

HyperStore Installation Guide

Version 6.2



Contents

Chapter 1. HyperStore Installation Introduction	1
Chapter 2. Preparing Your Environment	3
2.1. DNS Set-Up	4
2.2. Load Balancing	6
Chapter 3. Preparing Your Nodes	9
3.1. Host Hardware Requirements	10
3.2. OS, Network Interfaces, and Disks	10
3.3. Prerequisite Packages Installation	15
3.4. Pre-Install Check	18
Chapter 4. Installing a New HyperStore System	21
Chapter 5. Getting Started with a Newly Installed HyperStore System	25
Chapter 6. Upgrading an Existing HyperStore System	29
6.1. Before Upgrading Your System	29
6.2. Upgrading Your System	30
6.3. After Upgrading Your System	31
Chapter 7. HyperStore Installation Reference	33
7.1. Installation Troubleshooting	34
7.2. HyperStore Listening Ports	35
7.3. Outbound Internet Access	38
7.4. File System Requirements	38
7.5. Data Directory Mount Point List (fslist.txt)	41
7.6. Cluster Survey File (survey.csv)	42
7.7. cloudianInstall.sh Command Line Options	43

Chapter 1. HyperStore Installation Introduction

This installation documentation is for Clouidian HyperStore^R version 6.2. For more information about the 6.2 release, see "What's New in HyperStore 6.2" in the *HyperStore Administrator's Guide*.

The main part of this installation documentation describes how to do a **fresh installation** of Clouidian HyperStore 6.2.

If you are **upgrading** to 6.2 from an existing HyperStore system, you can go directly to "**Upgrading an Existing HyperStore System**" (page 29).

If you do not yet have the HyperStore 6.2 package, you can obtain it from the Clouidian FTP site *ftp.cloudian.com*. You will need a login ID and password (available from Clouidian Support). Once logged into the FTP site, change into the *Clouidian_HyperStore* directory and then into the *cloudian-6.2* sub-directory. If you are installing on to bare host machines, download both the **.iso* file (with which you can install and configure CentOS on your hosts) and the **.bin* file (the HyperStore product package). If you are installing HyperStore on to hosts that already have CentOS or RHEL installed -- or are upgrading an existing HyperStore system -- download only the **.bin* file.

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 by emailing *cloudian-license@cloudian.com*. If you are using the ISO, the ISO package includes a very low capacity license that enables you to install HyperStore and verify basic functionality, prior to acquiring a full evaluation or production license.

This page left intentionally blank

Chapter 2. Preparing Your Environment

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

- **"DNS Set-Up"** (page 4)
- **"Load Balancing"** (page 6)

2.1. DNS Set-Up

For your HyperStore system to be accessible to external clients, you must configure your DNS name servers with entries for the HyperStore service endpoints. **Clouddian, Inc. 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. You can use the *configure-dnsmasq* option when you launch the HyperStore installation script as described later in this document.

The table that follows shows the DNS entries that you must configure on your name servers to resolve HyperStore service endpoints. By default the HyperStore system automatically derives service endpoint values from your organization's domain, which you will supply when you run the HyperStore interactive installer. The table shows the default format of each service endpoint. The default S3 endpoint formats are consistent with the format that Amazon uses for its S3 endpoints.

If you do not want to use the default S3 or CMC or Admin service endpoint formats, the HyperStore system allows you to specify custom endpoint values during the installation. If you intend to create custom endpoints rather than accepting the default endpoints, configure DNS entries to resolve those custom endpoint values rather than the default-formatted endpoint values shown below. Make a note of the custom endpoints for which you have configured DNS entries, so that later you can correctly specify those endpoints when you perform the HyperStore installation.

In a production environment, each of these service endpoints should resolve to the IP addresses of two or more [load balancers](#) (configured for fail-over), with the load balancers in turn distributing request traffic across all the nodes in the cluster. In a multi-region system, the regional S3 endpoints should resolve to local region load balancers which distribute traffic across local region nodes; while the Admin and CMC service endpoints should all resolve to load balancers in the default service region, which in turn distribute traffic to nodes in the default service region. For background information on service regions see "Service Regions Feature Overview" in the "Major Features" section of the HyperStore Administrator's Guide.

Note In an evaluation or testing environment, using round-robin DNS is an acceptable alternative to using load balancers. See the Example that follows the table.

DNS Entry	Default Format and Example	Description
S3 service endpoint (one per service region)	<i>s3-<region>.<your-domain></i> <i>s3-tokyo.enterprise.com</i>	This is the service endpoint to which S3 client applications will submit requests. The <i><region></i> segment indicates the HyperStore service region. You must choose a service region name for your HyperStore installation, even if you intend to have only one service region.

DNS Entry	Default Format and Example	Description
		<p>The region name must be lower case with no dots, dashes, underscores, or spaces. Make sure that the region name you supply when doing the installation matches the region name you used in your DNS configuration.</p> <p>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 <code>s3-tokyo.enterprise.com</code> and <code>s3-osaka.enterprise.com</code>.</p> <p>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.</p> <p>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.</p>
S3 service endpoint wildcard (one per service region)	<code>*.s3-<region>.<your-domain></code> <code>*.s3-tokyo.enterprise.com</code>	This S3 service endpoint wildcard entry is necessary to resolve S3 requests pertaining to a specific storage bucket (which is nearly all S3 requests).
S3 static website service endpoint (one per service region)	<code>s3-website-<region>.<your-domain></code> <code>s3-website-tokyo.enterprise.com</code>	This S3 service endpoint is used for buckets configured as static websites.
S3 static website endpoint wildcard (one per service region)	<code>*.s3-website-<region>.<your-domain></code> <code>*.s3-website-tokyo.enterprise.com</code>	This S3 static website endpoint wildcard entry is necessary to make S3 requests resolvable, for buckets configured as static websites.
Admin Service endpoint (one per entire system)	<code>s3-admin.<your-domain></code> <code>s3-admin.enterprise.com</code>	This is the service endpoint for HyperStore's Admin API. The Cloudbian Management Console accesses this RESTful HTTP API, and you can also access the API directly with a command line tool

DNS Entry	Default Format and Example	Description
		such as <i>cURL</i> or a client application of your own creation.
Cloudian Management Console (CMC) domain (one per entire system)	<i>cmc.<your-domain></i> <i>cmc.enterprise.com</i>	The CMC is HyperStore's web-based console for making S3 requests (such as creating storage buckets or uploading objects) or performing system provisioning and administration tasks.

2.1.0.1. Example

Below is an example of a round-robin DNS configuration for a three-node HyperStore system that has only one service region. The organization's top-level domain is *enterprise.com*. Note that HyperStore uses a peer-to-peer architecture in which the S3 Service, Admin Service, and CMC all run on every HyperStore node, and thus in the example all of the nodes are part of the round-robin for each service endpoint.

```
s3-tokyo.enterprise.com IN A 10.1.1.1
                        10.1.1.2
                        10.1.1.3
*.s3-tokyo.enterprise.com IN A 10.1.1.1
                               10.1.1.2
                               10.1.1.3
s3-website-tokyo.enterprise.com IN A 10.1.1.1
                                     10.1.1.2
                                     10.1.1.3
*.s3-website-tokyo.enterprise.com IN A 10.1.1.1
                                       10.1.1.2
                                       10.1.1.3
s3-admin.enterprise.com IN A 10.1.1.1
                             10.1.1.2
                             10.1.1.3
cmc.enterprise.com IN A 10.1.1.1
                       10.1.1.2
                       10.1.1.3
```

2.2. Load Balancing

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

For high availability it's 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 IP 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 detailed guidance on load balancing set-up, request a copy of the *HyperStore Load Balancing Best Practice Guide* from your Clodian Sales Engineering representative.

