



Hewlett Packard
Enterprise

HPE Helion CloudSystem 10.0 Update 2 ESXi Installation Guide

Abstract

The information in this guide is for administrators using HPE Helion CloudSystem Software version 10.0 Update 2, who are assigned to configure and provision compute resources for deployment and use in virtual data centers.

Part Number: 5200-2466a
Published: March 2017
Edition: 2

Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise 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. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise 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.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Acknowledgments

Microsoft® and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Red Hat™ is a registered trademark of Red Hat, Inc. in the United States and other countries.

VMware™ is a registered trademark of VMware, Inc. in the United States and/or other jurisdictions.

Contents

HPE Helion CloudSystem Installation Overview.....	7
Audience.....	7
Before you begin.....	7
ESXi Environment Installation.....	8
ESXi Management hypervisor supported configurations.....	9
Process overview.....	9
Installation prerequisites.....	9
Installing the operating system.....	10
Configure networks accessible to the server.....	10
Set network interface settings on the ESXi Management hypervisor.....	11
Configure the management cluster and add the ESXi hypervisors.....	11
Download and prepare installation images and tools.....	13
Prepare the deployment platform.....	14
Create the Deployer virtual appliance.....	14
Install vCenter OVF tools package.....	17
Changing virtual appliance disk size.....	18
Preparing Helion CloudSystem Enterprise.....	19
Updating the MySQL JDBC driver name in the configuration file.....	19
Deploy the cloud from the installation UI.....	21
Installation UI considerations.....	21
Fields in the installation UI.....	21
Initiate the deployment action from the installation UI.....	30
Access configuration files and log files.....	30
Shut down the Helion CloudSystem installation UI after deployment.....	31
Manually distribute virtual machines to support HA failover.....	31
Deploy the cloud from the CLI.....	33
Helion CloudSystem installation configuration file.....	33
Configure the installation configuration file.....	34
Fields in the installation configuration file.....	35
Validate the installation configuration file.....	44
Deploy Helion CloudSystem virtual appliances.....	44
Modify Affinity rules.....	46
Troubleshoot installation errors.....	48
Using <code>csdeploy</code> log files to find information on deployment errors.....	48
Viewing all failed deployment jobs and tasks.....	48
Reset the environment after a failed deployment.....	48

Configure the cloud after deployment.....	49
Access Helion CloudSystem virtual appliances.....	50
Access Helion CloudSystem virtual appliances.....	51
Access appliance virtual machine portals.....	51
Access the Operations Console on the Management appliance.....	51
Access the Logging Dashboard (Kibana) on the Management appliance.....	51
Access the HA Proxy on the Management appliance.....	51
Access the OpenStack user portal on the Cloud controller appliance.....	52
Access HPE CSA, Marketplace portal, and HPE OO on the Enterprise appliance.....	52
Identify Helion CloudSystem virtual appliances.....	54
Deployer virtual appliance name and IP address allocation.....	54
Management appliance names and IP address allocation.....	54
Monitoring appliance names and IP address allocation.....	55
Database appliance names and IP address allocation.....	55
Foundation appliance names and IP address allocation.....	56
Enterprise appliance names and IP address allocation.....	57
VIP address assignment.....	57
Update Helion CloudSystem.....	59
Update Helion CloudSystem to 10.0 Update 2.....	60
(Optional) Update HPE CSA to 4.60 patch 2.....	63
Support and other resources.....	65
Websites.....	66
Software technical support and update service.....	67
Registration for Software technical support and update service.....	67
Using your Software technical support and update service.....	67
Consulting and installation services.....	67
HPE Datacenter Care for Cloud.....	67
Management security.....	67
Join the discussion.....	67
Related documents.....	68
Support and other resources.....	70
Accessing Hewlett Packard Enterprise Support.....	70
Accessing updates.....	70
Customer self repair.....	70

Remote support.....	71
Warranty information.....	71
Regulatory information.....	71
Documentation feedback.....	72

Appendix..... 73

Migrating standard vSwitches to distributed vSwitches.....	73
Install Octavia load balancer.....	73
Installing the Octavia load balancer.....	74
Add Enterprise to a Foundation-only deployment.....	76
Adding Enterprise to a Foundation-only deployment.....	76
Update service endpoints with the FQDN on Helion CloudSystem Enterprise appliances.....	77
Disk requirements.....	78
Disk requirements (Enterprise-only).....	82

Revision history

Edition	Publication date	Helion CloudSystem version	Edition notes
1	February 2017	10.0 Update 2	Initial edition
2	March 2017	10.0 Update 2	Corrected procedure (Optional) Update HPE CSA to 4.60 patch 2

HPE Helion CloudSystem Installation Overview

Audience

This guide is intended for experienced system administrators with a working knowledge of the following concepts.

- Top of Rack (ToR) switches for networking
- CLI commands for Windows and Linux
- VMware vCenter functionality
- Secure shell (ssh) and secure copy (scp) commands

If you plan to use the OpenStack CLI and APIs to manage some of the cloud resources from the command line, it is helpful to have experience with OpenStack technologies such as Nova, Glance, Cinder, Keystone, and Neutron.

Before you begin

HPE Helion CloudSystem is a robust product that offers a flexible cloud installation. Make sure that you understand the known issues associated with the build and the resource requirements as well as the supported hardware, software, and tools before you start the installation.

Refer to the **[Enterprise Information Library](#)** to find:

- *HPE Helion CloudSystem 10.0 Update 2 Release Notes*
- *HPE Helion CloudSystem 10.0 Update 1 Support Matrix*

Customizing the network infrastructure

This installation guide is written according to Helion CloudSystem best practices for network configuration. To customize the network configuration, review the details of each individual network in the *HPE Helion CloudSystem 10.0 Network Planning Guide* to learn more about network flexibility in the overall network infrastructure. During installation, apply your customized network choices to the installation configuration file.

See the *HPE Helion CloudSystem 10.0 Network Planning Guide* in the **[Enterprise Information Library](#)**.

Using the command text file

The file `cs10_02_esx_commands.txt` is attached to the PDF. The file includes all commands used in the procedures in this guide. You can copy and paste the commands from the text file directly to your command line.

To reveal the text file, click the paper clip icon on the left side of the PDF.

ESXi Environment Installation

ESXi Management hypervisor supported configurations

Helion CloudSystem supports three HA configurations:

- **Physical HA:**

Management cluster with three hosts (recommended)

Helion CloudSystem virtual appliances are deployed in trios and distributed across three physical servers. This configuration offers high availability and performance benefits for robust production use.

- **Virtual HA:**

Management cluster with a single host

Helion CloudSystem virtual appliances are deployed in trios on a single physical server. This configuration is supported for small development, test, or production environments where physical high availability is not a requirement. Contact HPE Support for assistance if you want to expand to a three host HA configuration after deployment is complete.

- **Non-HA:**

Management cluster with a single host

Seven singleton virtual appliances are deployed on a single physical server. You cannot expand to an HA configuration at a later time.

For virtual appliance resource requirements, see the *HPE Helion CloudSystem 10.0 Update 1 Support Matrix* in the [Enterprise Information Library](#).

Process overview

Procedure

1. [Installation prerequisites](#) on page 9
2. [Installing the operating system](#) on page 10
3. [Configure networks accessible to the server](#) on page 10
4. [Set network interface settings on the ESXi Management hypervisor](#) on page 11
5. [Configure the management cluster and add the ESXi hypervisors](#)
6. [Download and prepare installation images and tools](#) on page 13

Installation prerequisites

Before you begin the installation, make sure that the following prerequisite steps are complete for the environment.

Prerequisites

- An ESXi ISO is downloaded from the VMware website.

Download can be obtained from <https://my.vmware.com/en/web/vmware/downloads>.

- If you are using HPE hardware and need a customized image, the physical server is configured with a supported ESXi operating system.

The HPE Customized image for the supported version of ESXi is available on the [HPE ProLiant servers with VMware](#) website.

See the *HPE Helion CloudSystem 10.0 Update 1 Support Matrix* in the [Enterprise Information Library](#) for a list of supported versions of the ESXi operating system.

- You understand the networking and storage configuration of the server.
- A large disk is presented to the server that meets the storage requirements outlined in the **Virtual appliance resource requirements** section of the *HPE Helion CloudSystem 10.0 Update 1 Support Matrix* in the [Enterprise Information Library](#).

Set the device path policy to **round-robin 0** for each host if you have multi-path configured in the management host. See [Modifying path information for ESXi/ESX hosts](#).

Installing the operating system

Install a supported version of the ESXi operating system on each Management hypervisor in the management cluster.

For a list of supported operating systems, see the *HPE Helion CloudSystem 10.0 Update 1 Support Matrix* in the [Enterprise Information Library](#).

Procedure

1. Using the iLO IP address for the ESXi Management hypervisor blade, log in to the remote console.
2. Mount the ISO image.
3. Power on the blade.
4. Configure the settings:
 - a. Time zone
 - b. Language
 - c. Hostname
 - d. Select local installation drive

NOTE:

Follow the operating system recommendations for boot storage.

Configure networks accessible to the server

In the network infrastructure that leads up to the physical server, configure three trunks: Management trunk, Data trunk, and Storage trunk.

Management trunk

The Control plane networks make up the Management trunk. Control plane networks include:

- CONF Network

This untagged network must be able to communicate with vCenter.

If defining teamed NICs (bonding) on the CONF network in the installation configuration file or the installation UI, make sure to manually add all interfaces in the host vSwitch in vCenter before deploying Helion CloudSystem.

- Data Center Management Network
- Cloud Management Network
- Consumer Access Network

This network provides cloud users with public access to appliance consoles. If you do not include this network in the control plane, the Data Center Management Network will provide both public and administrative access.

- External Network

This network is optional. Add the External Network if Helion CloudSystem Foundation is included in the cloud environment and floating IP addresses are required.

Data trunk

The Tunnel Underlay Network or the VLAN Data Trunk support data requirements in the cloud. You can have one Tunnel Underlay Network in a cloud environment, and multiple VLAN Data Trunks. Best practice is to put each data network in a separate trunk.

Storage trunk

The Storage trunk includes all storage networks required to support the compute storage requirements of the cloud environment. You can add all storage networks to one Storage trunk.

Set network interface settings on the ESXi Management hypervisor

Configure the Management Network that vCenter uses to manage the ESXi Management hypervisor.

Procedure

1. Using the iLO address for the ESXi Management hypervisor blade, log in to the iLO remote console.
2. Press **F2**.
3. Select **Configure Management Network**.
4. In the VLAN section choose one of the following options:
 - a. If you want the host to have a separate CONF IP address, leave the VLAN value empty to indicate an untagged network.
 - b. If you are using any other network, set the VLAN ID for that network.

IMPORTANT:

The network where vCenter is configured must be able to communicate with the network configured with the ESXi hosts.

5. Configure an IP address for the Management Network.

If you combined the CONF and Management Network, then the IP address must be outside the range reserved on the CONF Network.
6. Save and restart the network.
7. To verify that the network is configured properly, ping the IP address.
8. Repeat this procedure for each Management hypervisor that you plan to include in the management cluster.

Configure the management cluster and add the ESXi hypervisors

Prerequisites

- You have administrator privileges to log in to vCenter.
- One large datastore is created.

If a trio of management hosts is configured in the cluster, all hosts share the volume.

Storage requirements depend on the type of installation you are planning. Refer to the **Management hypervisor hardware requirements** tables in the *HPE Helion CloudSystem 10.0 Update 1 Support Matrix* in the **Enterprise Information Library**.

- Your vCenter name does not contain spaces.

- DRS is enabled in vCenter to manage cluster resources for multiple hosts, and is configured according to VMware best practices.
- All hypervisors in the cluster are installed and configured consistently so that the NICs are set to the same speed, network configuration and naming conventions.

Procedure

1. Log in to vCenter with the administrator user name and password.

NOTE:

Do not use unsupported vCenter password characters. You can find a list of unsupported characters on the [VMware Knowledge Base](#).

Use a local administrator account (not a domain administrator account).

2. In the data center that supports Helion CloudSystem, create the management cluster.

NOTE:

The character “/” is not supported in the cluster or data center name.

3. Add hosts to the cluster.

- a. Right-click on the cluster and select **Add Host**.
- b. Enter the IP address or FQDN of the ESXi hypervisor that you configured in the previous section.
- c. Enter the credentials.
- d. Take all default options.
- e. Repeat until all hosts are added.

4. Verify that you have the following minimum networking pieces configured for each host that you added to the cluster:

- The CONF network is in the same trunk as the vCenter Management Network and the CONF network can communicate with vCenter.
- If using the default **vNetwork Standard Switch (vSS)**, you have:
 - a VM Network, which is untagged.
 - vMotion
 - a vCenter Management Network

See [Set network interface settings on the ESXi Management hypervisor](#).

- If using a **vNetwork Distributed Switch (vDS)**, you have:
 - migrated the default vSS to vDS

See [Migrating standard vSwitches to distributed vSwitches](#).

 - established the relationship between the CONF network and the vCenter Management Network in one of two ways:

Option one — CONF network is separate from the vCenter Management Network

- Create the vDS and assign the default dvPortGroup to the vCenter Management Network (tagged).
- Create a second dvPortGroup for the CONF network (untagged).

Option two — CONF network and the vCenter Management Network are the same

- Use the default dvPortGroup (untagged) for the CONF network.

5. Configure each host in the cluster to use an external time server.

See [Configuring Network Time Protocol \(NTP\) on ESX/ESXi hosts using the vSphere Client](#).

! **IMPORTANT:**

All hosts must be pointing to the same NTP server that is used during the installation of Helion CloudSystem.

6. Verify that the cluster is well configured.
 - a. vMotion is configured on all hosts in the cluster.
 - b. DRS is enabled and set to automatic in the cluster.
 - c. VMware EVC is configured correctly based on the type of host in the cluster.

Download and prepare installation images and tools

The Helion CloudSystem release kit is downloaded from Software Depot at <http://www.hpe.com/support/softwaredepot>. You will need the following release files to perform the procedures in this guide.

- **HPE Helion CloudSystem Enterprise ESX 10.0.2, Feb 2017:** Contains the OVA image for the Enterprise virtual appliance.
- **HPE Helion CloudSystem Foundation ESX 10.0.2, Feb 2017:** Contains the OVA images for the Helion CloudSystem Foundation virtual appliances.
- **HPE Helion CloudSystem Tools 10.0.2, Feb 2017:** Contains the following tools components:
 - `csstart`: Command line to install the Helion CloudSystem Deployer virtual appliance
 - `octavia-amphora-haproxy-guest-image.tgz`: Octavia image used when creating Load Balancers
 - `oo-content-10.50.zip`: HPE Operations Orchestration (OO) content jar files
 - `installer-10.50.69-win64.exe`: Installs OO Central and Studio 10.50
 - `upgrade-10.50-to-10.51.zip`: Upgrades OO Central, Studio, and Remote Action Server (RAS) from 10.50 to 10.51

NOTE:

You must run the same version of OO Central and OO Studio. OO Central 10.51 is installed by default. To use OO Studio, first install OO Studio 10.50, then upgrade to 10.51.

For information about installing OO Studio, see the "Prepare HPE Operations Orchestration for Helion CloudSystem Enterprise" chapter in the *HPE Helion CloudSystem 10.0 Update 1 Administrator Guide* at [Enterprise Information Library](#).

To upgrade OO Studio:

- Unzip the upgrade file and move the `upgrade` folder and its contents to the location where Studio is installed (typically `C:\Program Files\Hewlett-Packard\HP Operations Orchestration`).
- From `C:\Program Files\Hewlett-Packard\HP Operations Orchestration\upgrade\10.51\bin`, run `apply-upgrade.bat`. Ensure that Studio is closed, then apply the upgrade.

-
- **MySQL Connector/J-JDBC driver, Feb 2017:** Contains an updated driver for the Monitoring and Enterprise appliances. This driver must be installed on the first Management appliance before deploying Helion CloudSystem virtual appliances.

NOTE:

Use the following links to verify the integrity of the files after a download.

- [HPE GPG or RPM Signature Verification](#)
 - [Microsoft MD5 verification](#)
-

Prepare the deployment platform

The Deployer virtual appliance is the platform used to install and configure Helion CloudSystem.

Create the Deployer virtual appliance

Prerequisites

- You have administrator privileges to log in to vCenter.
- vCenter has a security certificate.

See the [VMware Knowledge Base](#) for instructions on implementing certificates with vSphere.

