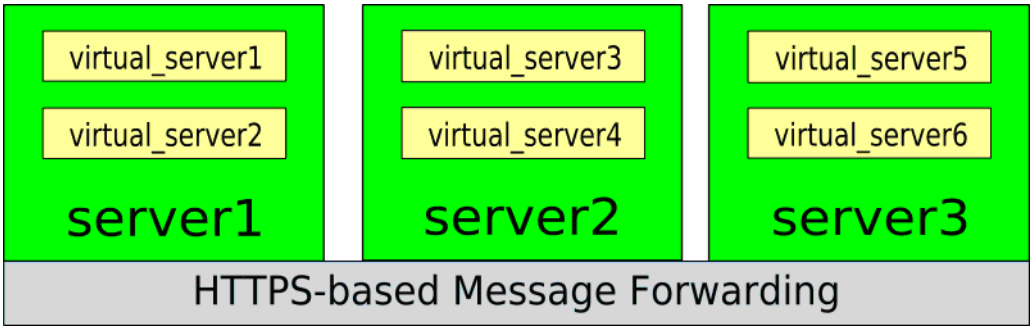

Setup Details

This section describes the details about the basic server pooling setup which includes two physical servers and one virtual interface.

This setup includes two instances of the HPOM management server (M1, M2), each with their own HPOM database. Each instance and the database are fully online all the time. In a backup server scenario with buffered message forwarding set up between them, each message is forwarded from an active server to a backup server.

A virtual interface (V) belongs only to one physical server at a time and could be managed by any load balancing software.

All managed nodes have M1, M2 and V defined as responsible managers. Each responsible manager has all rights including the action execution and the configuration deployment. Responsible managers are defined in the HPOM configuration directory with the `allnodes` file. This file is used to generate the `mgrconf` policy, which must be deployed to all nodes.

M1 and M2 entries in the `allnodes` file are required for the server-to-agent communication (configuration deployment, action execution). A virtual interface is required for the agent-to-server communication (message sending). For more information, refer to the *HP Operations Manager HTTPS Agent Concepts and Configuration Guide*, Chapter "Environments with Multiple OVO Management Servers (MoM)".

All managed nodes have also virtual interfaces defined as primary managers.

In the past, the `opcragt -primmgr` command was required for enabling the deployment of templates from a backup server. With new HTTPS agents, it is no longer needed. For more information about differences between using DCE and HTTPS agents, refer to the *HP Operations Manager HTTPS Agent Concepts and Configuration Guide*, Chapter "Backward Compatibility and the Differences between OVO 7 and OVO 8".

Server pooling is not supported for the DCE agents. Server pooling in the mixed DCE and HTTPS agents environment is possible, but is not recommended, because it involves the increased risk of error occurrence (for example, accidentally calling `opcragt -primmgr` for an HTTPS node).

The configuration deployment can be performed from both servers, but the configuration data on both servers should be kept synchronized.

M1 and M2 are set up to allow HTTPS-based buffered message forwarding from one server to another. This is defined with the `msgforw` file in the HPOM configuration directory. After failover, you do not have to synchronize messages between the two databases, because they are already synchronized. When the failed physical server starts up again, it receives all missed messages. For more details, refer to the chapter "Configuring Management Server Templates" in the *HPOM Administrator's Reference* guide.

All Java GUIs should connect to the virtual interface. Username and password for each user should be identical on both servers, so that Java GUIs can automatically reconnect in case of a failover.

When a switch occurs, the virtual interface is transferred from the primary manager to the backup server. Java GUIs show a small delay because they have to reconnect, which is done without user intervention. Agents also reconnect automatically without user intervention. The delay in message processing is reduced, because the agents practically do not buffer messages any more at the primary manager downtime. This solution also provides the database redundancy. If the primary database becomes corrupt, a forced switchover of the virtual interface can take place.

Figure 4. Configuration Setup Schema

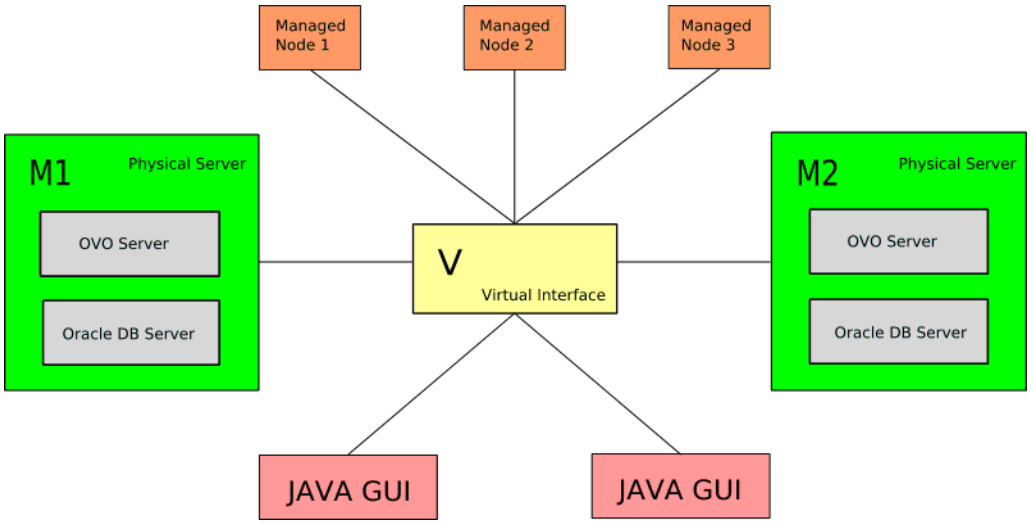

Configuring Server Pooling in HPOM for UNIX Environments

This section describes how to set up two physical servers with one virtual interface in HPOM for UNIX environments. With the procedure described below, you can easily set up more than two physical servers, and more than one virtual interface.

