

Clouddian HyperIQ and LogInsight User's Guide

Version 2.2

This page left intentionally blank.

Confidentiality Notice

The information contained in this document is confidential to, and is the intellectual property of, Cloudian, Inc. Neither this document nor any information contained herein may be (1) used in any manner other than to support the use of Cloudian software in accordance with a valid license obtained from Cloudian, Inc, or (2) reproduced, disclosed or otherwise provided to others under any circumstances, without the prior written permission of Cloudian, Inc. Without limiting the foregoing, use of any information contained in this document in connection with the development of a product or service that may be competitive with Cloudian software is strictly prohibited. Any permitted reproduction of this document or any portion hereof must be accompanied by this legend.

This page left intentionally blank.

Contents

What's New in Version 2.2	9
Chapter 1. HyperIQ and LogInsight Introduction	11
Chapter 2. Installing or Upgrading HyperIQ	13
2.1. Requirements	13
2.2. Installing HyperIQ as a Virtual Appliance	14
2.2.1. Deploying the HyperIQ VM	14
2.2.2. Installing the HyperIQ Application	15
2.3. Upgrading a HyperIQ Virtual Appliance	19
2.3.1. Starting From HyperIQ on Rocky Linux, Upgrading to HyperIQ 2.2	19
2.3.2. Starting From HyperIQ on Ubuntu with Two Disks, Upgrading to HyperIQ 2.2 on Rocky Linux	20
2.3.3. Starting From HyperIQ on Ubuntu with One Disk, Upgrading to HyperIQ 2.2 on Rocky Linux	23
2.3.4. Starting From HyperIQ on Ubuntu, Upgrading to HyperIQ 2.2 on Ubuntu	25
2.4. Installing or Upgrading HyperIQ in an Existing Docker Environment	26
2.5. Using an Existing Grafana Deployment to Display HyperIQ Dashboards	27
2.6. Uninstalling HyperIQ	27
Chapter 3. Getting Started With HyperIQ	29
3.1. Connecting to the HyperIQ UI	29
3.2. Viewing HyperIQ Dashboards and Panels	29
3.2.1. Filtering the Data in a Dashboard	31
3.2.2. Changing the Time Range for a Dashboard	31
3.2.3. Using an Panel Legend to Focus the Data in a Graph Panel	32
3.2.4. Viewing a Larger Version of a Panel	33
3.2.5. Viewing a Description of the Data in a Panel	33
3.2.6. Opening the Cloudian Management Console	34
3.2.7. Opening LogInsight	34
3.2.8. For More Information About Working with Dashboards and Panels	34
3.3. Adding HyperIQ Users	34
3.3.1. Adding a HyperIQ User	34
3.3.2. Assigning Permissions for a Specific Dashboard	36
Chapter 4. Administering HyperIQ	39
4.1. Using the iqmgr Tool to Manage HyperIQ	39

4.1.1. Updating HyperIQ to Monitor Additional HyperStore Resources	39
4.1.2. Changing HyperIQ Configuration Settings	40
4.1.3. Managing HyperIQ Users	42
4.1.4. Additional iqmgr Commands	43
4.2. Using Docker Commands	44
4.3. Configuring LDAP Authentication for HyperIQ Users	44
4.4. Alerts and Notifications	46
4.4.1. Reviewing Pre-Configured Alert Rules	47
4.4.2. Setting Up Notification Channels	52
4.4.3. Customizing Alert Thresholds	54
4.4.4. Silencing or Pausing Alerts	55
4.4.5. Customized Alerts and HyperIQ Upgrades	57
Chapter 5. Installing LogInsight	59
5.1. Requirements	59
5.2. Setting Up Your HyperStore Data Feed for LogInsight	59
5.2.1. Creating a Bucket For LogInsight Input Files	60
5.2.2. Configuring HyperStore to Write to the LogInsight Bucket	62
5.2.3. Special Requirements If You Have Multiple HyperStore Clusters	63
5.3. Installing LogInsight	64
5.3.1. Deploying the LogInsight Virtual Appliance	64
5.3.2. Running the LogInsight Installer	65
5.4. Starting Or Stopping LogInsight	68
5.5. Linking from the HyperIQ UI to the LogInsight UI	68
5.6. Upgrading LogInsight	69
Chapter 6. Getting Started With LogInsight	71
6.1. LogInsight Data Ingestion	71
6.2. Connecting to LogInsight Web UI	71
6.3. Viewing LogInsight Dashboards and Panels	72
6.3.1. Filtering the Data in a Dashboard	73
6.3.2. Changing the Time Range for a Dashboard	74
6.3.3. Using a Panel Legend to Focus the Data in a Graph Panel	74
6.3.4. Viewing a Larger Version of a Panel	75

6.3.5. Viewing a Description of the Data in a Panel	75
6.4. LogInsight Dashboards At a Glance	76
6.4.1. HyperStore Hardware and Software Dashboards	77
6.4.2. HyperStore S3 Dashboards	79
6.5. Adding LogInsight View-Only Users	81
6.6. LogInsight Alerts	82

This page left intentionally blank.

What's New in Version 2.2

The main new features and enhancements in the Clouidian HyperIQ 2.2 release are:

- New panels --
 - The *Hardware Info* dashboard's "CPU/Mem/Load Details" section now includes a "Swap Usage" panel.
 - The *S3 Analytics* dashboard's "S3 Transaction Duration and Size" section now includes an "Object Size Distribution" panel.
- New alerts -- New pre-configured alerts have been added: *HDDDiskTempHigh*, *NVMeDiskTempHigh*, *NodeFanLowRPM*, and *NodeFanNoRPM*. For more information see **"Reviewing Pre-Configured Alert Rules"** (page 47).
- Improved visibility and alerting for bonded interfaces -- If you are using bonded interfaces in your HyperStore system, the HyperIQ *Hardware Info* dashboard's "Network Traffic Details" section will now show status information for the slave (physical) interfaces that are bonded and the existing HyperIQ alert *NodeInterfaceDown* will now fire if a slave interface goes down.
- New way of handling your customized alerts during upgrades -- For more information see **"Customized Alerts and HyperIQ Upgrades"** (page 57).
- Renamed installation command-line option for monitoring File or Search services -- The installation command-line option that you use if you want HyperIQ to monitor HyperStore File Service nodes or HyperStore Search Service nodes has been renamed: it used to be `--auxiliary-nodes` and is now called `--monitor-pool-nodes`. If you are upgrading a HyperIQ system for which you had used the `--auxiliary-nodes` command-line option during the system's original installation, use the `--monitor-pool-nodes` option during upgrade. For more information about upgrade see **"Upgrading a HyperIQ Virtual Appliance"** (page 19).
- Grafana upgrade -- HyperIQ now uses Grafana version 11.4. Previously HyperIQ used Grafana 9.5.
- New HyperIQ partner product, LogInsight -- LogInsight is a Clouidian product that works in tandem with HyperIQ to provide comprehensive visibility into your HyperStore system's status, performance, and usage. While HyperIQ prioritizes real-time reporting and alerting based on HyperStore key performance metrics, LogInsight performs deep-dive analysis of HyperStore system, application, and transaction logs from the past few months of operation. With its ability to uncover meaningful trends, correlations, and anomalies in massive amounts of log data, LogInsight helps you identify and redress looming system shortfalls or problems before they materialize, and efficiently troubleshoot problems that do occur.

Like HyperIQ, LogInsight is designed to run on a dedicated virtual machine separate from your HyperStore cluster, and is an entirely on-premise solution that does not require external internet connectivity.

Contact your Clouidian Sales representative today to learn more about LogInsight.

Note For more granular HyperIQ release details including bug fixes please see the release notes.

This page left intentionally blank

Chapter 1. HyperIQ and LogInsight Introduction

HyperIQ and LogInsight are Clouidian products that together provide you comprehensive visibility into your HyperStore system's real-time status, performance, and usage, as well as a view into meaningful patterns and trends from recent weeks and months. Equipped with this information you can manage your HyperStore system not just reactively, but also proactively to stay ahead of the curve as service conditions and usage patterns evolve.

Although HyperIQ and LogInsight can each be used as stand-alone products, they are best used in combination:

- **HyperIQ** provides you a real-time view into your HyperStore system, as metrics relating to host hardware and operating systems, applications and services, and S3 user behavior are ingested into HyperIQ every few minutes. A powerful tool for operational intelligence, HyperIQ shows you exactly how things are going for your HyperStore system and for your users, at this moment. HyperIQ also supports alerting based on key performance indicators, deliverable through notification channels such as email, Slack, and MS Teams.
- **LogInsight** ingests data once a day and can provide as much as a 90-day look-back capability. LogInsight performs deep-dive analysis of system, application, and transaction logs to detect correlations, trends, and anomalies with implications for current and future system performance and service usage. As a potent tool for predictive maintenance, capacity planning, and multivariate analysis, LogInsight helps you identify and redress looming system shortfalls or problems before they materialize, and efficiently troubleshoot problems that do occur.

HyperIQ and LogInsight have in common that they:

- Both are available from Clouidian as OVAs and are intended to run on virtual machines separate from your HyperStore cluster.
- Both are entirely on-premise systems that do not require external internet connectivity.
- Both use Grafana-based web front-ends to deliver dynamic and flexible visualizations.

If you are going to use both products you will install each product on a separate virtual machine, and each product has its own web UI.

The rest of this document covers these topics:

- **"Installing or Upgrading HyperIQ"** (page 13)
- **"Getting Started With HyperIQ"** (page 29)
- **"Administering HyperIQ"** (page 39)
- **"Installing LogInsight"** (page 59)
- **"Getting Started With LogInsight"** (page 71)

This page left intentionally blank

Chapter 2. Installing or Upgrading HyperIQ

The chapter covers the following HyperIQ installation and upgrade topics:

- **"Requirements"** (page 13)
- **"Installing HyperIQ as a Virtual Appliance"** (page 14)
- **"Upgrading a HyperIQ Virtual Appliance"** (page 19)
- **"Installing or Upgrading HyperIQ in an Existing Docker Environment"** (page 26)
- **"Using an Existing Grafana Deployment to Display HyperIQ Dashboards"** (page 27)
- **"Uninstalling HyperIQ"** (page 27)

2.1. Requirements

The table below describes what you will need to install and use HyperIQ.

Requirement	Description
Virtualization or containerization platform	<p>HyperIQ is available in two forms:</p> <ul style="list-style-type: none">• A virtual appliance, for installation on VMware, Hyper-V, or VirtualBox.• A containerized application, for installation in an existing Docker environment running on Linux (requires docker 19.03 or newer and docker-compose 1.17 or newer; Kubernetes is not currently supported)
HyperIQ product package	<p>On the Cloudian Support portal are two HyperIQ 2.2 package files, <i>hyperiq-2.2.ova</i> and <i>hyperiq-2.2.bin</i>. Which file(s) to download depends on what you want to do:</p> <ul style="list-style-type: none">• Newly install HyperIQ 2.2 as a virtual appliance: <i>hyperiq-2.2.ova</i>• Upgrade from an existing HyperIQ VM that's still running Ubuntu and has two virtual disks (an OS disk and a dedicated HyperIQ data disk) to HyperIQ 2.2 on Rocky Linux with two virtual disks: <i>hyperiq-2.2.ova</i>• Upgrade from an existing HyperIQ VM that's still running Ubuntu and has just one virtual disk to HyperIQ 2.2 on Rocky Linux with two virtual disks: <i>hyperiq-2.2.ova</i> and also <i>hyperiq-2.2.bin</i>• Upgrade from an existing HyperIQ VM to HyperIQ 2.2 without migrating the OS (this retains your current HyperIQ VM's OS, whether Ubuntu or Rocky Linux): <i>hyperiq-2.2.bin</i>• Install or upgrade to HyperIQ 2.2 in your own Docker environment: <i>hyperiq-2.2.bin</i> <div><p>Note If your existing HyperIQ VM is still running Ubuntu, migrating to Rocky Linux as part of your HyperIQ 2.2 upgrade is recommended (and instructions are in the "Upgrading a HyperIQ Virtual Appliance" (page 19) section). The OVA packages from HyperIQ versions 2.1 and later have been based on Rocky Linux. Starting with (future release) HyperIQ 2.3, running HyperIQ on Ubuntu will no longer be supported.</p></div>

Requirement	Description
HyperStore system(s)	<p>The most common topology is to have one HyperIQ instance monitoring one HyperStore system (which may be a multi-region system). Other supported topologies are:</p> <ul style="list-style-type: none"> • Having one HyperIQ instance monitoring multiple HyperStore systems. • Having two HyperIQ instances monitoring a HyperStore system (for HyperIQ high availability)
HyperStore license for "HyperIQ Enterprise" (only required for S3 Analytics)	<p>To view S3 Analytics data in HyperIQ, your HyperStore system license must include "HyperIQ Enterprise" support. To view other types of monitoring data in HyperIQ -- such as data pertaining to HyperStore system health and performance -- no special license is required (any HyperStore license supports using HyperIQ for this system health data).</p>
Network access between HyperIQ and HyperStore	<p>Your network must allow:</p> <ul style="list-style-type: none"> • Each HyperIQ host to access each HyperStore host on ports 22 and 9999 • Each HyperStore host to access each HyperIQ host on port 3443 • If your HyperStore system includes HyperStore File Service, HyperStore Search Service, and/or HyperBalance hosts, HyperIQ must be able to access those hosts on ports 22 and 9999. Those hosts do not need to access HyperIQ on port 3443 or any other port.

2.2. Installing HyperIQ as a Virtual Appliance

This section describes how to do a fresh installation of HyperIQ 2.2 as a virtual appliance. It covers the two stages of the HyperIQ virtual appliance setup:

- **"Deploying the HyperIQ VM"** (page 14)
- **"Installing the HyperIQ Application"** (page 15)

Note If you are using the HyperBalance load balancer with your HyperStore system and **if you want HyperIQ to monitor HyperBalance** (as well as monitoring your HyperStore system), before installing HyperIQ you must enable the Prometheus exporter on each of your HyperBalance nodes. On each HyperBalance node, log in to the HyperBalance UI, go to **Cluster Configuration -> Layer 7 -- Advanced Configuration** and then check the "Enable Prometheus Exporter" checkbox.

2.2.1. Deploying the HyperIQ VM

Use your virtualization platform console to deploy a new virtual machine (VM) from the HyperIQ OVA file that you downloaded from the Cloudbian Support portal.

For example, for VMware ESXi:

1. From the device on which you downloaded the HyperIQ OVA file, connect to your ESXi Server console and log in.
2. On the left side of the ESXi console, select **Virtual Machines**.
3. Click **Create / Register VM**. This launches the **New virtual machine** creation wizard.

4. For the VM creation type select “Deploy a virtual machine from an OVF or OVA file”, then click **Next**.
5. Enter your desired name for the virtual machine, then click in the file selection area of the wizard screen and browse (on your local device) to the HyperIQ OVA file that you downloaded. After selecting the OVA file, click **Next**.
6. In the remaining screens of the VM creation wizard, accept the default values (or make adjustments if you wish).
7. In the wizard's **Ready to complete** screen, review the attributes of the new VM and then click **Finish**.

Although the wizard will close and the new VM will appear in the **Virtual machines** list right away, it will take some time -- about 10-15 minutes depending on your environment -- for creation of the VM to complete. You can track the progress by opening the **Recent tasks** panel at the bottom of the **Virtual machines** list.

Note When deploying a VM from the HyperIQ OVA on VMware, the OS type will automatically default to "Other Linux 64 bit" and the boot firmware will automatically default to EFI. If you are using a different virtualization platform (Hyper-V or VirtualBox) and if these two settings are not automatically selected by default, you will need to explicitly choose these settings as you deploy the VM.

The HyperIQ VM that is created has these specs:

- 2 CPU cores
- 8GB RAM
- Two disks:
 - 50GB disk for the OS and application
 - 150GB disk for HyperIQ data (mounted as `/var/lib/docker`)
- OS = Rocky Linux 8.10, 64-bit
- hostname = Defaults to *hyperiq* but is configurable during HyperIQ application setup (as described further below)

IMPORTANT ! The HyperIQ VM stores its data locally, on the 150GB disk, without any replication. **In a production environment you should take appropriate precautions to ensure data durability**, such as performing regular backups of the VM (or at least regular backups of the HyperIQ data disk).

2.2.2. Installing the HyperIQ Application

If the HyperIQ VM is not already powered on, power it on now. Then log in to the VM with these credentials:

- User name: **cloudian**
- Password: **Cloudian1!**

You will be prompted to create a new password for the 'cloudian' user. The change password prompt will require you first to enter again your "current UNIX password" -- which is *Cloudian1!* -- and then your desired new password. Remember the new password, since you will need it whenever you want to log in to the HyperIQ virtual machine. (Note that this password for logging into the HyperIQ host machine is a different password than the password for logging into the HyperIQ web UI. You will set a HyperIQ UI password later in this installation procedure.)

Once you're logged in, an interactive HyperIQ installer will lead you through the process of installing the HyperIQ application and integrating it with your HyperStore cluster. The command syntax for launching the installer is:

```
$ iqinstall [--replace-ips 'PRIVATE_IP1=PUBLIC_IP1,PRIVATE_IP2=PUBLIC_IP2...']  
[-k FULL_PATH_TO_KEY_FILE] [--monitor-pool-nodes] [--with-managed-services] [--secondary-hiq]
```

Options:

- `--replace-ips 'PRIVATE_IP1=PUBLIC_IP1,PRIVATE_IP2=PUBLIC_IP2...'`

You must use this option if in your HyperStore system you configured your installation "survey file" (or "topology file" in HyperStore 8.0 and later) with IP addresses from a private back-end network -- which is an atypical survey (or topology) file configuration -- and your HyperIQ node cannot access that private network. With this option you specify, for each HyperStore node, the private IP address and its corresponding public IP address (which HyperIQ can access). This example is for a three-node HyperStore system:

```
$ iqinstall --replace-ips '10.10.210.2=10.10.1.212,10.10.210.3=10.10.1.213,10.10.210.4=10.10.1.214'
```

- `-k FULL_PATH_TO_KEY_FILE`

If you want HyperIQ to use an SSH key that you have already placed on the HyperIQ host to make SSH connections to HyperStore nodes -- rather than using an SSH password that you can supply during the install process as described further below -- launch the installer with the `-k` option. For example you can copy the SSH key from your HyperStore Configuration Master node (in that node's installation staging directory) over to your HyperIQ node, and then specify that file with the `-k` option.

If you have multiple SSH keys -- such as if you are having HyperIQ monitor multiple HyperStore clusters and you have different SSH keys for the different clusters; or if you are having HyperIQ monitor a HyperStore cluster that includes auxiliary nodes (HyperStore File Service nodes or HyperStore Search Service nodes) and you have a different key for the auxiliary nodes than for the rest of your HyperStore nodes -- you can specify multiple `-k` flags each with a different key, like `-k PATH_TO_KEY_FILE1 -k PATH_TO_KEY_FILE2`.

Note The option of using an SSH key rather than a password is supported only for *root* user access to HyperStore nodes. For accessing HyperStore as the *sa_admin* user (in the scenario where you've enabled the HyperStore Shell and disabled root access to your HyperStore nodes), only password authentication is supported currently.

- `--monitor-pool-nodes`

Include this option if your HyperStore system includes a HyperStore File Service pool and/or a HyperStore Search Service pool and you want HyperIQ to monitor the nodes in those pools (as well as monitoring your HyperStore object storage nodes).

If your HyperStore system version is 8.1.x and you are using the HyperStore Firewall, HyperIQ monitoring of File Service and/or Search Service pool nodes is not supported. The HyperStore Firewall in that version of HyperStore does not support providing the access to the pool nodes that would be required for HyperIQ to monitor them.

If your HyperStore system version is 8.2.x and you are using the HyperStore Firewall, HyperIQ monitoring of File Service and/or Search Service pool nodes is supported -- but only if you make the following changes to the HyperStore Firewall configuration, using the *hsctl* tool on your HyperStore Config Controller node (you can do this before or right after installing your HyperIQ application):


```
# Define a custom service for the HyperStore Firewall:
$ hsctl config create firewall.services.custom hiq_pool_nodes
$ hsctl config set firewall.services.custom.hiq_pool_nodes.ports='tcp/9999'

# EITHER allow ANY host to connect to port 9999:
$ hsctl config set firewall.services.custom.hiq_pool_nodes.allow=ANY

# OR define an ACL for the HyperIQ instance and allow only it to connect to port 9999
# (in the second command replace HIQ_NODE_IPADDR with the actual HIQ IP address):
$ hsctl config create firewall.acl hiq_instance
$ hsctl config set firewall.acl.hiq_instance.addrs='HIQ_NODE_IPADDR/32'
$ hsctl config set firewall.services.custom.hiq_pool_nodes.allow=hiq_instance

# Apply your configuration changes:
$ hsctl config apply firewall
```

- *--with-managed-services*

Include this option only if instructed to do so by Cloudian Support. This will be applicable only if you are a Support Plus, Premier Support, or HyperCare customer, and you have already received and installed (in your HyperStore system) the special HyperStore license file associated with these Support programs.

- *--secondary-hiq*

Include this option if this is a secondary HyperIQ instance, monitoring a HyperStore system that is also being monitored by another HyperIQ instance that you have already installed. This option facilitates a high availability (HA) HyperIQ deployment, whereby two separate HyperIQ instances monitor the same HyperStore system so that if one HyperIQ instance becomes unavailable the other HyperIQ instance can still be used. Note that although this installation option refers to a "secondary" HyperIQ, in fact both HyperIQ instances will be equal in their role and operation -- neither instance will be in a "primary" role and neither instance will be limited to a back-up or fail-over role. Put differently, the two HyperIQ instances will be deployed in an "Active / Active" configuration. (The only difference is in regard to certain aspects of metrics log processing; having just one of the two HyperIQ instances be installed as "*--secondary-hiq*" avoids an undesirable duplication of these tasks.)

Note If you are going to have two different HyperIQ instances monitoring the same HyperStore system, do not use the same hostname for both HyperIQ instances.

When you launch the installer you will be prompted to run the IP networking setup for your HyperIQ host. Answering yes to this prompt opens the Network Manager Text User Interface (*nmtui*). Using this standard Linux tool you can edit the host's Ethernet connection interface *ens32* (for example if you wish you can manually specify an IP address, gateway, and DNS servers -- by default these are automatically assigned by DHCP). If you edit the connection interface remember to apply the change by using the "Activate a connection" menu options to deactivate the connection and then activate it again. Also with the *nmtui* tool you can change the hostname if you wish (the default is *hyperiq*). For more information about the *nmtui* tool see any reliable online source.

After you've exited the *nmtui* dialog the installation proceeds to set up the core HyperIQ components as Docker containers. This includes containerized instances of the mature open source technologies **Prometheus** and **Grafana**, which HyperIQ uses for the collection, analysis, and visualization of HyperStore monitoring data. This automated setup process may take a few minutes.