Procedure

1. Log in to vCenter using the VMware vSphere Client or a web browser.
2. Select **Inventory > Hosts and Clusters** and select the management cluster.
3. Select **File > Deploy OVF template**.
4. Follow the screen prompts.
 - a. Point to the Deployer virtual appliance image, **HPE-Helion-CloudSystem-foundation-10.0.2.ova**, which was unzipped on your staging environment. Give the virtual machine a unique name that is not currently used in vCenter.
 - b. Best practice is to select thin provisioning. Installation is much faster with this option.
5. Verify that the virtual machine was created. The machine is listed on the left side of the screen under the **VMs and Templates** view.

Do NOT boot the virtual machine.
6. Add Network adapters.
 - a. Right-click the Deployer virtual machine and select **Edit Settings**.
 - b. Select **Add new device type**, then select **Ethernet adapter**.
 - c. Select the default adapter type **VMXNET3**.
 - d. In the Network Connection field, perform one of the following options:
 - I. Select the **VM Network** for VSS.
 - II. Select the **dvPortGroup** for vDS. This port is a dvPortGroup that has the untagged CONF network.
 - e. Make sure that the box **Connect at power-on** is checked.
 - f. Click **Finish**, then click **OK**.
7. Power on the Deployer virtual machine.

View: vSphere Standard Switch vSphere Distributed Switch

Networking

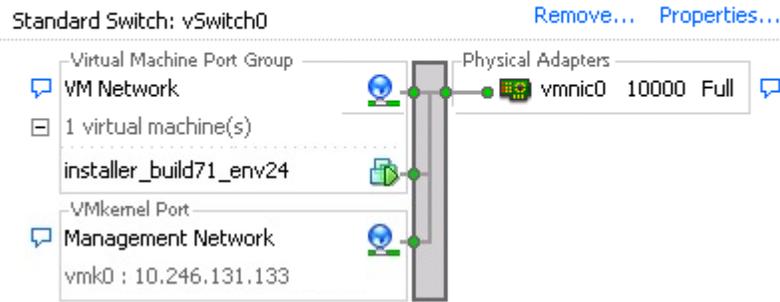


Figure 1: Sample vCenter network VSS configuration before installing Helion CloudSystem

View: vSphere Standard Switch vSphere Distributed Switch

Networking

▼ Distributed Switch: MGMT-DVS-07 Manage Virtual Adapters... Manage Physical Adapters...

MGMT-DVS-07 ⓘ

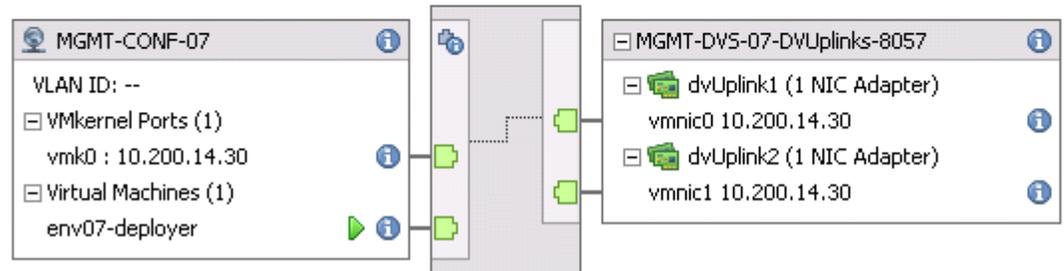


Figure 2: Sample vCenter network VDS configuration before installing Helion CloudSystem

8. Open the Deployer virtual machine console and log on using stack credentials.

For the stack account, use stack/stack.

9. Add the CONF network in the Deployer.

- a. Open the `cs-setup-esx.yml` file for editing:


```
sudo vi /etc/csdeploy/cs-setup-esx.yml
```
- b. Add the environment details to the configuration file.

Configuration field	Explanation
Hypervisor Details	
host-prefix	Enter a prefix with a maximum of eight characters.

Table Continued

Configuration field	Explanation
<pre>dns-settings: nameservers: - ip-address-1 - ip-address-2</pre>	<p>(Optional) Add a DNS server that will resolve host names. To add a server, replace “- ip-address-1” with the IP address of the first DNS server and repeat that process if you need a second DNS server.</p> <p>If you are using servers with a host-name or FQDN, a load balancer with an FQDN, or vCenter with an FQDN, you must define a DNS IP address.</p>
type:	Enter esx to indicate that you are using an ESXi Management hypervisor.
vcenter-hostname:	Enter the IP address or hostname of the vCenter appliance hosting the ESXi Management hypervisor
vcenter-port	(Optional) Enter a port number.
vcenter-verify-ssl	<p>Enter the appropriate option for your environment:</p> <p>True: You have a certificate installed in vCenter for secure connections.</p> <p>False: No certificate is installed in vCenter.</p>
vcenter-username:	<p>Enter the user name of the vCenter appliance.</p> <p>If a special character is used in the user name, the value must be between single quotes.</p>
vcenter-password:	<p>Enter the password for the vCenter appliance.</p> <p>If a special character is used in the password, the value must be between single quotes.</p>
vcenter-cluster-path:	<p>Enter the path in the format:</p> <pre><datacenter_name>/host/<cluster_name></pre>
vcenter-datastore:	<p>Enter the name of the datastore.</p> <p>Example: SanDisk:env23-mgmt</p> <p>If a special character is used in the datastore name, the value must be between single quotes.</p>

Table Continued

Configuration field	Explanation
Control Plane Networking Details	
<pre>--name: CONF types: - hlm members: - cidr: addresses: - gateway-ip:</pre>	<p>This section defines the required Cloud Configuration Network that functions as both a Helion Lifecycle Management Network (hlm) and a PXE network.</p> <ul style="list-style-type: none"> Enter a CIDR value. <p>Example:</p> <pre>10.254.0.0/16</pre> <ul style="list-style-type: none"> The address range is required for CONF. Enter a range of IP addresses that will be assigned to the virtual appliances on the CONF network. The first IP address must be that of the Deployer virtual appliance. <p>Example:</p> <pre>10.254.xxx.xxx-10.254.xxx.xxx</pre> <ul style="list-style-type: none"> (Optional) Enter a valid gateway IP address if you need a route out of the CONF network. <hr/> <p>ⓘ IMPORTANT:</p> <p>IP addresses from pre-existing compute nodes, or compute nodes you plan to install, must be outside the CONF/DCM informed range during the Helion CloudSystem installation.</p> <hr/>

c. Save the changes.

10. Download and read the license agreement:

```
csdeploy save-eula --path /home/stack/
```

11. To configure a CONF IP address, run `csdeploy` and accept the license agreement:

```
csdeploy setup --eula-accepted
```

ⓘ IMPORTANT:

Running this command implies acceptance of the EULA.

12. Verify that the Deployer appliance is ready.

You can access the Deployer virtual appliance from the IP address on the CONF network.

13. Verify that you can reach vCenter and the Management host:

```
ping <vCenter_IP>
```

ⓘ IMPORTANT:

During deployment, the CONF network should be isolated. It should only be accessed by the person responsible for deploying Helion CloudSystem.

Install vCenter OVF tools package

The OVF tools package is required to support Helion CloudSystem on an ESXi Management hypervisor.

Procedure

1. Download the VMware Open Virtualization Format (OVF) Tool for Linux-64-bit:

<https://my.vmware.com/web/vmware/details?downloadGroup=OVFTOOL410&productId=490>

2. Using stack credentials, log in to the Deployer virtual appliance.
3. Switch to the `thirdparty` directory:
`cd /home/stack/stage/thirdparty`
4. Copy the VMware tools package to the `/home/stack/stage/thirdparty` folder on the Deployer virtual machine.
5. Open the `thirdparty.conf` file and make sure that the `ovftool_installer` references the same package that was added to the `home/stack/stage/thirdparty` folder.
6. Make the file executable.
`chmod +x <VMware_bundle_name>`
7. Change the owner of the file.
`chown stack:stack <VMware_bundle_name>`

Changing virtual appliance disk size

Prerequisites

- The Deployer virtual appliance is deployed on an ESXi or KVM management hypervisor.

Changes to virtual appliance disk size must be completed before deploying the remaining Helion CloudSystem virtual appliances. Once the changes are made, you cannot revert to the original disk size. The bare hLinux disk size cannot be changed.

After changing the disk size, make sure to update the memory and vCPU resources to the corresponding size. You can find a list of size requirements for each resource in the *HPE Helion CloudSystem 10.0 Update 1 Support Matrix* in the [Enterprise Information Library](#).

Procedure

1. Using administrator credentials, SSH in to the Deployer virtual appliance.
2. Open the installation configuration file in `/etc/csdeploy/examples`.
3. Scroll to the appliance-settings section of the file.
4. Uncomment the appliances-settings section of the installation configuration file.
5. Update the disk sizes and the corresponding memory and vcpus requirements.

See [Disk requirements](#) or [Disk requirements \(Enterprise-only\)](#).

Example

```
# Optional: Specify the amount of vcpus, memory and disk sizes to install in
the control-plane appliances.
# The default values are as commented below.
# Memory and Disk sizes are in GigaBytes.
MANAGEMENT-ROLE:
  memory: 8
  vcpus: 4
  disk1-management: 100
MONITORING-ROLE:
  memory: 32
  vcpus: 12
  disk1-monitoring: 150
  disk2-kafka: 150
  disk3-elasticsearch: 150
  disk4-vertica: 150
DB_RABBIT-ROLE:
  memory: 8
```

```

vcpus: 8
disk1-db-rabbit: 100
disk2-mysql: 200
CONTROLLER-ROLE:
memory: 24
vcpus: 8
disk1-controller: 100
NETWORK-ROLE:
memory: 32
vcpus: 8
ENTERPRISE-ROLE:
memory: 32
vcpus: 8
disk-enterprise: 411

```

6. Save the changes to the file.
7. Deploy the remaining Helion CloudSystem virtual appliances from the command line.

Preparing Helion CloudSystem Enterprise

Helion CloudSystem Enterprise requires two additional prerequisites before deployment.

- A MySQL JDBC driver must be downloaded and the image copied to the Deployer virtual appliance. The driver is available from the **MySQL Connector/J-JDBC driver, Feb 2017** kit in the HPE Software Depot at <http://www.hpe.com/support/softwaredepot> (search for Helion CloudSystem Software).
Customers must download this package, even if Helion CloudSystem is preinstalled or a technical services agent is engaged to assist with the installation.
- An **HPE-Helion-CloudSystem-enterprise-10.0.2.ova** must be added to the Deployer virtual appliance images directory.

Procedure

1. Download the Enterprise release package.
2. Unzip the release package and save the `HPE-Helion-CloudSystem-enterprise-10.0.2.ova` image to a local machine.
3. Rename the image to `enterprise.ova`, then copy it to the folder `/home/stack/helion/images` on the Deployer virtual appliance.
4. Download the **MySQL Connector/J-JDBC driver, Feb 2017** release package.
5. Unzip `libmysql-java_5.1.32-1_all.deb` and save it to the local machine.
6. Copy it to the folder `/home/stack/stage/thirdparty` on the Deployer virtual appliance.
7. If you are updating from Helion CloudSystem 10.0 or 10.0 Update 1, log in to **each** Enterprise appliance and clear the `configure_enterprise.log` file.

Example (from the first Enterprise appliance)

```
stack@cse-ccp-ea-m1-clm:~$ > /home/stack/configure_enterprise.log
```

Updating the MySQL JDBC driver name in the configuration file

The MySQL JDBC driver name must match the name listed in the `thirdparty.conf` file. If you changed the driver name, update the name in the `thirdparty.conf` file.

Procedure

1. From the Deployer appliance, open the `thirdparty.conf` file:

```
sudo vi /home/stack/stage/thirdparty/thirdparty.conf
```

2. Modify the value of `lib_mysql_java_file_name` to match the MySQL JDBC driver file name.
3. Save the changes.

Deploy the cloud from the installation UI

Installation UI considerations

- To define a custom size for virtual appliance memory or if you want to install multiple versions of Helion CloudSystem and configure virtual router IDs, you must deploy from the CLI.
- Define a network for hlm (CONF), admin (Data Center Management Network), internal (Cloud Management Network), and public (Consumer Access Network) to support control plane requirements. If needed, you can combine the admin and public functions into the Data Center Management Network. To make additional network role combinations, deploy from the CLI.
- Storage networks are not configured with the installation UI. If you require storage, use the `csoperate` CLI to add storage networks after deploying Helion CloudSystem. See the *HPE Helion CloudSystem 10.0 Command Line Interface Guide* in the [Enterprise Information Library](#).

Fields in the installation UI

After deploying the Deployer virtual appliance, the installation UI is accessed from the URL `http://<deployer-ip>:3000`.

Follow the onscreen prompts in the installation UI. An explanation of the fields in the UI is provided in the following table.

Table 1: Fields in the installation UI

Control plane hypervisor	Explanation
Upload configuration file	If you have an existing YAML configuration file with Helion CloudSystem values, you can load the file here. The values are automatically populated in the UI fields.
vCenter IP address or hostname	IP address or hostname of the vCenter appliance hosting the ESXi Management hypervisor NOTE: This information is automatically added to the UI.
Verify SSL	To toggle SSL verification on and off, click the square. <ul style="list-style-type: none">• Green square: SSL is enabled• No square: SSL is disabled NOTE: This information is automatically added to the UI.
vCenter username	To log in to vCenter, enter the user name. NOTE: This information is automatically added to the UI.

Table Continued

Control plane hypervisor	Explanation
vCenter password	To log in to vCenter, enter the password. <hr/> NOTE: This information is automatically added to the UI. <hr/>
vCenter datastore	Enter the name of the datastore <hr/> NOTE: This information is automatically added to the UI. <hr/>
vCenter cluster path	Enter the path in the format datacenter/host/cluster_name <hr/> NOTE: This information is automatically added to the UI. <hr/>
Cloud Configuration	Explanation
Hostname Prefix	Enter a prefix with a maximum of eight characters.
NTP Servers	At least one NTP server is required in the configuration. Add the NTP server IP addresses, separated by commas. <hr/> ! IMPORTANT: Sync the Management hypervisor and the NTP server before deployment is initiated. <hr/>
Stack Account Password	Enter the password that is required to log in to the virtual appliances using the stack account. Click the eye icon to reveal the text and verify that you entered the correct password. <hr/> ! IMPORTANT: The password must be a complex password of 12-128 alphanumeric characters and optional special characters ~!@#%&^*()_-=+{}[] ;,:.~?/. Include at least one uppercase letter, one lowercase letter, and one digit from 0-9. <hr/>
DNS Servers	(Optional) Add a DNS server that will resolve host names. Add the DNS server IP addresses, separated by commas. If you are using servers with a host-name or FQDN or a load balancer with an FQDN, you must define a DNS IP address.
Enable HA	To toggle HA on and off, click the square. <ul style="list-style-type: none"> • Green square: HA is enabled • No square: HA is disabled
Cloud Selection	Explanation

Table Continued

Control plane hypervisor	Explanation
Foundation	To add or remove Foundation, click the square. <ul style="list-style-type: none"> • Green square: Foundation is installed • No square: Foundation is not installed
Enterprise	To add or remove Enterprise, click the square. <ul style="list-style-type: none"> • Green square: Enterprise is installed • No square: Enterprise is not installed
Glance Disk Size	Enter the size of the Glance disk in GB. For glance disk size requirements, see the <i>HPE Helion CloudSystem 10.0 Update 1 Support Matrix</i> in the <u>Enterprise Information Library</u> .
OO Admin Password	Enter a password consisting of 12-20 characters with alphanumeric characters and allowed special characters ~!@#\$\$%*_ -+= () {} [] ; , . ^ ` ? / . This field is required for Enterprise. The password set here is for the HPE Operations Orchestration portal. Click the eye icon to reveal the text and verify that you entered the correct password.
Bond Configuration	Explanation
Create bond	You can define the bond mode and the interfaces that support that type of bonding. Supported modes include virtual-port, ip-hash, srcmac-hash, explicit-order, nic-load (for vDS). <ul style="list-style-type: none"> • Enter the bond name. • Select a bond mode from the drop-down list. • Check the devices that are bonded for a network. Select at least two devices.
Networking	Explanation
Network Routing Configuration	<ul style="list-style-type: none"> • DVR: Supports RHEL KVM and HPE Linux compute hosts. <hr/> <p>! IMPORTANT:</p> <p>If you select DVR, then all import, provision, and activate actions for ESXi compute resources are disabled. To re-enable these features after deploying Helion CloudSystem, you must tear down the environment and redeploy with the CVR option.</p> <hr/> <ul style="list-style-type: none"> • CVR: Supports ESXi, RHEL KVM, HPE Linux, and Hyper-V compute hosts.

