Veritas CloudPoint 1.0 Administrator's Guide

Veritas CloudPoint Administrator's Guide

Last updated: 2017-09-13

Document version: 1.0 Rev 6

Legal Notice

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. VERITAS TECHNOLOGIES LLC SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

Veritas Technologies LLC
500 E Middlefield Road
Mountain View, CA 94043

http://www.veritas.com

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

https://www.veritas.com/support

You can manage your Veritas account information at the following URL:

https://my.veritas.com

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)	CustomerCare@veritas.com
Japan	CustomerCare_Japan@veritas.com

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

doc.feedback@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

http://www.veritas.com/community/

Contents

Chapter 1

Deploying CloudPoint

This chapter includes the following topics:

About the deployment approach

CloudPoint is distributed as a Docker container image that is built on an Ubuntu 16.04 Server Long Term Support (LTS) base image. The image contains the following:

- All CloudPoint services
- MongoDB
- RabbitMQ server
- A web server to support the CloudPoint graphical user interface (GUI)

This deployment approach has the following advantages:

- All the packages and scripts you need to deploy CloudPoint are in the container itself.

- There are minimal installation requirements.
- Deployment requires only a few commands.

Supported snapshots

The following table shows the types of snapshots you can take for each supported application, cloud, or storage array.

Snapshot type	Supported assets
Application	Oracle 12c
Host	AWS EC2 instances Azure virtual machine Google virtual machine
Disk	Hitachi Data Systems G-Series HPE 3PAR

Note: File system level snapshots are not supported. On the File Systems pages, when you select **Create Snapshot**, CloudPoint takes a disk level snapshot of the file system.

Deciding where to run CloudPoint

You can deploy CloudPoint in the following ways:

- Deploy CloudPoint on-premises and manage on-premises assets.
- Deploy CloudPoint on-premises and manage assets in one or more clouds.
- Deploy CloudPoint in a cloud and manage assets in that cloud.
- Deploy CloudPoint in a cloud and manage assets in multiple clouds.

System requirements

Table 1-1 Supported applications, operating systems, and platforms

Category	Support
Applications	Oracle 12c* single node; CloudPoint has been verified on Oracle 12c and Oracle 12cR1 Linux native file systems: ext2, ext3, ext4, and XFS
Operating system	Red Hat Enterprise Linux (RHEL) 7.x; CloudPoint has been verified on RHEL 7.1, 7.2, and 7.3
Cloud platforms	Amazon Web Services Microsoft Azure Google Cloud
Storage platforms	Hitachi Data Systems (HDS) HPE 3PAR

Table 1-2 Minimum system requirements

Cloud vendor	Requirements
Amazon Web Services	Elastic Compute Cloud (EC2) instance type: t2.medium vCPUs: 2 Memory (GB): 4 with a solid-state drive (SSD) for the root disk Storage: 50 GB Elastic Block Store (EBS) volume with encryption for the snapshot asset database
Microsoft Azure	Virtual machine type: Standard_DS2_v2 CPU cores: 2 Memory (GB): 7 with an SSD for the root disk Storage: 50 GB Premium SSD for the snapshot asset database

Table 1-2 Minimum system requirements *(continued)*

Cloud vendor	Requirements
Google Cloud	Virtual machine type: Ubuntu 16.04 Server LTS instance vCPUs: 2 Memory (GB): 7.5 with a standard persistent disk Storage: 50 GB SSD persistent disk for the snapshot asset database
x86 physical host	Operating system: Ubuntu 16.04 Server LTS CPUs: Single-socket, multi-core Memory: 10 GB with an addition 50 GB for the snapshot asset database

CloudPoint also has the following space requirements.

Table 1-3 Space considerations

Item	Space requirements
CloudPoint docker container	< 1 GB
On-host agent and plug-ins	~ 20 MB

Creating and mounting a volume to store CloudPoint data

Before you deploy CloudPoint, you should create and mount a volume to store CloudPoint data.

See "Creating a volume and file system to store CloudPoint data (Amazon AWS)" on page 12.

See "Creating and mounting a volume to store CloudPoint data (Google Cloud)" on page 12.

See "Creating and mounting a volume to store CloudPoint data (Microsoft Azure)" on page 12.

Creating a volume and file system to store CloudPoint data (Amazon AWS)

To create a volume and file system to store CloudPoint data (Amazon AWS)

1. On the EC2 dashboard, click **Volumes > Create Volumes**.
2. Follow the instructions on the screen and specify the following:
   - Volume type: General Purpose SSD
   - Size: 50 GB
3. Use the following instructions to create a file system and mount the device to `/cloudpoint` on the instance host.

   http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-using-volumes.html

Creating and mounting a volume to store CloudPoint data (Google Cloud)

To create and mount a volume to store CloudPoint data (Google Cloud)

◆ Create the disk for the virtual machine, initialize it, and mount it to `/cloudpoint`.

https://cloud.google.com/compute/docs/disks/add-persistent-disk

Creating and mounting a volume to store CloudPoint data (Microsoft Azure)

To create and mount a volume to store CloudPoint data (Microsoft Azure)

1. Create a new disk and attach it to the virtual machine.

   https://docs.microsoft.com/en-us/azure/virtual-machines/linux/attach-disk-portal

   You should choose the managed disk option.

   https://docs.microsoft.com/en-us/azure/virtual-machines/linux/attach-disk-portal#use-azure-managed-disks
2. Initialize the disk and mount it to `/cloudpoint`. For details, see the section "Connect to the Linux VM to mount the new disk" in the following link:

   https://docs.microsoft.com/en-us/azure/virtual-machines/linux/add-disk

Deploying CloudPoint using a Docker image

Veritas distributes a Docker image with CloudPoint already installed. The image is located on the Veritas' customer portal, MyVeritas.

https://my.veritas.com

Before you complete the steps in this section, make sure that you are familiar with CloudPoint installation requirements.

See "System requirements" on page 10.

To deploy CloudPoint using a Docker image

1. Create the instance or prepare the physical host to install CloudPoint.
   - If you deploy CloudPoint in a public cloud, do the following:
     - Choose an Ubuntu 16.04 Server LTS instance image that meets CloudPoint installation requirements.
     - Add sufficient storage to the instance to meet the installation requirements.
   - If you deploy CloudPoint on-premises, do the following:
     - Install Ubuntu 16.04 Server LTS on a physical x86 server.
     - Add sufficient storage to the server to meet the installation requirements.
2. Install Docker for Ubuntu. Enter the following:

   ```
   sudo apt-get install docker.io
   ```

3. If you have not done so already, create a volume and file system on the host to store CloudPoint metadata and asset metadata. The specific steps depend on your cloud environment.

   See "Creating a volume and file system to store CloudPoint data (Amazon AWS)" on page 12.

   See "Creating and mounting a volume to store CloudPoint data (Microsoft Azure)" on page 12.

   See "Creating and mounting a volume to store CloudPoint data (Google Cloud)" on page 12.
4. Download the CloudPoint image (`Veritas_CloudPoint_1.0.1_IE.img`) from MyVeritas.

   If necessary, copy the downloaded image to the machine on which you will deploy CloudPoint.

5 Load the image. Enter the following:

```
# sudo docker load -i /home/ubuntu/Veritas_CloudPoint_1.0.1_IE.img
```

6 On the instance or physical host, make sure that the following ports are open:

443	CloudPoint user interface uses this port as the default HTTPS port.
5671	The RabbitMQ server uses this port for communications. This port must be open to support multiple agents.

Note: If the instance is in a cloud, configure this information under **Security Group > Inbound**.

7 Run the CloudPoint container. Enter the following:

```
# sudo docker run --name container_name -it -d \
-v full_path_to_volume_name \
-p host_port:container_port \
-p host_port:container_port loaded_image
```

For example:

```
# sudo docker run --name CloudPoint -it -d -v /cloudpoint:/cloudpoint \
-p 443:443 -p 5671:5671 veritas/cloudpoint:1.0.1
```

Note: If you do not specify the volume as `-v` *`full_path_to_volume_name/full_path_to_volume_name`*, the container writes to the Docker host file system.

8 Open your browser and point it to the host on which CloudPoint is installed.

```
https://ubuntu_docker_host_name
```

The configuration screen is displayed and the host name is added to the list of hosts on which to configure CloudPoint.

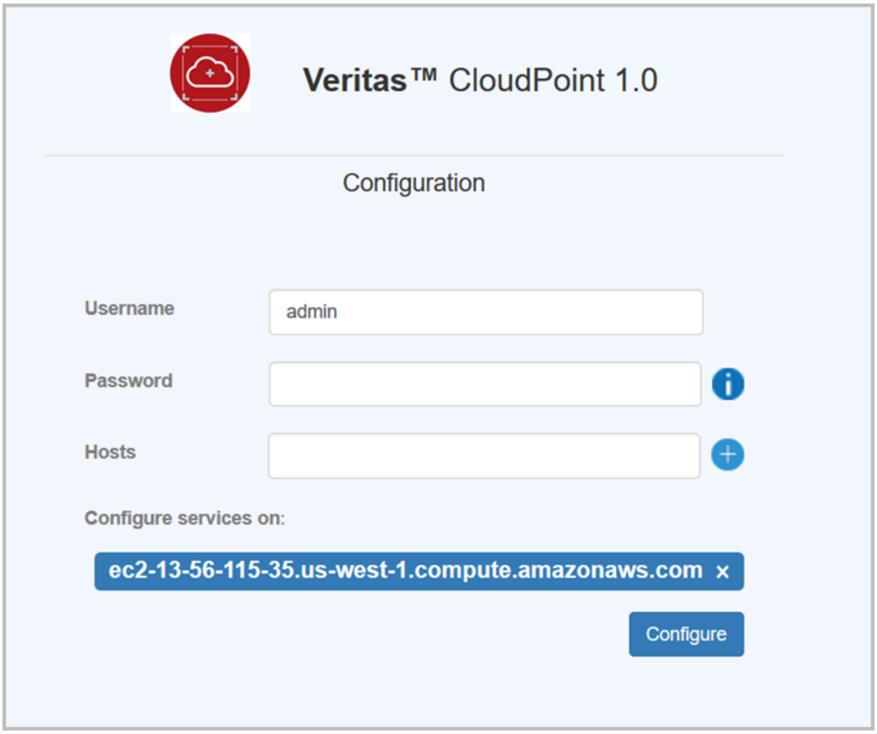

Veritas™ CloudPoint 1.0

Configuration

Username: admin

Password:

Hosts:

Configure services on:

ec2-13-56-115-35.us-west-1.compute.amazonaws.com ×

The default user name is `admin` and the default password is `cloudpoint`.

9 (Optional) If you want to add more hosts, enter the URL in the **Host** field and click **+** to add it to the list of hosts to configure.

Note: Typically only one host is configured.

---

10 Click **Configure**. The CloudPoint login screen is displayed.

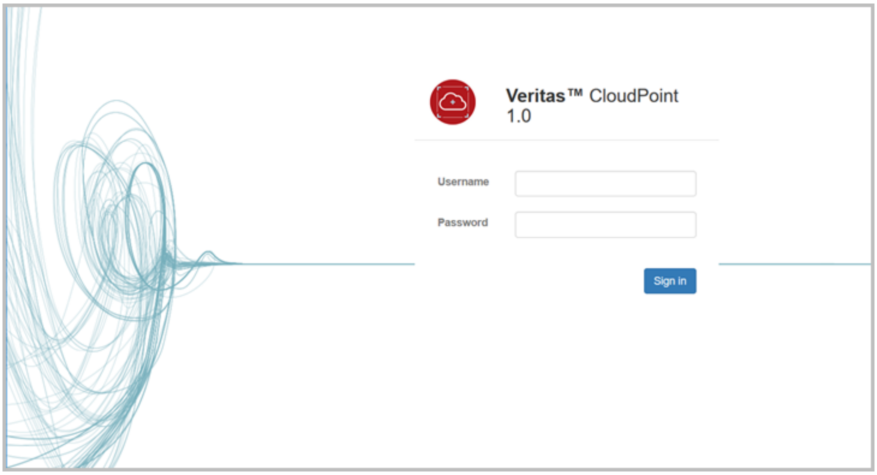

Note: It may take a few minutes for all CloudPoint services to start.