Next will be a series of prompts for information that HyperIQ needs to integrate with your HyperStore cluster. This includes:

- A **unique name for your HyperStore cluster**. Only lower case letters, dashes, and underscores are allowed. Here, "cluster" means an entire HyperStore system (which depending on your deployment topology, may encompass multiple data centers or service regions). The name that you give to your cluster is only for purposes of organizing data within HyperIQ and does not need to match any name that you've configured within HyperStore. Note that if you are connecting HyperIQ to multiple HyperStore clusters (multiple independent HyperStore systems that you want to monitor), enter a name for one of the clusters here, and then later in the installation process -- after HyperIQ has been successfully integrated with the first cluster-- you will be given the opportunity to integrate HyperIQ with an additional cluster or clusters. If you do integrate HyperIQ with multiple HyperStore clusters you will need to give each HyperStore cluster a different name.

Note If this is a secondary HyperIQ instance -- monitoring a HyperStore cluster that is also being monitored by another HyperIQ instance that you've already installed -- use the same HyperStore cluster name that you did when you installed the first HyperIQ instance. The installer does not enforce this as a requirement, but it will be more logical for HyperIQ users if the HyperStore cluster has the same name in the UIs of both HyperIQ instances.

- The **IP address of one of your HyperStore nodes** within the cluster (it can be any node, and the IP address must be reachable by HyperIQ), and the user name (either *root* or *sa_admin*) and **SSH password** for connecting to that node. Note that the password prompt will not appear if you're using an SSH key instead. From the one HyperStore IP address that you provide, the HyperIQ installer will automatically obtain the IP addresses of all of the rest of your HyperStore nodes. The installer then installs Prometheus "exporters" on every HyperStore node. After this step completes, the installer finalizes the configuration of Prometheus and Grafana on the HyperIQ VM.
- The **IP address of your HyperBalance load balancer**, if you are using HyperBalance with your HyperStore cluster and you would like HyperBalance statistics to appear in HyperIQ. (See the Note about HyperBalance at the beginning of "**Installing HyperIQ as a Virtual Appliance**" (page 14).)

Lastly, the installer prompts you to create a **password for logging into the HyperIQ web UI**. For the password you can use alphanumeric characters, and also special characters if you wish. Remember this password, since you will need it whenever you want to access the HyperIQ web UI. You can also use this password to access the Prometheus web UI.

After the installer completes it returns you to the terminal prompt. Next you can use a browser to access the HyperIQ UI as described in "**Getting Started With HyperIQ**" (page 29).

Note If you have a very large HyperStore cluster it's possible that you may encounter copy timeout errors (displayed in your terminal) as the HyperIQ set-up process tries to concurrently copy package files to multiple HyperStore nodes through your network. In that case, launch the HyperIQ installer again but this time add the command option *-t 5* (for example *iqinstall -t 5*). This reduces the number of concurrent threads used during HyperIQ set-up (so that 5 threads are used rather than the default of 10) which will make the aggregate set-up time somewhat longer but should eliminate the copy time-out problem.

Note For its access to port 9999 on your HyperStore nodes, HyperIQ uses HTTPS (TLS v1.2) with Basic Authentication. The HyperIQ setup tool automatically configures these security mechanisms using a self-signed certificate for TLS and a randomly generated password for Basic Authentication.

2.3. Upgrading a HyperIQ Virtual Appliance

In recent HyperIQ releases, in the OVA for new installs there were updates made not only to the HyperIQ software but also to the underlying virtual machine attributes:

- The HyperIQ 2.1 OVA switched to Rocky Linux as the OS, versus Ubuntu in previous HyperIQ OVA versions
- The HyperIQ 2.0 OVA switched to using two virtual disks (one disk for the OS and application, and one disk for HyperIQ data), versus one virtual disk in previous HyperIQ OVA versions. It also switched to Ubuntu 22.04 as the OS, versus Ubuntu 18.04 in previous HyperIQ OVA versions.

The recommended virtual machine configuration, as implemented in the HyperIQ OVA from version 2.1 and later, is Rocky Linux as the OS, and two disks including a dedicated HyperIQ data disk.

The procedure for upgrading a HyperIQ virtual machine to HyperIQ version 2.2 depends on your starting point - in particular, whether you've already migrated to the recommended configuration with Rocky Linux and two disks. If your current HyperIQ VM is still on Ubuntu, the procedure also depends on whether you want to migrate to Rocky Linux as part of the upgrade to HyperIQ 2.2, or remain on Ubuntu. Migrating to Rocky Linux now is recommended, since starting with the next HyperIQ major release -- HyperIQ 2.3 -- running HyperIQ on Ubuntu will no longer be supported.

Follow whichever of the below procedures is applicable to your upgrade path. (If you're unsure of the OS and/or the number of disks for your current HyperIQ VM, you can log in to the VM and use -- for example -- the `hostnamectl` command to check the OS distribution and `lsblk` to check the drive information.)

Note Each of these upgrade path procedures will retain your HyperIQ monitoring data, and also any custom dashboards that you've created and any group or user configurations and alerting configurations.

- **"Starting From HyperIQ on Rocky Linux, Upgrading to HyperIQ 2.2"** (page 19)
- **"Starting From HyperIQ on Ubuntu with Two Disks, Upgrading to HyperIQ 2.2 on Rocky Linux"** (page 20)
- **"Starting From HyperIQ on Ubuntu with One Disk, Upgrading to HyperIQ 2.2 on Rocky Linux"** (page 23)
- **"Starting From HyperIQ on Ubuntu, Upgrading to HyperIQ 2.2 on Ubuntu"** (page 25)

2.3.1. Starting From HyperIQ on Rocky Linux, Upgrading to HyperIQ 2.2

If your current HyperIQ virtual machine has Rocky Linux as its OS, this necessarily means that it's running HyperIQ version 2.1 and that it has two disks (an OS and application disk, and a dedicated HyperIQ data disk). To upgrade to HyperIQ 2.2, follow the procedure below. This procedure will upgrade the HyperIQ software to version 2.2 but will not make any changes to the virtual machine -- it will continue to use the same operating system and the same number of disks as your existing HyperIQ VM.

Download the *hyperiq-2.2.bin* file (**not** the *.ova* file) from the Cloudian Support portal, put it into a working directory on your existing HyperIQ virtual appliance, and run these commands:

```
$ chmod +x hyperiq-2.2.bin
$ ./hyperiq-2.2.bin install [--replace-ips 'PRIVATE_IP1=PUBLIC_IP1,PRIVATE_IP2=PUBLIC_IP2...']
[-k KEY_FILENAME] [--monitor-pool-nodes] [--with-managed-services] [--secondary-hiq]
```

When upgrading, use the same installation command-line options that you did when you installed your existing HyperIQ application. For descriptions of the command options, see the **"Installing the HyperIQ Application"** (page 15) section.

Note In HyperIQ 2.1 and earlier, the installation command-line option *--monitor-pool-nodes* was called *--auxiliary-nodes*. So if you used the *--auxiliary-nodes* option when installing an existing HyperIQ 2.1 or older application, use the *--monitor-pool-nodes* option when running the *install* command to upgrade to HyperIQ 2.2.

After running the *install* command with the applicable options, follow the prompts to complete the upgrade.

2.3.2. Starting From HyperIQ on Ubuntu with Two Disks, Upgrading to HyperIQ 2.2 on Rocky Linux

If your current HyperIQ virtual machine has Ubuntu as its OS and has two disks (an OS and application disk, and a dedicated HyperIQ data disk), you can upgrade to HyperIQ 2.2 running on Rocky Linux (and still with two disks) by following the steps in this section.

This procedure upgrades your HyperIQ software to version 2.2, while also migrating to the Rocky Linux OS that is used in the HyperIQ 2.2 OVA. It accomplishes this by having you create a new VM from the HyperIQ 2.2 OVA file and then transfer the HyperIQ data disk from your existing old HyperIQ VM (running Ubuntu) to the new HyperIQ 2.2 VM (running Rocky Linux). After confirming a successful migration and successful operation of your new HyperIQ 2.2 VM, you will be able to delete your old HyperIQ VM.

1. In your hypervisor, create a backup of your old HyperIQ VM.
2. Log in to your old HyperIQ VM as the *cloudian* user and confirm that its data disk is mounted at */var/lib/docker*. For example (emphasis added):

```
cloudian@hyperiq:~$ lsblk
NAME        MAJ:MIN RM  SIZE RO TYPE MOUNTPOINTS
sda          8:0    0  48.8G  0 disk
├─sda1       8:1    0  48.7G  0 part /
├─sda14      8:14   0    4M    0 part
└─sda15      8:15   0   106M  0 part /boot/efi
sdb          8:16   0 146.5G  0 disk
└─sdb1       8:17   0 146.5G  0 part /var/lib/docker
```

3. Shut down your old HyperIQ VM.
4. In your hypervisor create a **new** VM using the HyperIQ 2.2 OVA file. (For instructions, see **"Deploying the HyperIQ VM"** (page 14). Do not yet perform the procedure for **"Installing the HyperIQ Application"** (page 15); you will do this as a later step as described further below.) When you've finished creating the new VM, log in to the new VM as the *cloudian* user. The default initial password is *Cloudian1!* -- you will be prompted to change this to a password of your choosing.

- ```
[cloudian@hyperiq ~]$ lsblk
```
- | NAME   | MAJ:MIN | RM | SIZE   | RO | TYPE | MOUNTPOINT      |
|--------|---------|----|--------|----|------|-----------------|
| sda    | 8:0     | 0  | 146.5G | 0  | disk |                 |
| └─sda1 | 8:1     | 0  | 146.5G | 0  | part | /var/lib/docker |
| sdb    | 8:16    | 0  | 48.8G  | 0  | disk |                 |
| └─sdb1 | 8:17    | 0  | 512M   | 0  | part | /boot/efi       |
| └─sdb2 | 8:18    | 0  | 1G     | 0  | part | /boot           |
| └─sdb3 | 8:19    | 0  | 7.9G   | 0  | part | [SWAP]          |
| └─sdb4 | 8:20    | 0  | 39.5G  | 0  | part | /               |

```
/etc/fstab
...
...
...
UUID=27cb40a9-cdf5-43db-a1f7-f7fd8b87023f /var/lib/docker ext4 defaults 1 1
```

6. Shut down the new HyperIQ 2.2 VM.
7. In your hypervisor:
  - a. Detach the (empty) data disk from your new HyperIQ 2.2 VM.
  - b. Make a copy of the (populated) data disk from your old HyperIQ VM, and attach it to the new HyperIQ 2.2 VM.
8. Restart the new HyperIQ 2.2 VM and log in as the *cloudian* user.
9. Determine where the newly attached data disk partition is. For example:

```
[root@hyperiq cloudian]# lsblk
```

| NAME   | MAJ:MIN | RM | SIZE   | RO | TYPE | MOUNTPPOINT |
|--------|---------|----|--------|----|------|-------------|
| sda    | 8:0     | 0  | 146.5G | 0  | disk |             |
| └─sda1 | 8:1     | 0  | 146.5G | 0  | part |             |
| sdb    | 8:16    | 0  | 48.8G  | 0  | disk |             |
| └─sdb1 | 8:17    | 0  | 512M   | 0  | part | /boot/efi   |
| └─sdb2 | 8:18    | 0  | 1G     | 0  | part | /boot       |
| └─sdb3 | 8:19    | 0  | 7.9G   | 0  | part | [SWAP]      |
| └─sdb4 | 8:20    | 0  | 39.5G  | 0  | part | /           |

10. Run `sudo -s` to elevate to root.
11. Run the following command to add the new data disk to `/etc/fstab`, replacing `DISK` with the full device name for your data disk's partition. (This command is all one line.)

For example, from Step 9 the data disk partition is *sda1*, the device name would be */dev/sda1*, and so the command would be:

21

12. Reboot the HyperIQ 2.2 VM, then log in as the *cloudian* user and confirm that the new data disk was mounted as */var/lib/docker*. For example:

```
[cloudian@hyperiq ~]$ lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda 8:0 0 146.5G 0 disk
└─sda1 8:1 0 146.5G 0 part /var/lib/docker
sdb 8:16 0 48.8G 0 disk
├─sdb1 8:17 0 512M 0 part /boot/efi
├─sdb2 8:18 0 1G 0 part /boot
├─sdb3 8:19 0 7.9G 0 part [SWAP]
└─sdb4 8:20 0 39.5G 0 part /
```

You now have a new HyperIQ 2.2 VM, running Rocky Linux 8.10, and using a data disk that is populated with data from your pre-existing HyperIQ deployment.

13. Next, install the HyperIQ 2.2 application. When performing the new installation:

- Use the same *iqinstall* command-line options that you previously used when you set up your old HyperIQ application, **plus** use the command-line option ***--os-migration*** (which is specific to the upgrade path that you are implementing and is **mandatory** for this upgrade path -- **if you omit this option the migration will fail**). For example, if for your old installation you ran the installer with no command-line options:

```
$ iqinstall
```

For this new VM run the installer as:

```
$ iqinstall --os-migration
```

Note that the command-line option that was called *--auxiliary-nodes* in HyperIQ 2.1 and earlier is now called *--monitor-pool-nodes*. So if for your old installation you ran the installer as:

```
$ iqinstall --auxiliary-nodes
```

For this new VM run the installer as:

```
$ iqinstall --monitor-pool-nodes --os-migration
```

**IMPORTANT !** Using *--os-migration* on the command-line is mandatory for this upgrade path, regardless of what other command-line options you use with *iqinstall*.

- During the interactive installation use the **same HyperStore cluster name** and the **same HyperStore node IP address** that you previously used when you set up your old HyperIQ application.
- You can use the same hostname as you did for your old HyperIQ VM or a different hostname, whichever you prefer (the default is *hyperiq*; the HyperIQ application installation gives you the chance to edit this and other IP networking settings).

For more details about the installation process see **"Installing the HyperIQ Application" (page 15)**.

To verify the success of the migration, log in to the HyperIQ 2.2 web UI on the new VM (if you don't recall how to do so, see **"Getting Started With HyperIQ" (page 29)**). Verify that the panels are populated with your HyperStore monitoring data, including data from prior to your upgrade to HyperIQ 2.2. Also verify that any customizations you had made in your old HyperIQ instance -- such as custom dashboards, users, or alerts -- have migrated over to the new HyperIQ instance.

After verifying that all your HyperIQ data and your customizations have successfully migrated over to the new HyperIQ instance, as a precaution you can keep your old HyperIQ VM (not running) for a week or so -- just to

confirm that things are going smoothly with your new HyperIQ 2.2 VM -- and then you can delete the old HyperIQ VM.

### 2.3.3. Starting From HyperIQ on Ubuntu with One Disk, Upgrading to HyperIQ 2.2 on Rocky Linux

If your current HyperIQ virtual machine has Ubuntu as its OS and has just one disk, you can upgrade to HyperIQ 2.2 running on Rocky Linux and having two disks (an OS and application disk, and a dedicated HyperIQ data disk) by following the steps below.

This procedure upgrades your HyperIQ software to version 2.2, while also migrating to the Rocky Linux OS that is used in the HyperIQ 2.2 OVA. It accomplishes this by having you create a new VM from the HyperIQ 2.2 OVA file and then copy the Docker volume data from your existing old HyperIQ VM (running Ubuntu) to the new HyperIQ 2.2 VM (running Rocky Linux). After confirming a successful migration and successful operation of your new HyperIQ 2.2 VM, you will be able to delete your old HyperIQ VM.

To perform this procedure you must first obtain from Cloudian **both** of these HyperIQ 2.2 packages:

- *hyperiq-2.2.ova*
- *hyperiq-2.2.bin*

After obtaining both of those package files, proceed as follows:

1. Working with your **old VM**:
  - a. In your hypervisor, create a backup of your old HyperIQ VM.
  - b. Upgrade your old HyperIQ VM to HyperIQ version 2.2, by following the procedure in **"Starting From HyperIQ on Ubuntu, Upgrading to HyperIQ 2.2 on Ubuntu"** (page 25). This will upgrade your HyperIQ application to version 2.2 without making any changes to the VM attributes (it will still be running Ubuntu as its OS, and will have just one disk). When you have completed that procedure, return to this procedure and continue to Step 2.
2. Creating and working with the **new VM**:
  - a. In your hypervisor create a new VM using the HyperIQ 2.2 OVA file. (For instructions, see **"Deploying the HyperIQ VM"** (page 14). Do not yet perform the procedure for **"Installing the HyperIQ Application"** (page 15); you will do this as a later step as described further below.) When you've finished creating the new VM, log in to the new VM as the *cloudian* user. The default initial password is *Cloudian1!* -- you will be prompted to change this to a password of your choosing. When deploying this new HyperIQ 2.2 VM, **use the same hostname as your old HyperIQ VM but a different IP address**. You can run *sudo nmtui* on the VM to configure its hostname and IP address. This new HyperIQ VM must be reachable by your old HyperIQ VM through your network (since later in this procedure you will copy data from the old VM to the new VM, through the network). When you have finished deploying the new HyperIQ VM leave the new VM up and running, remain logged in to the new VM, and return to this procedure to continue to Step 2b.
  - b. While logged in to the new HyperIQ VM, run the following command to create Docker volumes that will be used by HyperIQ (this command is all one line):

```
for vol in hyperiq_alertmanager_config hyperiq_alertmanager_data hyperiq_etcd_data hyperiq_grafana_config hyperiq_grafana_data hyperiq_iqmgr_config
```

```
hyperiq_nginx_config hyperiq_postgres_data hyperiq_prometheus_config
hyperiq_prometheus_data hyperiq_remotewriter_config hyperiq_remotewriter_
data; do docker volume create $vol; done
```

- c. Verify that all the Docker volumes listed in the step above (12 volumes in total) have been created:

```
docker volume ls
```

- d. Stop the Docker daemon on the new VM:

```
sudo systemctl stop docker
```

- e. Set a password for the *root* user on the new HyperIQ VM (you will need this *root* password in Step 3c when you copy Docker volume data from the old VM to the new VM):

```
sudo passwd root
```

3. Working with your **old VM**:

- a. Logged in to your old HyperIQ VM, stop the HyperIQ Docker containers:

```
cd ~/.config/cloudian/hyperiq/; docker-compose stop
```

- b. Stop the Docker daemon on the old VM:

```
sudo systemctl stop docker
```

- c. Copy Docker volume data from the old VM to the new VM. In the command below replace *NEW-HYPERIQ-IP* with the IP address from the new HyperIQ VM. (This command is all one line.)

```
sudo rsync -av --exclude 'metadata.db' --exclude 'backingFsBlockDev'
/var/lib/docker/volumes/ root@NEW-HYPERIQ-IP:/var/lib/docker/volumes/
```

When prompted, supply the *root* password for the new VM.

4. On the **new VM**, install the HyperIQ 2.2 application. When performing the new installation:

- Use the same *iqinstall* command-line options that you previously used when you set up your old HyperIQ application, **plus** use the command-line option ***--os-migration*** (which is specific to the upgrade path that you are implementing and is **mandatory** for this upgrade path -- **if you omit this option the migration will fail**). For example, if for your old installation you ran the installer with no command-line options:

```
$ iqinstall
```

For this new VM run the installer as:

```
$ iqinstall --os-migration
```

Note that the command-line option that was called *--auxiliary-nodes* in HyperIQ 2.1 and earlier is now called *--monitor-pool-nodes*. So if for your old installation you ran the installer as:

```
$ iqinstall --auxiliary-nodes
```

For this new VM run the installer as:

```
$ iqinstall --monitor-pool-nodes --os-migration
```

**IMPORTANT !** Using *--os-migration* on the command-line is mandatory for this upgrade path, regardless of what other command-line options you use with *iqinstall*.



- During the interactive installation use the **same HyperStore cluster name** and the **same host-name** and the **same HyperStore node IP address** that you previously used when you set up your old HyperIQ application.

For more details about the installation process see **"Installing the HyperIQ Application" (page 15)**.

To verify the success of the migration, log in to the HyperIQ 2.2 web UI on the new VM (if you don't recall how to do so, see **"Getting Started With HyperIQ" (page 29)**). Verify that the panels are populated with your HyperStore monitoring data, including data from prior to your upgrade to HyperIQ 2.2. Also verify that any customizations you had made in your old HyperIQ instance -- such as custom dashboards, users, or alerts -- have migrated over to the new HyperIQ instance.

Next, log in to any HyperStore node and verify that in this file, the IP address is correctly the IP address of your new HyperIQ VM:

```
cat /etc/hyperiq/grafana-agent/agent.yml
```

After verifying that all your HyperIQ data and your customizations have successfully migrated over to the new HyperIQ instance, as a precaution you can keep your old HyperIQ VM (not running) for a week or so -- just to confirm that things are going smoothly with your new HyperIQ 2.2 VM -- and then you can delete the old HyperIQ VM.

### 2.3.4. Starting From HyperIQ on Ubuntu, Upgrading to HyperIQ 2.2 on Ubuntu

Remaining on Ubuntu when you upgrade to HyperIQ 2.2 is not recommended, but is allowed. Note that **starting with (future release) HyperIQ 2.3, running HyperIQ on Ubuntu will no longer be supported**. If when HyperIQ 2.3 comes out you haven't yet migrated to Rocky Linux, you will need to do so when you upgrade to HyperIQ 2.3.

To upgrade to HyperIQ 2.2 while remaining on Ubuntu, follow the procedure below. This procedure will upgrade the HyperIQ software to version 2.2 but will not make any changes to the virtual machine -- it will continue to use the same operating system and the same number of virtual disks as your existing HyperIQ VM.

Download the *hyperiq-2.2.bin* file (**not** the .ova file) from the Cloudian Support portal, put it into a working directory on your HyperIQ virtual appliance, and run these commands:

```
$ chmod +x hyperiq-2.2.bin
$./hyperiq-2.2.bin install [--replace-ips 'PRIVATE_IP1=PUBLIC_IP1,PRIVATE_IP2=PUBLIC_IP2...']
[-k KEY_FILENAME] [--monitor-pool-nodes] [--with-managed-services] [--secondary-hiq]
```

When upgrading, use the same installation command-line options that you did when you installed your existing HyperIQ application. For descriptions of the command options, see the **"Installing the HyperIQ Application" (page 15)** section.

**Note** In HyperIQ 2.1 and earlier, the installation command-line option *--monitor-pool-nodes* was called *--auxiliary-nodes*. So if you used the *--auxiliary-nodes* option when installing an existing HyperIQ 2.1 or older application, use the *--monitor-pool-nodes* option when running the *install* command to upgrade to HyperIQ 2.2.

After running the *install* command with the applicable options, follow the prompts to complete the upgrade.

## 2.4. Installing or Upgrading HyperIQ in an Existing Docker Environment

**Note** If you are using the HyperBalance load balancer with your HyperStore system and **if you want HyperIQ to monitor HyperBalance** (as well as monitoring your HyperStore system), before installing HyperIQ you must enable the Prometheus exporter on each of your HyperBalance nodes. On each HyperBalance node, log in to the HyperBalance UI, go to **Cluster Configuration -> Layer 7 -- Advanced Configuration** and then check the "Enable Prometheus Exporter" checkbox.

**Note** If you previously installed HyperIQ in an existing Docker environment and now you want to **upgrade** to HyperIQ version 2.2, follow the steps in this section. The upgrade procedure is the same as the procedure for doing a fresh installation.