Table Continued

Control plane hypervisor	Explanation
<p>CONF Network</p> <p>(click ..., then Edit to see network details)</p>	<p>By default, the CONF network is added to the Networking section. These fields define the required Cloud Configuration Network (CONF) that functions as both a Helion Lifecycle Management Network (hlm) and a PXE network.</p> <ul style="list-style-type: none"> • Vlan ID is not specified since on an untagged network. • Interface is either the physical interface name or the name of the bond. • One large address range is required for CONF. You cannot extend the address range after deployment. • CIDR value <p>Example: 10.200.24.0/23</p> <ul style="list-style-type: none"> • (Optional) Enter a valid gateway IP address if you need a route out of the CONF network. <hr/> <p>! IMPORTANT:</p> <p>CONF IP addresses from pre-existing compute nodes, or compute nodes you plan to install, must be inside the large address range identified for CONF.</p>
<p>CLM Network</p> <p>(click ..., then Edit to see network details)</p>	<p>By default, the Cloud Management Network (CLM) is added to the Networking section. Edit these fields to define the required internal network that communicates with all Helion CloudSystem virtual appliances.</p> <ul style="list-style-type: none"> • In the VLAN ID field, enter the VLAN ID you plan to use for the CLM network. • In the Interface field, select the physical interface name. <p>If this network is bonded, enter the bond name here instead of the physical interface name.</p> <ul style="list-style-type: none"> • In the CIDR field, update the CIDR value if you plan to use a value different from the default 192.168.0.0/23. • (Optional) Specify an address range and/or gateway IP address. <hr/> <p>! IMPORTANT:</p> <p>If a gateway IP address is defined for this network, then use the addresses field to define a range of IP addresses. The range must exclude the gateway IP address.</p> <p>If you do not specify the address range, the network consumes the entire CIDR.</p>

Table Continued

Control plane hypervisor	Explanation
DCM Network (click ..., then Edit to see network details)	<p>By default, the Data Center Management Network (DCM) is added to the Networking section. Edit these fields to define the required admin network that connects the Deployer appliance, Management appliance, Enterprise appliance, Cloud controller, and Network nodes to hardware and hardware management tools.</p> <ul style="list-style-type: none"> • If combining admin and public roles on this network, select both the DCM and CAN network types. • In the VLAN ID field, enter the VLAN ID you plan to use for the DCM network. • In the Interface field, select the physical interface name. <p>If this network is bonded, enter the bond name here instead of the physical interface name.</p> <ul style="list-style-type: none"> • In the Address field, enter an optional range of IP addresses that will be assigned to the virtual appliances on the DCM network. Multiple IP address ranges are supported. <p>Example: 192.168.219.100–192.168.219.224</p> <ul style="list-style-type: none"> • In the CIDR field, enter a CIDR value. <p>Example: 192.168.216.0/25</p> <ul style="list-style-type: none"> • In the Gateway IP field, enter a gateway IP address. <hr/> <p>! IMPORTANT:</p> <p>DCM IP addresses from pre-existing compute nodes, or compute nodes you plan to install, must be outside the range identified for the DCM.</p>

Table Continued

Control plane hypervisor	Explanation
<p>CAN Network</p> <p>(click ..., then Edit to see network details)</p>	<p>By default, the Consumer Access Network (CAN) is added to the Networking section. Edit these fields to define the optional public network that connects to Helion CloudSystem virtual appliance consoles and portals. Delete the CAN if it is not needed in your configuration.</p> <ul style="list-style-type: none"> • If combining admin and public roles on the DCM, then delete this network in the UI. • In the VLAN ID field, enter the VLAN ID you plan to use for the CAN network. • In the Interface field, select the physical interface name. <p>If this network is bonded, enter the bond name here instead of the physical interface name.</p> <ul style="list-style-type: none"> • In the Address field, enter an optional range of IP addresses that will be assigned to the virtual appliances on the CAN network. Multiple IP address ranges are supported. <p>Example: 30.30.30.10–30.30.30.16</p> <ul style="list-style-type: none"> • In the CIDR field, enter a CIDR value. <p>Example: 30.30.0.0/29</p> <ul style="list-style-type: none"> • In the Gateway IP field, enter an optional gateway IP address for the network. <hr/> <p>NOTE:</p> <p>You can remove the CAN network and combine the network function with a different network. Add the network type “public” to the configuration for the network taking the place of CAN.</p>

Table Continued

Control plane hypervisor	Explanation
<p>Neutron external network (click ..., then Edit to see network details)</p>	<p>The External network is not automatically added to the Networking section. If you want to include this network in your configuration, click Create Network. You do not need an External network in an Enterprise-only deployment.</p> <p>You can add the External Network after deploying Helion CloudSystem using the <code>csoperate</code> CLI commands.</p> <ul style="list-style-type: none"> • Enter a name for the network. • From Network types, select the Neutron external network. • In the VLAN ID field, enter the VLAN ID you plan to use for the External network. • In the Interface field, select the physical interface name (interface1). <p>If this network is bonded, enter the bond name here instead of the physical interface name.</p> <hr/> <p>! IMPORTANT:</p> <p>Do not enter an IP address or IP address range for this network. The installation UI allows you to enter an IP address or range. If you enter a value, the installation will display a validation error.</p>
<p>VLAN Data Trunk (click Create Network)</p>	<p>The VLAN Data Trunk is not automatically added to the Networking section.</p> <hr/> <p>! IMPORTANT:</p> <p>If you plan to use compute resources in a Helion CloudSystem Foundation environment, you must have a VLAN Data Trunk or Tunnel underlay network in the configuration.</p> <hr/> <ul style="list-style-type: none"> • Enter a name for the network. • From Network types, select VLAN underlay network (VUL). • In the Interface field, select the physical interface name. You cannot use the same interface for VxLAN and VLAN. • Enter the optional Segmentation ID range you plan to use to support Tenant networks. <p>Example: "866:868"</p> <hr/> <p>NOTE:</p> <p>If you choose not to add the VLAN Data Trunk during the initial deployment, you can add it later using the <code>csoperate</code> CLI commands.</p>

Table Continued

Control plane hypervisor	Explanation
Tunnel underlay network (click Create Network)	<p>The Tunnel Underlay Network is not automatically added to the Networking section.</p> <hr/> <p>! IMPORTANT:</p> <p>If you plan to use compute resources in a Helion CloudSystem Foundation environment, you must have a VLAN Data Trunk or Tunnel underlay network in the configuration.</p> <hr/> <ul style="list-style-type: none"> • Enter a name for the network. • From Network types, select Tunnel underlay network (TUL). • In the VLAN ID field, enter the VLAN ID you plan to use for the Tunnel Underlay Network. • In the Interface field, select the physical interface name. <p>If this network is bonded, enter the bond name here instead of the physical interface name.</p> <ul style="list-style-type: none"> • In the Addresses field, enter an optional range of IP addresses that will be assigned to the virtual appliances on the network. Multiple IP address ranges are supported. <p>Example: 192.168.97.100–192.168.103.224</p> <ul style="list-style-type: none"> • In the CIDR field, enter a CIDR value. <p>Example: 192.168.97.0/25</p> <ul style="list-style-type: none"> • In the Gateway IP field, enter the optional gateway IP address for the network. • In the Segmentation ID Range field, enter the optional Segmentation ID range you want to use to support Tenant Networks. <p>Example: 5000:9000</p> <hr/> <p>NOTE:</p> <p>If you choose not to add the Tunnel Underlay Network during the initial deployment, you can add it later using the <code>csoperate</code> CLI commands.</p> <hr/>
Load Balancer	Explanation

Table Continued

Control plane hypervisor	Explanation
Public	<p>These fields capture the endpoint for public access to appliance portals and management consoles, as well as certificate details.</p> <ul style="list-style-type: none"> In the VIP field, enter an IP address in the range for the Consumer Access Network (CAN). <hr/> <p>NOTE:</p> <p>If you combined the internal and public network roles on the Data Center Management Network, the first IP address is reserved for the internal VIP. Make sure that the public VIP and admin VIP are unique IP addresses.</p> <hr/> <ul style="list-style-type: none"> Enter an optional Fully Qualified Domain Name. The endpoint is the public FQDN for accessing services from the Consumer Access Network. If this field is blank, the IP address specified in the VIP field will be used for the endpoints. Enter an optional email address for the administrator who is managing certificates. If you plan to use a customer-specific certificate, paste the certificate text in the Certificate box. <p>The certificate file is saved in <code>/home/stack/helion/my_cloud/config/tls/certs/</code>.</p> <hr/> <p>NOTE:</p> <p>If you leave the certificate field empty, then Helion CloudSystem will generate a self-signed certificate.</p> <p>For information about certificates, see the "Manage certificates" chapter in the <i>HPE Helion CloudSystem 10.0 Update 1 Administrator Guide</i> at Enterprise Information Library.</p> <hr/>
Admin	<p>These fields capture the endpoint for admin access to appliance portals and management consoles, as well as certificate details.</p> <ul style="list-style-type: none"> In the VIP field, enter an IP address in the range for the Data Center Management Network (DCM). <hr/> <p>NOTE:</p> <p>If you combined the internal and public network roles on the Data Center Management Network, the first IP address is reserved for the internal VIP. Use an IP address that is outside the range for the internal and admin networks. Make sure that the public VIP and admin VIP are unique IP addresses.</p> <hr/> <ul style="list-style-type: none"> Enter an optional Fully Qualified Domain Name. The endpoint is the public FQDN for accessing services from the Data Center Management Network. If this field is blank, the IP address specified in the VIP field will be used for the endpoints.

Control plane hypervisor	Explanation
	<ul style="list-style-type: none"> • Enter an optional email address for the administrator who is managing certificates. • If you plan to use a customer-specific certificate, paste the certificate text in the Certificate box. <p>The certificate file is saved in <code>/home/stack/helion/my_cloud/config/tls/certs/</code>.</p> <hr/> <p>NOTE:</p> <p>If you leave the certificate field empty, then Helion CloudSystem will generate a self-signed certificate.</p> <p>For information about certificates, see the "Manage certificates" chapter in the <i>HPE Helion CloudSystem 10.0 Update 1 Administrator Guide</i> at Enterprise Information Library.</p>

Initiate the deployment action from the installation UI

After a successful deployment, you can go back in to the installation UI, but the fields cannot be edited and you cannot deploy again or restore default values.

Procedure

1. After all the required information is added in the UI fields, click **Deploy Helion CloudSystem**.
2. Verify the validation messages.
 - a. If all information is correct, the deployment starts automatically.
 - b. If a warning is displayed during validation, click **View** to see the warning details. You can click **Continue** to ignore the warning or click **Close** to go back and update the deployment details. If you edit fields, you must click **Deploy Helion CloudSystem** again.
 - c. If an error is discovered during validation, click **View** to see the error details, then fix the error and deploy again.
 - d. If you want to download a copy of the configuration file to use for a new deployment, click **Download Configuration File**.
 - e. To restore default settings to the UI fields, click the **Restore defaults** button.

Access configuration files and log files

During deployment, you can access log files in the UI by clicking the **View Log** button. The actual log files are saved on the Deployer virtual appliance in the `/var/log/cloudsystem-installer-ui/` directory.

The UI also generates a configuration file. This file is saved at `/opt/stack/service/cloudsystem-installer-ui/venv/model/cs-install-options.yml`.



TIP:

Use the `csdeploy export-config` command to export this configuration file after running the installation UI, in case you want to reuse the file for new deployments from the CLI.

For a detailed explanation of this command, see the *HPE Helion CloudSystem 10.0 Command Line Interface Guide* in the [Enterprise Information Library](#).

If you want to move the `cs-install-options.yml` file, encrypt the file before moving it.

Deployment errors in log files

All information about deployment activities, including stack traces, is sent to the log file `/var/log/csdeploy/csdeploy.log`.

Once the log file size exceeds 100 MB, it is compressed to save disk space.

See [Using log files to find information on deployment errors](#).

Shut down the Helion CloudSystem installation UI after deployment

After successfully deploying Helion CloudSystem, it is important to shut down the installation UI.

Procedure

1. Using stack credentials, log in to the Deployer virtual appliance.
2. Shut down the installation service:

```
sudo systemctl stop cloudsystem-installer-ui
```
3. Disable the service permanently:

```
sudo systemctl disable cloudsystem-installer-ui
```

The service log is available in the directory `/var/log/cloudsystem-installer-ui/service.log`. This log does not contain passwords. Consider removing this log for security purposes.

Manually distribute virtual machines to support HA failover

Prerequisites

- Helion CloudSystem was successfully installed.
- You have administrative privileges in vCenter.

If the DRS feature is not supported in the ESXi management host, you must manually redistribute the virtual machines across all three hosts in the management cluster.

Procedure

1. Using administrative credentials, log in to vCenter.
2. Expand the cluster and host to reveal the associated virtual machines.
3. Select a virtual machine, right-click and click **Migrate > Host**.
4. Redistribute the hosts in the following order:
 - **ESXi host1**
 - ccp-cc-m1
 - ccp-db-m1
 - ccp-mgmt-m1
 - ccp-mon-m1
 - ccp-nn-m1
 - deployer
 - **ESXi host2**
 - ccp-cc-m2
 - ccp-db-m2
 - ccp-mgmt-m2

- ccp-mon-m2
- ccp-nn-m2
- **ESXi host3**
 - ccp-cc-m3
 - ccp-db-m3
 - ccp-mgmt-m3
 - ccp-mon-m3
 - ccp-nn-m3

Deploy the cloud from the CLI

After deploying the Deployer virtual appliance, you must configure the installation configuration file, and then use the `csdeploy` CLI commands to initiate the deployment process. During deployment, Helion CloudSystem virtual appliances are created and the networks are attached.

Process overview:

Procedure

1. [Configure the installation configuration file](#) on page 34
2. [Deploy Helion CloudSystem virtual appliances](#) on page 44

Helion CloudSystem installation configuration file

The installation configuration file contains all the characteristics of the Management hypervisor and the cloud network infrastructure.

There are several prepopulated configuration files in the Deployer virtual appliance directory `/etc/csdeploy/examples`.



TIP:

When you open an example file, you will see the vCenter settings and the CONF network settings are already populated in the file. These values are pulled from the `cs-setup-esx.yml` file.

Choose the sample configuration file that best represents the cloud environment you plan to create. You will edit this configuration file in the following procedure.



IMPORTANT:

Do not edit any other `.yaml` files in the Deployer virtual appliance unless directed to do so in an installation procedure. Modifying the `.yaml` input models could cause configurations to become out of sync.

- **cs-install-options-ESX_Hyper-V-hlinux-RHEL-Vlan-CVR-VSA-3PAR-vmrk-swift.yml**

This sample configuration file supports:

- Management networks: CONF, DCM, CLM, CAN, EXT
- Compute resources: ESXi, Hyper-V, HPE Linux, RHEL KVM
- Routing: CVR
- Data network: VLAN Data Trunk
- Storage networks: iSCSI, OBS, OBP

- **cs-install-options-Enterprise-Only.yml**

This sample configuration file supports:

- Management networks: CONF, DCM, CLM, CAN, EXT

- **cs-install-options-RHEL-hlinux-Vxlan-DVR-VSA-3PAR.yml**

This sample configuration file supports:

- Management networks: CONF, DCM, CLM, CAN, EXT
- Compute resources: HPE Linux, RHEL KVM
- Routing: DVR
- Data network: Tunnel Underlay Network
- Storage networks: iSCSI

- **cs-install-options-ccp-only.yml**

This sample configuration file supports:

- Management networks: CONF, DCM, CLM, CAN, EXT
- Data network: Tunnel Underlay Network and VLAN Data Trunk

- **cs-install-options-hlinux-RHEL-Vlan+Vxlan-CVR-Ceph.yml**

This sample configuration file supports:

- Management networks: CONF, DCM, CLM, CAN, EXT
- Compute resources: HPE Linux, RHEL KVM
- Routing: CVR
- Data network: Tunnel Underlay Network and VLAN Data Trunk
- Storage networks: OSDI, OSDC

- **cs-install-options-hlinux-RHEL-Vlan-CVR-VSA-3PAR.yml**

This sample configuration file supports:

- Management networks: CONF, DCM, CLM, CAN, EXT
- Compute resources: HPE Linux, RHEL KVM
- Routing: CVR
- Data network: VLAN Data Trunk
- Storage networks: iSCSI

- **cs-install-options-poc.yml**

This sample configuration file supports:

- Management networks: CONF, DCM (with roles internal, admin and public), EXT
- Data network: Tunnel Underlay Network and VLAN Data Trunk

Configure the installation configuration file

 **IMPORTANT:**

Make sure that you understand network definitions and the supported bonding configurations before working with the configuration file.

You must have a network type defined for hlm, admin, and internal in the configuration. If you want to combine networks, review the network characteristics defined in the *HPE Helion CloudSystem 10.0 Network Planning Guide* in the [Enterprise Information Library](#).