The following configuration changes are made to make this solution operable:

- All HPOM servers are configured in a backup server scenario with buffered message forwarding set up between them.
- A virtual interface and all physical servers are added as responsible managers in the responsible manager policy file.
- A virtual interface is set as a primary manager on all managed nodes.

Installing the HPOM for UNIX Management Server

Install the HPOM for UNIX management server as a standalone server on each physical server (M1 and M2).

For instructions on how to install the HPOM for UNIX management server, refer to *HP OpenView Operations for UNIX Installation Guide*.

Configuring Management Server Nodes for HPOM for UNIX

To set up certificate servers in the HPOM for UNIX environment, refer to the *HP Operations Manager HTTPS Agent Concepts and Configuration Guide*, Chapter "Security in Manager of Manager (MoM) Environments".

Configuring the Virtual Interface for HPOM for UNIX

Create new OV resource group on both physical servers, as follows:

1. Create a new OV resource group for the V virtual interface. On both management servers, enter the following command:

   ```
   /opt/OV/lbin/xpl/init_ovrg.sh virt
   ```

2. Create a new `OvCoreId` for the virtual interface. On the M1 physical server, enter the following:

   ```
   /opt/OV/bin/ovcoreid -create -ovrg virt
   /opt/OV/bin/ovcoreid -ovrg virt > /tmp/virt.coreid
   ```

   Copy the `/tmp/virt.coreid` file from the M1 to the same location on the M2. On the physical server M2, execute the following command:

   ```
   /opt/OV/bin/ovcoreid -set `cat /tmp/virt.coreid` -ovrg virt
   ```

3. Issue a new certificate for the V virtual interface. On the physical server M1, enter the following:

   ```
   /opt/OV/bin/ovcm -issue -file /tmp/virt.cert -name
   <V_virtual_interface> -pass virt \
   -coreid `cat /tmp/virt.coreid`
   ```

where:
<V_virtual_interface> is the name of the V virtual interface.

4 Import the new certificate into the keystore. On the physical server M1, enter the following command:

```
/opt/OV/bin/ovcert -importcert -ovrg virt -file \ /tmp/virt.cert -pass
virt
```

To verify the imported certificate, use the following command:

```
/opt/OV/bin/ovcert -list -ovrg virt
```

The certificate and Trusted Certificates should be listed.

Copy the `/tmp/virt.cert` file from the M1 to the same location on the M2. On the physical server M2, enter the following:

```
/opt/OV/bin/ovcert -importcert -ovrg virt -file \ /tmp/virt.cert -pass
virt
```

To verify the imported certificate, use the following command:

```
/opt/OV/bin/ovcert -list -ovrg virt
```

The certificate and Trusted Certificates should be listed.

5 Bind the V virtual interface's IP address to the new OV resource group `virt`. On both physical servers, enter the following command:

```
/opt/OV/bin/ovconfchg -ovrg virt -ns bbc.cb \
-set SERVER_BIND_ADDR <V_IP_address> -set SERVER_PORT 383
```

where:
<V_IP_address> is the IP address of the V virtual interface.

6 Add the V virtual interface to the Node Bank. On both physical servers, enter the following:

```
/opt/OV/bin/OpC/utils/opcnode -add_node \
node_name=<V_virtual_interface> \
net_type=NETWORK_IP mach_type=<machine_type> \
group_name=<node_group_name>
```

where:
<V_virtual_interface> is the fully qualified domain name (FQDN) of the V virtual interface.
<machine_type> is `MACH_BBC_HPUX_PA_RISC` (for HP-UX PA-RISC) or `MACH_BBC_HPUX_IPF32` (for HP-UX Itanium) or `MACH_BBC_SOL_SPARC` (for Solaris Sparc).
<node_group_name> is `hp_ux` (for HP-UX) or `solaris` (for Solaris).

7 Set `OvCoreId` for the V virtual interface in the Node Bank. On both physical servers, enter the following:

```
/opt/OV/bin/OpC/utils/opcnode -chg_id \ node_name=<V_virtual_interface>
id=`cat /tmp/virt.coreid`
```

where:
<V_virtual_interface> is the name of the V virtual interface.

8 Activate the V virtual interface on the physical server M1. Enter the following:

```
ifconfig <V_network_interface>:1 plumb          (Solaris only)
ifconfig <V_network_interface>:1 inet <V_IP_address> \ netmask
<V_netmask_value> up
```

where:
<V_network_interface> is the network interface used for the V virtual interface's IP address (for example, `lan0` for HP-UX or `hme0` for Solaris).
<V_IP_address> is the IP address of the V virtual interface.
<V_netmask_value> is the netmask value for the V virtual interface.