If the Docker environment in which you are installing HyperIQ is a VM dedicated solely to HyperIQ, the following VM configuration is recommended:

- 2 CPU cores
- 8GB RAM
- Two disks: 50GB root disk and a 150GB disk mounted to `/var/lib/docker` (for storing all HyperIQ data)
- OS = Rocky Linux 8.10, 64-bit

If you are installing HyperIQ in a shared Docker environment, it is recommended that the CPU, RAM, and disk resources listed above be available to HyperIQ.

To install HyperIQ version 2.2 in a Docker environment:

1. On the host machine on which you want to install HyperIQ, log in as a user who has permissions to run `docker` commands. In a working directory on the host machine, download the HyperIQ installer (`hyperiq-2.2.bin`) from the Cloudian Support portal.
2. Change into the working directory in which you've downloaded the HyperIQ installer file, if you're not already in that directory.
3. Run the following command to make the installer file executable:

```
chmod +x hyperiq-2.2.bin
```

4. Start the installation:

```
./hyperiq-2.2.bin install [--replace-ips 'PRIVATE_IP1=PUBLIC_IP1,PRIVATE_IP2=PUBLIC_IP2...']
[-k FULL_PATH_TO_KEY_FILE] [--monitor-pool-nodes] [--with-managed-services] [--secondary-hiq]
```

For description of these install command line options and for description of the remainder of the installation process please see **"Installing the HyperIQ Application"** (page 15), starting from that section's description of the installation command-line options. Note that:

- If you are **upgrading** your HyperIQ application in your own existing Docker environment, use the same installation command-line options that you did when you installed your existing HyperIQ application.

**Note** In HyperIQ 2.1 and earlier, the installation command-line option `--monitor-pool-nodes` was called `--auxiliary-nodes`. So if you used the `--auxiliary-nodes` option when installing an existing

HyperIQ 2.1 or older application, use the `--monitor-pool-nodes` option when running the `install` command to upgrade to HyperIQ 2.2.

- After completing the HyperIQ installation (or upgrade), leave the (`hyperiq-2.2.bin` file in the working directory in which you downloaded it -- you will need this file if at some point in the future you want to uninstall HyperIQ.

## 2.5. Using an Existing Grafana Deployment to Display HyperIQ Dashboards

To use an existing Grafana deployment to display HyperIQ dashboards you must be running **Grafana 9.x** or newer. HyperIQ 2.x is developed and tested on Grafana version 11.4. This procedure requires installing and running a HyperIQ instance, from which you will import dashboards.

1. Install HyperIQ on a host as described in **"Installing HyperIQ as a Virtual Appliance"** (page 14) or **"Installing or Upgrading HyperIQ in an Existing Docker Environment"** (page 26).
2. In your existing Grafana UI, go to **Configuration -> Data Sources**. Click **Add data source**.
3. Select Prometheus as the data source type. When configuring the data source:
  - The data source name must be *Prometheus*
  - The data source URL by default is `http://<HyperIQ-IPAddress>:3000/prometheus`
  - Basic Auth must be enabled, with user name *cloudian\_admin* and the HyperIQ/Prometheus web UI password that you set during HyperIQ installation.
4. After Saving the data source configuration, select the **Dashboards** tab for the new data source and import the HyperIQ dashboards.

## 2.6. Uninstalling HyperIQ

Uninstalling HyperIQ will stop and remove Docker containers, images, networks and volumes, and the HyperIQ logs and tools directory. **Your HyperIQ data will be lost** (you will lose accumulated HyperStore monitoring data stored in the HyperIQ database, from the period during which you've had HyperIQ running).

To uninstall HyperIQ, log in to the HyperIQ host and run the appropriate command for your HyperIQ deployment type as indicated below.

**Note** By default, the uninstall operation will also remove the Prometheus exporters from all your HyperStore nodes (these exporters were installed on your HyperStore nodes during the HyperIQ setup process). If you want to leave these exporters in place on your HyperStore nodes, use the `--keep-exporters` option.

HyperIQ virtual appliance:

```
iquninstall [--keep-exporters]
```

HyperIQ installed in an existing Docker environment:

Change into the directory in which you downloaded the `hyperiq-2.2.bin` file before you installed

HyperIQ, and then run:

```
./hyperiq-2.2.bin uninstall [--keep-exporters]
```

# Chapter 3. Getting Started With HyperIQ

This chapter covers the following topics:

- **"Connecting to the HyperIQ UI"** (page 29)
- **"Viewing HyperIQ Dashboards and Panels"** (page 29)
- **"Adding HyperIQ Users"** (page 34)

## 3.1. Connecting to the HyperIQ UI

To connect to the HyperIQ UI point a browser to either of the following URLs:

*http://<HyperIQ IP address>:3000*

*https://<HyperIQ IP address>:3443* (this uses a self-signed TLS certificate)

HyperIQ comes with two pre-configured users:

- **cloudian\_admin** -- This user has full privileges including the ability to view and edit dashboards, and to create new HyperIQ users. The password for this user is the HyperIQ UI password that you specified during HyperIQ set-up.
- **org\_editor** -- This user has the ability to view and edit dashboards, but not to create new users. The default password for this user is *admin*. When logging in as this user you will be prompted to change the user's password, which you can do within the flow of the login process.

If you are logging in shortly after completing the installation of HyperIQ, note that it may take several minutes for HyperIQ to become populated with HyperStore monitoring data.

## 3.2. Viewing HyperIQ Dashboards and Panels

When you've logged in to HyperIQ for the first time you are at the **HyperIQ Main** dashboard. This is one of several HyperIQ dashboards:

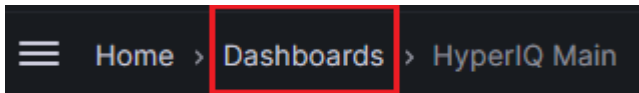
- **HyperIQ Main** -- HyperStore host OS level metrics such as CPU and memory utilization, storage capacity consumption and availability, and front-end and back-end network traffic rates.
- **Hardware Info** -- Intelligent Platform Management Interface (IPMI) sensor data such as hardware component temperatures, fan speeds, and power usage. This data is available in HyperIQ only if your HyperStore host hardware supports IPMI and each HyperStore host has *ipmitool* installed. For example, all HyperStore Appliance machines support IPMI and have *ipmitool* installed.
- **HyperIQ Self Monitoring** -- Dashboard for monitoring the activity and status of the HyperIQ node itself.
- **Services** -- HyperStore service level metrics such as up/down status and resource consumption for each service that runs on HyperStore hosts, such as the S3 Service, the Admin Service, Cassandra, and the Redis Credentials DB and Redis QoS DB.
- **Cassandra Status -- Basic** -- Basic Cassandra metrics such as service status, keyspace size, and resource consumption. (Cassandra is the NoSQL distributed database technology on which HyperStore's Metadata DB is built, for storing object metadata and system metadata.)

- **Cassandra Status -- Advanced** -- Advanced Cassandra metrics including keyspace operations latencies; details regarding Memtables, SSTables, and compaction; and more. Useful if support engineers are deep-diving into system performance issues.
- **Disk Info** -- Usage information for data disks and metadata disks on your HyperStore nodes.
- **Operation Status** -- Status metrics for HyperStore cluster operations such as repair, cleanup, and rebalance.
- **S3 Analytics** -- HyperStore S3 usage metrics categorized by user, user group, bucket, and S3 operation type (including cross-region replication if applicable). The S3 Analytics dashboard is available only if your HyperStore system license includes support for the Enterprise version of HyperIQ. For information about acquiring a HyperStore license that supports the Enterprise version of HyperIQ, contact Cloudian Support.

To switch between dashboards, in the upper part of the UI click **HyperIQ Dashboards**, and then select a dashboard from the list that displays.

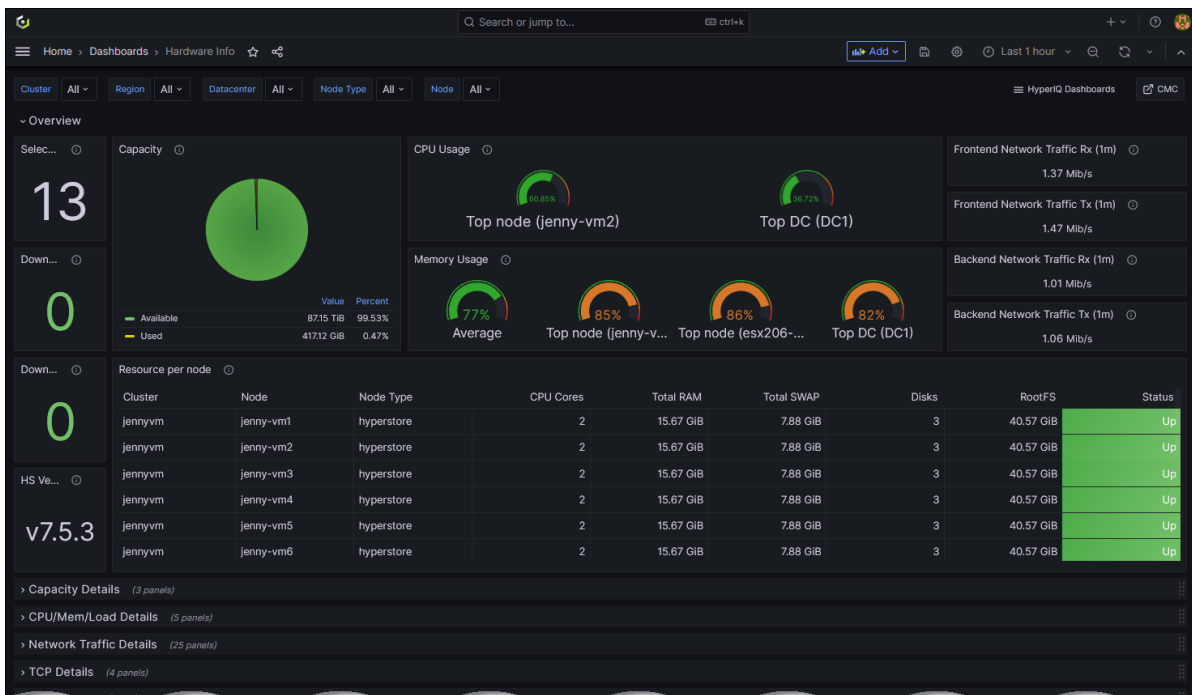


You can also switch between dashboards by clicking **Dashboards** in top left of the UI and then selecting a dashboard.

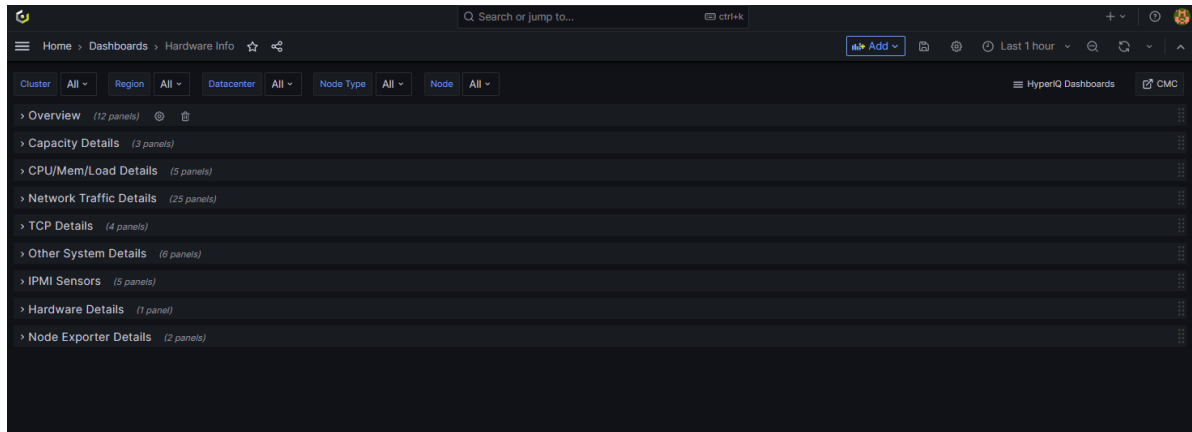


Each dashboard consists of a set of "panels", with each panel providing a graphical representation of one metric (such as CPU Utilization or Top S3 Operations). In each dashboard, panels are grouped into "rows" that work like folders. When you first access a dashboard, the first panel row is open so that the panels from that row are all displayed. To see the full list of panel rows available in a dashboard, collapse the open row.

**Hardware Info** dashboard with **Overview** panel row open:




**Hardware Info** dashboard with all panel rows collapsed:

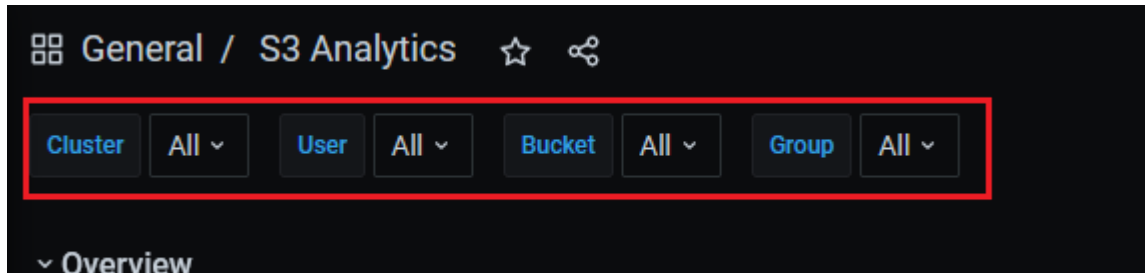


### 3.2.1. Filtering the Data in a Dashboard

For each dashboard, in the top left of the HyperIQ UI you can filter the data in ways relevant to that dashboard. For example in the **HyperIQ Main** dashboard you can filter data by **Cluster**, **Region**, **Datacenter**, **Node Type**, and **Node**. The Cluster filter is relevant only if you are using HyperIQ to monitor more than one HyperStore system. The Node Type filter is relevant only if you are using HyperIQ to monitor HyperStore File nodes and/or HyperStore Search nodes as well as regular HyperStore object storage nodes.



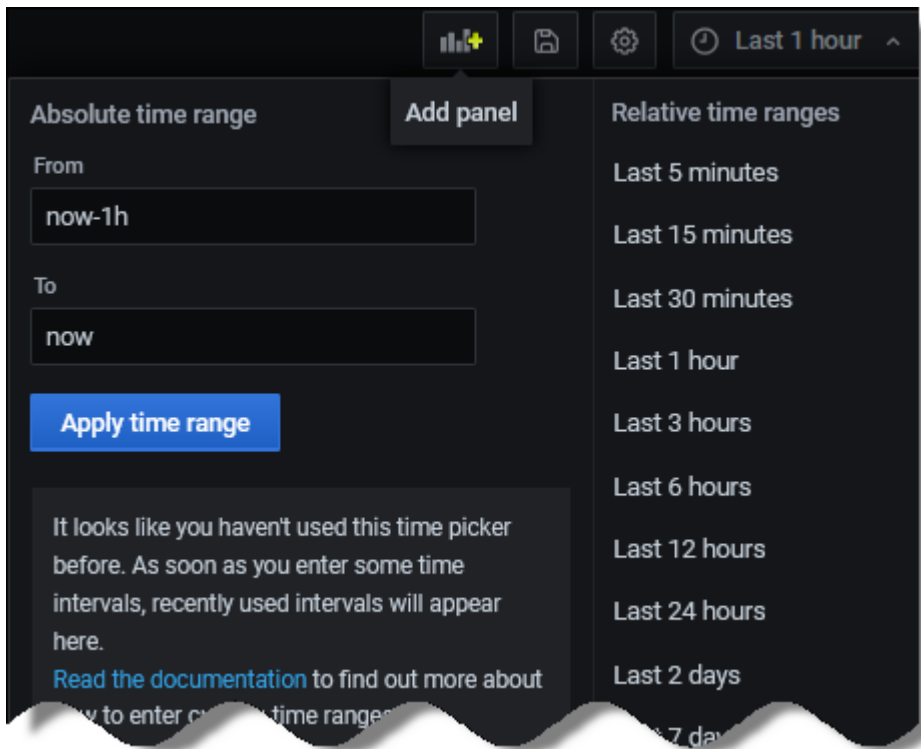
As a second example, in the **S3 Analytics** dashboard you can filter data by **Cluster**, **User**, **Bucket**, and **Group**.



If you filter a dashboard's data using the filters on the top left of the HyperIQ UI, your chosen filter applies to **all panels in the dashboard**.

### 3.2.2. Changing the Time Range for a Dashboard


In the top right of the HyperIQ UI, you can change the time period for which a dashboard's panels display historical data. By default, such panels display data for the past hour. If you change the time period the change applies to all time period sensitive panels on the dashboard.



**Note** By default HyperIQ stores HyperStore monitoring data for only 30 days and so HyperIQ does not support looking back further than 30 days into your HyperStore system's history. The retention period -- and therefore the length of your view back into HyperStore history -- is configurable as described in **"Using the iqmgr Tool to Manage HyperIQ"** (page 39).


### 3.2.3. Using an Panel Legend to Focus the Data in a Graph Panel

Most graph panels include a legend that indicates the entities for which data has been graphed -- for example, the nodes for which data has been graphed. In this image data has been graphed for four nodes and those nodes are listed in the legend:



You can click on an entity in the legend to have the graph show the data for just that entity:





To have the graph return to showing data for all the entities, deselect the selected entity by clicking it again.


Most legends are in the form of tables that show some statistical highlights of the data (such as max, mean, and last value for each entity). You can click a column heading to sort the table by that column's data.

### 3.2.4. Viewing a Larger Version of a Panel

To enlarge a panel, hover your cursor on the panel and press the letter "v" on your keyboard. To exit the enlarged view, press the Esc key.

### 3.2.5. Viewing a Description of the Data in a Panel


For brief information about the data in a panel, hold your cursor over the "i" to the right of the panel name.



**Note** In the current release, not all panels have an "i" icon, and for those that do the thoroughness of the descriptive text varies from panel to panel.

### 3.2.6. Opening the Clodian Management Console

In the upper part of the UI you can click the **CMC** link to open the Clodian Management Console for your HyperStore system, in a separate browser tab.



**Note** For the CMC link to work, your HyperStore system's CMC service endpoint must be configured as an entry in your environment's DNS set-up.

### 3.2.7. Opening LogInsight

In the upper part of the UI you can click the **LogInsight** link to open the Clodian LogInsight web UI, if you have installed and set up LogInsight. For more information see **"Installing LogInsight"** (page 59) (including the post-installation step **"Linking from the HyperIQ UI to the LogInsight UI"** (page 68)).



### 3.2.8. For More Information About Working with Dashboards and Panels

HyperIQ is built on Grafana data visualization technology and exposes much of the Grafana native functionality. For more information about Grafana capabilities see the Grafana official documentation or community resources, which are available on the web.

## 3.3. Adding HyperIQ Users


The pre-configured user *cloudian\_admin* is a "Super Admin" user with full HyperIQ permissions, and is the only user who can add new HyperIQ users. When you are logged in as the *cloudian\_admin* user you can create new HyperIQ users and then assign each user one of three roles with differing permission levels:

- **Admin:** Can view, edit, and create dashboards and panels; and manage existing users (such as by changing a user's role or deleting a user). Cannot add new users.
- **Editor:** Can view, edit, and create dashboards and panels. Cannot add or manage users. The pre-configured user *org\_editor* has the Editor permission level.
- **Viewer:** Can view dashboards and panels. Cannot edit or create dashboards and panels, or add or manage users.

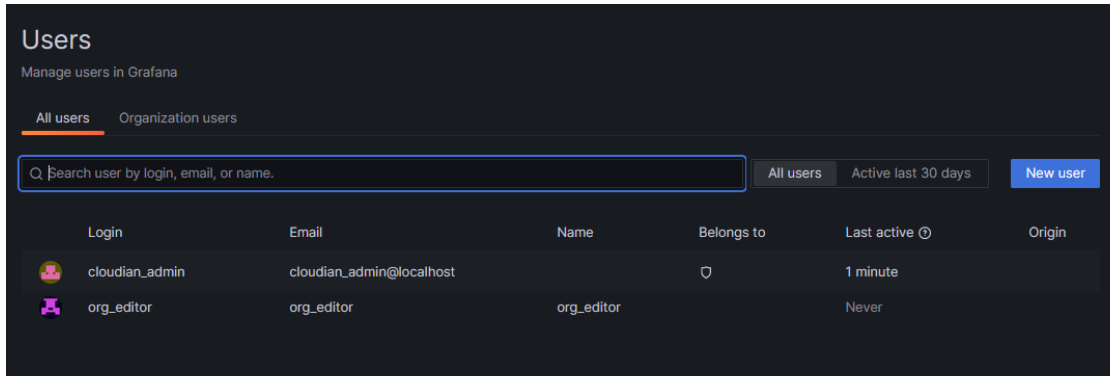
### 3.3.1. Adding a HyperIQ User

**Note** New users created in the way described below can only log into HyperIQ's Grafana UI (to access dashboards) -- they cannot log into the Prometheus UI that is described in the **"Administering HyperIQ"** (page 39) section of this document. If you want to create a user who can log into the HyperIQ Grafana UI and also the Prometheus UI, don't create the user through the HyperIQ Grafana UI -- instead use the method described in **"Using the iqmgr Tool to Manage HyperIQ"** (page 39). Once you've created the user you can use the Grafana UI to manage the user, including changing the user's role if you wish.

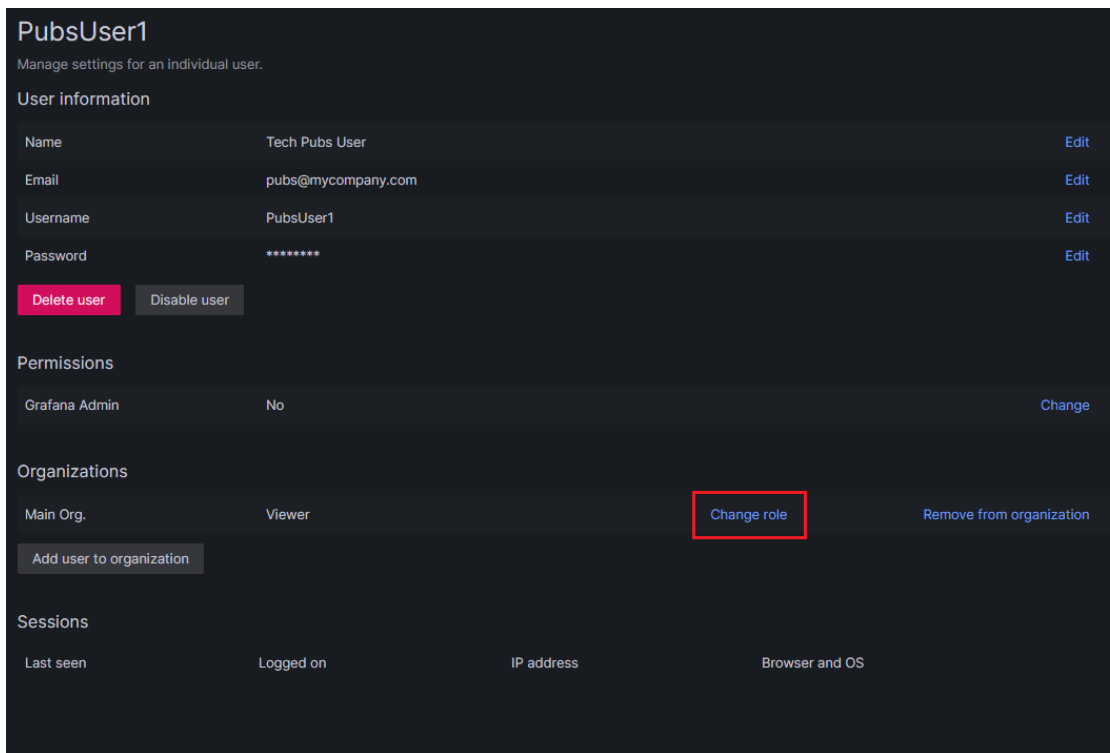
1. In the upper left of the HyperIQ UI, click the menu icon.