Procedure

1. Using stack credentials, log in to the Deployer virtual appliance.
The default stack credentials are `stack/stack`. You will change these credentials later in the installation configuration file.
2. Switch to the examples directory:

```
cd /etc/csdeploy/examples
```
3. Open the example file that most closely resembles your cloud requirements:

```
vi cs-install-options-ESX_Hyper-V-hlinux-RHEL-Vlan-CVR-VSA-3PAR-vmrk-swift.yml
```
4. Create a copy of the configuration file and save it as `cs-install-options.yml` in the directory `/etc/csdeploy`.
5. Add restrictive permissions to the file, such as 640, to protect the sensitive data that will be added to the file.
6. Update the configuration file.

See [Fields in the installation configuration file](#).

7. Save the installation configuration file.
8. If you extracted the configuration file and edited it with a text editor, paste the YAML in YAML Lint to validate the structure. Do not paste from YAML Lint back into the configuration file. Just use this as a validation.

See [YAML Lint](#).

9. Save the file to the `/etc/csdeploy` directory.

Fields in the installation configuration file

Table 2: Fields in the installation configuration file

Configuration field	Explanation
host-prefix	Enter a prefix with a maximum of eight characters.
enable-ha	<ul style="list-style-type: none"> • True: HA is enabled • False: HA is not enabled
enable-dvr	<ul style="list-style-type: none"> • True: DVR is used <p>Only RHEL KVM and HPE Linux compute hosts are supported with DVR.</p> <hr/> <p>! IMPORTANT:</p> <p>If you select DVR, then all import, provision, and activate actions for ESXi compute resources are disabled. To re-enable these features after deploying Helion CloudSystem, you must tear down the environment and redeploy with the CVR option.</p> <hr/> <ul style="list-style-type: none"> • False: CVR is used <p>ESXi, RHEL KVM, HPE Linux, and Hyper-V compute hosts are supported with CVR.</p>
<pre>dns-settings: nameservers: - ip-address-1 - ip-address-2</pre>	<p>(Optional) Add a DNS server that will resolve the hostnames. To add a server, replace “- ip-address-1” with the IP address of the first DNS server and repeat that process if you need a second DNS server.</p> <p>If you are using servers with a host-name or FQDN or a load balancer with an FQDN, you must define a DNS IP address.</p>
<pre>ntp-servers: -</pre>	<p>At least one NTP server is required in the configuration. To add a server, enter the IP address after the dash. If the host name is resolved by a name server, then you can optionally add a host name.</p> <p>Example: – 10.254.72.64</p> <hr/> <p>! IMPORTANT:</p> <p>The Management hypervisor and the NTP server must be in sync before the deployment is initiated.</p>

Table Continued

Configuration field	Explanation
appliance-sizings	<p>The virtual CPU and memory requirements for each type of Helion CloudSystem virtual appliance. Minimum requirements are prepopulated in the file.</p> <ul style="list-style-type: none"> • MANAGEMENT-ROLE 4 vCPU, 16 GB memory • MONITORING-ROLE 4 vCPU, 16 GB memory • DB-RABBIT-ROLE 4 vCPU, 16 GB memory • CONTROLLER-ROLE 4 vCPU, 16 GB memory • NETWORK-ROLE 4 vCPU, 8 GB memory • ENTERPRISE-ROLE 8 vCPU, 32 GB memory
appliance-ssh-password:	<p>Enter the password that is required to log in to the virtual appliances using the stack account.</p> <p>The password must be a complex password of 12-128 alphanumeric characters and optional special characters ~!@#%&^*()_-=+{}[] ;,:. `?/. Include at least one uppercase letter, one lowercase letter, and one digit from 0-9.</p> <p>If a special character is used in the password, the value must be between single quotes.</p>
type:	<p>Enter esx to indicate that you are using an ESXi Management hypervisor.</p>
vcenter-hostname:	<p>Enter the IP address or the FQDN of the vCenter appliance.</p> <p>If a special character is used in the host name, the value must be between single quotes.</p>
vcenter-verify-ssl:	<p>True: You have a certificate installed in vCenter for secure connections.</p> <p>False: No certificate is installed in vCenter.</p>
vcenter-username:	<p>Enter the user name for the vCenter appliance.</p> <p>If a special character is used in the user name, the value must be between single quotes.</p>
vcenter-password:	<p>Enter the password for the vCenter appliance.</p> <p>If a special character is used in the password, the value must be between single quotes.</p>

Table Continued

Configuration field	Explanation
vcenter-cluster-path:	Enter the path in the format: <code><datacenter_name>/host/<cluster_name></code>
vcenter-datastore:	Enter the name of the datastore Example: SanDisk:env23-mgmt If a special character is used in the datastore name, the value must be between single quotes.
<pre> bonds: bond0: devices: -- vmnic1 -- vmnic2 mode: virtual-port bond1: devices: -- vmnic3 -- vmnic4 mode: nic-load </pre>	You can define the bond/team mode and the interfaces that support that type of bonding. <ul style="list-style-type: none"> • Enter the bond name in the bondx field (bond0, bond1, bond2). • List the interfaces that are bonded for a network. If you plan to use bonded or teamed NICs , replace the physical interface name in the network iface field with the bond name.
Network Configuration	
<pre> --name: CONF types: - hlm members: - cidr: addresses: - iface: gateway-ip: </pre>	This network is the required Cloud Configuration Network that functions as both a Helion Lifecycle Management Network (hlm) and a PXE network. <ul style="list-style-type: none"> • vlanid must be 0 for an untagged network. • Enter a CIDR value. Example: 10.200.24.0/23 • The address range is required for CONF. Enter one large range of IP addresses that will be assigned to the virtual appliances on the CONF network. This range must include the IP address of the Deployer virtual appliance. • The iface is the physical interface name. If using bonding, the iface is the bond name. • (Optional) Enter a valid gateway IP address if you need a route out of the CONF network. <hr/> <p>! IMPORTANT: CONF IP addresses from pre-existing compute nodes, or compute nodes that you plan to install, must be inside the range identified for CONF.</p> <hr/>

Table Continued

Configuration field	Explanation
<pre>--name: DCM types: - admin members: - vlanid: cidr: addresses: - gateway-ip: iface:</pre>	<p>This section defines the required admin network that connects the Deployer appliance, Management appliance, Enterprise appliance, Cloud controller, and Network node to hardware and hardware management tools.</p> <ul style="list-style-type: none"> Update the VLAN ID with the VLAN ID you plan to use for the DCM network. Enter a CIDR value. Example: 192.168.216.0/24 Enter a list of single IP addresses or a range of IP addresses that will be assigned to the virtual appliances on the DCM network. Example: 192.168.219.100–192.168.219.224 Enter the gateway IP. The iface is the physical interface name. If using bonding, the iface is the bond name. <hr/> <p>! IMPORTANT: DCM IP addresses from pre-existing compute nodes, or compute nodes you plan to install, must be outside the range identified for DCM.</p> <hr/>
<pre>--name: CLM types: - internal members: - vlanid: cidr: iface:</pre>	<p>This section defines the required internal network that communicates with all Helion CloudSystem virtual appliances.</p> <ul style="list-style-type: none"> Update the VLAN ID with the VLAN ID you plan to use for the CLM network. Enter a CIDR value. Example: 192.168.0.0/23 (Optional) Enter a list of single IP addresses or a range of IP addresses that will be assigned to the virtual appliances on the CLM network. Example: 192.168.28.xx–192.168.28.xx The iface is the physical interface name. If using bonding, the iface is the bond name. (Optional) Specify a gateway IP address. <hr/> <p>! IMPORTANT: If a gateway IP address is defined for this network, then use the addresses field to define a range of IP addresses. The range must exclude the gateway IP address. If you do not specify the address range, the network consumes the entire CIDR.</p> <hr/>

Table Continued

Configuration field	Explanation
<pre>--name: CAN types: - public members: - vlanid: cidr: addresses: - iface:</pre>	<p>These fields define the optional Consumer Access Network (CAN).</p> <ul style="list-style-type: none"> Update the VLAN ID with the VLAN ID you plan to use for the CAN network. Enter a CIDR value. Example: 30.30.0.0/28 Enter a list of single IP addresses or a range of IP addresses that will be assigned to the virtual appliances on the CAN network. Example: 30.30.30.2–30.30.30.12 The iface is the physical interface name. If using bonding, the iface is the bond name.
<pre>--name: External types: - external members: - vlanid: iface:</pre>	<p>These fields define the External network. You only need an External network if you are including Helion CloudSystem Foundation in your cloud. You can comment out the External network if it is not needed in your configuration.</p> <p>You can add the External Network after deploying Helion CloudSystem using the <code>csoperate</code> CLI commands.</p> <ul style="list-style-type: none"> Update the VLAN ID with the VLAN ID you plan to use for the External network. The iface is the physical interface name. If using bonding, the iface is the bond name.
<pre>--name: CloudVxLAN types: - tul members: - vlanid: cidr: addresses: - gateway-ip: iface: range:</pre>	<p>These fields define the Tunnel Underlay Network.</p> <hr/> <p>! IMPORTANT: If you plan to use compute resources in a Helion CloudSystem Foundation environment, you must have a VLAN Data Trunk or Tunnel Underlay Network in the configuration.</p> <hr/> <ul style="list-style-type: none"> Update the VLAN ID with the VLAN ID you plan to use for the Tunnel Underlay Network. Enter a CIDR value. Example: 192.168.97.0/25 Enter a list of single IP addresses or a range of IP addresses that will be assigned to the virtual appliances on the EXT network. Example: 192.168.97.100–192.168.103.224 Enter the optional gateway-IP. The iface is the physical interface name. If using bonding, the iface is the bond name. Define the range you want to use. Example: 5000:9000

Table Continued

Configuration field	Explanation
<pre>--name: CloudVLAN types: - vul members: - iface: range:</pre>	<p>These fields define the VLAN Data Trunk.</p> <hr/> <p>! IMPORTANT:</p> <p>If you plan to use compute resources in a Helion CloudSystem Foundation environment, you must have a VLAN Data Trunk or Tunnel Underlay Network in the configuration.</p> <hr/> <ul style="list-style-type: none"> The iface is the physical interface name. If using bonding, the iface is the bond name. Define the range you want to use. Example: "866:868, 901:915"
Storage Networking Details	
<pre>--name: iSCSI types: - bls members: - vlanid: cidr: iface:</pre>	<p>If you plan to use HPE StoreVirtual VSA or HPE 3PAR block storage, add the configuration details for the iSCSI network.</p> <ul style="list-style-type: none"> Update the VLAN ID with the VLAN ID you plan to use for the iSCSI Network. Enter a CIDR value. Example: 192.168.50.0/24 Enter the physical interface name. If using bonding, the iface is the bond name.
<pre>--name: OBS types: - obs members: - vlanid: cidr: iface:</pre>	<p>One of two networks required to support OpenStack Swift object storage.</p> <ul style="list-style-type: none"> Update the VLAN ID with the VLAN ID you plan to use for the OBS Network. Enter a CIDR value. Example: 192.168.40.0/25 Enter the physical interface name. If using bonding, the iface is the bond name.
<pre>--name: OBP types: - obp members: - vlanid: cidr: iface:</pre>	<p>One of two networks required to support OpenStack Swift object storage.</p> <ul style="list-style-type: none"> Update the VLAN ID with the VLAN ID you plan to use for the OBP Network. Enter a CIDR value. Example: 192.168.50.0/25 Enter the physical interface name. If using bonding, the iface is the bond name.

Table Continued

Configuration field	Explanation
<pre>--name: OSDI types: - osdi members: - vlanid: cidr: iface:</pre>	<p>One of two networks required to support OpenStack Ceph block and object storage.</p> <ul style="list-style-type: none"> Update the VLAN ID with the VLAN ID you plan to use for the OSDI Network. Enter a CIDR value. Example: 192.168.30.0/24 Enter the physical interface name. If using bonding, the iface is the bond name.
<pre>--name: OSDC types: - osdc members: - vlanid: cidr: iface:</pre>	<p>One of two networks required to support OpenStack Ceph block and object storage.</p> <hr/> <p>! IMPORTANT:</p> <p>If you plan to use Ceph storage for KVM compute nodes, then define the OSDC Network in the configuration file during deployment or with the <code>csoperate</code> CLI after deployment. The OSDC Network must be defined before activating compute nodes. If you activate compute nodes before defining this network, the compute nodes will not be able to access it.</p> <hr/> <ul style="list-style-type: none"> Update the VLAN ID with the VLAN ID you plan to use for the OSDC Network. Enter a CIDR value. Example: 192.168.40.0/24 Enter the physical interface name. If using bonding, the iface is the bond name.
Load Balancer Details	
vrid:	<p>The optional vrid (Virtual Router Identifier) must be an integer value from 1 to 253. If a value is supplied, then the next two numbers will also be used when configuring the cloud. For example, if 7 is entered in the vrid field, then 7, 8 and 9 will be used.</p>

Table Continued

Configuration field	Explanation
Public settings	<p>These fields capture the endpoint for public access to appliance portals and management consoles, as well as certificate details.</p> <ul style="list-style-type: none"> In the vip field, enter an IP address in the range for the Consumer Access Network, or the network configured to provide public access. <hr/> <p>NOTE:</p> <p>If you combined the internal and public network roles on a single network, do not use the first IP address in the range. The first IP address is reserved for the internal VIP.</p> <hr/> <ul style="list-style-type: none"> If the certificate: field is left blank, Helion CloudSystem will generate a self-signed certificate named <code>cs-public-cert</code>. If a certificate name is entered in this field, it must match a user-supplied certificate that is saved in <code>/home/stack/helion/my_cloud/config/tls/certs/</code>. <p>For information about certificates, see the "Manage certificates" chapter in the <i>HPE Helion CloudSystem 10.0 Update 1 Administrator Guide</i> at Enterprise Information Library.</p> <ul style="list-style-type: none"> In the certificate-email field, enter an optional email address for the administrator who is managing certificates. In the fqdn field, enter an optional Fully Qualified Domain Name. This endpoint is the public FQDN for accessing services from the Consumer Access Network. If this field is blank, the IP address will be used for the endpoints. <hr/> <p>NOTE:</p> <p>If using an FQDN, you must have a DNS installed that contains an entry for the FQDN/VIP. This information must be present before initiating the deployment.</p> <hr/>

Table Continued

Configuration field	Explanation
Admin settings	<p>These fields capture the endpoint for admin access to appliance portals and management consoles, as well as certificate details.</p> <ul style="list-style-type: none"> In the vip field, enter an IP address in the range for the Data Center Management Network, or the network configured to provide admin access. <hr/> <p>NOTE:</p> <p>If you combined the internal and admin network roles on a single network, do not use the first IP address in the range. The first IP address is reserved for the internal VIP. Use an IP address that is outside the range for the internal and admin networks. Make sure that the public VIP and admin VIP are unique IP addresses.</p> <hr/> <ul style="list-style-type: none"> If the certificate: field is left blank, Helion CloudSystem will generate a self-signed certificate named <code>cs-admin-cert</code>. If a certificate name is entered in this field, it must match a user-supplied certificate that is saved in <code>/home/stack/helion/my_cloud/config/tls/certs/</code>. <p>For information about certificates, see the "Manage certificates" chapter in the <i>HPE Helion CloudSystem 10.0 Update 1 Administrator Guide</i> at Enterprise Information Library.</p> <ul style="list-style-type: none"> In the certificate-email field, enter an email address for the administrator who is managing certificates. In the fqdn field, enter an optional Fully Qualified Domain Name. This endpoint is the public FQDN for accessing services from the Data Center Management Network. If this field is blank, the IP address will be used for the endpoints. <hr/> <p>NOTE:</p> <p>If using an FQDN, you must have a DNS installed that contains an entry for the FQDN/VIP. This information must be present before initiating the deployment.</p> <hr/>
glance-disk-size	<p>Enter the size of the Glance disk in GB. For glance disk size requirements, see the <i>HPE Helion CloudSystem 10.0 Update 1 Support Matrix</i> in the Enterprise Information Library.</p>
oo-admin-password	<p>Enter a password consisting of 12-20 characters with alphanumeric characters and allowed special characters <code>~!@#\$%*_ -+= () {} [] ; , . ^ ` ? /</code>.</p> <p>This field is required for Enterprise. It is the password required to access the HPE Operations Orchestration portal.</p> <hr/> <p>NOTE:</p> <p>If a special character is used in the password, it must be inside single quotes.</p> <hr/>

Validate the installation configuration file

Before deploying Helion CloudSystem virtual appliances, you can validate the integrity of the configuration file.