Note For a non-production environment round-robin DNS is an acceptable alternative to using load balancers. For more information see "**DNS Set-Up**" (page 4).

This page left intentionally blank

Chapter 3. Preparing Your Nodes

This guide to preparing your nodes for HyperStore software installation covers these topics:

- **"Host Hardware Requirements"** (page 10)
- **"OS, Network Interfaces, and Disks"** (page 10)
- **"Prerequisite Packages Installation"** (page 15)
- **"Pre-Install Check"** (page 18)

3.1. Host Hardware Requirements

The table below shows the recommended and minimum host hardware specifications for a HyperStore system installation.

Recommended for production systems	<ul style="list-style-type: none"> • 1 CPU, 8 cores • 64GB • 2 x 300GB SSD (for RAID-1 mirrored hosting of the OS as well as Cassandra and Redis databases storing system metadata) • 12 x 4TB HDD (for <i>ext4</i> file systems storing object data) (JBOD, no RAID) • 2x10GbE Ports <div style="background-color: #e0f0e0; padding: 10px; margin-top: 10px;"> <p>Note If you plan to use erasure coding for object data storage, 2 CPUs per node is recommended. Also, be aware that the higher your erasure coding <i>m</i> value (such as with $k+m = 9+3$ or $8+4$), the higher the need for Cassandra metadata storage capacity. Consult with your Cloudbian representative to ensure that you have adequate Cassandra storage capacity to support your desired <i>m</i> value.</p> </div>
Minimum for production systems	<ul style="list-style-type: none"> • 1 CPU, 8 cores • 32GB RAM • 2 x 160GB SSD (for RAID-1 mirrored hosting of the OS as well as Cassandra and Redis databases storing system metadata) • 12 x 2TB HDD (for <i>ext4</i> file systems storing object data) (JBOD, no RAID) • 1x1GbE Port
Minimum for an evaluation install	<p>HyperStore software can be installed on a single host that has just one data drive. The host should have at least 1GB of hard drive space, at least 16GB RAM, and preferably at least 8 processor cores. If you install HyperStore software on a host with less resources than this, the install script will display a warning about the host having less than recommended resources. If you try to install HyperStore software on a host with less 100MB hard drive space or less than 2GB RAM, the installation will abort.</p>

3.2. OS, Network Interfaces, and Disks

This topic describes how to install an OS, configure network interfaces, and set up data disks on each of your host machines using the HyperStore ISO file and the HyperStore system setup tool that's included within the ISO. (For information about acquiring the ISO file see "**HyperStore Installation Introduction**" (page 1).)

OS, Network Interface, and Disk Requirements if you are using the software-only version of HyperStore

If you are using the software-only version of HyperStore and not the ISO:

- **For operating system**, your host machines must be running RHEL 6.x or CentOS 6.x (64 bit). Cloudbian, Inc. recommends using version 6.8 or a later 6.x version.
- **For network interface configuration:**
 - Use static IPv4 addresses. Do not use DHCP. Do not use IPv6.
 - If you use *iptables* or *SELinux* make sure that there are no restrictions at all on internal communication between nodes.
 - On each host the following listening ports must be open: 7; 22; 80; 123; 443; 6379; 6380; 7000; 7199; 8140; 8443; 8888; 9078; 9160; 16080; 16443; 18081; 19050; 19070; 19080; 19081; 19082; 19083; 19084; 19090; 19443. For details about port usage, see the installation Reference topic "**HyperStore Listening Ports**" (page 35).
- **For disk setup**, if your disks are not configured yet Cloudbian, Inc. recommends using the HyperStore system setup tool to do the disk setup on each host. To do so without the ISO you must first complete the "**Prerequisite Packages Installation**" (page 15) task, which will distribute the setup tool to each of your hosts. Then perform the "Setting Up the Disks and Mount Points" task described below, for each host. **If your disks are already configured** and you are not using the HyperStore system setup tool for disk setup:
 - Your disk setup must meet the requirements described in the Installation Reference topic "**File System Requirements**" (page 38).
 - Before installing HyperStore you must manually create a default mount point list for the cluster as described in the Installation Reference topic "**Data Directory Mount Point List (fslist.txt)**" (page 41).

3.2.1. Installing the OS

Follow these steps to install CentOS 6.8 on a host machine using the HyperStore ISO.

1. Burn the HyperStore ISO file on a CD or DVD. Use a "burn image" operation, not a simple copy operation.
2. Have the host machine boot from the ISO CD or DVD. This triggers the installation of the CentOS operating system.

More detail on this step

- a. Insert the CD or DVD into the host machine, then start (or restart) the host machine and use your hardware-specific hot key to enter the BIOS setup utility. Then use the BIOS setup utility to change the machine's boot order so that it boots first from the CD/DVD drive. When you save your changes, the machine will restart and boot from the ISO file.
 - b. When the CentOS Installer welcome screen appears select "Install or upgrade an existing system" and press enter. If you do not select within 5 seconds, the "Install or upgrade an existing system" option automatically executes. The OS install will typically take less than 10 minutes to complete.
 - c. After the install has completed, the CD/DVD tray ejects and the system begins to reboot. Remove the CD/DVD from the tray when it ejects so that the machine does not boot from the CD/DVD again.
3. After CentOS installs and reboots, log in as *root*. The initial password is *password*. After logging in proceed to the "Configuring the Network Interfaces" task described below.

3.2.2. Configuring the Network Interfaces

Follow these steps to configure network interfaces and set the local time zone for a node, using the HyperStore setup tool. You should already be logged into the node as described above.

1. Change into the `/root/CloudianTools` directory. Then launch the `system_setup.sh` tool.

```
[root]# ./system_setup.sh
```

This displays the tool's main menu.

```
System Setup
1) Configure Networking
2) Change timezone
3) Setup Disks
5) Change root Password
   Please Change Password

D) Download Cloudian HyperStore Binary File
   Please Download or place Cloudian HyperStore Binary file in '/root/CloudianPackages'

S) Script Settings
A) About system_setup.sh

X) Exit

Choice:
```

Note If you wish you can enter "5" to change the root password for the node now; or alternatively you will be able to change the root password for all your nodes at once later in the installation procedure.

Note Ignore the "Please download or place Cloudian HyperStore binary file..." message for now; you will do this later in the installation procedure.

