

Veritas CloudPoint 2.1.2 Administrator's Guide

Linux

Veritas CloudPoint Administrator's Guide

Last updated: 2019-06-28

Document version: 2.1.2 Rev 0

Legal Notice

Copyright © 2019 Veritas Technologies LLC. All rights reserved.

Veritas, the Veritas Logo, Veritas InfoScale, and NetBackup are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third-party software for which Veritas is required to provide attribution to the third party ("Third-Party Programs"). Some of the Third-Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the third-party legal notices document accompanying this Veritas product or available at:

<https://www.veritas.com/licensing/process>

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Veritas Technologies LLC and its licensors, if any.

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.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on premises or hosted services. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Veritas Technologies LLC
2625 Augustine Drive
Santa Clara, CA 95054

<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

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The document version appears on page 2 of each guide. The latest documentation is available on the Veritas website:

<https://sort.veritas.com/documents>

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:

cloudpointdocs@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Chapter 1	Getting started with CloudPoint	9
	About CloudPoint	9
	What kinds of assets can you protect?	10
	Understanding your CloudPoint license	11
Section 1	Installing and configuring CloudPoint	16
Chapter 2	Preparing for installation	17
	About the deployment approach	17
	Deciding where to run CloudPoint	18
	Meeting system requirements	19
	Creating an instance or preparing the physical host to install CloudPoint	24
	Installing Docker	25
	Creating and mounting a volume to store CloudPoint data	25
	Verifying that specific ports are open on the instance or physical host	26
Chapter 3	Deploying CloudPoint	28
	Installing CloudPoint	28
	Configuring CloudPoint from your browser and signing in	31
	Verifying that CloudPoint installed successfully	36
Chapter 4	Using plug-ins to discover assets	37
	About plug-ins	37
	Determining the types of plug-ins and agents to install	38
Chapter 5	Configuring off-host plug-ins	40
	Configuring an off-host plug-in	40
	AWS plug-in configuration notes	43
	Configuring AWS permissions for CloudPoint	45

	Dell EMC Unity array plug-in configuration notes	48
	Google Cloud Platform plug-in configuration notes	49
	Google Cloud Platform permissions required by CloudPoint	49
	Configuring a GCP service account for CloudPoint	51
	Preparing the GCP service account for plug-in configuration	52
	HPE 3PAR plug-in configuration notes	53
	Microsoft Azure plug-in configuration notes	54
	Configuring permissions on Microsoft Azure	55
	Nutanix plug-in configuration notes	57
	Pure Storage FlashArray plug-in configuration notes	57
	Huawei OceanStor array plug-in configuration notes	58
Chapter 6	Configuring the on-host agents and plug-ins	59
	About agents	59
	Preparing to install the Linux-based on-host agent	61
	Optimizing your Oracle database data and metadata files	62
	Oracle plug-in configuration notes	62
	Preparing to install the Windows-based on-host agent	63
	About the installation and configuration process	63
	Downloading and installing the on-host agent	64
	Configuring the Linux-based on-host agent	67
	MongoDB plug-in configuration notes	69
	Configuring the Windows-based on-host agent	70
	Configuring the Windows-based agent on a host if an agent has been previously installed	72
	Configuring the on-host plug-in	73
	Configuring VSS to store shadow copies on the originating drive	74
	Enabling the Microsoft SQL plug-in on the Windows host	75
	Running the Windows agent as a service	76
Chapter 7	Protecting assets with CloudPoint's agentless feature	78
	About the agentless feature	78
	Prerequisites for the agentless configuration	79
	Granting password-less sudo access to host user account	79
	Configuring the agentless feature	80

Section 2	Configuring users	82
Chapter 8	Setting up email and adding users	83
	Configuring the CloudPoint sender email address	83
	About adding users to CloudPoint	86
	Adding AD users to CloudPoint using LDAP	87
	Adding users to CloudPoint manually	89
	Deleting a user from CloudPoint	91
Chapter 9	Assigning roles to users for greater efficiency	93
	About role-based access control	93
	Displaying role information	94
	Creating a role	94
	Editing a role	98
	Deleting a role	99
Section 3	Protecting and managing data	100
Chapter 10	User interface basics	101
	Signing in to CloudPoint	101
	Focusing on an asset type	102
	Navigating to your assets	103
	Using the action icons	105
Chapter 11	Protecting your assets with policies	106
	About policies	106
	Creating a policy	107
	Assigning a policy to an asset	111
	Listing policies and displaying policy details	113
	Editing a policy	115
	Deleting a policy	116
Chapter 12	Replicating snapshots for added protection	119
	About snapshot replication	119
	Requirements for replicating snapshots	120
	Configuring replication rules	120
	Editing a replication rule	122
	Deleting a replication rule	123

Chapter 13	Managing your assets	124
	Creating a snapshot manually	124
	Displaying asset snapshots	128
	Replicating a snapshot manually	130
	About snapshot restore	132
	About single file restore (granular restore)	135
	Single file restore requirements and limitations	136
	Single file restore support on Linux	137
	Single file restore limitations on Linux	137
	Single file restore support on Windows	137
	Single file restore limitations on Windows	137
	Restoring a snapshot	138
	Additional steps required after a SQL Server disk-level snapshot restore to a new location	140
	Restoring individual files within a snapshot	141
	Deleting a snapshot	145
Chapter 14	Monitoring activities with notifications and the job log	147
	Working with notifications	147
	Using the Job Log	148
Chapter 15	Indexing and classifying your assets	152
	About indexing and classifying snapshots	152
	Configuring classification settings using VIC	154
	Indexing and classifying snapshots	154
	Statuses for indexing and classification	155
Chapter 16	Protection and disaster recovery	157
	About protection and disaster recovery	157
	Backing up CloudPoint	158
	Restoring CloudPoint	161
Section 4	Maintaining CloudPoint	164
Chapter 17	CloudPoint logs	165
	CloudPoint logs	165

Chapter 18	Troubleshooting CloudPoint	168
	Restarting CloudPoint	168
	Docker may fail to start due to a lack of space	169
	Some CloudPoint features do not appear in the user interface	170
Chapter 19	Upgrading CloudPoint	172
	About CloudPoint upgrades	172
	Supported upgrade path	172
	Preparing to upgrade CloudPoint	172
	Upgrading CloudPoint	173
Chapter 20	Working with your CloudPoint license	179
	Displaying CloudPoint license and protection information	179
	Upgrading your CloudPoint license	180
Section 5	Reference	184
Chapter 21	Storage array support	185
	Dell EMC Unity arrays	185
	Dell EMC Unity array plug-in configuration parameters	185
	Supported Dell EMC Unity arrays	186
	Supported CloudPoint operations on Dell EMC Unity arrays	186
	Hewlett Packard Enterprise (HPE) 3PAR array	188
	3PAR array plug-in configuration parameters	188
	Supported 3PAR arrays	189
	Supported CloudPoint operations on 3PAR array assets	189
	Pure Storage FlashArray	190
	Pure Storage FlashArray plug-in configuration parameters	190
	Supported Pure Storage FlashArray models	190
	Supported CloudPoint operations on Pure Storage FlashArray models	191
	Huawei OceanStor arrays	192
	Huawei OceanStor array plug-in configuration parameters	192
	Supported Huawei OceanStor arrays	193
	Supported CloudPoint operations on Huawei OceanStor array	194
Chapter 22	Working with CloudPoint using APIs	197
	Accessing the Swagger-based API documentation	197

