Quick-Start Guide:
Deploying Your Cloudian HyperStore Hybrid Storage Service

Contents

Intended Audience

The purpose of this document is to help a new user deploy a 3-node Cloudian storage cluster in your datacenter for use with the Cloudian HyperStore Hybrid Cloud Service from AWS Marketplace. This document is a quick start guide and does not replace the existing HyperStore Installation Guide or HyperStore Administrator's Guide. Refer to those guides for more details and for larger deployments.

Setting Up the AWS Proxy Server

A required part of the Cloudian HyperStore Hybrid Cloud Service is an AWS Proxy Server that runs on AWS EC2. Your on-premise HyperStore system will transmit system usage data to the AWS Proxy Server which will relay the usage data to the AWS Metering Service for billing purposes. Follow these steps to deploy your AWS Proxy Server.

1. **From the Cloudian HyperStore Hybrid Storage Service product page, launch an EC2 instance**.

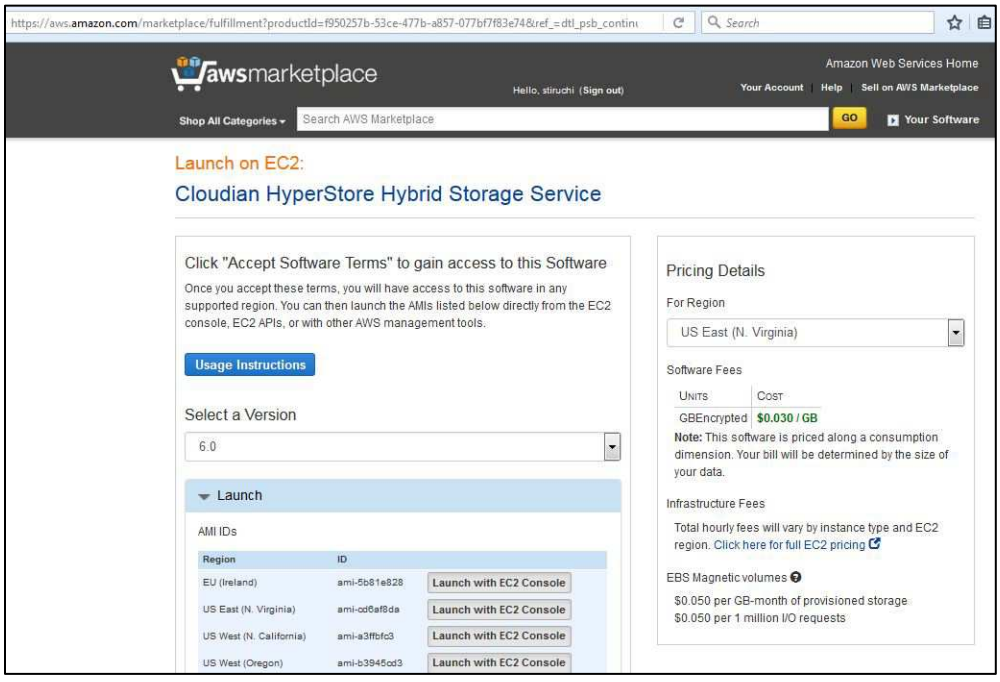

2. **During the initial configuration and launch of your AWS Proxy Server instance, apply these required attributes to the instance**. (Note: If you have already completed the initial launch of your instance, then in the EC2 console's **Instance** section you can apply these attributes and then reboot the instance.)

   IAM Role

   The instance must have assigned to it an IAM Role that supports AWS Marketplace Metering Full Access. During initial instance configuration, in the **Configure Instance** panel, select to Create a New IAM Role. Give the new role any name and then attach to it the policy "AWSMarketplaceMeteringFullAccess". Be careful to select "**AWSMarketplaceMeteringFullAccess**", not just "AWSMarketplaceFullAccess". (Note: If you assigned an insufficient IAM Role to the instance during initial configuration, you can edit the Role through the EC2 Console's **Instance** section.)

   Security Group Firewall Rules

   The instance's Security Group firewall rules must include the following:

   - Custom TCP Rule for port 17081 (this will be accessed by your HyperStore cluster)
   - Custom TCP Rule for port 8888 (this allows you to access the AWS Proxy Server's console through a browser)

   You can set these Security Group rules during initial instance configuration or subsequently through the EC2 Console's **Instance** section.