2. From the menu item list, enter "1" for Configure Networking. This displays the Networking configuration menu.

```
System Setup » Networking

Interface  IP Address      State  Type      Mode  Master  Speed
1) eth0    192.168.0.20/24 Up      Ethernet --    --      1 Gb/s
2) eth1    --              Down   Ethernet --    --      1 Gb/s
3) eth2    --              Down   Ethernet --    --      1 Gb/s
4) eth3    --              Down   Ethernet --    --      1 Gb/s

Select a number from the list above to edit an interface's configuration

D) Change Domain Name (<unset>)
H) Change Hostname (cloudian-node1)

B) Create Bond Interface
U) Create VLAN Interface

N) Restart Networking
R) Refresh Interface Details

P) Return to the previous menu

Choice: _
```

3. Enter "**d**" and follow the prompts to specify the domain to which this host belongs (for example "enterprise.com").
4. Enter "**h**" and follow the prompts to set the host's hostname.
5. Configure each of the host's physical interfaces:
 - a. From the interface list at the top of the display, enter the number for the interface that you want to configure (for example "**1**"). This displays the menu for configuring the selected interface.

```
System Setup » Networking » Interface Configuration
Hostname: hyperstore1
Interface: eth0

  1) Show Configuration
  2) Edit Configuration
  3) Disable

  R) Restart

  D) Delete Configuration

  P) Return to the previous menu

What would you like to do? _
```

- b. Enter "**2**" for Edit Configuration and follow the prompts to configure the interface settings (IP address, netmask, and so on.). Use **static** IPv4 addresses. Do not use DHCP. Do not use IPv6. To accept a default at a prompt press Enter. After completing the settings enter "**yes**" to save the settings. Then enter "**yes**" again to restart the interface so as to apply the settings.
- c. After a successful restart of the interface, you are brought back to the configuration menu for that interface. Enter "**p**" to return to the main Networking configuration menu.
- d. **Repeat Steps a, b, and c for each of the host's physical interfaces.**

Note: When configuring a second interface, at the end of the configuration sequence you will be asked whether that interface's IP address is the one you will use for your cluster survey file (which you will create later in the HyperStore installation procedure). In the survey file you will use the **public** address for each host, not the private/internal address. So answer "**no**" if you're configuring an internal interface.

6. Optionally, from the Networking configuration menu you can also:
 - Enter "**b**" for Create Bond Interface, and follow the prompts to create and configure a bond interface.
 - Enter "**v**" for Create VLAN Interface, and follow the prompts to create and configure a VLAN interface.
7. After you've completed your networking configuration for the node, from the Networking configuration menu enter "**n**" to Restart Networking for the node. After the networking restart successfully completes, return to the setup tool's main menu.

- From the main menu, enter "2" for Change Timezone and follow the prompts to specify the time zone in which the node is located. After returning to the setup tool's main menu, proceed to "Setting Up the Disks and Mount Points" as described below.

3.2.3. Setting Up the Disks and Mount Points

Follow these steps to configure the data disks and mount points for a node, using the HyperStore setup tool. You should already be at the main menu of the setup tool as described above.

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

```
System Setup » Setup Disks
Selected Disks: sdc sdd sde sdf sdg sdh sdi sdj sdk sdl sdm sdn

  Device  Size  Dependencies      Device  Size  Dependencies
1) sda    256G  6                 2) sdb    256G  6
3) sdc     4T    0                 4) sdd     4T    0
5) sde     4T    0                 6) sdf     4T    0
7) sdg     4T    0                 8) sdh     4T    0
9) sdi     4T    0                10) sdj     4T    0
11) sdk     4T    0                12) sdl     4T    0
13) sdm     4T    0                14) sdn     4T    0

C) Configure Selected Disks
T) Toggle Selection for all disks
R) Refresh Disks
P) Return to the previous menu

Toggle devices by selecting the number listed beside device
Choice: _
```

- 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 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) disks 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" and follow the prompts to have the tool configure the selected disks.
- After you've configured all the data disks on the node, exit the setup tool and log out of the node.

Next, using the CD or DVD that you created with the HyperStore ISO file on it, **complete the "Installing the OS", "Configuring Network Interfaces", and "Setting Up the Disks and Mount Points" tasks for each of the other nodes.**

Note If you're setting node passwords as you go along -- rather than waiting until later in the installation procedure when you can do them all at once -- it's recommended to set the same root password on each node. Otherwise the pre-installation cluster validation tool described later in the procedure will not be fully functional.

Note When configuring disks, the setup script invokes the `ext4` "lazy initialization" feature by which a background, low-priority process of disk initialization will continue for up to two to three hours after you've used the script to format the disks. You can continue with the installation and you can use the HyperStore system after it's installed, but you should wait for a few hours after the installation before running any performance tests.

When you are done with these setup tasks for all your nodes, proceed to "**Prerequisite Packages Installation**" (page 15).

3.3. Prerequisite Packages Installation

This section describes how to install software packages that must be on all of your nodes before you install HyperStore. Working from a single node you will be able to install the packages across your whole cluster.

1. Choose any one of your nodes to serve as the "Puppet Master" node, and log into the node. From this node you will orchestrate the HyperStore installation, and throughout the life of your HyperStore system you will use this node to orchestrate cluster configuration (leveraging the [Puppet](#) software that's integrated with HyperStore).
2. Load the HyperStore product package and your Cloudian license file into the installation staging directory on the Puppet Master node:
 - **If you installed the OS from the ISO**, a directory `/root/CloudianPackages` already exists for you to use as the installation staging directory. Copy your Cloudian license file into that directory, or you can use the temporary low-capacity license file that's already in that directory. If you already have the HyperStore product package (the `*.bin` file) you can copy it into the staging directory also. If you don't have the package yet you can launch the HyperStore setup tool (`/root/CloudianTools/system_setup.sh`) and use main menu option "d" for downloading the HyperStore package file into your staging directory.
 - **If you did not install the OS from the ISO**, create an installation staging directory `/root/CloudianPackages` and copy the HyperStore product package and your Cloudian license file into that directory. If you don't yet have the package or a license, you can acquire them as described in the "**HyperStore Installation Introduction**" (page 1).

Note You can use a different directory as the installation staging directory if you wish, but if you do so you will need to adjust the HyperStore setup tool settings as described further below. Remember that the directory must persist throughout your HyperStore system's lifetime.

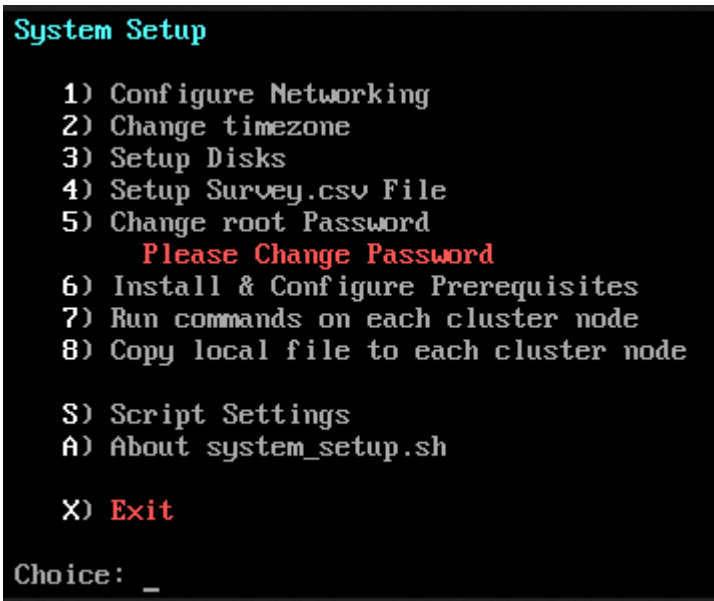
3. Change into the installation staging directory and then run the commands below to unpack the HyperStore package:

```
[root]# chmod +x CloudianHyperStore-6.2.bin
[root]# ./CloudianHyperStore-6.2.bin <license-file-name>
```

4. Launch the `system_setup.sh` tool. If you installed the OS from the HyperStore ISO, launch the `system_setup.sh` instance in the `/root/CloudianTools` directory. If you did not install the OS from the ISO, launch the `system_setup.sh` instance that's now in your installation staging directory.

```
[root]# ./system_setup.sh
```

This displays the tool's main menu.

