HP Operations Manager

For the Windows operating system

Software Version: 8.16 and higher

High Availability Through Server Pooling

Document Release Date: August 2010

Software Release Date: August 2010

Legal Notices

Warranty

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.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

For information about third-party license agreements, see the license-agreements directory on the product installation media.

Copyright Notices

Acknowledgements

This product includes software developed by the JDOM Project (http://www.jdom.org/).
This product includes cryptographic software written by Eric Young (eay@cryptsoft.com).
This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (http://www.openssl.org/).
This product includes software written by Tim Hudson (tjh@cryptsoft.com).
This product includes software developed by the Apache Software Foundation (http://www.apache.org/).
This product includes an interface of the 'zlib' general purpose compression library, which is Copyright © 1995-2002 Jean-loup Gailly and Mark Adler.

Trademark Notices

Adobe®, Acrobat®, and PostScript® are trademarks of Adobe Systems Incorporated.
AMD is a trademark of Advanced Micro Devices, Inc.
HP-UX Release 10.20 and later and HP-UX Release 11.00 and later (in both 32 and 64-bit configurations) on all HP 9000 computers are Open Group UNIX 95 branded products.
Intel®, Itanium®, and Pentium® are trademarks of Intel Corporation in the U.S. and other countries.
Java™ is a US trademark of Sun Microsystems, Inc.
Microsoft®, Windows®, Windows NT®, Windows® XP, and Windows Vista® are U.S. registered trademarks of Microsoft Corporation.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
UNIX® is a registered trademark of The Open Group.

Support

You can visit the HP Software Support web site at:

http://www.hp.com/go/hpsoftwaresupport

HP Software Support Online provides an efficient way to access interactive technical support tools. As a valued support customer, you can benefit by using the support site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests

- Download software patches
- Manage support contracts
- Look up HP support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training
- Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract.

To find more information about access levels, go to:

http://h20230.www2.hp.com/new_access_levels.jsp

For more information about HP Passport, go to:

http://h20229.www2.hp.com/passport-registration.html

Table of Contents

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 using the command `opcragt -primmgr`. (Note that in large environments this switch may take some time to complete.)

In a server pooling scenario, HPOM management servers are also configured identically, but 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.

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 allows 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, which allows the implementation of load balancing. **Load balancing** in general terms refers to spreading a workload among multiple computers. Load balancing is often achieved by a load balancing software or hardware device.

In the context of HPOM, load balancing refers to the concept of switching the responsibility for a group of managed nodes or switching a Java GUI from one management server to another; for example, in the following situations:

- The load from incoming messages is too high.
- The number of managed nodes is too high for one management server. With a second management server added, you can split the load by directing half of the managed nodes or Java GUIs to the second management server.

Server pooling is not designed for dynamic, short term load balancing. This is because the agent may run into a timeout and may start buffering messages. After a successful switchover, it will establish a new connection. Load balancing can be used, however, for longer term, manual load balancing.

HPOM does not support any load balancing software installed on the management server. To move some of the virtual interfaces to other, less used physical servers, you use commands such as `ovbbccb`, `netstat`, and `ifconfig`.

If you are using a load balancer in your network environment, you can set up the virtual IP address on the load balancer instead of on a management server. The load balancer forwards the data to a management server according to its rules. Nodes that communicate with the management server using outbound-only connections are not supported together with load balancers.

Comparison of Cluster, Backup Server, and Server Pooling

The following table compares HPOM cluster, backup server, and server pooling environments.