11 Sign in to CloudPoint user interface with the default user name and password.

See "Signing in to CloudPoint" on page 18.

12 On the CloudPoint user interface, verify that key components are installed and that the deployment is successful.

See "Verifying that CloudPoint was deployed successfully" on page 16.

13 Configure plug-ins.

See "Configuring an off-host plug-in" on page 52.

Verifying that CloudPoint was deployed successfully

After you deploy CloudPoint, and sign into the interface, verify that following components are installed successfully:

- Off-host agent
- Off-host plug-ins

To verify that CloudPoint was deployed successfully

1 Verify that the off-host agent is installed. Do the following:

- On the left side of any CloudPoint page, click **Agents**.
- On the **Agent** details page, verify that the Linux-based off-host agent is installed and that it is **online**. You can click the host name link for detailed information about the agent.

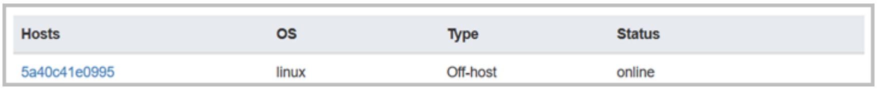

Hosts	OS	Type	Status
5a40c41e0995	linux	Off-host	online

2 Verify that the off-host plug-ins are installed. Do the following:

- On the left side of any CloudPoint page, click **Plugins**.
- On the **Plugins** details page, verify that the off-host plug-ins are installed and **online**.

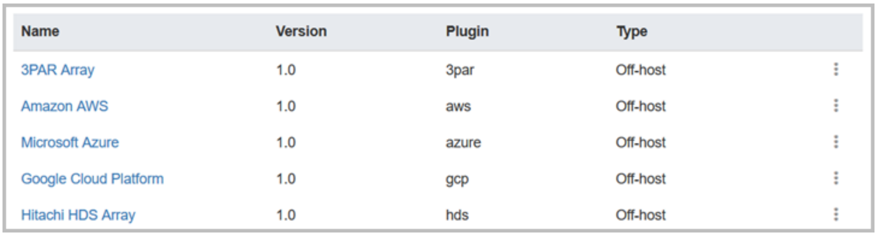

Name	Version	Plugin	Type	
3PAR Array	1.0	3par	Off-host	⋮
Amazon AWS	1.0	aws	Off-host	⋮
Microsoft Azure	1.0	azure	Off-host	⋮
Google Cloud Platform	1.0	gcp	Off-host	⋮
Hitachi HDS Array	1.0	hds	Off-host	⋮

Deploying the on-host plug-ins

On-host plug-ins are deployed on the same host at the application (or file system) you want to snapshot. CloudPoint supports the following on-host plug-ins:

- Oracle Database, for Oracle applications
- Linux FS/Devices, for Linux file systems

User interface basics

This chapter includes the following topics:

Signing in to CloudPoint

After you configure CloudPoint, the sign in screen is automatically displayed. It is also displayed any time you point your browser to the URL of the host running CloudPoint.

The first time you sign in to CloudPoint, have the default credentials ready. The default user name is `admin` and the password is `cloudpoint`. After you sign in, you can change your credentials.

To sign in to CloudPoint

1 On the sign in screen, enter your CloudPoint user name and password.

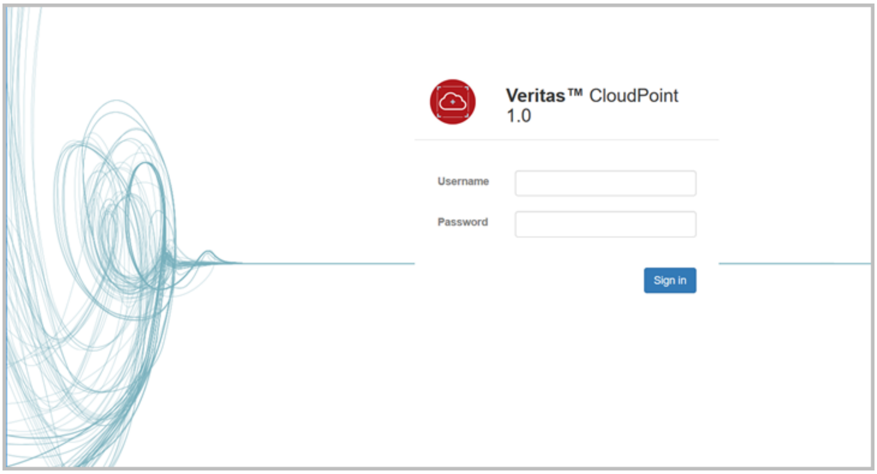

2 Click **Sign in**.

Note: If this is the first time you have signed in to CloudPoint, verify that CloudPoint was installed successfully.

See "Verifying that CloudPoint was deployed successfully" on page 16.

About the CloudPoint dashboard

The CloudPoint dashboard gives you can overview of your assets and how they are protected, as well as plug-in and snapshot data. From the dashboard, you can easily navigate to more detailed information.

When you log into CloudPoint, the dashboard displays by default. You can also access it by clicking **Dashboard** on the left side of any CloudPoint page.

Displaying information on assets, protection levels, and snapshot data

The CloudPoint dashboard displays the following information about applications, disks, file systems, and hosts:

- The total number of assets discovered in that category and the number of assets protected under a policy.

For example, if the **Applications** area of the dashboard displays 15/20 in the **Protected** column, CloudPoint has discovered 20 applications and 15 of them are protected with a policy.

- The snapshot count.

The top right corner of each asset area displays a down arrow. Click the down arrow to expand the asset area and display snapshot information on each cloud vendor or on-premises assets.

To get detailed information on the snapshots for a particular asset type, click the **Snapshot** link, in that part of the dashboard. For example, to see application details, click **Snapshot** at bottom of the **Applications** area.

The bottom of the dashboard displays a summary of snapshot activity across all asset types for the last day, week, and month. It displays the number of successfully created snapshots and the number that have failed.

Displaying information on plug-ins

CloudPoint supports the following types of plug-ins:

- **Off-host plug-ins** run separately from the instance or host where the application runs. Examples of off-host plug-ins are AWS, Azure, and Google plug-ins for cloud environments, and the HDS G-series and HP 3PAR plug-ins for arrays.
- **On-host plug-ins** run on the same instance or host as the application itself. An on-host plug-in discovers the application and its underlying storage. It also plays a key role in taking and restoring snapshots. When you take a snapshot of an application, the on-host plug-in quiesces the application and its under storage stack before the snapshot. It unquiesces them after the snapshot completes. The on-host plug-in also invokes the restore operation. Examples of on-host plug-ins are the Oracle plug-in and Linux file system plug-in.

The **Plugins** area of the CloudPoint dashboard displays the number of off-host and on-host plug-ins configured in your environment as well as the total number available.

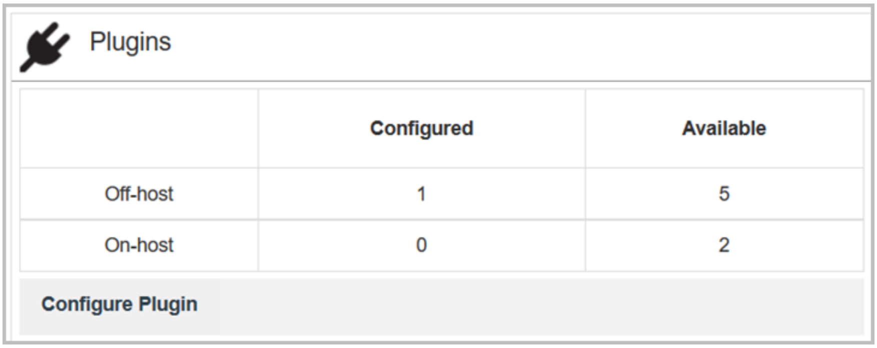

	Configured	Available
Off-host	1	5
On-host	0	2

To display a list of your plug-ins and to configure them, click **Configure Plugin** at the bottom of the **Plugins** area.

See “Configuring an off-host plug-in” on page 52.

See “Configuring an on-host plug-in” on page 55.

Getting quick information from user interface icons

The top of every CloudPoint page includes the following icons. Click an icon to display a screen with status or important information on CloudPoint operations. After you view a screen, click anywhere outside the screen to close it.

Table 2-1 CloudPoint icons

Click this icon ...	To display ...
🕒	Displays recent CloudPoint activity, including creating, restoring, and deleting snapshots.
?	CloudPoint online Help. The online Help displays information on CloudPoint deployment and administration.
👤	Information related to your CloudPoint account, including the following: ■ Admin user name ■ Role ■ User group (if any) ■ Domain Use this screen to change your CloudPoint password and to sign out of CloudPoint

Chapter 3

Protecting your assets with policies

This chapter includes the following topics:

About policies

A policy lets you define and automate how you protect your snapshot data. You can then assign the policy to your assets to ensure their regular and consistent protection.

- A policy includes information such as the following:
  - The storage level of the snapshot; that is whether it is a disk snapshot or a host snapshot
  - The snapshot type; for example, whether it is a clone or a copy on write
  - How many snapshots are retained before earlier ones are deleted
  - How often the snapshot is taken

- How long the snapshot schedule recurs

You can assign more than one policy to an asset. For example, you can create a policy that snapshots assets weekly, and another than snapshots assets daily. You can associate an asset with both of them.

See “Creating a policy” on page 23.

See “Assigning a policy to an asset” on page 24.

See “Listing policies” on page 25.

Creating a policy

To create a policy

1. On the left side of any CloudPoint screen, click **Policies**.
2. On the **Policies** screen, click **Create Policy**.
3. The **Create Policy** page is displayed.

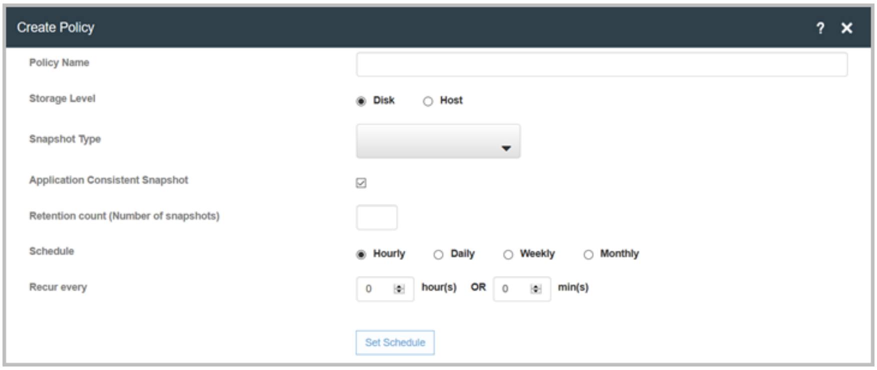

Policy Name	
Storage Level	● Disk ○ Host
Snapshot Type	▾
Application Consistent Snapshot	☑
Retention count (Number of snapshots)	
Schedule	● Hourly ○ Daily ○ Weekly ○ Monthly
Recur every	0 hour(s) OR 0 min(s)

On the **Create Policy** page, enter the following:

Policy name	The policy name
Storage Level	The storage level that you snapshot (either disk or host)
Snapshot Type	The snapshot type (either a clone or copy on write)
Application Consistent Snapshot	Whether you take an application consistent snapshot or a crash-consistent snapshot. An application-consistent snapshot is recommended for taking snapshots of database applications.

Retention

How many snapshot versions to keep for each asset associated with this policy.