A terminal window showing the 'System Setup' main menu. The menu items are listed in a numbered list from 1 to 8, followed by 'S) Script Settings', 'A) About system_setup.sh', and 'X) Exit'. The 'Exit' option is highlighted in red. Below the menu, the prompt 'Choice: _' is visible.

```
System Setup

1) Configure Networking
2) Change timezone
3) Setup Disks
4) Setup Survey.csv File
5) Change root Password
   Please Change Password
6) Install & Configure Prerequisites
7) Run commands on each cluster node
8) Copy local file to each cluster node

S) Script Settings
A) About system_setup.sh

X) Exit

Choice: _
```

Note If you are using an installation staging directory other than `/root/Cloud-ianPackages`, from the setup tool's main menu enter "s" for Script Settings. Then in the Script Settings menu enter "8" for Staging Directory, and follow the prompts to specify your staging directory. Then return to the setup tool's main menu and continue with the steps below.

- From the setup tool's main menu, enter "4" for Setup Survey.csv File and follow the prompts to create a cluster survey file with an entry for each of your HyperStore nodes. For each node you will enter a region name, hostname, public IP address, data center name, and rack name. (In the hostname you can use periods, dashes, and underscores but not spaces. In the region, data center, and rack names do not use periods, underscores, or spaces). Optionally you can also specify a node's internal interface name, only if the node is using a different internal interface than the rest of the nodes in the cluster (if all nodes use the same internal interface you can leave this value empty).

```

System Setup » Survey File » Add Entry

No Entries Found
Region  Hostname  IP Address  Datacenter  Rack  Interface
-----  -
Lines in red are commented out in the survey file.

Region Name: region1
Hostname: hyperstore10
Attempting auto IP resolution for hyperstore10 ... Done
IP Address: 192.168.2.18
Data Center Name: DC1
Rack name: RAC1
Internal Interface (optional):

Adding entry to /root/CloudianPackages/survey.csv ... Done

Would you like to add another entry? (Yes/No) [Yes] █

```

After adding an entry for each node, return to the setup tool's main menu. (For more information about the survey file, see the Installation Reference topic "**Cluster Survey File (survey.csv)**" (page 42)).

6. If you have not yet changed the root password for your nodes, do so now by entering "5" from the main menu and following the prompts. Recall that the current root password (if you haven't changed it) is "password". It's recommended to use the same password for each node. Otherwise the pre-installation cluster validation tool described later in the procedure will not be fully functional.
7. 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 either for the whole cluster (if, as recommended, each node has the same root password) or for each node in turn (if the nodes have different passwords). When the prerequisite installation completes for all nodes, return to the setup tool's main menu.

After successfully installing HyperStore prerequisites, from the setup tool's main menu proceed to the "**Pre-Install Check**" (page 18).

Note If you did not use the ISO to install the OS on your nodes but you would like to use the HyperStore system setup script to configure the data disks and mount points on each node, then **before doing the Pre-Install Check** do the following on each of your nodes one at a time, including the Puppet Master:

1. Log into the node and launch the `system_setup.sh` script which will be in the installation staging directory (the "Configure HyperStore Prerequisites" operation that you completed above pushes the script out to the same directory on each node):

```
[root]# <installation_staging_directory>/system_setup.sh
```

2. From the script's main menu select option "3" to Setup Disks. For more information see "**Setting Up the Disks and Mount Points**" (page 14).

3.4. Pre-Install Check

This section describes the final installation preparation step: running an automated check to verify that your cluster now meets all HyperStore requirements for hardware, prerequisite packages, and network connectivity.

1. Upon completing the "**Prerequisite Packages Installation**" (page 15) task, you should be on the Puppet Master node and viewing the main menu of the HyperStore setup tool. From the main menu enter "r" for Run Pre-Installation Checks. This displays the Pre-Installation Checklist menu.

```
System Setup » Pre-installation Checklist

1) Quiet Mode: Disabled
2) Skip Network Check: False
3) Create Log: Disabled
4) Zombie Mode: Disabled
5) Force sync NTP: False

Script Settings
6) Staging Directory: /root/CloudianPackages
7) Survey File: %STAGING_DIRECTORY%/survey.csv

H) Display help information

R) Run Pre-Install Checks

P) Return to the previous menu

Choice: █
```

2. From the Pre-Installation Checklist menu enter "r" for Run Pre-Install Checks. The script then 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 check, proceed to **"Installing a New HyperStore System"** (page 21).

This page left intentionally blank

Chapter 4. Installing a New HyperStore System

This section describes how to do a fresh installation of HyperStore 6.2 software, after "**Preparing Your Environment**" (page 3) and "**Preparing Your Nodes**" (page 9). From your Puppet Master node you can install HyperStore software across your whole cluster.

1. On your Puppet Master node, change into your installation staging directory. Then launch the HyperStore installation script as follows:

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

Note If you have not configured your DNS environment for HyperStore (see "**DNS Set-Up**" (page 4)) 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.

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

For more script launch options, see the Installation Reference topic "**cloudianInstall.sh Command Line Options**" (page 43).

When you launch the installer the main menu displays:

```
Cloudian HyperStore(R) 6.2 Installation/Configuration
-----
0 ) Run Pre-Installation checks
1 ) Install Cloudian HyperStore
2 ) Cluster Management
3 ) Upgrade From 6.1.1 to 6.2
4 ) Advanced Configuration Options
5 ) Uninstall Cloudian HyperStore
6 ) Help
x ) Exit

Choice: █
```

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 "**Pre-Install Check**" (page 18) -- 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.

- From the installer main menu, enter "1" for Install Clouidian 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 the following cluster configuration information:

- The name of the **internal interface** that your nodes will use by default for internal cluster communications. For example, *eth1*.
- The "**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, like "DC1:3,DC2:2". The default is 3 replicas per DC.
- 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 you can enter 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. See "**DNS Set-Up**" (page 4) for details about HyperStore service endpoints.

At the conclusion of the installation an "Install Clouidian 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:

```
Install Clouidian HyperStore
-----

a ) Specify Nodes, Check Connectivity
b ) Specify Cluster Configuration
c ) Review Cluster Configuration
d ) Install Packages and Configure Nodes
e ) Load Schema and Start Services [OK]
f ) Install Third-party Prerequisite Packages
g ) Install Clouidian HyperStore Packages
x ) Return to Main Menu

Choice: █
```

After seeing that the status is OK, return to the installer's main menu.

Note The "Install Clouidian HyperStore" sub-menu supports re-executing specific installation operations on specific nodes. This may be helpful if the installer interface indicates that an operation has failed.

-
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. You can then further exercise your HyperStore system using the Cloudfian Management Console, as described in **"Getting Started with a Newly Installed HyperStore System"** (page 25).

Note For troubleshooting information, see the Installation Reference topic **"Installation Troubleshooting"** (page 34).

