Hitachi Content Platform Gateway Linux Cluster Setup with SAN Storage

Release Version 4.1.4

Linux Failover Clustering for Virtual and Physical Servers

The objective of this document is to cover the setup of Linux Failover Clustering with two Hitachi Content Platform Gateway nodes running on VMware or physical machines.

MK-HCPG009-00

June 2021

© 2020, 2021 Hitachi Vantara LLC. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including copying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., or Hitachi Vantara LLC (collectively "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara LLC at https://support.hitachivantara.com/en_us/contact-us.html.

Notice: Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara LLC.

By using this software, you agree that you are responsible for:

1. Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals; and
2. Verifying that your data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

Notice on Export Controls. The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

Hitachi and Lumada are trademarks or registered trademarks of Hitachi, Ltd., in the United States and other countries.

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, Lotus, MVS, OS/390, PowerHA, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, the Microsoft Corporate Logo, MS-DOS,

Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or website are properties of their respective owners.

Copyright and license information for third-party and open source software used in Hitachi Vantara products can be found at https:// www.hitachivantara.com/en-us/company/legal.html.

Table of Contents

Chapter 1 Overview

This is an overview of the Linux HA cluster configuration. The two gateways are in a cluster and have a shared cache, only one gateway is active. This configuration does not support both sites being active at the same time. In this configuration, the databases on the gateways are replicated to each other. The Cluster network is used for management of the nodes and user access to the Cluster shares. The Replication network is a private network used for the database replication. The Storage network is used for accessing the shared storage. The IP addresses in this diagram are for illustrative purposes only. Use IP addresses provided by the customer.

Figure 1.1 HA Configuration

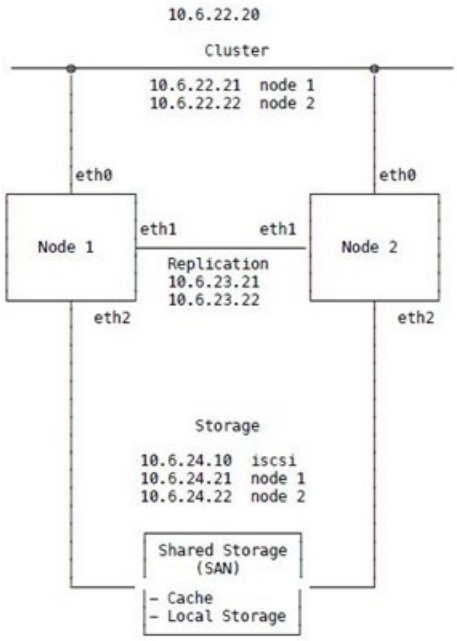

WARNING: The first time the cluster fails over from the Primary node (Node 1) to the Secondary node (Node 2) login to the HCP Gateway UI, navigate to the Shares page and select the Start button on all the Shares. After that, the shares will start automatically on the node that is the active node in the cluster.

Chapter 2 Node Configuration

Note: The configuration script must be run from the ESXi or physical console. Do not run the script from an ssh or putty session.

Node 1

1. Login to the Node 1 as user **vault** with password **0rgan1c**.
2. Add the IP address and hostname information to the configure.txt file /home/vault/configure.txt. Refer to the **Example Configure.txt** chapter for the details.
3. Configure the network by issuing the command **sudo configure 1**
4. Change the cluster password by issuing the **sudo passwd hacluster**
5. If prompted for the **vault** user password, enter **0rgan1c**.
6. For the **hacluster** user, enter the default password **0rgan1c@Cluster**
7. Reboot the server by issuing the command **sudo reboot**

Node 2

1. Login to the Node 2 as user **vault** with password **0rgan1c**
2. Copy the configure.txt file /home/vault/configure.txt from Node 1 to Node 2.
3. Configure the network by issuing the command **sudo configure 2**
4. Change the cluster password by issuing the **sudo passwd hacluster**
5. If prompted for the **vault** user password, enter **0rgan1c**
6. For the **hacluster** user, enter the default password **0rgan1c@Cluster**
7. Reboot the server by issuing the command **sudo reboot**

Chapter 3 SSH

Note: These steps can be run from the console or an ssh or putty session.

Node 1

1. Login to the Node 1 as user **vault** with password **0rgan1c**
2. Change to the root user by issuing the command **sudo -i**
3. Generate an SSH key by issuing the command **ssh-keygen**
4. Accept defaults for all of the prompts.

   Generating public/private rsa key pair.