Getting started with CloudPoint

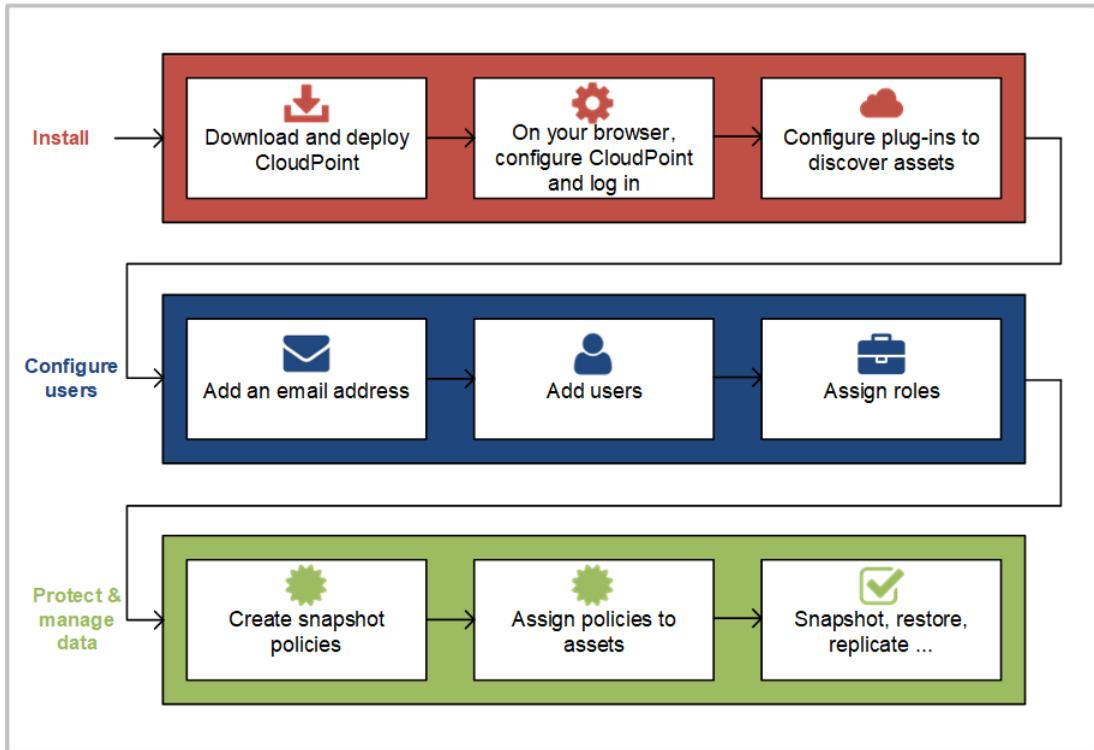
This chapter includes the following topics:

- [About CloudPoint](#)
- [What kinds of assets can you protect?](#)
- [Understanding your CloudPoint license](#)

About CloudPoint

Before you work with CloudPoint, it's helpful to have an overview. The following figure traces your path through CloudPoint, from installation and configuration through to data protection. Knowing this process makes getting started much easier.

Figure 1-1 Your path through CloudPoint



As you review the figure, keep in mind the following.

- Some of these tasks may only take a few minutes. You can be up and running with CloudPoint quickly.
- If you are managing a small environment and intend to only have one administrator, you can skip the steps on configuring users.
- The CloudPoint features you can use vary depending on the type of license you have. Also, some features may be in a technical preview stage. You should not use those features in a production environment. Any technical preview features are identified as such.

What kinds of assets can you protect?

CloudPoint offers snapshot-based data protection for your cloud or on-premises assets.

The following table shows the types of assets CloudPoint protects. The specific assets you can protect depends on the type of CloudPoint license you have.

Table 1-1 Supported assets

Category	Supported assets
Applications	<ul style="list-style-type: none">■ Amazon Relational Database Service (RDS) applications and Aurora database clusters■ MongoDB Enterprise Edition 3.6■ Microsoft SQL 2014 and 2016■ Oracle 12c, 12c R1
Disks	<ul style="list-style-type: none">■ Dell EMC Unity arrays■ Hitachi Data Systems G-Series arrays■ HPE 3PAR arrays■ Nutanix■ Pure Storage FlashArray arrays■ Huawei OceanStor arrays
File systems	File systems supported by the following operating systems: <ul style="list-style-type: none">■ Linux■ Windows 2012 and 2016
Hosts	<ul style="list-style-type: none">■ AWS EC2 instances■ Azure virtual machines■ Google virtual machines■ Nutanix virtual machines■ VMware virtual machines <p>Note: VMware VMs are supported only up to CloudPoint 2.1 release. Starting with CloudPoint 2.1.2, the CloudPoint plug-in for VMware has been deprecated. Support for VMware virtual machines is no longer available.</p>

Refer to the CloudPoint requirements for a more specific list of supported assets.

See [“Meeting system requirements”](#) on page 19.

Understanding your CloudPoint license

Your CloudPoint license determines the CloudPoint features you can use, the amount and kind of data you can protect, and the time for which you can continue to protect that data. CloudPoint offers various licensing options to choose from depending on your requirement.

If you want to try CloudPoint

If you want to explore CloudPoint features and functionality, the following options are available:

- **Freemium**

The Freemium license is a perpetual free license that does not expire and gives you a chance to try out a subset of the CloudPoint features in your on-premise or preferred cloud environment. This license lets you protect up to 10 TB of front-end terabyte data (FETB).

- **Evaluation**

The Evaluation license is a time-bound license that is valid for 60 days and allows you to try out all of the CloudPoint features in your on-premise or preferred cloud environment. This license lets you protect up to 1000 TB of front-end terabyte data (FETB).

After installing CloudPoint when you begin the initial CloudPoint configuration, you are presented with a choice to activate one of these licensing options. You must pick one of the licenses to complete the configuration and begin using CloudPoint.

If you want to buy CloudPoint

If you need more advanced features, you can upgrade the Freemium or Evaluation license to unlock the bundle that is right for you. CloudPoint offers the following types of paid licenses:

- **Enterprise**

The Enterprise edition is a fully-featured offering that is available as a perpetual as well as a subscription-based license. This is a fully loaded license that allows you to use all of the CloudPoint features in your preferred on-premise or cloud environments. This includes features such as application-consistent snapshots of your Oracle database, SQL Server, or MongoDB workloads, indexing, and classification.

- **On-prem**

The on-prem edition is a subset of the Enterprise edition and is available as a perpetual as well as a subscription-based license. This is a fully loaded license that allows you to use all of the CloudPoint features in your on-premise environment. This license enables you to discover the storage arrays and take application-consistent snapshots of the workloads deployed on them.

Unlike the Enterprise license, you cannot use this license to protect assets in your cloud environment.

- **In-cloud**

The in-cloud edition is also a subset of the Enterprise edition and is available as a perpetual as well as a subscription-based license. This is a fully loaded license that allows you to use of the CloudPoint features in your preferred cloud

environment. This license enables you to connect and discover the assets in Amazon Web Services (AWS), Google Cloud Platform (GCP), and Microsoft Azure cloud environments and take application-consistent snapshots of the workloads running on them.

Unlike the Enterprise license, you cannot use this license to protect assets in your on-premise environment.

More about perpetual and subscription licenses

All the paid licenses are available in both perpetual and subscription-based options. The perpetual licenses do not expire and are metered based on the FETB storage capacity. If the storage utilization reaches the maximum entitled capacity, you have to upgrade the license and expand the storage capacity to continue protecting newer assets.

Subscription-based licenses are time-bound and have to be renewed (or upgraded in case of an Evaluation license) at the end of their expiry date. Time-bound licenses are metered based on the amount of FETB storage capacity and the subscription validity period. If the license expires or the storage utilization has reached the maximum entitled limit, you have to renew or upgrade the license and expand the storage capacity to continue protecting newer assets.

The Enterprise edition (both perpetual and subscription-based license) is also available based on the number of workload instances that you want to protect.

The following table summarizes what each license provides.