This page left intentionally blank

Chapter 5. Getting Started with a Newly Installed HyperStore System

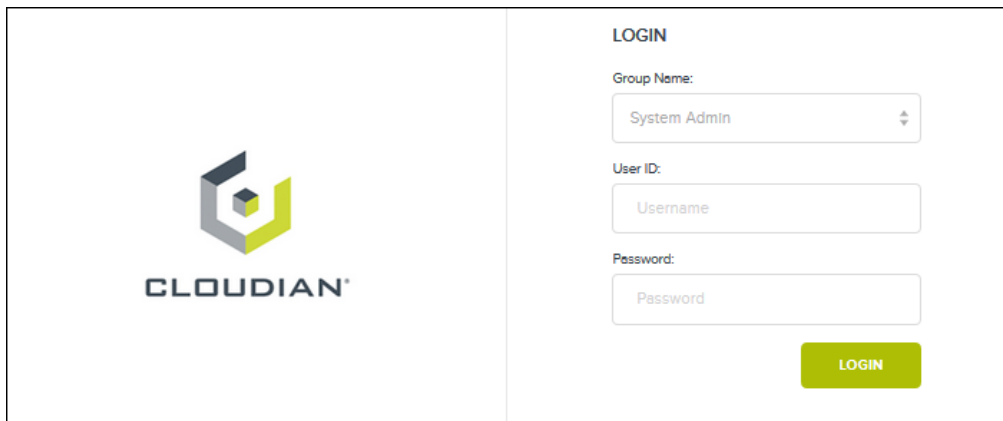
Note For complete HyperStore documentation see the HyperStore Help, which you can access through the CMC's "Help" button or by opening the file `<installation_staging_directory>/doc/HyperStoreHelp/HyperStoreHelp.html` on your Puppet master node.

Once a newly installed HyperStore system is up and running, you can use the Cloudian Management Console (CMC) to take the system for a test drive. Follow these steps:

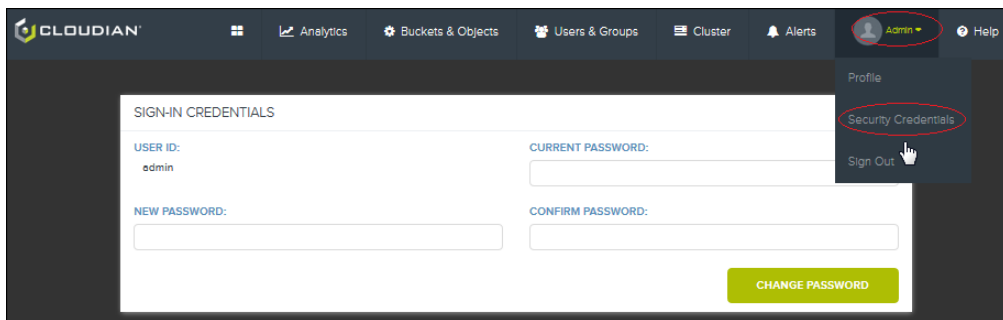
1. Point a browser to `http://<CMC_host_IP_address>:8888/Cloudian`

Since the CMC runs on all of your HyperStore nodes by default, you can use any node's IP address.

2. The connection will automatically switch over to SSL and you will get a certificate warning. Follow the prompts to add an exception for the certificate and accept it. You should then see the CMC's login screen.



3. Log in with the system administrator user ID `admin` and default password `public`.
4. In the upper right corner of the CMC UI, hold your cursor over your user ID ("admin"). From the drop-down menu that displays, select **Security Credentials** to open the **Sign-In Credentials** interface. For security, change the system administrator password.



5. Select the **Cluster** tab, then **Storage Policies**, then **Create Storage Policy**. This opens the **Create New Policy** interface. Create and **Save** a default storage policy for your system. A storage policy is a method of storing and protecting S3 object data and object metadata. Leave the "Group Visibility" unspecified so that this policy is visible to all groups. (At a later time you can edit this policy; create additional policies; and assign a different policy to be the system default policy if you wish).

STORAGE POLICIES
+ CREATE STORAGE POLICY

CREATE NEW POLICY

Policy Name

Policy Description

NUMBER OF DATACENTERS


DATA DISTRIBUTION SCHEME

Replicas Within Single Datacenter
 EC Within Single Datacenter

OBJECT

↓

REPLICA



NUMBER OF REPLICAS

DATACENTER ASSIGNMENT

REGION	DATACENTER	REPLICA	LOCAL EC
sfo-del-region1	DC1	1 of 3	
		2 of 3	disable
		3 of 3	

CONSISTENCY SETTING

CONSISTENCY LEVEL	READ	WRITE
ALL	<input type="checkbox"/>	<input type="checkbox"/>
QUORUM	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
ONE	<input type="checkbox"/>	

GROUP VISIBILITY

Please select a Group ADD

Size Threshold

Compression Type

Server-Side Encryption

SAVE
CANCEL

For detail about the available options when you configure a new storage policy see [Add a Storage Policy](#).

6. Create a regular S3 service user account that will enable you to test the system's S3 object storage services.

Note The system administrator role cannot have its own S3 storage service account.

- a. Select **Users & Groups** → **Manage Groups** → **New Group** to open the **Add New Group** interface. Create a new user group.

ADD NEW GROUP Active Group

Group Name: * Group Description:

Rating Plan:
Please select a Rating Plan

Enable LDAP Authentication

CANCEL SAVE

- b. Select **Users & Groups** → **Manage Users** → **New User** to open the **Add New User** interface. Create a new regular user, assigned to the user group that you created in the previous step. Make a note of the group name, user ID, and the password that you assign to the user so that you will be able to log in as that user.

ADD NEW USER Active User

User ID: * User Type: User Group Name: * Please select a Group

Password: * Confirm Password: *

More

CANCEL SAVE

7. Log out of the CMC, and then log back in as the new user that you created. The CMC now displays only the functions that are available to regular service users, including the **Buckets & Objects** interface.

BUCKETS + ADD NEW BUCKET

ADD NEW BUCKET

Bucket Name Region region1 Storage Policy *HSFSPolicy_region1

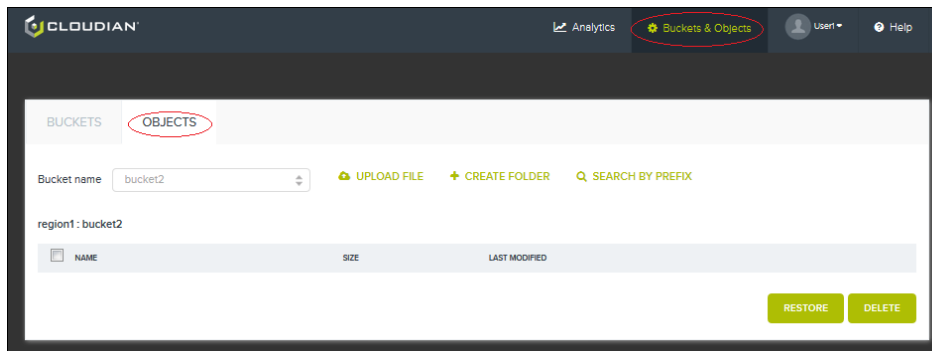
Storage Policy Description
HSFSPolicy_region1

CANCEL CREATE

Note As an alternative to logging out of the CMC, you can remain logged in as an administrator and use the **Manage Users** page to retrieve the user that you just created. Then click the "View User Data" link for that user. The **Buckets & Objects** interface will display for you as if you were that user.

8. Experiment with the CMC **Buckets & Objects** interface. For example:

- a. Add a new storage bucket by entering a bucket name and clicking **Create**. (You must create a bucket before you can upload any objects.)
- b. Use the **Objects** tab to create a folder in the bucket by entering a folder name and clicking **Create Folder**.



- a. Open the folder then upload a file into it by using the **Upload File** function.
- b. Verify that the uploaded file appears in the folder contents list, then verify that you can download the file by clicking on the file name.