Note: An asset may have more total snapshots than the number specified here. If an asset is associated with multiple policies, it has snapshots with each policy. Also, the snapshots you create manually do not count toward the retention total. Manual snapshots are not automatically deleted.

Schedule

Select how often a snapshot is taken: hourly, daily, weekly, or monthly.

Recur every

Based on what you specified for the schedule, determine how long you want the schedule recur.

- For an hourly schedule, specify the number of minutes or hours.
- For a daily schedule, specify the time of day and click **Set Start Time**.
- For a weekly schedule, select the time of day, the day of the week, and click **Set Start Time**.
- For a monthly schedule, select the time of day and the day of the month the snapshot occurs. To schedule snapshots on multiple days, click **Add Day** and specify another day. C lick **Set Start Time**.

When you complete these fields, CloudPoint updates the page to indicate when the policy runs.

Click **Set Schedule**.

4 Click **Next**.

5 On the submission page click **Finish**.

6 Verify that the policy has been created. Note the new entry on the **Policies** page.

Assigning a policy to an asset

After you create a policy, you assign it to one or more assets. For example, you can create a policy to create weekly snapshots and assign the policy to all your database applications. Also, an asset can have more than one policy. For example, in addition to weekly snapshots, you can assign a second policy to your database applications to snapshot them once a month.

When you complete the steps in this section, keep in mind the following:

- The steps for assigning a policy are the same regardless of the type of asset you assign it to.
- Also use these steps when you want to change the policy that is associated with an asset.

To assign a policy to an asset

1. Determine the type of asset you want to project.
2. On the left side of any CloudPoint page, click an asset name to display a list of those assets (**Applications**, **Hosts**, **File Systems**, or **Disks**).
3. In the list of assets, find the one you want to protect, click the vertical ellipsis in the table row, and select **Manage Policy**.
4. On the **Manage Policies** screen assign one or more policies to the asset. In the **Available Policies** column, click the policy you want to assign. The **Policy Details** area of the screen summarizes the policy. Click the greater than sign (**>**) to move it to the **Selected Policies** column. Repeat this step for as many policies as you want to add.

   (Optional) To remove a policy from an asset, select it in the **Selected Policies** column, and click the less than (**<**) sign.
5. When you are done assigning policies, click **Next**.

Listing policies

To list policies

◆ On the left side of any CloudPoint screen, click **Policies**.

The Policies page displays the following information for each policy:

- Policy ID
- Name
- How many snapshots are retained for each asset associated with the policy
- Snapshot type
- Protection level

From this screen, you can do the following:

- Create a new policy
- Display information about a specific policy

See “About policies” on page 22.

See “Creating a policy” on page 23.

See “Displaying policy details” on page 26.

Displaying policy details

To display policy details

1. On the left side of any CloudPoint screen, click **Policies**.
2. On the **Policies** page, click the link of the policy whose details you want to view.

   The **Policies** page updates to display the following:

   - The **Policy Details** area displays the following:
     - Policy ID
     - Name
     - Snapshot schedule
     - How many snapshots are retained for each asset
     - The protection level

   The lower part of the screen lists the assets (applications, hosts, file systems, or disks) associated with the policy. This information includes:

   - Asset name
   - Asset type
   - Vendor logo

See “About policies” on page 22.

See “Creating a policy” on page 23.

Modifying a policy

To modify a policy

1. On the left side of any CloudPoint screen, click **Policies**.
2. On the **Policies** page, locate the policy you want to delete, click the vertical ellipsis and select **Modify**.
3. On the **Create Policy** screen, modify the policy values. The steps for doing this are the same as creating a new policy.

See “About policies” on page 22.

See “Creating a policy” on page 23.

Deleting a policy

To delete a policy

1. On the left side of any CloudPoint screen, click **Policies**.
2. On the **Policies** page, locate the policy you want to delete, click the vertical ellipsis and select **Delete**.
3. On the **Delete Policy** screen, click **Next**.

See “About policies” on page 22.

See “Creating a policy” on page 23.

Chapter 4

Working with applications

This chapter includes the following topics:

Displaying application details

To display application details

1. On the left side of any CloudPoint page, click **Applications**.
2. On the **Applications** page, find the application you are interested in and click its link.
3. Review the **Application Details** page. The top of the page displays the following:
   - Name
   - Region
   - Hosts
   - Snapshot count
   - Plug-in that is used to discover information
   - Disks the application uses
   - Vendor
   - Snapshot policies that are assigned to the application

The bottom of the page displays the following details on each snapshot:

- Name
- Type
- Creation time

On the **Application Details** page, you can also restore or delete snapshots.

Note: If the **Applications** page indicates that an application cannot be snapshot, if may be because the on-host plug-in is configured in a different zone (or region) than the corresponding agent. Both the plug-in and the agent must be configured in the same zone (or region).

See “Configuring an on-host plug-in” on page 55.

See “Configuring an off-host plug-in” on page 52.

See “Restoring an application snapshot” on page 31.

See “Deleting an application snapshot” on page 33.

Taking an application snapshot manually

You can schedule regular snapshots using a CloudPoint policy. However, you can also take snapshots on demand from the CloudPoint or using a RESTful API.

To take an application snapshot manually

1. On the left side of any CloudPoint page, click **Applications**.
2. On the **Applications** page, locate the application you want to snapshot in the table. On the right side of the application's table row, click the vertical ellipsis and select **Create Snapshot**.

3 On the **Create Snapshot** page, specify the following:

Snapshot name

A name 32 characters or less.

Note: In Google Cloud, the snapshot name can only contain letters, numbers, and hyphens. The name should begin and end with a letter.

Description

This field is optional. Enter text here to help you remember the details of the snapshot.

Protection level

Disk - takes a snapshot of the disks the application uses.

Host - takes a snapshot of all the disks associated with the instance.

Snapshot type

Clone - creates a copy of the instance.

Application-consistent snapshot

If you specify an application-consistent snapshot, the snapshot includes the data in memory as well as I/O transactions in process. Because the snapshot is application aware, when you restore the snapshot, you do not need to restore the database application.

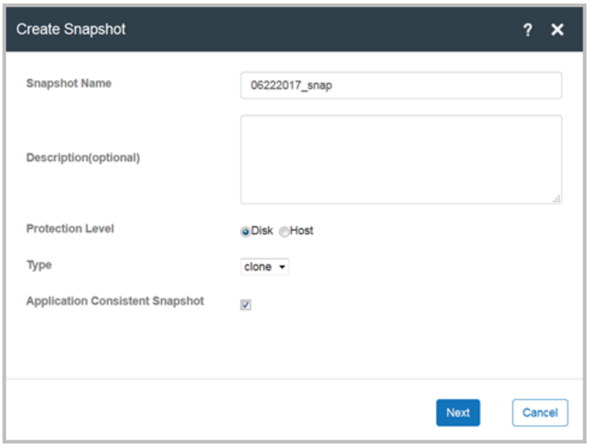

Snapshot Name: 06222017_snap

Description(optional):

Protection Level: ● Disk ○ Host

Type: clone ▾

Application Consistent Snapshot: ☑

Click **Next**.

4 On the snapshot submission page, click **Finish**.

5 Return to the **Applications** page, and click on the application's link. The **Application Details** page is displayed, and the new snapshot appears in the list.

Restoring an application snapshot

To restore an application snapshot manually

1. On the left side of any CloudPoint page, click **Applications**.
2. On the **Applications** page, locate the application whose snapshot you want to restore, and click its link.
3. On the **Application Details** page, locate the snapshot, click the vertical ellipsis and select **Restore**.

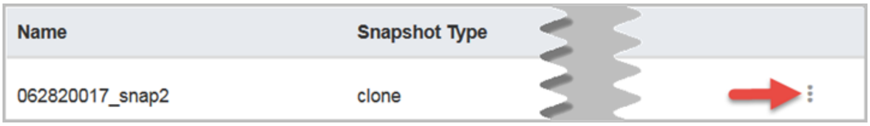

Name	Snapshot Type
062820017_snap2	clone

4 On the **Restore Snapshot** screen, specify how you want to restore the snapshot.

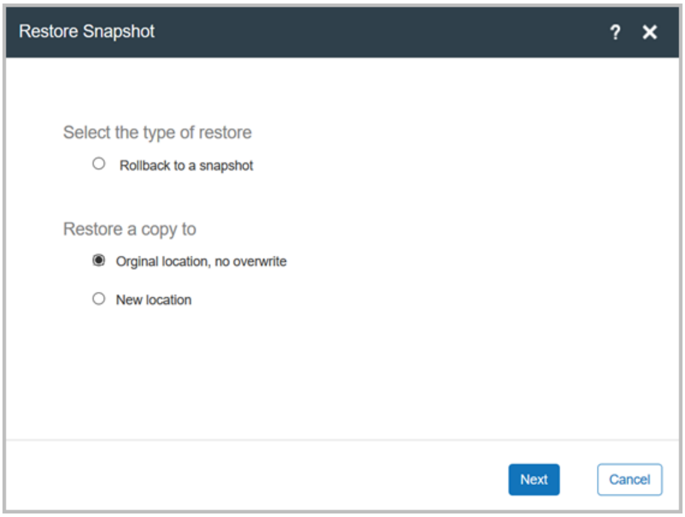

Restore Snapshot

Select the type of restore

○ Rollback to a snapshot

Restore a copy to

● Orginal location, no overwrite

○ New location

Rollback to a snapshot	**Note:** Currently, this option is only supported for array snapshots.
Original location, no overwrite	Restore the snapshot to the original location, but the current data is also preserved.
New location	If you select this option, CloudPoint displays a drop-list of available hosts.

5 Click **Next**.

After the application snapshot is restored, shutdown and restart the Oracle database instance.

See "Shutting down and restarting the Oracle database after restoring an Oracle application snapshot" on page 32.

Shutting down and restarting the Oracle database after restoring an Oracle application snapshot

After you restore an Oracle application snapshot, there are several follow-up steps you must perform on the Oracle database.

To shut down and restart the Oracle database after restoring an Oracle application snapshot

1. Log in to the database instance.
2. Enter the following:

   ```
   SQL> alter database end backup
   ```

3. Bring down the database. Enter the following:

   ```
   SQL> shutdown
   ```

4. Bring up the database with a new metafile and data files. Enter the following:

   ```
   SQL> startup
   ```

Give the database some time to read the new data.

If the database does not come up, contact your database administrator to investigate the problem.

Deleting an application snapshot

To delete an application snapshot

1. On the left side of any CloudPoint page, click **Applications**.
2. On the **Applications** page, click the application whose snapshot you want to delete.
3. On the **Application Details** page, locate the snapshot, click the vertical ellipsis and select **Delete**.

   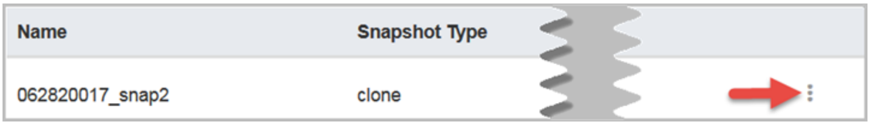

Name	Snapshot Type
062820017_snap2	clone

4. On the **Delete Snapshot** confirmation screen, click **Next**.
5. On the **Delete Snapshot** submission screen, click **OK**.
6. Return to the **Application Details** page and verify that the snapshot has been removed.

Chapter 5

Working with hosts

This chapter includes the following topics:

Displaying host details

To display host details

1. On the left side of any CloudPoint page, click **Hosts**.
2. On the **Hosts** page, find the host you are interested in and click its link.
3. Review the **Host Details** page. The top of the page displays the following:
    - Name
    - Region
    - Host type, either virtual or physical
    - State
    - Snapshot count
    - Plug-in that is used to discover information
    - Private IP address
    - Number of devices that are connected to the host
    - Vendor
    - Snapshot policies that are assigned to the host

- Public IP address
- Host ID

The bottom of the page lists the snapshots that are associated with the host. The following information is displayed for each snapshot:

- Name
- Type
- Creation time

On the **Host Details** page, a vertical ellipsis appears at the end of each table row. Select the ellipsis to restore or delete the snapshot.