3. **In the EC2 Console's Instance section, make a note of the following information for your instance**:
   - The Instance Public IP address. You will need this when installing your on-premise HyperStore cluster and also when accessing the AWS Proxy Server console to monitor your cluster usage. In this example, the Public IP address is *54.200.9.207*
   - The instance ID. You will need this when accessing the AWS Proxy Server console. In this example, the instance ID is *i-2a8e5ab3*

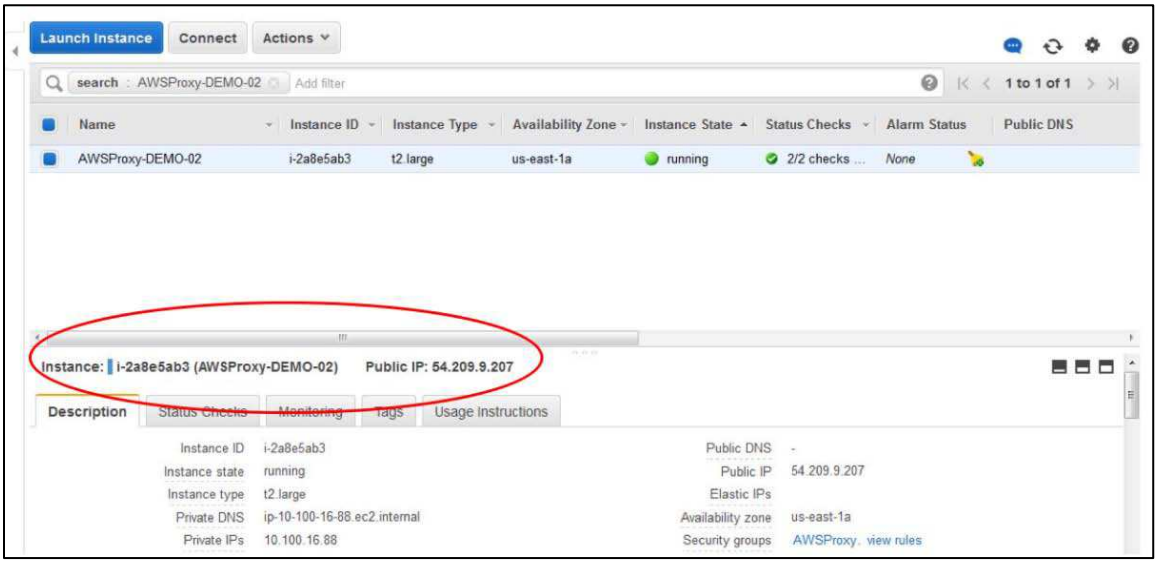

You have successfully deployed an AWS Proxy Server.

Node Requirements for Your On-Premise Cloudian Cluster

3 Intel compatible servers each with the following minimum configuration:

- Processor: 1 CPU, 8 cores, 2.4GHz
- Memory: 32GB
- Disk: 12 x 2TB HDD, 2 x 160GB SSD (12 drives for data, 2 drives for OS/Metadata)
- RAID: RAID-1 recommended for the OS/Metadata, JBOD for the Data Drives
- Network: 2x10GbE Port

Example Cloudian HyperStore Setup

In the remainder of this document, we will be installing a 3-node Cloudian HyperStore cluster.
Our servers are:

- cloudiannode1 IP address: 192.168.2.41
- cloudiannode2 IP address: 192.168.2.42
- cloudiannode3 IP address: 192.168.2.43

We are using our DNS server for service endpoint resolution. The domain is: dom.local

Preparing Cluster Node Hardware

1. **Choose any one of your nodes and use the Cloudian HyperStore ISO for OS installation .** This will install CentOS on along with the HyperStore software binary files. An ISO is a bootable image which will deploy the OS software, related packages and Cloudian software onto a node.

   Follow these steps to install CentOS from the HyperStore Appliance ISO file:

   a. Burn the HyperStore Appliance ISO file on to either a CD or a DVD. **Be sure to use a "burn image" operation**, not a simple copy operation.

   b. Insert the CD or DVD into the host machine.

   c. Start the host machine (or restart it if it's already running) and use your hardware-specific hot key to enter the BIOS setup utility.

   d. Use the BIOS setup utility to change the machine's boot order so that it boots first from the CD/DVD drive. When you save your changes, the machine will restart and boot from the ISO file. After a few moments, the CentOS Installer welcome screen will briefly appear.

   e. Select the first option in the welcome menu, "Install or upgrade an existing system", and press enter. If you do not make a selection within 5 seconds, the "Install or upgrade an existing system" option automatically executes.