9 Start the `virt` OV resource group on the physical server M1:

`/opt/OV/bin/ovbbccb -start virt`

Configuring the Primary Manager for HPOM for UNIX

To create and configure a responsible manager file for managed nodes, follow the procedure below:

1 Add the physical servers and the virtual interface as responsible managers to the `allnodes` file. First, copy the file using the following command:

`cp /etc/opt/OV/share/conf/OpC/mgmt_sv\ /tmpl_respmgrs/backup-server / etc/opt/OV/share/conf/OpC\ /mgmt_sv/work_respmgrs/allnodes`

Modify the `allnodes` file to contain the physical servers and the virtual interface (M1, M2 and V). Example of the `allnodes` file:

```
#
# Responsible Manager Configurations for a backup server
#
RESPMGRCONFIGS
        RESPMGRCONFIG
        DESCRIPTION "responsible mgrs"
             SECONDARYMANAGERS
                     SECONDARYMANAGER
                     NODE IP 0.0.0.0 "serv1.bbn.hp.com"
                     SECONDARYMANAGER
                     NODE IP 0.0.0.0 "serv2.bbn.hp.com"
                     SECONDARYMANAGER
                     NODE IP 0.0.0.0 "virt.bbn.hp.com"
              ACTIONALLOWMANAGERS
                     ACTIONALLOWMANAGER
                     NODE IP 0.0.0.0 "serv1.bbn.hp.com"
                     ACTIONALLOWMANAGER
                     NODE IP 0.0.0.0 "serv2.bbn.hp.com"
                     ACTIONALLOWMANAGER
                     NODE IP 0.0.0.0 "virt.bbn.hp.com"
```

2 Verify whether you correctly modified the `allnodes` file, using the following command:

`/opt/OV/bin/OpC/opcmomchk /etc/opt/OV/share/conf/OpC\ /mgmt_sv/work_respmgrs/allnodes`

3 Copy the `allnodes` file to the configuration directory. Enter the following:

`cp /etc/opt/OV/share/conf/OpC/mgmt_sv/work_respmgrs\ /allnodes /etc/opt/OV/share/conf/OpC/mgmt_sv\ /respmgrs/allnodes`

4 Distribute the responsible manager template. On the physical server M1, execute the following command:

`/opt/OV/bin/OpC/opcragt -distrib -templates \ <M1_server_name>`

where:
<M1_server_name> is the name of the physical server M1, on which you are distributing the template.

To verify the deployed templates, use the following command:

/opt/OV/bin/ovpolicy -list

This should return the following message:

```
mgrconf            "OVO authorization"          enabled    1
```

5 Repeat the procedure on the physical server M2. Copy the `/etc/opt/OV/share/conf/OpC/mgmt_sv/respmgrs/allnodes` file from the M1 to the same location on the M2. On the physical server M2, execute the following command:

/opt/OV/bin/OpC/opcragt -distrib -templates \ *<M2_server_name>*

where:
<M2_server_name> is the name of the physical server M2, on which you are distributing the template.

To verify the deployed templates, use the following command:

/opt/OV/bin/ovpolicy -list

This should return the following message:

```
mgrconf            "OVO authorization"          enabled    1
```

NOTE
The responsible manager policy (`mgrconf`) must also be deployed to all managed nodes. See Configuring Managed Nodes for HPOM for UNIX on page 13 for more information.

Configuring Message Forwarding for HPOM for UNIX

To create and configure buffered message forwarding between two physical servers, follow the procedure bellow:

1. Create and modify the `msgforw` template. Copy the file using the following command:

   `cp /etc/opt/OV/share/conf/OpC/mgmt_sv/tmpl_respmgrs\ /msgforw /etc/opt/OV/share/conf/OpC/mgmt_sv/respmgrs\ /msgforw`

   Modify the `msgforw` file to contain both physical servers. Example of the `msgforw` file:

   ```
   TIMETEMPLATES

   RESPMGRCONFIGS
      RESPMGRCONFIG
      DESCRIPTION "msg-forwarding specification"
         MSGTARGETRULES
            MSGTARGETRULE
            DESCRIPTION "all messages"
               MSGTARGETRULECONDS
               MSGTARGETMANAGERS
                  MSGTARGETMANAGER
                     TIMETEMPLATE "$OPC_ALWAYS"
                     OPCMGR IP 0.0.0.0 "serv1.bbn.hp.com"
                     MSGCONTROLLINGMGR
                  MSGTARGETMANAGER
                     TIMETEMPLATE "$OPC_ALWAYS"
                     OPCMGR IP 0.0.0.0 "serv2.bbn.hp.com"
                     MSGCONTROLLINGMGR
   ```

   Verify that you correctly modified the `msgforw` file, using the following command:

   `/opt/OV/bin/OpC/opcmomchk \ /etc/opt/OV/share/conf/OpC/mgmt_sv/respmgrs/msgforw`

2. Activate buffered message forwarding on the physical server M1 by setting the variable `OPC_HTTPS_MSG_FORWARD` to `TRUE`:

   `/opt/OV/bin/ovconfchg -ovrg server -ns opc -set \ OPC_HTTPS_MSG_FORWARD TRUE`

3. Repeat the procedure on the physical server M2. Copy the `/etc/opt/OV/share/conf/OpC/mgmt_sv/respmgrs/msgforw` file from the M1 to the same location on the M2. On the physical server M2, execute the following command:

   `/opt/OV/bin/ovconfchg -ovrg server -ns opc -set \ OPC_HTTPS_MSG_FORWARD TRUE`

4. Activate the `msgforw` template on the management server. Enter the following on both physical servers:

   `/opt/OV/bin/OpC/opcsv -start`

   To verify the `msgforw` template, start the GUI and open a message browser on both physical servers, then send a message from the physical server M1:

   `/opt/OV/bin/OpC/opcmsg a=test o=test msg_t=test`

   The message should be displayed in the message browser on both physical servers.