Table 1-2 CloudPoint licensing options

	Freemium	Evaluation	Enterprise	On-prem	In-cloud
Cloud support					
<ul style="list-style-type: none"> ■ Amazon Web Services (AWS) ■ Google Cloud Platform (GCP) ■ Microsoft Azure 	✓	✓	✓	X	✓
Storage Array support					
<ul style="list-style-type: none"> ■ Hitachi Data System (HDS) ■ Dell EMC ■ HPE ■ Pure Storage ■ Huawei ■ NetApp ■ Nutanix 	✓	✓	✓	✓	X
Workload support					

Table 1-2 CloudPoint licensing options (*continued*)

	Freemium	Evaluation	Enterprise	On-prem	In-cloud
Host protection	✓	✓	✓	X	✓
Disks and Volumes	✓	✓	✓	✓	✓
<ul style="list-style-type: none"> ■ Microsoft Windows file system ■ Linux file systems 	X	✓	✓	✓	✓
<ul style="list-style-type: none"> ■ Microsoft SQL Server ■ Oracle database ■ MongoDB 	X	✓	✓	✓	✓
Feature support					
<ul style="list-style-type: none"> ■ Automated assets discovery ■ Manual and policy-based snapshots ■ Role-based access control (RBAC) ■ Snapshot search and recovery ■ Snapshot replication ■ Agentless ■ Active Directory (AD) integration 	✓	✓	✓	✓	✓
<ul style="list-style-type: none"> ■ Application-consistent snapshots ■ Indexing and classification ■ Granular recovery (Single File Restore) 	X	✓	✓	✓	✓
License meter					
Perpetual	✓	X	✓	✓	✓
Subscription (12, 24, and 36 months)	X	60 days	✓	✓	✓
Per Front End Terabyte (FETB)	10 TB	1000 TB	✓	✓	✓
Per 10-pack instance bundle	X	X	✓	X	X
Supportability					
VOX community	✓	✓	✓	✓	✓
Veritas essential support	X	X	✓	✓	✓

If CloudPoint is integrated with Veritas NetBackup

If you are integrating CloudPoint with Veritas NetBackup, the existing CloudPoint licenses and the corresponding features described in the table earlier will work as is. A separate license is not required.

See “[Displaying CloudPoint license and protection information](#)” on page 179.

See “[Upgrading your CloudPoint license](#)” on page 180.

The licensing options and feature entitlements described here are indicative and are subject to change. For the latest information on CloudPoint licensing, pricing, and procurement, contact your Veritas sales representative or refer to the following:

<https://www.veritas.com/product/backup-and-recovery/cloudpoint/buy>

Installing and configuring CloudPoint

- [Chapter 2. Preparing for installation](#)
- [Chapter 3. Deploying CloudPoint](#)
- [Chapter 4. Using plug-ins to discover assets](#)
- [Chapter 5. Configuring off-host plug-ins](#)
- [Chapter 6. Configuring the on-host agents and plug-ins](#)
- [Chapter 7. Protecting assets with CloudPoint's agentless feature](#)

Preparing for installation

This chapter includes the following topics:

- [About the deployment approach](#)
- [Deciding where to run CloudPoint](#)
- [Meeting system requirements](#)
- [Creating an instance or preparing the physical host to install CloudPoint](#)
- [Installing Docker](#)
- [Creating and mounting a volume to store CloudPoint data](#)
- [Verifying that specific ports are open on the instance or physical host](#)

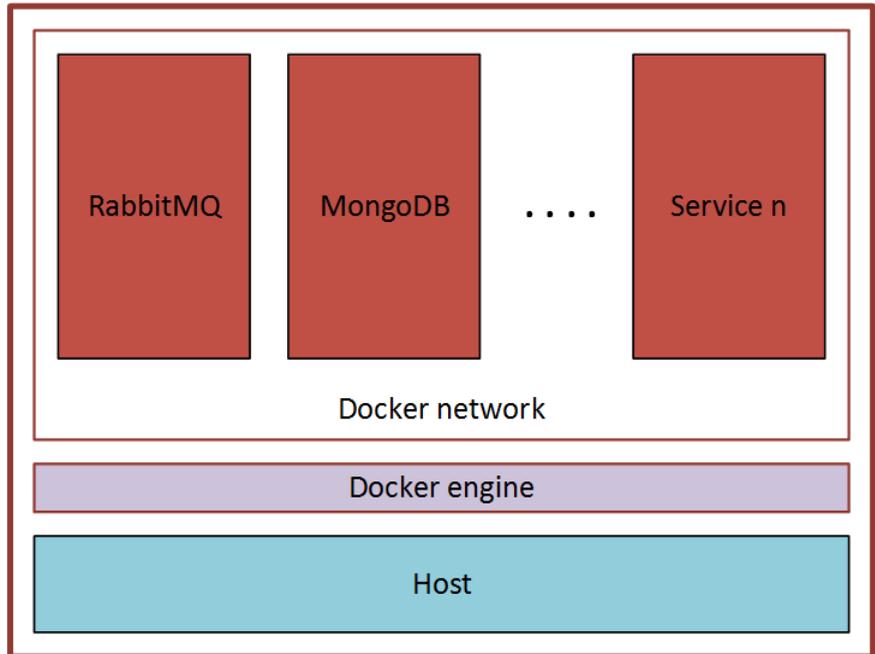
About the deployment approach

CloudPoint is distributed as a Docker image that is built on an Ubuntu 16.04 Server Long Term Support (LTS) base image or RHEL 7.5.

CloudPoint uses a micro-services model of installation. When you load and run the Docker image, CloudPoint installs each service as an individual container in the same Docker network. All containers securely communicate with each other using RabbitMQ.

Two key services are RabbitMQ and MongoDB. RabbitMQ is CloudPoint's message broker, and MongoDB stores information on all the assets CloudPoint discovers. The following figure shows CloudPoint's micro-services model.

Figure 2-1 CloudPoint's micro-services model



This deployment approach has the following advantages:

- CloudPoint has minimal installation requirements.
- Deployment requires only a few commands.

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.

Veritas recommends that you deploy CloudPoint in the same location as that of the assets that you wish to protect. If you wish to protect assets in a cloud, deploy the CloudPoint host instance in the same cloud environment. Similarly, if you wish to protect on-premise assets, deploy the CloudPoint host in the same on-premise environment.

If you install CloudPoint on multiple hosts, we strongly recommend that each CloudPoint instance manage separate resources. For example, two CloudPoint instances should not manage the same AWS account or the same Azure subscription. The following scenario illustrates why having two CloudPoint instances manage the same resources creates problems:

- CloudPoint instance A and CloudPoint instance B both manage the assets of the same AWS account.
- On CloudPoint instance A, the administrator takes a snapshot of an AWS virtual machine. The database on CloudPoint instance A stores the virtual machine's metadata. This metadata includes the virtual machine's storage size and its disk configuration.
- Later, on CloudPoint instance B, the administrator restores the virtual machine snapshot. CloudPoint instance B does not have access to the virtual machine's metadata. It restores the snapshot, but it does not know the virtual machine's specific configuration. Instead, it substitutes default values for the storage size configuration. The result is a restored virtual machine that does not match the original.

Meeting system requirements

CloudPoint host requirements

The host on which you install CloudPoint must meet the following requirements.