See “Restoring a host snapshot” on page 37.

See “Deleting a host snapshot” on page 39.

Taking a host snapshot manually

Note: Currently, you cannot take a snapshot of a Google Cloud virtual machine; however, you can take a disk snapshot on all clouds that CloudPoint supports.

To take a host snapshot manually

1 On the left side of any CloudPoint page, click **Hosts**.

2 On the **Hosts** page, locate the host you want to snapshot in the table. On the right side of the host's table row, click the vertical ellipsis and select **Create Snapshot**.

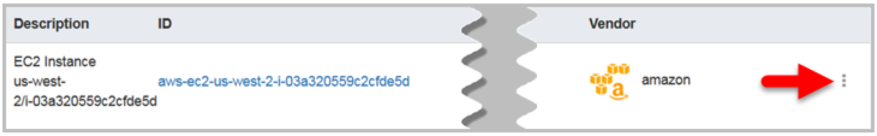

Description	ID	Vendor	
EC2 Instance us-west-2/i-03a320559c2cfde5d	aws-ec2-us-west-2-i-03a320559c2cfde5d	amazon	⋮

3 On the **Create Snapshot** page, specify the following:

Snapshot Name	A name 32 characters or less.
Description	This field is optional. Enter text here to help you remember the details of the snapshot.
Type	Clone - creates a copy of the instance (virtual machine).
Application Consistent Snapshot	If you specify an application-consistent snapshot, the snapshot includes the data in memory as well as I/O transactions in process. Because the snapshot is application aware, when you restore the snapshot, you do not need to restore the database application.

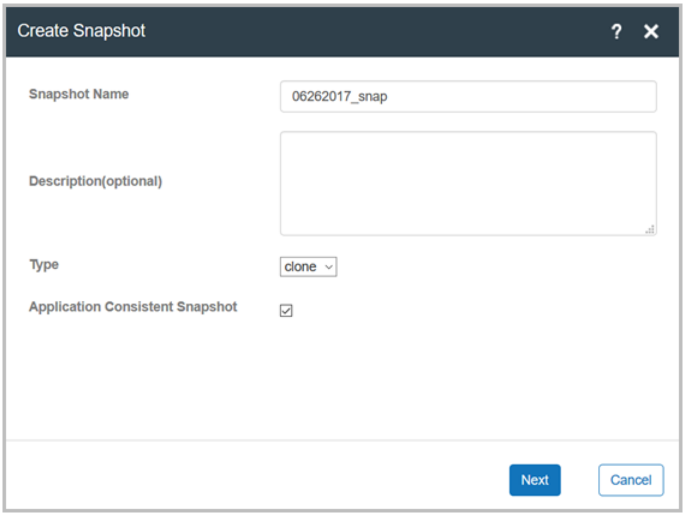

Create Snapshot

Snapshot Name	06262017_snap
Description(optional)	
Type	clone
Application Consistent Snapshot	☑

Click **Next**.

4 On the snapshot submission page, click **Finish**.

5 Verify that the snapshot has been created. Return to the **Hosts** page, and click on the host's link. The **Host Details** page is displayed, and the new snapshot appears in the list.

Restoring a host snapshot

To restore a host snapshot manually

1 On the left side of any CloudPoint page, click **Hosts**

2 On the **Host** page, click on the host whose snapshot you want to restore.

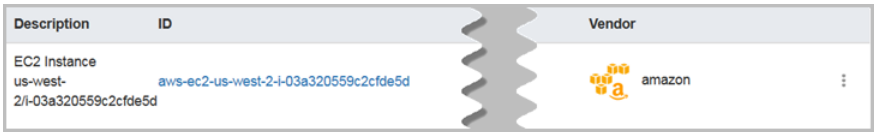

Description	ID	Vendor
EC2 Instance us-west-2/i-03a320559c2cfde5d	aws-ec2-us-west-2-i-03a320559c2cfde5d	amazon

3 On the **Host Details** page, locate the snapshot, click the vertical ellipsis and select **Restore**.

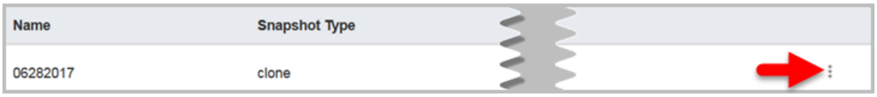

Name	Snapshot Type
06282017	clone

4 On the **Restore Snapshot** screen, specify how you want to restore the snapshot.

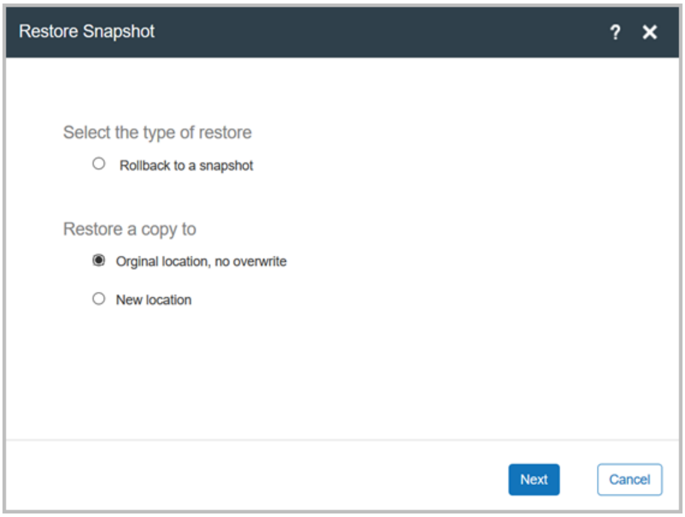

Restore Snapshot

Select the type of restore

○ Rollback to a snapshot

Restore a copy to

● Orginal location, no overwrite

○ New location

Rollback to a snapshot

Note: Currently, this option is only supported for array snapshots.

Original location, no overwrite

Restore the snapshot to the original location, but the current data is also preserved.

Note: This option is not supported for array snapshots.

New location

If you select this option, CloudPoint displays a drop-list of available hosts.

Note: This option is not supported for array snapshots.

5 Click **Next**.

Deleting a host snapshot

To delete a host snapshot manually

1. On the left side of any CloudPoint page, click **Hosts**.
2. On the **Hosts** page, click the host whose snapshot you want to delete.

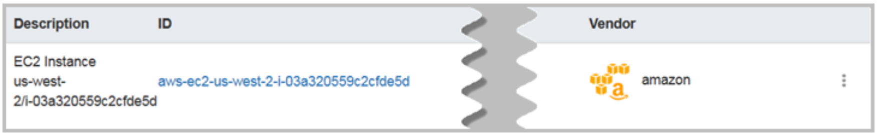

Description	ID	Vendor
EC2 Instance us-west-2/i-03a320559c2cfde5d	aws-ec2-us-west-2-i-03a320559c2cfde5d	amazon

3. On the **Hosts Details** page, locate the snapshot, click the vertical ellipsis and select **Delete**

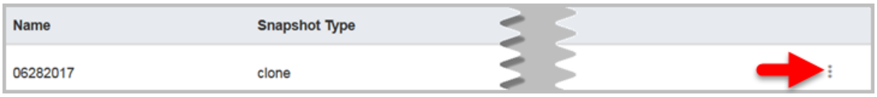

Name	Snapshot Type
06282017	clone

4. On the **Delete Snapshot** confirmation screen, click **Next**.
5. On the **Delete Snapshot** submission screen, click **OK**.
6. Return to the **Host Details** page and verify that the snapshot has been removed.

Chapter 6

Working with file systems

This chapter includes the following topics:

Displaying file system details

To display file system details

1. On the left side of any CloudPoint page, click **File Systems**.
2. On the **File Systems** page, find the file system you are interested in and click its link.
3. Review the **File System Details** page. The top of the page displays the following:
    - Type
    - Path
    - Snapshot policies that are assigned to the file system
    - Snapshot count
    - Plug-in that is used to discover information
    - Device
    - File system ID

    The bottom of the page displays the following details on each snapshot:

- Name
- Type
- Creation time

On the **File System Details** page, you can also restore or delete snapshots.

See “Restoring a file system snapshot” on page 42.

See “Deleting a file system snapshot” on page 43.

Taking a file system snapshot manually

To take a file system snapshot manually

1. On the left side of any CloudPoint page, click **File Systems**.
2. On the **File Systems** page, locate the file system you want to snapshot in the table. On the right side of the file system's table row, click the vertical ellipsis and select **Create Snapshot**.
3. On the **Create Snapshot** page, specify the following:

Snapshot name	A name 32 characters or less. **Note:** If you take a snapshot of a Google Cloud asset, the name can include a hyphen, but no other special characters are allowed.
Description	This field is optional. Enter text here to help you remember the details of the snapshot.
Snapshot type	Disk - takes a snapshot of the disks the instance uses. Host - takes a snapshot of the entire instance.
Application consistent snapshot	If you specify an application-consistent snapshot, the snapshot includes the data in memory as well as I/O transactions in process. Because the snapshot is application aware, when you restore the snapshot, you do not need to restore the database application.

   Click **Next**.

4. On the snapshot submission page, click **Finish**
5. Return to the **File Systems** page, and click on the file system's link. The **File System Details** page is displayed, and the new snapshot appears in the list.

Restoring a file system snapshot

To restore a file system snapshot

1. On the left side of any CloudPoint page, click **File Systems**.
2. On the **File Systems** page, click on the file system whose snapshot you want to restore.
3. On the **File System Details** page, locate the snapshot, click the vertical ellipsis and select **Restore**.
4. On the **Restore Snapshot** screen, specify how you want to restore the snapshot.

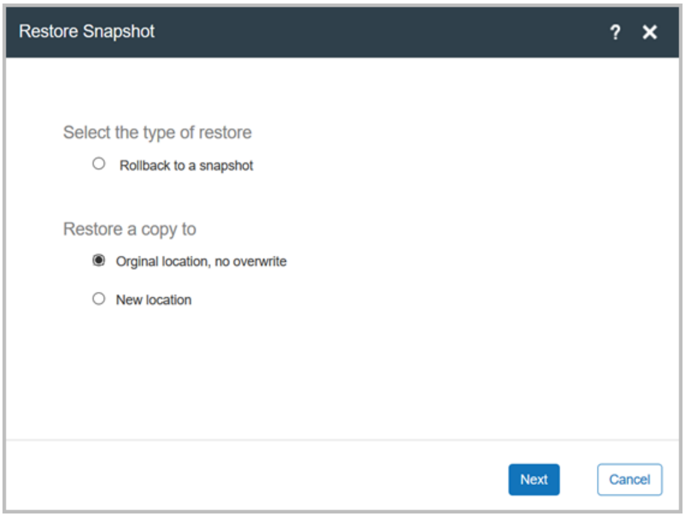

Restore Snapshot

Select the type of restore

○ Rollback to a snapshot

Restore a copy to

● Orginal location, no overwrite

○ New location

Rollback to a snapshot	**Note:** Currently, this option is not supported.
Original location, no overwrite	Restore the snapshot to the original location, but the current data is also preserved.
New location	If you select this option, CloudPoint displays a drop-list of available hosts.

5. Click **Next**.

Deleting a file system snapshot

To restore a file system snapshot

1. On the left side of any CloudPoint page, click **File Systems**.
2. On the **File Systems** page, click on the file system whose snapshot you want to delete.
3. On the **File System Details** page, locate the snapshot, click the vertical ellipsis and select **Delete**.
4. On the **Delete Snapshot** confirmation screen, click **Next**.
5. On the **Delete Snapshot** submission screen, click **OK**.
6. Return to the **File System Details** page and verify that the snapshot has been removed.