Configuring Managed Nodes for HPOM for UNIX

An agent profile maintained on the HPOM for UNIX management server is a list of the configuration settings, which is copied to the HTTPS agent upon installation. Any settings defined in the `bbc_inst_defaults` file are also added to the agent profile. To learn more about the agent profile, refer to the *HP Operations Manager HTTPS Agent Concepts and Configuration Guide*, Chapter "Agent Profile". This section shows how to configure both newly installed agents and the already existing agents.

If you want to perform a new agent installation on the managed nodes, follow the procedure bellow:

1 Instruct each managed node that its primary manager is the V virtual interface. Place an entry in the `bbc_inst_defaults` file on both physical servers. The file is located in the following directory:

`/etc/opt/OV/share/conf/OpC/mgmt_sv/`

Add the namespace and the primary manager specification in the `bbc_inst_defaults` file on the physical server M1 as follows:

```
[eaagt]
OPC_PRIMARY_MGR=<V_virtual_interface>
```

where:
<V_virtual_interface> is the name of the V virtual interface.

2 Repeat the procedure on the the physical server M2.

Add the namespace and the primary manager specification to the `bbc_inst_defaults` file as follows:

```
[eaagt]
OPC_PRIMARY_MGR=<V_virtual_interface>
```

where:
<V_virtual_interface> is the name of the V virtual interface.

3 Install an HTTPS agent on each managed node. To verify the deployed templates, use the following command on the physical server M1:

```
/opt/OV/bin/ovpolicy -list -host <node>
```

where:
<node> is the name of the managed node.

This should return the following message:

```
mgrconf          "OVO authorization"          enabled    1
```

On existing HTTPS agents, follow the procedure below:

1 Distribute the responsible manager template. On the physical server M1, execute the following command:

```
/opt/OV/bin/OpC/opcragt -distrib -templates \ <managed_nodes>
```

where:
<managed_nodes> are the names of managed nodes that belong to the physical server M1.

To verify the deployed templates, use the following command on the physical server M1:

```
/opt/OV/bin/ovpolicy -list -host <node>
```

where:
<node> is the name of the managed node.

This should return the following message:

```
mgrconf          "OVO authorization"          enabled    1
```

2 Repeat the procedure on the physical server M2. Execute the following command:

/opt/OV/bin/OpC/opcragt -distrib -templates \ *<managed_nodes>*

where:
<managed_nodes> are the names of managed nodes that belong to the physical server M2.

To verify the deployed templates, use the following command on the physical server M2:

/opt/OV/bin/ovpolicy -list -host *<node>*

where:
<node> is the name of the managed node.

This should return the following message:

```
mgrconf            "OVO authorization"          enabled    1
```

3 Instruct each managed node that its primary manager is the V virtual interface.

For each managed node that belongs to physical server M1 (M2), use the following command on physical server M1 (M2):

**/opt/OV/bin/OpC/opcragt -set_config_var **
eaagt:OPC_PRIMARY_MGR=*<virtual_interface>* *<node>*

where:
<virtual_interface> is the name of the V virtual interface, and
<node> is the name of the managed node that belongs to physical server M1 (M2).

The alternative way to set `OPC_PRIMARY_MGR` is to configure the `mgrconf` policy with the `MSGTARGETRULE` rule, so that the virtual interface is used as a target for all messages. For more on this, refer to the *HPOM Administrator's Reference* guide.

The following is an example of the relevant part of the `allnodes` file:

```
MSGTARGETRULECONDS
   MSGTARGETMANAGERS
      MSGTARGETMANAGER
         TIMETEMPLATE "$OPC_ALWAYS"
         OPCMGR IP 0.0.0.0 "virt.bbn.hp.com"
```

CAUTION

Do *not* perform the `opcragt -primmgr` commands on the physical server. If you do it, the `OPC_PRIMARY_MGR` entry on the HTTPS agent side is overwritten with the address of the physical server, where `opcragt -primmgr` was executed. The HTTPS agent will send messages to that physical server (unless `MSGTARGETRULE` is defined in the `allnodes` file), and in the case of a switchover, the HTTPS agent will not send messages to the new physical server with the virtual interface's IP address.

Moving a Virtual Interface to another HPOM for UNIX Physical Server

When you need to move a virtual interface from one physical server to another, follow the procedure described below:

To disable the V virtual interface on the physical server M1, enter the following:

1 Stop the `virt` OV resource group on the physical server M1. Enter the following:

```
/opt/OV/bin/ovbbccb -stop virt
```

2 Stop the V virtual interface on the physical server M1. Enter the following:

```
ifconfig <V_network_interface>:1 inet 0.0.0.0 down
ifconfig <V_network_interface>:1 unplumb     (Solaris only)
```

where:
<V_network_interface> is the network interface used for the V virtual interface's IP address (for example, `lan0` for HP-UX or `hme0` for Solaris).

To enable the V virtual interface on the physical server M2, enter the following:

1 Start the V virtual interface on the physical server M2. Enter the following:

```
ifconfig <V_network_interface>:1 plumb     (Solaris only)
ifconfig <V_network_interface>:1 inet <V_IP_address> \ netmask
<V_netmask_value> up
```

where:
<V_network_interface> is the network interface used for the V virtual interface's IP address (for example, `lan0` for HP-UX or `hme0` for Solaris).
<V_IP_address> is the IP address of the V virtual interface.
<V_netmask_value> is the netmask value for the V virtual interface.

2 Start `virt` OV resource group on the physical server M2. Enter the following:

```
/opt/OV/bin/ovbbccb -start virt
```

Server Pooling in an HPOM for UNIX Cluster Environment

You can use a cluster to automate the monitoring and switching of the virtual interfaces. Consider the following aspects:

1 Server pooling is *not* supported in case you have an existing HPOM for UNIX installation in a cluster, as described in *HP OpenView Operations Installation Guide*, Chapter "Installing OVO in a Cluster Environment".