Table 2-1 Operating system and processor requirements for CloudPoint host

Category	Requirement
Operating system	<ul style="list-style-type: none"> ■ Ubuntu 16.04 Server LTS ■ Red Hat Enterprise Linux (RHEL) 7.5
Processor architecture	x86_64 / AMD64 / 64-bit processors

Table 2-2 System requirements for the CloudPoint host

Host on which CloudPoint is installed	Requirements
Amazon Web Services (AWS) instance	<ul style="list-style-type: none"> ■ Elastic Compute Cloud (EC2) instance type: t2.large ■ vCPUs: 2 ■ RAM: 8 GB ■ Root disk: 64 GB with a solid-state drive (GP2) ■ Data volume: 50 GB Elastic Block Store (EBS) volume of type GP2 with encryption for the snapshot asset database; use this as a starting value and expand your storage as needed.
Microsoft Azure VM	<ul style="list-style-type: none"> ■ Virtual machine type: D2s_V3 Standard ■ CPU cores: 2 ■ RAM: 8 GB ■ Root disk: 64 GB SSD ■ Data volume: 50 GB Premium SSD for the snapshot asset database; storage account type Premium_LRS; set Host Caching to Read/Write.
Google Cloud Platform (GCP) VM	<ul style="list-style-type: none"> ■ Virtual machine type: n1-standard-2 ■ vCPUs: 2 ■ RAM: 8 GB ■ Boot disk: 64 GB standard persistent disk, Ubuntu 16.04 Server LTS ■ Data volume: 50 GB SSD persistent disk for the snapshot asset database with automatic encryption
Nutanix VM	<ul style="list-style-type: none"> ■ Virtual machine type: 64-bit with a CloudPoint supported operating system ■ vCPUs: 8 ■ RAM: 8 GB or more ■ Root disk: 64 GB with a standard persistent disk ■ Data volume: 50 GB for the snapshot asset database
VMware VM	<ul style="list-style-type: none"> ■ Virtual machine type: 64-bit with a CloudPoint supported operating system ■ vCPUs: 8 ■ RAM: 8 GB or more ■ Root disk: 64 GB with a standard persistent disk ■ Data volume: 50 GB for the snapshot asset database

Table 2-2 System requirements for the CloudPoint host (*continued*)

Host on which CloudPoint is installed	Requirements
Physical host (<i>x86_64 / AMD64</i>)	<ul style="list-style-type: none"> ■ Operating system: A 64-bit CloudPoint supported operating system ■ CPUs: x86_64 (64-bit), single-socket, multi-core, with at least 8 CPU count ■ RAM: 8 GB or more ■ Boot disk: 64 GB ■ Data volume: 50 GB for the snapshot asset database

Disk space requirements

CloudPoint uses the following file systems on the host to store all the container images and files during installation:

- / (root file system)
- /var

The /var file system is further used for container runtimes. Ensure that the host on which you install CloudPoint has sufficient space for the following components.

Table 2-3 Space considerations for CloudPoint components

Component	Space requirements
CloudPoint Docker containers	5 GB
CloudPoint on-host agent and plug-ins	350 MB

Additionally, CloudPoint also requires a separate volume for storing CloudPoint data. Ensure that you create and mount this volume to /cloudpoint on the CloudPoint host.

Table 2-4 Space consideration for CloudPoint data volume

Volume mount path	Size
/cloudpoint	50 GB or more

Applications, operating systems, cloud, and storage platforms supported by CloudPoint agents and plug-ins

CloudPoint supports the following applications, operating systems, cloud, and storage platforms.

These assets are supported irrespective of how you configure CloudPoint, whether using the CloudPoint cloud or storage agents and plugins (earlier known as off-host plug-ins), or using the CloudPoint application configuration plugins (earlier known as on-host plug-ins), or using the CloudPoint agentless feature.

Table 2-5 Supported applications, operating systems, cloud, and storage platforms

Category	Support
Applications	<ul style="list-style-type: none"> ■ File systems <ul style="list-style-type: none"> ■ Linux native file systems: ext2, ext3, ext4, and XFS ■ Microsoft Windows: NTFS <p>For granular restore (single file restore (SFR)) support, refer to the following: See “Single file restore requirements and limitations” on page 136.</p> <ul style="list-style-type: none"> ■ Microsoft SQL 2014 and SQL 2016 ■ MongoDB Enterprise Edition 3.6 See “MongoDB plug-in configuration notes” on page 69. ■ Oracle 12c, Oracle 12c R1 CloudPoint has been verified on Oracle single node configurations. See “Oracle plug-in configuration notes” on page 62.
VMware	<p>vSphere 6.0 and later</p> <p>Note: VMware VMs are supported only up to CloudPoint 2.1 release. Starting with CloudPoint 2.1.2, the CloudPoint plug-in for VMware has been deprecated. Support for VMware virtual machines is no longer available.</p>
Operating systems on supported assets	<ul style="list-style-type: none"> ■ Red Hat Enterprise Linux (RHEL) 7.4, 7.5 Oracle has been verified on RHEL 7.1, 7.2, and 7.3 ■ Windows Server 2012 and Windows Server 2016

Table 2-5 Supported applications, operating systems, cloud, and storage platforms (*continued*)

Category	Support
Cloud platforms	<ul style="list-style-type: none"> ■ Amazon Web Services (AWS) If you wish to protect applications, the applications must be hosted on a t2.large or a higher specification AWS instance type. CloudPoint currently does not support applications that are running on t2.medium or a lower instance type. ■ Microsoft Azure If you wish to protect applications, the applications must be hosted on a D2s_V3 Standard or a higher specification Azure virtual machine type. ■ Google Cloud Platform (GCP) If you wish to protect applications, the applications must be hosted on a n1-standard-2 or a higher specification GCP virtual machine type. ■ Nutanix Acropolis Hypervisor (AHV)
Storage platforms	<ul style="list-style-type: none"> ■ Dell EMC Unity arrays See “Dell EMC Unity arrays” on page 185. ■ Hewlett Packard Enterprise (HPE) 3PAR arrays <ul style="list-style-type: none"> ■ Model: HP_3PAR 8200 ■ Firmware: 3.1.3 firmware ■ Software: HP 3PAR Management Console 4.5.0 See “Hewlett Packard Enterprise (HPE) 3PAR array” on page 188. ■ Pure Storage FlashArray See “Pure Storage FlashArray” on page 190. ■ Huawei OceanStor arrays See “Huawei OceanStor arrays” on page 192.

Table 2-6 CloudPoint compatibility lists

Compatibility list	Document link
Cloud Application Compatibility List (SCL)	
Hardware Compatibility List (HCL)	

Supported browsers

CloudPoint supports the following browsers for accessing the CloudPoint user interface.

Table 2-7 Supported browsers

Browser	Versions
Google Chrome	57.0.2987 or higher
Mozilla Firefox	52.0.0 or higher

Note: CloudPoint only runs on desktop devices. Mobile devices are not supported.

CloudPoint time zone

Ensure that the time zone settings on the host where you wish to deploy CloudPoint are as per your requirement and synchronized with a public NTP server.

By default, CloudPoint uses the time zone that is set on the host where you install CloudPoint. The timestamp for all the entries in the logs are as per the clock settings of the host machine.

However, the date and time for the operations and tasks in the CloudPoint user interface (UI) might reflect the browser time that corresponds to the local system from where the browser is launched.

Creating an instance or preparing the physical host to install CloudPoint

If you deploy CloudPoint in a public cloud, do the following:

- Choose an Ubuntu 16.04 Server LTS or RHEL 7.5 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 or RHEL 7.5 on a physical x86 server.
- Add sufficient storage to the server to meet the installation requirements.

Installing Docker

Note: CloudPoint supports Docker version 18.03 and later.

Table 2-8 Installing Docker

Platform	Description
Docker on Ubuntu	Refer to the following documentation for instructions on installing Docker on Ubuntu: https://docs.docker.com/install/linux/docker-ce/ubuntu/#set-up-the-repository
Docker on RHEL	Refer to the Docker documentation for detailed information: https://docs.docker.com/install/linux/docker-ee/rhel/#prerequisites Before installing CloudPoint, you must enable the shared mounts. To enable shared mounts <ol style="list-style-type: none"> 1 In the <code>docker.service</code> system unit file, modify the parameter MountFlags=slave to MountFlags=shared. 2 Save and close the unit file and then verify the change using the following command: <pre># cat /usr/lib/systemd/system/docker.service grep MountFlags</pre> The output resembles the following: <pre>MountFlags=shared</pre> 3 Reload the daemon using the following command: <pre># sudo systemctl daemon-reload</pre> 4 Restart the docker service using the following command: <pre># sudo systemctl restart docker</pre>

Creating and mounting a volume to store CloudPoint data

Before you deploy CloudPoint in a cloud environment, you must create and mount a volume of at least 50 GB to store CloudPoint data. The volume must be mounted to `/cloudpoint`.