   Enter file in which to save the key (/root/.ssh/id_rsa): **<enter>**

   Enter passphrase (empty for no passphrase): **<enter>**

   Enter same passphrase again: **<enter>**

   Your identification has been saved in /root/.ssh/id_rsa

   Your public key has been saved in /root/.ssh/id_rsa.pub

   The key fingerprint is:

   SHA256:6cBHpax8rdJuRj55ZRo9sWuBa061pAp68Ex2/3wWFSw root@hcpg-1

   The key's random art image is:

```
+---[RSA 3072]----+
|        .  .  |
|     . o   E o |
|      +    . .|
|    o o o  .  .|
|     = S .ooo . |
|   . o*o.o+B..  |
|    *o++oo*.+ . |
|    .++*o*.o o  |
|   .. oo=.ooo   |
+----[SHA256]-----+
```

5. Copy the ssh key to Node 2, the default name of Node 2 is hcpg-2, by issuing the command **ssh-copy-id root@hcpg-2** Answer **yes** when prompted to continue connecting. If you used a different name for Node 2, replace it in the command above.

   /usr/bin/ssh-copy-id: INFO: Source of key(s) to be installed: "/root/.ssh/id_rsa.pub"

   The authenticity of host 'hcpg-2 (10.6.22.22)' can't be established.

   ECDSA key fingerprint is SHA256:DJ/EfAf4hOaeooyrSUUgxoX80x6AdHwUtOzWU2Lu3x4.

   Are you sure you want to continue connecting (yes/no/[fingerprint])? **yes**

   /usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter out any that are already installed

   /usr/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompted now it is to install the new keys

   root@hcpg-2's password:

   Number of key(s) added: 1

   Now try logging into the machine, with: "ssh 'root@hcpg-2'"

   and check to make sure that only the key(s) you wanted were added

6. Verify that root can access Node 2 using the hostname by issuing the command **ssh root@hcpg-2**
7. Logout out of Node 2 by issuing the command **exit**
8. Verify that root can access Node 2 using the eth0 IP Address by issuing the command **ssh root@10.6.22.22** Note that this IP address is here for illustrative purposes only, use the IP addresses supplied by the customer.
9. Logout out of Node 2 by issuing the command **exit**

Node 2

1. Login to the Node 2 as user **vault** with password **0rgan1c**
2. Change to the root user by issuing the command **sudo -i**
3. Generate an SSH key by issuing the command **ssh-keygen**
4. Accept defaults for all of the prompts.

   Generating public/private rsa key pair.

Enter file in which to save the key (/root/.ssh/id_rsa): **<enter>**

Enter passphrase (empty for no passphrase): **<enter>**

Enter same passphrase again: **<enter>**

Your identification has been saved in /root/.ssh/id_rsa

Your public key has been saved in /root/.ssh/id_rsa.pub

The key fingerprint is:

SHA256:6cBHpax8rdJuRj55ZRo9sWuBa061pAp68Ex2/3wWFSw root@hcpg-2

The key's random art image is:

```
+---[RSA 3072]----+
|        .  .  |
|     . o   E o |
|      +    . .|
|    o o o  .  .|
|     = S .ooo . |
|   . o*o.o+B..  |
|    *o++oo*.+ . |
|    .++*o*.o o  |
|   .. oo=.ooo   |
+----[SHA256]-----+
```

5. Copy the ssh key to Node 1, default name of Node 1 is hcpg-1, by issuing the command **ssh-copy-id root@hcpg-1** Answer **yes** when prompted to continue connecting.

/usr/bin/ssh-copy-id: INFO: Source of key(s) to be installed: "/root/.ssh/id_rsa.pub"

The authenticity of host 'hcpg-1 (10.6.22.21)' can't be established.

ECDSA key fingerprint is SHA256:DJ/EfAf4hOaeooyrSUUgxoX80x6AdHwUtOzWU2Lu3x4.

Are you sure you want to continue connecting (yes/no/[fingerprint])? **yes**

/usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter out any that are already installed.

/usr/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompted now it is to install the new keys

root@hcpg-1's password:

Number of key(s) added: 1

Now try logging into the machine, with: "ssh 'root@hcpg-1'"

and check to make sure that only the key(s) you wanted were added.

6. Verify that root can access Node 1 using the hostname by issuing the command **ssh root@hcpg-1**
7. Logout out of Node 1 by issuing the command **exit**
8. Verify that root can access Node 1 using the eth0 IP Address by issuing the command **ssh root@10.6.22.21** Note that this IP address is here for illustrative purposes only, use the IP addresses supplied by the customer.
9. Logout out of Node 1 by issuing the command **exit**

Chapter 4 Storage

Note: These steps are run on both Node 1 and Node 2. Change the IP address to match your iSCSI server. The file path in Step 2 will match your environment and will not match the example provided below. These steps can be run from the console or an ssh or putty session. Identify the IP address of the customer's ISCSI server.