Chapter 6. Upgrading an Existing HyperStore System

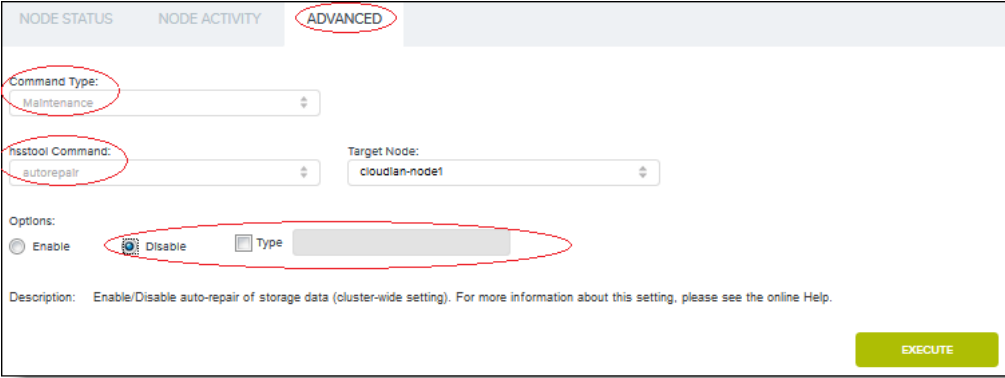
The instructions that follow are for upgrading to HyperStore version 6.2 **from version 6.1 or newer**.

IMPORTANT: If you are currently on a version older than 6.1, do not use the procedure described here. Instead contact Clouidian Support to request instructions for your particular upgrade path.

6.1. Before Upgrading Your System

Before upgrading your HyperStore system, log into the CMC and do the following:

1. Use the CMC's [Node Advanced](#) page to temporarily **shut down the scheduled auto-repair feature**, so that no new auto-repairs kick off while you are performing the upgrade. On the Node Advanced page execute the Maintenance command *autorepair* with the Disable option. The target node can be any node in the cluster. Leave the "Type" option unspecified so that all auto-repair types (Replicas, EC, and Cassandra) are disabled.



The screenshot shows the 'Node Advanced' page in the CMC. The 'ADVANCED' tab is selected. The 'Command Type' dropdown is set to 'Maintenance'. The 'fishtool Command' dropdown is set to 'autorepair'. The 'Target Node' dropdown is set to 'cloudian-node1'. Under 'Options', the 'Disable' radio button is selected, and the 'Type' field is empty. A description at the bottom reads: 'Enable/Disable auto-repair of storage data (cluster-wide setting). For more information about this setting, please see the online Help.' An 'EXECUTE' button is located at the bottom right.

2. Use the CMC's [Repair Status](#) page to **verify that no repairs or rebalancing operations of any type are currently in progress** in your cluster. All nodes should have an "All Clear" status (green cube icon). If any type of repair (including proactive repair) or rebalance is in progress, wait until it completes before you proceed with the upgrade.
3. Use the CMC's [Dashboard](#) page to **verify that no services are currently down** in your cluster -- the dashboard's "Cluster Health" section should show a green circle and status should say "Healthy". If the circle is orange then one or more services are down. In that case, check the CMC's [Alerts](#) page for alerts regarding downed services. **Resolve any errors** identified by alerts, working with Clouidian Support if necessary. [Restart any services](#) that are down.

IMPORTANT: If you are using Xen, Amazon EC2, or Logical Volume Manager (LVM), contact Clouidian Support before upgrading your system.

6.2. Upgrading Your System

To perform the upgrade to HyperStore version 6.2:

1. On your Puppet master node, create a **new staging directory**, with a different name than the directory that you used for your existing cluster. From here forward you will be working in the new staging directory.

Note: Do not make the new staging directory a sub-directory under the staging directory of your existing HyperStore version.

2. Download the HyperStore 6.2 package. (If you need download access information see the introductory note at the top of this document). Place the package into your new staging directory.
3. Copy your HyperStore license file into the new staging directory. (Your current license file is located under `/etc/cloudian-<your-current-version>-puppet/modules/baselayout/files` and ends with suffix `.lic`. Copy this file to the new staging directory.)
4. In your new staging directory, run the following command to unpack the HyperStore package:

```
root# ./CloudianHyperStore-6.2.bin <license-file-name>
```

5. In your new staging directory, launch the installer:

```
[root]# ./cloudianInstall.sh
```

```
Cloudian HyperStore(R) 6.2 Installation/Configuration
-----
0 ) Run Pre-Installation checks
1 ) Install Cloudian HyperStore
2 ) Cluster Management
3 ) Upgrade From 6.1.1 to 6.2
4 ) Advanced Configuration Options
5 ) Uninstall Cloudian HyperStore
6 ) Help
x ) Exit

Choice: █
```

6. From the installer main menu enter "3" for "Upgrade from <your version number> to 6.2". Then at the prompt, confirm that you wish to continue with the automated upgrade.

The upgrade script will first check your configuration template files (`*.erb` files) from your current HyperStore system to determine whether you made any customizations to those settings (changes from the default values). If you have made such configuration changes, the installer creates a text file that lists those changes and prompts you to review the text file. If you want to carry your customizations forward to your upgraded system, you must **manually edit the new version's configuration templates** (under `/etc/cloudian-<new-version-#>-puppet/modules`) before continuing with the upgrade process. You do

not need to do a Puppet push -- just edit the configuration templates and then resume the automated upgrade process. After the upgrade successfully completes, proceed to Post-Upgrade Steps.

Note Customizations that you may have made to the configuration file *common.csv* are handled differently. The installer detects such customizations and automatically applies the same customizations to the new version's *common.csv* file, without you having to do anything.

Note If the upgrade fails and rolls back, you can get details about the failure from the *cloudian-installation.log* and the *cloudian-upgrade.log*, both of which are generated in your staging directory. The upgrade process also generates an *upgrade-logNconfig*.tgz* "S.O.S" tar file (which packages together multiple upgrade-related files) that you can provide to Cloudian Support in the event of upgrade problems.

6.3. After Upgrading Your System

After all HyperStore nodes have been upgraded, verify that all services are running and that the HyperStore version is now 6.2:

1. After the automated upgrade completes, you should be taken back to the main menu of the HyperStore installer. The first post-upgrade step is to confirm that all your HyperStore services are up and running:
 - a. From the installer's main menu select "Cluster Management".
 - b. From the Service Management sub-menu that displays select "Manage Services".
 - c. At the "Select a service to manage:" prompt, select All Services.
 - d. At the "Enter command" prompt, type *status*.
All services on all nodes should then indicate that they are running.
2. Next, confirm that the HyperStore software version is correct:
 - a. Still on the Service Management menu, at the "Select a service to manage:" prompt select the S3 Service.
 - b. At the "Enter command" prompt, type *version*.
On all nodes the S3 version should indicate version 6.2.
After confirming the version you can exit the installer.
3. Use the CMC to check on your upgraded cluster:
 - Use the functions under the **Cluster** tab to confirm the health of your cluster and individual nodes.
 - Use the functions under the **Users & Groups** tab to confirm that you can retrieve your user groups and users.
 - If you retrieve a user in the **Manage Users** page and access their data (or if you log out of the CMC as system admin and then log back in as a regular user), you will see the **Buckets & Objects** page. Here you can confirm that you can successfully download or upload objects.