      - The console will display progress information as the packages are installed. There will be a pause for a couple minutes after the last package is installed, and then post-installation scripts will run. Altogether the system installation will typically take less than 10 minutes to complete.

f. After the install has completed, the CD/DVD tray ejects and the system begins to reboot. Remove the CD/DVD from the tray when it ejects so that the machine does not boot from the CD/DVD again.

CentOS 6.8 and HyperStore prerequisites are now installed on the host, and the HyperStore binaries and host set-up tool have been placed on the host.

2. Configure the node's basic OS and network settings

a. Log in to the node as *root*. The default root password is *password*. It is our best practice recommendation to update your root access for production environments, which we will do later in the process.

b. cd to */root/CloudianTools* directory.

c. Run *configure_appliance.sh* to perform the following tasks:

i. Change the default root password

Although it is possible to configure a different root password for each node in the cluster, this would prevent the network validation tool running, so bear this in mind when setting passwords.

ii. Change time zone

iii. Configure network. For this task the interactive script will ask you to provide the information below:

- Hostname -- Hostname of the machine. Do not use capital letters in the hostname, and do not use *localhost* as the hostname.
- Network Interface -- Network interface device to which these configurations will be applied; for example, "eth0".
- IP Address -- IP (v4) address of the network interface being configured. This must be a static IP address (not DHCP). Do not use the loopback address (127.0.0.1). Do not use IPv6.
- Netmask -- Netmask for the subnet in which this host machine will be deployed.
- Default Gateway - IP address of the default gateway device for the subnet in which this host machine will be deployed.
- Primary DNS server -- IP address of the DNS server that this host will use.
- Secondary DNS server - IP address of the secondary DNS server that this host will use, if any.
- Domain -- Domain to which this host will belong. This is the fully qualified domain name (FQDN) of the host, minus the hostname segment. For example:
  - If the host's hostname is *cloudiannode1* and its FQDN will be *cloudiannode1.enterprise.com*, then specify *enterprise.com* as the domain.
  - If the host's hostname is *cloudiannode1* and its FQDN will be *cloudiannode1.west.enterprise.com*, then specify *west.enterprise.com* as the domain.

3. Format and mount available disks for Cloudian S3 data storage

The script will automatically configure and mount available disks as *ext4* file systems, for storage of HyperStore S3 objects. The mount points will be named as */cloudian1*, */cloudian2*, and so on. Console output will inform you of the progress for each disk as the script creates a partition, creates

an *ext4* file system, and mounts the file system. The image below highlights one of the mount points in the results of a *df –k* command run after the disk formatting has completed.

```
[root@cloudiannode1 CloudianPackages]# df -k
Filesystem           1K-blocks    Used Available Use% Mounted on
/dev/mapper/vg0-root  94459756 2314420  87340320   3% /
tmpfs                  4030620       0   4030620   0% /dev/shm
/dev/sda1               487652   30212    431840   7% /boot
/dev/sdb1            203014840   60684 202954156   1% /cloudian1
```

If the node has 10 or more drives, then during the ISO install the two smallest drives were mirrored and used for the OS. Those drives will also store Cassandra data, Cassandra commit logs, and Redis data. The rest of the drives are considered available for the *configure_appliance.sh* script to mount as *ext4* file systems for S3 object storage.

4. **Reboot the node.**
   Rebooting is mandatory in order to apply the network configuration settings that you specified.

   If necessary for your particular hardware, on reboot enter the BIOS utility and set the boot order back to booting first from the hard drive rather than the CD/DVD drive.

5. **Repeat this procedure on every node in the HyperStore cluster**. Use the same HyperStore ISO CD/DVD for each host.

Installing the Cloudian HyperStore Cluster

1. **Under */root/CloudianPackages* on one of your nodes, run the following command to extract the package content on the node**:

   ./CloudianHyperStore-6.0.4-mms.bin cloudian_awsmms.lic

   Among other things, this unpacks the HyperStore product documentation into the directory */root/CloudianPackages/doc*. This includes an Installation Guide and an Administrator's Guide.

   Note: This node will be designated the Puppet master node for cluster configuration management.

2. **Pre-configure your HyperStore cluster with the information it needs to connect to your AWS Proxy Instance:**

   a) Change into this configuration directory:

   /etc/cloudian-6.0.4-puppet/modules/cloudians3/templates

   b) Then open the *mts.properties.erb* file in a text editor. In this file, find the *awsmmsproxy.host* property and set it to the public IP address of your AWS Proxy Server. For example, with our AWS Proxy Server example from earlier in this document it would be:

   awsmmsproxy.host=54.200.9.207

c) Save your change and close the file.