Verifying that specific ports are open on the instance or physical host

Table 2-9 Volume creation steps for each supported cloud vendor

Vendor	Procedure
Amazon Web Services (AWS)	<ol style="list-style-type: none"> 1 On the EC2 dashboard, click Volumes > Create Volumes. 2 Follow the instructions on the screen and specify the following: <ul style="list-style-type: none"> ■ Volume type: General Purpose SSD ■ Size: 50 GB 3 Use the following instructions to create a file system and mount the device to <code>/cloudpoint</code> on the instance host. http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-using-volumes.html
Google Cloud Platform	<ul style="list-style-type: none"> ◆ Create the disk for the virtual machine, initialize it, and mount it to <code>/cloudpoint</code>. https://cloud.google.com/compute/docs/disks/add-persistent-disk
Microsoft Azure	<ol style="list-style-type: none"> 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 <code>/cloudpoint</code>. 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

Verifying that specific ports are open on the instance or physical host

Make sure that the following ports are open on the instance or physical host.

Table 2-10 Ports used by CloudPoint

Port	Description
443	The CloudPoint user interface uses this port as the default HTTPS port.

Table 2-10 Ports used by CloudPoint (*continued*)

Port	Description
5671	The CloudPoint RabbitMQ server uses this port for communications. This port must be open to support multiple agents.

Keep in mind the following:

- If the instance is in a cloud, configure the ports information under required inbound rules for your cloud.
- If you configure SMTP on ports 25, 465, or 587, make sure that the ports are accessible from the CloudPoint host and necessary firewall rules are created to allow inbound and outbound communication on the ports.

Deploying CloudPoint

This chapter includes the following topics:

- [Installing CloudPoint](#)
- [Configuring CloudPoint from your browser and signing in](#)
- [Verifying that CloudPoint installed successfully](#)

Installing CloudPoint

The following figure shows where you are at in the CloudPoint installation and configuration process.

Figure 3-1 You are here in the installation and configuration process



Before you complete the steps in this section, make sure that you complete the following:

- Decide where to install CloudPoint.
See [“Deciding where to run CloudPoint”](#) on page 18.

Note: If you plan to install CloudPoint on multiple hosts, read this section carefully and understand the implications of this approach.

- Ensure that your environment meets system requirements.
See [“Meeting system requirements”](#) on page 19.

- Create the instance on which you install CloudPoint or prepare the physical host.
See “[Creating an instance or preparing the physical host to install CloudPoint](#)” on page 24.
- Install Docker.
See “[Installing Docker](#)” on page 25.
- Create and mount a volume to store CloudPoint data.
See “[Creating and mounting a volume to store CloudPoint data](#)” on page 25.
- Verify that specific ports are open on the instance or physical host.
See “[Verifying that specific ports are open on the instance or physical host](#)” on page 26.

Note: When you deploy CloudPoint, you may want to copy the commands below and paste them in your command line interface. If you do, replace the information in these examples that is different from your own: the product and build version, the download directory path, and so on.

To deploy CloudPoint

- 1 Download the CloudPoint image.

You can use the free edition or purchase a licensed version. Refer to the following for more information:

<https://www.veritas.com/product/backup-and-recovery/cloudpoint/buy>

The CloudPoint image name has the following format:

```
Veritas_CloudPoint_2.x.x_IE.img.gz
```

- 2 (Optional) If necessary, copy the downloaded image to the system on which you want to deploy CloudPoint.
- 3 Change directories to where you have downloaded the CloudPoint image.

4 Type the following command to load the image into Docker:

```
# sudo docker load -i Veritas_CloudPoint_2.x.x_IE.img.gz
```

For example:

```
# sudo docker load -i Veritas_CloudPoint_2.0.2_IE.img.gz
```

Messages similar to the following appear on the command line:

```
788ce2310e2f: Loading layer [=====>] 126.8 MB/126.8 MB
aa4e47c45116: Loading layer [=====>] 15.87 kB/15.87 kB
b3968bc26fbd: Loading layer [=====>] 14.85 kB/14.85 kB
c9748fbf541d: Loading layer [=====>] 5.632 kB/5.632 kB
2f5b0990636a: Loading layer [=====>] 3.072 kB/3.072 kB
d1348a46025a: Loading layer [=====>] 214.2 MB/214.2 MB
de54ad3327fe: Loading layer [=====>] 12.06 MB/12.06 MB
a8f411dfb821: Loading layer [=====>] 1.35 GB/1.35 GB
dc3db1bf7ffd: Loading layer [=====>] 25.6 kB/25.6 kB
e2344be00294: Loading layer [=====>] 25.6 kB/25.6 kB
Loaded image: veritas/flexsnap-cloudpoint:2.0.2.5300
```

Make a note of the loaded image name and version that appears on the last line of the output. The version represents the CloudPoint product version that is being installed. You will specify these details in the next step.

5 Type the following command to run the CloudPoint container:

```
# sudo docker run -it --rm
-v /<full_path_to_volume_name>:/<full_path_to_volume_name>
-v /var/run/docker.sock:/var/run/docker.sock
veritas/flexsnap-cloudpoint:<version> install
```

Replace the following parameters as per your environment:

Parameter	Description
<full_path_to_volume_name>	Represents the path to the CloudPoint data volume, which typically is /cloudpoint.
<version>	Represents the CloudPoint product version that you noted in the earlier step.

For example, if the CloudPoint version is 2.0.2.5300, the command syntax is as follows:

```
# sudo docker run -it --rm -v /cloudpoint:/cloudpoint
-v /var/run/docker.sock:/var/run/docker.sock
veritas/flexsnap-cloudpoint:2.0.2.5300 install
```

Note: This is a single command. Ensure that you enter the command without any line breaks.

In this step, CloudPoint does the following:

- Creates containers for each of the CloudPoint services.
- Runs the `flexsnap-api` container.
- Creates self-signed keys and certificates for `nginx`.
- Runs the `flexsnap-cloudpointconsole` container.

When these operations are completed, CloudPoint displays the following in the command prompt:

```
Please go to the UI and configure CloudPoint now.  
Waiting for CloudPoint configuration to complete .....
```

If you have difficulty with this step, note the following:

- 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.
 - If Docker fails to start, it may be because there is not enough space available for MongoDB.
See [“Docker may fail to start due to a lack of space”](#) on page 169.
- 6** This concludes the CloudPoint deployment process. The next step is to launch the CloudPoint user interface in your browser and complete the final configuration steps.
- See [“Configuring CloudPoint from your browser and signing in”](#) on page 31.

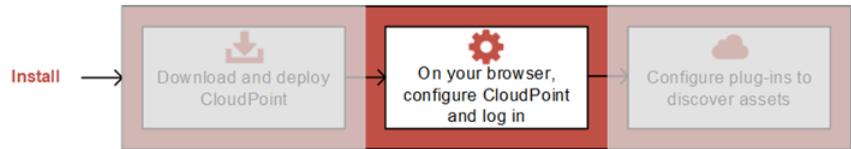
Note: If you ever need to restart CloudPoint, use the `docker run` command so that your environmental data is preserved.

See [“Restarting CloudPoint”](#) on page 168.

Configuring CloudPoint from your browser and signing in

The following figure shows where you are in the CloudPoint installation and configuration process.

Figure 3-2 You are here in the installation and configuration process



Before you complete the steps in this section, make sure that you have deployed CloudPoint on your instance or physical machine.

See [“Installing CloudPoint”](#) on page 28.

The final steps to configure CloudPoint are performed from a browser. Before you proceed, ensure that the browser is supported by CloudPoint.

See [“Meeting system requirements”](#) on page 19.

We recommend that you use Google Chrome.

To configure CloudPoint from your browser and sign in

- 1 Open your browser and enter the following URL in the address bar:

`https://<cloudpoint_hostFQDN>`

Here, `<cloudpoint_hostFQDN>` represents the Fully Qualified Domain Name (FQDN) of the host on which you installed CloudPoint.

The configuration screen is displayed.