2 The HPOM for UNIX management server must be installed as a standalone server on both physical servers. When installing HPOM for UNIX management servers in a cluster environment, the following question is displayed:

```
Configure OVO Server as HA resource group (y|n):
[y]
```

Enter **n**.

3 Instead of the `ifconfig` command, use the cluster specific commands to activate and deactivate the virtual interface, as described in the section Configuring the Virtual Interface for HPOM for UNIX on page 8. For example, you can use the `cmmodnet` command for HP Serviceguard.

4 Switching an IP address is a standard part of the cluster `HARG`[1], and you usually need only to specify the virtual IP address in the configuration scripts or settings.

5 Specify `/opt/OV/bin/ovbbccb -start virt` in your `HARG` start script and `/opt/OV/bin/ovbbccb -stop virt` in your `HARG` stop script. See Moving a Virtual Interface to another HPOM for UNIX Physical Server on page 15 for more information.

6 If you want to switch the virtual interface automatically when some server processes stop running, write a monitor script.

For example, you could write a script that calls `/opt/OV/bin/OpC/utils/ha/ha_mon_ovserver` script in a loop. The `ha_mon_ovserver` script will return 0 if all processes are running, otherwise it will return 1.

Migration to Server Pooling from an Existing HPOM for UNIX Backup Server Environment

The existing Backup Server environment could be easily turned into a server pooling environment. This could be performed in one step, or you could use a phased approach.

To migrate to server pooling, perform the following steps:

1 Configure a virtual interface. See Configuring the Virtual Interface for HPOM for UNIX on page 8 for more information.

2 Add the virtual interface to the existing responsible manager file and then distribute the changed responsible manager policy to all nodes.

3 Message forwarding between physical servers is probably already set as required. The HTTPS-based message forwarding is recommended.

4 Configure the agents on managed nodes to send their messages to the virtual interface. This can be performed in one step for all agents, or in phases for the limited number of agents. Use the procedure which describes how to change primary manager setting for the existing HTTPS agents from section Configuring Managed Nodes for HPOM for UNIX on page 13.

After you change the primary manager setting to the virtual interface on existing agents with the `opcragt` command, it is not necessary to restart the agents in order to start sending their messages to the virtual interface.

1. `HARG` is an HPOM for UNIX management server concept, and is equivalent to a package in HP Serviceguard, a service group in VERITAS Cluster System, and a resource group in Sun Cluster.

Configuring Server Pooling in HPOM for Windows Environments

This section describes how to set up two physical servers with one virtual interface in HPOM for Windows environments. With the procedure described below, you can easily set up more than two physical servers, and more than one virtual interface.

- All HPOM servers are configured in a backup server scenario with buffered message forwarding set up between them.
- A virtual interface and all physical servers are added as responsible managers in the responsible manager policy file.
- A virtual interface is set as a primary manager on all managed nodes.

Before starting to your installation and configuration tasks, be aware of the limitations for server pooling in HPOM for Windows environments stated below.

Limitations for Server Pooling in HPOM for Windows Environments

The following limitations apply to server pooling in HPOM for Windows environments:

- Connections to the virtual interface from the HPOM MMC console are not supported. Only web console sessions are supported.
- Physical nodes that are part of a clustered HPOM management server installation cannot be part of an HPOM server pool.
- Currently, clustered HPOM management servers cannot be added to an HPOM server pool. Investigations are underway to attempt to overcome this limitation.
- The following are *not* supported for this release of HPOM Server Pooling:
  - External DNS-based hostname redirects for the virtual interface (CNAME entries)
  - External IP mapping mechanisms
  - Route Health Injection

Installing the HPOM for Windows Management Server

Install the HPOM for Windows management server as a standalone server on each physical server (M1 and M2). For instructions on how to install the HPOM for Windows management server, refer to *HP Operations Manager for Windows Installation Guide*.

Configuring the Virtual Interface for HPOM for Windows

Create a new OV resource group on both physical servers, as follows:

1 Create a new OV resource group for the V virtual interface. On both HPOM for Windows management servers, enter the following commands:

```
md %OvDataDir%\shared\virt\conf\xpl\config
md %OvDataDir%\shared\virt\datafiles\xpl\config\jobs
```

where:
`virt` is the name of the virtual resource group.

2 Create a new OvCoreId for the virtual interface. On the M1 physical server, enter the following:

```
ovcoreid -create -ovrg virt
ovcoreid -ovrg virt > C:\tmp\virt_coreid.txt
```

Copy the `C:\tmp\virt_coreid.txt` file from the M1 to the same location on the M2. On the physical server M2, execute the following command:

```
ovcoreid -set <GUID from C:\tmp\virt_coreid.txt> -ovrg virt
```

where:
`GUID` is the GUID contained in `C:\tmp\virt_coreid.txt`.

3 Issue a new certificate for the V virtual interface. On the physical server M1, enter the following:

```
ovcm -issue -file C:\tmp\virt.cert -name \
<V_virtual_interface> -pass virt \
-coreid <GUID from C:\tmp\virt_coreid.txt>
```

where:
<V_virtual_interface> is the name of the V virtual interface.

4 Import the new certificate into the keystore. On the physical server M1, enter the following command:

```
ovcert -importcert -ovrg virt -file \
C:\tmp\virt.cert -pass virt
```

To verify the imported certificate, use the following command:

```
ovcert -list -ovrg virt
```

The certificate and Trusted Certificates should be listed.