3. **Return to the */root/CloudianPackages* directory, and create a copy of the *sample-survey.csv* file to a new file *survey.csv***:

```
[root@cloudiannode1 CloudianPackages]# cp sample-survey.csv survey.csv
[root@cloudiannode1 CloudianPackages]#
```

4. **Edit the new *survey.csv* file** to specify the region, the node name, IP address, datacenter name, and rack name of your Cloudian HyperStore S3 cluster:

```
# Survey file for multinode installation
#
# See your Cloudian HyperStore Installation Guide for
# detailed instructions on other options that can be
# used in this file.
#
# 1 Region, region1, with three hosts
region2,cloudiannode1,192.168.2.41,DC1,RAC1
region2,cloudiannode2,192.168.2.42,DC1,RAC1
region2,cloudiannode3,192.168.2.43,DC1,RAC1
```

In this example, our *survey.csv* file has the following configuration for the 3-node cluster:

- region2,cloudiannode1,192.168.2.41,DC1,RAC1
- region2,cloudiannode2,192.168.2.42,DC1,RAC1
- region2,cloudiannode3,192.168.2.43,DC1,RAC1

Note: You can optionally add a sixth comma-separated value to specify a node's NIC for internal cluster communication. If your nodes have multiple NICs, the install script (which you will run later in this procedure) will ask you to specify the default NIC for internal cluster communication. You only need to add the sixth field to *survey.csv* if not all nodes will use the same NIC for internal cluster communication.

5. **Validate that your nodes and network environment meet HyperStore requirements**.

The HyperStore package includes a *preInstallCheck.sh* script that you can use to assess your host nodes and network environment to ensure that they meet the technical requirements to support the next stage of HyperStore deployment. The script will output warnings and errors that may cause the software installation to fail if not rectified.

Note: The *preInstallCheck.sh* script is in alpha release. This script is part of our free use tools. It is distributed in the hope that it will be useful, but without any warranty or support. It is provided "AS IS" without any warranty of any kind.

Note: Node and environment requirements are described in the Cloudian HyperStore Installation Guide (under */root/CloudianPackages/doc*).

The *preInstallCheck.sh* script is in the directory */root/CloudianPackages/preinstall*. Change into that directory and launch the script as follows:

```
./preInstallCheck.sh
```

On launch, the script displays a menu.

At the prompt enter option **R** to run the script.

The script prompts you for the root user's password. The script only supports one password, so if different passwords have been used across nodes, this will prevent the script running effectively.

At the end of its run the script provides a summary that outlines any environmental issues that have been detected. "Warning" issues should be reviewed but don't necessarily require action. "Error" issues must be resolved before continuing with the install. In this example there are warnings but no errors.

When you're done reviewing the results, press any key to continue and then at the menu choose to exit the script. If you make any system changes to resolve errors found by the *preInstallCheck.sh* script, run the script again to verify that your environment meets HyperStore requirements before proceeding to the next step.

6. **Back in the */root/CloudianPackages* directory run *./cloudianInstall.sh***

   Note: If deploying in a non-production environment that lacks a DNS server, launch your Cloudian install with ***./cloudianInstall.sh configure-dnsmasq***. This uses the bundled open source domain resolution utility *dnsmasq* to resolve HyperStore service endpoints.

   When you run the install script it opens a menu. Select *'**Install Cloudian HyperStore**'*.
   When prompted, input *survey.csv* as your survey file name. You will also need to enter the root password. Continue with the setup by responding to the prompts. Press enter to accept a default.

   In the following screenshot, the information that we had provided in the *survey.csv* file is used in the Cloudian HyperStore cluster configuration. Note: In this setup we are using a DNS server for domain name resolution.

```
Configure cluster
-----------------


Use default network interface for all traffic.

Region [region2] Cassandra cluster name: Cloudianregion2

Please enter the service metadata replication strategy for region2 [DC1:3]:

NTP time server(s) for region region2:
Please enter your NTP time server(s) [0.centos.pool.ntp.org,1.centos.pool.ntp.org,2.centos.pool.


Cloudian HyperStore(R) S3 service endpoints are based on your desired top
level DNS domain name.  For example, yourcompany.com.
Please enter your top level domain name [mycloudianhyperstore.com]: dom.local

Service endpoints for region region2:
Region [region2] S3 service domain URL [s3-region2.dom.local]:
Region [region2] S3 Web site end point [s3-website-region2.dom.local]:

S3 Admin service endpoint: s3-admin.dom.local

Domain name of your Cloudian Management Console service [cmc.dom.local]: cmc-region2.dom.local
Cloudian Management Console service endpoint: cmc-region2.dom.local


Configurations saved to file [./CloudianInstallConfiguration.txt].

Your existing Puppet manifest files are backed up in ./manifests.20160802112106

Install Cloudian HyperStore(R) packages
---------------------------------------
```

   Notice that at the prompts we provided our domain *dom.local* and we revised the CMC service endpoint to *cmc-region2.dom.local*, and for the other service endpoints we accepted the defaults that the script provides. If you are using a DNS server, update your forward and reverse lookup zones to include the correct service names for your Cloudian HyperStore cluster.