Welcome to CloudPoint Initial Configuration

Admin Account Setup

Username *

@veritas.com

Password *

Confirm Password *

Host information

Hosts *

Hostnames

ec2-13-56-228-125.us-west-1.compute.amazonaws.com

Help us improve CloudPoint by automatically sending your usage information to Veritas.

I agree to the terms and conditions of the [EULA](#).

Configure

- 2 In the Admin Account Setup section, enter a username and password. They are configured as the CloudPoint administrator username and password.

The user name should meet the following requirement:

- A valid email address
If you forget the admin password, you can configure CloudPoint to send instructions for restoring the password to this email address.

The admin password should meet the following requirements:

- At least six characters
- No spaces
- No & (ampersand) character

- 3 If you want to add additional host names, then in the Host Information section, enter the name in the **Hosts** field and click **+**.

The specified name is added to the list of host names to use for configuring CloudPoint. The name (CloudPoint host FQDN) that you used to launch the initial configuration screen earlier is added to the list by default.

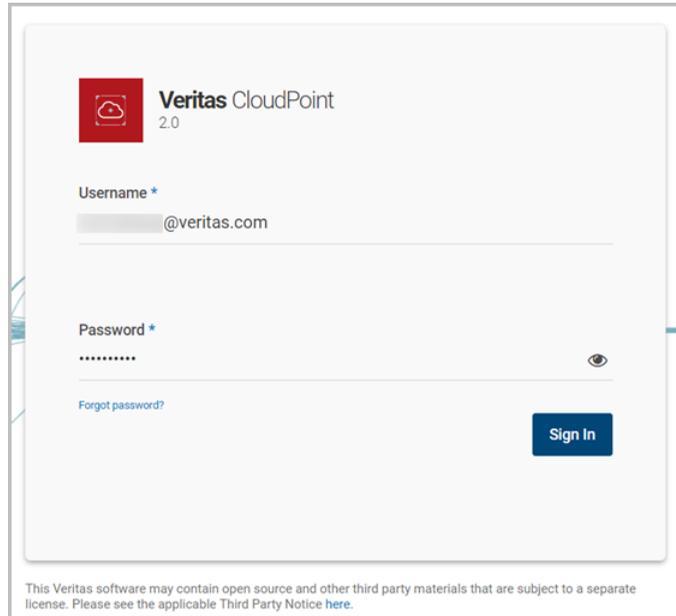
The host name is used to generate a server certificate for the CloudPoint host. If you connect to the host using different names (for example, *myserver*, *myserver.mydomain*, or *myserver.mydomain.mycompany.com*), then ensure that you add all the names here if you want to enable CloudPoint access using those names.

Note: The names you add in this field must point to the same CloudPoint host. Typically only one hostname is configured and it is generally the Fully Qualified Domain Name (FQDN) of the host.

- 4 Select **Help us improve CloudPoint by automatically sending your usage information to Veritas** to enable the Telemetry service. When enabled, this service collects your CloudPoint usage information and shares it with Veritas anonymously.
- 5 Read the End User License Agreement and then select the I agree to the terms and conditions option.
- 6 Click **Configure** to begin the initial configuration process.
An installation status screen is displayed as Veritas CloudPoint configures the remaining services. This process can take a few minutes.
- 7 After the installation completes, click **Refresh browser**. If you see the login screen, it confirms that CloudPoint is installed and configured successfully.

- 8 On the login screen, enter the CloudPoint administrator username and password and then click **Sign In**.

The username and password are the same that you specified on the initial configuration screen in step 2 earlier.



The coffee screen is displayed. The coffee screen provides a quick high level overview of your CloudPoint environment. After you configure CloudPoint to protect your assets, you can use this coffee screen to get a quick update on the overall protection status.

- 9 Your next step is to configure one or more plug-ins. On the coffee screen, click **Manage cloud and arrays**.



Plug-ins are the software modules that discover assets in your cloud or on-premise environment.

See [“Verifying that CloudPoint installed successfully”](#) on page 36.

Verifying that CloudPoint installed successfully

Verify that CloudPoint installed successfully by doing one of the following on the physical machine or instance command line:

- Verify that the success message is displayed.

```
Configuration complete at time Mon Jan 22 at 29:11:02 UTC 2018!
```

- Verify that the CloudPoint services are running and have UP status.

```
# sudo docker ps -a
```

The command output resembles the following:

CONTAINER ID	IMAGE	CREATED	STATUS
f4c70b6accff	veritas/flexsnap-cloudpointconsole:2.1.2.7542	6 hours ago	Up 6 hours
1cfe9f79f260	veritas/flexsnap-api:2.1.2.7542	6 hours ago	Up 6 hours
331c81a09ba2	veritas/flexsnap-api:2.1.2.7542	6 hours ago	Up 6 hours
4a2337b0af95	veritas/flexsnap-api:2.1.2.7542	6 hours ago	Up 6 hours
b4096679da38	veritas/flexsnap-api:2.1.2.7542	6 hours ago	Up 6 hours
27cd6a38d120	veritas/flexsnap-api:2.1.2.7542	6 hours ago	Up 6 hours
524dde7a1060	veritas/flexsnap-api:2.1.2.7542	6 hours ago	Up 6 hours
8bf5d31d948f	veritas/flexsnap-authorization-service:2.1.2.7542	6 hours ago	Up 6 hours
a1566d261f70	veritas/flexsnap-email-service:2.1.2.7542	6 hours ago	Up 6 hours
e8a4bd103b1f	veritas/flexsnap-identity-manager-service:2.1.2.7542	6 hours ago	Up 6 hours
52f26268ed26	veritas/flexsnap-licensing:2.1.2.7542	6 hours ago	Up 6 hours
da76eadf3c25	veritas/flexsnap-vic:2.1.2.7542	6 hours ago	Up 6 hours
4206a48a4d6b	veritas/flexsnap-telemetry:2.1.2.7542	6 hours ago	Up 6 hours
b54d1a6201e4	veritas/flexsnap-indexingsupervisor:2.1.2.7542	6 hours ago	Up 6 hours
9b0983c6418d	veritas/flexsnap-policy:2.1.2.7542	6 hours ago	Up 6 hours
6b3c14169321	veritas/flexsnap-scheduler:2.1.2.7542	6 hours ago	Up 6 hours
ba810e1f52f6	veritas/flexsnap-onhostagent:2.1.2.7542	6 hours ago	Up 6 hours
bbd1b1286e1a	veritas/flexsnap-agent:2.1.2.7542	6 hours ago	Up 6 hours
74b4742b589f	veritas/flexsnap-coordinator:2.1.2.7542	6 hours ago	Up 6 hours
8b9e22f8479d	veritas/flexsnap-mongodb:2.1.2.7542	6 hours ago	Up 6 hours (healthy)
8beead9166df	veritas/flexsnap-rabbitmq:2.1.2.7542	6 hours ago	Up 6 hours (healthy)
df3ebf833cfc	veritas/flexsnap-api-gateway:2.1.2.7542	6 hours ago	Up 6 hours
3710246dbd61	veritas/flexsnap-auth:2.1.2.7542	6 hours ago	Up 6 hours

Note: The number displayed in the image name (2.1.2.7542) represents the CloudPoint version. The version may vary depending on the actual product version being installed.

The command output displayed here is truncated to fit the view. The actual output may include additional details such as container names and ports used.

Using plug-ins to discover assets

This chapter includes the following topics:

- [About plug-ins](#)
- [Determining the types of plug-ins and agents to install](#)

About plug-ins

The following figure shows where you are in the CloudPoint installation and configuration process.

Figure 4-1 You are here in the installation and configuration process



If you have not completed the previous tasks, do so now.

See [“Installing CloudPoint”](#) on page 28.

See [“Configuring CloudPoint from your browser and signing in”](#) on page 31.

A CloudPoint plug-in is a low-level Python module that discovers assets in your environment and performs operations on them.

A plug-in has the following characteristics:

- A plug-in operates only on a particular asset type. For example, there is an AWS plug-in, a Pure Storage FlashArray plug-in, and so on.

