

HP Automation Insight



HP Vertica Cluster Load Balancing and Failover

Contents

| | |
|---|----------|
| Configure HP Vertica Nodes for Load Balancing | 2 |
| Before You Begin | 2 |
| Set Up VIP Using the Loopback Interface | 2 |
| Disable Address Resolution Protocol (ARP) | 3 |
| Install the HP Vertica IPVS Load Balancer Package..... | 4 |
| Install the HP Vertica IPVS Load Balancer Package on RHEL 5.x | 4 |
| Install the HP Vertica IPVS Load Balancer Package on RHEL 6.x | 4 |
| Configure the HP Vertica IPVS Load Balancer | 5 |
| About Public and Private IPs..... | 5 |
| Set Up the HP Vertica IPVS Load Balancer Configuration File on the Master and Slave Directors | 6 |

© Copyright 2014 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

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.

Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation. AMD is a trademark of Advanced Micro Devices, Inc. Intel and Xeon are trademarks of Intel Corporation in the U.S. and other countries. Oracle and Java are registered trademarks of Oracle and/or its affiliates.

This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>).

This document describes how to set up load balancing in your HP Vertica database cluster database for HP Automation Insight.

For complete information on configuring load balancing in your HP Vertica cluster, see the *HP Vertica Analytics Platform 6.1.x Administrator's Guide*.

Configure HP Vertica Nodes for Load Balancing

After you have installed your HP AI core server, the HP Vertica database cluster, and the data miner, you can set up your HP Vertica cluster to load balance and configure failover between the HP Vertica nodes. You must have at least 3 nodes in your HP Vertica cluster and a K-safety level of 1 to set up your HP Vertica cluster to load balance and set up a failover server.

In HP Vertica, load balancing supports multiple client connections through a single Virtual IP (VIP) address that all nodes in a cluster share. This is useful for balancing incoming client requests across nodes, as well as preventing node exclusion from clients in the case of node failure.

The IP Virtual Server (IPVS) provides network-protocol-level connection load balancing. When used with an HP Vertica database cluster, IPVS is installed on two database hosts.

IPVS is made up of the following components:

- The Virtual IP (VIP): The IP address that all client connections access.
- Real server IPs (RIP): The IP addresses of client network interfaces used for connecting database clients to the database engine.
- Cluster: A cluster of real HP Vertica servers (nodes).
- Virtual server: The single point of entry that provides access to a cluster, based on dynamic node selection.

Note: All IP addresses shown in this document are for illustrative purposes only. Please replace example IP addresses shown with actual IP addresses in your environment.

Complete the tasks in the following sections to set up a master and slave director and a third node for K-Safety.

Before You Begin

You must shut down the HP AI core server to disconnect it from the HP Vertica database before performing the configuration changes described in this document.

Set Up VIP Using the Loopback Interface

Perform the following steps to set up the loopback (lo) interface with an alias on each node. Setting up an alias on each node gives the node a virtual IP address.

1. Log in as root on the master director (node01):

```
$ su - root
```

2. Use a text editor to open the `ifcfg-lo` file:

```
[root@node01]# vi /etc/sysconfig/network-scripts/ifcfg-lo
```

3. Set up the loopback adapter with an alias for the Virtual IP address (VIP) by adding the following lines to the end of the `ifcfg-lo` file:

```
## vip device
DEVICE=lo:0
IPADDR=10.10.51.180
NETMASK=255.255.255.255
ONBOOT=yes
NAME=loopback
```

Note: When you add the above block to your file, be careful not to overwrite the 127.0.0.1 parameter, which is required for proper system operations.

4. Save the `ifcfg-lo` file.
5. Start the device:


```
[root@node01]# ifup lo:0
```
6. Repeat steps 1–5 on each node in the HP Vertica cluster.

Disable Address Resolution Protocol (ARP)

This procedure disables ARP (Address Resolution Protocol) for the VIP.

1. On the master director (node01), log in as root:


```
$ su - root
```
2. Use a text editor to open the `sysctl` configuration file:


```
[root@node01]# vi /etc/sysctl.conf
```
3. Add the following lines to the end of the file:

```
#LVS
net.ipv4.conf.eth1.arp_ignore = 1
net.ipv4.conf.eth1.arp_announce = 2
# Enables packet forwarding
net.ipv4.ip_forward = 1
```

Note: If the real IPs are configured on an interface other than `eth1`, replace the “eth1” in the above command with the appropriate interface.