2. In the drop-down menu that displays, select **Administration**. Then in the Administration menu select **Users and Access**, and then select **Users**. This displays the **Users** management interface:



3. Click **New user** to display the user creation interface. Enter the information for the new user and then click **Create user**.
4. When you finish creating the user, a **User information** screen for that user displays. Here you can change the user's role if you wish. By default newly created users have the Viewer role.



If at a future point in time you want to change a user's role you can do so by going to the **Users** management interface and clicking the user's row in the user list. This opens the **User information** screen for that user, and there you can change the user's role.

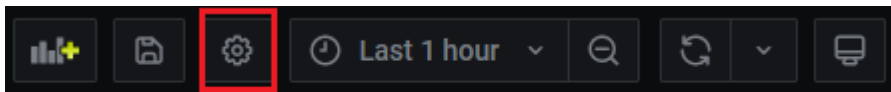
**Note** The **Users** management interface's **Invite** function -- for sending email invitations that enable recipients to become HyperIQ users if they wish to -- is not supported in HyperIQ.

### 3.3.2. Assigning Permissions for a Specific Dashboard

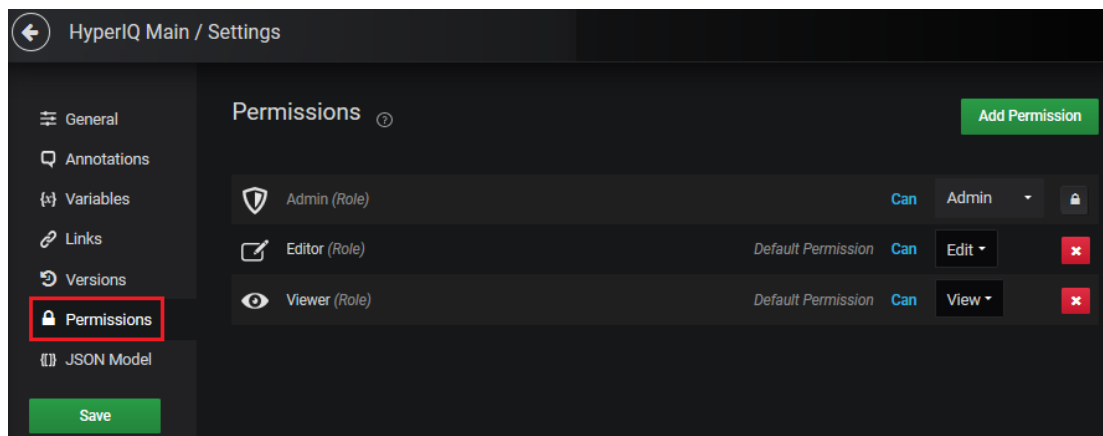
Depending on your HyperIQ use case, simply assigning users one of the three roles (Admin, Editor, or Viewer) may be sufficient for managing HyperIQ user permissions. However, if you wish you can implement more granular permissions, at the dashboard level.

An example would be if you want users with the Viewer role to be able to view only the S3 Analytics dashboard and not the other dashboards. To do so:

1. Go to the HyperIQ Main dashboard, and then in the upper right of the UI click the Dashboard Settings icon:



2. In the HyperIQ Main / Settings page select **Permissions**.



3. In the Viewer role row, click the red X at the right to remove the Viewer role from the list of roles that have permissions for this dashboard.
4. Repeat Steps 1 through 3 for each of the other dashboards **except** the S3 Analytics dashboard.

As a variation of the example above, you could set things up so that most users with the Viewer role can view only the S3 Analytics dashboard but certain chosen users could also have view-only permissions for the HyperIQ Main dashboard. To do this, follow the steps above, but for the HyperIQ Main dashboard, after removing the Viewer role from the permissions list click **Add Permission**. Then select the user to whom to grant permission, select View as the permission level, and click **Save**.

Permissions ⓘ

Add Permission

Add Permission For

User

gordon

Can

View

Save

Admin (Role)

Can

Admin

Editor (Role)

Can

Edit

This page left intentionally blank

# Chapter 4. Administering HyperIQ

This chapter covers the following topics:

- **"Using the *iqmgr* Tool to Manage HyperIQ"** (page 39)
- **"Using Docker Commands"** (page 44)
- **"Configuring LDAP Authentication for HyperIQ Users"** (page 44)
- **"Alerts and Notifications"** (page 46)

## 4.1. Using the *iqmgr* Tool to Manage HyperIQ

HyperIQ includes a management tool named *iqmgr*. The *iqmgr* tool is stored under this directory on your HyperIQ node:

```
$HOME/.config/cloudian/hyperiq
```

- If you installed HyperIQ as a virtual appliance, you can run *iqmgr* from any directory.

**Note** To log into the HyperIQ virtual appliance, use the username **cloudian** and the virtual appliance login password that you set when you installed the virtual appliance.

- If you installed HyperIQ in an existing Docker environment, you must change into the *\$HOME/.config/cloudian/hyperiq* directory to run *iqmgr*, or else include that directory path when running *iqmgr*. (Or for convenience you could, for example, create a symlink to the tool.)

With the *iqmgr* tool you can perform several types of HyperIQ management tasks:

- **"Updating HyperIQ to Monitor Additional HyperStore Resources"** (page 39)
- **"Changing HyperIQ Configuration Settings"** (page 40)
- **"Managing HyperIQ Users"** (page 42)
- **"Additional *iqmgr* Commands"** (page 43)

For some simple *iqmgr* commands you will specify command arguments, and simply execute the command. For *iqmgr* command that perform more complex tasks, when you launch the command you will be led through a series of interactive prompts.

### 4.1.1. Updating HyperIQ to Monitor Additional HyperStore Resources

**Update HyperIQ because you've made changes to the HyperStore cluster that HyperIQ has been monitoring:**

```
iqmgr setup [-k KEY_FILENAME] [--monitor-pool-nodes] [--with-managed-services] [--secondary-hiq]
```

Do this if you've changed your HyperStore cluster composition. For example, run *iqmgr setup* if you:

- Add a node or nodes to your HyperStore cluster
- Remove a node from your HyperStore cluster

- Add HyperStore File nodes and/or HyperStore Search nodes to your HyperStore cluster (and you want to monitor these "pool" nodes; note that in HyperIQ 2.1 and earlier this option was named "--auxiliary-nodes")
- Add a HyperBalance load balancer to your HyperStore cluster (and you want to monitor HyperBalance)

**Note** Before updating your HyperIQ setup to monitor HyperBalance you must enable the Prometheus exporter on each of your HyperBalance nodes. On each HyperBalance node, log into the HyperBalance UI, go to **Cluster Configuration -> Layer 7 -- Advanced Configuration** and then check the "Enable Prometheus Exporter" checkbox.

For descriptions of the command options see **"Installing the HyperIQ Application"** (page 15).

**Note** When prompted for whether you want to refresh the local *cloudianservicemap.json* file, answer **yes** (unless you have HyperStore on a private network to which HyperIQ is not attached and you've manually edited HyperIQ's copy of the service map file as described in the HyperIQ installation procedure).

**Have HyperIQ monitor an additional HyperStore cluster:**

```
iqmgr config cluster add [-k KEY_FILENAME] [--monitor-pool-nodes] [--with-managed-services]
[--secondary-hiq]
```

A "cluster" here means an entire HyperStore system, which may encompass one or more data centers and one or more service regions. So you would use this option if you want to have HyperIQ monitor an additional entire HyperStore system, as well as the HyperStore system that it's already been monitoring.

For descriptions of the command options see **"Installing the HyperIQ Application"** (page 15).

**Remove HyperIQ monitoring from a HyperStore cluster:**

```
iqmgr config cluster remove
```

### 4.1.2. Changing HyperIQ Configuration Settings

**Change the *cloudian\_admin* user's password for logging in to the HyperIQ Grafana UI:**

```
iqmgr config password
```

When you run this command you will be prompted to supply the new password. Do not include the password on the command line when launching the command.

This password is also used for logging in to the Prometheus UI.

**Note** Use this method for changing the *cloudian\_admin* user's password -- **not** the "Change Password" method that is available in the lower left of the HyperIQ UI. For other users (including the pre-configured *org\_editor* user), the Change Password method in the UI is the correct method to use.

**Configure SMTP integration for email notifications, or configure LDAP integration for the admin user:**

```
iqmgr config set <setting> <value>
```

Set a configuration setting to your desired value.

```
iqmgr config get <setting>
```



Get the current value of a setting.

```
iqmgr config unset <setting>
```

Revert a setting to having a null value.

```
iqmgr config apply
```

Apply a change that you've made with *iqmgr config set* or *iqmgr config unset*. This is mandatory -- using the *set* or *unset* commands will have no effect unless you follow it by using the *apply* command.

```
iqmgr config list
```

List all settings.

For setting details see **"Setting Up Notification Channels"** (page 52) and **"Configuring LDAP Authentication for HyperIQ Users"** (page 44).

**Configure for how many days to retain HyperStore monitoring data in the HyperIQ database:**

```
iqmgr config retention <n>d
```

Your *<n>* value must be an integer. For example, to retain day for 60 days:

```
iqmgr config retention 60d
```

After the retention period the data is automatically deleted. The default is 30 days.

**Note** The data retention period will determine how far back into time you can look when analyzing HyperStore statistics through HyperIQ. The longer the retention period, the more storage capacity is required for the HyperIQ database. Consult with Clodian Support if you have questions about HyperIQ storage sizing requirements for your use case and desired data retention period.

**Change the HyperIQ UI listening port numbers:**

```
iqmgr httpcfg --http-port <integer>
```

```
iqmgr httpcfg --https-port <integer>
```

By default the HTTP port is 3000 and the HTTPS port is 3443.

If you change the listening port numbers for the HyperIQ main UI the change will apply also to the Prometheus UI.

**IMPORTANT !** The HyperIQ HTTPS port is also used by HyperStore nodes, for transmitting data from HyperStore to HyperIQ. Make sure that your network allows each HyperStore node to access HyperIQ's HTTPS port.

**Configure HyperIQ to listen only on HTTPS and not on regular HTTP:**

```
iqmgr httpcfg -s
```

**Note** HyperIQ's HTTPS listener will accept connections from browsers using TLS versions 1.2 or 1.3. Connections using TLS 1.0 or 1.1 are not allowed.

**Create a New Self-Signed TLS Certificate for the HyperIQ HTTPS listener:**

```
iqmgr httpcfg --self-signed
```

HyperIQ automatically generates a self-signed TLS certificate to be used by the HyperIQ HTTPS listener. You can run the command above to trigger HyperIQ to generate a new self-signed certificate and use the new certificate for its HTTPS listen.

**Have the HyperIQ HTTPS listener use a TLS certificate that you provide:**

```
iqmgr httpcfg -c <full path to cert file> -k <full path to key file>
```

You must copy the certificate file and private key file on to the HyperIQ node before running the command above. The certificate and private key must be in .PEM format and the key must not have a passphrase.

**Change HyperIQ virtual appliance network settings:**

```
sudo iqmgr network config
```

This command for changing network settings is applicable only to the virtual appliance and not to HyperIQ installed in an existing Docker environment. The command opens the Network Manager Text UI (*nmtui*), a standard Linux tool for managing IP networking for a host machine.

To check the current network settings for the virtual appliance:

```
iqmgr network status
```

**Set the LogInsight URL for the HyperIQ UI's "LogInsight" link:**

```
iqmgr log-insight set <full-LogInsight-URL-including-port>
```

For more information see **"Linking from the HyperIQ UI to the LogInsight UI"** (page 68).

### 4.1.3. Managing HyperIQ Users

**Add a user, delete a user, or list users:**

```
iqmgr user add <username>
```

```
iqmgr user delete <username>
```

```
iqmgr user list
```

This way of adding users creates users who can access the HyperIQ Grafana UI **and also** the Prometheus UI (described in **"Alerts and Notifications"** (page 46)). If you want to create users who can access only the HyperIQ Grafana UI, don't use *iqmgr* -- instead use the Grafana UI's user management page to create the users as described in **"Adding a HyperIQ User"** (page 34).

Note that:

- When you run any of these commands you will be prompted to provide the *cloudian\_admin* user's password. To avoid being prompted for this password, you can include the argument *-g <password>* when running the command, where *<password>* is the *cloudian\_admin* user's password.
- When you run the *iqmgr user add* command you will be prompted to create (enter) a password for the new user. To avoid being prompted for this password, you can include the argument *-p <password>* when running the command, where *<password>* is the password to assign to the new user.
- Users that you create with the *iqmgr user add* command will have the Viewer role in the HyperIQ Grafana UI. (After creating users you can optionally revise their role as described in **"Getting Started With HyperIQ"** (page 29).)
- The *iqmgr user list* command returns a list of "Grafana users" and a list of "Nginx users" (Prometheus UI users). Users that you create through the *iqmgr user add* command will appear in both lists. Users that you create through the Grafana UI will appear only in the "Grafana users" list.

- If you want to delete a user who you had created through the *iqmgr user add* command, do it with the *iqmgr user delete* command so that the user is deleted as a Grafana user and also as a Prometheus UI user. If you instead delete such a user through the Grafana UI's user management page, the user will only be deleted as a Grafana user and will persist as a Prometheus UI user.

**Change the Prometheus UI password for a user other than the *cloudian\_admin* user.**

```
iqmgr user password <username>
```

Note that:

- When you run *iqmgr user password <username>* you will be prompted to enter the new password. Do not include the password on the command line when launching the command.
- This command does not change the user's Grafana password. To change the Grafana password for a user other than the *cloudian\_admin* user, use the user management page in the Grafana UI.
- To change the *cloudian\_admin* user's password for both the Grafana UI and Prometheus UI, use *iqmgr config password* as described earlier in this table.

#### 4.1.4. Additional iqmgr Commands

**Gather and package HyperIQ container logs:**

```
iqmgr logs
```

This gathers all HyperIQ related log files into a tarball named *hyperiq-logs-*<timestamp>*.tgz* under the current directory (the directory in which you run *iqmgr logs*). You may need to do this if you are working with Clodian Support to troubleshoot a HyperIQ issue.

**Copy a file to a specified container:**

```
iqmgr cp <source-filepath> <container>:<destination-path>
```

Note that:

- For this command and for the two container-related commands that follow below, you do not need to include "hyperiq\_" in the *<container>* name. For example, identify the Prometheus container as *prometheus*, not as *hyperiq\_prometheus*.
- If you have a HyperStore Single-Node system and you are logged into the HyperStore/HyperIQ node as the HyperStore Shell user *sa\_admin*, copying is only allowed between a container and the */home/sa\_admin* directory (or sub-directories under that directory)

**Copy a file from a specified container:**

```
iqmgr cp <container>:<source-filepath> <destination-path>
```

**Run a command inside a specified container:**

```
iqmgr exec <container> <command>
```

**Display the HyperIQ version number:**

```
iqmgr version
```

**Display *iqmgr* command help:**

```
iqmgr <command> --help
```

## 4.2. Using Docker Commands

HyperIQ is implemented as a containerized application that uses Docker for containerization and Docker Compose for container orchestration. This is the case even if you installed HyperIQ as a virtual appliance (Docker and Compose are packaged with the appliance).

HyperIQ consists of several containers including:

- *hyperiq\_grafana* (Grafana UI)
- *hyperiq\_prometheus* (Prometheus database)
- *hyperiq\_nginx* (web server used for Prometheus UI)
- *hyperiq\_iqmgr* (HyperIQ management service)

You can check on and manage the HyperIQ containers using standard *docker* and *docker-compose* commands. You can run such commands from the operating system command prompt of the HyperIQ host machine. For information about supported commands enter *docker --help* (or for *docker-compose* commands, enter *docker-compose --help*). You can also refer to Docker's online documentation for more detailed information and context.

For example, to list the currently running HyperIQ containers and their health status you can use the command ***docker ps***.

Another useful command is ***docker restart***, to restart a particular container. For example, to restart the Grafana UI container:

```
cloudian@hyperiq:~$ docker restart hyperiq_grafana
hyperiq_grafana
cloudian@hyperiq:~$
```

You should rarely need to access the interior of one of the HyperIQ containers, but it may be useful on occasion -- such as if you want to edit the main Grafana configuration file *grafana.ini*. Here the ***docker exec*** command is used to enter the Grafana container as *root*, then the *grafana.ini* file is edited using *vi*, then the container is exited, and then the container is restarted to apply the configuration change:

```
cloudian@hyperiq:~$ docker exec -u root -it hyperiq_grafana sh
/usr/share/grafana # vi /etc/grafana/grafana.ini
/usr/share/grafana # exit
cloudian@hyperiq:~$ docker restart hyperiq_grafana
hyperiq_grafana
cloudian@hyperiq:~$
```

HyperIQ also supports using ***docker-compose*** commands. To successfully run *docker-compose* commands you must be in the HyperIQ configuration directory (*\$HOME/.config/cloudian/hyperiq* on Linux). For a list of supported commands enter *docker-compose --help*.

## 4.3. Configuring LDAP Authentication for HyperIQ Users

To configure HyperIQ to support LDAP authentication when a user logs into the HyperIQ UI, use *iqmgr config set* commands to assign values to the settings listed below. Remember to subsequently run *iqmgr config apply*. (For an overview of *iqmgr config* command usage, see the introduction to "**Administering HyperIQ**" (page 39)).

When LDAP authentication is configured, if a new user logs into the HyperIQ UI with their LDAP-based login credentials and if HyperIQ successfully authenticates the user with your LDAP system, HyperIQ automatically creates a HyperIQ user account for the user (and then a HyperIQ administrator can see the user's account in the list of users, and optionally adjust the user's HyperIQ role and profile information). Each time the user subsequently logs into HyperIQ, HyperIQ will check with your LDAP system to validate the user's login credentials.

| LDAP Setting                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| /grafana/ldap/enabled         | Enable LDAP for HyperIQ -- true or false. Defaults to false.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| /grafana/ldap/allow_sign_up   | Allow new LDAP-authenticated Grafana users to be created - true or false. Defaults to true.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| /grafana/ldap/server_host     | Required if LDAP is enabled. LDAP server host                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| /grafana/ldap/search_base_dns | Required if LDAP is enabled. Array of base DN's to search, e.g. ["dc=your_domain,dc=com"]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| /grafana/ldap/bind_dn         | Required if LDAP is enabled. Search user bind DN template, e.g. '%s@your_domain.com'                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| /grafana/ldap/bind_password   | Search user bind password. If the password contains # or ; you have to wrap it with triple quotes, e.g. ""#password;""<br><br><div> <p><b>Note</b></p> <ul style="list-style-type: none"> <li>* For the bind password setting do not include the password value on the command line. Instead just run <code>iqmgr config set /grafana/ldap/bind_password</code> and then you will be prompted to provide the bind password.</li> <li>* Setting <code>bind_password</code> can be skipped if you are able to set a <code>bind_dn</code> that matches all possible users, for example <code>"cn=%s,o=users,dc=grafana,dc=org"</code>. For more information on Grafana LDAP settings see <a href="https://grafana.com/docs/grafana/latest/setup-grafana/configure-security/configure-authentication/ldap/">https://grafana.com/docs/grafana/latest/setup-grafana/configure-security/configure-authentication/ldap/</a>.</li> </ul> </div> |
| /grafana/ldap/server_port     | LDAP server port. Defaults to 389 (or to 636 if 'use_ssl' is set to 'true')                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| /grafana/ldap/search_filter   | User search filter, default: '(sAMAccountName=%s)'                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| /grafana/ldap/use_ssl         | Set to true if connections to the LDAP server should use an encrypted TLS connection (either with STARTTLS or LDAPS). Defaults to false.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| /grafana/ldap/start_tls       | If set to true, use LDAP with STARTTLS instead of LDAPS. Defaults to false.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| /grafana/ldap/client_cert     | If using LDAPS for connecting to LDAP server: Path in <code>hyperiq_grafana</code> container to your client certificate (for example <code>/home/my.crt</code> )<br><br><div> <p><b>Note</b> For this and for the client private key and root</p> </div>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

| LDAP Setting                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                               | <p>CA certificate settings below, you must copy the certificate and key files into the <i>hyperiq_grafana</i> container, so that the files are at the location(s) that you specify with these settings. You can copy files into the container with:</p> <pre>docker cp &lt;local-path-to-file&gt; &lt;container&gt;:&lt;destination directory&gt;</pre> <p>or</p> <pre>iqmgr cp &lt;local-path-to-file&gt; &lt;container&gt;:&lt;destination directory&gt;</pre> <p>For example:</p> <pre>docker cp ./ca.crt hyperiq_grafana:/home/grafana/</pre> <p>For more on <i>iqmgr cp</i> see <b>"Using the iqmgr Tool to Manage HyperIQ"</b> (page 39).</p> |
| /grafana/ldap/client_key      | If using LDAPS for connecting to LDAP server: Path in <i>hyperiq_grafana</i> container to your client private key (for example <i>/home/my.key</i> )                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| /grafana/ldap/root_ca_cert    | If using LDAPS for connecting to LDAP server: Path in <i>hyperiq_grafana</i> container to your root CA certificate (for example <i>/home/ca.crt</i> )                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| /grafana/ldap/ssl_skip_verify | If using LDAPS for connecting to LDAP server: Set to true if you want to skip SSL certificate validation (relevant only if "use_ssl" is set to true). Defaults to false.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## 4.4. Alerts and Notifications

**Note** In HyperIQ version 1.9.x and older, Prometheus's "Alertmanager" was used for HyperIQ alerts and notifications. In HyperIQ 2.0 and newer, Grafana is used for HyperIQ alerts and notifications, and Alertmanager is deprecated (but still functional). If you customized alert thresholds in HyperIQ version 1.9.x or older, and you have not yet migrated those customizations to HyperIQ's new Grafana based alerting system, instructions for migrating your customizations are available in the HyperIQ 2.0 User Guide (which you can obtain from Clodian if you don't already have it). Alertmanager will be removed in a future HyperIQ release.