Copy the `C:\tmp\virt.cert` file from the M1 to the same location on the M2. On the physical server M2, enter the following:

```
ovcert -importcert -ovrg virt -file \
C:\tmp\virt.cert -pass virt
```

To verify the imported certificate, use the following command:

```
ovcert -list -ovrg virt
```

The certificate and trusted certificates should be listed.

5 Bind the V virtual interface's IP address to the new OV resource group `virt`. On both physical servers, enter the following command:

```
ovconfchg -ovrg virt -ns bbc.cb \
-set SERVER_BIND_ADDR <V_IP_address> -set SERVER_PORT 383
```

where:
<V_IP_address> is the IP address of the V virtual interface.

6 To activate the V virtual interface on the physical server M1, assign the V virtual interface's IP address to the network interface adapter of M1.

7 Start the virt OV resource group on the physical server M1:

```
ovbbccb -start virt
```

Configuring the Primary Manager for HPOM for Windows

To configure the responsible managers for managed nodes, follow the procedure below:

1 In HPOM for Windows, create a new flexible management (agent-based) policy. Add the physical servers and the virtual interface (M1, M2, V) as responsible managers to the policy. You also need to add the core IDs of all responsible managers, which are returned by executing the following command on each physical system:

ovcoreid

The core ID of the virtual interface is the GUID that has been created and written to the file in Configuring the Virtual Interface for HPOM for Windows on page 17.

Example of the flexible management policy:

```
#
# Configuration policy
# defines action-allow managers
# messages are always send to the node's primary manager
#

RESPMGRCONFIGS

    RESPMGRCONFIG
    DESCRIPTION "Server Pool Members authorized for Management"
    SECONDARYMANAGERS
      SECONDARYMANAGER
         NODE IP 0.0.0.0 "serv1.ovowtest.dom" ID
         "6150dbc2-5e4c-7531-193f-858e0b716c94"
         DESCRIPTION "physical"
         SECONDARYMANAGER
         NODE IP 0.0.0.0 "serv2.ovowtest.dom" ID
         "8fee9552-fc01-7531-06dc-e78703acbf1b"
         DESCRIPTION "physical"
         SECONDARYMANAGER
         NODE IP 0.0.0.0 "virt.ovowtest.dom" ID
         "cf600f32-0aba-7532-1e83-f61cdcefb5d7"
         DESCRIPTION "virtual"
    ACTIONALLOWMANAGERS
         ACTIONALLOWMANAGER
         NODE IP 0.0.0.0 "serv1.ovowtest.dom" ID
         "6150dbc2-5e4c-7531-193f-858e0b716c94"
         DESCRIPTION "physical"
         ACTIONALLOWMANAGER
         NODE IP 0.0.0.0 "serv2.ovowtest.dom" ID
         "8fee9552-fc01-7531-06dc-e78703acbf1b"
         DESCRIPTION "physical"
         ACTIONALLOWMANAGER
         NODE IP 0.0.0.0 "virt.ovowtest.dom" ID
         "cf600f32-0aba-7532-1e83-f61cdcefb5d7"
         DESCRIPTION "virtual"
```

2 Verify whether you correctly modified the flexible management policy, using the **Check Syntax** button in the policy editor.

3 Save the policy. Specify a name like "OM Server Pool Authorization" to find the policy easier later. Close the Node Editor and assign the policy to a policy group in order to allow a download of the policy for a later configuration synch with other management servers. You can also use a new Policy Group for Server Pooling-specific agent policies.

4 Distribute the flexible management policy to the local physical server.

5 Remember to repeat the procedure on the physical server M2 later. It is useful to copy the flexible management policy to M2 by performing a configuration synch. The command line tool `ovpmutil` allows you to do this in an easy manor.

Configuring Message Forwarding for HPOM for Windows

To create and configure buffered message forwarding between two physical servers, follow the procedure below:

1 Create a new server-based flexible management policy. Add both physical servers as Message Target Managers to the policy. Example of the server-based flexible management policy:

```
#
# Server-Pool msg forwarding policy
#

TIMETEMPLATES
# none

RESPMGRCONFIGS
   RESPMGRCONFIG DESCRIPTION ""
   SECONDARYMANAGERS
   ACTIONALLOWMANAGERS

   MSGTARGETRULES
        MSGTARGETRULE DESCRIPTION "Forward all Messages to all Server Pool
        Members"
          MSGTARGETRULECONDS
          MSGTARGETMANAGERS
             MSGTARGETMANAGER
                TIMETEMPLATE "$OPC_ALWAYS"
                OPCMGR IP 0.0.0.0 "serv1.ovowtest.dom"
             MSGTARGETMANAGER
                TIMETEMPLATE "$OPC_ALWAYS"
                OPCMGR IP 0.0.0.0 "serv2.ovowtest.dom"
```

2 Verify whether you correctly modified the server-based flexible management policy, using the **Check Syntax** button in the policy editor.

3 Save the policy. Specify a name like "OM Server Pool Message Synch" to find the policy easier later. Close the Node Editor and assign the policy to a policy group in order to allow a download of the policy for a later configuration synch with other management servers. You can also use a new Policy Group for Server Pooling-specific server policies.

4 Distribute the server-based flexible management policy to the local physical server.

5 Remember to repeat the procedure on the physical server M2 later. It is useful to copy the server-based flexible management policy to M2 by performing a configuration synch. The command line tool ovpmutil allows you to do this in an easy manor.