4. Save the `/etc/sysctl.conf` file.
5. Type the `ifconfig` command to verify that the interface is on the same subnet as the VIP:


```
[root@node01]# /sbin/ifconfig
```
6. Type the `ifconfig` command to verify that the loopback interface is up:

```
[root@node01]# /sbin/ifconfig lo:0
```

You should see output similar to the following:

```
lo:0 Link encap:Local Loopback
inet addr:10.10.50.100 Mask:255.255.255.255
```

```
UP LOOPBACK RUNNING MTU:16436 Metric:1
```

If you do not see `UP LOOPBACK RUNNING`, bring up the loopback interface:

```
[root@node01]# /sbin/ifup lo
```

7. Type the following command to commit your changes to the kernel from the configuration file:

```
[root@node01]# /sbin/sysctl -p
```

8. Repeat steps 1–7 on all nodes in the HP Vertica cluster.

Install the HP Vertica IPVS Load Balancer Package

Before you install the IP Virtual Server (IPVS) Load Balancer package, make sure that the HP Vertica Analytics Database package for your version of Red Hat Enterprise Linux (5.x or 6.x) is downloaded and installed.

Perform the following steps to install the HP Vertica IPVS Load Balancer package for Red Hat Enterprise Linux 5 or Red Hat Enterprise Linux 6.

Note: The examples below use node01 for the master director and node02 for the slave director.

Install the HP Vertica IPVS Load Balancer Package on RHEL 5.x

1. Log in as root on the master director (node01):

```
$ su - root
```

2. Copy the IPVS Load Balancer package for RHEL 5 to the Vertica server from the HP AI core server. Mount the ISO on disk 1 to an existing directory on the HP AI core server, such as `/mnt`.
3. Use the `scp` command to secure copy the IPVS Load Balancer RPM from the HP AI core server to any of the Vertica directories on the master server, such as `/tmp` or `/ai`.

```
# scp root@ai-core-server:/mnt/VerticaIPVSLoadBalancer-6.1-0.RHEL5.x86_64.rpm /tmp
```

4. Change directory to the location of the downloaded file:

```
# cd /tmp
```

or

```
# cd /ai
```

5. Type the `rpm -Uvh` command to install or upgrade the IPVS Load Balancer package:

```
# rpm -Uvh VerticaIPVSLoadBalancer-6.1-0.RHEL5.x86_64.rpm
```

6. Repeat steps 1–4 on the slave director (node02).

Install the HP Vertica IPVS Load Balancer Package on RHEL 6.x

1. Log in as root on the master director (node01):

```
$ su - root
```

Copy the IPVS Load Balancer package for RHEL 6 to the Vertica server from the HP AI core server. Use the `scp` command to secure copy the IPVS Load Balancer RPM from the HP AI core server to any of the Vertica directories on the master server, such as `/tmp` or `/ai`.

```
# scp root@ai-core-
server:/opt/HP/CBI/installcontent/VerticaIPVSLoadBalancer-6.1-
0.RHEL6.x86_64.rpm /tmp
```

2. Change directory to the location of the downloaded file:

```
# cd /tmp
```

or

```
# cd /ai
```

3. Type the following command as root to add the IPVS Load Balancer package to the kernel of the master server:

```
# /sbin/modprobe ip_vs
```

4. Type the following command to verify that ip_vs is loaded:

```
# lsmod | grep ip_vs
```

5. Type the `rpm -Uvh` command to install or upgrade the IPVS Load Balancer package:

```
# rpm -Uvh VerticaIPVSLoadBalancer-6.1-0.RHEL6.x86_64.rpm
```

6. Repeat steps 1–5 on the slave director (node02).

Configure the HP Vertica IPVS Load Balancer

The HP Vertica IPVS Load Balancer package contains the `configure-keepalived.pl` script in the `/sbin` directory. If you run the `configure-keepalived.pl` script with no arguments, it prints a usage summary similar to the following:

```
--ripiips      | Comma separated list of HP Vertica nodes; public IPs (e.g.,
                | 10.10.50.116, etc.)
--priv_ips     | Comma separated list of HP Vertica nodes; private IPs
                | (e.g., 192.168.51.116, etc.)
--riport      | Port on which HP Vertica runs. Default is 5433
--iface       | Public ethernet interface HP Vertica is configured to use
                | (e.g., eth0)
--emailto     | Address that should get alerts (e.g., user@server.com)
--emailfrom   | Address that mail should come from (e.g., user@server.com)
--mailserver  | E-mail server IP or hostname (e.g., mail.server.com)
--master      | If this director is the master (default), specify --master
--slave       | If this director is the slave, specify --slave
--authpass    | Password for keepalived
--vip         | Virtual IP address (e.g., 10.10.51.180)
--delayloop   | Seconds keepalived waits between healthchecks. Default is 2
--algo        | Sets the algorithm to use: rr, wrr, lc (default), wlc,
                | lbic, lblcr, dh, sh, sed, nq
--kind        | Sets the routing method to use. Default is DR.
--priority    | By default, master has priority of 100 and the backup
                | (slave) has priority of 50
```

About Public and Private IPs

If your cluster uses private interfaces for spread cluster communication, you must use the `--priv_ips` switch to enter the private IP addresses that correspond to the public IP addresses (or RIPv6). The IPVS keepalived daemon uses the private IP addresses to determine when a node has left the cluster.