This section covers the following topics:

- **"Reviewing Pre-Configured Alert Rules"** (page 47)
- **"Setting Up Notification Channels"** (page 52)
- **"Customizing Alert Thresholds"** (page 54)

- **"Silencing or Pausing Alerts"** (page 55)
- **"Customized Alerts and HyperIQ Upgrades"** (page 57)

#### 4.4.1. Reviewing Pre-Configured Alert Rules

HyperIQ comes with a variety of pre-configured alert rules, covering the main conditions and events that merit your attention as a HyperStore system administrator. To review the alert rule details, in the upper left of the HyperIQ UI click the menu icon:



Then in the menu select **Alerting -> Alert rules**. This displays the **Grafana-managed** alert rules, organized into categories:

|                                           |           |    |  |
|-------------------------------------------|-----------|----|--|
| 106 rules 104 normal 2 recording          |           |    |  |
| Grafana-managed <span>Export rules</span> |           |    |  |
| > Cassandra Alerts > cassandra_alerts     | 16 normal | 1m |  |
| > Disk Alerts > disk_alerts               | 9 normal  | 1m |  |
| > Node Alerts > node_alerts               | 21 normal | 1m |  |
| > Prometheus Alerts > prometheus_alerts   | 8 normal  | 1m |  |
| > S3 Alerts > s3_alerts                   | 10 normal | 1m |  |
| > Service Alerts > service_alerts         | 3 normal  | 1m |  |

**Note** Lower in the screen is a section of "Data source-managed" alert rules, which are legacy from when HyperIQ used Alertmanager (see the Note about this directly under the **"Alerts and Notifications"** (page 46) section heading in this document). Ignore the "Data source-managed" alert rules.

In the **Grafana-managed** alert rules section you can select a category and then within a category select individual alert rules to examine the details of those rules.

Also, the tables that follow below in this section of the documentation provide summaries of the alert rules.

Note that:

- By default, firing alerts (if any) will display in the "HyperIQ Main" dashboard but no notifications will be sent. If you want notifications to be sent -- to one or more email addresses for example -- follow the instructions in **"Setting Up Notification Channels"** (page 52).
- Each HyperIQ alert has a unique alert code. For alerts that may appear both in HyperIQ and in HyperStore's Cloudian Management Console (CMC), the same alert code is used by both applications. For example, if the Cassandra Service is down on a node, the alert in the CMC and the alert in HyperIQ both have the alert code "SS000002".
- By default, the pre-configured S3 alerts -- alerts based on your users' utilization of HyperStore's S3 Service -- are disabled ("paused"). This is because there are no universally appropriate thresholds for such alerts, and instead such thresholds should be based on your system's expected S3 workload and your preferences regarding what levels of service utilization merit alerts. For information about setting custom thresholds for pre-configured alert rules see **"Customizing Alert Thresholds"** (page 54) and for unpausing rules see **"Silencing or Pausing Alerts"** (page 55).

**Note** Enabling S3 alert rules is supported **only if your HyperStore license supports HyperIQ Enterprise**.

HyperIQ's pre-configured alert rules are summarized in the tables below.

- **"Cassandra Alerts"** (page 48)
- **"Disk Alerts"** (page 49)
- **"Node Alerts"** (page 50)
- **"Prometheus Alerts"** (page 51)
- **"S3 Alerts"** (page 51)
- **"Service Alerts"** (page 52)

#### 4.4.1.1. Cassandra Alerts

| Alert Rule                                       | Alert Code | Default Threshold                                                                       |
|--------------------------------------------------|------------|-----------------------------------------------------------------------------------------|
| CassandraCompactionPendingLimitExceededCritical  | CA000001   | Pending compactions on a node have been exceeding 1000                                  |
| CassandraCompactionPendingLimitExceededHigh      | CA000002   | Pending compactions on a node have been exceeding 500                                   |
| CassandraDiskAvailableSpaceLimitExceededCritical | CA000003   | Available space on a Cassandra disk has fallen below 30%.                               |
| CassandraDiskAvailableSpaceLimitExceededHigh     | CA000004   | Available space on a Cassandra disk has fallen below 50%.                               |
| CassandraDroppedMessagesPerSecondLimitExceeded   | CA000005   | Dropped internal Cassandra messages have exceeded 10 per second                         |
| CassandraHighHeapUsage                           | CA000006   | Cassandra heap memory usage on a node has exceeded 80%                                  |
| CassandraCriticalHeapUsage                       | CA000007   | Cassandra heap memory usage on a node has exceeded 90%                                  |
| CassandraHighCPUUsage                            | CA000008   | Cassandra's CPU usage as a percentage of total CPU usage on a node has exceeded 80%.    |
| CassandraCriticalCPUUsage                        | CA000009   | Cassandra's CPU usage as a percentage of total CPU usage on a node has exceeded 90%.    |
| CassandraKeyspacePartitionTooLarge               | CA000010   | During the past 24 hours the average maximum partition size on a node has exceeded 1GB. |
| CassandraLoadRatioNotBalanced                    | CA000011   | Cassandra load on a node differs from other nodes by more than two standard deviations  |
| CassandraKeyspaceReadLatencyHigh                 | CA000012   | During the past hour the average of the 99.9th percentile read latency                  |



| Alert Rule                                 | Alert Code | Default Threshold                                                                                                                                                                                                                                                         |
|--------------------------------------------|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                            |            | measurements for an object metadata keyspace on a node has exceeded 1 second.                                                                                                                                                                                             |
| CassandraTombstoneScannedCountMedium       | CA000013   | During the past hour the total number of tombstones scanned during key-space read operations on a node has exceeded 75,000.                                                                                                                                               |
| CassandraTombstoneScannedCountHigh         | CA000014   | During the past hour the total number of tombstones scanned during key-space read operations on a node has exceeded 100,000.                                                                                                                                              |
| CassandraTombstoneErrorThresholdExceeded   | CA000015   | The number of tombstones scanned in a single query exceeded 100,000 and so the query failed. Threshold is configurable by HyperStore setting <i>cassandra.tombstone.threshold.fail</i> (HyperStore 8.x) or <i>cassandra_tombstone_failure_threshold</i> (HyperStore 7.x). |
| CassandraTombstoneWarningThresholdExceeded | CA000016   | The number of tombstones scanned in a single query exceeded 50,000. Threshold is configurable by HyperStore setting <i>cassandra.tombstone.threshold.warn</i> (HyperStore 8.x) or <i>cassandra_tombstone_warn_threshold</i> (HyperStore 7.x).                             |

#### 4.4.1.2. Disk Alerts

| Alert Rule                  | Alert Code | Default Threshold                                                     |
|-----------------------------|------------|-----------------------------------------------------------------------|
| NodeDiskWillFillIn30Days    | DS000001   | A disk is projected to be full in 30 days at current write rate       |
| NodeUnusualDiskReadLatency  | DS000002   | Read latency on a disk has been exceeding 250ms                       |
| NodeUnusualDiskWriteLatency | DS000003   | Write latency on a disk has been exceeding 250ms                      |
| RootDiskOutOfSpace          | DS000004   | Available space on a data disk has dropped below 10%                  |
| DiskSMARTHealthStatus       | DS000005   | SMART healthcheck has failed for a disk                               |
| HDDDiskTempHigh             | DS000006   | An HDD's temperature has exceeded 50° Celsius for 5 minutes           |
| NVMeDiskTempHigh            | DS000007   | An NVMe Disk's temperature has exceeded 70° Celsius for 5 minutes     |
| DataDiskOutOfSpace          | GS000004   | Available space on a data disk has dropped below 20%                  |
| NodeDiskFailure             | GS000005   | Completed writes on a disk are declining, could indicate disk failure |

## 4.4.1.3. Node Alerts

| Alert Rule                      | Alert Code | Default Threshold                                                                                                                                                                                            |
|---------------------------------|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NodeOutOfMemory                 | ND000001   | Available memory on a node fell below 10%                                                                                                                                                                    |
| HighCPUUsage                    | GS000006   | Average CPU usage on a node has exceeded 90% over a five minute period                                                                                                                                       |
| NodeHighInboundNetworkTraffic   | NS000007   | Average inbound traffic on a node has exceed 70% of NIC receiving capacity over a two minute period                                                                                                          |
| NodeHighOutboundNetworkTraffic  | NS000008   | Average outbound traffic on a node has exceed 70% of NIC sending capacity over a two minute period                                                                                                           |
| NodeDown                        | GS000001   | n/a                                                                                                                                                                                                          |
| NodeOutOfDiskSpace              | GS000003   | Total available disk space on a node is less than 10% of disk capacity                                                                                                                                       |
| NodeOutOfNodesHigh              | ND000002   | A disk has less than 20% of inodes free                                                                                                                                                                      |
| NodeOutOfNodesCritical          | ND000003   | A disk has less than 10% of inodes free                                                                                                                                                                      |
| NodeRAIDDiskFailure             | ND000004   | n/a                                                                                                                                                                                                          |
| NodeKernelVersionDeviations     | ND000005   | n/a                                                                                                                                                                                                          |
| NodeHyperStoreVersionDeviations | ND000006   | n/a                                                                                                                                                                                                          |
| NodeNetworkInterfaceFlapping    | ND000007   | n/a                                                                                                                                                                                                          |
| NodeInterfaceDown               | ND000008   | n/a<br><br><b>Note</b> If bonded interfaces are being used on your HyperStore nodes, this alert rule will monitor (and can be triggered by) the individual slave interfaces as well as the bonded interface. |
| NodeHighTCPInErrs               | ND000009   | TCP In errors have increased by over 100 in the last 30 minutes on a node                                                                                                                                    |
| NodeHighTCPRetransHigh          | ND000010   | More than 5% of TCP Out segments in the last 10 minutes have been retransmitted on a node                                                                                                                    |
| NodeNetworkBondDegraded         | ND000011   | n/a                                                                                                                                                                                                          |
| NodeTimeOutOfSyncWithHIQ        | ND000012   | A node's time is out of sync with HyperIQ by more than 5 seconds                                                                                                                                             |
| NodeTimeNotSyncedWithNTP        | ND000013   | A node's time is not sync'd with an NTP server                                                                                                                                                               |
| NodePSUNotActive                | ND000014   | n/a                                                                                                                                                                                                          |
| NodeFanLowRPM                   | ND000015   | A node's fan speed is below 600 RPM (but above 0) for 5 minutes                                                                                                                                              |
| NodeFanNoRPM                    | ND000016   | A node's fan speed is 0 RPM for 5 minutes                                                                                                                                                                    |
| NodeCPUTempHigh                 | ND000017   | A node's CPU temperature has been over 80° Celsius for 15 minutes                                                                                                                                            |

## 4.4.1.4. Prometheus Alerts

| Alert                                | Alert Code | Default Threshold                                                                 |
|--------------------------------------|------------|-----------------------------------------------------------------------------------|
| PrometheusTargetMissing              | PR000001   | n/a                                                                               |
| PrometheusAllTargetsMissing          | PR000002   | n/a                                                                               |
| PrometheusConfigurationReloadFailure | PR000003   | n/a                                                                               |
| PrometheusTooManyRestarts            | PR000004   | Prometheus has restarted more than twice in the last 10 minutes on any node       |
| PrometheusConfigurationReloadFailure | AM000001   | n/a                                                                               |
| PrometheusRuleEvaluationSlow         | PR000005   | Prometheus rule evaluation took more time than the scheduled interval on any node |
| PrometheusNotificationsDropped       | PR000006   | In the last 5 minutes, Prometheus notifications have been dropped on any node     |
| PrometheusNotificationFailing        | AM000002   | If any Prometheus notifications fail to be sent                                   |

## 4.4.1.5. S3 Alerts

With the exception of the "S3LogLinesIgnored" alert rule, all S3 analytics alert rules:

- Are disabled (paused) by default, and you can only enable (unpause) them if your HyperStore license supports HyperIQ Enterprise.
- Have placeholder thresholds that you must customize to levels appropriate to your HyperStore use case. The placeholder threshold values that you should review and customize are shown in bold text in the table below.

For information about setting custom thresholds see **"Customizing Alert Thresholds"** (page 54) and for unpausing rules see **"Silencing or Pausing Alerts"** (page 55).

| Alert Rule                        | Alert Code | Default Threshold                                                                                                                    |
|-----------------------------------|------------|--------------------------------------------------------------------------------------------------------------------------------------|
| S3LogLinesIgnored                 | S3000001   | n/a                                                                                                                                  |
| User30DayDownloadLimitExceeded    | S3000002   | A user's downloads during the past 30 days has exceeded <b>100,000</b> bytes.                                                        |
| User30DayUploadLimitExceeded      | S3000003   | A user's uploads during the past 30 days has exceeded <b>100,000</b> bytes.                                                          |
| GETTransactionsPerSecond          | NS000001   | S3 GET transaction count for the whole HyperStore system during the past 15 minutes has exceeded <b>498</b> transactions per second. |
| PUTTransactionsPerSecond          | NS000002   | S3 PUT transaction count for the whole HyperStore system during the past 15 minutes has exceeded <b>500</b> transactions per second. |
| DELETEDTransactionsPerHour        | S3000004   | Total DELETE transaction count during the past hour has exceeded <b>5000</b> transactions.                                           |
| TotalSystemBytecountLimitExceeded | S3000005   | Total net bytes of object data stored in the HyperStore system exceeds <b>5,000,000</b> . This is based on <i>net</i> bytes of       |

| Alert Rule                     | Alert Code | Default Threshold                                                                                                                        |
|--------------------------------|------------|------------------------------------------------------------------------------------------------------------------------------------------|
|                                |            | object data, which excludes overhead from replication or erasure coding.                                                                 |
| TotalSystemObjectLimitExceeded | S3000006   | Total number of objects stored in the HyperStore system exceeds <b>500,000</b>                                                           |
| GETThroughputLimitExceeded     | NS000003   | S3 GET transaction throughput for the whole HyperStore system during the past 15 minutes has exceeded <b>5,000,000</b> bytes per second. |
| PUTThroughputLimitExceeded     | NS000004   | S3 PUT transaction throughput for the whole HyperStore system during the past 15 minutes has exceeded <b>5,000,000</b> bytes per second. |
| 503sLast10MinLimitExceeded     | S3000007   | Number of HTTP 503 (Service Unavailable) errors in the last 10 minutes has exceeded <b>30</b> .                                          |

#### 4.4.1.6. Service Alerts

| Alert Rule                    | Alert Code | Default Threshold                                                                               |
|-------------------------------|------------|-------------------------------------------------------------------------------------------------|
| ServiceDown:                  |            |                                                                                                 |
| "admin"                       | SS000001   | n/a                                                                                             |
| "cassandra"                   | SS000002   | n/a                                                                                             |
| "hyperstore"                  | SS000003   | n/a                                                                                             |
| "redis-qos"                   | SS000004   | n/a                                                                                             |
| "redis-credentials"           | SS000005   | n/a                                                                                             |
| "cloudians3"                  | SS000007   | n/a                                                                                             |
| "iq_s3"                       | SA000003   | n/a                                                                                             |
| "iq_node"                     | SA000004   | n/a                                                                                             |
| "iq_service"                  | SA000005   | n/a                                                                                             |
| "iq_admin"                    | SA000006   | n/a                                                                                             |
| "iq_jmx"                      | SA000007   | n/a                                                                                             |
| "ntpd"                        | SA000008   | n/a                                                                                             |
| "cmc"                         | SA000009   | n/a                                                                                             |
| CassandraExceptions           | SA000001   | If there has been an increase in Cassandra storage exceptions on any node in the past 5 minutes |
| EnterpriseLicenseNotInstalled | SA000002   | n/a                                                                                             |

#### 4.4.2. Setting Up Notification Channels

To set up **email notifications** so that firing alerts are sent to one or more system administrator email addresses:


1. Log into the HyperIQ host and use the *iqmgr* command line tool to configure the information that HyperIQ needs to connect to your SMTP server (for background information on using *iqmgr* see **"Using**

the **iqmgr Tool to Manage HyperIQ**" (page 39)):


```
iqmgr config set /grafana/alerting/smtp/host <value>
iqmgr config set /grafana/alerting/smtp/username <value>
iqmgr config set /grafana/alerting/smtp/password
iqmgr config set /grafana/alerting/smtp/insecure_skip_verify <value>
iqmgr config set /grafana/alerting/smtp/from <value>
iqmgr config apply
```

Setting descriptions:

- `.../host` -- SMTP host with port number; for example `smtp.enterprise.org:587`
  - `.../username` -- User name for SMTP authentication (optional)
  - `.../password` -- Password for SMTP authentication (optional). For this setting just enter `iqmgr config set /grafana/alerting/smtp/password` (with no value) and then you will be prompted to enter the desired password.
  - `.../insecure_skip_verify` -- Set to `true` to skip validating the SMTP server's SSL/TLS certification (optional, defaults to false). You should set this to `true` if your SMTP server uses a self-signed certificate.
  - `.../from` -- SMTP 'From' address for HyperIQ to use when sending email notifications
2. Log into the HyperIQ UI and set an Alerting "Contact Point" for your system administrator email address(es):
    - a. In the upper left of the UI select the menu icon:



- b. Select **Alerting -> Contact points**, then select to edit the "grafana-default-email" contact point:



- c. In the **grafana-default-email** screen that displays, select Edit to specify the desired administrator email address(es):
      - i. In the "Addresses" field, enter the administrator email address. If you wish you can enter multiple addresses separated by semi-colons.
      - ii. Click **Test** to send a predefined test email to your specified address(es). This will use the SMTP server information that you configured in Step 1 to send the test email. If the test fails with an SMTP error, on the HyperIQ host command line double check that your SMTP settings are correct (`iqmgr config get /grafana/alerting/smtp`).
      - iii. If you do not want HyperIQ to send notifications when an alert is **resolved**, select "Notification settings" and then select the "Disable resolved message" checkbox.
      - iv. Click **Save contact point**.

For more information about setting up notification channels -- such as notifications via **Slack or Microsoft Teams** -- access the Grafana 11.4 documentation online (<https://grafana.com/docs/grafana/v11.4/>) and in the table of contents go to **Alerting -> Configure Notifications**. In short, you can set up a contact point for Slack


or MS Teams (or other contact types supported in the alerting configuration interface) and then integrate the contact point into the pre-configured notification policies.

### 4.4.3. Customizing Alert Thresholds

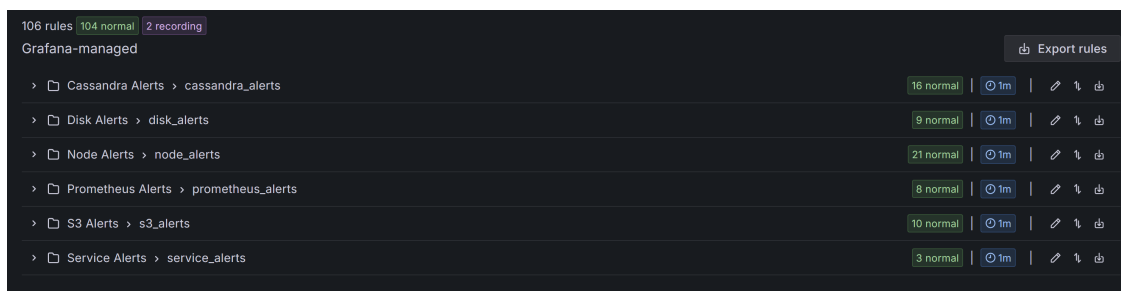
Many alerts are based on the crossing of a particular threshold, such as memory usage exceeding a certain percentage of available RAM or disk space availability falling below a certain percentage. The default thresholds for HyperIQ alerts are summarized in **"Reviewing Pre-Configured Alert Rules"** (page 47). If you wish you can customize the threshold used by an alert rule rather than keeping the default threshold. In the case of S3 alerts, customizing is **required** since the default thresholds are simply placeholders that are not appropriate to a production environment.

To customize an alert threshold:

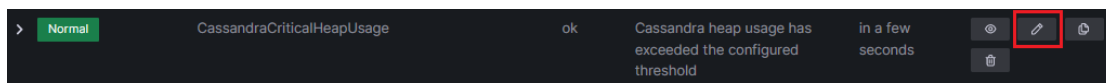
1. In the upper left of the HyperIQ UI select the menu icon:



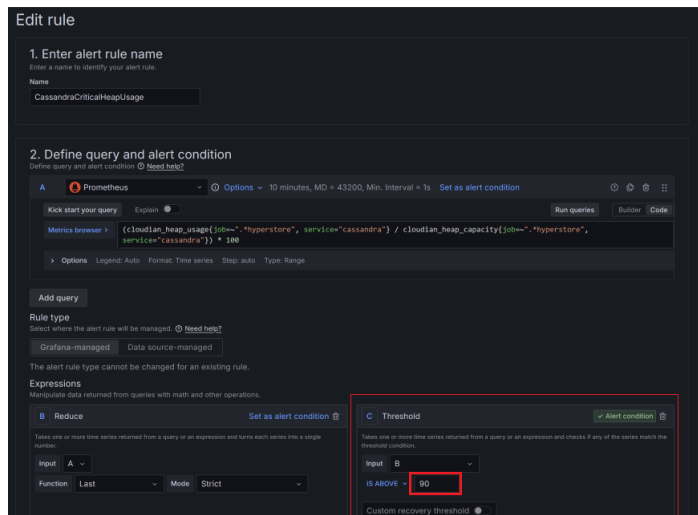
2. Then in the menu select **Alerting -> Alert rules**. This displays the Grafana based alert rules, organized into categories:



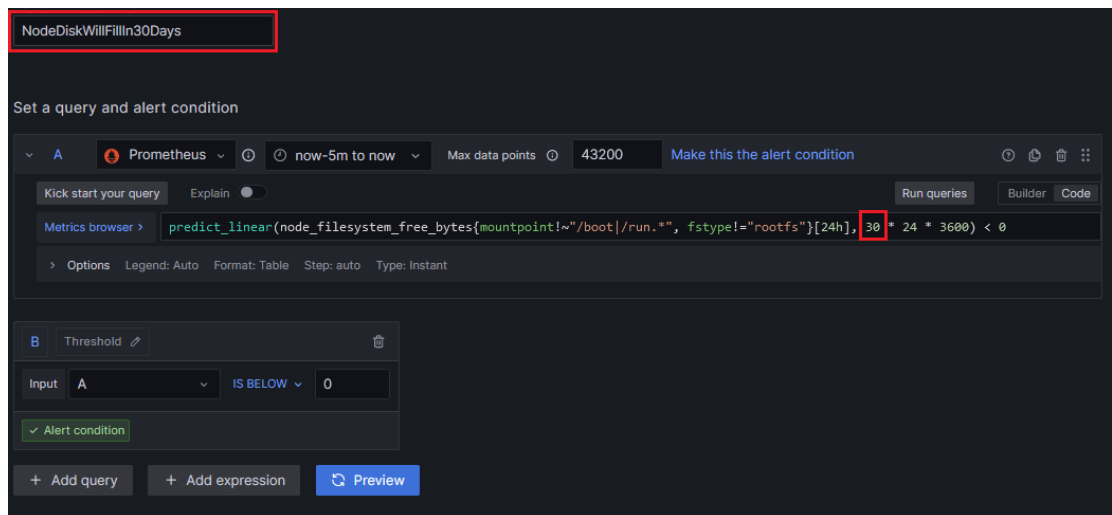
3. Select a category and then select to edit the alert rule for which you want to customize the threshold; for example:



4. In the **Edit rule** screen, edit the alert rule's Threshold:



Note that while in most cases what you may want to customize for an alert rule is the "Threshold" value, in other cases you may want to edit the query associated with the rule -- for example, if you wanted to change the "NodeDiskWillFillIn30Days" alert rule into a "NodeDiskWillFillIn60Days" alert rule:



5. In the upper right of the **Edit rule** screen click **Save rule and exit**.


#### 4.4.4. Silencing or Pausing Alerts

If you "silence" an alert then no notifications will be sent out for that alert for a time duration of your choosing. Any firing alerts that are silenced will still appear in the "Firing Alerts" panel in the HyperIQ Main dashboard, but no notifications will be sent for those alerts.