Working with disks

This chapter includes the following topics:

Displaying disk details

To display disk details

1. On the left side of any CloudPoint page, click **Disks**.
2. On the next page, find the disk you are interested in and click its link.
3. Review the **Volume Details** page. The top of the page displays the following:
   - Name
   - Region
   - Size
   - Snapshot count
   - Plug-in that is used to discover information
   - DevPath
   - Vendor
   - Snapshot policies that are assigned to the disk
   - Disk ID

The bottom of the page lists the snapshots that are associated with the disk. The following information is displayed for each snapshot:

- Name
- Type
- Creation time

On the **Volume Details** page, a vertical ellipsis appears at the end of each table row. Select the ellipsis to restore or delete the snapshot.

See "Restoring a disk snapshot" on page 47.

See "Deleting a disk snapshot" on page 48.

Taking a disk snapshot manually

To take a disk snapshot manually

1. On the left side of any CloudPoint page, click **Disks**.
2. On the next page, locate the volume you want to snapshot in the table. On the right side of the volume's table row, click the vertical ellipsis and select **Create Snapshot**.

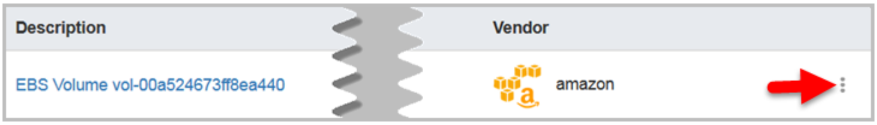

Description	Vendor
EBS Volume vol-00a524673ff8ea440	amazon

3 On the **Create Snapshot** page, specify the following:

Snapshot name

A name 32 characters or less.

Note: If you are taking a snapshot of a Google Cloud asset, the name can include a hyphen, but no other special characters are allowed.

Description

This field is optional. Enter text here to help you remember the details of the snapshot.

Snapshot type

Clone/Copy on write (COW) - creates a copy of the disk.

Application consistent snapshot

If you specify an application-consistent snapshot, the snapshot includes the data in memory as well as I/O transactions in process. Because the snapshot is application aware, when you restore the snapshot, you do not need to restore the database application.

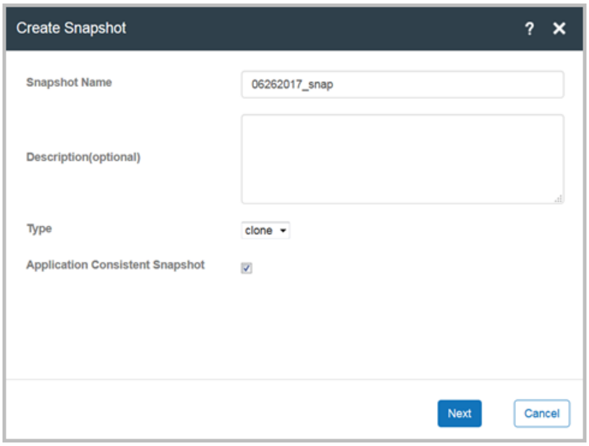

Create Snapshot

Snapshot Name: 06262017_snap

Description(optional):

Type: clone

Application Consistent Snapshot: ☑

4 Click **Next**.

5 On the snapshot submission page, click **Finish**.

6 Return to the **Volume Details** page and verify that the new snapshot appears in the list.

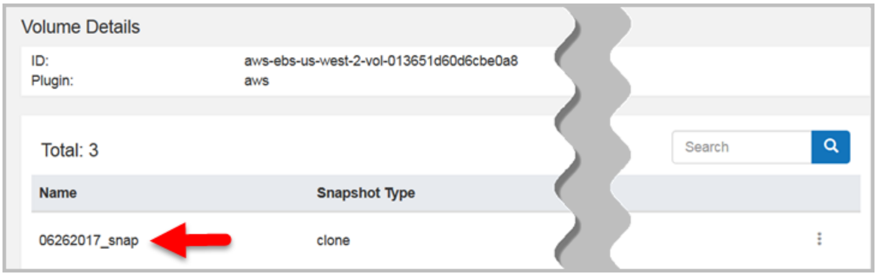

Volume Details

ID: aws-ebs-us-west-2-vol-013651d60d6cbe0a8
Plugin: aws

Total: 3

Name	Snapshot Type
06262017_snap	clone

Restoring a disk snapshot

Before you restore a disk snapshot, verify the following:

- The target instance must not be running.
- If you restore a disk snapshot to an instance that has a volume with the same UUID as the snapshot, the instance does not boot. Before you restore the snapshot, make sure it has a unique UUID.

To restore a disk snapshot manually

1 On the left side of any CloudPoint page, click **Disk**.

2 Locate the disk whose snapshot you want to restore, and click its link.

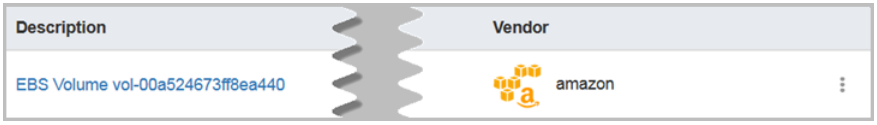

Description	Vendor
EBS Volume vol-00a524673ff8ea440	amazon

3 On the **Volume Details** page, locate the snapshot, click the vertical ellipsis and select **Restore**.

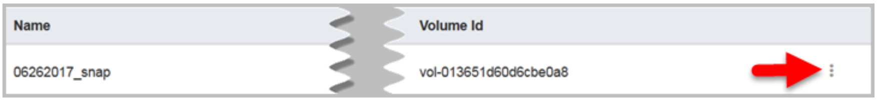

Name	Volume Id
06262017_snap	vol-013651d60d6cbe0a8

4 On the **Restore Snapshot** screen, specify how you want to restore the snapshot.

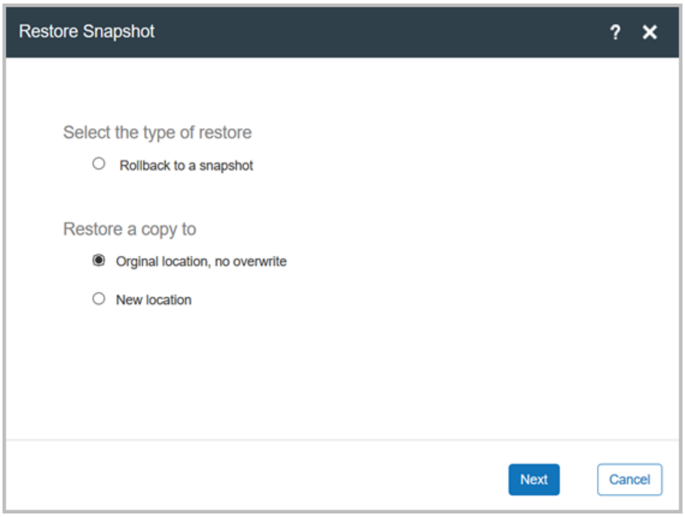

Select the type of restore

○ Rollback to a snapshot

Restore a copy to

● Orginal location, no overwrite

○ New location

Rollback to a snapshot

Note: Currently, this option is not supported.

Original location, no overwrite

Restore the snapshot to the original location, but the current data is also preserved.

New location

If you select this option, CloudPoint displays a drop-list of available hosts.

5 Click **Next**.

Deleting a disk snapshot

To delete a disk snapshot manually

1 On the left side of any CloudPoint page, click **Disk**

2 On the next page, click the disk whose snapshot you want to delete.

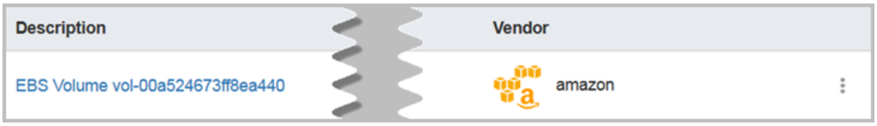

Description	Vendor
EBS Volume vol-00a524673ff8ea440	amazon

3 On the **Volume Details** page, locate the snapshot, click the vertical ellipsis and select **Delete**.

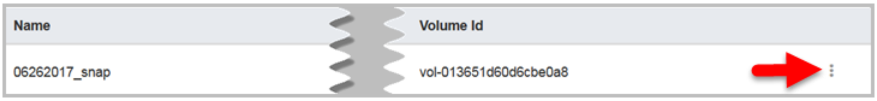

Name	Volume Id	
06262017_snap	vol-013651d60d6cbe0a8	⋮

4 On the **Delete Snapshot** confirmation screen, click **Next**.

5 On the **Delete Snapshot** submission screen, click **OK**.

6 Return to the **Volume Details** page and verify that the snapshot has been removed.

Chapter 8

Working with plugins

This chapter includes the following topics:

About plug-ins

A CloudPoint plug-in is a low-level module that discovers assets in your environment and performs operations on them. Plug-ins are implemented as Python modules.

A plug-in only operates on a particular type of data source. For example, there is an AWS plug-in, a Hitachi Data Systems (HDS) G-Series array plug-in, and so on.

You can run multiple instances of a plug-in to gather information from multiple sources within a particular type of source. For example, you can deploy a separate AWS plug-in for each AWS account.

You can also run multiple instances of a plug-in for the same data source but in separate processes or hosts for load-balancing or high availability purposes.

There are two types of plug-ins.

- **On-host plug-ins** run on the same instance or host as the application itself. An on-host plug-in discovers the application and its underlying storage. It also plays a key role in taking and restoring snapshots. When you take a snapshot of an application, the on-host plug-in quiesces the application and its under storage stack before the snapshot. It unquiesces them after the snapshot completes.

The on-host plug-in also invokes the restore operation. Examples of on-host plug-ins are the Oracle plug-in and Linux file system plug-in.

- **Off-host plug-ins** run separately from the instance or host where the application runs. Examples of off-host plug-ins are AWS, Azure, and Google plug-ins for cloud environments, and the HDS G-series and HP 3PAR plug-ins for arrays.

Each plug-in is wrapped in an agent.

See “About agents” on page 60.

See “Determining the types of plug-ins and agents to install” on page 51.

Determining the types of plug-ins and agents to install

When do you need to install off-host plug-ins and on-host agents and plug-ins?

You need to install off-host plug-ins to discover the virtual machines, hosts, and disks and to manage their protection. After you install and configure off-host plug-ins, you can take crash-consistent snapshots of the virtual machines and disks that the plug-ins manage. The virtual machines can run any operating system. You do not have to install on-host agents or plug-ins to take crash-consistent snapshots.

However, to discover applications and file systems and protect them with application-consistent snapshots, you must install an on-host agent and one or more on-host plug-ins. (The snapshots can be at the host of disk level.)

The CloudPoint user interface displays the plug-ins available for deployment. The CloudPoint container includes the following plug-ins:

- Off-host plug-ins:
  - Amazon AWS
  - Microsoft Azure
  - Google Cloud
  - Hitachi Data Systems (HDS) G-Series array
  - HP 3PAR array
- On-host plug-ins:
  - Oracle
  - Linux file systems ext2, ext3, ext4, and XFS

See “Configuring an off-host plug-in” on page 52.

See “Configuring an on-host plug-in” on page 55.

Configuring an off-host plug-in

At a minimum, you must configure off-host plug-ins to create crash-consistent snapshots of your assets. However, If you want to create application-consistent snapshots of your assets, you must also configure the appropriate on-host plug-ins.

To complete the steps in this section, make sure that you have the required parameters for the plug-in that you want to configure.

To configure off-host plug-ins

1. On the right side of any CloudPoint screen, click **Plug-ins**.
2. Make sure that you understand the configuration fields that are required for the plug-in and that you have this information ready.

   See "Amazon Web Services plug-in configuration notes" on page 52.

   See "Google Cloud plug-in configuration notes" on page 53.

