Cloudian HyperStore Installation Guide

Version 7.5.2

This page left intentionally blank.

This page left intentionally blank.

Contents

Chapter 1. HyperStore Installation Introduction

This documentation describes how to do a **fresh installation** of Cloudian HyperStore 7.5.2.

> **Note** For instructions on **upgrading** to 7.5.2 from an older HyperStore version see "Upgrading Your HyperStore Software Version" in the ***Cloudian HyperStore Administrator's Guide***.

To do a fresh installation you need the **HyperStore product package**. To obtain the HyperStore 7.5.2 package:

- If you have an active Support contract you can download the package from the Cloudian Support Portal (https://cloudian-support.force.com).
- If you are doing an evaluation and do not have a Support contract you can obtain the package from your Cloudian sales representative.

To install and run HyperStore software you need a **HyperStore license file** — either an evaluation license or a production license. If you do not have a license file you can obtain one from your Cloudian sales representative or by registering for a free trial on the Cloudian website.

This page left intentionally blank

Chapter 2. Preparing Your Environment

Before installing HyperStore, Cloudian recommends that you prepare these aspects of your environment:

- **"DNS Set-Up"** (page 10)
- **"Load Balancing"** (page 13)

2.1. DNS Set-Up

Subjects covered in this section:

For your HyperStore system to be accessible to external clients, you must configure your DNS name servers with entries for the HyperStore service endpoints. **Cloudian recommends that you complete your DNS configuration prior to installing the HyperStore system.** This section describes the required DNS entries.

> **Note** If you are doing just a small evaluation and do not require that external clients be able to access any of the HyperStore services, you have the option of using the lightweight domain resolution utility *dnsmasq* which comes bundled with HyperStore -- rather than configuring your DNS environment to support HyperStore service endpoints. If you're going to use *dnsmasq* you can skip ahead to **"Host Hardware and OS Requirements"** (page 15). During installation of HyperStore software you can use the *configure-dnsmasq* option if you want to use *dnsmasq* for domain resolution. Details are in the software installation procedure.

2.1.1. HyperStore Service Endpoints

HyperStore includes a variety of services each of which is accessible to clients by way of a web service endpoint. On your name servers you will need to configure a DNS entry for each of these service endpoints.

By default the HyperStore system uses a standard format for each service endpoint, building on two values that are specific to your environment:

- Your organization's domain (for example *enterprise.com*)
- The name or names of your HyperStore service region or regions (for example *boston* for a single-region system, or *boston* and *chicago* for a multi-region system). Only lower case alphanumeric characters and dashes are allowed in region names.

During HyperStore installation you will supply your domain and your service region names, and the interactive installer will show you the default service endpoints derived from the domain and region names. During installation you can accept the default endpoints or specify custom endpoints instead. The table that follows below is based on the default endpoint formats.

> **Note**
> * Including the string "s3" in your domain or in your region name(s) is not recommended. By default HyperStore generates S3 service endpoints by prepending an "*s3-*" prefix to your *<region-name>.<domain>* combination. If you include "s3" within either your domain or your region name, this will result in two instances of "s3" in the system-generated S3 service endpoints, and this may cause S3 service requests to fail for some S3 clients.
> * If you specify custom endpoints during installation, do not use IP addresses in your endpoints.
> * HyperStore by default derives the S3 service endpoint(s) as *s3-<regionname>.<domain>*. However

HyperStore also supports the format *s3.<regionname>.<domain>* (with a dot after the "s3" rather than a dash) if you specify custom endpoints with this format during installation.

The table below shows the default format of each service endpoint. The examples show the service endpoints that the system would automatically generate if the domain is *enterprise.com* and the region name is *boston*.

Service Endpoint	Default Format and Example	Description
S3 service endpoint (**one per service region**)	*s3-<regionname>.<domain>* *s3-boston.enterprise.com*	This is the service endpoint to which S3 client applications will submit requests. If you are installing a HyperStore system across multiple service regions, each region will have its own S3 service endpoint, and therefore you must create a DNS entry for each of those region-specific endpoints — for example *s3-boston.enterprise.com* and *s3-chicago.enterprise.com*.
S3 service endpoint wildcard (**one per service region**)	**.s3-<regionname>.<domain>* **.s3-boston.enterprise.com*	This S3 service endpoint wildcard entry is necessary to resolve virtual-hosted-style S3 requests, wherein the bucket name is specified as a sub-domain -- for example *bucket1.s3-boston.enterprise.com* and *bucket2.s3-boston.enterprise.com* and so on.
S3 static website endpoint (**one per service region**)	*s3-website-<regionname>.<domain>* *s3-website-boston.enterprise.com*	This S3 service endpoint is used for buckets configured as static websites. **Note** Your S3 static website endpoint cannot be the same as your S3 service endpoint. They must be different, or else the website endpoint will not work properly.
S3 static website endpoint wildcard (**one per service region**)	**.s3-website-<regionname>.<domain>* **.s3-website-boston.enterprise.com*	This S3 static website endpoint wildcard entry is necessary to resolve virtual-hosted-style S3 requests, wherein the bucket name is specified as a sub-domain, for buckets configured as static websites.
Admin Service endpoint (**one per entire system**)	*s3-admin.<domain>* *s3-admin.enterprise.com*	This is the service endpoint for HyperStore's Admin API. The Cloudian Management Console accesses this API, and you can also access this API directly with a third party client (such as a command line tool like *cURL*).
IAM Service endpoint (**one per entire**	*iam.<domain>* *iam.enterprise.com*	This is the service endpoint for accessing HyperStore's implementation of the Identity and Access Management API.

Service Endpoint	Default Format and Example	Description
system)		
STS Service endpoint (**one per entire system**)	*sts.<domain>* *sts.enterprise.com*	This is the service endpoint for accessing HyperStore's implementation of the Security Token Service API. **Note** Resolve the STS endpoint to the same address as the IAM endpoint, or use CNAME to map the STS endpoint to the IAM endpoint.
SQS Service endpoint (**one per entire system**)	*s3-sqs.<domain>* *s3-sqs.enterprise.com*	This is the service endpoint for accessing HyperStore's implementation of the Simple Queue Service (SQS) API. **Note** The SQS Service is disabled by default. For information about enabling this service, see the SQS section of the ***Cloudian HyperStore AWS APIs Support Reference***.
Cloudian Management Console (CMC) endpoint (**one per entire system**)	*cmc.<domain>* *cmc.enterprise.com*	The CMC is HyperStore's web-based console for performing system administrative tasks. The CMC also supports actions such as creating storage buckets or uploading objects into buckets.

2.1.2. Configuring Resolution of Service Endpoints

IMPORTANT ! Cloudian Best Practices suggest that a highly available load balancer be used in production environments where consistent performance behavior is desirable. For environments where a load balancer is unavailable, other options are possible. Please consult with your Cloudian Sales Engineer for alternatives.

For a production environment, in your DNS configuration each HyperStore service endpoint should resolve to the virtual IP address(es) of two or more load balancers that are configured for high availability. For more detail see **"Load Balancing"** (page 13).

2.1.3. Using Customized Service Endpoints

If you do not want to use the default service endpoint formats, the HyperStore system allows you to specify custom endpoint values during the installation process. If you intend to create custom endpoints, configure DNS entries to resolve the custom endpoint values that you intend to use, rather than the default-formatted endpoint values shown in the **"HyperStore Service Endpoints"** (page 10) table. Make a note of the custom endpoints for which you configure DNS entries, so that later you can correctly specify those custom endpoints when you perform the HyperStore installation.

If you want to use a **custom S3 endpoint** that does not include a region string, the installer allows you to do so. Note however that if your S3 endpoints lack region strings the system will not be able to support the region name validation aspect of AWS Signature Version 4 authentication for S3 requests (but requests can still succeed without the validation).

If you want to use **multiple S3 endpoints per service region** -- for example, having different S3 endpoints resolve to different data centers within one service region -- the installer allows you to do this. For this approach, the recommended syntax is *s3-<regionname>.<dcname>.<domain>* -- for example *s3-boston.dc1.enterprise.com* and *s3-boston.dc2.enterprise.com*.

Note If you want to change HyperStore service endpoints after the system has already been installed, you can do so as described in the "Changing S3, Admin, CMC, or IAM Service Endpoints" section of the ***Cloudian HyperStore Administrator's Guide***. If you change any endpoints, be sure to update your DNS configuration.

2.2. Load Balancing

IMPORTANT ! Cloudian recommends that a highly available load balancer be used in production environments where consistent performance behavior is desirable. For environments where a load balancer is unavailable, other options are possible. Please consult with your Cloudian Sales Engineer for alternatives. The discussion below assumes that you are using a load balancer.

HyperStore uses a peer-to-peer architecture in which each node in the cluster can service requests to the S3, Admin, CMC, IAM, STS, and SQS service endpoints. In a production environment you should use load balancers to distribute S3, Admin, CMC, IAM, STS, and SQS service endpoint requests evenly across all the nodes in your cluster. In your DNS configuration the S3, Admin, CMC, IAM, STS, and SQS service endpoints should resolve to the virtual IP address(es) of your load balancers; and the load balancers should in turn distribute request traffic across all your nodes. Cloudian recommends that you **set up your load balancers prior to installing the HyperStore system**.

For high availability it is preferable to use two or more load balancers configured for failover between them (as versus having just one load balancer which would then constitute a single point of failure). The load balancers could be commercial products or you can use open source technologies such as **HAProxy** (load balancer software for TCP/HTTP applications) and **Keepalived** (for failover between two or more load balancer nodes). If you use software-defined solutions such as these open source products, for best performance you should install them on dedicated load balancing nodes -- not on any of your HyperStore nodes.

For a **single-region HyperStore system**, for each service configure the load balancers to distribute request traffic across all the nodes in the system.

For a **multi-region HyperStore system**:

- Configure each region's S3 service endpoint to resolve to load balancers in that region, which distribute traffic across all the nodes within that region.
- Configure the Admin, IAM, STS, SQS, and CMC service endpoints to resolve to load balancers in the **default service region**, which distribute traffic to all the nodes in the default service region. (You will specify a default service region during the HyperStore installation process. For example, you might have service regions *boston* and *chicago*, and during installation you can specify that *boston* is the default service region.)

For detailed guidance on load balancing set-up, request a copy of the ***HyperStore Load Balancing Best Practice Guide*** from your Cloudian Sales Engineering representative.

Note The HyperStore S3 Service supports **PROXY Protocol** for incoming connections from a load balancer. This is disabled by default, but after HyperStore installation is complete you can enable it by configuration if you wish. For more information search on "s3_proxy_protocol_enabled" in the ***Cloudian HyperStore Administrator's Guide***.

Note For information about how to perform health checks of HyperStore's HTTP(S) services such as the S3 Service and the CMC, see the "Checking HTTP(S) Responsiveness" section in the ***Cloudian HyperStore Administrator's Guide***.

Chapter 3. Preparing Your Nodes

This section covers these topics to help you select and prepare your HyperStore host machines:

- **"Host Hardware and OS Requirements"** (page 15)
- **"Preparing Your Nodes For HyperStore Installation"** (page 18)

3.1. Host Hardware and OS Requirements

Subjects covered in this section:

- ***"Hardware Requirements"*** *(page 15)*
- ***"Operating System Requirements"*** *(page 16)*
- ***"Note: Automatic Exclusions to OS Package Updates"*** *(page 18)*

3.1.1. Hardware Requirements

The table below shows the recommended and minimum specifications for individual host machines in a HyperStore system. In general your hosts should either by HyperStore Appliances or hosts with specifications comparable to HyperStore Appliances.

	HSA-1600 Appliance		HSA-4200 Appliance (Per Node)		Recommended VM Spec		Minimum VM Spec (Test/Dev/POC)	
Component	**Description**	**Qty**	**Description**	**Qty**	**Description**	**Qty**	**Description**	**Qty**
Processors	Intel Xeon 5218R	2	Intel Xeon 5218R	2	20 core CPU	2	8 core CPU	2
RAM	32GB (128GB)	4	32GB (256GB)	8	128GB	1	128GB	1
Flash Metadata Tier	NVMe Storage 1920GB	2	NVMe Storage 1920GB	2	SSD/NVMe Storage 1920GB	2	SSD Storage 960GB	2
HDD Capacity Tier	7200RPM SAS HDD -- Up to 14TB per disk	12	7200RPM SAS HDD -- Up to 16TB per disk	30	7200RPM SAS HDD -- Up to 14TB per disk	4-12	7200RPM SAS HDD -- Up to 14TB per disk	4-12
Networking	10/25GbE	2	10/25GbE	2	10GbE	2	10GbE	2

Note that:

- The "Minimum VM Spec" is applicable for test, development, or evaluation environments, and is not recommended for running in production.
- For production environments the recommended specifications in the table are a starting point. For guidance on scaling a HyperStore deployment to meet your workload requirements, consult with your Cloudian sales representative.

- For virtualization environments, running HyperStore on VMware ESXi and vSphere is supported so long as the VMs have specs meeting or exceeding those in the table. However, avoid KVM or Xen as there are known problems with running HyperStore in those virtualization environments. For more guidance on deploying HyperStore on VMware, ask your Cloudian representative for the "Best Practices Guide: Virtualized Cloudian HyperStore on VMware vSphere and ESXi".

3.1.2. Operating System Requirements

To install HyperStore 7.5.2, on each host machine the operating system must be **CentOS Minimal version 7.4 or newer 7.x** (or Red Hat Enterprise Linux 7.4 or newer 7.x). HyperStore 7.5.2 does **not** support installation on:

- Older versions of CentOS/RHEL
- CentOS/RHEL 8.x or newer
- Other types of Linux distribution
- Non-Linux operating systems

If you have not already done so, install CentOS Minimal 7.4 or newer 7.x (or RHEL 7.4 or newer 7.x) in accordance with your hardware manufacturer's recommendations.

> **Note** Cloudian recommends using CentOS/RHEL **7.7 or newer 7.x**.