- The following types of plug-ins are available:
 - An **off-host plug-in** runs separately from the instance or host on which the application runs.
For example, the CloudPoint AWS, Microsoft Azure, and Google plug-ins are off-host plug-ins for cloud environment. Similarly, the CloudPoint Pure Storage FlashArray and Dell EMC plug-ins are off-host plug-ins for storage arrays.
 - An **on-host plug-in** runs 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 underlying storage before taking the snapshot. It unquiesces them after the snapshot completes. The on-host plug-in also helps in the restore operation to mount a file system and bring up the application.
For example, the CloudPoint Oracle plug-in, the Linux file system plug-in, and the Microsoft Windows plug-in are examples of on-host plug-ins.
- You can run multiple instances of a plug-in to gather information from multiple sources within a particular type of asset. 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.
- Each plug-in is wrapped in an agent.

See [“About agents”](#) on page 59.

See [“Determining the types of plug-ins and agents to install”](#) on page 38.

Determining the types of plug-ins and agents to install

To determine the types of plug-ins and agents to install, use the following guidelines:

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

- CloudPoint supports the following off-host plug-ins:
 - Amazon AWS
 - Dell EMC Unity Array
 - Google Cloud Platform
 - Hewlett-Packard Enterprise 3PAR array
 - Huawei OceanStor array
 - Microsoft Azure
 - Nutanix Acropolis Hypervisor (AHV)
 - Pure Storage FlashArray

Note: NetApp and Hitachi Data Systems (HDS) storage arrays are not supported in this release even though the CloudPoint plug-ins for NetApp and HDS may be visible in the CloudPoint user interface (UI).

- CloudPoint supports the following on-host plug-ins:
 - Linux file systems ext2, ext3, ext4, and XFS
 - Microsoft Windows
 - Oracle database
 - MongoDB
 - Microsoft SQL

Configuring off-host plug-ins

This chapter includes the following topics:

- [Configuring an off-host plug-in](#)
- [AWS plug-in configuration notes](#)
- [Dell EMC Unity array plug-in configuration notes](#)
- [Google Cloud Platform plug-in configuration notes](#)
- [HPE 3PAR plug-in configuration notes](#)
- [Microsoft Azure plug-in configuration notes](#)
- [Nutanix plug-in configuration notes](#)
- [Pure Storage FlashArray plug-in configuration notes](#)
- [Huawei OceanStor array plug-in configuration notes](#)

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.

The steps to configure an off-host plug-in are the same, regardless of the particular asset. Only the configuration parameters vary.

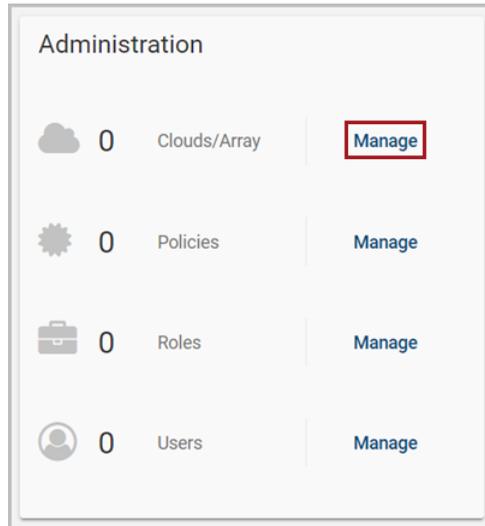
Before you complete the steps in this section, make sure that you gather the information you need to configure your particular plug-in.

See [“AWS plug-in configuration notes”](#) on page 43.

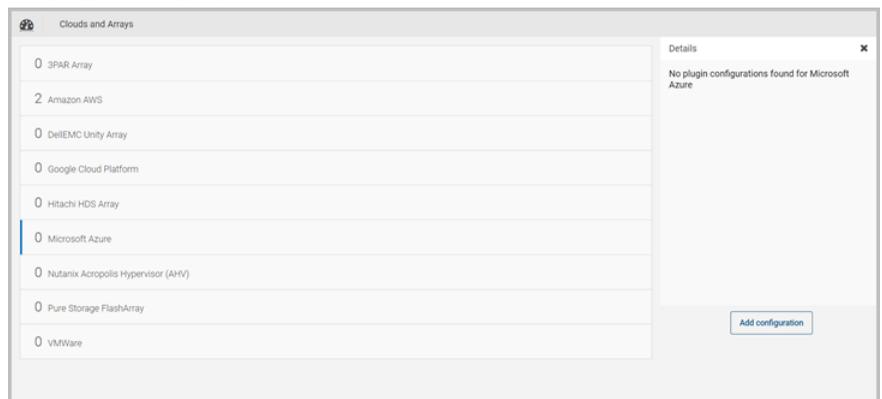
- See [“Dell EMC Unity array plug-in configuration notes”](#) on page 48.
- See [“Google Cloud Platform plug-in configuration notes”](#) on page 49.
- See [“HPE 3PAR plug-in configuration notes”](#) on page 53.
- See [“Microsoft Azure plug-in configuration notes”](#) on page 54.
- See [“Nutanix plug-in configuration notes”](#) on page 57.
- See [“Pure Storage FlashArray plug-in configuration notes”](#) on page 57.
- See [“Huawei OceanStor array plug-in configuration notes”](#) on page 58.

To configure an off-host plug-in

- 1 On the dashboard, in the **Administration** widget, locate **Clouds/Array**, and click **Manage**.

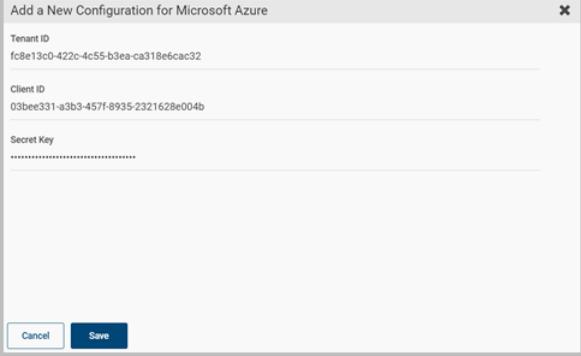


- 2 On the **Clouds and Arrays** page, select the plug-in to configure. (This example configures an Azure plug-in. When you select the plug-in, the **Details** page for the plug-in is displayed.



- 3 On the **Details** page, click **Add configuration**.

- 4 On the **Add a New Configuration** page, enter the configuration parameters you gathered for the plug-in. This Azure example specifies the **Tenant ID**, **Client ID**, and **Secret Key**.



Add a New Configuration for Microsoft Azure

Tenant ID
fb8e13c0-422c-4c55-b3ea-ca318e6cac32

Client ID
03bee331-a3b3-457f-8935-2321628e004b

Secret Key
.....

Cancel Save

Note: If you configure a Google Cloud plug-in, make sure you that you format the private key data properly before you enter it in the **Private Key** field.

See [“Google Cloud Platform plug-in configuration notes”](#) on page 49.

- 5 After you complete the configuration screen, click **Save**.

After you configure the plug-in, return to the dashboard. The statistics for applications, hosts, file systems, and disks are updated as appropriate. This update indicates the new plug-in has discovered assets.

AWS plug-in configuration notes

The Amazon Web Services (AWS) plug-in lets you create, restore, and delete snapshots of the following assets in an Amazon cloud:

- Elastic Compute Cloud (EC2) instances
- Elastic Block Store (EBS) volumes
- Amazon Relational Database Service (RDS) instances
- Aurora clusters

Note: Before you configure the AWS plug-in, make sure that you have configured the proper permissions so CloudPoint can work with your AWS assets.

The following information is required for configuring the CloudPoint plug-in for AWS:

Table 5-1 AWS plug-in configuration parameters

CloudPoint configuration parameter	AWS equivalent term and description
Access key	The access key ID, when specified with the secret access key, authorizes CloudPoint to interact with the AWS APIs.
Secret key	The secret access key.
Regions	One or more AWS regions in which to discover cloud assets.

