

HP Operations Manager

For the UNIX and Linux operating systems

Software Version: 9.1x

High Availability Through Server Pooling

Document Release Date: March 2014

Software Release Date: November 2013



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.

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 Notice

© Copyright 2008-2014 Hewlett-Packard Development Company, L.P.

Trademark Notices

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

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 registered trademark of Oracle and/or its affiliates.

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

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

UNIX® is a registered trademark of The Open Group.

Support

Visit the HP Software Support Online web site at: <http://www.hp.com/go/hpsoftwaresupport>

This web site provides contact information and details about the products, services, and support that HP Software offers.

HP Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support web 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 register for an HP Passport ID, go to:

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

To find more information about access levels, go to:

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

HP Software Solutions Now accesses the HPSW Solution and Integration Portal Web site. This site enables you to explore HP Product Solutions to meet your business needs, includes a full list of Integrations between HP Products, as well as a listing of ITIL Processes. The URL for this Web site is

<http://h20230.www2.hp.com/sc/solutions/index.jsp>

Contents

Contents	3
Chapter 1: Introduction to server pooling	5
Load balancing	6
Comparison of cluster, backup server, and server pooling	7
Requirements	9
Chapter 2: Setup details	10
Agent-based flexible management setup	10
Message forwarding setup	11
Configuration synchronization	11
Connecting Java GUIs	11
Connecting HPOM Configuration Value Packs and Administration UIs	11
Switchover behavior	12
Chapter 3: Server pooling scenarios	13
Scenario 1	13
Scenario 2	13
Scenario 3	14
Chapter 4: Configuring server pooling	15
Install the management server	15
Configure management server nodes	15
Add the virtual node to the Node Bank	16
Configure the virtual interface	17
Optional: Create an OV resource group for the virtual interface	19
Configure the primary manager	22
Exchange trusted certificates	24
Configure message forwarding	25
Configure managed nodes	27
Move a virtual interface to another physical server	31
To disable the V virtual interface on the physical server M1	31
To enable the V virtual interface on the physical server M2	32

Migrate from existing backup server environments	34
To migrate to server pooling	34
Automate switchover of the virtual interface	35
Server-based message forwarding	37
Forwarding messages to all servers in a backup server configuration	38
Forwarding messages to the virtual interface of a server pool	39
Load balancing	41
Management server—agent load balancing	41
Management server—user interface load balancing	42
Configure the load balancer	43
Configure HPOM	44

Chapter 1: Introduction to 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.

Tip: For information about configuring other management servers outside of the server pool to forward their messages to the server pool, see ["Server-based message forwarding" on page 37](#).

Load balancing

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, 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.

HP has tested the F5 BIG-IP load balancer in conjunction with server pooling. Other load balancers are supported but have not been tested. For details, see ["Load balancing" on page 41](#).

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 GUIs	Automatic	Manual It is possible to automate failover by providing a list of backup servers (OPC_JGUI_BACKUP_SRV).	Manual Quick failover of GUIs by moving the virtual IP address (can be automated).
Failover of agents	Automatic	Manual (or scripted) failover to new server with <code>opcragt -primmgr</code> (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 (using <code>opccfgdwn</code> and <code>opccfgupld</code>).	Need to regularly synchronize the configuration (using <code>opccfgdwn</code> and <code>opccfgupld</code>).
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 (unless a load balancer is used), 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).		

Feature	Cluster	Backup server (hot standby)	Server pooling (hot standby - switch of virtual IP address)
Load balancing		Backup server can be used to share the GUI load (both servers are fully operational management servers).	Backup server can be used to share the GUI load (both servers are fully operational management servers).
Hardware cost	Higher hardware cost (special hardware is needed to avoid 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:

<http://support.openview.hp.com/selfsolve/document/KM323488>.

- All physical servers must be located in the same subnet, unless a load balancer is used. With a load balancer, the physical servers can be in different subnets and thus geographically remote.
- *For server-based message forwarding to the virtual interface of a server pool.* All management servers must be running one of the following versions of HPOM:
 - 9.1x (with management server patch 9.10.200 or higher)

For more information, see "[Server-based message forwarding](#)" on page 37.

- One or more virtual interfaces.
- Environments with managed nodes with IPv6 addresses.

If your environment includes managed nodes with IPv6 addresses, make sure the following requirements are met:


- You need an IPv6 address for the virtual interface that the IPv6 only managed nodes know.
- You must make sure that the IPv6 address of the virtual interface is known on both management servers, for example by adding the following second line to the `/etc/hosts` file:

```
<V_IPv4_address> <FQDN> <hostname> # IPv4 virtual IP address  
<V_IPv6_address> <FQDN> <hostname> # IPv6 virtual IP address
```

Chapter 2: 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.