Your Cloudian HyperStore S3 Cloud Storage is now up and running. You have successfully deployed a 3-node HyperStore cluster.

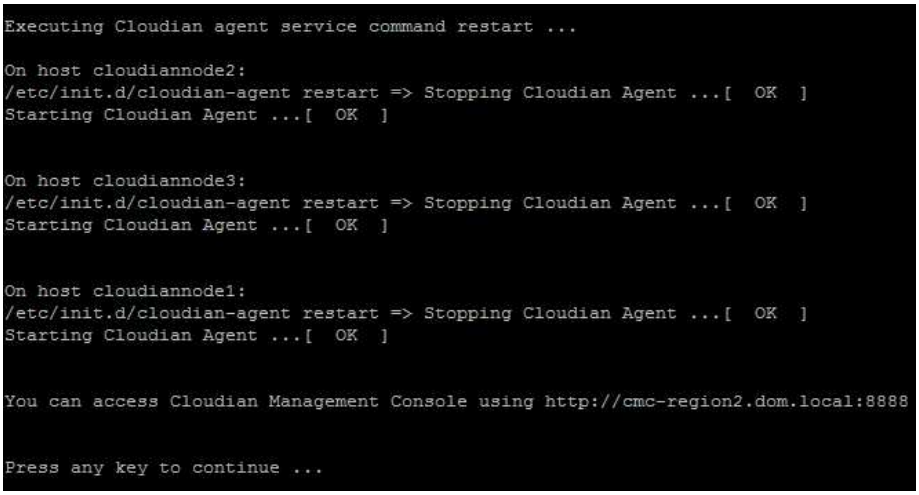

Completing Your Cloudian HyperStore Setup

With a browser access your Cloudian Management Console (CMC) at *http://<CMC_service_endpoint>:8888/Cloudian*. The default *System Admin* group user id is *admin* and the default password is *public*. It is our best practice recommendation to update your default admin password in a production environment (after you have logged into the CMC, you can change your password by selecting **Security Credentials** under your login name).

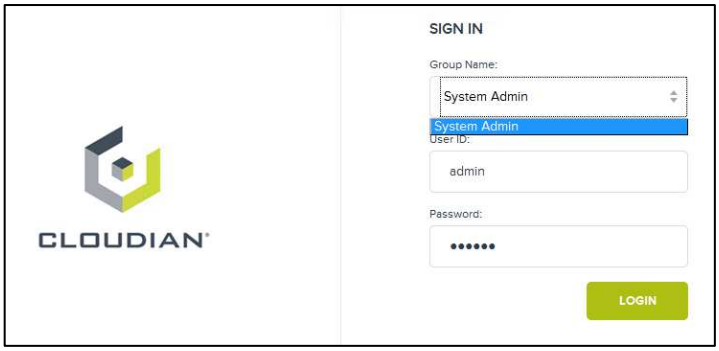

Note: If you do have a DNS server in your environment, your Windows/Linux client that will access your Cloudian HyperStore must have mappings of IP addresses to host names. For a non-production setup that does not have a DNS server, update your Windows/Linux client host file to include the name resolution of your Cloudian HyperStore cluster. For example, include the following in your hosts file:

- The CMC service endpoint (such as *cmc-region2.dom.local*)
- The S3 service endpoint (such as *s3-region2.dom.local*)
- The FQDN for the storage bucket that you will create (such as *bucket1.s3-region2.dom.local*). You can update the bucket name resolution when you or other users create additional buckets.

In the Cloudian Management Console we will create a storage policy, a user group, and an individual user.