Procedure

1. Verify that the configuration file is saved in the path `/etc/csdeploy/cs-install-options.yml`.
2. Using the stack credentials defined in the installation configuration file, log in to the Deployer virtual appliance.
3. Run the validate command:

```
csdeploy validate --config /etc/csdeploy/cs-install-options.yml --add-foundation --add-enterprise
```

NOTE:

If you are only installing Helion CloudSystem Enterprise, remove `--add-foundation` from the command.

If you are only installing Helion CloudSystem Foundation, remove `--add-enterprise` from the command.

4. Resolve all errors.

Warnings will not prevent the deployment action from completing.

NOTE:

If the validation returns the error "Encountered a parsing error", check the section in the installation configuration file that is referenced in the error for an element that was removed or edited incorrectly. Fix the error, then run `csdeploy` again.

Deploy Helion CloudSystem virtual appliances

Prerequisites

- The installation configuration file was validated and all errors were resolved.
See [Validate the installation configuration file](#).
- If you want to use the Octavia load balancer, it is already installed.
See [Installing the Octavia load balancer](#).

Procedure

1.  **IMPORTANT:**
Do not run the commands in this procedure as the root user, unless indicated.
1. Run `csdeploy`, referencing the configuration file you created (`cs-install-options.yml`).
 - a. **Foundation only installation**

```
csdeploy start --config /etc/csdeploy/cs-install-options.yml --add-foundation 2>&1 | tee /home/stack/csdeploy-foundation.log
```
 - b. **Enterprise only installation**

```
csdeploy start --config /etc/csdeploy/cs-install-options.yml --add-enterprise 2>&1 | tee /home/stack/csdeploy-enterprise.log
```
 - c. **Foundation and Enterprise installation**

```
csdeploy start --config /etc/csdeploy/cs-install-options.yml --add-foundation --add-enterprise 2>&1 | tee /home/stack/csdeploy-foundent.log
```

2. Verify that the PLAY RECAP message shows a summary section with the message `Installation Completed Successfully!`

NOTE:

Installation times vary according to the type of hardware used to deploy Helion CloudSystem. It may take several hours for this message to display.

3. Verify the network configuration in vCenter.

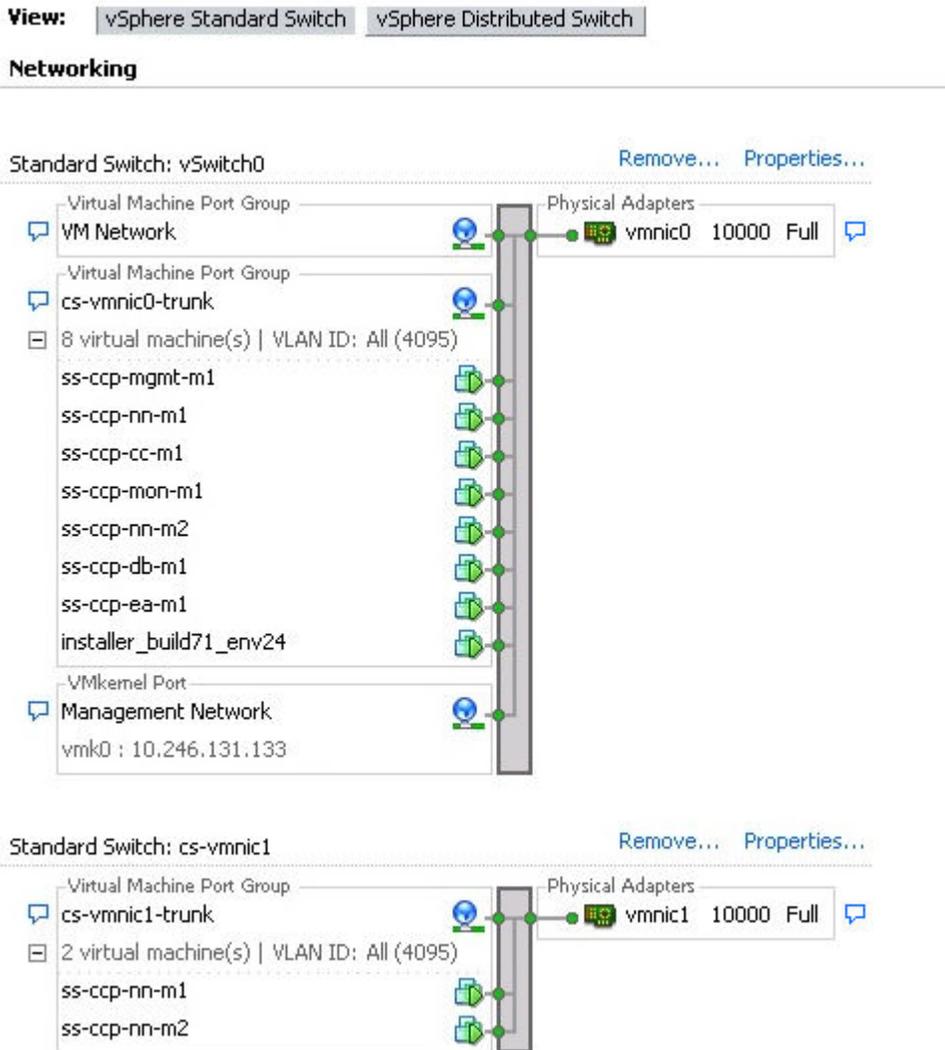


Figure 3: Sample vCenter with VSS after deploying Helion CloudSystem

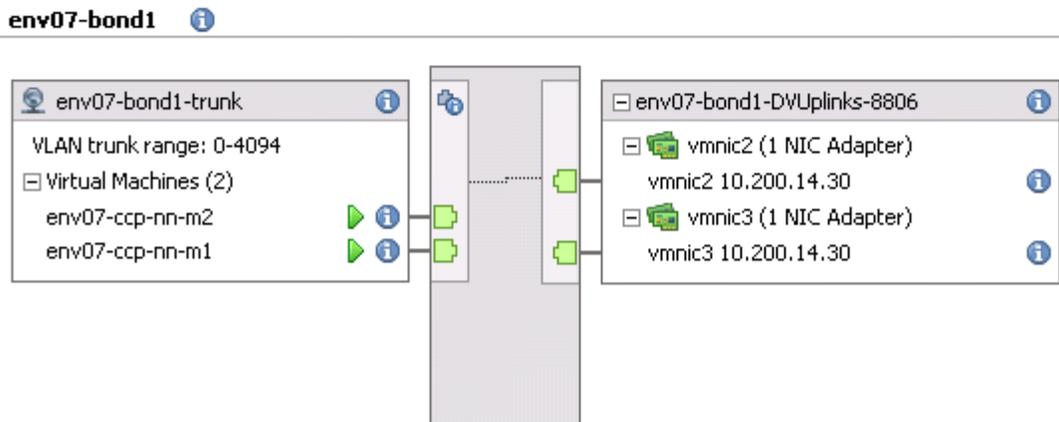
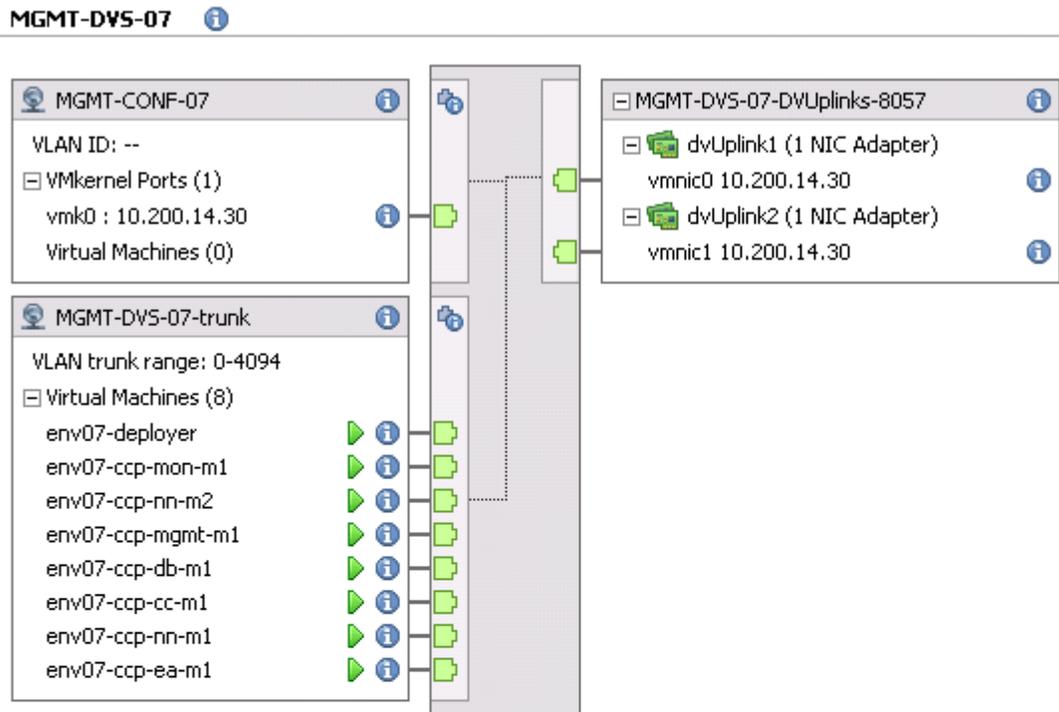


Figure 4: Sample vCenter with VDS after deploying Helion CloudSystem

If you decide to add Enterprise to Helion CloudSystem after deploying the virtual appliances, see [Adding Enterprise to a Foundation-only deployment](#).

If the deployment failed, see [Reset the environment after a failed deployment](#).

Modify Affinity rules

Helion CloudSystem creates all DRS rules during installation. You can modify the DRS rules post deployment.

Procedure

1. To modify the rules from the vSphere client, select

Management cluster > Edit settings > vSphere DRS > Rules.

2. To modify the rules from the vSphere Web Client, select

Management cluster > Manage tab > Setting tab > Configuration > VM/Host Rules.

Troubleshoot installation errors

Using `csdeploy` log files to find information on deployment errors

Prerequisites

- The initial Helion CloudSystem deployment action failed.
- A log file was saved to `/var/log/csdeploy/csdeploy.log`.

The `csdeploy` log files contain all messages related to the deployment process. Use these logs to find errors that occurred during deployment.

Procedure

1. From the Deployer virtual appliance, find the log files in `/var/log/csdeploy/csdeploy.log`.
2. Open each log file and look for `CSDEPLOY START begin` and `CSDEPLOY START end`.

All messages are included in this section of the log.

3. Resolve any errors listed in the start section of the log file, then retry the deployment.

Viewing all failed deployment jobs and tasks

If an error occurs while deploying Helion CloudSystem, use the `csdeploy job` command to find the details of the errors.

Procedure

1. Using the stack credentials, log in to the Deployer virtual appliance.
2. Request a list of all `csdeploy` jobs:
`csdeploy job`
3. Find the `csdeploy` job ID and then request a list of all failed jobs:
`csdeploy job --id <ID_number>|grep Failed`
4. Find the failed task ID and then request a list of task details:
`csdeploy task --id <ID_number>`
5. Resolve deployment errors.

Reset the environment after a failed deployment

If the deployment fails at any point, you can reset the environment to its original state and try the deployment again.

Procedure

1. In vCenter, remove all Helion CloudSystem virtual machines.
2. Delete the port groups and switches created during deployment.
3. Delete any DRS rules created during deployment.
4. Retry the installation.

Configure the cloud after deployment

After Helion CloudSystem is deployed:

- In Foundation, configure resources including networks, storage, and compute hosts

See "Resource configuration in Helion CloudSystem" in the *HPE Helion CloudSystem 10.0 Update 1 Administrator Guide* at [Enterprise Information Library](#).

- In Enterprise, configure providers and optionally upload additional content

See "HPE CSA providers and content" in the *HPE Helion CloudSystem 10.0 Update 1 Administrator Guide* at [Enterprise Information Library](#).

NOTE:

If you plan to use **Capsule HPE OneView - Topology** or **Capsule HPE ICsp - Topology** in a new installation of Helion CloudSystem 10.0 Update 2, see "Installing HPE OneView and HPE ICsp Topology Components" in the *HPE Helion CloudSystem 10.0 Update 1 Administrator Guide*.

Access Helion CloudSystem virtual appliances

Access Helion CloudSystem virtual appliances

To access portals and user interfaces hosted on Helion CloudSystem virtual appliances, use the URL and login credentials for each UI.

Each URL contains a VIP address for admin and public access to portals and user interfaces.

Determine VIP addresses for admin and public access of portals

Both HA and non-HA configurations will have a VIP address for the DCM and CAN networks.

VIP addresses are displayed in the **Summary** section of the PLAY RECAP message after a successful installation. You can also find VIP addresses in the installation logs.

Access appliance virtual machine portals

To access appliance virtual machines, SSH in to the virtual machine.

Credentials:

- default user name: stack
- password: (defined during deployment in the configuration file field `appliance-ssh-password`)

Access the Operations Console on the Management appliance

Use one of the following options to access the Operations Console.

- From the Data Center Management Network (admin): `https://DCM-VIP:9095`
- From the Consumer Access Network (public): `https://CAN-VIP:9095`

Credentials:

- User name: admin
- password: (find it from the Deployer appliance using the following command)

```
sudo grep -r OS_PASSWORD ~/keystone.osrc
```

Access the Logging Dashboard (Kibana) on the Management appliance

Use one of the following options to access the Logging Dashboard.

- From the Operations Console, select **Menu > General > Logging**
- From the Data Center Management Network (admin): `https://DCM-VIP:5601`

Credentials:

- User name: kibana
- password: (find it from the Deployer appliance using the following command)

```
sudo grep -r logging_kibana_password ~/scratch/ansible/next/hos/ansible/group_vars/
```

Access the HA Proxy on the Management appliance

The HA Proxy works in conjunction with the keepalived load balance configuration to provide high availability support for the environment. The HA Proxy is hosted on the Management appliances.

Access the HA Proxy with the IP address on the Data Center Management Network:

- From the first Management appliance: `http://MA1-DCM-IP:1993`
- From the second Management appliance: `http://MA2-DCM-IP:1993`
- From the third Management appliance: `http://MA3-DCM-IP:1993`

Credentials:

- User name: admin
- password: (find it from any Management appliance using the following command)

```
sudo grep -r auth /etc/haproxy/haproxy.cfg
```

Access the OpenStack user portal on the Cloud controller appliance

Use one of the following options to access the OpenStack user portal:

- Administrators access the OpenStack user portal from the Data Center Management Network (admin): `https://DCM-VIP`
- Cloud users access the OpenStack user portal from the Consumer Access Network (public): `https://CAN-VIP`

Credentials:

- domain: default
- User name: admin
- password: (find it from the Deployer appliance using the following command)

```
sudo grep -r OS_PASSWORD ~/keystone.osrc
```

Access HPE CSA, Marketplace portal, and HPE OO on the Enterprise appliance

ⓘ IMPORTANT:

If you configured the load balancer with FQDNs instead of VIP addresses, update the service endpoints before accessing portals on the Enterprise appliance. See [Update service endpoints with the FQDN on Helion CloudSystem Enterprise appliances](#).

Access the **HPE CSA** portal:

- From the Data Center Management Network (admin): `https://DCM-VIP:8444/csa`
- From the Consumer Access Network (public): `https://CAN-VIP:8444`

Credentials:

- User name: admin
- password: cloud

Access the **Marketplace portal**:

- From the Consumer Access Network (public): `https://CAN-VIP:8089`

NOTE:

If you try to access the Marketplace Portal from `https://DCM-VIP:8089`, it will automatically redirect you to the public endpoint on the Consumer Access Network.

Credentials:

- User name: consumer
- password: cloud

Access HPE Operations Orchestration:

- From the Data Center Management Network (admin): `https://DCM-VIP:9091`
- From the Consumer Access Network (public): `https://CAN-VIP:9091`

Credentials:

- User name: administrator
- password: (set in the installation UI or the `cs-install-options.yml` file in the field `oo-admin-password`).

Identify Helion CloudSystem virtual appliances

To perform maintenance or troubleshooting on Helion CloudSystem virtual appliances, use the VM name and the appliance name. It is also important to understand how IP addresses are allocated to appliances on the various networks.

Deployer virtual appliance name and IP address allocation

To perform maintenance or troubleshooting on the Deployer virtual appliance, use the VM name and the appliance name.

Appliance name: Deployer

VM name: <cloud.host-prefix>-ccp-c0-m1-<network>

Table 3: Deployer IP address assignment

Appliance name	CONF assignment	DCM assignment	CLM assignment	CAN assignment	iSCSI assignment
Deployer	IP-01	IP-01	IP-02	N/A	N/A

Management appliance names and IP address allocation

To perform maintenance or troubleshooting on the Management virtual appliances, use the VM name and the appliance name.