   See "Hewlett-Packard Enterprise 3PAR plug-in configuration notes" on page 54.

   See "Hitachi Data Systems plug-in configuration notes" on page 54.

   See "Microsoft Azure plug-in configuration notes" on page 53.
3. On the **Plug-ins** page, click the link of the plug-in you are interested in. CloudPoint displays a detailed description of the plug-in, including the following:
4. Perform the preconfiguration steps that are listed in the plug-in description.
5. On the CloudPoint interface, select **Configure** to configure the plug-in.
6. On the **Configure Plug-in** screen, enter the required parameters for the specific plug-in. For a cloud plug-in, use the drop-down list to select the cloud region. Click **Next**.
7. On the **Configure Plug-in** submission screen, click **OK**.
8. Verity that the plug-in is configured. Navigate to the **Hosts** page and verify that CloudPoint has discovered all the hosts in the region.

See "Configuring an on-host plug-in" on page 55.

Amazon Web Services plug-in configuration notes

Amazon Web Services (AWS) plug-in lets you take snapshots of Elastic Compute Cloud (EC2) instances and Elastic Block Store (EBS) volumes in an Amazon cloud.

To configure the plug-in, you must specify the following information:

- Access key
- Secret key

- Cloud region

This plug-in lets you create, delete, and restore snapshots at the instance level and the disk level.

See “Configuring an off-host plug-in” on page 52.

Google Cloud plug-in configuration notes

This plug-in lets you create, delete, and restore disk snapshots in all zones where Google Cloud is present. In a future release, you will be able to perform instance snapshots as well.

To configure the plug-in, you need the following information:

projectId	Client ID that used for operations
clientId	ID of the project from which the resources are managed
clientEmail	Email of the clientId mentioned above
privateKeyId	ID of the private key
privateKey	Actual private key **Note:** You must enter this key without quotes (neither single quotes nor double quotes). Do not enter any spaces or return characters at the beginning or end of the key.
zones	List of zones in which the plug-in operates

You also need the`google-api-python-client` Python library.

See “Configuring an off-host plug-in” on page 52.

Microsoft Azure plug-in configuration notes

The Microsoft Azure plug-in lets you take snapshot of virtual machines and managed disks in an Azure cloud.

To configure the plug-in, you must specify the following information:

- Client ID
- Secret ID
- Tenant ID

This plug-in lets you create, delete, and restore snapshots at the virtual machine level and the managed disk level.

To have the Azure plug-in scan and manage all the resources in Azure, do the following:

- Use the portal to create an Azure Active Directory application for the Azure plug-in
- Assign service principal to a role to access resources.

For more details, follow the steps in the following Azure documentation:

https://docs.microsoft.com/en-us/azure/azure-resource-manager/resource-group-create-service-principal-portal

The Azure plug-in has the following limitations:

- The current release of the plug-in does not support snapshots of blobs.
- CloudPoint currently supports creating and restoring snapshots of Azure-managed disks and virtual machines that are backed by managed disks.

See "Configuring an off-host plug-in" on page 52.

Hewlett-Packard Enterprise 3PAR plug-in configuration notes

This plug-in enables you to handle snapshots of disks on 3PAR Array.

To configure the plug-in, you must specify the following information:

- Array IP address
- Username
- Password

The `python-3parclient` library must be installed.

This plug-in lets you create and delete snapshots of the clone and copy-on-write (COW) types. You can restore a COW snapshot, but not a clone snapshot.

See "Configuring an off-host plug-in" on page 52.

Hitachi Data Systems plug-in configuration notes

This plug-in enables you to take disk snapshots of a Hitachi Data Systems (HDS) array.

To configure this plug-in, you must specify the following information. All the parameters are required.

CloudPoint configuration parameter	Description

Baseurl	The base URL for accessing the HDS REST API Server. The URL has the following format: `protocol://host-name:port-number/ConfigurationManager`
IP	The IP address of the HDS REST API Server
Username	The user name of the HDS REST API Server
Password	The password to access the HDS REST API Server

This plug-in lets you create disk snapshots of the copy-on-write (COW) type. You can also delete and restore snapshots.

See “Configuring an off-host plug-in” on page 52.

Configuring an on-host plug-in

Before you perform the steps in this section, do the following:

- Download and install an on-host agent.
  See “Downloading and installing an agent” on page 62.
- Install any prerequisites for the plug-in and meet any vendor-specific requirements.
  See “Preparing to install the Oracle CloudPoint plug-in” on page 56.
- Make sure that you understand the configuration fields that are required for the plug-in and that you have this information ready.
  See “Oracle plug-in configuration notes” on page 58.

  Note: The Linux plug-in does not require any configuration parameters.

Note: To take a snapshot of an Oracle instance, you must configure the plug-in in the same zone (or region) as the corresponding on-host agent. For example, when you configure the on-host AWS/Google plug-in it must be in the same zone as the on-host AWS/Google agent. If the plug-in is configured in a different region, the **Applications** page on the user interface displays that the applications cannot be snapshot.

To configure an on-host plug-in

1. On the left side of any CloudPoint page, click **Agents**.
2. On the **Agents** page, locate an on-host agent, select the vertical ellipsis, and click **Configure**.
3. From the drop-down list, select the type of plug-in you want to configure.
4. Specify the plug-in specific configuration values.
5. Click **Submit**.

   After a period of time, the plug-in discovers the relevant assets, and the CloudPoint user interface is updated.

Preparing to install the Oracle CloudPoint plug-in

To complete the steps in this section, you need root privileges.

To prepare your environment for the Oracle plug-in, you must install the following:

- Required RPMs
- Python cx_Oracle library

To install the required RPMs

1. Make sure you have an Oracle single sign-on.
2. Download the following RPMs from the Oracle website.
   - oracle-instantclient12.2-basic x86_64 RPM
   - oracle-instantclient12.2-devel x86_64 RPM

   The RPMs should be for the Linux x86_64 architecture and be version 12.2 or higher.
3. Install the RPMs.

   ```
   # rpm -ivh oracle-instantclient12.2-basic.*.*
   # rpm -ivh oracle-instantclient12.2-devel.*.*
   ```

To install the Python cx_Oracle library

1. Make sure that PIP is installed on this RHEL machine. PIP is a package management system that installs and manages packages written in Python. You can verify this by entering `pip` on the command line.

   If you do not have PIP installed, enter the following commands:

   ```
   # yum install python-pip -y
   # pip install --upgrade pip
   ```

2. Install the library.

   ```
   # pip install cx_Oracle
   ```

   Note: If the installation fails, do not continue.

In addition, to installing these components, do the following:

- Configure the agent.
  See "Configuring an agent" on page 62.
- Optimize your Oracle database data and meta files.
  See "Optimizing your Oracle database data and metadata files" on page 58.

Creating an Oracle database user

When you configure the CloudPoint Oracle-plug-in, you must specify the name and non-blank password of the Oracle database user for the instance. If a user does not exist, create one using steps in this section.

To complete the steps in this section, you must have root privileges.

To create an Oracle database user

1. Log in to the database instance.
2. At the SQL prompt, run the following queries to create a user with a password. The password cannot be NULL.

   ```
   SQL> create user user-name identified by password;
   SQL> grant dba to user-name;
   ```

3. Verify that the user was created properly. Enter the following:

   ```
   SQL> connect user-name/password as sysdba
   ```

   If any problems occur, contact your Oracle database administrator.

4. Verify that the root user can logon to the Oracle user without a password. Make sure that you are logged on as the root user and enter the following:

   ```
   # su - oracle-user
   ```

   If you were able to logon to the Oracle user, the command line appears as follows:

   ```
   oracle-user#
   ```

   If you do not see this result, make *oracle-user* passwordless for the root user.

Optimizing your Oracle database data and metadata files

CloudPoint takes disk snapshots. For better backup and recovery, you should optimize your Oracle database data and metadata files.

Each Oracle database instance has a control file. The control file contains information about managing rather operating the database for each transaction. For faster and efficient backup and recovery, Oracle recommends that you put the control file in the same file system as the database redo log file. If the database control file resides on the file system which is created on top of the boot disk or root disk, contact your database administrator to move the control file to the appropriate location. For more information on control files and how to move them, contact your database administrator, or see the Oracle documentation.

https://docs.oracle.com/cd/B10500_01/server.920/a96521/control.htm#3545

After you use a snapshot to restore an application, do not perform any operations. Allow some time for Oracle to read new data and bring up the database. If the database does not come up, contact the database administrator to determine the cause of the problem.

Oracle plug-in configuration notes

When you configure the Oracle plug-in, have the following information ready. The plug-in needs this information to operate on Oracle databases. You can gather many of these values from the file `/etc/oratab`.

- Oracle user name
- Instance name

- Instance name home path
- Database user name

  Note: If a database user is not present, create one and provide a non-blank password.

  See “Creating an Oracle database user” on page 57.

- Database password
  The password cannot be NULL. The Oracle Python Library `cx_Oracle` does not support a NULL password.

After you assemble this information, you can configure the Oracle plug-in.

See “Configuring an on-host plug-in” on page 55.

When you configure the Oracle plug-in, also keep in mind the following:

- Each database you configure should be up and running. Verify that the `pmon` process is running for the database instance.
- Each database instance you add should have a specified user name and password. Outside of the scope of CloudPoint, verify that the user name and password work.

Listing plug-ins

To list plug-ins

◆ On the left side of any CloudPoint screen, click **Plug-ins**.

The **Plug-ins** page displays the following information for each plug-in:

- Name
- Version
- Vendor or file system that is associated with the plug-in
- Plug-in type. An on-host plug-in is installed on the instance (or physical host) as the application. An off-host plug-in is not.

From this page, you can select and configure a plug-in.

Chapter 9

Working with agents

This chapter includes the following topics:

About agents

Agents are wrappers around the CloudPoint plug-ins and perform several functions common to most or all plug-ins. Agents do the following:

- Translate between the message protocol and the plug-in interface.
- Ensure secure communication between the plug-ins and rest of the CloudPoint components.
- Provide a common implementation of certain tasks such as polling for asset changes (if the plug-in does not support pushing updates).
- Handle authentication.

There are two types of agents: on-host agents and off-host agents. An on-host agent must be installed and configured on a host where an application is running. The on-host agent manages one or more on-host plug-ins. You need on-host agents and on-host plug-ins to take snapshots of an Oracle application or a Linux file system.

In contrast, off-host agents and off-host plug-ins do not need a separate host on which to run. You use off-host agents and off-host plug-ins to take snapshots of public cloud assets and on-premises storage arrays.

On the CloudPoint console, the **Agents** page lists the hosts on which agents are running. The **Type** column indicates whether the agent is an on-host agent or off-host agent.

See "About plug-ins" on page 50.

The following table shows you the type of agent required for each type of asset snapshot.

Table 9-1 Asset types and the type of agent required

Asset type and vendors	On-host agent required	Off-host agent required
Application ▪ Oracle	x	
File system ▪ Linux	x	
Public cloud (host snapshot or disk snapshot) ▪ Amazon AWS ▪ Google Cloud ▪ Microsoft Azure		x
On-premises storage array ▪ Hewlett-Packard Enterprise (HPE) 3PAR ▪ Hitachi Data Systems (HDS) G-Series		x

Agent dependencies

Before you install an agent, make sure that you install the following dependency.

- `python2-pika` package. Enter the following commands:

```
# yum install \
https://dl.fedoraproject.org/pub/epel/epel-release-latest-7.noarch.rpm -y

# yum install python2-pika -y
```

Downloading and installing an agent

To complete the steps in this section, you need root privileges.

To download and install an agent

1. Make sure you install all agent dependencies.

   See “Agent dependencies” on page 61.