Below, see these additional requirements related to host systems on which you intend to install HyperStore:

- **"Partitioning of Disks Used for the OS and Metadata Storage"** (page 16)
- **"Host Firewall Services Must Be Disabled"** (page 17)
- **"Python 2.7.x is Required"** (page 17)
- **"Do Not Mount /tmp Directory with 'noexec'"** (page 17)
- **"root User umask Must Be 0022"** (page 18)

3.1.2.1. Partitioning of Disks Used for the OS and Metadata Storage

For the disks used for the OS and metadata storage -- typically two mirrored SSDs as noted in the hardware requirements table above -- **do not accept the default partition schemes offered by CentOS/RHEL**:

- By default CentOS/RHEL allocates a large portion of disk space to a */home* partition. This will leave inadequate space for HyperStore metadata storage.
- By default CentOS/RHEL proposes using LVM. Cloudian recommends using standard partitions instead.

Cloudian recommends that you manually create a partition scheme like this:

For Software RAID

- 1x 1G as */boot*, Device Type RAID1, label *boot*, fs: *ext4*
- 1x 8G as SWAP, Device Type RAID1, label *swap*
- 1x remaining space as / , Device Type RAID 1, label *root*, fs: *ext4*

For Hardware SUDO RAID with UEFI

- 1x 1G as */boot/efi*, label *efi*
- 1x 1G as */boot*, label *boot*, fs: *ext4*
- 1x 8G as SWAP, label *swap*
- 1x remaining space as */* , label *root*, fs: *ext4*

3.1.2.2. Host Firewall Services Must Be Disabled

To install HyperStore the following services **must be disabled on each HyperStore host machine**:

- *firewalld*
- *iptables*
- *SELinux*

To disable ***filewalld***:

```
# systemctl stop firewalld
# systemctl disable firewalld
```

RHEL/CentOS 7 uses *firewalld* by default rather than the *iptables* service (*firewalld* uses *iptables* commands but the *iptables* service itself is not installed on RHEL/CentOS by default). So you do not need to take action in regard to ***iptables*** unless you installed and enabled the *iptables* service on your hosts. If that's the case, then disable the *iptables* service.

To disable **SELinux**, edit the configuration file */etc/selinux/config* so that *SELINUX=disabled*. Save your change and then restart the host.

HyperStore includes a built-in firewall service (a HyperStore-custom version of the *firewalld* service) that is configured to protect HyperStore internal services while keeping HyperStore public services open. In fresh installations of HyperStore 7.2 or later, the HyperStore firewall is enabled by default upon the completion of HyperStore installation. In HyperStore systems originally installed as a version older than 7.2 and then later upgraded to 7.2 or newer, the HyperStore firewall is available but is disabled by default. After installation of or upgrade to HyperStore 7.2 or later, you can enable or disable the HyperStore firewall by using the installer's Advanced Configuration Options. For instructions see the "HyperStore Firewall" section in the ***Cloudian HyperStore Administrator's Guide***.

> **Note** For information about HyperStore port usage see **"HyperStore Listening Ports"** (page 41).

3.1.2.3. Python 2.7.x is Required

The HyperStore installer **requires Python version 2.7.x**. The installer will abort with an error message if any host is using Python 3.x. To check the Python version on a host:

```
# python --version
Python 2.7.5
```

3.1.2.4. Do Not Mount */tmp* Directory with 'noexec'

The */tmp* directory on your host machines must not be mounted with the '*noexec*' option. If the */tmp* directory is mounted with '*noexec*', you will not be able to extract the HyperStore product package file and the HyperStore

installer (installation script) will not function properly.

3.1.2.5. root User umask Must Be 0022

On hosts on which you will install HyperStore, the *root* user umask value must be '0022' (which is the default on Linux hosts). If the *root* user umask is other than '0022' the HyperStore installation will abort.

3.1.3. Note: Automatic Exclusions to OS Package Updates

As part of HyperStore installation, the HyperStore installation script will install prerequisites including Puppet, Facter, Ruby, and Salt on your HyperStore host machines. If you subsequently use *yum update* or *yum upgrade* to update your OS packages, HyperStore automatically excludes Puppet, Facter, Ruby, and Salt related packages from the update. This is to ensure that only the correct, tested versions of these packages are used together with HyperStore. After HyperStore installation, this auto-exclusion is configured in the */etc/yum/pluginconf.d/versionlock.list* file on your host machines. You can review that file if you wish to see specifically which packages are "locked" at which versions, but do not remove any entries from the lock list.

3.2. Preparing Your Nodes For HyperStore Installation

Subjects covered in this section:

Note These instructions assume that you have already configured basic networking on each of your nodes. In particular, each node must already be configured with a hostname and IP address, and the nodes must be able to reach each other in the network.

3.2.1. Installing HyperStore Prerequisites

After verifying that your nodes meet HyperStore's **hardware and OS requirements**, follow these steps to install and configure HyperStore prerequisites on all of your nodes:

1. Log into one of your nodes as *root*. This will be the node through which you will orchestrate the HyperStore installation for your whole cluster. Also, as part of the HyperStore installation, **Puppet** configuration management software will be installed and configured in the cluster, and this HyperStore node will become the **Configuration Master node** for purposes of ongoing cluster configuration management. Note that the Configuration Master node must be one of your HyperStore nodes. It cannot be a separate node outside of your HyperStore cluster.
2. On the node that you've logged into, download or copy the HyperStore product package (*CloudianHyperStore-7.5.2.bin* file) into a working directory. Also copy your Cloudian license file (**.lic* file) into that same directory. Pay attention to the license file name since you will need the license file name in the next step.

   Note The license file must be your cluster-wide license that you have obtained from Cloudian, not a license for an individual HyperStore Appliance machine (not a *cloudian_appliance.lic* file).

3. In the working directory run the commands below to unpack the HyperStore package:

```
# chmod +x CloudianHyperStore-7.5.2.bin
# ./CloudianHyperStore-7.5.2.bin <license-file-name>
```

This creates an **installation staging directory** named */opt/cloudian-staging/7.5.2*, and extracts the HyperStore package contents into the installation staging directory.

> **Note** The installation staging directory must persist for the life of your HyperStore system. Do not delete the installation staging directory after completing the install.

4. Change into the installation staging directory:

```
# cd /opt/cloudian-staging/7.5.2
```

5. In the installation staging directory, launch the *system_setup.sh* tool:

```
# ./system_setup.sh
```

This displays the tool's main menu.

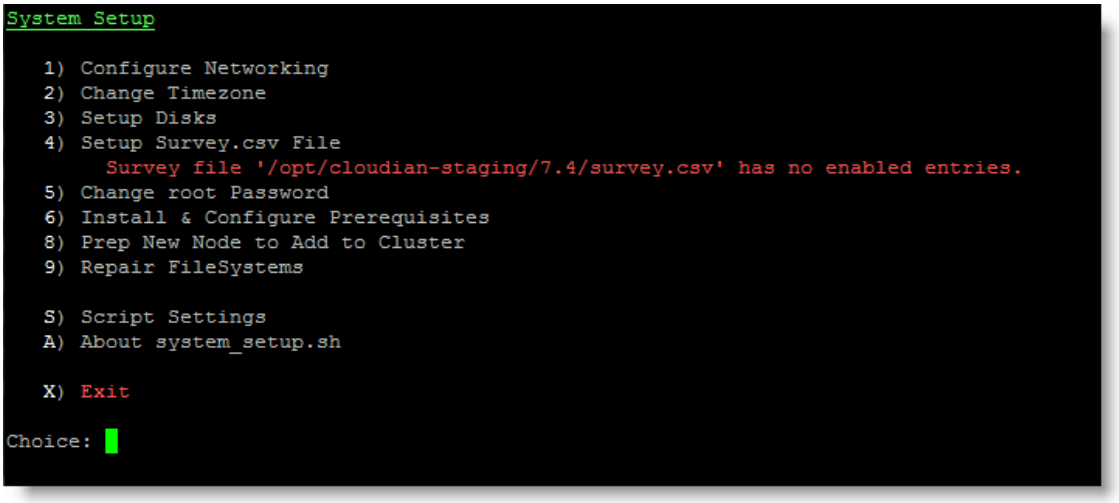

6. From the setup tool's main menu, enter "**4**" for **Setup Survey.csv File** and follow the prompts to create a system survey file with an entry for each of your HyperStore nodes (including the Configuration Master node). For each node you will enter a region name, hostname, public IPv4 address, data center name, and rack name. For each node you can also optionally enter an internal interface name.

   - For each node's hostname, specify the node's **short hostname** (as would be returned if you ran the *hostname -s* command on the node) -- not an FQDN.

     > **Note** Do not use the same short hostname for more than one node in your entire HyperStore system. Each node must have a unique short hostname within your entire HyperStore system, even in the case of nodes in different data centers or service regions that have different domains.

   - For the region, data center, and rack name the only allowed character types are ASCII alpha-numerical characters and dashes. For the region name letters must be lower case. **Do not include the string "s3"** in the region name.

- Make sure the region name matches the region string that you use in your S3 endpoints in your **"DNS Set-Up"** (page 10).
- Within a data center, **use the same "rack name" for all of the nodes**, even if some nodes are on different physical racks than others.

- For each node, you can optionally specify the name of the interface that the node uses for internal cluster communications.
  - For each node for which you do **not** specify an internal interface name here in the survey file, HyperStore will use a default internal interface name that you will supply later in the HyperStore installation process.
  - For each node for which you **do** specify an internal interface name here in the survey file, HyperStore will use that internal interface name for that node. The node-specific internal interface name in the survey file overrides the default internal interface name that you will supply later in the HyperStore installation process.

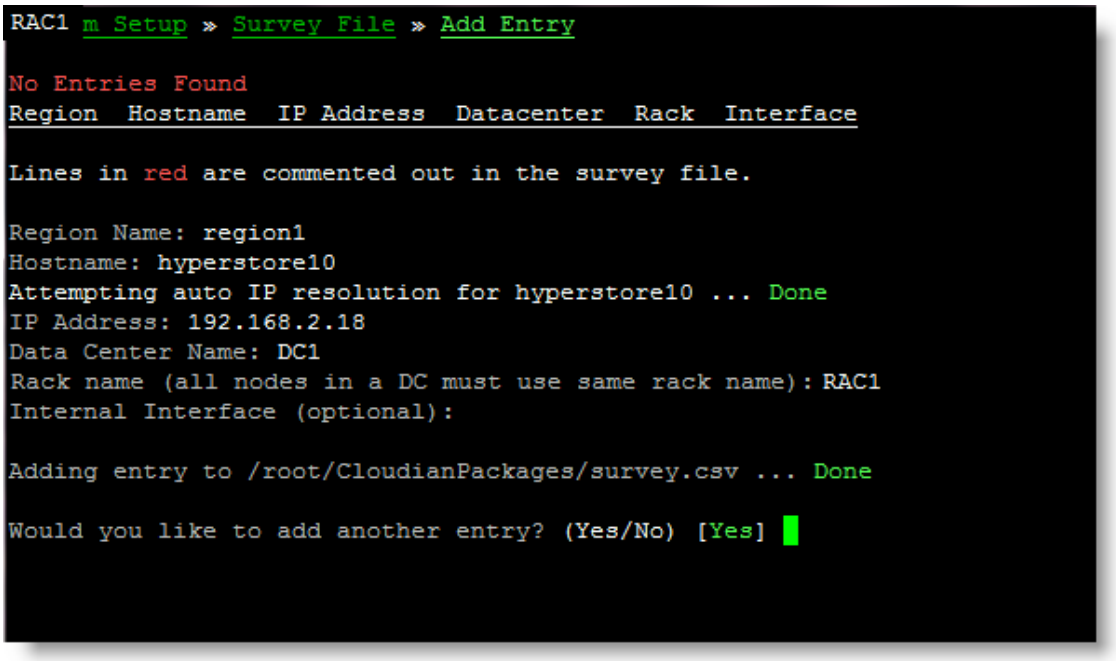

After you've added an entry for each node, return to the setup tool's main menu.

> **Note** Based on your input at the prompts, the setup tool creates a survey file named *survey.csv*, in your installation staging directory. This file must remain in your staging directory -- do not delete or move it. For more information about the contents of the survey file, see the Installation Reference topic **"survey.csv (Cluster Survey File)"** (page 50).

7. If you want to change the root password for your nodes, do so now by entering "**5**" for **Change Root Password** and following the prompts. It's recommended to use the same root password for each node. Otherwise the pre-installation cluster validation tool described later in the procedure will not be fully functional.

> **Note** If your host machines are "hardened" HyperStore Appliances -- which have the HyperStore Shell already enabled and the root password disabled -- then the "Change Root Password" option will not appear in the setup tool's main menu.

8. Back at the setup tool's main menu enter "**6**" for **Install & Configure Prerequisites**. When prompted about whether you want to perform this action for all nodes in your survey file enter "**yes**". The tool will connect to each of your nodes in turn and install the prerequisite packages. You will be prompted to provide the root password for the cluster nodes (unless an SSH key is present, in which case that will be used rather than a password). When the prerequisite installation completes for all nodes, return to the setup tool's main menu.

> **Note** If *firewalld* is running on your hosts the setup tool prompts you for permission to disable it. And if *Selinux* is enabled on your hosts, the tool automatically disables it without prompting for permission (or more specifically, changes it to "permissive" mode for the current running session and changes the configuration so it will be disabled for future boots of the hosts). For information on why these services must be disabled on HyperStore host machines see **"Operating System Requirements"** (page 16).

After the prerequisite installation completes for all nodes (as indicated by console messages from the setup tool), return to the setup tool's main menu and proceed to **"Configuring Network Interfaces, Time Zone, and Data Disks"** (page 21).

3.2.2. Configuring Network Interfaces, Time Zone, and Data Disks

After **"Installing HyperStore Prerequisites"** (page 18), you should be at the main menu of the *system_setup.sh* tool. Next follow these steps to configure network interfaces (if you haven't already fully configured them), set the time zone, and configure data disks on each node in your HyperStore cluster.

1. From the system setup tool's main menu, complete the setup of the Configuration Master node itself:
   a. From the system setup tool's main menu, enter "**1**" for **Configure Networking**. This displays the Networking configuration menu.