If you "pause" an alert rule, then that alert rule is disabled until you unpause the alert rule. While paused, an alert rule is not evaluated, and no alerts will fire from that alert rule.

##### To silence an alert:

1. In the upper left of the HyperIQ UI select the menu icon:



- Then in the menu select **Alerting -> Silences**. Then in the **Silences** screen click **Create Silence** (or **Add Silence** if there are already active or recently expired silences in the screen). This displays the **Silence Alert Rule** screen:

- Specify the Duration for which you want the silence to last -- for example *30m* or *4h* or *1d*.
- Specify one or more matching labels to identify which alert you want to silence. Some examples:

Silence all alerts from a particular alert rule:

```
alertname = <alertname>
```

or

```
alert_code = <alert code> (alert codes can be obtained from alert notifications or from the edit alert rule screen or from the "Reviewing Pre-Configured Alert Rules" (page 47) section of this document).
```

Silence all alerts from a particular node:

```
nodename = <hostname>
```

Silence all alerts from a particular alert rule on a particular node (uses two 'matchers'):

```
alertname = <alertname>
```

```
nodename = <hostname>
```

Silence all alerts from a whole HyperStore system (use with caution and only briefly):

```
cluster = <cluster name>
```


- Click **Save silence** to start the silence.

The silence will appear in the main **Silences** screen, and the status of the silence will be "Active". When the duration that you specified for the silence completes, the silence's status will switch to "Expired". You have the option of expiring a silence early -- before its configuration duration period completes -- by clicking **Unsilence**. Five days after expiring, a silence will be automatically deleted from the **Silences** screen.

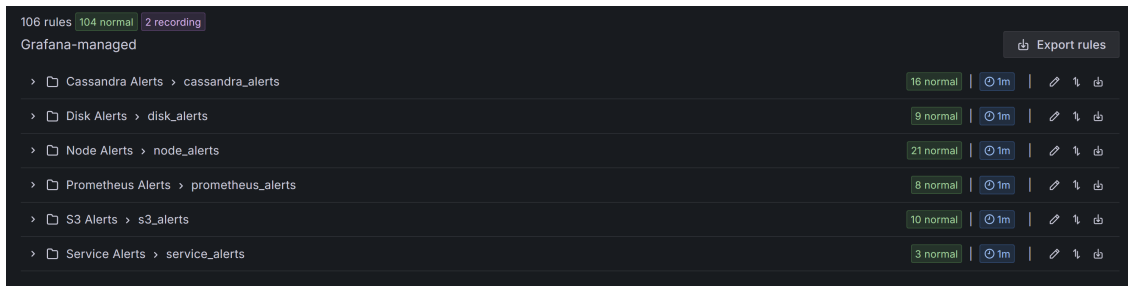
#### To pause an alert rule:

- In the upper left of the HyperIQ UI select the menu icon:

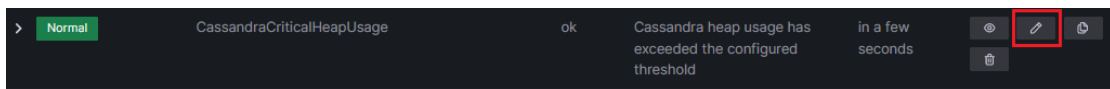




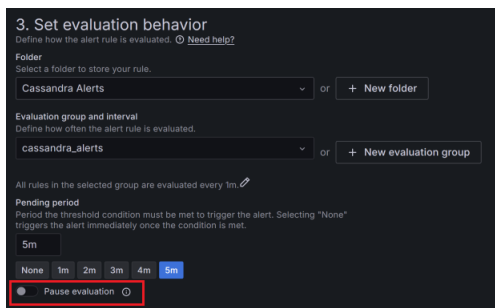
- Then in the menu select **Alerting -> Alert rules**. This displays the Grafana based alert rules, organized into categories:



- Select a category and then select to edit the alert rule that you want to pause; for example:



- In the **Edit rule** screen, in the "Set evaluation behavior" section of the screen, use the "Pause evaluation" control to pause the alert rule:



- In the upper right of the **Edit rule** screen click **Save rule and exit**.

In the alert rule list, the rule's State will now indicate that the rule is paused:



The alert rule will remain paused until you unpaue it.

**To unpaue an alert rule** repeat the steps above except use the "Pause evaluation" control to unpaue the alert rule.

#### 4.4.5. Customized Alerts and HyperIQ Upgrades

Starting with upgrades to HyperIQ 2.2 (upgrading to HyperIQ 2.2 from earlier HyperIQ versions), if you have customized a Grafana-based alert in your existing HyperIQ system, during upgrade that alert will be renamed as `<alertname>_CUSTOM_<existing-version>` and a corresponding "new" alert named just `<alertname>` will be created -- with its default attributes, absent your customization -- but will be disabled ("paused", in Grafana alerting parlance).

For example if you've had HyperIQ 2.1.1, and you customized the *CassandraHeapHighUsage* alert, and then you upgrade to HyperIQ 2.2, during upgrade your customized *CassandraHeapHighUsage* alert will be renamed to *CassandraHeapHighUsage\_CUSTOM\_2.1.1*, and a new, separate *CassandraHeapHighUsage* alert will be created in the system -- with its default attributes, absent your customization -- and this alert will be disabled (paused). Both alerts will exist in your system, but the one with the default attributes will be paused (unless you unpause it).

During the upgrade, terminal output will inform you of any such alert name changes for your system. This will occur only if you have customized alerts in your existing system.

"Customizing" an alert in this context means that you have done any of the following:

- Changing the definition of an alert (such as changing its threshold to a value different than the default threshold).
- Pausing an alert that is enabled by default.
- Enabling an alert that is disabled (paused) by default.

During the upgrade, terminal output will also inform you if any alerts have had their definition changed as a result of the upgrade. This can occur if Clouddian concludes, based on field experience, that an alert's default attributes (for example, the alert threshold) should be changed as part of the upgrade.

## Chapter 5. Installing LogInsight

This chapter covers the following LogInsight set-up topics:

- **"Requirements"** (page 59)
- **"Setting Up Your HyperStore Data Feed for LogInsight"** (page 59)
- **"Installing LogInsight"** (page 64)
- **"Starting Or Stopping LogInsight"** (page 68)
- **"Linking from the HyperIQ UI to the LogInsight UI"** (page 68)
- **"Upgrading LogInsight"** (page 69)

### 5.1. Requirements

The table below describes what you will need to install and use LogInsight.

| Requirement                                                                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LogInsight OVA                                                               | LogInsight is distributed as an Open Virtual Appliance (OVA) file named <i>loginsight-external-&lt;version&gt;.ova</i> . You can obtain the OVA file from your Cloudbian representative.                                                                                                                                                                                                                                                                                                                                          |
| Virtualization platform                                                      | You can deploy the LogInsight virtual appliance on either VMware (ESXi 5.5 or newer, or Workstation 10.x or newer) or VirtualBox. Though it may work on other virtualization platforms, Cloudbian's internal testing and verifications of LogInsight have been primarily on VMware and VirtualBox.                                                                                                                                                                                                                                |
| HyperStore version                                                           | The HyperStore version requirement depends on whether or not you are using the HyperStore Smart Support system (which automatically sends daily HyperStore system information to Cloudbian Support). Using Smart Support is highly recommended. <ul style="list-style-type: none"> <li>• If you are using Smart Support, then to use LogInsight you must have HyperStore version 8.1.1 or newer.</li> <li>• If you are not using Smart Support, then to use LogInsight you can have any HyperStore 7.x or 8.x version.</li> </ul> |
| HyperStore license for "HyperIQ Enterprise" (only required for S3 analytics) | To view S3 log analytics data in LogInsight, your HyperStore system license must include "HyperIQ Enterprise" support. To view other types of log analytics data in HyperIQ -- such as data pertaining to HyperStore system health and performance - - no special license is required (any HyperStore license supports using LogInsight for this system health data).                                                                                                                                                             |
| Network access between LogInsight and HyperStore                             | LogInsight must be able to access your HyperStore system's S3 Service endpoint, to read from the LogInsight input file bucket that you set up as described in <b>"Setting Up Your HyperStore Data Feed for LogInsight"</b> (page 59). This is the only interaction between LogInsight and your HyperStore system.                                                                                                                                                                                                                 |

### 5.2. Setting Up Your HyperStore Data Feed for LogInsight

At a high level, LogInsight operates as follows:

- Once a day your HyperStore system writes a comprehensive bundle of logs, configuration files, and command output files to an S3 bucket that you create when you set up LogInsight.
- Once a day LogInsight ingests the latest bundle from that S3 bucket, parses and analyzes all the different types of data in the bundle, and stores the resulting analyzed data in its local database on the LogInsight host.
- You access LogInsight's web UI to review visualizations of the system status, performance, and usage data, as well as any alerts resulting from the crossing of pre-configured status, performance, or usage thresholds.

This section covers the following topics for setting up the daily HyperStore data feed that will be ingested by LogInsight:

- **"Creating a Bucket For LogInsight Input Files"** (page 60)
- **"Configuring HyperStore to Write to the LogInsight Bucket"** (page 62)
- **"Special Requirements If You Have Multiple HyperStore Clusters"** (page 63)

**Note** If you have multiple HyperStore clusters (multiple independent HyperStore systems) that you want to monitor with LogInsight, review the **"Special Requirements If You Have Multiple HyperStore Clusters"** (page 63) section **before performing the steps in the first two sections.**

### 5.2.1. Creating a Bucket For LogInsight Input Files

To create a new, dedicated S3 bucket to which your HyperStore system will write LogInsight input files, log into the CMC as a system admin user and complete the steps below.

1. In the CMC's **Users & Groups** → **Manage Groups** page, create a new group in which you will subsequently (in Step 2) create a new user who will own the LogInsight input file storage bucket. For example, you could call the new group *LogInsightGroup*.
2. In the CMC's **Users & Groups** → **Manage Users** page, within the group that you created in Step 1, create a new regular user (not an admin user) who will own the LogInsight input file storage bucket. For example, you could call the new user *LogInsightBucketOwner*. Remember the CMC login password that you specify for the user -- you will need this password in the next step below.
3. Log out of the CMC, then log back in as the user that you created in Step 2.
4. In the CMC's **Buckets & Objects** → **Buckets** page, create a bucket in which the HyperStore system can write LogInsight input files. It's best to name the bucket in a way that identifies the HyperStore cluster in which you are creating this bucket -- for example, *production-cluster*. Later when you start using LogInsight, this bucket's name will appear in LogInsight user interface as part of the labeling scheme for organizing monitoring data. In the LogInsight UI the data for each of your service regions will be labeled as *<LogInsightbucketname>=<regionname>* -- for example *production-cluster=region1* and *production-cluster=region2*.

**Note** This user account should not own any buckets other than the LogInsight bucket.

5. While still logged into the CMC as the LogInsight bucket owner, go to **Bucket Properties -> Lifecycle Policy** for the bucket you created, and configure a lifecycle policy that expires (deletes) "Current Version" objects 7 days after creation. This is to prevent a build-up of unneeded old data in the bucket. If you need more instruction see the HyperStore Help, available through the CMC's **Help** link.

**Note** Once LogInsight is in operation, once each day it will retrieve and parse the most recent day's bundle of files from the LogInsight bucket. LogInsight then stores the parsed and processed data in its own local database (on the LogInsight host) for 90 days. So technically, only the most recent day's file bundle is needed from the LogInsight bucket, and not older file bundles. However, retaining the file bundles in the LogInsight bucket for 7 days before they are auto-deleted is a reasonable precaution.

6. While still logged into the CMC as the LogInsight bucket owner, go to **<MyAccount> → Security Credentials** and make a note of the following information that is available in that page:
  - For the region in which you created the bucket, the S3 Service endpoint (either the HTTP endpoint or the HTTPS endpoint, whichever you want LogInsight to use when it accesses the bucket).
  - The user's S3 access key.
  - The user's S3 secret key.

You will need this information when you perform the HyperStore configuration steps described in "**Configuring HyperStore to Write to the LogInsight Bucket**" (page 62); and again when you are installing and setting up LogInsight (as described in "**Installing LogInsight**" (page 64)). Handle this information carefully, particularly the user's security key.

Note that:

- After you complete the configuration steps in the "**Configuring HyperStore to Write to the LogInsight Bucket**" (page 62) section below, your HyperStore system will start writing a bundle of logs, configuration files, command output files, and system statistics files to the bucket once a day. For more details about what will be written to this bucket, in your HyperStore user documentation see the overview of the Smart Support feature. (In the case of LogInsight, your system will be uploading the file bundle to your specified local bucket, rather than to Cloudian Support).
- The writing of the LogInsight file bundle is implemented by a HyperStore system cron job which by default runs at 2:13AM UTC time. If you want to change the time that this cron job runs, see your HyperStore user documentation for information about changing the configuration of system cron job schedules.

**Note** The writing of the LogInsight file bundle will occur within an hour of the time specified for the cron job execution. Because in most customer environments this mechanism is also used to upload daily Smart Support data to Cloudian Support, a random wait time (not exceeding one hour) is built into the upload process so that not all customer environments are uploading to Cloudian Support at the same time. This random wait time applies also if the system is writing files to a local LogInsight bucket.

- LogInsight will read from this bucket once a day, at a time that you configure when you install LogInsight.
- It is recommended that after you finish configuring HyperStore to write to the LogInsight bucket, you proceed promptly to installing and starting LogInsight. If you delay the installation of LogInsight for multiple days, HyperStore will write to the LogInsight bucket each day -- but when LogInsight comes online it will only retrieve the file bundle from the most recent day.

## 5.2.2. Configuring HyperStore to Write to the LogInsight Bucket

The procedure for configuring HyperStore to write LogInsight input files to the target bucket depends on whether or not you are using HyperStore's "Smart Support" system, which automatically sends HyperStore system metrics to Cloudian Support once a day. Using Smart Support is highly recommended. However, in some environments, regulatory restrictions or organizational policies may preclude the use of Smart Support.

Customers using Smart Support must have HyperStore version 8.1.1 or newer to use LogInsight. Customers not using Smart Support -- "dark site" customers -- can use LogInsight with any HyperStore 7.x or 8.x version.

- **"Steps for Smart Support Customers with HyperStore 8.1.1 or Newer"** (page 62)
- **"Steps for 'Dark Site' Customers with HyperStore 7.x"** (page 62)
- **"Steps for 'Dark Site' Customers with HyperStore 8.x"** (page 63)

### 5.2.2.1. Steps for Smart Support Customers with HyperStore 8.1.1 or Newer

Log in to your HyperStore Config Controller node and run the following *hsctl* commands:

```
hsctl config set hyperIQ.phoneHome.enabled=true
hsctl config set hyperIQ.phoneHome.url=s3-endpoint-for-loginsight-bucket-region
hsctl config set hyperIQ.phoneHome.bucket=name-of-loginsight-bucket
hsctl config set hyperIQ.phoneHome.auth.accessKey=s3-access-key-of-loginsight-bucket-owner
hsctl config set hyperIQ.phoneHome.auth.secretKey=s3-secret-key-of-loginsight-bucket-owner
hsctl config apply all
```

For example:

```
hsctl config set hyperIQ.phoneHome.enabled=true
hsctl config set hyperIQ.phoneHome.url=https://s3-region1.enterprise.com:443
hsctl config set hyperIQ.phoneHome.bucket=production.cluster
hsctl config set hyperIQ.phoneHome.auth.accessKey=29fca755942b39a6e7b0
hsctl config set hyperIQ.phoneHome.auth.secretKey=8h5MisqcPriS/yR5glhKeWShrqVQ6s8eDlGAdTy1
hsctl config apply all
```

No service restart is necessary.

### 5.2.2.2. Steps for 'Dark Site' Customers with HyperStore 7.x

1. Log in to your HyperStore Puppet Master node and complete the steps below. In the configuration file paths referenced in these steps, replace *<version>* with your HyperStore 7.x version, such as 7.5.3.
2. With a text editor open this configuration file:

```
/etc/cloudian-<version>-puppet/modules/cloudians3/templates/mts.properties.erb
```

3. Find the setting *phonehome.enabled* and set it to *true*:

```
phonehome.enabled=true
```

Then save and close the file.

4. With a text editor open this configuration file:

```
/etc/cloudian-<version>-puppet/manifests/extdata/common.csv
```

5. Find these *phonehome\_\** settings and configure them as follows:

```
phonehome_uri,s3-endpoint-for-loginsight-bucket-region
```

```
phonehome_bucket,name-of-loginsight-bucket
phonehome_access_key,s3-access-key-of-loginsight-bucket-owner
phonehome_secret_key,s3-secret-key-of-loginsight-bucket-owner
```

For example:

```
phonehome_uri,https://s3-region1.enterprise.com:443
phonehome_bucket,production.cluster
phonehome_access_key,29fca755942b39a6e7b0
phonehome_secret_key,8h5MisqcPriS/yR5glhKeWShrqVQ6s8eDlGAdTyl
```

Then save and close the file.

6. Change into the installation directory (*/opt/cloudian-staging/<version>*), then launch the HyperStore installer:

```
./cloudianInstall.sh
```

7. At the installer main menu enter **2** for "Cluster Management", then at the Cluster Management menu enter **b** for "Push Configuration Settings to Cluster". At the prompt, push your configuration changes out to all of your HyperStore hosts.

No service restart is necessary.

### 5.2.2.3. Steps for 'Dark Site' Customers with HyperStore 8.x

Log in to your HyperStore Config Controller node and run the following *hsctl* commands:

```
hsctl config set common.phoneHome.enabled=true
hsctl config set common.phoneHome.url=s3-endpoint-for-loginsight-bucket-region
hsctl config set common.phoneHome.bucket=name-of-loginsight-bucket
hsctl config set common.phoneHome.auth.accessKey=s3-access-key-of-loginsight-bucket-owner
hsctl config set common.phoneHome.auth.secretKey=s3-secret-key-of-loginsight-bucket-owner
hsctl config apply common
```

For example:

```
hsctl config set common.phoneHome.enabled=true
hsctl config set common.phoneHome.url=https://s3-region1.enterprise.com:443
hsctl config set common.phoneHome.bucket=production.cluster
hsctl config set common.phoneHome.auth.accessKey=29fca755942b39a6e7b0
hsctl config set common.phoneHome.auth.secretKey=8h5MisqcPriS/yR5glhKeWShrqVQ6s8eDlGAdTyl
hsctl config apply common
```

No service restart is necessary.

### 5.2.3. Special Requirements If You Have Multiple HyperStore Clusters

If you have a single HyperStore system that has multiple service regions, no special configuration is required and you can simply follow the steps in **"Creating a Bucket For LogInsight Input Files"** (page 60) and **"Configuring HyperStore to Write to the LogInsight Bucket"** (page 62), as described above.

However, if you have multiple independent HyperStore systems (also known as clusters), and if you want your one LogInsight instance to show status and statistics for each of your HyperStore clusters, then when performing the steps described in the sections above you need to take these additional actions:

- When **Creating a Bucket For LogInsight Input Files**: Create just one group and one user, in just one of your clusters, but under that user account create *one bucket for each of your HyperStore clusters*. Name

each of the buckets in a way that indicates the HyperStore cluster for which the bucket will store LogInsight input files. To reiterate, there will be one bucket for each of your HyperStore clusters, but all of those buckets will exist under one user account in just one of the clusters.

- When **Configuring HyperStore to Write LogInsight Input Files to the Target Bucket**: Configure the specified "phone home" related settings (described in "**Configuring HyperStore to Write to the LogInsight Bucket**" (page 62)) separately in each of your HyperStore clusters. In each cluster the configuration will point to the bucket that you created for that cluster. These buckets will all be owned by the same user account and will reside together in one of your clusters, as noted in the first bullet point.

For example, if you have three independent HyperStore clusters named Cluster1, Cluster2, and Cluster3, then in Cluster1 you could create a group named LogInsightGroup and a user named LogInsightBucketOwner. And under that user's account you could create three buckets named "insightbucket-cluster1", "insightbucket-cluster2", and "insightbucket-cluster3". Then when doing the phone home setting configuration, you would configure Cluster1's phone home settings to point to bucket "insightbucket-cluster1" (which resides in Cluster1); and configure Cluster2's phone home settings to point to bucket "insightbucket-cluster2" (which also resides in Cluster1); and configure Cluster3's phone home settings to point to bucket "insightbucket-cluster3" (which also resides in Cluster1).

## 5.3. Installing LogInsight

This section covers the following LogInsight installation topics:

- "**Deploying the LogInsight Virtual Appliance**" (page 64)
- "**Running the LogInsight Installer**" (page 65)

### 5.3.1. Deploying the LogInsight Virtual Appliance

Use your virtualization platform console to deploy a new virtual machine (VM) from the LogInsight OVA file.

For example, for VMware ESXi:

1. From the computer on which you downloaded the LogInsight OVA file, connect to your ESXi Server console and log in.
2. On the left side of the ESXi console, select **Virtual Machines**.
3. Click **Create / Register VM**. This launches the **New virtual machine** creation wizard.
4. For the VM creation type select "Deploy a virtual machine from an OVF or OVA file", then click **Next**.
5. Enter your desired name for the virtual machine, then click in the file selection area of the wizard screen and browse (on your local computer) to the LogInsight OVA file that you downloaded. After selecting the OVA file, click **Next**.
6. In the remaining screens of the VM creation wizard, accept the default values (or make adjustments if you wish).
7. In the wizard's **Ready to complete** screen, review the attributes of the new VM and then click **Finish**.

Although the wizard will close and the new VM will appear in the **Virtual machines** list right away, it will take some time -- about 10-15 minutes depending on your environment -- for creation of the VM to complete. You can track the progress by opening the **Recent tasks** panel at the bottom of the **Virtual machines** list.

The LogInsight VM that is created has these specs:



- 6 vCPU cores, 8GB RAM, a 50GB disk for the OS and a 150GB disk for LogInsight data
- OS = Rocky Linux version 8.10, 64-bit
- hostname = *loginsight*

**IMPORTANT !** The LogInsight VM stores its data on the 150GB disk without any replication. **In a production environment you should take appropriate precautions to ensure data durability**, such as performing regular backups of the VM (or at least regular backups of the LogInsight data disk).

### 5.3.2. Running the LogInsight Installer

**Note** The LogInsight installer does not need an outgoing internet connection. All of the needed resources are in the product package that you extracted earlier in this set-up procedure.

If the LogInsight VM is not already powered on, power it on now. Then log into the VM's console with these credentials:

- User name: **cloudian**
- Password: **Cloudian1!**

You will be prompted to create a new password for the *cloudian* user. The change password prompt will require you first to enter again the current password (*Cloudian1!*) and then your desired new password. Remember the new password, since you will need it whenever you want to log into the LogInsight virtual machine. (The LogInsight web UI has a different password as discussed in the next chapter of this document.)