Agent-based 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 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). An entry for the virtual interface is required for the agent-to-server communication (message sending). For more information, see the *HP Operations Manager Administrator's Reference*, section "Flexible Management Configuration".

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 the `msgforw` file in the HPOM configuration directory. 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 the section “Flexible Management Configuration” in the *HP Operations Manager Administrator’s Reference*.

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.

You can upload all configuration changes (using `opccfgupld`) at runtime.

Connecting Java GUIs

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.

Connecting HPOM Configuration Value Packs and Administration UIs

HPOM Administration UIs can connect to the virtual interface. In some situations you may want to always connect to the physical server M1, for example in a backup server scenario. In an environment with a primary and a backup server, you may want to continue your configuration tasks on the primary server, even when the virtual interface has switched to the backup server. If you do not use the MIDAS Administration UI, configuration changes are not automatically synchronized between the management servers.

If you are maintaining HPOM with the MIDAS Administration UI, you can connect to any physical management server because your configuration changes are automatically synchronized with all management servers.

Switchover behavior

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.

Chapter 3: 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

The following figure 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


The following figure 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.



Scenario 3

The following figure 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`.



Chapter 4: 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 HPOM management server as a standalone server on each physical server (M1 and M2). See ["Install the management server" below](#).
2. Set up certificate servers in the HPOM environment. See ["Configure management server nodes" below](#).
3. Add the virtual node to the Node Bank. See ["Add the virtual node to the Node Bank" on the next page](#).
4. Configure the virtual interface. See ["Configure the virtual interface" on page 17](#).
5. Add a virtual interface and all physical servers as responsible managers in the responsible manager policy file. See ["Configure the primary manager" on page 22](#).
6. Exchange the servers' trusted certificates. See ["Exchange trusted certificates" on page 24](#).
7. Configure all HPOM servers in a backup server scenario with buffered message forwarding set up between them. See ["Configure message forwarding" on page 25](#).
8. Set the virtual interface as a primary manager on all managed nodes. See ["Configure managed nodes" on page 27](#).

If a physical server with a virtual interface is temporarily inaccessible for any reason, you can switch the virtual interface to another physical server (see ["Move a virtual interface to another physical server" on page 31](#)).

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 Installation Guide*.

Configure management server nodes

To set up certificate servers in the HPOM environment, see the *HP Operations Manager Administrator's Reference*, section "Security in Flexible Management Environments".

Add the virtual node to the Node Bank

Add the virtual node to the Node Bank as described below:

1. Create a new UUID for the V virtual interface, for example using the `uuidgen` command line tool.
2. 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)

MACH_BBC_HPUX_IPF32 (for HP-UX Itanium)

MACH_BBC_SOL_SPARC (for Solaris Sparc)

MACH_BBC_LX26RPM_X64 (for Linux)

<node_group_name> is `hp_ux` (for HP-UX), `solaris` (for Solaris), or `linux` (for Linux)

3. Set the UUID 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=<UUID>
```

where:

<UUID> is the UUID generated in step 1.

4. *Java GUI applets only.* Set the virtual node name for use by Java GUI applets started by a CGI script:

```
/opt/OV/bin/ovconfchg -ovrg server -ns opc -set VIRT_MGMT_SERVER <V_virtual_
interface>
```

<V_virtual_interface> is the fully qualified domain name (FQDN) of the V virtual interface.

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. See ["Optional: Create an OV resource group for the virtual interface" on page 19](#).

You can skip this section if a load balancer is used, in which case you configure the virtual IP address on the load balancer. See ["Load balancing" on page 41](#).

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

HP-UX and Linux management servers, IPv4 addresses only:

```
ifconfig <V_network_interface>:1 inet <V_IPv4_address> netmask <V_netmask_IPv4_value> up
```

HP-UX management servers, IPv4 and IPv6 addresses:

```
ifconfig <V_network_interface>:1 inet <V_IPv4_address> netmask <V_netmask_IPv4_value> up
ifconfig <V_network_interface>:2 inet6 <V_IPv6_address> prefix <V_netmask_IPv6_suffix> up
```

Linux management servers, IPv4 and IPv6 addresses:

```
ifconfig <V_network_interface>:1 inet <V_IPv4_address> netmask <V_netmask_IPv4_value> up
ifconfig <V_network_interface> inet6 add <V_IPv6_address>/<V_netmask_IPv6_suffix>
```

Solaris management servers, IPv4 addresses only:

```
ifconfig <V_network_interface>:1 plumb
ifconfig <V_network_interface>:1 inet <V_IPv4_address> netmask <V_netmask_IPv4_value> up
```

Solaris management servers, IPv4 and IPv6 addresses:

```
ifconfig <V_network_interface>:1 plumb
ifconfig <V_network_interface>:1 inet <V_IPv4_address> netmask <V_netmask_IPv4_value> up
ifconfig <V_network_interface>:2 inet6 plumb
ifconfig <V_network_interface>:2 inet6 <V_IPv6_address>/<V_netmask_IPv6_suffix> up
```

where:

<V_network_interface> is the network interface used for the V virtual interface's IP address

(for example, `lan0` for HP-UX, `hme0` for Solaris, or `eth0` for Linux).

`<V_IPv[4|6]_address>` is the IP address of the V virtual interface.

`<V_netmask_IPv4_value>` is the netmask value for the V virtual interface.

`<V_netmask_IPv6_suffix>` is the subnet mask of the IPv6 address of the V virtual interface.

2. If your nodes communicate with the management server using outbound-only communication, continue with ["Optional: Create an OV resource group for the virtual interface" on the next page](#). Otherwise continue with ["Configure the primary manager" on page 22](#).

Optional: 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:

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

```
/opt/OV/sbin/xpl/ovinit.sh -dependencies OvOUSvr -ovrg virt
```

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:

```
/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.

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

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

```
/opt/OV/sbin/xpl/ovinit.sh -dependencies OvOUSvr -ovrg virt
```

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

4. 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.

5. 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 on the physical server M1:

```
/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 on the physical server M2:

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

The certificate and trusted certificates should be listed.

6. On both physical servers, stop the processes:

```
/opt/OV/bin/ovc -kill
```

7. Bind the V virtual interface's IP address to the new OV resource group `virt`. On both physical servers, enter one of the following commands depending on whether your environment includes managed nodes with IPv6 addresses:

IPv4 addresses only:

```
/opt/OV/bin/ovconfchg -ovrg virt -ns bbc.cb -set SERVER_BIND_ADDR <V_IPv4_address> -set SERVER_PORT 383
```

IPv4 and IPv6 addresses:

```
/opt/OV/bin/ovconfchg -ovrg virt -ns bbc.cb -set SERVER_BIND_ADDR "<V_IPv4_address,V_IPv6_address>" -set SERVER_PORT 383
```

where:

`<V_IPv[4|6]_address>` is the IP address of the V virtual interface.

8. *Linux only.* Bind the communication broker to the local IP address of the physical server. On both physical servers, enter one of the following commands depending on whether your environment includes managed nodes with IPv6 addresses:

IPv4 addresses only:

```
/opt/OV/bin/ovconfchg -ns bbc.cb -set SERVER_BIND_ADDR <Local_IPv4_address>
```

IPv4 and IPv6 addresses:

```
/opt/OV/bin/ovconfchg -ns bbc.cb -set SERVER_BIND_ADDR "<Local_IPv4_
address,Local_IPv6_address>"
```

where:

<Local_IPv[4|6]_address> is the IP address of the physical server.

9. On both physical servers, start the processes:

```
/opt/OV/bin/ovc -start
```

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

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

Configure the primary manager

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 policy. 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 policy.

To verify the deployed policies, use the following command:

```
/opt/OV/bin/ovpolicy -list
```

This should return the following message:

```
mgrconf "OV0 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 policy.

To verify the deployed policies, use the following command:

```
/opt/OV/bin/ovpolicy -list
```

This should return the following message:

```
mgrconf "OV0 authorization" enabled 1
```

Note: The responsible manager policy (`mgrconf`) must also be deployed to all managed nodes. See ["Configure managed nodes" on page 27](#) for more information.

Exchange trusted certificates

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

1. On server M1, execute the following command:

```
/opt/OV/bin/ovcert -exporttrusted -file /tmp/M1.cer -ovrg server
```

2. Transfer the file /tmp/M1.cer to server M2.

3. On server M2, execute the command:

```
/opt/OV/bin/ovcert -importtrusted -file /tmp/M1.cer -ovrg server
```

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

```
/opt/OV/bin/ovcert -exporttrusted -file /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 /tmp/M1_M2.cer to server M1.

6. On server M1, execute the command:

```
/opt/OV/bin/ovcert -importtrusted -file /tmp/M1_M2.cer -ovrg server
```

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

```
/opt/OV/bin/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:

```
/opt/OV/bin/ovcert -updatetrusted
```


Configure message forwarding

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

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. Copy the `/etc/opt/OV/share/conf/OpC/mgmt_sv/respmgrs/msgforw` file from the M1 to the same location on the M2.
3. Activate the `msgforw` template on both physical servers. Call `ovconfchg` without any parameters, use the following command:

```
/opt/OV/bin/ovconfchg
```

4. Verify the `msgforw` template:
 - a. Start the GUI and open a message browser on both physical servers.
 - b. Send a message from the physical server M1:

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

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

Configure managed nodes

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

1. 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 `bbc_inst_defaults` file on both physical servers. The file is located in the following directory:

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

- b. Add the namespace and the primary manager specification to the `bbc_inst_defaults` file on both physical servers as follows:

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

where:

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

2. 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 `mgrconf` policy, enter the following:

```
/opt/OV/bin/OpC/opctmp1dwn <node>
```

where:

<node> is the name of the managed node.