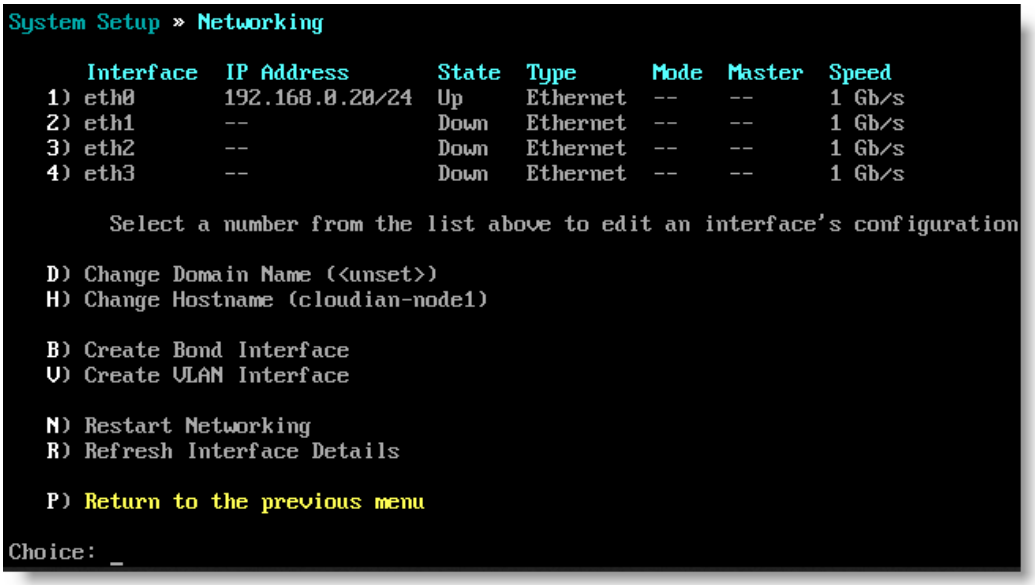

Here you can review the current network interface configuration for the Configuration Master node, and if you wish, perform additional configuration such as configuring an internal/back-end interface. When you are done with any desired network interface configuration changes for this node, return to the setup tool's main menu.

> **Note** When setting/changing a node's hostname, if you enter a hostname that includes upper case letters the setup tool automatically converts the hostname to entirely lower case letters.

b. From the setup tool's main menu, enter "**2**" for **Change Timezone** and set the time zone for this node.

c. From the setup tool's main menu, enter "**3**" for **Setup Disks**. This displays the Setup Disks menu.

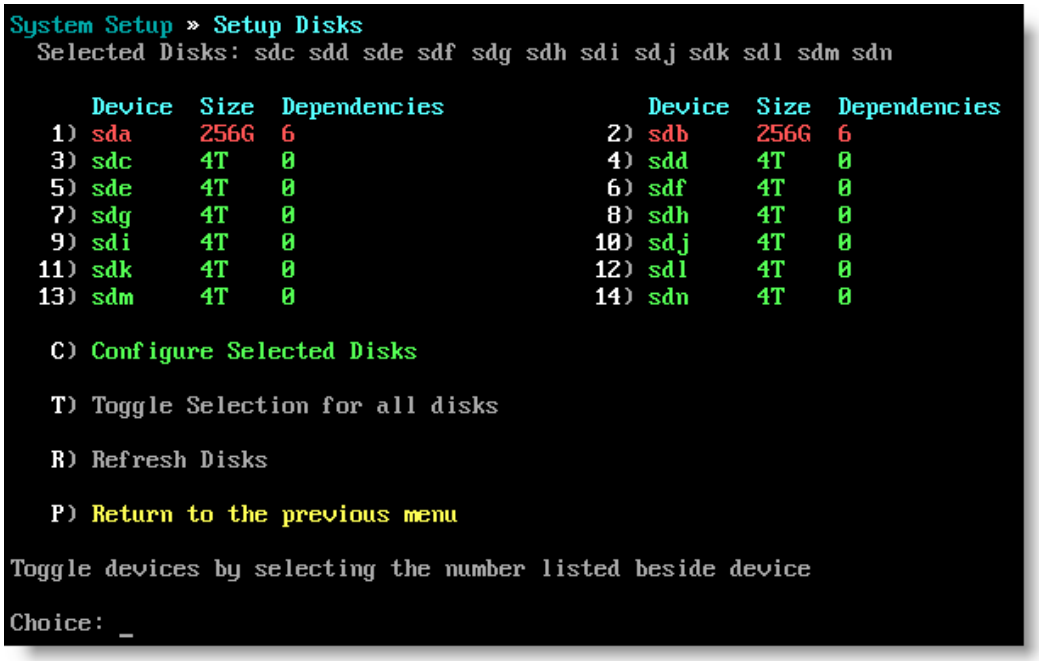

From the list of disks on the node select the disks to format as HyperStore data disks, for storage of S3 object data. By default the tool automatically selects all disks that are not already mounted and do not contain a */root*, */boot* or *[swap]* mount indication. Selected disks display in green font in the disk list. The tool will format these disks with *ext4* file systems and assign them mount points */cloudian1*, */cloudian2*, */cloudian3*, and so on. You can toggle (select/deselect) a disk by entering at the prompt the disk's number from the displayed list (such as "**3**"). Once you're satisfied with the selected list in green font, enter "**c**" for **Configure Selected Disks** and follow the prompts to have the tool configure the selected disks.

> **IMPORTANT !** Cloudian recommends using the HyperStore system setup tool to format and mount your data disks. **If you have already formatted and mounted your data disks using third party tools**, then instead of using the disk configuration instructions in this section follow the guidelines and instructions in **"File System Requirements"** (page 46).

2. Next, complete the setup of the other nodes in your cluster:
    a. From the setup tool's main menu enter "**8**" for **Prep New Node to Add to Cluster**.
    b. When prompted enter the IP address of one of the remaining nodes (the nodes other than the Configuration Master node), and then enter the login password for the node.
    c. Using the node preparation menu that displays:
        i. Review and complete network interface configuration for the node.
        ii. Set the time zone for the node.
        iii. Configure data disks for the node. Then return to the system setup tool's main menu.
    d. Repeat Steps "a" through "c" for each of the remaining nodes in your installation cluster.

After you've prepared all your nodes and returned to the setup tool's main menu, proceed to **"Running the Pre-Install Checks Script"** (page 23).

3.2.3. Running the Pre-Install Checks Script

Follow these steps to verify that your cluster now meets all HyperStore requirements for hardware, prerequisite packages, and network connectivity.

1. At the setup tool's main menu enter "**r**" for **Run Pre-Installation Checks**. This displays the Pre-Installation Checklist menu.

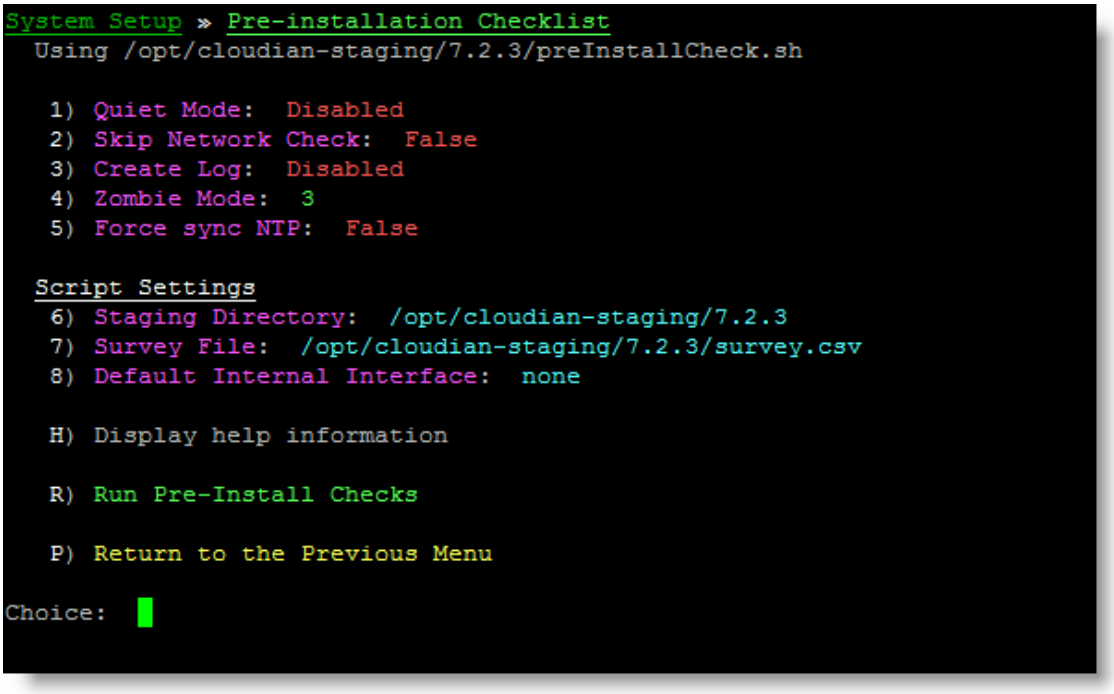

2. From the Pre-Installation Checklist menu enter "**r**" for **Run Pre-Install Checks**. After prompting you for a cluster password the script checks to verify that your cluster meets all requirements for hardware, prerequisite packages, and network connectivity.

> **Note** The script only supports your providing one root password, so if some of your nodes do not use that password the script will not be able to check them and you may encounter errors during HyperStore installation if requirements are not met.

At the end of its run the script outputs to the console a list of items that the script has evaluated and the results of the evaluation. You should review any “Warning” items but they don’t necessarily require action (an example is if the hardware specs are less than recommended but still adequate for the installation to proceed). **You must resolve any “Error” items before performing the HyperStore software installation**, or the installation will fail.

When you’re done reviewing the results, press any key to continue and then exit the setup script. **If you make any system changes to resolve errors found by the pre-install check, run the pre-install check again afterward** to verify that your environment meets HyperStore requirements.

After your cluster has successfully passed the pre-install checks, proceed to **"Installing a New HyperStore System"** (page 25).

Chapter 4. Installing a New HyperStore System

This section describes how to do a fresh installation of HyperStore 7.5.2 software, after **"Preparing Your Environment"** (page 9) and **"Preparing Your Nodes For HyperStore Installation"** (page 18). From your Configuration Master node you can install HyperStore software across your whole cluster.

Note Applicable to software-only customers: By default HyperStore appends the existing */etc/ssh/sshd_config* file on HyperStore host machines, adding a HyperStore-managed section to the file. It does so in order to facilitate FIPS support and support of the HyperStore Shell feature. If you do not want HyperStore to append the */etc/ssh/sshd_config* file on HyperStore host machines, then before performing the installation steps below edit */etc/cloudian-7.5.2-puppet/manifests/extdata/common.csv* on the Configuration Master node so that *sshdconfig_disable_override* is set to *true* (by default it is *false*). Then proceed with the installation.

1. On your Configuration Master node, in your installation staging directory (*/opt/cloudian-staging/7.5.2*), launch the HyperStore installation script as follows:

   ```
   [7.5.2]# ./cloudianInstall.sh -s survey.csv
   ```

   Note **If you have not configured your DNS environment** for HyperStore (see **"DNS Set-Up"** (page 10)) and you want to instead use the included *dnsmasq* utility to resolve HyperStore service endpoints, launch the install script with the *configure-dnsmasq* option as shown below. This is not appropriate for production systems.

   [7.5.2]# ./cloudianInstall.sh -s survey.csv configure-dnsmasq

   For more script launch options, see the Installation Reference topic **"cloudianInstall.sh"** (page 52).

   When you launch the installer the main menu displays:

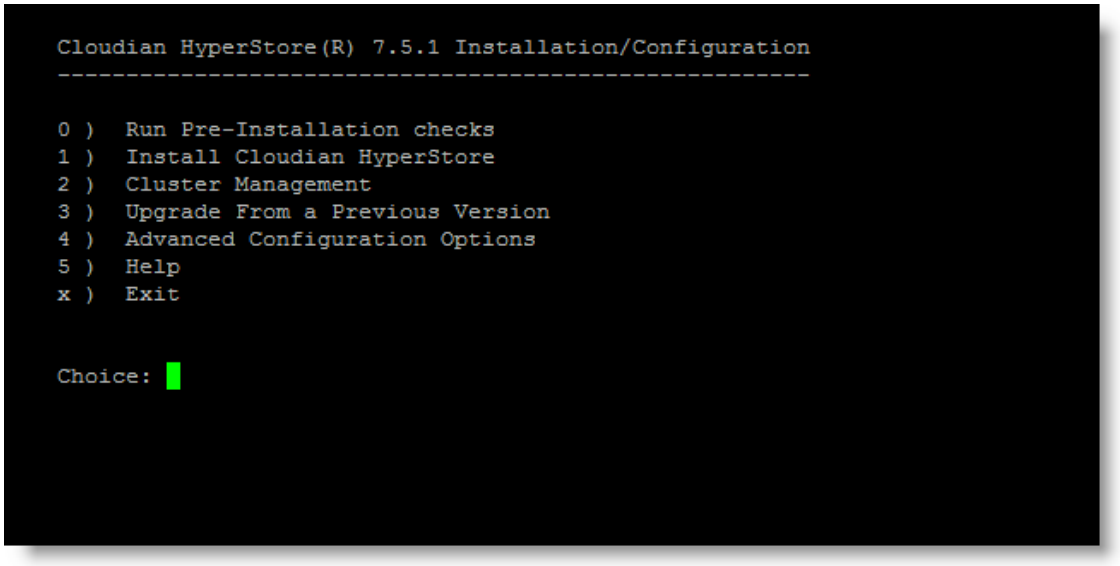

Note The installer menu includes an item "0" for Run Pre-Installation Checks. This is the same pre-installation check that you already ran from within the *system_setup.sh* tool as described in **"Running the Pre-Install Checks Script"** (page 23) -- so you can ignore this option in the installer menu. If you did **not** run the pre-install check already, then do so from the installer menu before proceeding any further.

2. From the installer main menu, enter "**1**" for Install Cloudian HyperStore. Follow the prompts to perform the HyperStore installation across all the nodes in your cluster survey file (which you created earlier during the node preparation task).

   During the HyperStore installation you will be prompted to provide cluster configuration information including the following:

   - The name of the **internal interface** that your nodes will use by default for internal cluster communications. For example, *eth1*.

     Note The system will use this default internal interface name for all nodes for which you did not specify an internal interface name in your cluster survey file (which you created during the **"Installing HyperStore Prerequisites"** (page 18) procedure). If in the survey file you specified internal interface names for some or all of your nodes, the system will use those internal interface names for those nodes, rather than the default internal interface name.