- Appliance name: Management-m1
VM name: <cloud.host-prefix>-ccp-mgmt-m1-<network>
- Appliance name: Management-m2
VM name: <cloud.host-prefix>-ccp-mgmt-m2-<network>
- Appliance name: Management-m3
VM name: <cloud.host-prefix>-ccp-mgmt-m3-<network>

Table 4: Management appliance IP address assignment

Appliance name	CONF assignment	DCM assignment	CLM assignment	CAN assignment	iSCSI assignment
Management-m1	IP-02	IP-02	IP-03	IP-01	IP-01
Management-m2	IP-03	IP-03	IP-04	IP-02	IP-02
Management-m3	IP-04	IP-04	IP-05	IP-03	IP-03

! **IMPORTANT:**

When performing maintenance on a virtual machine hosting the Management appliances, make sure that **Management-m1** is always the last node shutdown and the first node restarted.

Monitoring appliance names and IP address allocation

To perform maintenance or troubleshooting on the Monitoring virtual appliances, use the VM name and the appliance name.

- Appliance name: `Monitoring-m1`
VM name: `<cloud.host-prefix>-ccp-mon-m1-<network>`
- Appliance name: `Monitoring-m2`
VM name: `<cloud.host-prefix>-ccp-mon-m2-<network>`
- Appliance name: `Monitoring-m3`
VM name: `<cloud.host-prefix>-ccp-mon-m3-<network>`

Table 5: Monitoring appliance IP address assignment

Appliance name	CONF assignment	DCM assignment	CLM assignment	CAN assignment	iSCSI assignment
Monitoring-m1	IP-05	IP-05	IP-06	N/A	N/A
Monitoring-m2	IP-06	IP-06	IP-07	N/A	N/A
Monitoring-m3	IP-07	IP-07	IP-08	N/A	N/A

! **IMPORTANT:**

When performing maintenance on a virtual machine hosting the Monitoring appliances, make sure that **Monitoring-m1** is always the last node shutdown and the first node restarted.

Database appliance names and IP address allocation

To perform maintenance or troubleshooting on the Database virtual appliances, use the VM name and the appliance name.

- Appliance name: `Database-m1`
VM name: `<cloud.host-prefix>-ccp-db-m1-<network>`
- Appliance name: `Database-m2`
VM name: `<cloud.host-prefix>-ccp-db-m2-<network>`
- Appliance name: `Database-m3`
VM name: `<cloud.host-prefix>-ccp-db-m3-<network>`

Table 6: Database appliance IP address assignment

Appliance name	CONF assignment	DCM assignment	CLM assignment	CAN assignment	iSCSI assignment
Database-m1	IP-08	IP-08	IP-09	N/A	N/A
Database-m2	IP-09	IP-09	IP-10	N/A	N/A
Database-m3	IP-10	IP-10	IP-11	N/A	N/A

! **IMPORTANT:**

When performing maintenance on a virtual machine hosting the Database appliances, make sure that **Database-m1** is always the last node shutdown and the first node restarted.

Foundation appliance names and IP address allocation

To perform maintenance or troubleshooting on the Foundation virtual appliances, use the VM name and the appliance name.

- Appliance name: CloudController-m1
VM name: <cloud.host-prefix>-ccp-cc-m1-<network>
- Appliance name: CloudController-m2
VM name: <cloud.host-prefix>-ccp-cc-m2-<network>
- Appliance name: CloudController-m3
VM name: <cloud.host-prefix>-ccp-cc-m3-<network>
- Appliance name: NetworkNode-m1
VM name: <cloud.host-prefix>-ccp-nn-m1-<network>
- Appliance name: NetworkNode-m2
VM name: <cloud.host-prefix>-ccp-nn-m2-<network>
- Appliance name: NetworkNode-m3
VM name: <cloud.host-prefix>-ccp-nn-m3-<network>

Table 7: Foundation appliance IP address assignment

Appliance name	CONF assignment	DCM assignment	CLM assignment	CAN assignment	iSCSI assignment
CloudController-m1	IP-11	IP-11	IP-12	N/A	IP-04
CloudController-m2	IP-12	IP-12	IP-13	N/A	IP-05
CloudController-m3	IP-13	IP-13	IP-14	N/A	IP-06

Table Continued

Appliance name	CONF assignment	DCM assignment	CLM assignment	CAN assignment	iSCSI assignment
NetworkNode-m1	IP-14	IP-14	IP-15	N/A	N/A
NetworkNode-m2	IP-15	IP-15	IP-16	N/A	N/A
NetworkNode-m3	IP-16	IP-16	IP-17	N/A	N/A

ⓘ IMPORTANT:

When performing maintenance on a virtual machine hosting the cloud controllers, make sure that **CloudController-m1** is always the last node shutdown and the first node restarted.

When performing maintenance on a virtual machine hosting the network nodes, make sure that **NetworkNode-m1** is always the last node shutdown and the first node restarted.

Enterprise appliance names and IP address allocation

To perform maintenance or troubleshooting on the Enterprise virtual appliances, use the VM name and the appliance name.

- Appliance name: Enterprise-m1
VM name: <cloud.host-prefix>-ccp-ea-m1-<network>
- Appliance name: Enterprise-m2
VM name: <cloud.host-prefix>-ccp-ea-m2-<network>
- Appliance name: Enterprise-m3
VM name: <cloud.host-prefix>-ccp-ea-m3-<network>

Table 8: Enterprise appliance IP address assignment

Appliance name	CONF assignment	DCM assignment	CLM assignment	CAN assignment	iSCSI assignment
Enterprise-m1	IP-17	IP-17	IP-17	IP-04	N/A
Enterprise-m2	IP18	IP18	IP-18	IP-05	N/A
Enterprise-m3	IP-19	IP-19	IP-19	IP-06	N/A

ⓘ IMPORTANT:

When performing maintenance on a virtual machine hosting the Enterprise appliances, make sure that **Enterprise-m1** is always the last node shutdown and the first node restarted.

VIP address assignment

VIP addresses are specified during deployment when configuring the load balancer settings.

Table 9: VIP address assignment

CONF assignment	DCM assignment	CLM assignment	CAN assignment	iSCSI assignment
N/A	IP-20	IP-01	IP-04	N/A

Update Helion CloudSystem

Update Helion CloudSystem to 10.0 Update 2

Prerequisites

- Helion CloudSystem 10.0 or 10.0 Update 1 is installed.
- You downloaded the Helion CloudSystem 10.0 Update 2 package from the **HPE Software Depot**.
- You verified that the Operations Console does not display any error messages for services or compute resources.

In the Operations Console, navigate to **Home > Alarm Dashboard**.

The update process applies all new changes in the update ISO. You can run the update process multiple times without changing the final result of the update process, beyond the initial update action.

You cannot roll back the update process. You can, however, edit configuration files and reapply them.

Procedure

1. Using stack credentials, SSH in to the Deployer virtual appliance.
2. Copy the update ISO from the update package to `/home/stack`.
3. Create a folder to mount the update ISO.

```
mkdir /media/cdrom
```

4. Mount the Helion CloudSystem update ISO.

```
sudo mount /home/stack/cs-update-10.0.2.iso /media/cdrom/
```

5. Run the included preparation script from the mounted ISO:

```
cd /media/cdrom/hos
./prepare-cs-init.bash
```

6. To update the Deployer virtual appliance, change directories and run the initialization script included in the update kit.

```
cd /home/stack
./cs-init.bash
```

Expected Result: (no error, no failed or unreachable)

PLAY RECAP

demo-ccp-c0-m1-clm	: ok=4	changed=1	unreachable=0	failed=0
demo-ccp-cc-m1-clm	: ok=4	changed=3	unreachable=0	failed=0
demo-ccp-db-m1-clm	: ok=4	changed=3	unreachable=0	failed=0
demo-ccp-ea-m1-clm	: ok=4	changed=3	unreachable=0	failed=0
demo-ccp-mgmt-m1-clm	: ok=4	changed=3	unreachable=0	failed=0
demo-ccp-mon-m1-clm	: ok=4	changed=3	unreachable=0	failed=0
demo-ccp-nn-m1-clm	: ok=4	changed=3	unreachable=0	failed=0
demo-ccp-nn-m2-clm	: ok=4	changed=3	unreachable=0	failed=0

7. Prepare the new Ansible tree for the update.
 - a. Resolve any merge conflicts that may have arisen. Re-do any changes made to service configuration jinja (.j2) templates (for example, `~/helion/my_cloud/config/cs-ace/cs-ace-py.j2`), then commit the changes.

```
cd ~/helion
git status
```

```
git add -A
git commit -m "Update to CS 10.0.2"
```

b. Run the config-processor.

```
cd ~/helion/hos/ansible/
ansible-playbook -i hosts/localhost config-processor-run.yml \
-e encrypt="" -e rekey=""
```

```
Expected Result: (no error, no failed or unreachable)
Total: ----- 21.48s
localhost      : ok=28  changed=19  unreachable=0  failed=0
```

c. Run the ready-deployment script:

```
ansible-playbook -i hosts/localhost ready-deployment.yml
```

```
Expected Result: (no error, no failed or unreachable)
Total: ----- 10.17s
localhost      : ok=23  changed=13  unreachable=0  failed=0
```

8. If Helion CloudSystem Enterprise is installed, run the following update playbook.

```
cd ~/scratch/ansible/next/hos/ansible
ansible-playbook -i hosts/verb_hosts cs-enterprise-pre-upgrade.yml
```

9. Run the remaining update playbook.

```
ansible-playbook -f 5 -i hosts/verb_hosts hlm-upgrade.yml 2>&1 \
| tee /home/stack/upgrade.log
```

10. Unmount and remove the update ISO.

```
sudo umount /media/cdrom/
rm -rf /home/stack/cs-update-10.0.2.iso
```

11. (Optional) If you are planning to add Helion CloudSystem Enterprise at a later time, copy the 10.0.2 Enterprise file to /home/stack/helion/images/.

12. (Optional) If you have ESXi compute resource and previously loaded the hlm-shell-vm.ova to vCenter, delete the template from the Deployer appliance and upload a new one. You will find the new hlm-shell-vm.ova in the Deployer virtual appliance after the update.

13. Reboot the Helion CloudSystem control plane, HPE Linux compute hosts, Nova proxy, and OVSvApp to load the new kernel.

14. Verify that the update completed successfully.

a. In the Operations Console:

- Navigate to the help menu and verify that version 10.0.2 is displayed.
- Navigate to the Alarms dashboard and make sure that no critical alarms are displayed.

b. In all Helion CloudSystem virtual appliances:

- Verify that running the csversion CLI displays 10.0.2.

Example

```
stack@cs10-ccp-c0-m1-clm:~$ csversion -b
```

```
HPE Helion CloudSystem Deployer Appliance v10.0.2
Powered by: HPE Helion OpenStack cchos-3.0.4 (build 01-1754)
```

Build information for installer:

Version: 10.0.2-RC2-50-20170208T161516

Coordinated Build: build.4 (qa-build.4-build-RC2-50)

Time of build: 2017-02-08T16:15:16Z

Tag/Branch: branch:hp/cs/10.0.2-release

c. If Helion CloudSystem Enterprise is installed:

- Verify that the OpenStack content, base utility, and OpenStack common utility content pack versions display as 16.09.
- Log on to all Enterprise virtual appliances and make sure that you see the file `/usr/local/hpe/csa/jboss-as/standalone/deployments/csa.war/WEB-INF/lib/apis-04.60.0001-SNAPSHOT.jar`.

15. (Optional) Update HPE CSA on the Enterprise appliances to 4.60 patch 2.

16. After the update completes successfully and before you create a support dump (if necessary), HPE recommends that you delete the following log files from the Deployer appliance to reduce the possibility that sensitive information is included in the support dump archive.

- `/var/log/csdeploy/*`
- `/var/log/cloudsystem-installer-ui/*`

17. Back up the control plane appliances.

HPE recommends that you back up the control plane after an update, because backups from previous versions of Helion CloudSystem should not be restored on appliances running a later version.

For backup procedures, see the *HPE Helion CloudSystem 10.0 Update 1 Administrator Guide* at **Enterprise Information Library**.

(Optional) Update HPE CSA to 4.60 patch 2

Prerequisites

- Helion CloudSystem 10.0 Update 2 is installed.
- You have a recent backup of the Enterprise appliance.

See the "Backup, restore, and recover Helion CloudSystem appliance data" chapter in the *HPE Helion CloudSystem 10.0 Update 1 Administrator Guide* at [Enterprise Information Library](#).

- You stopped new subscription creation and subscription modification.

Process overview

- Download and copy HPE CSA 4.6 patch 2 to each Enterprise appliance
- Stop services
- Install the patch on each Enterprise appliance
- Start services
- Clean up and perform final validation

Procedure

1. Download and copy HPE CSA 4.6 patch 2 to the Enterprise appliances.
 - a. Download Cloud Service Automation 4.60 patch 2 for Linux from <https://softwaresupport.hpe.com/km/KM02729766>. A valid HPE Passport account is required.
 - b. Unzip the downloaded file. The unzipped file is named HPE_CSA_Patch_04.60.0002.bin.
 - c. Log in to the Deployer appliance using stack credentials.
 - d. Using scp or winscp, copy the patch to the Deployer appliance.
 - e. From the Deployer appliance, copy the patch to each of the three Enterprise appliances (<cs>-ccp-ea-m1-clm, <cs>-ccp-ea-m2-clm, and <cs>-ccp-ea-m3-clm).

In the following example, replace <cs> with the cloud prefix you set during Helion CloudSystem deployment.

```
scp HPE_CSA_Patch_04.60.0002.bin stack@<cs>-ccp-ea-m1:/home/stack
scp HPE_CSA_Patch_04.60.0002.bin stack@<cs>-ccp-ea-m2:/home/stack
scp HPE_CSA_Patch_04.60.0002.bin stack@<cs>-ccp-ea-m3:/home/stack
```

- f. Delete the patch file from the Deployer appliance.
2. Stop services on the Enterprise appliance.
 - a. SSH to the Deployer appliance using stack credentials.
 - b. Execute the following "for loop" to stop csa and mpp services.

Do not use the csoperate command to stop all services on the Enterprise appliances, because the MySQL service must be running during the patch installation. The for loop may take some time to complete.

In the following example, replace <cs> with the cloud prefix you set during Helion CloudSystem deployment.

```
for ea in <cs>-ccp-ea-m1-clm <cs>-ccp-ea-m2-clm <cs>-ccp-ea-m3-clm; \
do ssh stack@$ea "sudo -u root service csa stop;
sudo -u root service mpp stop"
done
```

3. Install the patch on each Enterprise appliance.

- a. From the first Enterprise appliance, switch to the root user by entering `sudo -i`.
- b. Move the downloaded patch file to the `csauser` directory and set privileges on the file.

```
mv /home/stack/HPE_CSA_Patch_04.60.0002.bin /home/csauser/
cd /home/csauser/
chown csauser.csagrps HPE_CSA_Patch_04.60.0002.bin
chmod +x HPE_CSA_Patch_04.60.0002.bin
```

- c. Obtain the CSA database password.

```
cat /var/lib/heat-cfn/heat-cfn-init-data |grep csa_db_password
```

- d. Open the `/home/csauser/patch.properties` file.
- e. Add the following lines and the CSA database password to the `patch.properties` file.

```
#CSA DB user password
REPORTING_DB_SELECTED=no
CSA_REPORTING_DB_USER_NAME=csadbuser
CSA_DB_PASSWORD=<CSA database password obtained in step 3c>
```

- f. Install the patch as `csauser`.

```
cd /home/csauser/
sudo -u csauser ./HPE_CSA_Patch_04.60.0002.bin \
-i silent -f patch.properties
```

- g. Verify that the patch was installed successfully by checking for errors in the following log files:

- `/usr/local/hpe/csa/_CSA_4_60_2_installation/Logs/install.log`
- `/usr/local/hpe/csa/_CSA_4_60_2_installation/Logs/csa_InstallPatch.log`

- h. Delete the patch file from the Enterprise appliance.

```
rm /home/csauser/HPE_CSA_Patch_04.60.0002.bin
```

- i. Repeat steps 3a through 3h on the second and third Enterprise appliances.

4. Start services on the Enterprise appliance.

- a. SSH to the Deployer appliance using stack credentials.
- b. Execute the following "for loop" to start `csa` and `mpp` services.