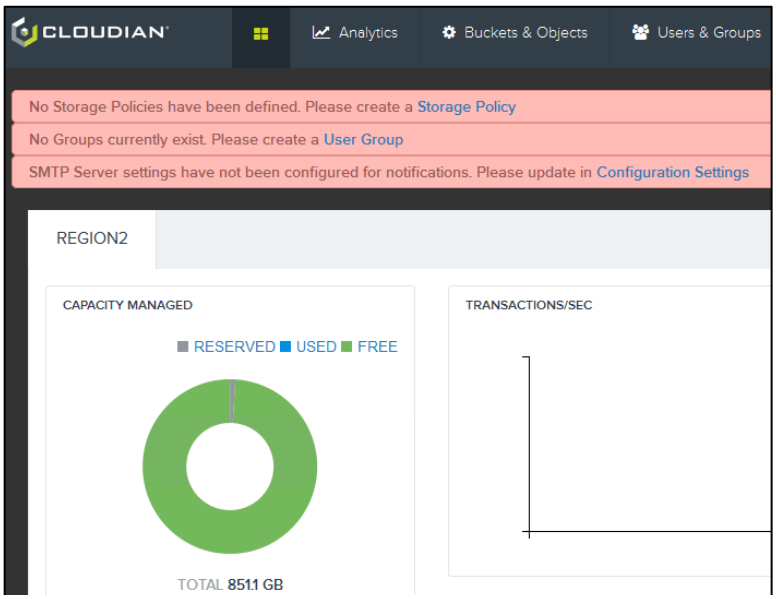

Select **Cluster -> Storage Policy** to configure a storage policy. Our first storage policy will be created with the following settings:

- DATA DISTRIBUTION SCHEME: Replicas within a Single Datacenter
- NUMBER OF REPLICAS: 1 (note that in a production system this should be 2 or more)
- DATACENTER ASSIGNMENT: Not relevant in our single-DC environment
- DATA CONSISTENCY LEVEL: Strong
- METADATA CONSISTENCY LEVEL: Strong
- GROUP VISIBILITY: Unspecified (which makes this the default policy, visible to all groups)