   - The starting "**replication strategy**" that you want to use to protect system metadata (such as usage reporting data and user account information). The replication strategy you enter must be formatted as "<datacenter_name>:<replication_#>". For example, "DC1:3" means that in the data center named DC1, three instances of each system metadata object will be stored (with each instance on a different host). If you are installing HyperStore into multiple data centers you must format this as a comma-separated list specifying the replicas per data center -- for example "DC1:2,DC2:1". The default is 3 replicas per service region, and then subsequently the system automatically adjusts the system metadata replication level based on the storage policies that you create. For more on this topic see "Storage of System Metadata" in the ***Cloudian***

HyperStore Administrator's Guide.

- Your **organization's domain**. For example, *enterprise.com*. From this input that you provide, the installation script will automatically derive HyperStore service endpoint values. You can accept the derived endpoint values that the script presents to you, or optionally you can enter customized endpoint values at the prompts. For S3 service endpoint the default is to have one endpoint per service region, but you also have the option of entering multiple comma-separated endpoints within a service region -- if for example you want different data centers within the region to use different S3 service endpoints. If you want to have different S3 endpoints for different data centers within the same service region, the recommended S3 endpoint syntax is *s3-<region>.<dcname>.<domain>*. See **"DNS Set-Up"** (page 10) for more details about HyperStore service endpoints. Note that if you have multiple S3 service endpoints for your system, the first S3 service endpoint in your comma-separated list will be the "default" S3 service endpoint (and so will be used in public URLs if CMC users generate public URLs for some of their objects).

  IMPORTANT !
  * Do not use IP addresses in your service endpoints.
  * Including "s3" in the *<domain>* value is not recommended. By default HyperStore generates S3 service endpoints by prepending an "*s3-*" prefix to your *<region-name>.<domain>* combination. If you include "s3" within either your domain or your region name, this will result in two instances of "s3" in the system-generated S3 service endpoints, and this may cause S3 service requests to fail for some S3 clients.
  * HyperStore by default derives the S3 service endpoint(s) as *s3-<region-name>.<domain>*. However HyperStore also supports the format *s3.<regionname>.<domain>* (with a dot after the "s3" rather than a dash) if you specify a custom S3 endpoint with this format..
  * Your S3 static website endpoint cannot be the same as your S3 service endpoint. They must be different, or else the static website endpoint will not work properly.

- The **NTP servers** that HyperStore nodes should connect to for time synchronization. By default the public servers from the *pool.ntp.org* project are used. If you do not allow outbound connectivity from HyperStore hosts (and consequently public NTP servers cannot be reached) you must specify NTP server(s) within your environment that HyperStore hosts should connect to instead. The installation will fail if HyperStore hosts cannot connect to an NTP server.

At the conclusion of the installation an "Install Cloudian HyperStore" sub-menu displays, with indication of the installation status. If the installation completed successfully, the "Load Schema and Start Services" menu item should show an "OK" status:

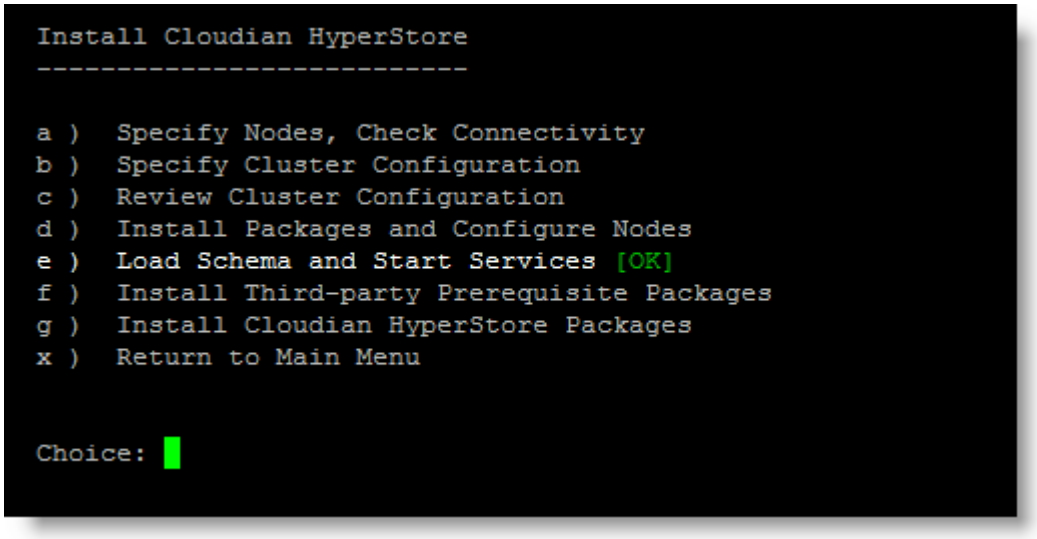

After seeing that the "Load Schema and Start Services" status is OK, return to the installer's main menu.

> **Note** The "Install Cloudian HyperStore" sub-menu supports re-executing specific installation operations on specific nodes or on all nodes. This may be helpful if the installer interface indicates that an operation has failed. If one of the operations in the menu indicates an error status, retry that operation by specifying the menu option letter at the prompt (such as "**e**" for "Load Schema and Start Services").

3. After installation has completed successfully, from the installer's main menu enter "**2**" for Cluster Management and then enter "**d**" for Run Validation Tests. This executes some basic automated tests to confirm that your HyperStore system is working properly. The tests include S3 operations such as creating an S3 user group, creating an S3 user, creating a storage bucket for that user, and uploading and downloading an S3 object.

After validation tests complete successfully, exit the installation tool.

For first steps to set up and try out your new HyperStore system, see "Getting Started with a New HyperStore System" in the ***Cloudian HyperStore Administrator's Guide***.

> **Note** For troubleshooting information, see the Installation Reference topic **"Installation Troubleshooting"** (page 40).

Chapter 5. Installing the Search Service and the File Service

This section covers these topics regarding optional services that you can add to an installed HyperStore system if you wish:

5.1. Search Service and File Service Installation Introduction

After you've installed your HyperStore system and verified that it's operating properly -- or anytime after you've put your HyperStore system into operation -- you can optionally install one or both of these additional HyperStore service components:

- **HyperStore Search Service** -- The HyperStore Search Service provides object metadata search functionality to your HyperStore users. Once you've installed the HyperStore Search Service in your HyperStore cluster, you can enable metadata search on a per storage policy basis. Users who create buckets that use a search-enabled storage policy can use the CMC to search for objects based on system-defined object metadata or user-defined object metadata. For more information about this feature, in the HyperStore Help (or the *HyperStore Administrator's Guide*) see *Setting Up Service Features -> Object Metadata and Search*.
- **HyperStore File Service** -- The HyperStore File Service provides SMB and NFS file services functionality to your HyperStore users. Once you've installed the HyperStore File Service in your HyperStore cluster, you can choose HyperStore users for whom to enable File Services functionality. Those users can then use the CMC to create and manage SMB and/or NFS file shares. The HyperStore File Service provides front end SMB and NFS access to these shares as well as a caching layer to speed performance, while the users' HyperStore buckets provide back end archival storage for the shares.For more information about this feature, in the HyperStore Help (or the *HyperStore Administrator's Guide*) see *Setting Up Service Features -> File Services (SMB/NFS)*.

The HyperStore Search Service and the HyperStore File Service install on HyperStore "**auxiliary nodes**". Auxiliary nodes are nodes that:

- You add to your HyperStore cluster after the core HyperStore system is already installed and running.
- Have lesser resource requirements than your core HyperStore system nodes (such as lesser disk space requirements).
- Run Rocky Linux as their operating system (unlike your core HyperStore system nodes, which run CentOS or RHEL; note that starting with HyperStore version 8.0 the core HyperStore system nodes will run Rocky Linux also.)
- Run K3s (a lightweight Kubernetes distribution) to orchestrate HyperStore Search Service or HyperStore File Service components as containerized applications. K3s is a largely transparent part of the HyperStore Search Service and HyperStore File Service application stacks and does not require explicit maintenance from you as the system administrator.

The HyperStore Search Service and the HyperStore File Service are independent of one another -- you can install one without the other if you wish.

For **development or testing** purposes, each of these services requires one auxiliary node. You cannot install both the HyperStore Search Service and the HyperStore File Service on the same auxiliary node.

For **production** purposes:

- The HyperStore Search Service requires three auxiliary nodes.
- The HyperStore File Service can run on one node or three nodes depending on your use case:
  - One node may be sufficient if only a small number of SMB or NFS clients will be accessing file shares and if high availability is not critical for your file services use case.
  - Three nodes are required if many SMB or NFS clients will be concurrently accessing file shares, and/or your use case requires high availability file services.

  Consult with your Cloudian representative to determine whether your production needs are best met by one HyperStore File node or three HyperStore File nodes.

Regardless of how many HyperStore Search nodes and HyperStore File nodes you have, you cannot install both the HyperStore Search Service and the HyperStore File Service on the same node.

In a production environment the HyperStore Search Service and HyperStore File Service can be installed without disruption to your HyperStore S3 storage service.

For more information see:

- **"Installing the HyperStore Search Service"** (page 30)
- **"Installing the HyperStore File Service"** (page 37)

5.2. Installing the HyperStore Search Service

Please read the **"Search Service and File Service Installation Introduction"** (page 29) before proceeding with this section.

Subjects covered in this section:

- ***"Prerequisites and High Level Workflow"** (page 30)*
- ***"Setting Up the HyperStore Search Host Machines"** (page 31)*
- ***"Integrating the HyperStore Search Nodes Into Your HyperStore Cluster"** (page 33)*
- ***"Installing the HyperStore Search Service Software"** (page 37)*
- ***"Completing the Set-Up of the Metadata Search Feature"** (page 37)*

5.2.1. Prerequisites and High Level Workflow

The prerequisites for installing the HyperStore Search Service are:

- You've already completed the installation of HyperStore and your HyperStore system is up and running.
- You've obtained the HyperStore Rocky Linux OVA from Cloudian. (Contact Cloudian to obtain the OVA if you haven't yet done so.)

> **Note** You must use the HyperStore Rocky Linux OVA from Cloudian to create your VMs. The OVA has the correct set of packages to support the HyperStore Search Service. **Do not use the generic Rocky Linux ISO** from the Rocky Linux organization.

- You have a virtualization platform on which you can create VMs. For production deployments use an enterprise grade virtualization platform such as VMware ESXi (version 6.7 or newer).
- The hosts on which you will install the HyperStore Search Service have internet access.

At a high level the workflow for installing and setting up the HyperStore Search Service is as follows:

1. Set up either one or three VMs to serve as auxiliary nodes, using the HyperStore Rocky Linux OVA.
2. On your HyperStore Configuration Master node, use the HyperStore installer to configure the integration of the HyperStore Search auxiliary nodes into your HyperStore cluster. (In this same step you can also configure the integration of HyperStore File nodes into HyperStore, if you are installing both the HyperStore Search Service and the HyperStore File Service concurrently.)
3. Use the HyperStore installer to install HyperStore Search software on the auxiliary nodes.
4. Complete the set-up of the Metadata Search feature.

Each of these steps are described below.

5.2.2. Setting Up the HyperStore Search Host Machines

The HyperStore Search Service is designed to run on either:

- One virtual machine (for development, testing, and evaluation only)
- Three virtual machines (for production)

The HyperStore Search Service does not support running on two hosts or on more than three hosts.

The HyperStore Search host(s) must be in one of your HyperStore data centers. If you are going to run the HyperStore Search Service on three host machines, all three hosts must be in the same data center.

To set up a VM for hosting the HyperStore Search Service:

1. Use your virtualization platform console to deploy a new virtual machine (VM) from the HyperStore Rocky Linux OVA file that you obtained from Cloudian. (For example, in VMware ESXi this is the “Deploy a virtual machine from an OVF or OVA file” option.) It may take some time for the creation of the VM to complete. The VM that is created has these basic specs:
    - 4 vCPUs
    - 16 GB RAM
    - 2 virtual disks, one at 250GB (for OS and software) and one at 300GB (for search index storage)
    - OS = Rocky Linux 8.7
2. After VM creation has completed and the VM is running, use the virtualization platform console to log into the VM will user name *root* and password *HSApplianceLinux8*
3. Change into the directory */root/CloudianTools*
4. Launch the Cloudian system setup tool:

```
./system_setup.sh
```

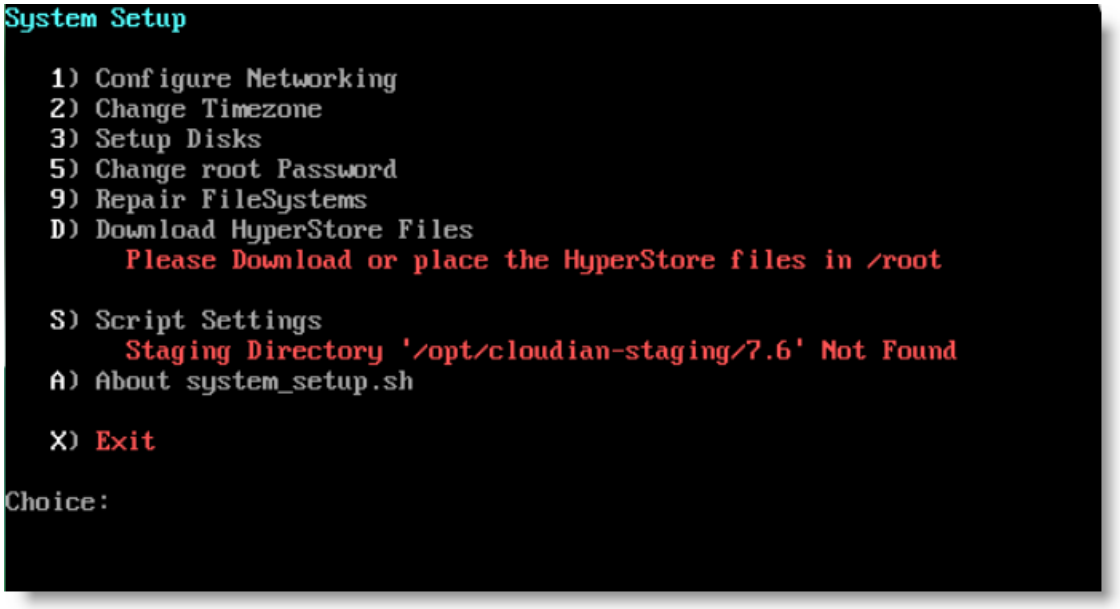

Note Ignore the red messages that display under options "D" and "S".