Feature	Cluster	Backup server (hot standby)	Server pooling (hot standby - switch of virtual IP address)
Failover of consoles	Automatic	Manual	Manual Quick failover of consoles by moving the virtual IP address (can be automated).
Failover of agents	Automatic	Manual (or scripted) failover to new server with `opcragt -primmgr` (this may take some time to complete for large numbers of managed nodes).	Manual Quick failover of agents by moving the virtual IP address (can be automated).
Configuration synchronization	Not necessary	Need to regularly synchronize the configuration .	Need to regularly synchronize the configuration .
Disaster recovery	Cluster nodes must be close together, which means no disaster recovery.	Backup server can be located remotely. Continuous operation is possible even when the primary site is completely unavailable.	All servers must be in the same subnet, which means they need to be close together and thus do not provide for disaster recovery.
Data corruption	Data corruption is possible (if data is corrupted on the shared disk, it is corrupted on all cluster nodes).		
Load balancing		Backup server can be used to share the consolesload (both servers are fully operational management servers).	Backup server can be used to share the consoles load (both servers are fully operational management servers).
Hardware cost	Higher hardware cost (special hardware is needed to provide for no single point of failure) .		

Requirements

The following requirements apply:

- Two or more physical servers.

  For details of the HPOM versions supported for server pooling, see the support matrix at HP Software Support Online.

- All physical servers *must* be located in the same subnet.

  Special network configuration is required for Windows Server 2008 environments.

- One or more virtual interfaces.
- HTTPS agents.

  Server pooling is *not* supported with DCE agents.

Limitations

The following limitations apply:

- 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.
- Clustered HPOM management servers cannot be added to an HPOM server pool.
- 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

Setup details

A basic server pooling setup includes two physical servers and one virtual interface:

- You need 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. The management servers are configured as backup servers with buffered message forwarding set up between them. Each message is forwarded from an active server to a backup server.
- A virtual interface (V), which belongs to only one physical server at a time.

Flexible management setup

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 an agent-based flexible management policy, which must be deployed to all nodes.

M1 and M2 entries in the flexible management policy are required for the server-to-agent communication (configuration deployment, action execution). An entry for the virtual interface is required for the agent-to-server communication (message sending).

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

You can deploy the configuration from both servers (M1 and M2) if you keep the configuration data on both servers synchronized.

Message forwarding setup

M1 and M2 are set up to allow HTTPS-based buffered message forwarding from one server to another. This is defined with a server-based flexible management policy. After failover, you do not need 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, see Server-based Flexible Management Policies.

Configuration synchronization

The node configuration in the database should be identical on all servers. In addition, node groups, policies, policy groups, and tools that are used to manage the nodes should be identical. It is therefore recommended that you periodically exchange the configuration between the management servers.

Connecting Web consoles

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

Switchover behavior

When a switch occurs, the virtual interface is transferred from the primary manager to the backup server. Web consoles 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.

Server pooling scenarios

The following scenarios 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.

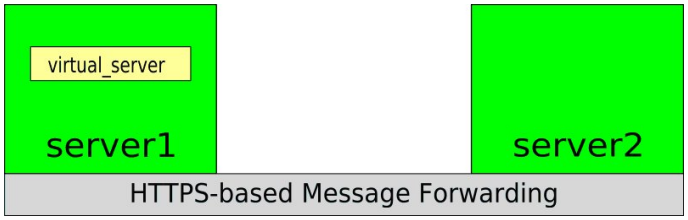

Scenario 2

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 Web consoles connected to the `virtual_server2`.

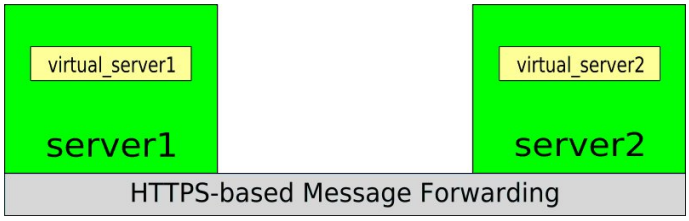

Scenario 3

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 you detect a high server load on the physical server `server1`, switch one of the virtual interfaces from that server to another server. When the load on the physical server `server1` decreases, switch the virtual interface back to the `server1`.