2. On the left side of any CloudPoint page, click **Agents**.
3. On the **Agents** page, click **Download Agent RPM**.
4. From the command line, install the agent RPM. Use the following syntax:

   ```
   # rpm -ivh CloudPoint_agent_RPM_name
   ```

   For example:

   ```
   # rpm -ivh VRTScloudpoint-agent-1.0.1-RHEL7.x86_64.rpm
   ```

After you install the agent, continue with any steps required to prepare the CloudPoint plug-in.

See “Preparing to install the Oracle CloudPoint plug-in” on page 56.

Configuring an agent

To complete the steps in this section, you need root privileges.

To configure an agent

1. On the agent host, create the file `/etc/flexsnap.conf` with the following lines:

   ```
   [global]
   target = CloudPoint public name or IP address
   ```

2. Configure the agent with the token. Do the following:

   - On the left side of any CloudPoint user interface page, click **Agents**.
   - On the **Agents** page, click **Get Token**. You will use this token when you configure the agent on the command line. The token is valid for one minute.
   - Click **Copy Token**.
   - If the agent service has been configured on this host before, remove the `keys` directory. Enter the following command:

     ```
     # rm -rf /opt/VRTScloudpoint/keys
     ```

- Copy the token and start the `flexsnap-agent`. Enter the following:

  ```
  # flexsnap-agent copied_token
  ```

  Note: If you encounter an error, contact Veritas Customer Support.

3 Enable the agent service. Enter the following:

```
# systemctl enable flexsnap-agent
```

4 Run the agent service. Enter the following:

```
# systemctl start flexsnap-agent
```

5 Verify that the agent is configured. On the CloudPoint user interface, click **Agents** and verify that a new agent is running. The **Host** column displays the new agent name, the **Type** column indicates that it is an **On-host** agent, and the **Status** column indicates that it is **online**.

CloudPoint logs

This chapter includes the following topics:

CloudPoint logs

CloudPoint maintains the following logs to monitor activity and troubleshoot issues. The logs are stored on the path `/cloudpoint/logs`. CloudPoint retains multiple versions of each log, with a number appended to the log name; for example, `flexsnap-agent.log.2`.

Table 10-1 CloudPoint logs

Log	Description
`flexsnap-agent.log`	The log for the service that runs one or more plug-ins. These plug-ins discover assets and perform asset management tasks such as creating, restoring, and deleting snapshots.
`flexsnap-api.log`	The log for the service that translates RESTful API requests into JSON-formatted requests. These requests are sent to the coordinator.
`flexsnap-auth.log`	The log for the authentication service. It records authentication requests coming through RabbitMQ when other services connect. Typically, you do not need to examine this file. This log is primarily for support use.
`flexsnap-coordinator.log`	The log for the service that manages a database of assets. The coordinator also routes requests from the PI service to the appropriate agents.

Chapter 11

Upgrading CloudPoint

This chapter includes the following topics:

Upgrading CloudPoint

When you upgrade CloudPoint, all the snapshot data and configuration data from your previous version is maintained in the external `/cloudpoint` volume. This information is external to the container and the image, so it is preserved during the upgrade.

In the following upgrade steps, you replace the container that runs your current version of CloudPoint with a new container.

To upgrade CloudPoint

1 Download the latest CloudPoint image (for example, `Veritas_CloudPoint_1.0.1_IE.img`) from MyVeritas.

If necessary, copy the downloaded image to the machine on which you will deploy CloudPoint.

2 Load the latest CloudPoint image. For example, enter the following:

```
# sudo docker load -i /home/ubuntu/Veritas_CloudPoint_1.0.1_IE.img
```

3 Verify that the latest image is loaded successfully.

```
# sudo docker images
REPOSITORY           TAG       IMAGE ID       CREATED        SIZE
veritas/cloudpoint   1.0.1     0b998143f820   5 hours ago    999.2 MB
veritas/cloudpoint   latest    0b998143f820   5 hours ago    999.2 MB
veritas/cloudpoint   1.0       980dc5de806d   2 weeks ago    999.3 MB
```

4 Stop the current running container. Enter the following:

```
# sudo docker stop CloudPoint
```

5 Remove the old container. Enter the following:

```
# sudo docker rm CloudPoint
```

6 Run the latest container image. Enter the following:

```
# sudo docker run --name CloudPoint -it -d -v /cloudpoint:/cloudpoint \
-p 443:443 -p 5671:5671 veritas/cloudpoint:latest
```

Note: If you encounter a problem running the new container, start CloudPoint from the old image and contact Veritas Customer Support.

7 Verify that the latest container is running. Enter the following:

```
# sudo docker ps
```

8 If the new container is running and you can successfully log in to the CloudPoint web console, you can remove the old CloudPoint image. For example, enter the following:

```
# sudo docker rmi image_ID_of_older_image
```

Chapter 12

API reference

This chapter includes the following topics:

About APIs

Most CloudPointAPI calls return a single asset or a list of assets. An asset has the following fields:

Field	Description
`id`	The asset's unique identifier.
`type`	The asset type; for example, a disk, a file system, or a host.
`plugin`	The name of the plug-in (or plug-ins) that provide information about the asset.
`vendor`	The name of the asset's vendor; for example, Amazon or Microsoft. This field is optional.

Field	Description
`parentId`	A list of asset IDs for the parents of the asset. For example, the parent ID of a file system is the ID of the disk or volume under the file system. For snapshots, this field is optional.
`snapMethods`	A list of snapshot methods that this asset supports. If the asset does not support snapshots, the list is empty.

Snapshot assets also have the following fields:

Field	Description
`snapSourceId`	The ID of the original asset of which this is a snapshot.
`snapType`	The snapshot type; for example, copy-on-write or clone.
`restoreMethods`	A list of the restore methods that the snapshot supports.
`attachment`	A description of where a restore operation can attach a newly-created asset.

An asset may have additional fields to better describe it. The additional fields vary depending on the asset. In the REST API, you can search on any of the fields in an asset.

There is also a `_links` field, which provides links to related URLs in the REST API.

Sample API requests

This section shows how to perform common CloudPoint operations using the CloudPoint APIs. Use this section to learn how the syntax of a CloudPoint API maps to the actual curl command you specify on the command line.

Authenticating a user

Before you can perform an tasks with the CloudPoint APIs, a user must be authenticated. To perform user authentication using curl, enter the following command:

```
curl --cookie-jar path-to-cookie file -g -k -X \
POST -H "Content-Type: application/json" -d '{"username":username,\
"password":password}' https://hostname/api/rest/authenticate
```

This command returns the cookie and token for the user. The cookie is saved in the *path-to-cookie* file.

The token and the path to the cookie must be passed to the curl command to authenticate the user for all requests.

The following example uses curl with the token:

```
curl --cookie path-to-cookie file -g -k \
https://hostname/flexsnap/api/v1/assets/ \
--header 'Authorization: Bearer getaccesstoken' \
--header 'Content-Type: application/json'
```

GET requests

To list all assets, use the `GET /flexsnap/api/v1/assets/` API. The API request has the following format.

```
curl --cookie path-to-cookie file -g -k \
https://hostname/flexsnap/api/v1/assets/ \
--header 'Authorization: Bearer getaccesstoken' \
--header 'Content-Type: application/json'
```

Note: Before you use this API, make sure a user is authenticated and that you have a cookie and token for the user.

See “Authenticating a user” on page 68.

To get a specific asset, enter the following:

```
curl --cookie path-to-cookie file -g -k \
https://hostname/flexsnap/api/v1/assets/asset_id \
--header 'Authorization: Bearer getaccesstoken' \
--header 'Content-Type: application/json'
```

POST requests

To create a snapshot, use the `POST /flexsnap/api/v1/assets/assetid/snapshots/` API.

This API has the following request parameters:

`snapType`	The type of snapshot to create. You must choose this from the asset's list of `snapMethods`.
`name`	The name of the snapshot to create.
`description`	(Optional) A description of the snapshot.

Note: Before you use this API, make sure a user is authenticated and that you have a cookie and token for the user.

See "Authenticating a user" on page 68.

The following example creates a snapshot:

```
curl --cookie <path-to-cookie file> -X \
POST -H "Content-Type: application/json" -d \
'{"snapType": value from snap methods for the asset,\
 "name": snapshot name, "description": description for snapshot}' \
https://hostname/flexsnap/api/v1/assets/asset_id/snapshots/ \
-k --header 'Authorization: Bearer getaccesstoken'
```

DELETE request

To delete a snapshot, use the `DELETE /flexsnap/api/v1/assets/assetid /snapshots/snapid` API.

Note: Before you use this API, make sure a user is authenticated and that you have a cookie and token for the user.

See "Authenticating a user" on page 68.

The following example deletes a snapshot:

```
curl --cookie path to cookie file -X  \
DELETE https://hostname/flexsnap/api/v1/assets/ \
asset_id/snapshots/snap_id -k \
--header "Authorization: Bearer getaccesstoken"
```

PUT request

To restore a snapshot, use the `PUT /flexsnap/api/v1/assets/assetid` API.

This API has the following request parameter:

`snapid`	The asset ID of the snapshot to restore.

Note: Before you use this API, make sure a user is authenticated and that you have a cookie and token for the user.

See "Authenticating a user" on page 68.

The following example restores a snapshot:

```
curl --cookie path-to-cookie file -X \
PUT -H "Content-Type: application/json" -d '{"snapid": <snap_id>}’ \
https://hostname/flexsnap/api/v1/assets/asset_id -k \
--header 'Authorization: Bearer getaccesstoken’
```

PATCH request

To modify an existing policy, use the `PATCH {"op": opstr, "asset": assetid} /flexsnap/api/v1/policies/policy_id` API.

This API has the following request parameter:

`opstr`	Either `add` or `remove`.

Note: Before you use this API, make sure a user is authenticated and that you have a cookie and token for the user.

See “Authenticating a user” on page 68.

The following example modifies a policy:

```
curl --cookie path-to-cookie file -X \
PATCH -H "Content-Type: application/json" \
-d ‘{"op": add or remove, "asset": assetid}’ \
https://hostname/flexsnap/api/v1/policies/policy_id \
-k --header ‘Authorization: Bearer getaccesstoken’
```

Asset APIs

Asset APIs let you create, restore and delete snapshots of assets such as applications, hosts, and disks. You can also use asset APIs to list a group of asset objects, such as a list of hosts, or display information on a single asset.

List assets

```
GET /flexsnap/api/v1/assets/
```

Response: A list of asset objects.

Get a specific asset

```
GET /flexsnap/api/v1/assets/assetid
```

Response: The requested asset object.

List snapshots of an asset

```
GET /flexsnap/api/v1/assets/assetid/snapshots/
```

Response: A list of asset objects.

List a specific snapshot

```
GET /flexsnap/api/v1/assets/assetid/snapshots/snapid
GET /flexsnap/api/v1/assets/snapid
```

Response: The requested asset object.

Note: Because a snapshot is an asset, you can access it directly through the `/assets/` resource.

Create a snapshot

```
POST /flexsnap/api/v1/assets/assetid/snapshots/
```

This API has the following request parameters:

`snapType`	The type of snapshot to create. This must be chosen from the asset's `snapMethods`.
`name`	The name of the snapshot to create.
`description`	A description of the snapshot. This is optional.

Response: The asset object of the newly-created snapshot.

When the workflow is integrated, creating a snapshot responds with the following instead:

```
{
'status': status,
'error': error text,
'taskid': id of workflow task
}
```

status	Either pending, running, successful, or failed.
error	If the task failed, this field is the text of an exception.
taskid	The ID of the task dispatched to create the snapshot. Use this field to query for status.

Restore a snapshot

Depending on the asset type, you can restore a snapshot in the following ways:

- Restore the snapshot to the original asset, overwriting the asset's data.
- Restore a snapshot to its original location.
- Restore a snapshot to a new location.

Restore to the original asset, overwriting the asset's data

```
PUT /flexsnap/api/v1/assets/assetid
```

This API has the following request parameter:

snapid	The asset ID of the snapshot to restore.