5. From the system setup tool main menu, complete the following tasks:

   1) Configure Networking. Select this main menu option to move to the Networking sub-menu. From the Networking sub-menu

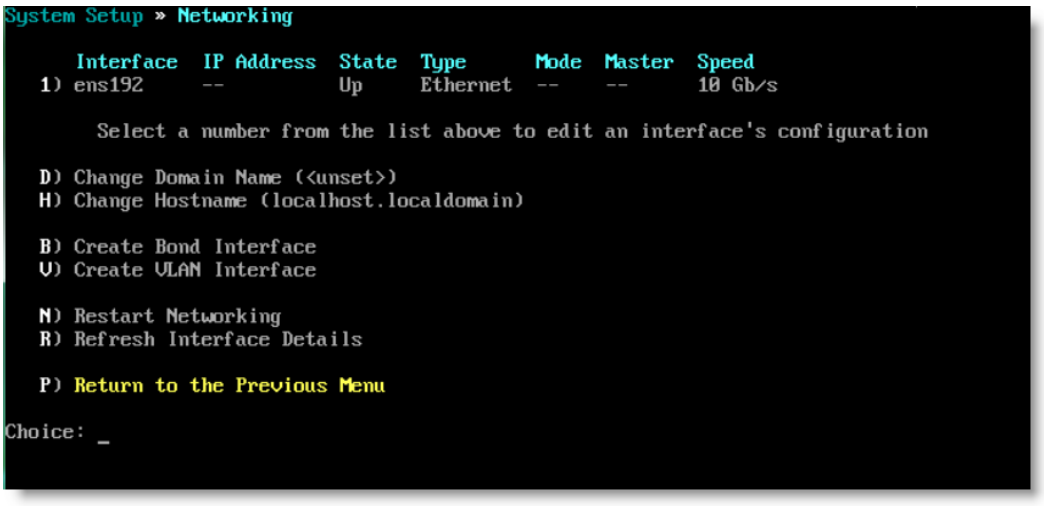

   From the Networking sub-menu, first edit the network interface configuration (follow the prompts to set a static IPv 4 address, a network mask, and optionally a default gateway and a DNS address; and when prompted answer yes to restart the interface to apply the configuration). Then back at the Networking sub-menu complete options "D" and "H" to set the domain name and host-name for this host. Then return to the system setup main menu.

   2) Change Timezone. Select this main menu option and then follow the prompts to set the time zone in which this host resides.

5) Change root Password. Select this main menu option and then follow the prompts to change the root password for this host. (Recall that the default root password is *HSApplianceLinux8*, but for security you should change this to a password of your choosing.)

Do not perform these tasks from the main menu:

3) Setup Disks. **Do not do this.** Disk setup will be executed automatically when you install HyperStore Search Service software later in this procedure.

9) Repair FileSystems. Do not do this.

D) Download HyperStore Files. Do not do this.

Note This is the same setup tool that is used on HyperStore core nodes, and so some menu tasks -- those noted above -- are not applicable when setting up auxiliary nodes for hosting the HyperStore Search Service.

If you are going to install the HyperStore Search Service on three hosts, **complete the steps above for each of the three hosts.**

Note that if you set the same root password for each of the three hosts you will have fewer prompts to answer when you use the HyperStore interactive installer to complete the HyperStore Search Service configuration and installation procedures.

5.2.3. Integrating the HyperStore Search Nodes Into Your HyperStore Cluster

After you've set up your HyperStore Search host machines, use the HyperStore installer to integrate the HyperStore Search hosts into your HyperStore system's configuration as described further below. You will need the following information:

- Root password for SSH access to the HyperStore Search hosts. This is used only during the installation.
- Name of the HyperStore service region and data center in which the HyperStore Search hosts reside. These must match a region name and data center name from your HyperStore system configuration. (If you don't recall the region and data center names you can log into the CMC and find these names in the **Cluster -> Data Centers** page.)
- Hostname of each HyperStore Search host (short hostname, not FQDN)
- Physical IP address of each HyperStore Search host (IP address assigned to front end port of host)
- Two virtual IP addresses, both on the same network as the hosts' physical IP addresses:
  - One VIP to be used as the Kubernetes intranet IP address, for Kubernetes communications within the HyperStore Search Service cluster. (This IP address is not needed if you have only one HyperStore Search node)
  - One VIP to be used as the Search endpoint IP address. This cannot be the same as the Kubernetes intranet IP address.

    These two VIP addresses cannot be the same as any of your HyperStore Search host physical IP addresses.
- An FQDN to use for the Search service.
- A Search "prefix" to use for this HyperStore system. This is a lower case alphanumeric string of your choosing. Upper case letters, spaces, dashes, underscores, and special characters are not allowed.

This prefix will be used in the naming scheme for search indexes associated with this HyperStore system. (In the current HyperStore release, only one HyperStore system can use a given HyperStore Search cluster. But the use of a unique Search prefix per HyperStore system facilitates a possible enhancement in future HyperStore releases, whereby you could have multiple HyperStore systems access a single HyperStore Search cluster.)

Note If you are installing the HyperStore Search Service and the HyperStore File Service concurrently, you can integrate both sets of auxiliary nodes into your HyperStore system in this same procedure. For information you will need for HyperStore File node integration, see **"Installing the HyperStore File Service"** (page 37).

If you are installing the HyperStore Search Service after having previously installed the HyperFile Service, read **"Integrating HyperStore Search Nodes After Having Installed HyperStore File Previously"** (page 36) before proceeding with the steps below.

To integrate the HyperStore Search host nodes into your HyperStore cluster:

1. Log into your HyperStore Configuration Master node as *root*, then change into the installation staging directory (*/opt/cloudian-staging/7.5.2*). Once in the staging directory, launch the HyperStore installer:

   ```
   # ./cloudianInstall.sh
   ```

   This displays the installer's main menu:

   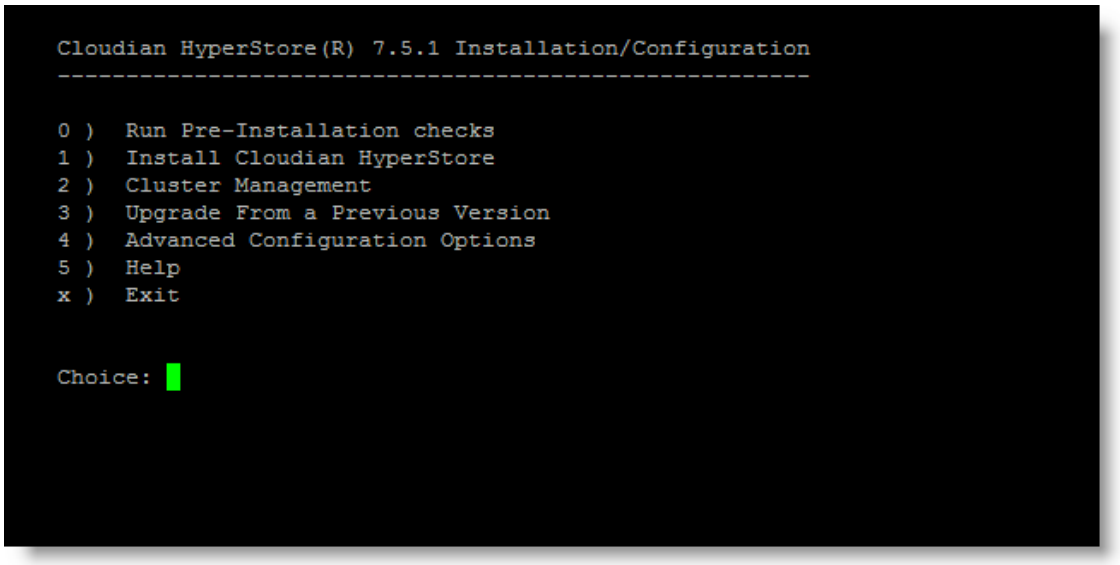

   *If you are using the **HyperStore Shell***

   If you are using the **HyperStore Shell (HSH)** as a **Trusted user**, from any directory on the Configuration Master node you can launch the installer with this command:

   ```
   $ hspkg install
   ```

   Once launched, the installer's menu options (such as referenced in the steps below) are the same regardless of whether it was launched from the HSH command line or the OS command line.

> **Note** Although you cannot use the HSH locally on auxiliary nodes, you can use it on your HyperStore Configuration Master when launching the installer to perform installer-driven set-up of auxiliary nodes.

2. Enter "1" for "Install Cloudian HyperStore". This displays the "Install Cloudian HyperStore" menu. (This menu is what displays if you've already installed your HyperStore system, which must be done before you can configure auxiliary node integration into the system.)

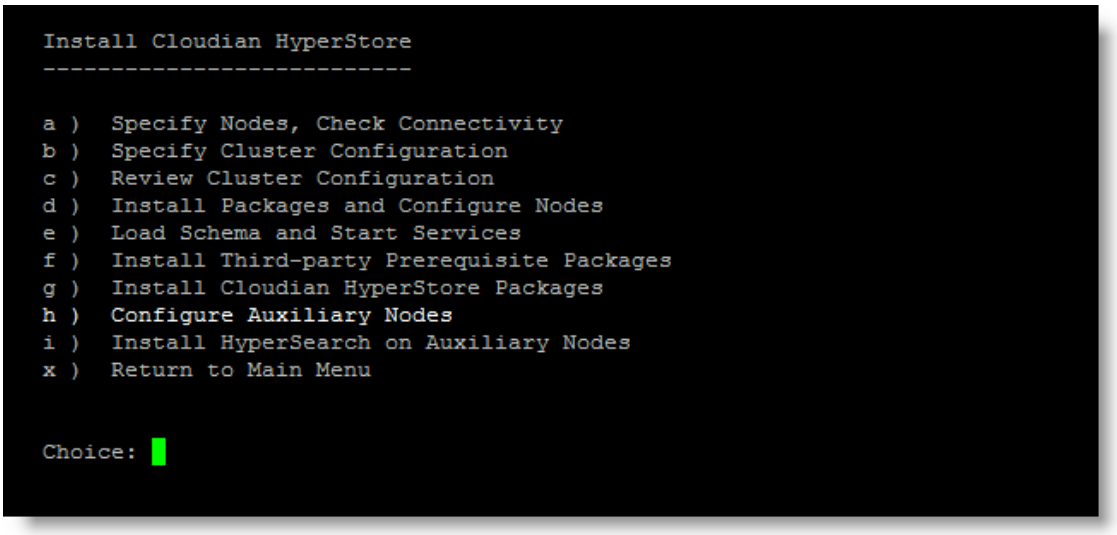

3. Enter "h" for "Configure Auxiliary Nodes". The installer will then prompt you for the information described in the introduction to this section, in this order:
   - First you will be prompted for the information required to create a dedicated survey file for your auxiliary nodes (nodes that will host either the HyperStore Search Service or the HyperFile Service). At the prompt for the survey file name accept the default name *survey_aux.csv* (by pressing enter; **do not change this survey file name** if you use the HyperStore Shell in your system). You will then be prompted to provide -- for each auxiliary host node in succession -- the service region name, hostname, host IP address, data center name, and host "role" (either *search* or *file*).
   - After completing the entries for the auxiliary survey file you will be prompted for the root password for your auxiliary node(s), and then for the search endpoint IP, the Kubernetes search intranet IP, the search FQDN, and the search prefix.
   - If you included File Services nodes among your entries for the auxiliary survey file, you will also now be prompted for the Kubernetes file intranet IP and the file endpoint IP.
   - When prompted *Would you like to push configurations now?*, answer *yes* and choose to push to all nodes.
4. Return to the installer's main menu and then enter "2" for "Cluster Management". From the Cluster Management menu enter "c" for "Manage Services", and then restart the Cloudian Management Console (CMC).

You've now completed the integration of auxiliary nodes into your HyperStore cluster.

Note The installer does not support editing the *survey_aux.csv* file once it's been created. If you make an error when creating the survey file via the installer, to correct the error you'll need to manually edit the *survey_aux.csv* file in the installation staging directory with a text editor (or with *hspkg config -e survey_aux.csv* if you are using the HyperStore Shell). After editing and closing the file, use the installer to again run the "h" for "Configure Auxiliary Nodes" option. This time you will not be prompted for survey file entries. At the other prompts (such as for the search endpoint IP) the default values shown at the prompts will be the values you already supplied -- you can accept these values if you do not wish to change them. When prompted *Would you like to push configurations now?*, answer *yes*.

As the next step, if you have not already done so, **create an entry in your environment's DNS system to resolve the Search FQDN to the Search endpoint IP address**. HyperStore will call the HyperStore Search Service by using the FQDN not the IP address, so the FQDN must be resolved by your DNS.

Once you've integrated the auxiliary nodes into your HyperStore cluster, and configured your DNS to resolve the Search FQDN, you can proceed to **"Installing the HyperStore Search Service Software"** (page 37).

5.2.3.1. Integrating HyperStore Search Nodes After Having Installed HyperStore File Previously

If you are integrating HyperStore Search nodes into your HyperStore system after having installed the HyperFile Service at a previous time, then **before performing the steps above** you must manually edit your existing auxiliary node survey file:

1. Log into your HyperStore Configuration Master node as *root*, then change into the installation staging directory (*/opt/cloudian-staging/7.5.2*).
2. Using your preferred text editor, open and edit the *survey_aux.csv* file. For each HyperStore Search node add a line to the auxiliary node survey file with this format:

   ```
   region-name,hostname,host-ip-address,data-center-name,role
   ```

   For a HyperStore Search node the role is *search*. For example:

   ```
   boston,tatum,10.50.41.33,DC1,search
   ```

   *If you are using the **HyperStore Shell***

   If you are using the **HyperStore Shell (HSH)** as a **Trusted user**, you can edit the auxiliary node survey file with this command

   ```
   $ hspkg config -e survey_aux.csv
   ```

   For more information on editing configuration files through the HSH see Using the HSH to Manage Configuration Files.

3. Save and close the edited auxiliary node survey file, then proceed with the steps in **"Integrating the HyperStore Search Nodes Into Your HyperStore Cluster"** (page 33) above. During those steps the installer will **not** prompt you for the information for creating a survey file; instead the installer will move directly to prompting you for the search endpoint IP and so on. See the steps above for details.