In the following example, replace `<cs>` with the cloud prefix you set during Helion CloudSystem deployment.

```
for ea in <cs>-ccp-ea-m1-clm <cs>-ccp-ea-m2-clm <cs>-ccp-ea-m3-clm; \
do ssh stack@$ea "sudo -u root service csa start;
sudo -u root service mpp start"
done
```

5. Clean up and perform final validation.

- a. Clear the browser cache.
- b. Launch the HPE CSA Console, log in, and verify that the version is updated to 04.60.0002.35. Make sure that you can perform operations.
- c. Launch the HPE CSA Marketplace, log in, and make sure that you can perform operations.

Support and other resources

Websites

General websites

Hewlett Packard Enterprise Information Library

www.hpe.com/info/EIL

Single Point of Connectivity Knowledge (SPOCK) Storage compatibility matrix

www.hpe.com/storage/spock

Storage white papers and analyst reports

www.hpe.com/storage/whitepapers

Locate Hewlett Packard Enterprise authorized resellers

In the United States, see the U.S. partner and store locator website:

www.hpe.com/support/service_locator

In other locations, see the Contact Hewlett Packard Enterprise website:

www.hpe.com/contact

HPE Datacenter Care for the Cloud

www.hpe.com/services/dccloud

For additional websites, see [Support and other resources](#).

Software technical support and update service

Purchased Helion CloudSystem products include one year of 24 x 7 HPE Software Technical Support and Update Service, and HPE Helion CloudSystem products with HPE OneView or ordered with a Converged System include three-year 24x7 Software Technical Support and Update Service. This service provides access to HPE technical resources for assistance in resolving software implementation or operation problems.

This service also provides access to software updates and reference manuals, either in electronic form or on physical media as they are made available from HPE. Customers who purchase an electronic license to use are eligible for electronic updates only.

With this service, Helion CloudSystem customers will benefit from expedited problem resolution as well as proactive notification and delivery of software updates. For more information, see the following website:

www.hpe.com/services/insight

Registration for this service takes place following online redemption of the license certificate.

Registration for Software technical support and update service

If you received a license entitlement certificate, registration for this service will take place upon online redemption of the license certificate/key.

Using your Software technical support and update service

Once registered, you will receive a service contract in the mail containing the Customer Service phone number and your Service Agreement Identifier (SAID). You will need your SAID when calling for technical support. Using your SAID, you can go to the HPE Support Center at www.hpe.com/support/hpesc to view your contract online.

Consulting and installation services

HPE provides a complete range of consulting, installation, and support services to ensure successful deployment and operation of your HPE Helion CloudSystem infrastructure. For more information, go to www.hpe.com/us/en/services/consulting/cloud.html.

HPE Datacenter Care for Cloud

Part of the HPE Helion Professional Services portfolio, HPE Datacenter Care for HPE Cloud offers a flexible relationship, one that will support your Helion CloudSystem hardware and software operating environment. For more information, go to www.hpe.com/services/dccloud.

Management security

HPE is proactive in its approach to the quality and security of all its management software. Sign up for the latest downloadable security updates at www.hpe.com/support/e-updates.

Join the discussion

The HPE Support Center at www.hpe.com/support/hpesc is a community-based, user-supported forum for HPE customers to participate in discussions among the customer community about HPE products.

Related documents

The latest versions of Helion CloudSystem manuals and white papers, including HPE CSA and Operations Orchestration (OO), can be downloaded from the Enterprise Information Library at www.hpe.com/info/CloudSystem/docs.

HPE Helion CloudSystem documents

- *HPE Helion CloudSystem 10.0 Update 1 Administrator Guide*
- *HPE Helion CloudSystem 10.0 Command Line Interface Guide*
- *HPE Helion CloudSystem 10.0 Update 2 ESXi Installation Guide*
- *HPE Helion CloudSystem 10.0 Update 2 KVM Installation Guide*
- *HPE Helion CloudSystem 10.0 Network Planning Guide*
- *HPE Helion CloudSystem 10.0 Update 2 Release Notes*
- *HPE Helion CloudSystem 10.0 Update 1 Support Matrix*
- *HPE Helion CloudSystem 10.0 Update 1 Troubleshooting Guide*

Online help for the Helion CloudSystem Operations Console is available by clicking the help control button (a question mark symbol) in the Operations Console user interface.

HPE Helion OpenStack documents

The latest version of Helion OpenStack information is viewable from HPE Helion Documentation at docs.hpcloud.com, including information about the following products:

- HPE Helion OpenStack
- HPE Helion Eucalyptus

HPE Insight Management documents

The latest versions of HPE Matrix Operating Environment manuals, white papers, and the HPE Insight Management Support Matrix can be downloaded from the Enterprise Information Library at www.hpe.com/info/matrixoe/docs including the following documents:

- *HPE Matrix Operating Environment Release Notes*
- *HPE Insight Management Support Matrix*
- *HPE Matrix Operating Environment Infrastructure Orchestration User Guide*
- *Cloud bursting with HPE CloudSystem Matrix infrastructure orchestration*

Third-party documents

OpenStack

- OpenStack documentation for the Liberty release is located at docs.openstack.org/liberty

With few exceptions (such as installation information), OpenStack documents are always set to display the current release stream on the OpenStack document website.

- *Cloud Administrator Guide*
- *Virtual Machine Image Guide*
- *API Quick Start*
- *Admin User Guide*
- *End User Guide*
- *Command reference*
 - Keystone commands
 - Glance commands
 - Neutron commands

- Nova commands
- Cinder commands

VMware

- [VMware vSphere documents](#)

Microsoft

- [Microsoft Windows Server documents](#)
- [Microsoft Hyper-V documents](#)

Red Hat

- [Red Hat Enterprise Linux 7 documents](#)

HPE 3PAR StoreServ documents

The latest versions of 3PAR StoreServ Storage manuals are available from the Storage tab at the top of the [**Enterprise Information Library**](#).

3PAR StoreServ Storage manuals can also be downloaded from the Hewlett Packard Enterprise Support Center, including the following documents:

- *HPE 3PAR StoreServ Storage Concepts Guide*
- *HPE 3PAR StoreServ Storage Troubleshooting Guide*

HPE StoreVirtual VSA documents

The latest versions of VSA StoreVirtual Storage manuals are available from the Storage tab at the top of the [**Enterprise Information Library**](#).

VSA StoreVirtual Storage manuals can also be downloaded from the Hewlett Packard Enterprise Support Center, including the following document:

- *HPE StoreVirtual Storage User Guide*

HPE ProLiant servers documents

The HPE Integrated Lights-Out QuickSpecs contain support information and are available from the QuickSpecs website:

[**www.hp.com/go/qs**](http://www.hp.com/go/qs)

The following ProLiant servers documentation is available:

- ProLiant BL BladeSystem servers:

[**www.hpe.com/info/blades**](http://www.hpe.com/info/blades)

- ProLiant DL, ML, and SL servers:

[**www.hpe.com/info/rackservers**](http://www.hpe.com/info/rackservers)

Support and other resources

Accessing Hewlett Packard Enterprise Support

- For live assistance, go to the Contact Hewlett Packard Enterprise Worldwide website:
<http://www.hpe.com/assistance>
- To access documentation and support services, go to the Hewlett Packard Enterprise Support Center website:
<http://www.hpe.com/support/hpesc>

Information to collect

- Technical support registration number (if applicable)
- Product name, model or version, and serial number
- Operating system name and version
- Firmware version
- Error messages
- Product-specific reports and logs
- Add-on products or components
- Third-party products or components

Accessing updates

- Some software products provide a mechanism for accessing software updates through the product interface. Review your product documentation to identify the recommended software update method.
- To download product updates:

Hewlett Packard Enterprise Support Center

www.hpe.com/support/hpesc

Hewlett Packard Enterprise Support Center: Software downloads

www.hpe.com/support/downloads

Software Depot

www.hpe.com/support/softwaredepot

- To subscribe to eNewsletters and alerts:

www.hpe.com/support/e-updates

- To view and update your entitlements, and to link your contracts and warranties with your profile, go to the Hewlett Packard Enterprise Support Center **More Information on Access to Support Materials** page:

www.hpe.com/support/AccessToSupportMaterials

ⓘ **IMPORTANT:**

Access to some updates might require product entitlement when accessed through the Hewlett Packard Enterprise Support Center. You must have an HPE Passport set up with relevant entitlements.

Customer self repair

Hewlett Packard Enterprise customer self repair (CSR) programs allow you to repair your product. If a CSR part needs to be replaced, it will be shipped directly to you so that you can install it at your convenience. Some parts do not qualify for CSR. Your Hewlett Packard Enterprise authorized service provider will determine whether a repair can be accomplished by CSR.

For more information about CSR, contact your local service provider or go to the CSR website:

<http://www.hpe.com/support/selfrepair>

Remote support

Remote support is available with supported devices as part of your warranty or contractual support agreement. It provides intelligent event diagnosis, and automatic, secure submission of hardware event notifications to Hewlett Packard Enterprise, which will initiate a fast and accurate resolution based on your product's service level. Hewlett Packard Enterprise strongly recommends that you register your device for remote support.

If your product includes additional remote support details, use search to locate that information.

Remote support and Proactive Care information

HPE Get Connected	www.hpe.com/services/getconnected
HPE Proactive Care services	www.hpe.com/services/proactivecare
HPE Proactive Care service: Supported products list	www.hpe.com/services/proactivecaresupportedproducts
HPE Proactive Care advanced service: Supported products list	www.hpe.com/services/proactivecareadvancedsupportedproducts

Proactive Care customer information

Proactive Care central	www.hpe.com/services/proactivecarecentral
Proactive Care service activation	www.hpe.com/services/proactivecarecentralgetstarted

Warranty information

To view the warranty for your product, see the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products* document, available at the Hewlett Packard Enterprise Support Center:

www.hpe.com/support/Safety-Compliance-EnterpriseProducts

Additional warranty information

HPE ProLiant and x86 Servers and Options	www.hpe.com/support/ProLiantServers-Warranties
HPE Enterprise Servers	www.hpe.com/support/EnterpriseServers-Warranties
HPE Storage Products	www.hpe.com/support/Storage-Warranties
HPE Networking Products	www.hpe.com/support/Networking-Warranties

Regulatory information

To view the regulatory information for your product, view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products*, available at the Hewlett Packard Enterprise Support Center:

www.hpe.com/support/Safety-Compliance-EnterpriseProducts

Additional regulatory information

Hewlett Packard Enterprise is committed to providing our customers with information about the chemical substances in our products as needed to comply with legal requirements such as REACH (Regulation EC No 1907/2006 of the European Parliament and the Council). A chemical information report for this product can be found at:

www.hpe.com/info/reach

For Hewlett Packard Enterprise product environmental and safety information and compliance data, including RoHS and REACH, see:

www.hpe.com/info/ecodata

For Hewlett Packard Enterprise environmental information, including company programs, product recycling, and energy efficiency, see:

www.hpe.com/info/environment

Documentation feedback

Hewlett Packard Enterprise is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback (**docsfeedback@hpe.com**). When submitting your feedback, include the document title, part number, edition, and publication date located on the front cover of the document. For online help content, include the product name, product version, help edition, and publication date located on the legal notices page.

Appendix

Migrating standard vSwitches to distributed vSwitches

Prerequisites

- One ESXi host in the cluster must use a NIC other than vmnic0/vmk0 to connect to the vCenter Management Network.
- vCenter is deployed and configured to use a NIC other than vmnic0 to connect to the vCenter Management Network.

By default, vCenter uses standard vSwitches (VSS). If you prefer to use distributed vSwitches (DVS), then you must migrate the standard vSwitches to distributed vSwitches.

Some VMware features required by Helion CloudSystem may require additional VMware licensing. Refer to VMware documentation for licensing requirements.

For additional information on migrating vSwitches, refer to the [VMware Knowledge Base Solution](#).

Procedure

1. Create a new distributed vSwitch.
The distributed vSwitch must have a unique name.
 2. Modify the **dvPortGroup** that is created to allow it to use the vCenter Management Network VLAN ID.
Make sure that the VLAN ID is added to the dvPortGroup you create before migrating from standard to distributed virtual switches.
 3. Add the hosts.
 - a. Select the hosts that will be part of the cluster.
 - b. Select the vmnic that has the vmk0 (this is usually vSwitch 0).
 - c. Migrate vmk0 to the dvPortGroup that was configured in step 2.
-
- NOTE:**
Do not migrate the VM network.
- d. Review the changes and select **Finish**.
4. To ensure that the configuration uses the new dvPortGroup, manually edit the vCenter virtual machine.
 5. Delete the standard vSwitch.
 6. Using HPE OneView or Virtual Connect, update the NIC configuration to match the Helion CloudSystem topology.

Install Octavia load balancer

Octavia is a software load balancing solution that supports a highly available control plane and data plane. The KEEPALIVED load balancer is installed on the Management appliance by default. Installing the Octavia load balancer is optional.

ⓘ **IMPORTANT:**

This version of Octavia, which is also included with HPE Helion OpenStack Load Balancer as a Service (LBaaS) v2, is **not recommended for use in a production environment**. If you want to use LBaaS v2 in production, use an external hardware load balancer such as F5, A10, Radware, or Citrix.

For more information on Octavia, see

- [HPE Helion OpenStack 3.0 documentation](#) and navigate to **HPE Helion OpenStack 3.0 > Operations Guide > Managing Networking > Configuring the Network Service > Load Balancer: Octavia Driver Administration**
- [OpenStack Octavia](#)

Installing the Octavia load balancer

Prerequisites

- You have a VLAN ID that can be assigned to a Provider Network. The network must be outside of the VLAN Data Trunk ranges defined in the `cs-install-options.yml` installation configuration file. The network must also be routed to the Data Center Management Network (admin).
- The environment has access to a KVM baremetal host (no ESXi compute clusters are currently supported).

These instructions must be run on the Deployer virtual appliance BEFORE running the `csdeploy` command that launches the remaining Helion CloudSystem virtual appliances.

Procedure

1. Log on to the Deployer virtual appliance.
2. Modify `neutron_config.yml`.
 - a. Open `/home/stack/helion/my_cloud/definition/data/neutron/neutron_config.yml`.
 - b. Uncomment all lines from `'configuration-data: '` to the end of the file.
 - c. Update the following values and define the Provider Network that will be used.

Refer to the [Helion OpenStack 3.0 documentation](#) and expand the help navigation **HPE Helion OpenStack 3.0 > Planning > Helion Lifecycle Manager Overview > Input Model > Configuration Objects** for more information on setting the `neutron_config.yml`.

Table 10: Fields in the `neutron_config.yml` file

Data section	Expected value
<pre>neutron_provider_networks: - name:</pre>	OCTAVIA-MGM-NET
<pre>provider: - network_type: vlan physical_network: segmentation_id: cidr: no_gateway: enable_dhcp: allocation_pools: - start: end: host_routes: - destination: nexthop:</pre>	<ul style="list-style-type: none"> • The <code>network_type</code> is vlan. • The <code>physical_network</code> is the lowercase name of a VLAN Data Trunk (vul) defined in the <code>cs-install-options.yml</code> file. • The <code>segmentation_id</code> is the VLAN ID assigned to the provider network. • CIDR is the CIDR assigned to the provider network. • The <code>allocation_pools</code> define the starting IP range and the ending IP range for the network. • The <code>host_routes</code> destination is the CIDR of the DCM network defined in <code>cs-install-options.yml</code>. • The <code>host_route</code> <code>nexthop</code> is the provider network gateway that is routed to the DCM network.

Example:

```
product:
  version: 2
configuration-data:
  - name: NEUTRON-CONFIG-CP1
    services:
      - neutron
    data:
      neutron_provider_networks:
        - name: OCTAVIA-MGMT-NET
          provider:
            - network_type: vlan
              physical_network: vlan
                segmentation_id: 969
            cidr: 10.246.165.128/25
            no_gateway: true
            enable_dhcp: true
            allocation_pools:
              - start: 10.246.165.135
                end: 10.246.165.200
            host_routes:
              - destination: 10.246.163.128/25
                nexthop: 10.246.165.129
```

d. Save and then close the file.

3. Modify the `control_plane_base.yml` file.

a. Open the file `/home/stack/helion/my_cloud/add_ons/clusters/control_plane_base.yml`.

b. Uncomment the following `configuration-data` lines:

Example:

```
configuration-data:
  - OCTAVIA-CONFIG-CP1
  - NEUTRON-CONFIG-CP1
```

c. Save and then close the file.

4. Modify the `control_plane_foundation.yml` file.

a. Open the file `/home/stack/helion/my_cloud/add_ons/clusters/control_plane_foundation.yml`

b. Uncomment everything from the two Octavia services to the service-components of the `CONTROLLER-ROLE`.

Example:

```
...
server-role: CONTROLLER-ROLE
service-components:
  ...
  - octavia-api
  - octavia-health-manager
```

c. Save and then close the file.