Exchange Trusted Certificates for HPOM for Windows

To enable server-based flexible management, you need to exchange the servers' trusted certificates.

1 On Server M1, execute the following command:

```
ovcert -exporttrusted -file C:\tmp\M1.cert -ovrg server
```

2 Transfer the file `C:\tmp\M1.cert` to server M2.

3 On Server M2, execute the command:

```
ovcert -importtrusted -file <M1_cert> -ovrg server
```

where:
<M1_cert> is the absolute filename of the certificate file that was transferred to Server M2.

4 To export the M2 trusted certificate, execute the following command:

```
ovcert -exporttrusted -file C:\tmp\M1_M2.cert -ovrg server
```

Note: The trusted Certificates of M1 and M2 have been already merged in step 3.

5 Transfer the file `C:\tmp\M1_M2.cert` to server M1.

6 On Server M1, execute the command:

```
ovcert -importtrusted -file <M1_M2_cert> -ovrg server
```

where:
<M1_M2_cert> is the absolute filename of the certificate file that was transferred from M2 to M1.

7 After the trusted certificates have been exchanged between both servers, execute the following command on both servers:

```
ovcert -updatetrusted
```

8 When you have completed all the changes in the trusted certificate configuration on the management servers, you need to synchronize these changes to the managed nodes. On all managed nodes, you need to execute the following command:

```
ovcert -updatetrusted
```

You can simplify this step by running the tool **Update trusted certificates** in the HPOM console. For more information, see the following topic in the HPOM for Windows Online Help: **Administering Your Environment → Configuring Certificates → Configure trusted certificates for multiple management servers**.

Configuring Managed Nodes for HPOM for Windows

An agent profile maintained on the HPOM management server is a list of the configuration settings, which is copied to the HTTPS agent upon installation. Any settings defined in the agent_install_defaults.cfg file are also added to the agent profile. To learn more about the agent profile, refer to the HPOM Online Help topic "Configure HTTPS agent installation defaults". This section shows how to pre-configure newly installed agents.

If you want to perform a new agent installation on the managed nodes, follow the procedure below:

1. Instruct each managed node that its primary manager is the V virtual interface. Place an entry in the `agent_install_defaults.cfg` file on both physical servers. The file is located in the following directory:

   `%OvDataDir%\shared\conf\pmad`

   If the file does not exist, you need to copy or rename the sample file that you find in the same directory. Add the namespace and the primary manager specification in the `agent_install_defaults.cfg` file on the physical server M1 as follows:

   `[eaagt]`
   `OPC_PRIMARY_MGR=<V_virtual_interface>`

   where:
   `<V_virtual_interface>` is the name of the V virtual interface.

2. To validate your changes to the agent install profile, execute the following command:

   `ovpmutil dnl prf /fqdn <name of any managed node> /d C:\tmp`

   The profile snippet created in `C:\tmp` should contain the following line:

   `eaagt:OPC_PRIMARY_MGR=<V_virtual_interface>`

3. Define the policy group that contains the flexible management policy to be auto-deployed to all new managed nodes. You configure this in the property dialog of the root node group in the Node Editor.

4. Install an HTTPS agent on each managed node.

On existing HTTPS agents, follow the procedure below:

1. Distribute the flexible management policy from Server M1 to all managed nodes that have M1 set as primary manager.

2. Repeat the procedure on the physical server M2.

3. Instruct each managed node that its primary manager is the V virtual interface.

   For each managed node that belongs to physical server M1 (M2), use the following command on physical server M1 (M2):

   `opcragt -set_config_var \`
   `eaagt:OPC_PRIMARY_MGR=<virtual_interface> <node>`

   where:
   `<virtual_interface>` is the name of the V virtual interface, and `<node>` is the name of the managed node that belongs to physical server M1 (M2).

The alternative way to set `OPC_PRIMARY_MGR` is to configure the flexible management policy with the `MSGTARGETRULE` rule, so that the virtual interface is used as a target for all messages. For more information, refer to the *HPOM Administrator's Reference Guide*.

The following is an example of the relevant part of the flexible management policy:

```
MSGTARGETRULECONDS
MSGTARGETMANAGERS
    MSGTARGETMANAGER
      TIMETEMPLATE "$OPC_ALWAYS"
      OPCMGR IP 0.0.0.0 "virt.ovowtest.dom" ID
      "cf600f32-0aba-7532-1e83-f61cdcefb5d7"
```

CAUTION

Do *not* perform the `opcragt -primmgr` commands on the physical server. If you do, the `OPC_PRIMARY_MGR` entry on the HTTPS agent side is overwritten with the address of the physical server, where **`opcragt -primmgr`** was executed. The HTTPS agent will send messages to that physical server (unless `MSGTARGETRULE` is defined in the flexible management policy), and in the case of a switchover, the HTTPS agent will not send messages to the new physical server with the virtual interface's IP address.

Moving a Virtual Interface to another HPOM for Windows Physical Server

When you need to move a virtual interface from one physical server to another, follow the procedure described below.

To disable the V virtual interface on the physical server M1, perform these steps:

1. Stop the `virt` OV resource group on the physical server M1. Enter the following:

   `ovbbccb -stop virt`

2. Stop the V virtual interface by de-assigning the virtual IP address from the physical server M1.

To enable the V virtual interface on the physical server M2, perform these steps:

1. Start the V virtual interface on the physical server M2 by assigning the virtual IP address to its network interface.
2. Start virt OV resource group on the physical server M2. Enter the following:

   `ovbbccb -start virt`

Additional Useful Information for HPOM for Windows

DNS Configuration

1. All management servers as well as the virtual interface should use static IP addresses.
2. All IP addresses, including those for the virtual interfaces, should appear in the Reverse Lookup zone on your DNS server. Create additional reverse pointer records if these entries are missing.