   Note After it prompts you for the search service values the installer will prompt you for the file service values. **For the file service, accept the existing values** (the values that you set when you installed HyperFile previously). These existing file service values will be shown at the installer prompts.

5.2.4. Installing the HyperStore Search Service Software

After using the HyperStore installer to integrate the HyperStore Search auxiliary nodes into your HyperStore cluster, use the installer to install HyperStore Search Service software on those auxiliary nodes. The installer will install and configure all needed software packages on to each node that you indicated as having the *search* role in your auxiliary node survey file. The software packages include K3s (a lightweight Kubernetes distribution) and OpenSearch version 1.3.8 (an open source search software suite).

To install HyperStore Search Service software:

1. If you do not already have the installer open and at the "Install Cloudian HyperStore" menu: Log into the HyperStore Configuration Master node as *root*, change into the installation staging directory (*/opt/cloudian-staging/7.5.2*), launch the installer (*./cloudianInstall.sh*), and at the installer's main menu enter "1" for "Install Cloudian HyperStore".
2. At the "Install Cloudian HyperStore" menu enter "i" for "Install HyperSearch On Auxiliary Nodes".

   This operation automatically installs and configures all needed software packages on each HyperStore Search Service auxiliary node, and starts the HyperStore Search Service and supporting services on those nodes.

   When prompted *Would you like to push configurations and restart S3 service?*, answer *yes*.

After the installer indicates a successful installation and restart of the S3 service, you can proceed to **"Completing the Set-Up of the Metadata Search Feature"** (page 37).

If the installer reports an HyperStore Search installation error:

- Double-check the accuracy of the entries in *survey_aux.csv* file in the installation staging directory.
- Confirm that internet connectivity from the HyperStore Search auxiliary nodes is healthy.

5.2.5. Completing the Set-Up of the Metadata Search Feature

After successfully installing the HyperStore Search Service software you can continue the set-up for the Metadata Search feature by performing these tasks:

- Use the CMC to configure alerts based on HyperStore Search Service status, if desired
- Use the CMC to enable Metadata Search for one or more storage policies
- Trigger the search indexing of already-existing objects in buckets that use those storage policies, if desired

For details see **Setting Up Service Features -> Object Metadata and Search -> Preparing the Metadata Search Feature** in either the *HyperStore Administrator's Guide* or the CMC's online Help.

5.3. Installing the HyperStore File Service

In the current release of HyperStore, a Cloudian Professional Services engagement is recommended for installing the HyperStore File Service. For more information about HyperStore File Service pricing and installation, contact your Cloudian representative.

This page left intentionally blank

Chapter 6. HyperStore Installation Reference

This section of the installation documentation provides reference information that you may find useful in some installation scenarios and circumstances.

6.1. Installation Troubleshooting

6.1.1. Installation Logs

When you run the HyperStore installer it generates the following logs that may be helpful for troubleshooting installation problems:

On the Configuration Master node (on which you're running the install script):

- *<installation-staging-directory>/cloudian-installation.log*
- */var/log/puppetserver/puppetserver.log*

On each Configuration Agent node (each node on which you're installing HyperStore):

- */tmp/puppet_agent.log*

Scanning these logs for error or warning messages should help you identify the stage at which the installation encountered a problem, and the nature of the problem. This information can further your own troubleshooting efforts, and also can help Cloudian Support pinpoint the problem in the event that you need assistance from Support.

Note When you use *system_setup.sh* to prepare your nodes for HyperStore installation, that tool writes its logging output to *system_setup.sh.log*, in the same directory as the *system_setup.sh* tool is located (typically your installation staging directory).

6.1.2. Debug Mode

Another potentially useful source of troubleshooting information is to run the installer in debug mode. In the installation staging directory:

```
# ./cloudianInstall.sh -d
```

For example, if you encounter an error while running the installer in regular (non-debug) mode, you can exit the installer menu and then launch the installer again in debug mode. You can then either re-execute the installation starting from the beginning, or re-execute the installation starting from the step that had previously failed. If you had partially run the installation, then when you subsequently select Install Cloudian HyperStore at the main menu a sub-menu will display to let you choose from among several installation tasks to run again.

When run in debug mode, the installer will write highly granular messages to both the console and the installation log (*cloudian-installation.log*).

6.1.3. Specific Issues

ISSUE: You encounter the following warnings:

```
Warning: Could not retrieve fact fqdn
Warning: Host is missing hostname and/or domain: cloudian-singlenode
```

Solution

As suggested by the warning messages, the domain part is missing for the host named "cloudian-singlenode". To resolve this edit the */etc/hosts* file or */etc/resolv.conf* file.

1. Edit the */etc/hosts* file and make sure the following entry exists:

   ```
   Ip-address   cloudian-singlenode.MyDomain.Com  cloudian-singlenode
   ```

   - *Ip-address* should be replaced with host's real IP address
   - *MyDomain.Com* should be replaced with your domain name of choice.

2. Edit the */etc/resolv.conf* file and make sure the following entry exists:

   ```
   Domain MyDomain.Com
   ```

   MyDomain.Com should be replaced with your domain name of choice.

Verify that the *facter fqdn* and *hostname –f* commands output ' cloudian-singlenode.MyDomain.Com' to the console.

ISSUE: Puppet is unable to propagate configuration settings from the Configuration Master node to the Configuration Agent nodes, and in the *puppet_agent.log* and/or *puppet_server.log* you see errors indicating certificate problems or access failures.

Solution

Try going to the installer's "Advanced Options" sub-menu and executing task [**h**] — "Remove Existing Puppet SSL Certificates". Then go back to the main menu and choose the appropriate action below, depending on what you were doing when you encounted the Puppet run failure:

- If you are doing the initial installation of your HyperStore cluster, choose "Install Cloudian HyperStore", then execute task "Install Packages and Configure Nodes [includes Run Puppet]".
- If you are performing post-installation configuration tasks, choose "Cluster Management", then execute task "Push Configuration Settings to the Cluster [Run Puppet]".

ISSUE: While working with the installation script, you get a console message indicating that Puppet access is locked.

Solution

The Puppet process can sometimes end up left in a "locked" state if a Puppet run is interrupted, such as by a Ctrl-<c> command or a host shutdown.

To unlock Puppet, go to the installer's "Advanced Options" sub-menu and execute task [**j**] — "Remove Puppet Access Lock". Then go back to the main menu and choose the appropriate Puppet-running action below, depending on what you were doing when you encountered the Puppet lock error:

- If you are doing the initial installation of your HyperStore cluster, choose "Install Cloudian HyperStore", then execute task "Install Packages and Configure Nodes [includes Run Puppet]".
- If you are performing post-installation configuration tasks, choose "Cluster Management", then execute task "Push Configuration Settings to the Cluster [Run Puppet]".

6.2. HyperStore Listening Ports

The HyperStore system uses the listening ports specified in the table below. Only the service ports for the CMC, S3, IAM, SQS, and Admin services -- the port numbers in italics in the "Listening Port" column -- should

be open to traffic originating from outside the HyperStore system. All other ports must be closed to traffic from outside the system, for system security.

Each HyperStore node includes a built-in HyperStore Firewall that implements port restrictions appropriate to a HyperStore cluster. The HyperStore Firewall is disabled by default in HyperStore systems that were originally installed as a version older than 7.2; and enabled by default in HyperStore systems that originally installed as version 7.2 or newer. You can enable/disable the firewall on all HyperStore nodes by using the installer's Advanced Configuration Options. For instructions see "HyperStore Firewall" in the ***Cloudian HyperStore Administrator's Guide***.

Note If you are installing HyperStore across **multiple data centers and/or multiple service regions**, the HyperStore nodes in each data center and region will need to be able to communicate with the HyperStore nodes in the other data centers and regions. This includes services that listen on the internal interface (such as Cassandra, the HyperStore Service, and Redis). Therefore you will need to configure your networking so that the internal networks in each data center and region are connected to each other (for example, by using a VPN).

Service	Listening Port	Interface(s) Bound To	Purpose
Cloudian Management Console (CMC)	*8888*	All	Requests from administrators' or end users' browsers over HTTP
	8443	All	Requests from administrators' or end users' browsers over HTTPS
S3 Service	*80*	All	Requests from the CMC or other S3 client applications over HTTP
	443	All	Requests from the CMC or other S3 client applications over HTTPS
	81	All	Requests relayed by an HAProxy load balancer using the PROXY Protocol (if enabled by configuration; see *s3_proxy_protocol_enabled* in *common.csv*)
	4431	All	Requests relayed by an HAProxy load balancer using the PROXY Protocol with SSL (if enabled by configuration)
	19080	Internal	JMX access
IAM Service and STS Service	*16080*	All	Requests from the CMC or other Identity and Access Management (IAM) or Security Token Service (STS) clients over HTTP

Service	Listening Port	Interface(s) Bound To	Purpose
			Note In the current HyperStore release, the STS Service uses the same listening ports as the IAM Service.
	16443	All	Requests from the CMC or other IAM or STS clients over HTTPS
	19084	Internal	JMX access
SQS Service	*18090*	All	Requests from Simple Queue Service (SQS) clients over HTTP
	18443	All	Requests from SQS clients over HTTPS (this is not supported in the current HyperStore release but will be in the future)
	19085	Internal	JMX access
Admin Service	*18081*	All	Requests from the CMC or other Admin API clients over HTTP
	19443	All	Requests from the CMC or other Admin API clients over HTTPS (Note: The CMC by default uses HTTPS to access the Admin Service) **IMPORTANT !** The Admin Service is intended to be accessed only by the CMC and by system administrators using other types of clients (such as cURL). Do not expose the Admin Service to a public network.
	19081	Internal	JMX access
Redis Monitor	9078	Internal	Communication between primary and backup Redis Monitor instances
	19083	Internal	JMX access

Service	Listening Port	Interface(s) Bound To	Purpose
HyperStore Service	19090	Internal	Data operation requests from the S3 Service
	19050	Internal	Communication between HyperStore Service instances
	19082	Internal	JMX access
Credentials DB and QoS DB (Redis)	6379	Internal	Requests to the Credentials DB from the S3 Service, HyperStore Service, or Admin Service; and communication between Credentials DB nodes
	6380	Internal	Requests to the QoS DB from the S3 Service, HyperStore Service, or Admin Service; and communication between QoS DB nodes
Metadata DB (Cassandra)	9042	Internal	Data operations requests from the S3 Service, HyperStore Service, or Admin Service, using CQL protocol
	9160	Internal	Data operations requests from the S3 Service, HyperStore Service, or Admin Service, using Thrift protocol
	7000	Internal	Communication between Cassandra instances
	7199	Internal	JMX access
Cloudian Monitoring Agent	19070	Internal	Requests from the Cloudian Monitoring Data Collector
Configuration Master	8140	Internal	On your Configuration Master node this port will service incoming requests from Puppet agents on your other HyperStore nodes
	4505	Internal	On your Configuration Master node this is the port to which Salt agents ("minions") establish a persistent connection so that the Master can publish to the minions.
	4506	Internal	Salt minions connect to this port on the Configuration Master as needed to send results to the Master, and to request files and

Service	Listening Port	Interface(s) Bound To	Purpose
			minion-specific data values.
SSH	22	All	The HyperStore installer accesses this SSH port on each node on which you are installing HyperStore software (during initial cluster install or if you subsequently expand your cluster)
NTP	123	All	NTP port for time synchronization between nodes
Echo	7	Internal	The Cloudian Monitoring Data Collector uses Echo (port 7) to check whether each node is reachable.
HyperIQ	9999	All	If you use Cloudian HyperIQ to monitor your HyperStore system, HyperIQ accesses port 9999 on each HyperStore node

6.3. Outbound Internet Access

The HyperStore installation process does not require outbound internet access. However, the following HyperStore features do access the internet once the system is in operation; and HyperStore does need access to NTP server(s) during the installation (see the "Pre-Configured ntpd" bullet point below). If you use forward proxying in your environment, after HyperStore installation you may want to set up forward proxying to support these HyperStore features:

- **Smart Support** — The Smart Support feature (also known as "Phone Home") securely transmits HyperStore daily diagnostic information to Cloudian Support over the internet. HyperStore supports configuring this feature to use an explicit forward proxy for its outbound internet access (after installation, the relevant settings in *common.csv* are *phonehome_proxy_host* and the other *phonehome_proxy_** settings that follow after it). To use a forward proxy with this feature you should configure your forward proxy to support access to **.s3-support.cloudian.com* (that is, to any sub-domain of *s3-support.cloudian.com*).
- **Auto-Tiering and Cross-Region Replication** — If you want to use either the auto-tiering feature or the cross-region replication feature (CRR), the S3 Service running on each of your HyperStore nodes requires outbound internet access. These features do not support configuring an explicit forward proxy, but you can use transparent forward proxying if you wish. (Setting up transparent forward proxying is outside the scope of this documentation.) For more information on these features see the "Auto-Tering Feature Overview" and "Cross-Region Replication Overview" sections in the ***Cloudian HyperStore Administrator's Guide***.
- **Pre-Configured ntpd** — Accurate, synchronized time across the cluster is vital to HyperStore service. In of your HyperStore data centers four of your HyperStore nodes are automatically configured to act as internal NTP servers. (If a HyperStore data center has only four or fewer nodes, then all the nodes in the data center are configured as internal NTP servers.) These internal NTP servers are configured to connect to external NTP servers — by default the public servers from the *pool.ntp.org* project. In order to

connect to the external NTP servers, the internal NTP servers must be allowed outbound internet access. This feature does not support configuring an explicit forward proxy, but you can use transparent forward proxying if you wish. (Setting up transparent forward proxying is outside the scope of this documentation.)

IMPORTANT ! If you do not allow HyperStore hosts to have outbound connectivity to the internet, then during the interactive installation process -- when you are prompted to specify the NTP servers that HyperStore hosts should connect to -- you must specify NTP servers within your environment, rather than the public NTP servers that HyperStore connects to by default. If HyperStore hosts cannot connect to any NTP servers, the installation will fail.

After HyperStore installation, to see which of your HyperStore nodes are internal NTP servers, log into the CMC and go to **Cluster → Cluster Config → Cluster Information**. On that CMC page you can also see your configured list of external NTP servers.

For more information on HyperStore's NTP set-up, see the "NTP Automatic Set-Up" section in the ***Cloudian HyperStore Administrator's Guide***.