1. Discover the iSCSI server by issuing the command **sudo iscsiadm -m discovery -t st -p 10.6.22.10** Note that you need to use the customer's ISCSI server IP address.
2. Edit the ISCSI configuration file by issuing the command **sudo vi /etc/iscsi/nodes/iqn.iscsi.server\:lun1/10.6.22.10\,3260\,1/default** Replace the 10.6.22.10 with the IP address of the customer's ISCSI server.
3. Add the authentication to the ISCSI configuration file
4. Near the top of the file

   Change:

   node.startup = manual

   To:

   node.startup = automatic

   Look for this line:

   node.session.initial_cmdsn <<<< add the text below before this line

   Note: use the correct user names and passwords for the customer's ISCSI server

   node.session.auth.authmethod = CHAP

   node.session.auth.username = iscsi-user

   node.session.auth.password = password

   node.session.auth.username_in = iscsi-target

   node.session.auth.password_in = secretpass

5. Restart the ISCSI services by issuing the command **sudo systemctl restart open-iscsi iscsid**
6. Check that the disk is in the block list

a. If the Linux Gateway was deployed from an OVA, check that the **sdc** disk is in the block list by issuing the command **lsblk**

```
NAME   MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda      8:0    0  100G  0 disk
├─sda1   8:1    0   96G  0 part /
├─sda2   8:2    0  3.5G  0 part [SWAP]
└─sda3   8:3    0  511M  0 part /boot/efi
sdb      8:16   0  100G  0 disk
└─sdb1   8:17   0  100G  0 part /var/lib/mysql
sdc      8:32   0    1T  0 disk
└─sdc1   8:33   0 1024G  0 part
sr0     11:0    1 1024M  0 rom
```

b. If the Linux Gateway was deployed from the Appliance ISO, check that the **sdb** disk is in the block list by issuing the command **lsblk**

```
NAME   MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda      8:0    0  100G  0 disk
├─sda1   8:1    0   110G  0 part /
├─sda2   8:2    0  3.5G  0 part [SWAP]
├─sda3   8:3    0  511M  0 part /boot/efi
└─sda4   8:4    0  10G  0 part /var/lib/mysql
sdb      8:16   0    1T  0 disk
└─sdb1   8:17   0 1024G  0 part
sr0     11:0    1 1024M  0 rom
```

7. Check the status of the ISCSI configuration by issuing the command **sudo iscsiadm -m session -o show**

tcp: [1] 10.6.22.10:3260,1 iqn.iscsi.server:lun1 (non-flash)

8. Reboot the node by issuing the command **sudo reboot**

Note: Steps 1 - 8 in this chapter need to be run on both nodes.

9. Login to the Node 1 as user **vault** with password **0rgan1c**
10. Create the /storage/reports folder by issuing the command **sudo mkdir /storage/reports**
11. Set the owner and permissions on the /storage/reports folder by issuing the commands **sudo chown wildfly:wildfly /storage/reports** and **sudo chmod 777 /storage/reports**

Chapter 5 Cluster

Note: These steps are run on Node 1. The script will ask for the cluster username and password from the **Configure System** chapter. When prompted enter the username **hacluster** and the password **0rgan1c@Cluster**

Note: If there is an error, the script will halt. After the error has been addressed, the script can be run again. The script will remove the existing and any partial cluster configuration before restarting the cluster build process.

Node 1

1. Login to the Node 1 as user **vault** with password **0rgan1c**
2. Change to the root user by issuing the command **sudo -i**
3. Configure the cluster by issuing the command **cluster** When prompted enter the username **hacluster** with the default password **0rgan1c@Cluster**

dos2unix: converting file /home/vault/configure.txt to Unix format...

Synchronizing state of corosync.service with SysV service script with /lib/systemd/systemd-sysv-install.

Executing: /lib/systemd/systemd-sysv-install enable corosync

Created symlink /etc/systemd/system/multi-user.target.wants/corosync.service → /lib/systemd/system/corosync.service.

Synchronizing state of pacemaker.service with SysV service script with /lib/systemd/systemd-sysv-install.

Executing: /lib/systemd/systemd-sysv-install enable pacemaker

Created symlink /etc/systemd/system/multi-user.target.wants/pacemaker.service → /lib/systemd/system/pacemaker.service.

Enter the Cluster username and password

Username: **hacluster**

Password: **0rgan1c@Cluster**

hcpg-2: Authorized

hcpg-1: Authorized

Note: These warnings about addresses can be ignored, they are not errors.