The output is located in `/var/opt/OV/share/tmp/OpC/<Long_hostname_of_node>`.

The `mgrconf` policy is stored in the files

`a4c90b9a-f15e-11d6-9032-001083fdff5e_data` and `a4c90b9a-f15e-11d6-9032-001083fdff5e_header.xml`. Transfer both files to the nodes that you plan to install and store them in the directory `/tmp/mgrconf`. (The downloaded policy is node-independent and can be used on any node.)

- ii. Create an agent profile. Any settings defined in the `bbc_inst_defaults` file are also added to the agent profile.
- iii. Manually install HTTPS agents with a profile on each managed node but do not yet run the `opcactivate` command to activate the agent. For more information on installing an agent manually, see the HP Operations Agent documentation.

- iv. Upload the `mgrconf`, enter the following:

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

For example:

```
ovpolicy -install -dir /tmp/mgrconf
```

- v. 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.

- Existing agent installations

To configure existing agent installations, follow the procedure below.

3. Verify the deployed policies. 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 policy. 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 policies, 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 policies, 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_<br>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 `allnodes` file with a MSGTARGETRULE rule, so that the virtual interface is used as a target for all messages. For more information, see the *HP Operations Manager Administrator's Reference*.

The following is an example of the relevant part of the `allnodes` file that you created in "Configure the primary manager" on page 22. Add the following lines at the end of the `allnodes` file and deploy the updated file to the nodes.

```
MSGTARGETRULECONDS
  MSGTARGETMANAGERS
    MSGTARGETMANAGER
```

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

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 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.

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.

You can skip this section if a load balancer is used, in which case you configure the load balancer to now point to the other physical server. See "[Load balancing](#)" on page 41.

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:

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

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

HP-UX management servers, IPv4 addresses only:

```
ifconfig <V_network_interface>:1 inet 0.0.0.0 down
```

HP-UX management servers, IPv4 and IPv6 addresses:

```
ifconfig <V_network_interface>:1 inet 0.0.0.0 down  
ifconfig <V_network_interface>:2 inet6 ::
```

Linux management servers, IPv4 addresses only::

```
ifconfig <V_network_interface>:1 down
```

Linux management servers, IPv4 and IPv6 addresses:

```
ifconfig <V_network_interface>:1 down  
ifconfig <V_network_interface> inet6 del <V_IPv6_address>/<V_netmask_IPv6_<br>suffix>
```

Solaris management servers, IPv4 addresses only:

```
ifconfig <V_network_interface>:1 unplumb
```

Solaris management servers, IPv4 and IPv6 addresses:

```
ifconfig <V_network_interface>:1 unplumb  
ifconfig <V_network_interface>:2 inet6 unplumb
```

where:

<V_network_interface> is the network interface used for the V virtual interface's IP address

(for example, `lan0` for HP-UX, `hme0` for Solaris, or `eth0` for Linux).
<V_IPv6_address> is the IP address of the V virtual interface.
<V_netmask_IPv6_suffix> is the subnet mask of the IPv6 address of the V virtual interface.

To enable the V virtual interface on the physical server M2

1. Start the V virtual interface on the physical server M1. Enter the following:

HP-UX and Linux management servers, IPv4 addresses only:

```
ifconfig <V_network_interface>:1 inet <V_IPv4_address> netmask <V_netmask_IPv4_value> up
```

HP-UX management servers, IPv4 and IPv6 addresses:

```
ifconfig <V_network_interface>:1 inet <V_IPv4_address> netmask <V_netmask_IPv4_value> up
ifconfig <V_network_interface>:2 inet6 <V_IPv6_address> prefix <V_netmask_IPv6_suffix> up
```

Linux management servers, IPv4 and IPv6 addresses:

```
ifconfig <V_network_interface>:1 inet <V_IPv4_address> netmask <V_netmask_IPv4_value> up
ifconfig <V_network_interface> inet6 add <V_IPv6_address>/<V_netmask_IPv6_suffix>
```