4. If prior to the upgrade you had made any customizations to the branding of the CMC interface, only your customized logos and customized application name will be retained after the upgrade. You will need to re-implement any changes that you had made to the browser tab title and/or the color scheme, by again following the instructions for **Rebranding the CMC UI**.

You are now done with upgrading to HyperStore 6.2.

Note At the end of the upgrade process, the HyperStore auto-repair feature is **automatically** set back to being enabled.

Chapter 7. HyperStore Installation Reference

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

- **"Installation Troubleshooting"** (page 34)
- **"HyperStore Listening Ports"** (page 35)
- **"Outbound Internet Access"** (page 38)
- **"File System Requirements"** (page 38)
- **"Data Directory Mount Point List (fslist.txt)"** (page 41)
- **"Cluster Survey File (survey.csv)"** (page 42)
- **"cloudianInstall.sh Command Line Options"** (page 43)

Note The install script's "Advance Configuration Options" are covered in the "System Configuration" section of the HyperStore Administrator's Guide.

7.1. Installation Troubleshooting

7.1.1. Installation Logs

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

On the Puppet master node (on which you're running the install script):

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

On each Puppet 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.

7.1.2. Debug Mode

Another potentially useful source of troubleshooting information is to run the installer in debug mode:

```
[root]# <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. The Install Cloudian HyperStore sub-menu will let you choose from:

- a) Specify Nodes, Check Connectivity
- b) Specify Cluster Configuration
- c) Install Packages and Configure Nodes [includes Run Puppet]
- d) Load Schema and Start Services

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

7.1.3. 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 *finger fqdn* and *hostname -f* commands output ' cloudian-singlenode.MyDomain.Com' to the console.

ISSUE: Puppet is unable to propagate configuration settings to the 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 [g] — "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 encountered 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 [i] — "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]".

7.2. HyperStore Listening Ports

The HyperStore system uses the listening ports specified in the table below. **These ports must be open on each host on which you run HyperStore.** Only the CMC Service and S3 Service (with

listening ports marked in italics below) should be exposed to traffic that ultimately originates from end users outside the data center. For all other listeners, the requests will originate only from other HyperStore components, internal to the cluster.

Service	Listening Port	Purpose
Clouddian Management Console (CMC)	<i>8888</i>	Requests from administrators' or end users' browsers via HTTP
	<i>8443</i>	Requests from administrators' or end users' browsers via HTTPS
S3 Service	<i>80</i>	Requests from the CMC or other S3 client applications via HTTP
	<i>443</i>	Requests from the CMC or other S3 client applications via HTTPS
	19080	JMX access
Admin Service	<i>18081</i>	Requests from the CMC or other Admin API clients via HTTP
	<i>19443</i>	Requests from the CMC or other Admin API clients via HTTPS (Note: The CMC by default uses HTTPS to access the Admin Service)
	19081	JMX access
IAM Service	<i>16080</i>	Requests from the CMC or other Admin API clients via HTTP
	<i>16443</i>	Requests from the CMC or other Admin API clients via HTTPS (Note: The CMC by default uses HTTPS to access the Admin Service)
	19084	JMX access
Redis Monitor	<i>9078</i>	Communication between primary and backup Redis Monitor instances
	19083	JMX access
HyperStore Service	19090	Data operation requests from the S3 Service
	19050	Communication between HyperStore Service instances
	19082	JMX access
Redis DBs	6379	Requests to the Redis Credentials DB from the S3 Service, HyperStore Service, or Admin Service; and communication between Redis Credentials instances
	6380	Requests to the Redis QoS DB from the S3 Service, HyperStore Service, or Admin Service; and communication between Redis QoS instances

Service	Listening Port	Purpose
Cassandra	9160	Data operations requests from the S3 Service, HyperStore Service, or Admin Service
	7000	Communication between Cassandra instances
	7199	JMX access
Cloudfian Monitoring Agent	19070	Requests from the Cloudfian Monitoring Data Collector
Puppet Master	8140	On your Puppet Master node (the HyperStore node from which you will manage cluster installation and configuration) this port will service incoming requests from Puppet agents on your other HyperStore nodes
SSH	22	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	NTP port for time synchronization between nodes
Echo	7	The Cloudfian Monitoring Data Collector uses this port on each node to check whether the node is reachable, if ICMP is unavailable.
ICMP	n/a	The Cloudfian Monitoring Data Collector uses ICMP to check whether each node is reachable.

7.2.1. Iptables and SELinux

If you use *iptables* or *SELinux* — or any type of firewall — on the nodes in your HyperStore cluster, make sure that there are **no restrictions on internal communication between nodes**. HyperStore nodes sometimes communicate with each other via JMX, and when they do, after initial connection establishment on the designated JMX port (see detail in the table above) a random port is used for continued communication. Therefore there cannot be any port restrictions on internal communication between HyperStore nodes. **Only external communication should be restricted.**

Note that if *iptables* is running or *SELinux* is enabled on hosts on which you are installing HyperStore, you will need to use the installation script launch option "force" — otherwise the installation will abort when it detects *iptables* running or *SELinux* enabled on a host.

7.2.2. 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 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).

7.3. 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. If you use forward proxying in your environment, after HyperStore installation you may want to set up forward proxying to support these HyperStore features:

- **Phone Home** — The Phone Home feature (also known as "Smart Support") 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 are `mts.properties.erb: phonehome.proxy.*`).
- **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.
- **Pre-Configured ntpd** — Accurate, synchronized time across the cluster is vital to HyperStore service. Two of your HyperStore nodes are automatically configured to act as local NTP servers (with one configured as the primary and the other as the secondary). All the other HyperStore nodes are automatically configured as clients to the two local NTP servers. The two nodes that act as local NTP servers are configured to connect to external NTP servers — the default public servers from the `pool.ntp.org` project. In order to connect to the public NTP servers at `pool.ntp.org`, the two local NTP servers must be allowed outbound internet access. This feature doesn't support configuring an explicit forward proxy, but you can use transparent forward proxying if you wish.

To see which of your HyperStore nodes are running as the primary and secondary local NTP servers, after HyperStore installation log into the CMC and go to **Cluster** → **Cluster Config** → **Cluster Information**.

7.3.1. 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 Admin 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 35) for HyperStore requirements regarding listening port access.

7.4. File System Requirements

You can use the HyperStore `system_setup.sh` tool to [configure the data disks and mount points on your nodes](#). The tool is part of the ISO and is also part of the HyperStore product package (the `.bin` file). If you do not use the system setup tool for disk setup, review the information below to make sure that your host machines meet HyperStore requirements.

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 Requirements](#)" (page 10)). On host machines with multiple hard drives:

- HyperStore will by default use the drive that the OS is on for storing service metadata (in Cassandra and Redis databases). If a host machine has 10 or more drives in total, Cloudian, Inc. recommends that you dedicate two drives to the OS (and service metadata) in a RAID-1 mirroring configuration. Preferably the OS 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. Format each of the 12 HDDs with *ext4* file systems and configure mount points such as */cloudian1*, */cloudian2*, */cloudian3* and so on.

HyperStore **does not support** XFS file systems; VirtIO disks; Logical Volume Manager (LVM); or Multipathing. For questions regarding these unsupported technologies, contact Cloudian Support:

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 or 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.