Certificate Handling

The command line tool ovcert is very helpful to examine the certificate key store of a management server. To view installed certificates, type:

`ovcert -list`

The output lists two key stores:

- Default key store
- Server key store

The default key store lists agent certificates. Certificates installed in the server key store are used by HPOM for Windows server components.

Each installed certificate is listed as a GUID. To view extended information about a certificate, execute the following command:

`ovcert -certinfo <GUID> [-ovrg server]`

Note: The parameter `-ovrg server` must be specified if the certificate referred to in <GUID> is installed in the server key store.

SSL Connections

The command line tool bbcutil allows testing connectivity and the correct certificate exchange between two nodes. To ping a remote system, execute the following command:

`bbcutil -ping https://<hostname>`

This ping command tests HPOM Communication Broker on the remote system. It will be reported as unavailable, if there is neither an HPOM Agent nor HPOM Server running on the system. An Ok result means that there are trusted certificates installed in the default key store. An SslError result indicates missing trusted certificates on one or both nodes (ping source and target).

Certificates of the server key store are not used by `bbcutil`.

You can also ping the network name of the virtual interface.

Configuration Synchronization

If you have more than two management servers in your server pool, it might become necessary to simplify the Server Pooling configuration synch between your physical servers. It is useful to store your configuration data in a single place that is either accessible from all physical servers (for example, a network file share), or on a removable media (for example, a floppy disk).

The Server Pool configuration includes:

- The physical server's data model, which includes
  - Server pooling-relevant policies for action-allowed managers (flexible management) and message forwarding (server-based flexible management)
  - Node configuration
  - User roles
  - Service model
  - Server configuration values
- File `<virt>_coreid.txt` that contains a virtual interface's Core ID
- File `<virt>.cert` that holds a virtual interface's certificate
- Trusted certificates of all physical servers (either the single certificate of each server or all certificates merged in one file)
- Default agent installation profile `agent_install_defaults.cfg`

You would typically create this store when you are finished configuring the first server.

1 Export the entire configuration by executing the following command:

ovpmutil cfg all dnl *<Config_Store_Dir>*

where:
<Config_Store_Dir> is a directory in your configuration store. If this directory does not exist, it will be created.

2 Copy the <virt>_coreid.txt and <virt>.cert files that have been created in the section Configuring the Virtual Interface for HPOM for Windows on page 17 to your configuration store. The coreid.txt file is used to reference the coreid when you want to create the virtual interface on other management servers. The .cert file holds the virtual interface's certificate, that needs to be imported after the virtual interface has been configured on a new server pool member.

3 Copy the .cert files that were created in the section Exchange Trusted Certificates for HPOM for Windows on page 21 to your configuration store. These certificates need to be imported to new server pool members to create a trust relationship between the servers and the managed agents.

4 Copy the default agent installation profile to the configuration store.

On the next HPOM for Windows server that needs to be configured for server pooling, perform the steps from the section Configuring the Virtual Interface for HPOM for Windows on page 17 to Configuring Managed Nodes for HPOM for Windows on page 22, and reuse the stored configuration instead of creating new artifacts.

1 Import the Server Configuration Data by executing the following command:

ovpmutil cfg all upl *<Config_Store_Dir>* /noautodeploy

where:
<Config_Store_Dir> is the directory in your configuration store that holds the configuration data.

2 Use the GUID from the <virt>_coreid.txt file to set the virtual interface's coreid. Import the certificate from <virt>.cert as described in the section Configuring the Virtual Interface for HPOM for Windows on page 17.

3 Exchange the trusted certificates from your Configuration Store with the one from the new Server Pool Member as described in section Exchange Trusted Certificates for HPOM for Windows on page 21. Please note that all Servers need to exchange their trusted certificates with all other Server Pool members.

4 Place the default agent installation profile from your configuration store into the directory described in section Configuring Managed Nodes for HPOM for Windows on page 22.

Server Pooling in an HPOM for Windows Cluster Environment

Nodes that run as a part of a clustered HPOM for Windows management server installation cannot be included into server pools.

Including clustered HPOM for Windows management servers into a Server Pool is currently not supported.

Migration to Server Pooling from Existing HPOM for Windows Backup Server Environments

The existing Backup Server environment could be easily turned into a server pooling environment. This could be performed in one step, or you could use a phased approach.

To migrate to server pooling, perform the following steps:

1. Configure a virtual interface. See Configuring the Virtual Interface for HPOM for Windows on page 17 for more information.
2. Add the virtual interface to the existing flexible management policy and then distribute the changed policy to all nodes.
3. Message forwarding between physical servers is probably already set as required. HTTPS-based message forwarding is recommended.
4. Configure the agents on managed nodes to send their messages to the virtual interface. This can be performed in one step for all agents, or in phases for a limited number of agents. See Configuring Managed Nodes for HPOM for Windows on page 22 for a description about how to change primary manager settings for the existing HTTPS agents.

   After you change the primary manager setting to the virtual interface on existing agents with the `opcragt` command, it is not necessary to restart the agents in order to start sending their messages to the virtual interface.

© 2008 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Intel® Itanium® is a trademark of Intel Corporastion in the U.S. and other countries.

Java™ is a US trademark of Sun Microsystems, Inc.

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

UNIX® is a registered trademark of The Open Group.

PDF only, October 2008