Solaris management servers, IPv4 addresses only:

```
ifconfig <V_network_interface>:1 plumb
ifconfig <V_network_interface>:1 inet <V_IPv4_address> netmask <V_netmask_IPv4_value> up
```

Solaris management servers, IPv4 and IPv6 addresses:

```
ifconfig <V_network_interface>:1 plumb
ifconfig <V_network_interface>:1 inet <V_IPv4_address> netmask <V_netmask_IPv4_value> up
ifconfig <V_network_interface>:2 inet6 plumb
ifconfig <V_network_interface>:2 inet6 <V_IPv6_address>/<V_netmask_IPv6_suffix> up
```

where:

<V_network_interface> is the network interface used for the V virtual interface's IP address (for example, `lan0` for HP-UX, `hme0` for Solaris, or `eth0` for Linux).
<V_IPv[4/6]_address> is the IP address of the V virtual interface.

<V_netmask_IPv4_value> is the netmask value for the V virtual interface.

<V_netmask_IPv6_suffix> is the subnet mask of the IPv6 address of the V virtual interface.

2. *Outbound-only communication only.* Start virt OV resource group on the physical server M2.
Enter the following:

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

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 17](#) for more information.
2. Add the virtual interface to the existing responsible manager file and then distribute the changed responsible manager file 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 27](#).

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.

Automate switchover of the virtual interface

You can use the following technologies to automate the monitoring and switching of the virtual interfaces:

- **Failover clusters**

When using a failover cluster such as HP Serviceguard or the Red Hat Cluster Suite to switch the virtual interface, consider the following aspects:

- Server pooling is *not* supported in case you have an existing HPOM installation in a cluster, as described in the *HP Operations Manager Installation Guide*, Chapter “Installing HPOM in a Cluster Environment”.
- The HPOM management server must be installed as a standalone server on both physical servers. When installing HPOM management servers in a cluster environment, the following question displays:

```
Run HPOM Server as a HA resource group {exit,back,?,y|n,"n"} ?
```

Enter n.

- Instead of the `ifconfig` command, use the cluster specific commands to activate and deactivate the virtual interface, as described in the section ["Configure the virtual interface" on page 17](#). For example, you can use the `cmmodnet` command for HP Serviceguard.
- Switching an IP address is a standard part of the cluster HARG, and you usually need only to specify the virtual IP address in the configuration scripts or settings.

HARG is an HPOM management server concept and is equivalent to a package in HP Serviceguard, a service group in VERITAS Cluster System, a resource group in Sun Cluster, and a service in Red Hat Cluster Suite.

- *Outbound-only communication only*. 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 ["Move a virtual interface to another physical server" on page 31](#) for more information.
- 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/utills/ha/ha_mon_ovserver` script in a loop. The `ha_mon_ovserver` script returns 0 if all processes are running, otherwise it returns 1.

- **HPOM High Availability (HA) Manager**

HA Manager is a light-weight solution that allows the configuration of an automatic failover of the virtual interface in a server pooling setup in a similar way as in a regular failover cluster. HA Manager is not an additional cluster software such as the HP Serviceguard or the Red Hat

Cluster Suite, but it is an alternative. It represents a cluster without special hardware with redundancy and a shared disk.

For more information about HA Manager in HPOM 9.1x, see the *HP Operations Manager High Availability Manager* white paper, which is available at <http://h20230.www2.hp.com/selfsolve/manuals>. HA Manager in HPOM 9.20 is described in the HP Operations Manager Administrator's Reference.

Server-based message forwarding

The default setup for server pooling consists of two or more management servers and a virtual interface combined to a server pool. Managed nodes (or Java GUIs) connect to the virtual interface rather than to the physical servers in the pool.

If your environment includes a hierarchy of management servers, you can configure the second (or higher) level of managers as members of a server pool, and configure the first level of managers to forward messages to the virtual interface of the server pool.

However, in some environments it is more appropriate to set up a backup server configuration, in which the first-level managers forward messages to both the primary and the backup server. HP recommends the following configurations:

- Two levels of management servers

If your flexible management environment includes only two levels of management servers, use a regular backup server setup and forward messages to both servers. For more information, see ["Forwarding messages to all servers in a backup server configuration" on the next page](#).