Upon login you will be in the */home/cloudian* directory. To install LogInsight, follow these steps:

1. Extract the LogInsight product package (which is in the */home/cloudian* directory):

```
tar -xvf loginsight-external-<version>.tar.gz
```

2. Change into the *loginsight-external-<version>* sub-directory that was created under the */home/cloudian* directory when you extracted the product package.
3. Run the command below to launch the LogInsight interactive installer. The installer will prompt you for the system configuration information that LogInsight requires, and then it will execute the installation. (If you want to know in advance what information you will need to provide, you can review the table below before launching the installer.)

```
./log-insight-ctl.sh install
```

The installer will prompt you for the following configuration information:

**Note** The installer will first ask if you want to use an existing *log-insight.env* file. This is applicable only if you are upgrading an existing LogInsight instance (see **"Upgrading LogInsight"** (page 69)). For a fresh installation, choose to "enter the required information now".

| Prompt for Information                            | Description                                                                                                                                                                                                                                                |
|---------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Whether to require HTTPS for LogInsight UI access | <p>Indicate whether you want the LogInsight web UI to require connecting clients to use HTTPS.</p> <p>If you choose to require HTTPS, the LogInsight UI will only accept client connections over HTTPS and will fail regular HTTP connection requests.</p> |

| Prompt for Information                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                           | <p>Also if you choose to require HTTPS:</p> <ul style="list-style-type: none"> <li>You will be prompted to provide the full path to the TLS certificate file and TLS certificate key file on the LogInsight virtual machine. You must provide these files and place them on the virtual machine yourself. LogInsight does not come pre-configured with a TLS certificate.</li> <li>After completing the LogInsight installation, you must complete an additional configuration step as described in <b>"Edit the docker-compose.yml File (Only If Requiring HTTPS for LogInsight UI Access)"</b> (page 67).</li> </ul> <p>If you choose not to require HTTPS, the LogInsight UI will only accept client connections over regular HTTP and will fail HTTPS connection requests.</p> |
| Domain                                    | DNS domain within which the LogInsight host resides, for example <i>enterprise.com</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| LogInsight URL                            | Root URL of the LogInsight web UI, including the transfer protocol ( <i>http://</i> or <i>https://</i> ), LogInsight host name ( <i>loginsight</i> ), and the Domain that you supplied at the preceding prompt. For example <i>http://loginsight.enterprise.com</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Time of LogInsight's daily data ingestion | The UTC time at which LogInsight will perform its daily ingestion of the files that HyperStore writes to the LogInsight input files bucket. This is the bucket that you created in the <b>"Setting Up Your HyperStore Data Feed for LogInsight"</b> (page 59) phase of the LogInsight set-up procedure. Note that HyperStore by default loads LogInsight ingest files to the bucket daily sometime between 2:13AM and 3:13AM UTC. You should set the LogInsight daily data ingestion to some time after that. By default (if you don't supply a time at the installer prompt) it is set to 09:30.                                                                                                                                                                                  |
| S3 access key                             | <p>The S3 access key belonging to the HyperStore user account that owns the LogInsight input files bucket.</p> <div> <p><b>Note</b> LogInsight will look for LogInsight input files in all buckets owned by this user account. Consequently, if you have multiple HyperStore systems and you've created one bucket per HyperStore system under this user account (as described in <b>"Setting Up Your HyperStore Data Feed for LogInsight"</b> (page 59) for customers having multiple HyperStore systems), LogInsight will retrieve input files from each of those buckets.</p> </div>                                                                                                                                                                                            |
| S3 secret key                             | The S3 secret key belonging to the HyperStore user account that owns the LogInsight input files bucket.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| S3 endpoint                               | The S3 service URI for the S3 service region in which the LogInsight input files bucket resides, including the transfer protocol ( <i>http://</i> or <i>https://</i> ) and the port (typically 80 or 443). For example <i>https://s3-region1.enterprise.com:443</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| S3 region                                 | The name of the S3 service region in which the LogInsight input files                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

| Prompt for Information | Description                                  |
|------------------------|----------------------------------------------|
|                        | bucket resides, for example <i>region1</i> . |

Your terminal will display status information as the installation proceeds, and then "Install Completed" when the installation is done. You can then proceed to starting LogInsight (as described in **"Starting Or Stopping LogInsight"** (page 68)) **unless** during the installation you chose to require HTTPS for accessing the LogInsight UI, in which case you need to complete the additional configuration steps in the next section below.

### 5.3.2.1. Edit the docker-compose.yml File (Only If Requiring HTTPS for LogInsight UI Access)

If during the LogInsight installation you chose to require HTTPS for access to the LogInsight UI, you need to do the following to complete your LogInsight configuration:

1. Change into the `/home/cloudian/loginsight-external-<version>/builds` directory, and then use the `vi` text editor to open the `docker-compose.yml` file.
2. Find this section of the file, and uncomment the lines in the two sub-sections that are prefaced with `"# Uncomment the lines below if HTTPS protocol is enabled"`. Do not uncomment those prefacing comments, only the lines that follow them. When uncommenting a line, just delete the `"#"` and do not otherwise alter the line. The lines to uncomment are highlighted in bold below.

```
services:
 li-grafana:
 image: li-grafana:${VERSION}
 container_name: li-grafana
 hostname: li-grafana
 tty: true
 stdin_open: true
 restart: always
 depends_on:
 - li-influxdb
 - li-postgres
 - li-log-insight-apisvr
 networks:
 - li-network
 volumes:
 - type: volume
 source: li-grafana-data
 target: /var/lib/grafana
 # Uncomment the lines below if HTTPS protocol is enabled
 #- type: bind
 # source: ${ENV_GF_SERVER_CERT_FILE}
 # target: /etc/grafana/cert/grafana.crt
 #- type: bind
 # source: ${ENV_GF_SERVER_CERT_KEY}
 # target: /etc/grafana/cert/grafana.key
 environment:
 - GF_SERVER_PROTOCOL=${ENV_GF_SERVER_PROTOCOL}
 - GF_SERVER_DOMAIN=${ENV_GF_SERVER_DOMAIN}
 - GF_SERVER_ROOT_URL=${ENV_GF_SERVER_ROOT_URL}
 - GF_PATHS_PLUGINS=/var/lib/grafana-plugins
 - GF_DASHBOARDS_DEFAULT_HOME_DASHBOARD_PATH=/etc/grafana/provisioning/dashboards/...
 # Uncomment the lines below if HTTPS protocol is enabled
```

```
#-GF_SERVER_CERT_FILE=/etc/grafana/cert/grafana.crt
#-GF_SERVER_CERT_KEY=/etc/grafana/cert/grafana.key
```

3. Save and exit the file.

Now you can proceed to starting LogInsight as described in the next section below.

## 5.4. Starting Or Stopping LogInsight

To start LogInsight, while in the `/home/cloudian/loginsight-external-<version>` directory run this command:

```
./log-insight-ctl.sh start
```

The terminal response will show the start-up of each of the six Docker containers that comprise LogInsight. Once all the containers are up and the terminal prompt becomes active again you can exit the LogInsight host and proceed to **"Getting Started With LogInsight"** (page 71).

If for some reason you need to stop LogInsight, while in the `/home/cloudian/loginsight-external-<version>` directory run this command:

```
./log-insight-ctl.sh stop
```

## 5.5. Linking from the HyperIQ UI to the LogInsight UI

Starting with HyperIQ version 2.2, the HyperIQ web UI has a "LogInsight" link that you can use to open the LogInsight web UI (in a separate browser tab). For this link to work, you must complete the following configuration set-up:

1. Log into the HyperIQ host as the `cloudian` user (using the host login password you established when you set up HyperIQ).
2. Run this `iqmgr` command to set the URL that will be associated with the HyperIQ UI's "LogInsight" link:

```
iqmgr log-insight set <full-LogInsight-URL-including-port>
```

For example:

```
iqmgr log-insight set http://loginsight.enterprise.com:3000
```

**Note** The port number for the LogInsight UI is always 3000, regardless of whether you are using HTTP or HTTPS for LogInsight UI access.

If you want to confirm that the configuration is correct you can do:

```
iqmgr log-insight show
```

As an alternative to opening the LogInsight UI from the HyperIQ UI, you can of course go directly to the LogInsight UI's URL from your browser.

For more information about logging into the LogInsight UI (including the UI user name and initial password) see **"Getting Started With LogInsight"** (page 71).

## 5.6. Upgrading LogInsight

To upgrade LogInsight to version 2.2 from a previous version (including the early version formerly known as "Insight Engine"), follow these steps:

1. Log into your existing LogInsight VM as the *cloudian* user. Upon login you will be in the */home/cloudian* directory.
2. Change into the */home/cloudian/loginsight-external-<existing-version>* sub-directory and then stop LogInsight:

```
./log-insight-ctl.sh stop
```

3. Change into the */home/cloudian* directory. Then download or copy the LogInsight 2.2 product package into this directory.
4. Extract the new LogInsight product package:

```
tar -xvf loginsight-external-<new-version>.tar.gz
```

The extraction will create a *loginsight-external-<new-version>* sub-directory under the */home/cloudian* directory.

5. Copy your existing LogInsight configuration file *log-insight.env* from your */home/cloudian/loginsight-external-<existing-version>/builds* directory to the new */home/cloudian/loginsight-external-<new-version>/builds* directory:
6. Change into the */home/cloudian/loginsight-external-<new-version>* directory, then run this command to install the new version of LogInsight:

```
./log-insight-ctl.sh install
```

Your terminal will display status information as the installation proceeds, and then "Install Completed" when the installation is done.

7. If you are requiring HTTPS for access to the LogInsight UI (if you configured your existing LogInsight version in that way and are continuing that configuration for your new LogInsight version) complete the step described in **"Edit the docker-compose.yml File (Only If Requiring HTTPS for LogInsight UI Access)"** (page 67), editing the *docker-compose.yml* file under the */home/cloudian/loginsight-external-<new-version>/builds* directory.
8. Change into the */home/cloudian/loginsight-external-<new-version>* directory and then start LogInsight:

```
./log-insight-ctl.sh start
```

This page left intentionally blank

# Chapter 6. Getting Started With LogInsight

This chapter covers the following topics to help get you started with the LogInsight web UI:

- **"LogInsight Data Ingestion"** (page 71)
- **"Connecting to LogInsight Web UI"** (page 71)
- **"Viewing LogInsight Dashboards and Panels"** (page 72)
- **"LogInsight Dashboards At a Glance"** (page 76)
- **"Adding LogInsight View-Only Users"** (page 81)
- **"LogInsight Alerts"** (page 82)

## 6.1. LogInsight Data Ingestion

After you've completed LogInsight installation, be aware that LogInsight will not be populated with any HyperStore data until both of the following have occurred:

1. HyperStore writes LogInsight input files to the bucket that you set up in the **"Setting Up Your HyperStore Data Feed for LogInsight"** (page 59) procedure. This occurs daily sometime between 2:13 AM and 3:13 AM UTC (by default).
2. LogInsight ingests files from the bucket. This occurs daily at a time that you configured during LogInsight installation (9:30 AM UTC by default).

After the first time that LogInsight's daily ingest from the input files bucket occurs, in the LogInsight UI you will be able to see back one day into your HyperStore system's history. After a second daily ingest occurs the next day, you will be able to see back two days into the history -- and so on. At maximum, after you've had LogInsight running for months, you will be able to see back 90 days into your HyperStore system's history.

## 6.2. Connecting to LogInsight Web UI

To connect to the LogInsight web UI point a browser to the following URL:

`http://<LogInsight-IP-address>:3000`


or

`https://<LogInsight-IP-address>:3000` (if you configured the LogInsight web UI for HTTPS support during the installation process)

Note that the same port number -- 3000 -- is used for both HTTP connections and HTTPS connections.

At the login screen, log in with username **cloudian\_admin** and (for your first login) the default password **admin**.

On first login you should be prompted to change the password to something other than the default. Also, at any time when logged into the UI you can change your password by clicking the user account icon in the upper right corner of the UI and then choosing **Change Password** from the drop-down menu.



### 6.3. Viewing LogInsight Dashboards and Panels

When you log in to LogInsight you are initially at the **Configuration and Status** dashboard.



The **Configuration and Status** dashboard is one of 16 LogInsight dashboards. Each dashboard consists of a set of "panels", with each panel providing a graphical representation of one metric (such as "CPU Utilization Per Node" or "S3 Operations Per Application"). For the full list of LogInsight dashboards and panels see **"LogInsight Dashboards At a Glance"** (page 76).

To switch between dashboards, in the upper right of the UI click **LogInsight**. This displays a list of all dashboards except the one that is currently open. Click a dashboard in the list to open that dashboard.

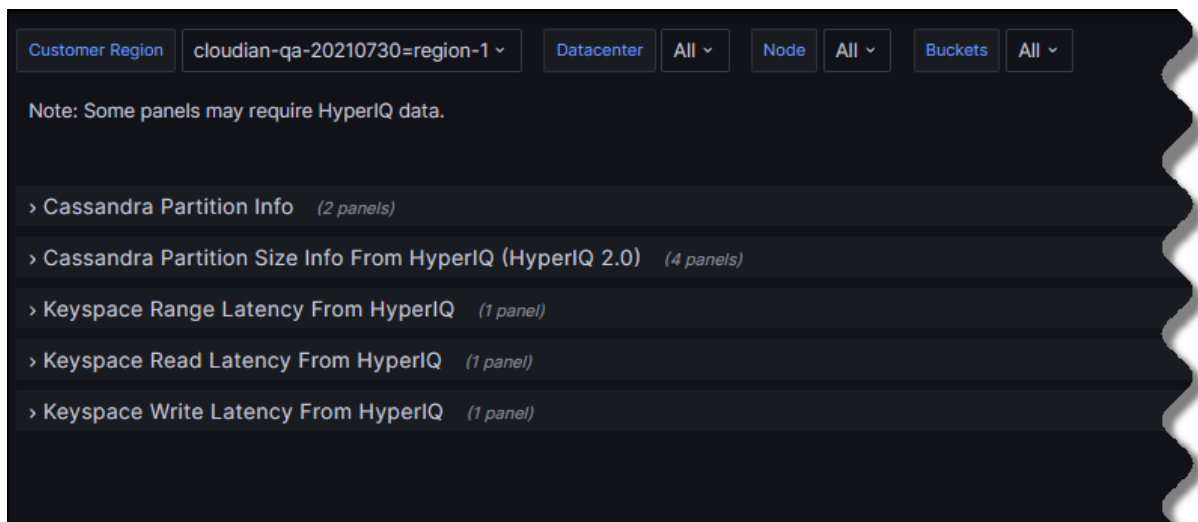




**Note** The S3 dashboards will be populated with data only if your HyperStore system has a "HyperIQ Enterprise" license.


In most dashboards, panels are grouped into folders. When you first access a dashboard, the first panel folder is open so that the panels from that row are all displayed. To see the full list of panel folders available in a dashboard, collapse the open folder.

For example, this image shows the **Cassandra** dashboard with all panel folders collapsed:



### 6.3.1. Filtering the Data in a Dashboard

For each dashboard, in the top left of the LogInsight UI you can filter the data in ways relevant to that dashboard. For example in the **Cassandra** dashboard you can filter data by **Customer Region**, **Datacenter**, **Node**, and **Buckets**.

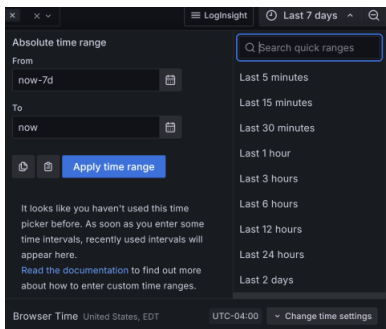


If you filter a dashboard's data using the filters on the top left of the LogInsight UI, your chosen filter applies to all panels in the dashboard.


**Note** "Customer Region" is a concatenation of the name of the LogInsight input file bucket (which you set up in the **"Setting Up Your HyperStore Data Feed for LogInsight"** (page 59) procedure) and the name of the service region. The "Customer Region" filter control is useful only if your HyperStore system has more than one service region.

### 6.3.2. Changing the Time Range for a Dashboard

In the top right of the LogInsight UI, you can change the time period for which a dashboard's panels display historical data. The default time period depends on which dashboard you are viewing. If you change the time period the change applies to all time period sensitive panels on the dashboard.

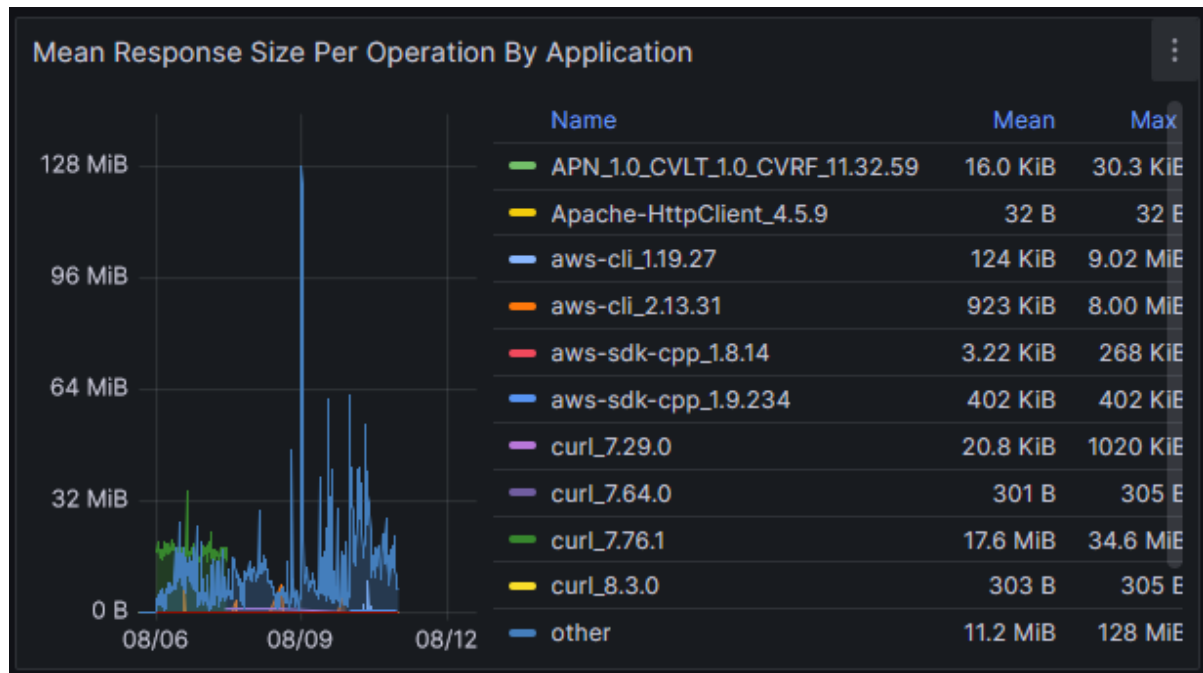


Note that for the S3 dashboards, in the upper left of the UI a "Retention Time" filter lets you choose between 7 Day, 30 Day, and 90 Day views of the data (with a default of 7 Day). For the S3 dashboards, the time range control in the upper right of the UI allows you to look at periods **shorter** than your selected "Retention Time" filter. Do not choose a time range in the upper right of the UI that is longer than your selected "Retention Time" filter, since the "Retention Time" filter sets an upper limit on how far back the view can go.

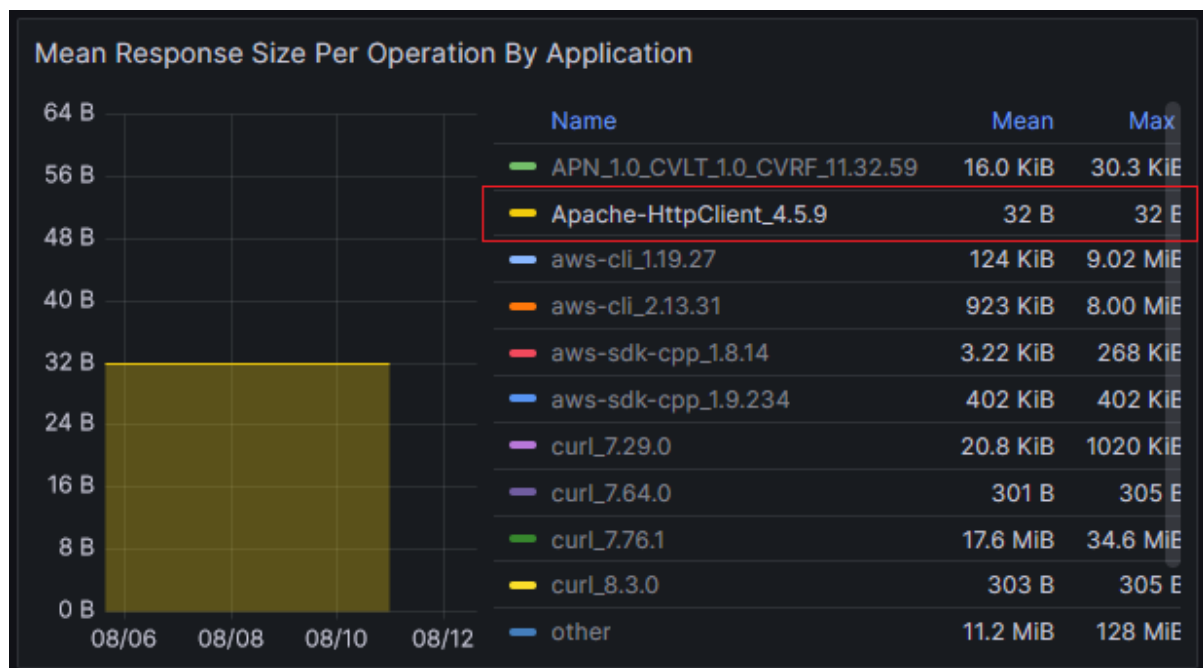


### 6.3.3. Using a Panel Legend to Focus the Data in a Graph Panel

Most graph panels include a legend that indicates the entities for which data has been graphed -- for example, the S3 applications for which data has been graphed:



You can click on an entity in the legend to have the graph show the data for just that entity:




To have the graph return to showing data for all the entities, deselect the selected entity by clicking it again.

#### 6.3.4. Viewing a Larger Version of a Panel

To enlarge a panel, hover your cursor on the panel and press the letter "v" on your keyboard. To exit the enlarged view, press the Esc key.

#### 6.3.5. Viewing a Description of the Data in a Panel

For brief information about the data in a panel, hold your cursor over the "i" to the right of the panel name.



**Note** In the current release, not all panels have an "i" icon, and for those that do the thoroughness of the descriptive text varies from panel to panel.

## 6.4. LogInsight Dashboards At a Glance

The two tables in this section list all of the LogInsight dashboards, and all of the panels available in each dashboard.