5. Run `csdeploy`.

See [Deploy Helion CloudSystem virtual appliances](#).

6. Activate the RHEL KVM compute host.

See **Activating an HPE Linux or RHEL KVM compute host** in the *HPE Helion CloudSystem 10.0 Update 1 Administrator Guide* in the **Enterprise Information Library**.

7. Extract the Octavia image from the **HPE Helion CloudSystem Tools** package in the release kit and copy it to the `/var/tmp` folder on the Deployer virtual appliance.
8. Run the playbook to load the amphora image:

```
cd ~/scratch/ansible/next/hos/ansible/ansible-playbook -i hosts/verb_hosts -e service_package=/var/tmp/octavia-amphora-haproxy-guest-image.tgz / service-guest-image.yml | tee /home/stack/amphora.log
```
9. Create the Octavia load balancer.
 - a. Log in to the Deployer virtual appliance.
 - b. Create a network and sub-net that will be used by Octavia to provision the amphora images.
 - c. Create an Octavia load balancer.

```
neutron lbaas-loadbalancer-create --name your-lb-name --provider octavia your-subnet-name
```

10. Verify that the load balancer was successfully created.

```
neutron lbaas-loadbalancer-list
```

Its state should change to **ACTIVE**.

Add Enterprise to a Foundation-only deployment

If you decide to add Enterprise to Helion CloudSystem after the virtual appliances are deployed, make sure that you prepared the Deployer virtual appliance for an Enterprise installation.

See **Preparing Helion CloudSystem Enterprise**.

Adding Enterprise to a Foundation-only deployment

Prerequisites

- A MySQL JDBC driver is downloaded and the image is copied to the Deployer virtual appliance.
- The **HPE-Helion-CloudSystem-enterprise-10.0.2.ova** is added to the Deployer virtual appliance images directory.

See **Preparing Helion CloudSystem Enterprise**.

Procedure

1. Using `stack` credentials defined during deployment, log in to the Deployer virtual appliance.
2. Open the `cs-install-options.yml` file:
 - a. If you deployed Helion CloudSystem from the CLI, the file will be in the path `/etc/csdeploy/cs-install-options.yml`.
 - b. If you deployed Helion CloudSystem from the installation UI, the file will be in the path `/opt/stack/service/cloudsystem-installer-ui/venv/model/cs_install_options.yml`.
3. Update the **ENTERPRISE-ROLE** vCPU and memory values in the **appliance-sizings** section of the configuration file.

- a. If you deployed Helion CloudSystem from the CLI, verify the values for the ENTERPRISE-ROLE and update them if needed.
- b. If you deployed Helion CloudSystem from the UI, add the ENTERPRISE-ROLE as the last entry in the appliance-sizings section and update the values.

```
#ENTERPRISE-ROLE:
#  vcpus: 8
#  memory: 32
```

4. In the **add-ons** section, update the `oo-admin-password`:
 - a. If you deployed Helion CloudSystem from the CLI, add the password to the `oo-admin-password` field.
 - b. If you deployed Helion CloudSystem from the UI, add the Enterprise section with the `oo-admin-password`.

```
# enterprise section required with --add-enterprise
enterprise:
  oo-admin-password: '<password>'
```

5. Save the file.

6. Run `csdeploy`:

- a. If you deployed Helion CloudSystem from the CLI, run:


```
csdeploy start --config /etc/csdeploy/examples/cs-install-options.yml --add-enterprise 2>&1 | tee /home/stack/csdeploy-enterprise.log
```
- b. If you deployed Helion CloudSystem from the installation UI, run:


```
csdeploy start --config /opt/stack/service/cloudsystem-installer-ui/venv/model/cs_install_options.yml --add-enterprise 2>&1 | tee /home/stack/csdeploy-enterprise.log
```

7. Verify that a message appears indicating a successful deployment.

Update service endpoints with the FQDN on Helion CloudSystem Enterprise appliances

If you configured the load balancers with FQDNs instead of VIP addresses, you must update the public endpoints on all Enterprise appliances. If you do not make this update, you will not be able to access the portals and consoles hosted by the Enterprise appliances.

Replace VIP addresses with FQDNs on Enterprise appliances

Procedure

1. Log in to the Deployer virtual appliance.
2. From the Deployer virtual appliance, SSH to the first Enterprise appliance.
3. Using root credentials, open the file `/usr/local/hpe/csa/jboss-as/welcome-content/index.html` for editing.
4. Replace the public VIP with the public FQDN:
`https://servername-public.domain.local:8444/csa/`
5. Using root credentials, open the file `/usr/local/hpe/csa/jboss-as/standalone/deployments/mpp.war/index.html` for editing.
6. Replace the public VIP with the public FQDN:
`https://servername-public.domain.local:8089/`
7. Using root credentials, open the file `/usr/local/hpe/csa/portal/node_modules/mpp-server/conf/mpp.json` for editing.
8. Update the following fields:

- a. Search for "idmProvider".
 - b. In "url", replace the public VIP with the public FQDN.
 - c. In "returnURL", Replace the public VIP with the public FQDN.
- Example:

```
"idmProvider": {
  "url": "https://servername-public.domain.local:8444",
  "returnUrl": "https://servername-public.domain.local:8089",
```

9. Using root credentials, open the file /usr/local/hpe/csa/jboss-as/standalone/deployments/csa.war/dashboard/config.json for editing.

10. Update the following fields:

- a. Search for "marketplace".
 - b. In "url", replace the public VIP with the public FQDN.
- Example:

```
"url": "https://servername-public.domain.local:8089"
```

11. Using root credentials, open the file /usr/local/hpe/csa/jboss-as/standalone/deployments/csa.war/WEB-INF/classes/csa.properties for editing.

- a. Search for "OOS_URL".
 - b. In "OOS_URL", replace the public VIP with the public FQDN.
- Example:

```
"OOS_URL": "https://servername-public.domain.local:9091"
```

12. If this is an HA configuration, repeat the preceding steps on each Enterprise appliance in the cluster.

13. Restart HPE CSA and Marketplace portal services on all Enterprise appliances:

- a. From the Deployer virtual appliance, stop services:

```
csoperate service stop enterprise-csa cs10cloud-ccp-enterprise
```

- b. Start services:

```
csoperate service start enterprise-csa cs10cloud-ccp-enterprise
```

Disk requirements

When installing Helion CloudSystem Foundation and Enterprise, or just Helion CloudSystem Foundation, use the following disk size guidelines.

If you are only installing Helion CloudSystem Enterprise, see [Disk requirements \(Enterprise-only\)](#).

Extra-small

The following table shows the disk size changes that were made to support the extra-small configuration requirements.

Virtual appliance name	filename	Disk size	Disk allocation
Deployer	N/A	Disk 0: 205 GB	Disk 0: root, installation files
Cloud controller	specs_controller.j2	Disk 0: 55 GB Disk 1: 25 GB	Disk 0: bare hLinux Disk 1: root, logs

Table Continued

Virtual appliance name	filename	Disk size	Disk allocation
Database-Rabbit	specs_db_rabbit.j2	Disk 0: 55 GB Disk 1: 25 GB Disk 2: 50 GB	Disk 0: bare hLinux Disk 1: root, logs Disk 2: Rabbit, MySQL
Enterprise	specs_enterprise.j2	Disk 0: 211 GB Disk 1: 5 GB	Disk 0: bare hLinux Disk 1: root
Management	specs_management.j2	Disk 0: 55 GB Disk 1: 25 GB	Disk 0: bare hLinux Disk 1: root, logs
Monitoring	specs_monitoring.j2	Disk 0: 55 GB Disk 1: 25 GB Disk 2: 50 GB Disk 3: 75 GB Disk 4: 50 GB	Disk 0: bare hLinux Disk 1: root, Crash, Zookeeper, logs Disk 2: Kafka Disk 3: Elastic Search Disk 4: VerticaDB
Network Node	specs_network.j2	Disk 0: 55 GB	Disk 0: bare hLinux

Small

The following table shows the disk size changes that were made to support the small configuration requirements.

Virtual appliance name	filename	Disk size	Disk allocation
Deployer	N/A	Disk 0: 205 GB	Disk 0: root, installation files
Cloud controller	specs_controller.j2	Disk 0: 55 GB Disk 1: 50 GB	Disk 0: bare hLinux Disk 1: root, logs

Table Continued

Virtual appliance name	filename	Disk size	Disk allocation
Database-Rabbit	specs_db_rabbit.j2	Disk 0: 55 GB Disk 1: 50 GB Disk 2: 100 GB	Disk 0: bare hLinux Disk 1: root, logs Disk 2: Rabbit, MySQL
Enterprise	specs_enterprise.j2	Disk 0: 211 GB Disk 1: 50 GB	Disk 0: bare hLinux Disk 1: root
Management	specs_management.j2	Disk 0: 55 GB Disk 1: 50 GB	Disk 0: bare hLinux Disk 1: root, logs
Monitoring	specs_monitoring.j2	Disk 0: 55 GB Disk 1: 100 GB Disk 2: 100 GB Disk 3: 150 GB Disk 4: 100 GB	Disk 0: bare hLinux Disk 1: root, Crash, Zookeeper, logs Disk 2: Kafka Disk 3: Elastic Search Disk 4: VerticaDB
Network Node	specs_network.j2	Disk 0: 55 GB	Disk 0: bare hLinux

Medium

The following table shows the disk size changes that were made to support the medium configuration requirements.

Virtual appliance name	filename	Disk size	Disk allocation
Deployer	N/A	Disk 0: 205 GB	Disk 0: root, installation files
Cloud controller	specs_controller.j2	Disk 0: 55 GB Disk 1: 100 GB	Disk 0: bare hLinux Disk 1: root, logs

Table Continued

Virtual appliance name	filename	Disk size	Disk allocation
Database-Rabbit	specs_db_rabbit.j2	Disk 0: 55 GB Disk 1: 100 GB Disk 2: 200 GB	Disk 0: bare hLinux Disk 1: root, logs Disk 2: Rabbit, MySQL
Enterprise	specs_enterprise.j2	Disk 0: 211 GB Disk 1: 200 GB	Disk 0: bare hLinux Disk 1: root
Management	specs_management.j2	Disk 0: 55 GB Disk 1: 100 GB	Disk 0: bare hLinux Disk 1: root, logs
Monitoring	specs_monitoring.j2	Disk 0: 55 GB Disk 1: 150 GB Disk 2: 150 GB Disk 3: 200 GB Disk 4: 150 GB	Disk 0: bare hLinux Disk 1: root, Crash, Zookeeper, logs Disk 2: Kafka Disk 3: Elastic Search Disk 4: VerticaDB
Network Node	specs_network.j2	Disk 0: 55 GB	Disk 0: bare hLinux

Large

The following table shows the disk size changes that were made to support the large configuration requirements.

Virtual appliance name	filename	Disk size	Disk allocation
Deployer	N/A	Disk 0: 205 GB	Disk 0: root, installation files
Cloud controller	specs_controller.j2	Disk 0: 55 GB Disk 1: 150 GB	Disk 0: bare hLinux Disk 1: root, logs

Table Continued

Virtual appliance name	filename	Disk size	Disk allocation
Database-Rabbit	specs_db_rabbit.j2	Disk 0: 55 GB Disk 1: 200 GB Disk 2: 400 GB	Disk 0: bare hLinux Disk 1: root, logs Disk 2: Rabbit, MySQL
Enterprise	specs_enterprise.j2	Disk 0: 211 GB Disk 1: 824 GB	Disk 0: bare hLinux Disk 1: root
Management	specs_management.j2	Disk 0: 55 GB Disk 1: 150 GB	Disk 0: bare hLinux Disk 1: root, logs
Monitoring	specs_monitoring.j2	Disk 0: 55 GB Disk 1: 200 GB Disk 2: 300 GB Disk 3: 300 GB Disk 4: 300 GB	Disk 0: bare hLinux Disk 1: root, Crash, Zookeeper, logs Disk 2: Kafka Disk 3: Elastic Search Disk 4: VerticaDB
Network Node	specs_network.j2	Disk 0: 55 GB	Disk 0: bare hLinux

Disk requirements (Enterprise-only)

When installing Helion CloudSystem Enterprise, use the following disk size guidelines.

Extra-small

The following table shows the disk size changes that were made to support the extra-small configuration requirements.

Virtual appliance name	filename	Disk size	Disk allocation
Deployer	N/A	Disk 0: 205 GB	Disk 0: root, installation files
Database-Rabbit	specs_db_rabbit.j2	Disk 0: 55 GB Disk 1: 25 GB Disk 2: 50 GB	Disk 0: bare hLinux Disk 1: root, logs Disk 2: Rabbit, MySQL
Enterprise	specs_enterprise.j2	Disk 0: 211 GB Disk 1: 5 GB	Disk 0: bare hLinux Disk 1: root
Management	specs_management.j2	Disk 0: 55 GB Disk 1: 25 GB	Disk 0: bare hLinux Disk 1: root, logs
Monitoring	specs_monitoring.j2	Disk 0: 55 GB Disk 1: 25 GB Disk 2: 50 GB Disk 3: 75 GB Disk 4: 50 GB	Disk 0: bare hLinux Disk 1: root, Crash, Zookeeper, logs Disk 2: Kafka Disk 3: Elastic Search Disk 4: VerticaDB

Small

The following table shows the disk size changes that were made to support the small configuration requirements.

Virtual appliance name	filename	Disk size	Disk allocation
Deployer	N/A	Disk 0: 205 GB	Disk 0: root, installation files
Database-Rabbit	specs_db_rabbit.j2	Disk 0: 55 GB Disk 1: 50 GB Disk 2: 100 GB	Disk 0: bare hLinux Disk 1: root, logs Disk 2: Rabbit, MySQL

Table Continued

Virtual appliance name	filename	Disk size	Disk allocation
Enterprise	specs_enterprise.j2	Disk 0: 211 GB Disk 1: 50 GB	Disk 0: bare hLinux Disk 1: root
Management	specs_management.j2	Disk 0: 55 GB Disk 1: 50 GB	Disk 0: bare hLinux Disk 1: root, logs
Monitoring	specs_monitoring.j2	Disk 0: 55 GB Disk 1: 50 GB Disk 2: 50 GB Disk 3: 75 GB Disk 4: 50 GB	Disk 0: bare hLinux Disk 1: root, Crash, Zookeeper, logs Disk 2: Kafka Disk 3: Elastic Search Disk 4: VerticaDB

Medium

The following table shows the disk size changes that were made to support the medium configuration requirements.

Virtual appliance name	filename	Disk size	Disk allocation
Deployer	N/A	Disk 0: 205 GB	Disk 0: root, installation files
Database-Rabbit	specs_db_rabbit.j2	Disk 0: 55 GB Disk 1: 50 GB Disk 2: 100 GB	Disk 0: bare hLinux Disk 1: root, logs Disk 2: Rabbit, MySQL
Enterprise	specs_enterprise.j2	Disk 0: 211 GB Disk 1: 200 GB	Disk 0: bare hLinux Disk 1: root

Table Continued

Virtual appliance name	filename	Disk size	Disk allocation
Management	specs_management.j2	Disk 0: 55 GB Disk 1: 50 GB	Disk 0: bare hLinux Disk 1: root, logs
Monitoring	specs_monitoring.j2	Disk 0: 55 GB Disk 1: 50 GB Disk 2: 50 GB Disk 3: 75 GB Disk 4: 50 GB	Disk 0: bare hLinux Disk 1: root, Crash, Zookeeper, logs Disk 2: Kafka Disk 3: Elastic Search Disk 4: VerticaDB

Large

The following table shows the disk size changes that were made to support the large configuration requirements.

Virtual appliance name	filename	Disk size	Disk allocation
Deployer	N/A	Disk 0: 205 GB	Disk 0: root, installation files
Database-Rabbit	specs_db_rabbit.j2	Disk 0: 55 GB Disk 1: 50 GB Disk 2: 100 GB	Disk 0: bare hLinux Disk 1: root, logs Disk 2: Rabbit, MySQL
Enterprise	specs_enterprise.j2	Disk 0: 211 GB Disk 1: 824 GB	Disk 0: bare hLinux Disk 1: root

Table Continued

Virtual appliance name	filename	Disk size	Disk allocation
Management	specs_management.j2	Disk 0: 55 GB Disk 1: 50 GB	Disk 0: bare hLinux Disk 1: root, logs
Monitoring	specs_monitoring.j2	Disk 0: 55 GB Disk 1: 50 GB Disk 2: 50 GB Disk 3: 75 GB Disk 4: 50 GB	Disk 0: bare hLinux Disk 1: root, Crash, Zookeeper, logs Disk 2: Kafka Disk 3: Elastic Search Disk 4: VerticaDB