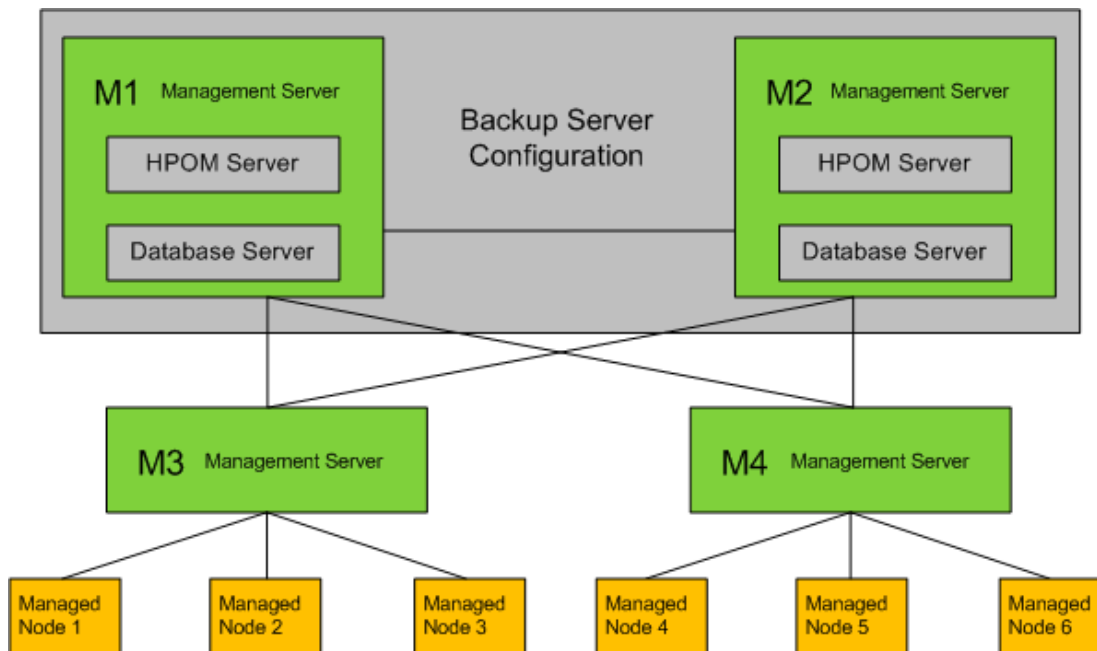
- Three or more levels of management servers

If your flexible management environment includes three or more levels of management servers, set up a server pool and forward messages to the virtual interface. For more information, see ["Forwarding messages to the virtual interface of a server pool" on page 39](#).

Forwarding messages to all servers in a backup server configuration

In the following figure, M3 and M4 forward their messages to the primary and backup server (M1 and M2) of a backup server configuration.

Forwarding messages in a backup server configuration



Setting up a backup server configuration has the following advantages:

- Because messages are sent to all management servers, the management servers always have the same set of messages. (No missing messages that were already delivered to the primary server but not forwarded to the standby before the virtual interface was switched.)
- The management servers always have the same set of messages. When the primary server goes down or in maintenance, the backup server already has the message. Messages for the primary server are buffered by the original server and sent when the primary server is up again (and at that time the primary server also receives message change events for those messages, if any).

Disadvantages:

- In a message-forwarding hierarchy with three levels or more, the toplevel manager receives all messages at least twice (all management servers send the same messages).

Although HPOM detects and discards duplicate messages, performance is reduced because twice the number of messages must be processed.

- The `msgforwfile` of the sending servers must contain all primary and backup management servers. In addition, all primary and backup servers must be set up as nodes in the node bank of the sending server.