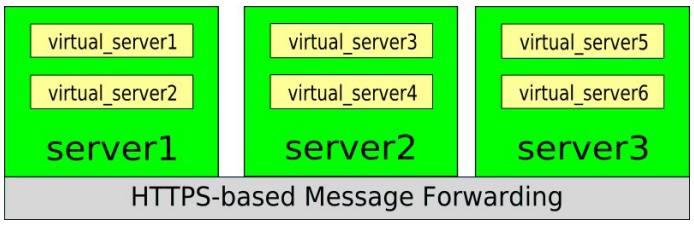

Configuring server pooling

The following procedure sets up two physical servers with one virtual interface. Using this procedure, you can easily set up more than two physical servers, and more than one virtual interface.

The following configuration changes are necessary:

1. "Install the management server" (on page 12).

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

2. "Configure the virtual interface" (on page 13).

   Configure the virtual interface.

3. "Configure the primary manager" (on page 15).

   Add a virtual interface and all physical servers as responsible managers in the responsible manager policy file.

4. "Configure message forwarding" (on page 17)

   Configure all HPOM servers in a backup server scenario with buffered message forwarding set up between them.

5. "Exchange trusted certificates" (on page 18).

   Exchange the servers' trusted certificates.

6. "Configure managed nodes" (on page 19).

   Set the virtual interface as a primary manager on all managed nodes.

If a physical server with a virtual interface is temporarily inaccessible for any reason, you can switch the virtual interface to another physical server.

Install the management server

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

Configure the virtual interface

Configure the virtual interface as described below. If your nodes communicate with the management server using outbound-only communication, you must also create a new OV resource group on both physical servers.

To configure the virtual interface

1. Activate the V virtual interface on the physical server M1 by adding the V virtual interface's IP address to the network interface adapter of M1 (using the netsh command, for example).
2. Invent a dummy core ID, for example `cf600f32-0aba-7532-1e83-f61cdcefb5d7`.
3. If your nodes communicate with the management server using outbound-only communication, continue with "To create an OV resource group for the virtual interface" (on page 13). Otherwise continue with "Configure the primary manager" (on page 15).

To create an OV resource group for the virtual interface

Perform the following steps only if your nodes communicate with the management server using outbound-only connections. Configure all outbound-only settings with the `-ovrg virt` option. This configures a new OV resource group named `virt`. See the *HPOM Firewall Concepts and Configuration Guide* for more information.

1. Create a new OV resource group for the V virtual interface. On the M1 physical server, enter the following command:

   ```
   md "%OvShareDir%\virt\conf\xpl\config"
   ```

   ```
   md "%OvShareDir%\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.
3. On the physical server M2, create the same OV resource group `virt` and set the same `OvCoreId`. Enter the following commands:

   ```
   md %OvShareDir%\virt\conf\xpl\config
   ```

   ```
   md %OvShareDir%\virt\datafiles\xpl\config\jobs
   ```

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

   where:
   `GUID` is the GUID contained in `C:\tmp\virt_coreid.txt`.
4. 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.

5. 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 on the physical server M1:

   ```
   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 on the physical server M2:

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

   The certificate and trusted certificates should be listed.

6. 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.

7. Activate the V virtual interface on the physical server M1 by adding the V virtual interface's IP address to the network interface adapter of M1.

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

   ```
   ovbbccb -start virt
   ```

Configure the primary manager

Agent-based flexible management policies enable you to configure managed nodes to send messages to different management servers based on time and message attributes.

To configure the primary manager

1. Create a new agent-based flexible management policy.
    a. Add the physical servers and the virtual interface (M1, M2, V) as responsible managers to the policy.
    b. Add the core IDs of all responsible managers, which are returned by executing the following command on each physical system:

       ```
       ovcoreid
       ```

    c. Add the core ID of the virtual interface. The core ID of the virtual interface is the dummy GUID created in "To configure the virtual interface" (on page 13). Alternatively, if your nodes communicate with the management server using outbound-only communication, the core ID of the virtual interface is the GUID that has been created and written to the file in "To create an OV resource group for the virtual interface" (on page 13).