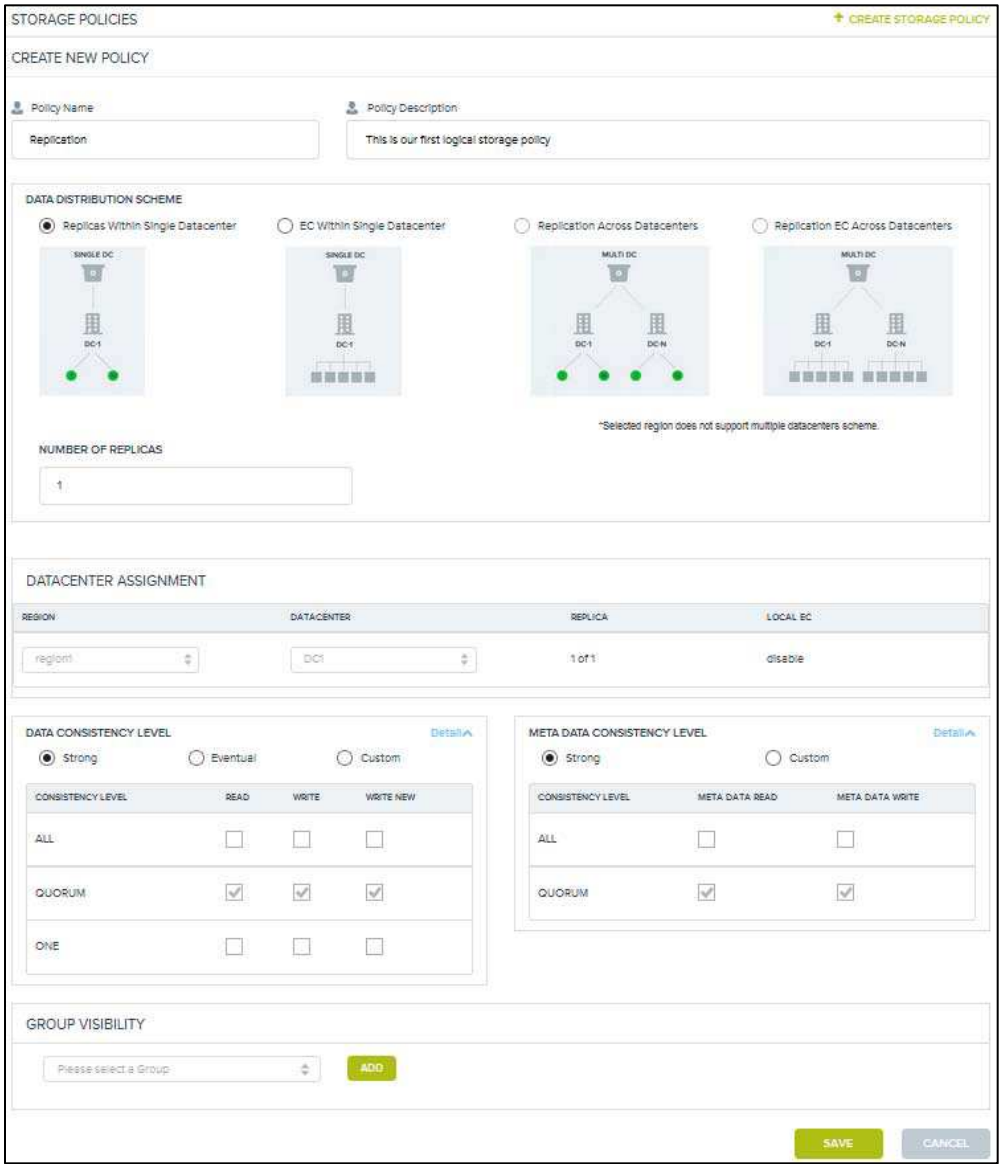

Click **Save** to complete the storage policy creation.

The next step in our setup is to create groups and users. Select **Users & Groups -> Manage Groups -> +New Group**. Create a new group then click **Save**. (Click **Help** for additional detail and options. For example, you can set quality of service controls for groups and users that you create.)

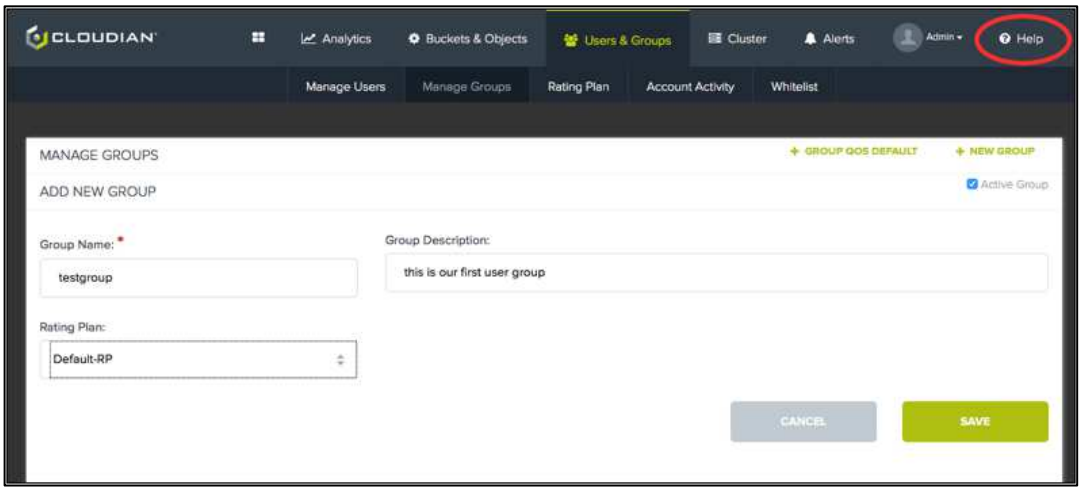

In our next step, we will create a user account. Select **Users & Groups -> Manage Users -> +New User**. Create a new user, and make a note of the user ID and password. Click on *More* for additional optional parameters that you can set for your user. Click **Save** when you're done**.**

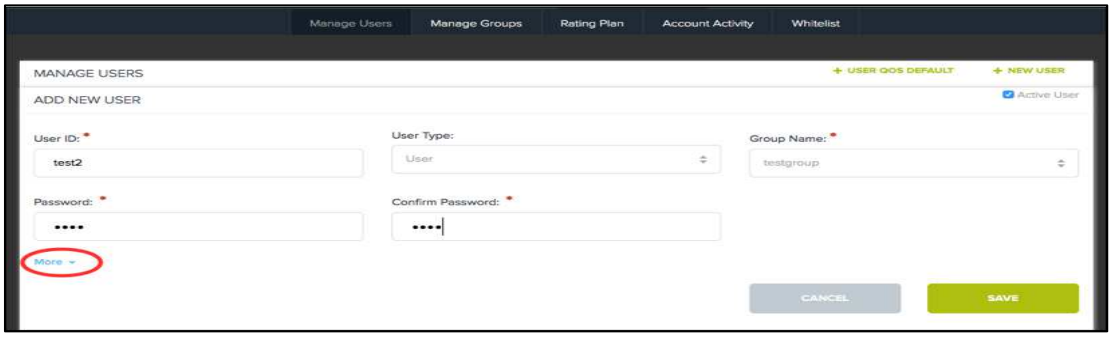

You have successfully completed the process of setting up a 3-node Cloudian HyperStore cluster, creating a storage policy, and creating a user group and a user. Sign out of the CMC as the administrator. As a user, log in with the account that was just created.