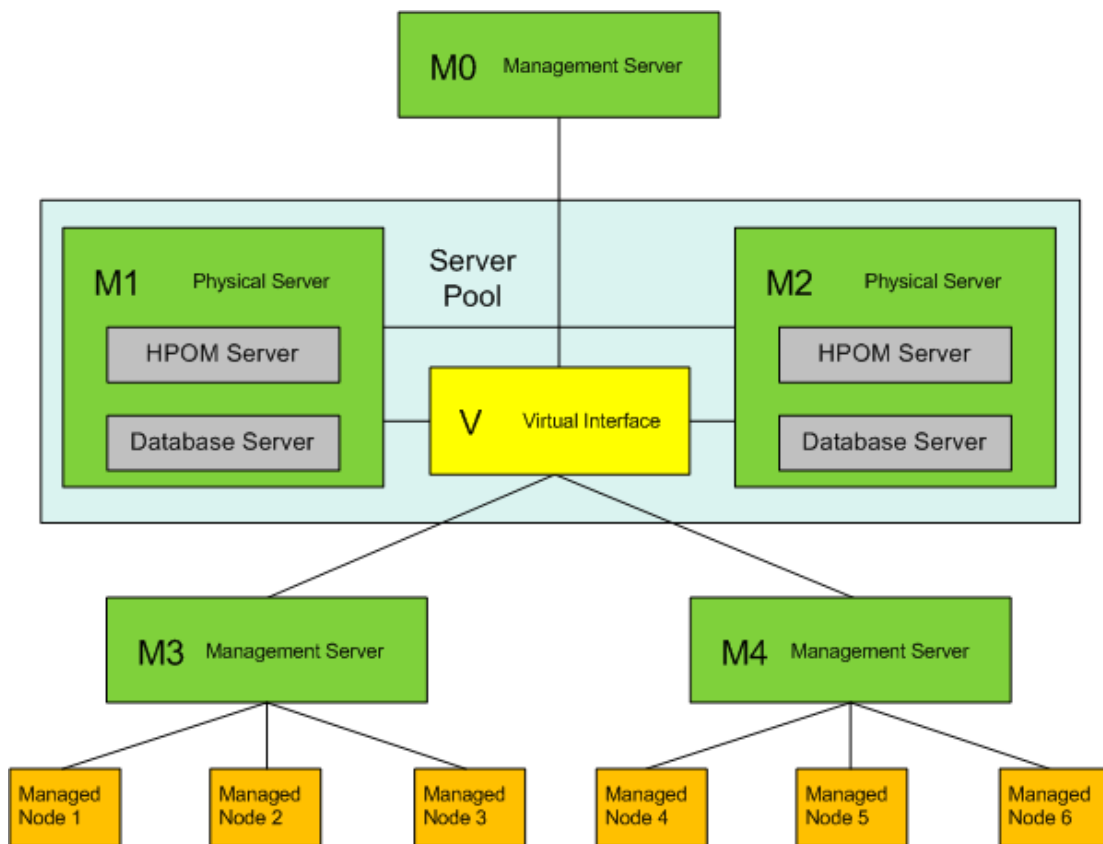
To configure server-based message forwarding to the physical servers, add all primary and backup servers to the `msgforw` file of the sending server.

Forwarding messages to the virtual interface of a server pool

The following figure shows five management servers. M0 is at the top of the message-forwarding hierarchy. M1 and M2 are part of the server pool. M3 and M4 forward their messages to the virtual interface.

Note: You can set up M0, M3, and M4 as regular management servers as shown in figure "Forwarding messages to the virtual interface" below, or also as members of a server pool or failover cluster.

Forwarding messages to the virtual interface



Forwarding messages to the virtual interface has the following advantages:

- Because the sending server has no knowledge of the physical servers in the pool, you can add other servers to the pool without having to update the configuration of the sending server.
- In a message-forwarding hierarchy with more than three levels, server pooling on the middle level avoids duplicate messages on the third level because the toplevel manager (M0 in the example "[Forwarding messages to the virtual interface](#)" on the previous page) receives all messages only once from the virtual interface.

Disadvantages:

- If the currently active physical server in the pool goes out of service, some messages that were already received but not yet processed and forwarded are delayed until that server is running again.
- In case of failure or maintenance, you must actively (or automatically) switch the virtual interface.

To configure server-based message forwarding to the virtual interface, add the virtual interface to the `msgforw` file of the sending server. (Do not add the physical servers to the `msgforw` file.)