   Example of the flexible management policy:

```
#
# Configuration policy
# defines action-allow managers
# messages are always sent 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. *Optional.* To check the policy's syntax, click **Check Syntax**. A message appears, which gives details of any errors.
3. Save the policy. Specify a name like "OM Server Pool Authorization" to find the policy easier later. Close the policy editor.

   Tip: Assign the policy to a policy group so that you can download the policy with ovpmutil. (You can then upload the policy on all management servers in the pool and thereby synchronize the configuration of all servers.)
4. Deploy the policy to the nodes that you want to configure.

Configure message forwarding

Server-based flexible management enables you to forward messages between multiple management servers.

To configure message forwarding

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
      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. *Optional.* To check the policy's syntax, click **Check Syntax**. A message appears, which gives details of any errors.

3. Save the policy. Specify a name like "OM Server Pool Message Synch" to find the policy easier later. Close the policy editor.

   Tip: Add the policy to a policy group so that you can download the policy with ovpmutil. (You can then upload the policy on all management servers in the pool and thereby synchronize the configuration of all servers.)

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

5. Make the same policy available on the physical server M2 and deploy it to M2.You can either create the same policy again on M2 or use the ovpmutil command line tool to download the policy on M1 and upload it on M2.

6. Exchange the node configurations between the management servers.

   Management servers immediately discard all messages from unknown nodes. Therefore, before a management server can accept forwarded messages, you must upload the appropriate node configurations.

   Tip: Alternatively, you can configure external nodes on management server M1 and M2, which can represent a range of nodes without the need to configure each individual managed node. This

option enables a management server to accept messages that originate from the nodes, but is not suitable if you need to manage the nodes from this management server. For example, you cannot start commands or launch tools on an external node.

Exchange trusted certificates

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

To exchange trusted certificates

1. On server M1, execute the following command:

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

2. Transfer the file `C:\tmp\M1.cer` to server M2.
3. On server M2, execute the command:

   ```
   ovcert -importtrusted -file <M1_cer> -ovrg server
   ```

   where:
   <M1_cer> 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.cer -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.cer` to server M1.
6. On server M1, execute the command:

   ```
   ovcert -importtrusted -file <M1_M2_cer> -ovrg server
   ```

   where:
   <M1_M2_cer> 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.

Configure managed nodes

To install new agents

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

1. Define the policy group that contains the flexible management policy to be auto-deployed to all new managed nodes. (This is the agent-based flexible management policy that you created in "Configure the primary manager" (on page 15). You configure auto-deployment in the property dialog of the root node group in the node editor.
2. Edit the agent installation default settings that you want the management server to apply when it installs agents remotely. (You can also use these settings for manual HTTPS agent installations by creating an agent profile.)

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

   a. 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`

   b. Add the namespace and the primary manager specification to the `agent_install_defaults.cfg` file on both physical serves as follows:

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

      where:
      <V_virtual_interface> is the name of the V virtual interface.
3. Install the agent software on the managed nodes:
   - Remote agent installations
     i. Install the agent remotely (using the agent installation defaults).
   - Manual agent installations
     i. On the physical server M1, download the agent-based flexible management policy, enter the following:

        `ovpmutil cfg pol dnl <download directory> /pnn <primary node name>`

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

        The output is located in *<download directory>*.

        Transfer the policy header and data files to the nodes that you plan to install and store them in a temporary folder.

     ii. Create an agent profile. Any settings defined in the `agent_install_defaults.cfg` file are also added to the agent profile.

     iii. To validate your changes to the agent installation defaults, execute the following command:

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

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

```
eaagt:
OPC_PRIMARY_MGR=<V_virtual_interface>
```

iv. Manually install HTTPS agents with a profile on each managed node but do not yet run the `opcactivate` command to activate the agent.

v. Upload the agent-based flexible management policy, enter the following:

```
ovpolicy -install -dir <directory>
```

vi. Activate the agent on the node, enter the following:

```
opcactivate -srv <V_virtual_interface>
```

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

To configure existing agents

On existing HTTPS agents, follow the procedure below:

1. Deploy 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 a `MSGTARGETRULE` rule, so that the virtual interface is used as a target for all messages.

The following is an example of the relevant part of the flexible management policy that you created in "Configure the primary manager" (on page 15). Add the following lines at the end of the flexible management policy and deploy the updated policy to the nodes.

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

Note: 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 theflexible 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.

Move a virtual interface to another 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

1. *Outbound-only communication only.* Stop the `virt` OV resource group on the physical server M1. Enter the following:

   ```
   ovbbccb -stop virt
   ```

2. Stop the V virtual interface by removing the virtual IP address from the physical server M1.

To enable the V virtual interface on the physical server M2

1. Start the V virtual interface on the physical server M2 by adding the virtual IP address to its network interface.
2. *Outbound-only communication only.* Start `virt` OV resource group on the physical server M2. Enter the following:

   ```
   ovbbccb -start virt
   ```

Server pooling in cluster environments

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

Including clustered HPOM management servers in a server pool is currently not supported.

Migrate from existing backup server environments

You can convert an existing backup server environment into a server pooling environment. This can be performed in one step, or you can use a phased approach.

To migrate to server pooling

1. Configure a virtual interface. See "Configure the virtual interface" (on page 13) for more information.
2. Add the virtual interface to the existing flexible management policy and then deploy the changed 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 that describes how to change primary manager setting for the existing HTTPS agents from section "Configure managed nodes" (on page 19).

   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.

Additional useful information

DNS configuration

- All management servers as well as the virtual interface should use static IP addresses.
- 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.

Network configuration on Windows Server 2008

In Windows Server 2008, when you activate an IP address using the netsh command-line utility, the resulting gratuitous ARP[1] message contains 0.0.0.0 as the sender's IP address. Most network devices regard these gratuitous ARPs as invalid and therefore cannot update their ARP tables. The ARP tables are not refreshed until the next valid ARP request arrives.

Therefore, in a Windows Server 2008 server pooling environment, when the virtual interface moves to another physical server, the agents on the managed nodes cannot reach physical server M1 any more and start buffering and the console cannot connect to the management server. As soon as the ARP tables are updated and the IP address of the virtual interface is associated with the MAC address of physical server M2, the connection problems are resolved.

To avoid interruptions, run the following commands on the computer that just acquired `<virtual_IP>` (for example, physical server M2) to manually update the ARP tables:

```
arp -d *
```

```
ping -S <virtual_IP> <target_name>
```

`<virtual_IP>` is the IP address of the virtual interface.

`<target_name>` can be any address on the subnet, but not one on the same computer. It can be any other computer or a gateway, or even a non-existing address.

The command `arp -d *` deletes all entries in the ARP tables of the computer where it is run. Before the computer can send an ICMP[2] ping, it must find out the MAC address of the target computer and therefore sends an ARP request. The command `ping -S <virtual_IP> <target_name>` sets the sender IP address of the ARP request to `<virtual_IP>`. All computers on the subnet receive this broadcast and update their ARP tables accordingly.

Tip: Add the commands to the server pooling scripts immediately after the netsh command that adds the IP address.

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

[1]Address Resolution Protocol
[2]Internet Control Message Protocol

The default key store lists agent certificates. Certificates installed in the server key store are used by HPOM 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 synchronization 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 USB flash drive).

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 "Configure the virtual interface" (on page 13) 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" (on page 18) 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 server that needs to be configured for server pooling, perform the steps from the section "Configure the virtual interface" (on page 13) to "Configure managed nodes" (on page 19), 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 "Configure the virtual interface" (on page 13).
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" (on page 18). 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 "Configure managed nodes" (on page 19).