Response: An asset object describing the restored asset.

Restore a snapshot to its original location

```
POST /flexsnap/api/v1/assets/
```

This API has the following request parameter:

snapid	The asset ID of the snapshot to restore.

Response: The asset object of the newly-created asset based on the snapshot.

This operation is used, for example, to restore an Amazon Elastic Block Store (EBS) snapshot to the instance where the volume originally was. EBS snapshots cannot be used to overwrite the data in an existing EBS volume. They can only be used to create a new EBS volume.

Restore a snapshot to a new location

`POST /flexsnap/api/v1/assets/`

This API has the following request parameters:

`snapid`	The asset ID of the snapshot to restore.
`dest`	The destination of the restore. The `dest` parameter value depends on the asset type. The snapshot's `attachment` field describes the type of attachment the restore operation supports. If the type is `name`, `dest` is the name of asset to create. Otherwise, it is an asset type; for example, a host. In this case, `dest` should be an asset ID. This operation is used, for example, to restore an EBS snapshot to a new Elastic Compute Cloud (EC2) instance.

Response: The asset object of the newly created asset based on the snapshot.

List restore targets

GET /flexsnap/api/v1/assets/*assetid*/snapshots/*snapid*/targets/

Response: List of asset objects which can be used as the destination when restoring the snapshot to a new location.

Delete a snapshot

DELETE /flexsnap/api/v1/assets/*assetid*/snapshots/*snapid*

Response: An object of type "success."

Agent management APIs

Agent management APIs let you configure CloudPoint agents and plug-ins, list agents and plug-ins, and display the return type of operations.

The information associated with an agent includes the following:

- Agentid: A unique identifier for the agent.
- Status: Status of agent whether it is running or not running.
- Host: Host system name where agent is running.
- Plugins: List of plug-ins configured on each agent.

- Target: Rabbitmq-server name.

The information associated with plug-in includes the following:

- PluginName : Name of plug-in.
- Configuration : Configuration for plug-in.

Agent APIs

GET /agents

Response: A list of agents with the following fields:

```
 [


    {
        "agentid": "<agent_id>",
        "hostname": "<the host on which it is running>",
        "lastMessage": <>,
        "onHost": true/false,
        "status": "online/offline"
    }
   ]
```

GET /agents/*agent_id*

Response: Information about a particular agent with the following fields:

```
{
        "agentid": "<agent_id>",
        "hostname": "<the host on which it is running>",
        "lastMessage": <>,
        "onHost": true/flase,
        "status": "online/offline"
    }
```

GET /agents/*agent_id*/plugins/

Response : A list of plug-ins configured on the particular agent ID.

```
[
    {
        "displayName": "<display name of the plugin>",
        "name": "plugin name",
        "version": "1.0"
    }
   ]
```

GET /agents/*agent_id*/plugins/*plugin_name*

Response: Information about the plug-in configured on the particular agent ID.

```
 {
        "displayName": "<display name of the plugin>",
        "name": "plugin name",
        "version": "1.0"
    }
```

Workflow APIs

List tasks

```
GET /flexsnap/api/v1/tasks/?
[status=status1,status2,...]
&run_since=n hours
&taskType=[operation namen1, ...]
&limit=[count]
&start_after=[count]
```

Arguments:

status can be "running", "successful", or "failed." Specify a comma-separated list to filter on or an OR statement of multiple status values.

run_since means list tasks started in the last number of hours.

limit and *start_after* are the output pagination limit and cursor.

Possible *operation name* options include:

```
"create-snapshot"
                "create-group-snapshot"
                "restore"
                "delete-snapshot"
                "delete-group-snapshots"
```

Response:

```
list of {
                  'taskid': <task id>,
                  'name': <task name>,
                  'status': <status>,
                  'progress': <integer between 0 and 100>,
                  'asset': <affected asset's id>,
                  'error': <error text>,
                  'ctime': <unix time>,
                  'mtime': <unix time>
       }
```

`unix time` is an integer representing the number of seconds since the epoch.

Show task properties, status, results

```
GET /flexsnap/api/v1/tasks/<task id>
```

Response:

```
{
'taskid': <task id>,
'name': <task name>,
'status': <status>,
'progress': <integer between 0 and 100>,
'asset': <affected asset's id>,
'error': <error text>
'ctime': <unix time>
'mtime': <unix time>
}
```

'name' describes the operation being performed.

'status' can be "pending", "running", "successful", or "failed."

'asset' is the id of the top level object that the task created, read, updated, or deleted.

'error' is the text of an exception if the task failed.

'unix time' is an integer representing the number of seconds since the epoch.

Delete tasks

```
DELETE /flexsnap/api/v1/tasks/task id
```

Arguments:

<task id> is the id of a task to delete.

Response:

```
{
'status': <status>,
'error': <error text>
}
```

```
DELETE /flexsnap/api/v1/tasks/?status=status&olderThan=number of days
```

Arguments:

status may be 'successful' or 'failed' in order to delete all completed tasks with the specified status.

olderThan delete all tasks older than the specified number of days.

Response:

```
{
'status': <status>,
'error': <error text>
}
```

Policy management APIs

Policy information includes the following:

- id : A string, a unique identifier assigned by the system for the policy.
- name: A string, user specified name for the policy.
- appConsist: “True” or "False", whether the snapshot will be application consistent.
- protectionLevel: "disk" or “host”, the level of protection.
- retentionCount: A numeric value, the number of snapshots to keep.
- snapTypePref: An ordered list of snap type such as COW or clone.
- schedule: A crontab format, the snapshot schedule for the policy, specified using the tags “minute”, “hour”, “mday”, “month”, and “wday.”
- tag: A string, any user specified information.
- assets: A list of assets that the policy applies to. You cannot modify this list directly. You must modify it through APIs.

List policies

```
GET /flexsnap/api/v1/policies/
```

Response: List of policies.

List a specific policy

```
GET /flexsnap/api/v1/policies/policyid
```

Response: The policy requested.

Create a policy

```
POST  /flexsnap/api/v1/policies/
```

This API has the following parameters:

Parameter	Required	Description	Data type
name	true	The snapshot name.	String: 2 to 32 characters
appConsist	true	Specifies whether the snapshot is application-consistent or crash-consistent.	Boolean: True or False
tag	true	A user-specified string with information about the policy; for example, the purpose of the policy.	String
snapTypePref	true	Specifies whether the snapshot is copy-on-write (COW) or clone.	cow, clone

Parameter		Required	Description	Data type
schedule	hour	true		Numeric: A value between 0-23, or *
	mday	true		Numeric: A value based on the number of days in the month – 1-28, 1-30, 1-31, or *
	minute	true		Numeric: 0-59
	month	true		Numeric: 1-12
	wday	true		Numeric: 1-7

Response: The policy is created.

Modify a policy

```
PUT /flexsnap/api/v1/policies/policyid
```

This API has the following parameters:

Parameter	Required	Description	Data type
name	true	The snapshot name.	String: 2 to 32 characters
appConsist	true	Specifies whether the snapshot is application-consistent or crash-consistent.	Boolean: True or False
tag	true	A user-specified string with information about the policy; for example, the purpose of the policy.	String

Parameter		Required	Description	Data type
snapTypePref		true	Specifies whether the snapshot is copy-on-write (COW) or clone.	cow, clone
schedule	hour	true		Numeric: A value between 0-23, or *
	mday	true		Numeric: A value based on the number of days in the month – 1-28, 1-30, 1-31, or *
	minute	true		Numeric: 0-59
	month	true		Numeric: 1-12
	wday	true		Numeric: 1-7

Response: The policy is modified.

Delete a policy

```
Delete /flexsnap/api/v1/policies/policyid
```

Response: An object of type "success."

Add or remove an asset to or from a policy

```
PATCH {"op": opstr,  "asset": "assetid} /flexsnap/api/v1/policies/policyid
```

The argument *opstr* is either “add” or “remove” (including the quotes).

Response: The policy is modified.

Apply a policy to an asset

```
PUT /flexsnap/api/v1/assets/assetid/policies/policyid
```

Response: The policy is modified.

Remove a policy from an asset

`DELETE /flexsnap/api/v1/assets/`*`assetid`*`/policies/`*`policyid`*

Response: The policy is modified.

List all policies applied to an asset

```
GET /flexsnap/api/v1/assets/<assetid>/policies/
```

Response: A list of policies applied to the given asset.

Example: creating a policy

```
curl -X POST -H "Content-type: application/json" \
-d '{"name":"SamplePolicy", "appConsist": false, \
"schedule":{"minute":"0", "hour":"2", \
"mday":"*", "month":"*", "wday":"*/4"}, \
"tag": "Backup at 02:00 on every Thursday"}' \
https://localhost/flexsnap/api/v1/policies/
```

The output is as follows:

```
{
    "appConsist": false,
    "assets": [],
    "id": "7539ff6d-e121-479f-9d73-175ac56214fe",
    "name": "SamplePolicy",
    "protectionLevel": "disk",
    "retentionCount": 5,
    "schedule": {
        "hour": "2",
        "mday": "*",
        "minute": "0",
        "month": "*",
        "wday": "*/4"
    },
    "snapTypePref": [
        "cow",
        "clone"
    ],
    "tag": "Backup at 02:00 on every Thursday"
}
```

Summary APIs

The following summary APIs are supported for the CloudPoint dashboard:

- Agent summary

  ```
  GET /flexsnap/api/v1/agents/summary
  ```

- Plug-in summary

  ```
  GET /flexsnap/api/v1/plugins/summary
  ```

- Asset summary

  ```
  GET /flexsnap/api/v1/assets/summary
  ```

Chapter 13

Storage array support

This chapter includes the following topics:

- Hitachi Data Systems (HDS) G-Series arrays
- Hewlett-Packard Enterprise (HPE) 3PAR array

Hitachi Data Systems (HDS) G-Series arrays

This section describes the following:

- The information you must supply to configure the Hitachi HDS Array plug-in
- The HDS G-series arrays that CloudPoint supports
- The CloudPoint operations you can perform on HDS G-series assets

Required G-Series array configuration information

CloudPoint parameter	Description
URL of HDS REST Server	The base URL for accessing the HDS REST API Server. The URL has the following format: `*protocol*://*host-name*:*port-number*/ConfigurationManager`
IP	The IP address of the HDS REST API Server
Username	The user name that is used to access the HDS REST API Server
Password	The password that is used to access the HDS REST API Server

Supported G-Series arrays

Array model	VSP G1000
Firmware version	80-01-21-XX/XX or later
Software development kit (SDK) required	Hitachi Configuration Manager

Supported CloudPoint operations on G-Series arrays

You can perform the following operations on disk assets:

- List all the disks.
- Create a thin image snapshot of a logical device (LDEV).
- Delete a thin image snapshot of an LDEV.
- Restore a thin image snapshot of an LDEV, overwriting the original object.

Hewlett-Packard Enterprise (HPE) 3PAR array

This section describes the following:

- The information you must supply to configure the 3PAR array plug-in
- The 3PAR arrays that CloudPoint supports
- The CloudPoint operations you can perform on 3PAR array assets

Required 3PAR array configuration information

CloudPoint configuration parameter	Description
Array IP Address	The IP address of the 3PAR array
Username	The user name that is used to log in to the array
Password	The password that is used to log in to the array

Supported 3PAR arrays

Array model	HP_3PAR 8200

Firmware version	3.1.3 firmware
Required software development kit	HP 3PAR Management Console 4.5.0
Library	hpe3parclient

Supported CloudPoint operations on 3PAR array assets

You can perform the following operations on 3PAR array assets:

- List all the disks.
- Create a copy-on-write (COW) virtual copy or clone (physical copy) snapshots of a volume.
- Delete a COW virtual copy or clone physical snapshots of a volume.
- Restore COW virtual copy snapshots of a volume, overwriting the original object.