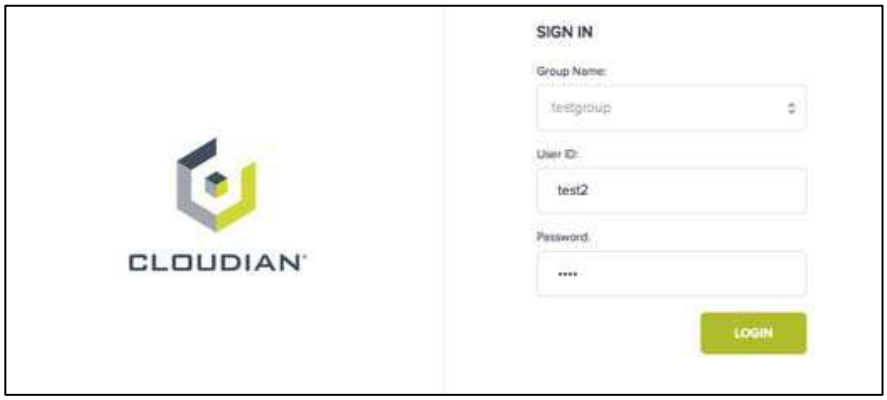

As the Cloudian HyperStore S3 cloud storage user, you can dynamically create your own bucket and start uploading objects.

The setup of a Cloudian S3 cloud storage solution is really that easy. Your users can start uploading, sharing files, pictures and also use it as a scale-out object storage for applications and analytics that support S3.

Accessing the AWS Proxy Console

You can access your AWS Proxy Server console from a supported web browser at the following URL:

http://<PUBLIC_IP_ADDRESS_OF_AWS_PROXY_INSTANCE>:8888/console

- In this example, it *is 'http://54.209.9.207:8888/console'*
- The login username is: *admin*
- The password is: *Instance ID*

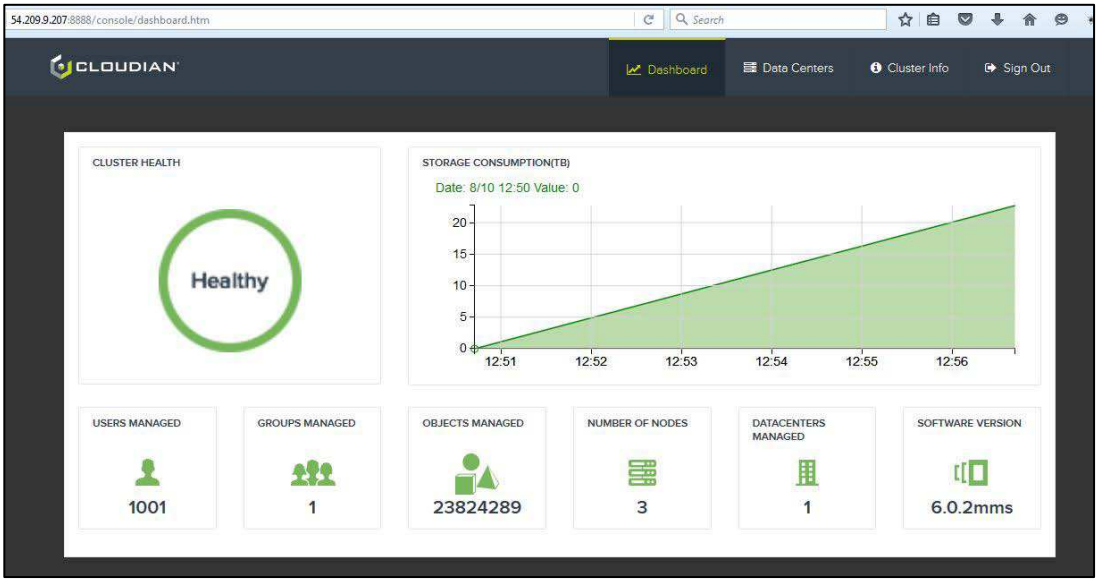

Note: The AWS Proxy Console provides limited, high level monitoring capabilities. For more complete monitoring and management of your HyperStore system, use your on-premise Cloudian Management Console. See your Cloudian HyperStore product documentation for detail.