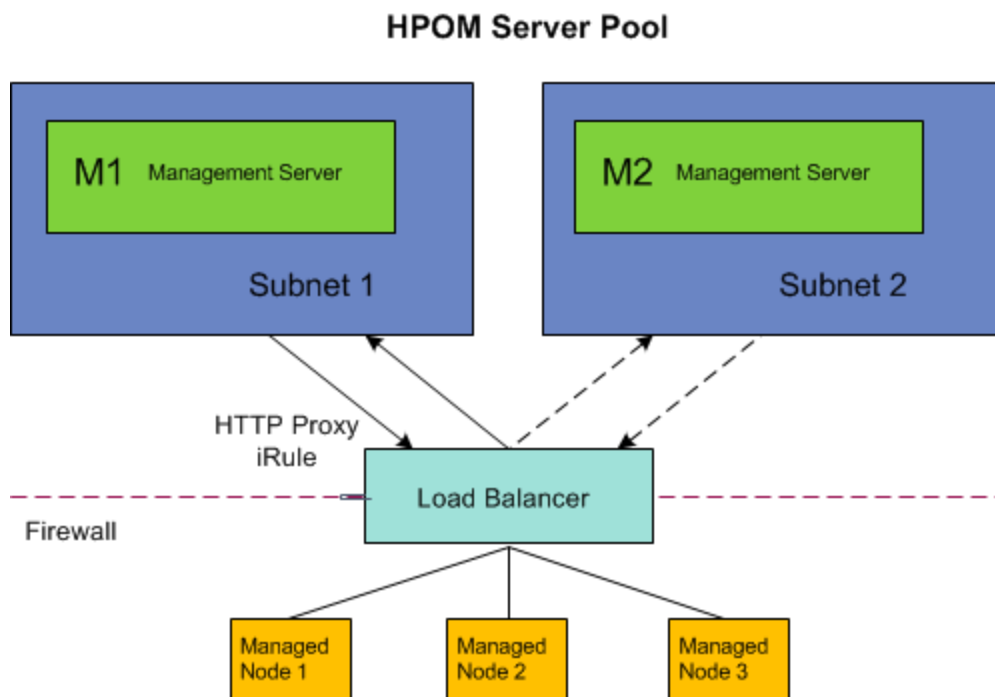
Load balancing

Load balancing divides the work load among several computers. As a result, system performance and availability increases. You can use a load balancer as a proxy between the managed nodes and the management server, between the user interface and the server, as well as between multiple management servers. The load balancer routes incoming traffic to the server depending on the load balancing method configured for the pool.

One of the advantages of using a load balancer in a server pooling setup is that the physical servers do not need to be in the same subnet and can thus be geographically remote.

Management server—agent load balancing

The following figure shows a load balancer that is configured to use the Least Connections load balancing method. The load balancer sends messages from the managed nodes to the management server with the fewest open connections (M1 in this example). M2 receives messages only when M1 already has open connections and is therefore considered busy.




If a firewall separates the management server and the managed nodes, set up an iRule to configure the load balancer as HTTP proxy. Data sent from M1 to the managed nodes (action requests, for example) is then routed by the load balancer to the appropriate managed nodes.

Note: Managed nodes send action responses to the physical management server. If the load balancer switches to another management server before an action response can be sent, the response is buffered on the node and delivered when the load balancer communicates with the server again.

Management server—user interface load balancing

You can also use load balancing to manage communication requests from the user interface to the management server.

In the following example, the user interfaces communicate with their management servers through dedicated load balancers. If one of the management servers is too busy or fails, the load balancer redirects the communication to the other management server.



Configure the load balancer

HP has tested the F5 BIG-IP load balancer in conjunction with server pooling. Other load balancers are supported but have not been tested.

The following steps describe how to configure the F5 BIG-IP load balancer. For details on how to configure other load balancers, see the documentation provided by the load balancer's manufacturer.

1. Use the BIG-IP Configuration Utility to set up a server pool:
 - a. Configure health monitoring of the pool. For example, monitor the pool using a TCP connection every 120 seconds (`tcp_120`).
 - b. Select the load balancing method that best suits your needs.

Note: To avoid issues with action responses being sent to the wrong physical server and issues with the order of messages, consider choosing one of the following recommendations:

- o Use the load balancer for failover and manual switching only. Configure the load balancer to direct all nodes to the active server.
 - o Alternatively, use source address and SSL persistence so that the same managed nodes always communicate with the same physical server.
 - c. Add the HPOM management servers as members to the pool.
2. Use the BIG-IP Configuration Utility to set up a virtual server:
 - a. In Service Port, type the port number on the management server you want the load balancer to use, for example 383 for the communication broker.
 - b. Assign the HPOM server pool to the virtual server.

Configure HPOM

- If a firewall separates the management server and the managed nodes, and only the load balancer is allowed to pass through the firewall, configure the management server to insert the virtual server in the action request fields of a message. The agent then sends the action response to the virtual server instead of to the physical management server from which the action request originated.

On each physical management server in the server pool, type the following commands:

```
ovconfchg -ovrg server -ns opc -set OPC_MGMTSV_IPADDR_ON_LOAD_BALANCER  
<virtual server IP address>
```

```
ovconfchg -ovrg server -ns opc -set OPC_MGMTSV_NAME_ON_LOAD_BALANCER <virtual  
server FQDN>
```

- Configure the agents to communicate with the virtual server:
 - a. On each managed node, use `ovconfchg` to set the `OPC_PRIMARY_MGR` parameter to the virtual server, for example:

```
ovconfchg -ns eaagt -set OPC_PRIMARY_MGR <virtual server FQDN>
```
 - b. *Optional.* Use `bbcutil -ping <virtual server FQDN>` to verify the connection.
 - c. *Optional.* HP Operations Agent 11.02 and higher by default set `AUTO_CONNECTION_CLOSE_INTERVAL` to 60 seconds. Configure a smaller value if required. For more information about the `AUTO_CONNECTION_CLOSE_INTERVAL` parameter, see the HP Operations Agent documentation.
- Configure the user interface to communicate with the virtual server.

When starting the user interface, connect to the virtual server instead of the management server. For example, in the Java GUI login screen, type the name of virtual server.