```
[root@hyperstore1 etc]# 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:

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

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

```
[root]# blkid
...
...
/dev/sdb1: UUID="a6fed29c-97a0-4636-afa9-9ba23e1319b4" TYPE=
E="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-D=<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:

```
[root]# mount -a
```

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

You Must Manually Create a Data Directory Mount Point List (fslist.txt)

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 Puppet Master node. For detail see "**Data Directory Mount Point List (fslist.txt)**" (page 41).

Mount Point Naming Guidelines

If you are installing HyperStore on multiple hosts that each have multiple disk drives, 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. For example, if each host has 12 disks for S3 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 S3 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, when you install HyperStore there is a way to configure the system to accommodate differences in mount point names across nodes -- this is described in the install instructions.

Option for Putting Cassandra Data on Dedicated Disks Rather Than the OS Disk

Regarding Cassandra data, another supported configuration — for a host with many drives — is to put your Cassandra data directory and Cassandra commit log directories each on dedicated disks, rather than on the OS disk. In this case you would have:

- OS drive (with Redis also).
- Cassandra data directory drive (mount point path **must** include `/cassandra`)
- Cassandra commit log directory drive (mount point path **must** include `/cassandra_commit`)
- Multiple drives for S3 object data (with mount points for example `/cloudian1`, `/cloudian2`, `/cloudian3` and so on).

In this configuration, where Cassandra data is on a different disk than the OS, it's advisable to use RAID-1 for the OS disk. It's not necessary to use RAID for a dedicated Cassandra disk.

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. Cloudbian, Inc. recommends that you set the reserved space to 0% for each drive on which you will store HyperStore data (S3 object data).

For each HyperStore data drive do the following.

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

root# tune2fs -l <device>

# Set Reserved block count to 0%:

root# tune2fs -m 0 <device>

# For example:

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

7.5. Data Directory Mount Point List (fslis.txt)

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 Puppet Master node. 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 `/cloudian8` -- you only need to create one mount point list file. If some nodes have a different set of mount points as 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 `fslis.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> <mount-point>
```

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 specify one mount point for metadata stored in Cassandra (must include the string "cassandra" in the mount point path) and/or one mount point for the Cassandra commit log (must include "cassandra_commit" in the mount point path). If you do not specify these Cassandra mount points in *fslist.txt*, by default the system automatically puts Cassandra data and commit log directories on the same disk on which the operating system and application files reside.

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-node5_fslist.txt* file and a *cloudian-node6_fslist.txt* file that specify mount points for two non-standard nodes.

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

7.6. Cluster Survey File (survey.csv)

During the **"Prerequisite Packages Installation"** (page 15) 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 add more nodes to your cluster.

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

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

Note The survey file **must** be kept in the installation staging directory, not in a different directory.

- *<region-name>* — The HyperStore system supports having multiple service regions with each region having its own independent storage cluster and own independent S3 object inventory, and with S3 clients 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. A region name can be a maximum of 52 characters, and no periods, underscores, spaces, or special characters are allowed within the name. For more information on regions see "Service Regions Feature Overview" in the "Major Features" section of the *HyperStore Administrator's Guide*.
- *<hostname>* — Hostname of a host machine on which you are installing HyperStore software. Be sure to include in your survey file one line for each node on which you are installing HyperStore. You can install to all your intended nodes at the same time, even nodes in a different data center or region. In the hostnames you can use periods (such as in FQDNs), dashes, and underscores -- but not spaces or special characters.

- *<ip4-address>* — Service network IP (v4) address that the hostname resolves to. Do not use IPv6.
- *<datacenter-name>* — Name of the data center in which the host machine is located. Maximum 256 characters, with no periods, underscores, spaces, or special characters allowed. If there is no official name you can just create one for purposes of HyperStore configuration.
- *<rack-name>* — Name of the server rack in which the host machine is located. Maximum 256 characters, with no periods, underscores, spaces, or special characters allowed. If there is no official name you can just create one for purposes of HyperStore configuration.
- *[<internal-interface>]* — Use this field only if 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.

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

```
region1,arcturus,10.10.2.1,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-vm7,10.10.1.33,DC1,RAC1
tokyo,cloudian-vm8,10.10.1.34,DC1,RAC1
tokyo,cloudian-vm9,10.10.1.35,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,hyperstore1,10.1.0.1,DC1,RAC1
boston,hyperstore2,10.1.0.2,DC1,RAC1
boston,hyperstore3,10.1.0.3,DC1,RAC1
boston,hyperstore4,10.2.0.1,DC2,RAC5
boston,hyperstore5,10.2.0.2,DC2,RAC5
chicago,hyperstore6,10.3.0.1,DC3,RAC2
chicago,hyperstore7,10.3.0.2,DC3,RAC2
chicago,hyperstore8,10.3.2.1,DC3,RAC7,eth2
chicago,hyperstore9,10.3.2.2,DC3,RAC7,eth2
```

7.7. cloudianInstall.sh Command Line Options

The HyperStore installation script *cloudianInstall.sh* resides in your installation staging directory on your Puppet Master node. Typically you would launch the script either like this:

```
[root]# ./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):

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

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

```
[root]# ./cloudianInstall.sh [-s <survey-filename>] [-c <config-filename>]
[-k <ssh-private-key-filename>] [-l <install-log-filename>] [-d] [-h]
[no-hosts] [configure-dnsmasq] [force]
```

Note If you use multiple options, on the command line place options that start with a "-" (which most options do) before options that do not (such as no-hosts).

After using the installer for product installation or subsequently for system configuration tasks, 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.

- `[-s <survey-filename>]` — Name of your cluster survey file. If you do not specify the survey filename argument, the script will prompt you for the file name during installation.
- `[-c <config-filename>]` — The first time you run `cloudianInstall.sh`, it creates in the installation staging directory a cluster configuration file `CloudianInstallationConfiguration.txt` based on your survey file input and your interactive input. If you are subsequently doing multiple installations for testing purposes, you can edit this file, rename it, and then use your customized version as input to `cloudianInstall.sh` runs by using the `-c <config-filename>` option (perhaps in combination with the `-b` option).
- `[-k <ssh-private-key-filename>]` — The Puppet master employs SSH for secure communication with the rest of your HyperStore installation nodes. Use the `-k <ssh-private-key-filename>` option if you want to use your own existing SSH authentication key pair rather than having the HyperStore install tool generate a key pair for you. If you want to use this option, then before running the `cloudianInstall.sh` script:
 - Copy your private key into the installation staging directory (where the install script resides).
 - Copy your public key to each host on which you plan to install HyperStore (see standard SSH set-up documentation for guidance); or copy your public key to the installation staging directory and the install script will copy the public key to your target nodes automatically.
- `[-l <install-log-filename>]` — By default installation logging information is written to `cloudian-installation.log` in the staging directory. You can use this option if you'd prefer that installation logging information be written to a different file. Include the path if you want the file location to be other than in the staging directory.
- `[-d]` — Turn on debugging output.
- `[-h]` — Display Help information for the tool. This option causes the tool to print a usage message and exit.
- `[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 4), 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. (See the System Configuration chapter of the HyperStore Administrator's Guide for instructions.)

- *[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 16GB
 - OS is Red Hat or CentOS
 - System Architecture is x86 64-bit
 - SELinux is disabled
 - Iptables is not running
 - Directory permissions are as needed
 - Hostnames of installation hosts are resolvable or installer is allowed to append to hosts' */etc/hosts* files
 - S3 service and Admin service endpoints (which you will specify during the interactive install) are resolvable -- which they should be so long as you either create DNS entries for HyperStore service endpoints or else use the *configure-dnsmasq* option when launching the installer