The IP host ID of the RIPs must correspond to the IP host ID of the private interfaces. For example, given the IP address mappings in Table 1, enter the IP addresses in the following order:

```
--ripips 10.10.50.116,10.10.50.117,10.10.50.118
--priv_ips 192.168.51.116,192.168.51.117,192.168.51.118
```

Table 1. Example Public and Private (Spread) IP Address Mapping

| Preconfigured IP | Node Assignment | Public IPs/Ethernet Interface (eth1) | Private IPs/Ethernet Interface (eth0) |
|---------------------|------------------------|--------------------------------------|---------------------------------------|
| VIP | Shared among all nodes | 10.10.50.115 | 192.168.51.115 |
| RIP master director | Cbi-Vertica-11 | 10.10.50.116 | 192.168.51.116 |
| RIP slave director | Cbi-Vertica-12 | 10.10.50.117 | 192.168.51.117 |
| RIP failover node | Cbi-Vertica-13 | 10.10.50.118 | 192.168.51.118 |

You must use IP addresses, not node names, or the `spread.pl` script may fail.

If you do not specify private interfaces, HP Vertica uses the public RIPs for the MISC check, as shown in Step 3 in the “Set Up the HP Vertica IPVS Load Balancer Configuration File on the Master and Slave Directors” section.

Set Up the HP Vertica IPVS Load Balancer Configuration File on the Master and Slave Directors

1. Log in as root to the master director (node01):

```
$ su - root
```

2. Run the HP-supplied `configure-keepalived.pl` script with the appropriate switches. The following is an example for the master director:

```
# /sbin/configure-keepalived.pl -ripips
10.10.50.116,10.10.50.117,10.10.50.118
--priv_ips 192.168.51.116,192.168.51.117,192.168.51.118
--riport 5433
--iface eth1
--emailto dbadmin@companyname.com
--emailfrom dbadmin@companyname.com
--mailserver mail.server.com
--master
--authpass password
--vip 10.10.51.180
--delayloop 2
--algo lc
--kind DR
--priority 100
```

The following is an example for the slave director:

```
# /sbin/configure-keepalived.pl -ripips
10.10.50.116,10.10.50.117,10.10.50.118
--priv_ips 192.168.51.116,192.168.51.117,192.168.51.118
--riport 5433
--iface eth0
--emailto dbadmin@companyname.com
--emailfrom dbadmin@companyname.com
--mailserver mail.server.com
--slave
--authpass password
--vip 10.10.51.180
--delayloop 2
--algo lc
--kind DR
--priority 50
```

Tip: Use a lower priority for the slave director `--priority` switch. HP recommends setting a priority of 50.

Note: Use the `--iface` switch to specify the interface for PIP.

Note: If you do not enter the `--priv_ips` switch, the public IP address of each node is monitored for keepalive checking.

Caution: The `--authpass password` switch must be the same on the master and slave directors.

3. Type the following command to check the `keepalived.conf` file to verify the private and public IP settings for the `--ripips` and `--priv_ips` switches and to verify that the `real_server` IP address is public:

```
# cat /etc/keepalived/keepalived.conf
```

An entry similar to the following should be found in the `keepalived.conf` file:

```
real_server 10.10.50.116 5433 {
    MISC_CHECK {
        misc_path "/etc/keepalived/check.pl 192.168.51.116"
```

```
    }  
}
```

4. Type the following command to start the spread script:

```
# /etc/init.d/spread.pl start
```

The `spread.pl` script rewrites the `check.txt` file to include only the remaining nodes in the event of a node failure. This prevents the virtual server from sending vsql requests to the failed node.

5. Type the following command to start keepalived on node01:

```
# /etc/init.d/keepalived start
```

6. If not already started, start `sendmail` to allow the directors to send mail messages:

```
# /etc/init.d/sendmail start
```

7. Include the following entries in the `/etc/rc.local` file:

```
/sbin/modprobe ip_vs  
rm -f /etc/keepalived/.spread.pl.lock  
/etc/init.d/spread.pl start  
/etc/init.d/spreadd restart  
/etc/init.d/keepalived restart
```

8. Repeat steps 1–7 on the slave director (node02).

Note: Replace the `--master` switch with the `--slave` switch in Step 2 when configuring the slave director.

Tip: Use a lower priority for the slave director `--priority` switch. HP recommends setting a priority of 50.

You must incorporate all command line commands in your startup scripts, for example in `/etc/rc.local`, if Keepalived and its associated services are not running after system startup or reboot.

If the server is restarted for any reason, you must remove the `.spread.pl.lock` file to cleanly restart the spread script. The following entries must be included in the `/etc/rc.local` file:

```
/sbin/modprobe ip_vs  
rm -f /etc/keepalived/.spread.pl.lock  
/etc/init.d/spread.pl start  
/etc/init.d/spreadd restart  
/etc/init.d/keepalived restart
```

For more information on configuring load balancing in your HP Vertica cluster, see the “Load Balancing” section of the *HP Vertica Analytics Platform 6.1.x Administrator’s Guide*.