No addresses specified for host 'hcpg-1', using 'hcpg-1'

No addresses specified for host 'hcpg-2', using 'hcpg-2'

Note: The next warnings about file found can be ignored

Warning: hcpg-1: Cluster configuration files found, the host seems to be in a cluster already

Warning: hcpg-2: Cluster configuration files found, the host seems to be in a cluster already

Note: The 'Destroying cluster' is not an error and can be ignored

Destroying cluster on hosts: 'hcpg-1', 'hcpg-2'...

hcpg-1: Successfully destroyed cluster

hcpg-2: Successfully destroyed cluster

Requesting remove 'pcsd settings' from 'hcpg-1', 'hcpg-2'

hcpg-2: successful removal of the file 'pcsd settings'

hcpg-1: successful removal of the file 'pcsd settings'

Sending 'corosync authkey', 'pacemaker authkey' to 'hcpg-1', 'hcpg-2'

hcpg-1: successful distribution of the file 'corosync authkey'

hcpg-1: successful distribution of the file 'pacemaker authkey'

hcpg-2: successful distribution of the file 'corosync authkey'

hcpg-2: successful distribution of the file 'pacemaker authkey'

Synchronizing pcsd SSL certificates on nodes 'hcpg-1', 'hcpg-2'...

hcpg-2: Success

hcpg-1: Success

Sending 'corosync.conf' to 'hcpg-1', 'hcpg-2'

hcpg-1: successful distribution of the file 'corosync.conf'

hcpg-2: successful distribution of the file 'corosync.conf'

Cluster has been successfully set up.

hcpg-1: Starting Cluster...

hcpg-2: Starting Cluster...

Note: The Warning about 'defaults' can be ignored

Warning: Defaults do not apply to resources which override them with their own defined values

3. Check the status of the cluster by issuing the command **pcs status**

Note: The cluster will initially take a short period of time to start, you may need to run this command more than once until all the resources are started and the cluster is operational.

Cluster name: hcpg

Stack: corosync

Current DC: hcpg-1 (version 2.0.1-9e909a5bdd) - partition with quorum

Last updated: Wed Mar 31 23:37:45 2021

Last change: Wed Mar 31 23:36:47 2021 by hacluster via crmd on hcpg-1

2 nodes configured

5 resources configured

Online: [hcpg-1 hcpg-2]

Full list of resources:

Resource Group: nfs-group

clusterip (ocf::heartbeat:IPaddr2): Started hcpg-1

nfsserver(lsb:nfs-kernel-server): Started hcpg-1

sam_storage (ocf::heartbeat:Filesystem): Started hcpg-1

sam_license (systemd:saml): Started hcpg-1

sam_service (systemd:sam.target): Started hcpg-1

Daemon Status:

corosync: active/disabled

pacemaker: active/disabled

pcsd: active/enabled

Chapter 6 Example Configuration

Refer to the Overview chapter for details of the IP addresses needed for the cluster. The IP addresses in this chapter are for illustrative purposes only. Use IP addresses provided by the customer.

```
# Cluster Name

# default name: hcpg

cluster.name=hcpg


# Cluster IP Information

cluster.virtual.ip.address=10.6.22.20


# DNS Nameservers

# example: nameservers=10.10.1.2, 10.10.2.2

nameservers=9.9.9.9


# Node 1 hostname

node1.hostname=hcpg-1


# Node 1 Cluster Address

node1.eth0.ip.address=10.6.22.21

node1.eth0.netmask=255.255.0.0

node1.eth0.gateway=10.6.0.1

node1.eth0.name=hcpg-1.dts-labs.com


# Node 1 Private Network
```

node1.eth1.ip.address=10.6.23.21

node1.eth1.netmask=255.255.0.0

Node 1 Storage Network

node1.eth2.ip.address=10.6.24.21

node1.eth2.netmask=255.255.0.0

node1.eth2.gateway=

node1.eth2.name=

Node 2 hostname

node2.hostname=hcpg-2

Node 2 Cluster Address

node2.eth0.ip.address=10.6.22.22

node2.eth0.netmask=255.255.0.0

node2.eth0.gateway=10.6.0.1

node2.eth0.name=hcpg-2.dts-labs.com

Node 2 Private Network

node2.eth1.ip.address=10.6.23.22

node2.eth1.netmask=255.255.0.0

Node 2 Storage Network

node2.eth2.ip.address=10.6.24.22

node2.eth2.netmask=255.255.0.0

node2.eth2.gateway=

node2.eth2.name=

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA
HitachiVantara.com | community.HitachiVantara.com

Contact Information
USA: 1-800-446-0744
Global: 1-858-547-4526
HitachiVantara.com/contact