6.3.1. Multi-DC Considerations

If you are installing HyperStore across multiple data centers and/or multiple service regions, the HyperStore nodes in each data center and region will need to be able to communicate with the HyperStore nodes in the other data centers and regions. This includes services that listen on the internal interface (such as Cassandra, the HyperStore Service, and Redis). Therefore you will need to configure your networking so that the internal networks in each data center and region are connected to each other (for example, by using a VPN). See **"HyperStore Listening Ports"** (page 41) for HyperStore requirements regarding listening port access.

6.4. File System Requirements

Subjects covered in this section:

Cloudian recommends that you use the HyperStore *system_setup.sh* tool to configure the disks and mount points on your HyperStore nodes, as described in **"Configuring Network Interfaces, Time Zone, and Data Disks"** (page 21). The tool is part of the HyperStore product package (when you extract the *.bin* file).

If you do not use the system setup tool for disk setup, use the information below to make sure that your hosts meet HyperStore file system requirements.

6.4.1. OS/Metadata Drives and Data Drives

Although it's possible to install HyperStore on a host with just a single hard drive, for a rigorous evaluation or for production environments each host should have multiple drives (see **"Host Hardware and OS Requirements"** (page 15)). On host machines with multiple hard drives:

- HyperStore will by default use the drive that the OS is on for storing system metadata (in the Metadata DB, the Credentials DB, and the QoS DB). Cloudian recommends that you dedicate two drives to the OS and system metadata in a RAID-1 mirroring configuration. Preferably the OS/metadata drives should be SSDs.
- You must format all other available hard drives with *ext4* file systems mounted on raw disks. These drives will be used for storing S3 object data. RAID is not necessary on the S3 object data drives.

For example, on a machine with 2 SSDs and 12 HDDs:

- Mirror the OS on the two SSDs. For more detailed recommendations for partitioning these disks see **"Partitioning of Disks Used for the OS and Metadata Storage"** (page 16).
- Format each of the 12 HDDs with *ext4* file systems and configure mount points such as */cloudian1*, */cloudian2*, */cloudian3* and so on.

> **Note** On the HDDs for storing object data, HyperStore **does not support** XFS file systems; VirtIO disks; Logical Volume Manager (LVM); or Multipathing. For questions regarding these unsupported technologies, contact Cloudian Support:

6.4.2. Mount Point Naming Guidelines

If you are installing HyperStore on multiple hosts that each have multiple disks for object data storage, use the same mount point naming scheme on each of your hosts. If all your hosts have the same number of disks, then they should all have the identical set of mount points for HyperStore object storage. For example, if each host has 12 disks for object storage, then on all your hosts you could name the mount points */cloudian1*, */cloudian2*, */cloudian3*, and so on up through */cloudian12*.

If in your installation cluster some hosts have more disks than others, use as much overlap in mount point naming as possible. For example, suppose that most of your hosts have 10 disks for storing object data while one host has 12 disks. In this scenario, all of the hosts can have mount points */cloudian1*, */cloudian2*, */cloudian3*, and so on up through */cloudian10*, while the one larger host has those same mount points plus also */cloudian11* and */cloudian12*.

> **Note** Although uniformity of mount point naming across nodes (to the extent possible) is desirable for simplicity's sake, the HyperStore installation does support a way to accommodate differences in the number or names mount points across nodes -- this is described in **"A Data Directory Mount Point List (fslist.txt) Is Required"** (page 49)..

6.4.3. Option for Putting the Metadata DB on Dedicated Drives Rather Than the OS Drives

Regarding the Metadata DB (built on Cassandra), another supported configuration is to put your Cassandra data on dedicated drives, rather than on the OS drives. In this case you would have:

- OS drives in RAID-1 configuration. The Credentials DB and QoS DB will also be written to these drives.
- Cassandra drives in RAID-1 configuration. On these drives will be written Cassandra data and also the Cassandra commit log.

  Note You must create a Cassandra data directory named as *<mountpoint>/cassandra* (for example *cassandradb/cassandra*) and a Cassandra commit log directory named as *<mountpoint>/cassandra_commit* (for example *cassandradb/cassandra_commit*).

- Multiple drives for S3 object data (with mount points for example */cloudian1*, */cloudian2*, */cloudian3* and so on), with no need for RAID protection.

6.4.4. You Must Use UUIDs in fstab

In your *fstab* file, **you must use UUIDs** to identify the devices to which you will mount HyperStore S3 object data directories. Do not use device names or LABELs.

If you are not using UUIDs in *fstab* currently, follow the instructions below to modify your *fstab* so that it uses UUIDs for the devices to which you will mount S3 object data directories (you do not need to do this for the OS/-metadata mount points).

As *root*, do the following:

1. Check whether your *fstab* is currently using UUIDs for your S3 object data drives. In the example below, there are two S3 object data drives and they are currently identified by device name, not by UUID.

   ```
   # cat /etc/fstab
   ...
   ...
   /dev/sdb1 /cloudian1  ext4  rw,noatime,barrier=0,data=ordered,errors=remount-ro 0 1
   /dev/sdc1 /cloudian2  ext4  rw,noatime,barrier=0,data=ordered,errors=remount-ro 0 1
   ```

2. Back up your existing fstab file:

   ```
   # cp /etc/fstab /etc/fstab.backup.<today's date>
   ```

3. Retrieve the UUIDs for your devices by using the *blkid* command.

   ```
   # blkid
   ...
   ...
   /dev/sdb1: UUID="a6fed29c-97a0-4636-afa9-9ba23e1319b4" TYPE="ext4"
   /dev/sdc1: UUID="rP38Ux-3wzO-sP3Y-2CoD-2TDU-fjpO-ffPFZV" TYPE="ext4"
   ```

4. Open *fstab* in an editor.
5. For each device that you are using for S3 object storage, replace the device name with *UUID="<UUID>"*, copying the device's UUID from the *blkid* response in the previous step. For example:

   ```
   # Original line

   /dev/sdb1 /cloudian1 ext4 rw,noatime,barrier=0,data=ordered,errors=remount-ro  0 1

   # Revised line

   UUID="a6fed29c-97a0-4636-afa9-9ba23e1319b4" /cloudian1  ext4  rw,noatime,barrier=0,
   data=ordered,errors=remount-ro   0 1
   ```

6. After editing *fstab* so that each device on which you will store S3 data is identified by a UUID, save your changes and close the *fstab* file.
7. Remount the host's file systems:

```
# mount -a
```

Repeat this process for **each host on which you will install HyperStore**.

6.4.5. A Data Directory Mount Point List (fslist.txt) Is Required

If you do not use the HyperStore *system_setup.sh* script to configure the data disks and mount points on your nodes, you must manually create a data directory mount point list file and place it in your installation staging directory on the Configuration Master node, as described below.

> **Note** If you use the *system_setup.sh* script to configure the disks and mount points on your nodes, the script creates the needed mount point list files automatically and you can ignore the instructions below.

If all your nodes have the same data mount points -- for example if all nodes have as their data mount points */cloudian1*, */cloudian2*, and so on through */cloudian12* -- you only need to create one mount point list file. If some nodes have a different set of mount points than do other nodes -- for example if some nodes have more data disks than other nodes -- you will need to create a default mount point list file and also a node-specific mount point list file for each node that differs from the default.

In your installation staging directory create a file named *fslist.txt* and in the file enter one line for each of your S3 data directory mount points, with each line using the format below.

```
<deviceName> <mountPoint>
```

Example of a properly formatted file (truncated):

```
/dev/sdc1 /cloudian1
/dev/sdd1 /cloudian2
...
```

> **Note** Use device names in your *fslist.txt* file, not UUIDs.

Optionally, you can also include an entry for the Cassandra data directory and an entry for the Cassandra commit log directory, if you do not want this data to be put on the same device as the operating system (see **"Option for Putting the Metadata DB on Dedicated Drives Rather Than the OS Drives"** (page 47)). If you do not specify these Cassandra directory paths in *fslist.txt*, then by default the system automatically puts Cassandra data and commit log directories on the same device on which the operating system resides.

Do not use symbolic links when specifying your mount points. The HyperStore system does not support symbolic links for data directories.

If some of your hosts have data directory mount point lists that differ from the cluster default, in the installation staging directory create a *<hostname>_fslist.txt* file for each such host. For example, along with the default *fslist.txt* file that specifies the mount points that most of your hosts use, you could also have a *cloudian-node11_fslist.txt* file and a *cloudian-node12_fslist.txt* file that specify mount points for two non-standard nodes that have hostnames *cloudian-node11* and *cloudian-node12*.

6.4.6. Reducing Reserved Space to 0% for HyperStore Data Disks

By default Linux systems reserve 5% of file system space for root user and system services. On modern large-capacity disks this can be a waste of a considerable amount of storage space. Cloudian recommends that you set the reserved space to 0% for each drive on which you will store HyperStore object data (S3 object data).

For each HyperStore data drive do the following.

```
### Check current "Reserved block count":

# tune2fs -l <device>

### Set Reserved block count to 0%:

# tune2fs -m 0 <device>

### For example:

# tune2fs -m 0 /dev/sdc1
```

6.5. survey.csv (Cluster Survey File)

During the **"Installing HyperStore Prerequisites"** (page 18) task you use the *system_setup.sh* script to create a cluster survey file which by default is named *survey.csv*. This file resides in your installation staging directory for the life of your HyperStore system. The survey file is automatically updated by the system if you subsequently use the CMC to add more nodes to your cluster; and it is automatically copied to your new installation staging directory when you execute a HyperStore version upgrade.

> **Note** The survey file must be kept in the installation staging directory, not in a different directory. Do not delete or move the survey file.

The survey file contains one line for each HyperStore host in your cluster (including the Configuration Master host), with each line using the format below.

```
<regionname>,<hostname>,<ip4-address>,<datacenter-name>,<rack-name>[,<internal-interface>]
```

- *<regionname>* — HyperStore service region in which the host is located. The HyperStore system supports having multiple service regions with each region having its own independent storage cluster and S3 object inventory, and with S3 application users able to choose a storage region when they create storage buckets. Even if you will have only one region you must give it a name. The maximum allowed length is 52 characters. The only allowed character types are lower case ASCII alphanumerical characters and dashes (a-z0-9 and dashes). Do not include the string "s3" in the region name. Make sure the region name matches the region string that you use in your S3 endpoints in your **"DNS Set-Up"** (page 10). For more information on regions see "Nodes, Data Centers, and Regions" in the Introduction section of the ***Cloudian HyperStore Administrator's Guide***.
- *<hostname>* — Short hostname of the host (as would be returned if you ran the *hostname -s* command on the host). This must be the node's short hostname, not an FQDN.

  > **Note** Do not use the same short hostname for more than one node in your entire HyperStore system. Each node must have a unique short hostname within your entire HyperStore system, even

in the case of nodes in different data centers or service regions that have different domains. For example, in your HyperStore system do not have two nodes with the same short hostname *vega* for which the FQDN of one is *vega.east.com* and the FQDN of the other is *vega.west.com*.

- *<ip4-address>* — IP address (v4) that the hostname resolves to. Do not use IPv6. This should be the IP address associated with the host's default, external interface -- not an internal interface.
- *<datacenter-name>* — Name of the data center in which the host machine is located. The maximum allowed length is 256 characters. The only allowed character types are ASCII alphanumerical characters and dashes (A-Za-z0-9 and dashes).
- *<rack-name>* — Name of the server rack in which the host machine is located. The maximum allowed length is 256 characters. The only allowed character types are ASCII alphanumerical characters and dashes (A-Za-z0-9 and dashes).

  Note Within a data center, use the same "rack name" for all of the nodes, even if some nodes are on different physical racks than others. For example, if you have just one data center, all the nodes must use the same rack name. And if you have two data centers named DC1 and DC2, all the nodes in DC1 must use the same rack name as the other nodes in DC1; and all the nodes in DC2 must use the same rack name as the other nodes in DC2.

- *[<internal-interface>]* — Use this field only for hosts that will use a different network interface for internal cluster traffic than the rest of the hosts in the cluster do. For example, if most of your hosts will use "eth1" for internal cluster traffic, but two of your hosts will use "eth2" instead, use this field to specify "eth2" for each of those two hosts, and leave this field empty for the rest of the hosts in your survey file. (Later in the installation procedure you will have the opportunity to specify the default internal interface for the hosts in your cluster -- the internal interface used by all hosts for which you do not specify the *internal-interface* field in your survey file.) If all of your hosts use the same internal network interface — for example if all hosts use "eth1" for internal network traffic — then leave this field empty for all hosts in the survey file.

  Note Cassandra, Redis, and the HyperStore Service are among the services that will utilize the internal interface for intra-cluster communications.

The example survey file below is for a single-node HyperStore installation:

```
region1,arcturus,65.10.2.1,DC1,RAC1
```

This second example survey file is for a three-node HyperStore cluster with just one service region, one data center, and one rack:

```
tokyo,cloudian-vm7,65.10.1.33,DC1,RAC1
tokyo,cloudian-vm8,65.10.1.34,DC1,RAC1
tokyo,cloudian-vm9,65.10.1.35,DC1,RAC1
```

This third example survey file below is for a HyperStore installation that spans two regions, with the first region comprising two data centers and the second region comprising just one data center. Two of the hosts use a different network interface for internal network traffic than all the other hosts do.

```
boston,hyperstore1,65.1.0.1,DC1,RAC1
boston,hyperstore2,65.1.0.2,DC1,RAC1
boston,hyperstore3,65.1.0.3,DC1,RAC1
```

```
boston,hyperstore4,66.2.0.1,DC2,RAC1
boston,hyperstore5,66.2.0.2,DC2,RAC1
chicago,hyperstore6,68.3.0.1,DC3,RAC1
chicago,hyperstore7,68.3.0.2,DC3,RAC1
chicago,hyperstore8,68.3.2.1,DC3,RAC1,eth2
chicago,hyperstore9,68.3.2.2,DC3,RAC1,eth2
```

6.6. cloudianInstall.sh

The *cloudianInstall.sh* tool (also known as "the installer") serves several purposes including:

- Installation of a HyperStore cluster (for detail see **"Installing a New HyperStore System"** (page 25))
- Implementing advanced, semi-automated system configuration changes (for detail see **"Installer Advanced Configuration Options"** (page 55))
- Pushing configuration file edits to the cluster and restarting services to apply the changes(for detail see "Pushing Configuration File Edits to the Cluster and Restarting Services" in the ***Cloudian HyperStore Administrator's Guide***)

The *cloudianInstall.sh* tool is in your installation staging directory on your Configuration Master node. To perform advanced configurations, or to push configuration file changes to the system and restart services, you would launch the tool simply like this, without using additional command line options:

```
# ./cloudianInstall.sh
```

6.6.1. Command Line Options When Using cloudianInstall.sh for HyperStore Cluster Installation

To perform a HyperStore cluster installation you typically would launch the script either like this:

```
# ./cloudianInstall.sh -s survey.csv
```

Or like this if you are not using your DNS environment to resolve HyperStore service endpoints and you want to use the bundled tool *dnsmasq* instead (which is not appropriate for production systems):

```
# ./cloudianInstall.sh -s survey.csv configure-dnsmasq
```

However the script does support additional command line options. The syntax is as follows:

```
# ./cloudianInstall.sh [-s <survey-filename>] [-k <ssh-private-key-filename>]
[-d] [-h] [no-hosts] [configure-dnsmasq] [no-firewall] [force] [uninstall]
```

Note If you use multiple options, on the command line place options that start with a "-" (such as *-s <survey-filename>* or *-d*) before options that do not (such as *no-hosts* or *configure-dnsmasq*).