- **"HyperStore Hardware and Software Dashboards"** (page 77)
- **"HyperStore S3 Dashboards"** (page 79)

### 6.4.1. HyperStore Hardware and Software Dashboards

| Dashboard                                                       | Panel Folders                                  | Panels                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------------------------------------------------|------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Alerts and Errors</b>                                        | Customer Region Alerts                         | Critical Alerts    High Alerts    Medium Alerts    Low Alerts    New Alerts    Still Firing Alerts    Still Firing > 1 Week Alerts    Still Firing > 1 Month Alerts    Customer Region Alerts                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|                                                                 | Error History                                  | Breakdown by Node for last 7 Days    Breakdown by Error Code for last 7 Days    Errors Over Time    Error Logs Last 30 Days                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Cassandra</b>                                                | Cassandra Partition Info                       | Cassandra Large Partition Current Size    Cassandra Large Partition Size                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|                                                                 | Cassandra Partition Size Info From HyperIQ 2.0 | Keyspace Max Partition Size (Bytes)    Keyspace Max Partition Size (Cells)    Cassandra Load Ratio    Cassandra Load                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|                                                                 | Keyspace Range Latency From HyperIQ            | Keyspace Range Latency                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|                                                                 | Keyspace Read Latency From HyperIQ             | Keyspace Read Latency                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|                                                                 | Keyspace Write Latency From HyperIQ            | Keyspace Write Latency                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Configuration and Status</b><br><br>(landing page dashboard) | Customer Region Status                         | Cluster Name    Total Bytes    Total Objects    Average Object Size    Total Users    Total Nodes    HyperStore Version    Capacity Managed    Total Filesystem Space    Total Filesystem Free    Total Filesystem Used    License Capacity Managed    License Storage Mode    License Max Storage Bytes    Stop Write Percentage    Start Write Percentage    License Expiration    Total Bytes (Over Time)    Total Objects (Over Time)    HSSTool Node Status    OOM Errors for Last 72 Hours    HS Disk Errors for Last 72 Hours    HS Disk Disables for Last 72 Hours    License Errors for Last 72 Hours    Nodes in Maintenance Mode    Critical Alerts    High Alerts    Medium Alerts    Low Alerts    Last Updated    Top 5 Node Filesystem Usage    SMART Disk Errors    Last Repair EC Operation for <DC-name>    Last Successful Repair EC Operation for <DC-name>    Phone Home Bundle Size |
|                                                                 | Customer Region Configuration                  | Region    Support Bucket    Datacenters    HyperStore Version    OS Version    Kernel Version    License Enforcement    License Expiration    License Storage Mode    License Max Storage    Maximum Tiered Storage    GDPR    Secure Deletes    Log Collect    HSH Enabled    HyperIQ Version    Object Lock Mode    Has HSSTool Info    Auto-Repair Replicas    Auto-Repair EC    Auto-Repair Cassandra    Has HyperIQ Metrics    Application                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

| Dashboard                    | Panel Folders                | Panels                                                                                                                                                                                                                                                                                                             |
|------------------------------|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                              |                              | Fingerprints    NTP Servers    Change Log                                                                                                                                                                                                                                                                          |
| <b>CPU and Memory</b>        | CPU and Memory Overview      | Node Memory / CPU Summary                                                                                                                                                                                                                                                                                          |
|                              | CPU Configuration            | Processor Details                                                                                                                                                                                                                                                                                                  |
|                              | CPU Usage                    | CPU Percentage Used by Node                                                                                                                                                                                                                                                                                        |
|                              | Memory Configuration         | Memory Details                                                                                                                                                                                                                                                                                                     |
|                              | Memory Errors                | EDAC CSRow Uncorrectable Errors    EDAC Uncorrectable Errors    EDAC CSRow Correctable Errors    EDAC Correctable Errors                                                                                                                                                                                           |
| <b>Disks and Filesystems</b> | Storage Overview             | Node Raid/Disk/Filesystem Summary                                                                                                                                                                                                                                                                                  |
|                              | RAID Info                    | Raids                                                                                                                                                                                                                                                                                                              |
|                              | Disk Info                    | Disk Details    Disk Status/Errors                                                                                                                                                                                                                                                                                 |
|                              | Disk Usage                   | IO Read By Node    IO Write By Node    Disk Space Available By Node    Disk Space Used By Node                                                                                                                                                                                                                     |
|                              | File System Info             | Cassandra Filesystem Details    HyperStore Filesystem Details                                                                                                                                                                                                                                                      |
|                              | File System Usage            | Cassandra Used Space By File System    HyperStore Used Space By File System                                                                                                                                                                                                                                        |
|                              | Cassandra File System Usage  | Cassandra Disk Space Used                                                                                                                                                                                                                                                                                          |
|                              | HyperStore File System Usage | Disk Space Used for <Hostname>                                                                                                                                                                                                                                                                                     |
|                              | Node File System Usage       | Node Disk Space Used for <DC-name>                                                                                                                                                                                                                                                                                 |
| <b>Networking</b>            | Network Interfaces           | Network Interface Cards                                                                                                                                                                                                                                                                                            |
|                              | Bonded Interfaces            | Bonded Interfaces                                                                                                                                                                                                                                                                                                  |
|                              | VLAN Interfaces              | VLAN Interfaces                                                                                                                                                                                                                                                                                                    |
|                              | Network Routes               | Routes                                                                                                                                                                                                                                                                                                             |
|                              | DNS Name Servers             | DNS Name Servers                                                                                                                                                                                                                                                                                                   |
|                              | Network Usage                | Total Network Traffic by Node (Rx)    Total Network Traffic by Node (Tx)    Frontend Network Traffic by Node (Rx) (HyperIQ Metrics)    Frontend Network Traffic by Node (Tx) (HyperIQ Metrics)    Backend Network Traffic by Node (Rx) (HyperIQ Metrics)    Backend Network Traffic by Node (Tx) (HyperIQ Metrics) |
| <b>Nodes</b>                 | no folders                   | Node Summary    HBA/Expander Summary                                                                                                                                                                                                                                                                               |

| Dashboard        | Panel Folders                   | Panels                                                                                                                                                                                  |
|------------------|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Operations       | Repair Info                     | Repair Overview    Repair Gantt Chart    Repair                                                                                                                                         |
|                  | Repair EC Info                  | Repair EC Overview    Repair EC Gantt Chart    Repair EC                                                                                                                                |
|                  | Repair Cassandra Info           | Repair Cassandra Overview    Repair Cassandra Gantt Chart    Repair Cassandra                                                                                                           |
|                  | Rebalance Info                  | Rebalance Overview    Rebalance Gantt Chart    Rebalance    Rebalance EC Gantt Chart    Rebalance EC                                                                                    |
|                  | Clean Up Info                   | Clean Up Overview    Clean Up Gantt Chart    Clean Up    Clean Up EC Gantt Chart    Clean Up EC                                                                                         |
|                  |                                 |                                                                                                                                                                                         |
| Services         | Service Status                  | Overall Status By Node    Admin Status By Node    Agent Status By Node    Cassandra Status By Node    HyperStore Status By node    Redis Credential Status By Node    S3 Status By Node |
|                  | Services Info                   | Service Interface Details for <DC-name>    Node Service Details for <DC-name>                                                                                                           |
|                  | Services' Java Heap Usage       | Admin Heap Used By Node    Cassandra Heap Used By Node    HyperStore Heap Used By Node    S3 Heap Used By Node                                                                          |
|                  | SSL Certificate Expiration      | Service SSL Certificate Expiration                                                                                                                                                      |
|                  |                                 |                                                                                                                                                                                         |
| Storage Policies | Customer Region Bucket Policies | Policy Counts    Replication (HSFS) Bucket Policy    EC Bucket Policy    Hybrid EC Bucket Policy    Bucket Policy Consistency Level                                                     |
|                  | Storage Policy by Bucket        | Storage Policy by Bucket    Redis Bucket Information    Bucket Partition Rules    Additional Redis Bucket Info                                                                          |

### 6.4.2. HyperStore S3 Dashboards

**Note** The S3 dashboards will be populated with data only if your HyperStore system has a "HyperIQ Enterprise" license.

| Dashboard    | Panel Folders | Panels                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Applications | no folders    | Summary Activity By Application    S3 Requests Per Application    Objects Deleted Per Application    Object Bytes Deleted Per Application    Top Applications for Bucket Listings    503 Error Count per Application    400 Error Count per Application    404 Error Count per Application    Mean Request Size Per Operation By Application    Mean Response Size Per Operation By Application    Total Request Size By Application    Total Response Size By Application    Mean Duration Per Operation By Application |

| Dashboard                       | Panel Folders                             | Panels                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Buckets</b>                  | Bucket Info                               | Summary Activity By Bucket and Application    Summary Activity By Bucket and Node    S3 Requests Per Bucket    Objects Deleted Per Bucket    Object Bytes Deleted Per Bucket    Top Buckets for Bucket Listings    503 Error Count per Bucket    400 Error Count per Bucket    404 Error Count per Bucket    Mean Request Size Per Operation By Bucket    Mean Response Size Per Operation By Bucket    Total Request Size By Bucket    Total Response Size By Bucket    Mean Duration Per Operation By Bucket |
|                                 | Bucket Operation Info By Bucket           | Operation Count Per Operation for bucket <bucketname>                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|                                 | Bucket Operation Info Over Time By Bucket | Bucket Operations Over Time for Bucket <bucketname>                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|                                 | Bucket Object Info                        | Object Size By Bucket Summary    MPU Part Size By Bucket Summary    Object Size Distribution                                                                                                                                                                                                                                                                                                                                                                                                                   |
|                                 | Bucket Object Distribution by Bucket      | Bucket <bucketname>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Customer Region Overview</b> | S3 Activity For Customer Region           | Total Customer Region Upload    Total Customer Region Download    Total Objects Deleted    Total Object Bytes Deleted    Active Users    Active Groups    Active Buckets    Active Applications    Top Buckets    Top Applications    Top Users    Top Groups    Total Bytes    Total Objects    Total Operation Count Per Operation Type    Total Status Count Per Operation    All S3 Requests    Count By Operation Type    Objects Deleted    Object Bytes Deleted                                         |
| <b>Groups</b>                   | no folders                                | Summary Activity By Group    S3 Requests Per Group    Objects Deleted Per Group    Object Bytes Deleted Per Group    Top Groups for Bucket Listings    503 Error Count Per Group    400 Error Count Per Group    404 Error Count Per Group    Mean Request Size Per Operation By Group    Mean Response Size Per Operation By Group    Total Request Size By Group    Total Response Size By Group    Mean Duration Per Operation By Group                                                                     |
| <b>Nodes</b>                    | no folders                                | Summary Activity By Node    Total Operation Count Per Node    Total Delete Operations Per Node    Total Objects Deleted Per Node    Total Object Bytes Deleted Per Node    Object Deletes Per Node    S3 Request By Node    Top Nodes for Bucket Listings    503 Error Count per Node    400 Error Count per Node    404                                                                                                                                                                                       |




| Dashboard    | Panel Folders | Panels                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|              |               | Error Count per Node    Mean Request Size Per Operation By Node    Mean Response Size Per Operation By Node    Total Request Size By Node    Total Response Size By Node    Mean Duration Per Operation By Node                                                                                                                                                                                                               |
| <b>Users</b> | no folders    | Summary Activity By User    S3 Requests Per User    Objects Deleted Per User    Object Bytes Deleted Per User    Top Users for Bucket Listings    503 Error Count per User    400 Error Count per User    404 Error Count per User    Mean Request Size Per Operation By User    Mean Response Size Per Operation By User    Total Request Size By User    Total Response Size By User    Mean Duration Per Operation By User |

## 6.5. Adding LogInsight View-Only Users

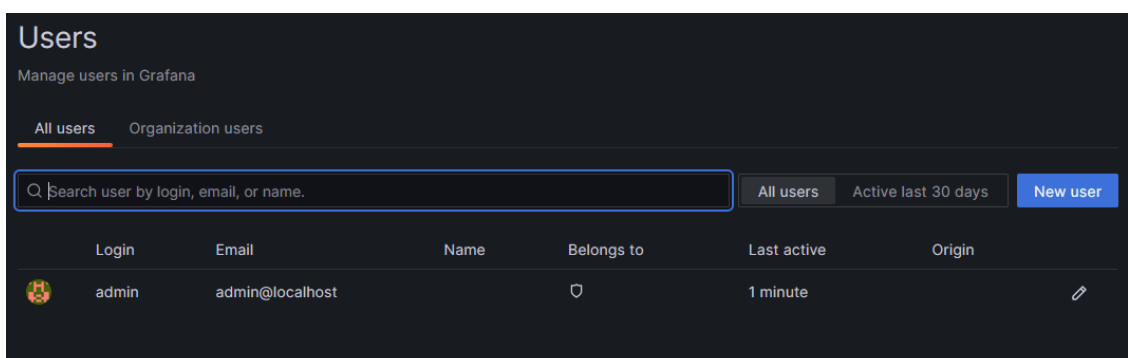
The pre-configured user *admin* is a "Super Admin" user with full LogInsight permissions, and is the only user who can add new LogInsight users. When you are logged in as the *admin* user, if you wish you can create new LogInsight users who can view LogInsight dashboards and panels. These users with "Viewer" level permissions will not be able to create additional users.

To create a view-only LogInsight user:

1. In the upper left of the UI, click the menu icon.



2. In the drop-down menu that displays, select **Administration**. Then in the Administration menu select **Users and Access**, and then select **Users**. This displays the **Users** management interface:



3. Click **New user** to display the user creation interface. Enter a Name, a Username (this will be the user ID for logging in to LogInsight web UI), and a password. Then click **Create user**.

**Note** The user will be able to change their password the first time they log in to LogInsight.

## 6.6. LogInsight Alerts

The table below lists LogInsight alerts and the HyperStore event or condition that triggers each alert. Note that:

- The **Alerts and Errors** dashboard displays alert counts and also a list of specific alerts that are currently firing
- The **Configuration and Status** dashboard's "Customer Region Status" panel folder displays alert counts.
- Alert notifications -- such as emails or Slack messages -- are not supported in the current version of LogInsight.

| Alert                                              | Description                                                               | Alert Message                                                                                          | Threshold | Severity |
|----------------------------------------------------|---------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------|-----------|----------|
| CustomerRegionLicenseExpirationApproachingRule     | Checking HyperStore license expiration date                               | Hyperstore license will expire in less than %d days (Expires %s)                                       | 90        | low      |
|                                                    |                                                                           |                                                                                                        | 60        | medium   |
|                                                    |                                                                           |                                                                                                        | 30        | high     |
|                                                    |                                                                           |                                                                                                        | 0         | critical |
| CustomerRegionCertificateExpirationApproachingRule | Checking HyperStore SSL certificates expiration date                      | SSL Certificate for service %s (URL: %s) will expire in less than %d days (Expires %s)                 | 90        | low      |
|                                                    |                                                                           |                                                                                                        | 60        | medium   |
|                                                    |                                                                           |                                                                                                        | 30        | high     |
|                                                    |                                                                           |                                                                                                        | 0         | critical |
| CustomerRegionWideVersionRule                      | Checking all the nodes in customer region for the same HyperStore version | The majority of customer region nodes are on version %s, while the following list of nodes differs: %s | 0         | critical |
| CustomerRegionVersionEndOfLifeRule                 | Checking HyperStore version to see if it is end of life                   | HyperStore version %s is no longer supported                                                           | 7.4.0     | medium   |
| NodesBalancedInCustomerRegionRule                  | Checking that all of the nodes in customer region are balanced            | The following nodes in data-center %s are not within plus or minus of %.2f standard deviation(s): %s   | 1         | medium   |

| Alert                                 | Description                                                                    | Alert Message                                                                                                | Threshold | Severity |
|---------------------------------------|--------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------|-----------|----------|
| CustomerRegionUsedSpacePercentageRule | Checking that the used space of the customer region does not exceed thresholds | The customer region used space percentage has exceeded the threshold %2f%%. Current percentage used is %2f%% | 70        | low      |
|                                       |                                                                                |                                                                                                              | 80        | medium   |
|                                       |                                                                                |                                                                                                              | 90        | high     |
| NodeFilesystemBalancedRule            | Checking that all of the filesystems in node are balanced                      | The following filesystem(s) are not within plus or minus of %2f standard deviation (s): %s                   | 1         | medium   |
| DiskSmartStatusRule                   | Checking that the disk's SMART status is ok or passed                          | Disk has a SMART Status of %s                                                                                | 0         | high     |
| FilesystemStatusRule                  | Checking that the filesystem's status is ok                                    | Filesystem has a Status of %s                                                                                | 0         | high     |
| FilesystemCassandraAvailableSpaceRule | Checking that the Cassandra filesystem's available space is sufficient         | Cassandra filesystem has an available percentage value of %d which is below the threshold of %d              | 70        | low      |
|                                       |                                                                                |                                                                                                              | 60        | medium   |
|                                       |                                                                                |                                                                                                              | 50        | medium   |
|                                       |                                                                                |                                                                                                              | 40        | critical |
| FilesystemAvailableSpaceRule          | Checking that the filesystem's (Non-Cassandra) available space is sufficient   | Filesystem has an available percentage value of %d which is below the threshold of %d                        | 30        | low      |
|                                       |                                                                                |                                                                                                              | 20        | medium   |
|                                       |                                                                                |                                                                                                              | 10        | high     |

| Alert                                   | Description                                                             | Alert Message                                                       | Threshold        | Severity |
|-----------------------------------------|-------------------------------------------------------------------------|---------------------------------------------------------------------|------------------|----------|
| CustomerRegionOOMErrorsRule             | Checking for OOME in the logs                                           | %d %s were found in the error logs                                  | OOME             | high     |
|                                         |                                                                         |                                                                     | HyperStore-OOME  | high     |
|                                         |                                                                         |                                                                     | Redismon-OOME    | high     |
|                                         |                                                                         |                                                                     | S3-OOME          | high     |
|                                         |                                                                         |                                                                     | Cassandra-OOME   | high     |
| CustomerRegionHsDiskErrorsRule          | Checking for HSDISKERROR in the logs                                    | "%d HSDISKERROR(s) were found in the error logs                     | HSDISKERROR      | high     |
| CustomerRegionMaintenanceModeErrorsRule | Checking for Maintenance Mode nodes                                     | %s is in Maintenance Mode                                           | Maintenance Mode | high     |
| CustomerRegionLicenseErrorsRule         | Checking for License errors in the logs                                 | %d License error(s) were found in the error logs                    | License          | high     |
| CustomerRegionMicron7450Rule            | Checking for Micron 7450 drives with certain firmware                   | The following nodes contain Micron 7450 drives with firmware %s: %s | E2MU110          | medium   |
| RaidFailedMembersRule                   | Checking for RAID's that have failed members                            | RAID has failed member drives                                       | 0                | critical |
| NetworkInterfaceLinkDownRule            | Checking for NIC that have ONBOOT=true and LINK-K=false                 | Network interface is offline                                        | 0                | high     |
| NetworkInterfaceBondDegradedRule        | Checking for bonds down slave interfaces                                | Bond interface is offline                                           | 0                | high     |
| ServiceRedisSlaveDataCenterRule         | Checking that the Redis Monitor slaves are in the same DC as the master | Redismon slave services are not in the same DC as the master ser-   | 0                | medium   |

| Alert                                          | Description                                                                                                 | Alert Message                                                                                                                                                                                              | Threshold  | Severity |
|------------------------------------------------|-------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|----------|
|                                                | ter Redis services                                                                                          | vice                                                                                                                                                                                                       |            |          |
| ServiceRedisCredSlaveCountRule                 | Checking that the Redis Credential service has enough slaves for the number of nodes. Value is the divisor. | Number of redis credential slave services is too low, recommended = %d, actual = %d                                                                                                                        | 5          | low      |
| ServiceRedisQosSlaveCountRule                  | Checking that the Redis QoS service has enough slaves for the number of nodes. Value is the divisor.        | Number of redis qos slave services is too low, recommended = %d, actual = %d                                                                                                                               | 5          | low      |
| ServiceCassandraLargePartitionUnder3NodesRule  | Checking for large Cassandra partitions                                                                     | Cassandra partition size is %.2f MB on less than 3 nodes, which is larger than the threshold of %.2f MB. Compaction may need to be run on these nodes. Node = %s, Keyspace = %s, Bucket = %s, Prefix = %s" | 1073741824 | medium   |
| ServiceCassandraLargePartition3NodesOrMoreRule | Checking for large cassandra partitions                                                                     | Cassandra partition size is %.2f MB on 3 or more nodes which is larger than the threshold of %.2f MB. Keyspace = %s, Bucket = %s, Prefix = %s                                                              | 1073741824 | high     |

| Alert                                 | Description                                                                                            | Alert Message                                                                                                                     | Threshold  | Severity |
|---------------------------------------|--------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------|------------|----------|
| EDACCorrectableMemoryErrorsRule       | Checking for correctable memory errors                                                                 | Correctable memory error count is %d which is larger than the threshold of %d. Node = %s, Controller = %s, CSRow = %s             | 150        | critical |
| EDACUncorrectableMemoryErrorsRule     | Checking for uncorrectable memory errors                                                               | Uncorrectable memory error count is %d which is larger than the threshold of %d. Node = %s, Controller = %s, CSRow = %s           | 0          | critical |
| VeeamHSVersionWithO365Rule            | Checking the customer region version is at least at the threshold level if the customer has Veeam O365 | Customer region has Veeam O365 (Veeam-7.0) but the Hyperstore version %s is below the threshold of %v                             | 7.5.2      | medium   |
| VeeamSinglePolicySmallObjectsRule     | Checking for customer regions with a single storage policy, small objects, and Veeam                   | Customer region has a single storage policy, veeam, and average object size of %.2f KB that is less than the threshold of %.2f KB | 1048576    | medium   |
| VeeamWrongHSVersionLargePartitionRule | Checking for customer regions with a HyperStore version less than the threshold, with Veeam bucket     | Bucket %s is on a Cassandra partition larger than %.2f and is running Veeam on a non-supported                                    | 1073741824 | medium   |

| Alert                                 | Description                                                                                                                                                      | Alert Mes-<br>sage                                                                                                                                                                           | Threshold  | Sever-<br>ity |
|---------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|---------------|
|                                       | ets on a large Cassandra partition                                                                                                                               | Hyperstore ver-<br>sion                                                                                                                                                                      |            |               |
| VeeamRightHSVersionLargePartitionRule | Checking for customer regions with a HyperStore version greater than or equal to the threshold, with Veeam O365 buckets on a large Cas-<br>sandra par-<br>tition | Bucket %s is on a Cas-<br>sandra par-<br>tition larger than %.2f and is running Veeam Office 365 on a sup-<br>ported Hyper-<br>store version. Check that bucket rules are set cor-<br>rectly | 1073741824 | mediu-<br>m   |
| NodeMaxS3ConcurrencyRule              | Checking for nodes with more con-<br>current trans-<br>actions than the threshold                                                                                | The following %s nodes (counts) exceeded the threshold of %d concurrent S3 trans-<br>actions in the last 7 days: %s                                                                          | 160        | mediu-<br>m   |
|                                       |                                                                                                                                                                  |                                                                                                                                                                                              | 1000       | mediu-<br>m   |
| NodeCassandraRunningOnHDDRule         | Checking for nodes that are running Cas-<br>sandra on spinning hard drives                                                                                       | Node %s has its Cassandra partition loc-<br>ated on non-SSD drive %s                                                                                                                         | HDD        | critical      |
| NodesS3OpsBalancedRule                | Checking that the nodes in a DC have bal-<br>anced S3 traffic                                                                                                    | Nodes in data-center %s are out of balance with a stand-<br>ard deviation greater than %.2f%s                                                                                                | 0.01       | low           |