*If you are using the **HyperStore Shell***

If you are using the HyperStore Shell (HSH) as a Trusted user, from any directory on the Configuration Master node you can launch the installer with this command:

```
$ hspkg install
```

The installer's options are the same regardless of whether it is launched from the HSH command line or the OS command line.

Note After using the installer, exit the installer when you're done. Do not leave it running. Certain automated system tasks invoke the installer and cannot do so if it is already running.

The supported command line options are:

- *[-s <survey-filename>]* — Name of your cluster survey file (including the full path to the file). If you do not specify the survey file name argument, the script will prompt you for the file name during installation.
- *[-k <ssh-private-key-filename>]* — The Configuration Master employs SSH for secure communication with the rest of your HyperStore installation nodes. By default the install script automatically creates an SSH key pair for this purpose. But if instead you would prefer to use your own existing SSH key pair for this purpose, you can use the installer's *-k <ssh-private-key-filename>* option to specify the name of the private key file (including the full path to the file). When you run the install script it will copy the private key and corresponding public key to the installation staging directory, and in the staging directory the key file will be renamed to *cloudian-installation-key*. Then from the staging directory, the public key file *cloudian-installation-key.pub* will be copied to each node on which you are installing HyperStore.
- *[-d]* — Turn on debugging output.
- *[-h]* — Display usage information for the install tool. This option causes the tool to print a usage message and exit.

  Note This usage information mentions more command line options than are described here in this Help topic. This is because the usage information includes installer options that are meant for HyperStore internal system use, such as options that are invoked by the CMC when you use the CMC to add nodes to your cluster or remove nodes from your cluster. You should perform such operations through the CMC, not directly through the installer. The CMC implements automations and sanity checks beyond what is provided by the install script alone.

- *[no-hosts]* — Use this option if you do not want the install tool to append entries for each HyperStore host on to the */etc/hosts* file of each of the other HyperStore hosts. By default the tool appends to these files so that each host is resolvable to the other hosts by way of the */etc/hosts* files.
- *[configure-dnsmasq]* — Use this option if you want the install tool to install and configure **dnsmasq**, a lightweight utility that can provide domain resolution services for testing a small HyperStore system. If you use this option the installer installs *dnsmasq* and automatically configures it for resolution of HyperStore service domains. If you did not create DNS entries for HyperStore service domains as described in **"DNS Set-Up"** (page 10), then you must use the *configure-dnsmasq* option in order for the system to be functional when you complete installation. Note that using *dnsmasq* is not appropriate in a production environment.

  Note If you do not have the installer install *dnsmasq* during HyperStore installation, and then later you decide that you do want to use *dnsmasq* for your already installed and running HyperStore system, do not use the *configure-dnsmasq* command line option when you re-launch the installer. Instead, re-launch the installer with no options and use the Installer Advanced Configuration Options menu to enable *dnsmasq* for your system.

- [*no-firewall*] — If this option is used, the HyperStore firewall will **not** be enabled upon HyperStore installation. By default the HyperStore firewall will be enabled upon completion of a fresh HyperStore installation. For more information about the HyperStore firewall see the "HyperStore Firewall" section in the

Cloudian HyperStore Administrator's Guide.

- *[force]* — By default the installer performs certain prerequisite checks on each node on which you are installing HyperStore and aborts the installation if any of your nodes fails a check. By contrast, if you use the *force* option when you launch the installer, the installer will output warning messages to the terminal if one or more nodes fails a prerequisite check but the installation will continue rather than aborting. The prerequisite checks that this feature applies to are:
  - CPU has minimum of 8 cores
  - RAM is at least 128GB
  - System Architecture is x86 64-bit
  - SELinux is disabled
  - firewalld is disabled
  - iptables is not running

> **Note** If you specify the *force* option when running the installer, the *force*option will "stick" and will be used automatically for any subsequent times the installer is run to install additional nodes (such as when you do an "Add Node" operation via the Cloudian Management Console, which invokes the installer in the background). To turn the *force*option off so that it is no longer automatically used when the installer is run to add more nodes, launch the installer and go to the Advance Configuration Options. Then choose option **t** for **Configure force behavior** and follow the prompts.

> **Note** Even if the *force* option is used the installer will abort if it detects an error condition on the host that will prevent successful installation.

- *[uninstall]* — If you use this option when launching the installer, the installer main menu will include an additional menu item -- "Uninstall Cloudian HyperStore".

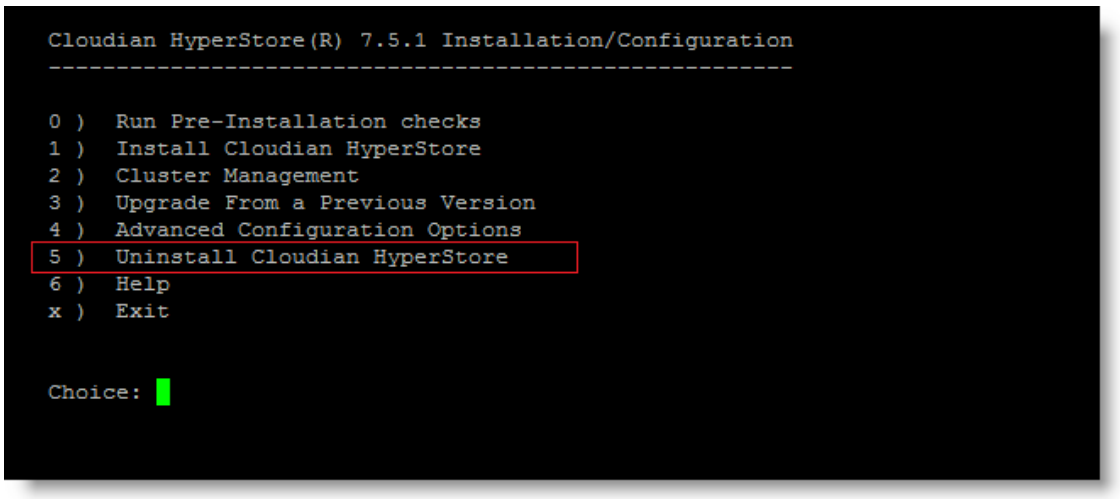

Use this menu option only if you want to **delete the entire HyperStore system, on all nodes, including any metadata and object data** stored in the system. You may want to use this Uninstall Cloudian HyperStore option, for example, after completing a test of HyperStore -- if you do not want to retain the test system.

IMPORTANT ! Do not use this option to uninstall a single node from a HyperStore system that you want to retain (such as a live production system). For instructions on removing a node from a HyperStore system see the "Removing a Node" section in the ***Cloudian HyperStore Administrator's Guide***.

6.7. system_setup.sh

The *system_setup.sh* tool is for setting up nodes on which you will install HyperStore software, either during initial cluster installation or during cluster expansion. For basic information about using *system_setup.sh*, change into the installation staging directory and run the following command:

```
# ./system_setup.sh --help
```

6.8. Installer Advanced Configuration Options

The HyperStore installation tool supports several types of advanced system configurations which can be implemented at any time after initial installation of the system. To access the advanced configuration options, on the Configuration Master node change into your installation staging directory (*/opt/cloudian-staging/7.5.2*) and launch the installer.

```
# ./cloudianInstall.sh
```

*If you are using the **HyperStore Shell***

If you are using the HyperStore Shell (HSH) as a Trusted user, from any directory on the Configuration Master node you can launch the installer with this command:

```
$ hspkg install
```

Once launched, the installer's menu options (such as referenced in the steps below) are the same regardless of whether it was launched from the HSH command line or the OS command line.

At the installer main menu's Choice prompt enter **4** for Advanced Configuration Options.

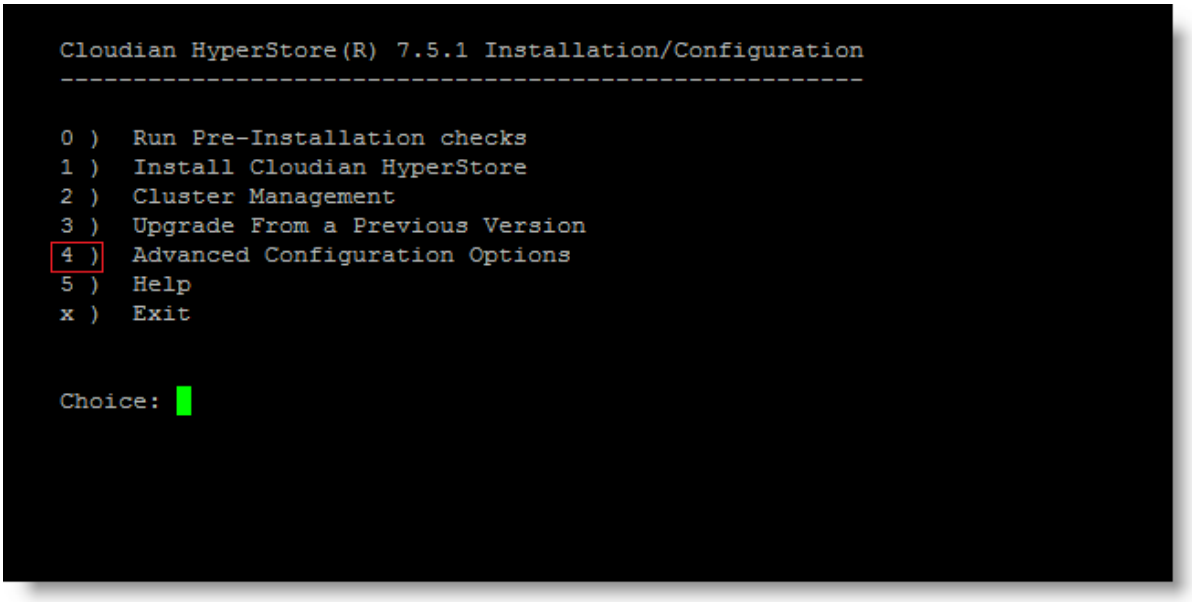

This opens the "Advanced Configuration Options" sub-menu.

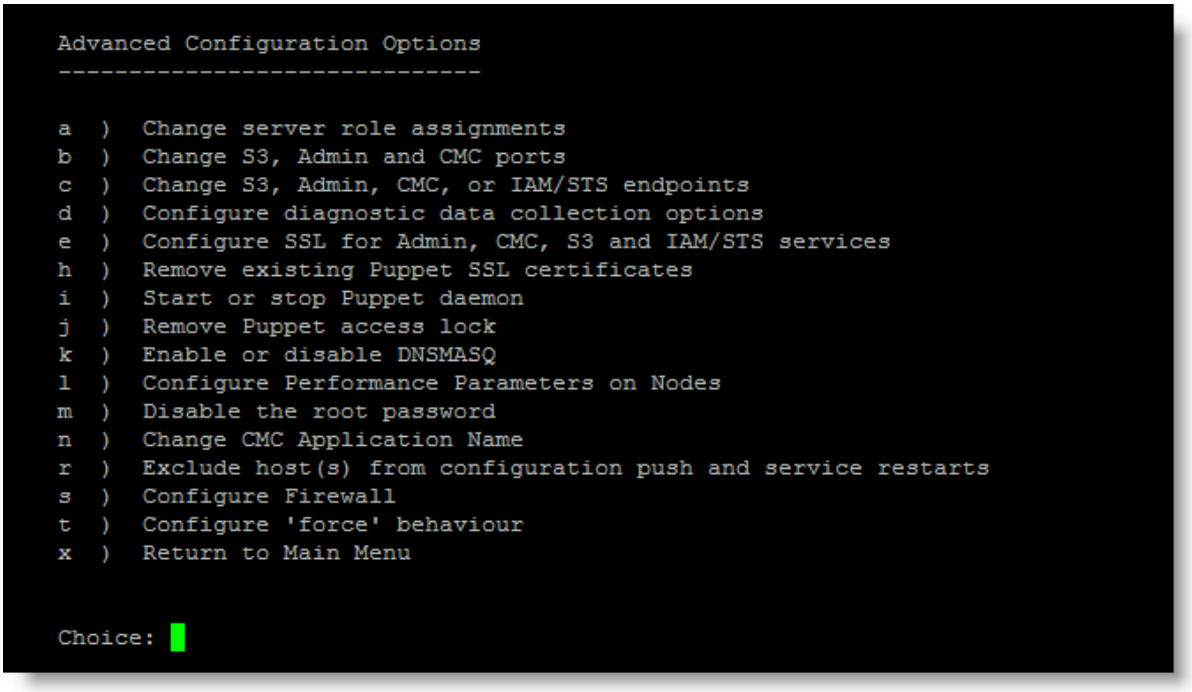

From this menu you can choose the type of configuration change that you want to make and then proceed through the interactive prompts to specify your desired settings.

For information about each of these options, see the "Reference -> Configuration Settings -> Installer Advanced Configuration Options" section of the ***Cloudian HyperStore Administrator's Guide***.

Note As a best practice, you should complete basic HyperStore installation first and confirm that things are working properly (by running the installer's Validation Tests, under the "Cluster Management" menu) before you consider using the installer's advanced configuration options.