Note: CloudPoint encrypts credentials using AES-256 encryption.

When CloudPoint connects to AWS, it uses the following endpoints. You can use this information to create a whitelist on your firewall.

- ec2.*.amazonaws.com
- sts.amazonaws.com
- rds.*.amazonaws.com
- kms.*.amazonaws.com

In addition, you must specify the following resources and actions:

- ec2.SecurityGroup.*
- ec2.Subnet.*
- ec2.Vpc.*
- ec2.createInstance
- ec2.runInstances

AWS plug-in considerations and limitations

This plug-in has the following limitations:

- You cannot delete automated snapshots of RDS instances and Aurora clusters through CloudPoint.
- All automated snapshot names start with the pattern `rds:..`

See [“Configuring an off-host plug-in”](#) on page 40.

Configuring AWS permissions for CloudPoint

To protect your Amazon Web Services (AWS) assets, CloudPoint must first have access to them. You must associate a permission policy with each CloudPoint user who wants to work with AWS assets.

Ensure that the user account or role is assigned the minimum permissions required for CloudPoint.

See [“AWS permissions required by CloudPoint”](#) on page 46.

To configure permissions on Amazon Web Services

- 1 Create or edit an AWS user account from Identity and Access Management (IAM).
- 2 Do one of the following.
 - To create a new AWS user account, do the following:
 - From IAM, select the **Users** pane and click **Add user**.
 - In the **User name** field, enter a name for the new user.
 - Select the **Access** type. This value determines how AWS accesses the permission policy. (This example uses Programmatic access).
 - Select **Next: Permissions**.
 - On the **Set permissions for username** screen, select **Attach existing policies directly**.
 - Select the previously created permission policy (shown below) and select **Next: Review**.
 - On the **Permissions summary** page, select **Create user**.
 - Obtain the **Access Key** and **Secret Key** for the newly created user.
 - To edit an AWS user account, do the following:
 - Select **Add permissions**.
 - On the **Grant permissions** screen, select **Attach existing policies directly**.
 - Select the previously created permission policy (shown below), and select **Next: Review**.
 - On the **Permissions summary** screen, select **Add permissions**.
- 3 To configure the AWS plug-in for the created or edited user, refer to the plug-in configuration notes.

See [“AWS plug-in configuration notes”](#) on page 43.

AWS permissions required by CloudPoint

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "EC2Backup",
      "Effect": "Allow",
      "Action": [
        "sts:GetCallerIdentity",
        "ec2:CreateSnapshot",
        "ec2:DescribeInstances",
        "ec2:DescribeInstanceStatus",
        "ec2:ModifySnapshotAttribute",
        "ec2:CreateImage",
        "ec2:CopyImage",
        "ec2:DescribeSnapshots",
        "ec2:DescribeVolumeStatus",
        "ec2:ModifySnapshotAttribute",
        "ec2:DescribeVolumes",
        "ec2:RegisterImage",
        "ec2:DescribeVolumeAttribute",
        "ec2:DescribeSubnets",
        "ec2:DescribeVpcs",
        "ec2:DeregisterImage",
        "ec2>DeleteSnapshot",
        "ec2:DescribeInstanceAttribute",
        "ec2:DescribeRegions",
        "ec2:ModifyImageAttribute",
        "ec2:DescribeAvailabilityZones",
        "ec2:ResetSnapshotAttribute",
        "ec2:DescribeHosts",
        "ec2:DescribeImages"
      ],
      "Resource": [
        "*"
      ]
    },
    {
      "Sid": "EC2Recovery",
      "Effect": "Allow",
      "Action": [
        "ec2:RunInstances",

```

```
    "ec2:AttachNetworkInterface",
    "ec2:DetachVolume",
    "ec2:AttachVolume",
    "ec2>DeleteTags",
    "ec2:CreateTags",
    "ec2:StartInstances",
    "ec2:RunInstances",
    "ec2:StopInstances",
    "ec2:CreateVolume",
    "ec2>DeleteVolume"
  ],
  "Resource": [
    "*"
  ]
},
{
  "Sid": "RDSBackup",
  "Effect": "Allow",
  "Action": [
    "rds:DescribeDBSnapshots",
    "rds:DescribeDBClusters",
    "rds:DescribeDBClusterSnapshots",
    "rds>DeleteDBSnapshot",
    "rds:CreateDBSnapshot",
    "rds:CreateDBClusterSnapshot",
    "rds:ModifyDBSnapshotAttribute",
    "rds:DescribeDBSubnetGroups",
    "rds:DescribeDBInstances",
    "rds:CopyDBSnapshot",
    "rds:CopyDBClusterSnapshot",
    "rds:DescribeDBSnapshotAttributes",
    "rds>DeleteDBClusterSnapshot",
    "rds:ListTagsForResource"
  ],
  "Resource": [
    "*"
  ]
},
{
  "Sid": "RDSRecovery",
  "Effect": "Allow",
  "Action": [
    "rds:ModifyDBInstance",
```

```

        "rds:ModifyDBClusterSnapshotAttribute",
        "rds:RestoreDBInstanceFromDBSnapshot",
        "rds:ModifyDBCluster",
        "rds:RestoreDBClusterFromSnapshot",
        "rds:CreateDBInstance",
        "rds:RestoreDBClusterToPointInTime",
        "rds:CreateDBSecurityGroup",
        "rds:CreateDBCluster",
        "rds:RestoreDBInstanceToPointInTime"
    ],
    "Resource": [
        "*"
    ]
},
{
    "Sid": "KMS",
    "Effect": "Allow",
    "Action": [
        "kms:ListAliases",
        "kms:ListKeys"
    ],
    "Resource": [
        "*"
    ]
}
]
}

```

Dell EMC Unity array plug-in configuration notes

Table 5-2 Dell EMC Unity array plug-in configuration parameters

CloudPoint configuration parameter	Description
Array IP Address	The IP address of the array.
Username	The username to access the array.
Password	The password to access the array.

Before you configure the plug-in, ensure that the specified user account has permissions to create, delete, and restore snapshots on the array.

For more information, see the [EMC Unity™ Quick Start Guide](#).

See “[Dell EMC Unity arrays](#)” on page 185.

See “[Configuring an off-host plug-in](#)” on page 40.

Google Cloud Platform plug-in configuration notes

The Google Cloud Platform plug-in lets you create, delete, and restore disk and host-based snapshots in all zones where Google Cloud is present.

Table 5-3 Google Cloud Platform plug-in configuration parameters

CloudPoint configuration parameter	Google equivalent term and description
Project ID	The ID of the project from which the resources are managed. Listed as <code>project_id</code> in the JSON file.
Client ID	The Client ID that is used for operations. Listed as <code>client_id</code> in the JSON file.
Client Email	The email address of the Client ID. Listed as <code>client_email</code> in the JSON file.
Private Key ID	The ID of the <code>private_key</code> . Listed as <code>private_key_id</code> in the JSON file.
Private Key	The private key. Listed as <code>private_key</code> in the JSON file. 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	A list of zones in which the plug-in operates.

See “[Google Cloud Platform permissions required by CloudPoint](#)” on page 49.

See “[Configuring a GCP service account for CloudPoint](#)” on page 51.

See “[Preparing the GCP service account for plug-in configuration](#)” on page 52.

See “[Configuring an off-host plug-in](#)” on page 40.

Google Cloud Platform permissions required by CloudPoint

Assign the following permissions to the service account that CloudPoint uses to access assets in the Google Cloud Platform:

```
compute.diskTypes.get
compute.diskTypes.list
compute.disks.create
compute.disks.createSnapshot
compute.disks.delete
compute.disks.get
compute.disks.list
compute.disks.setIamPolicy
compute.disks.setLabels
compute.disks.update
compute.disks.use
compute.globalOperations.get
compute.globalOperations.list
compute.images.get
compute.images.list
compute.instances.addAccessConfig
compute.instances.attachDisk
compute.instances.create
compute.instances.delete
compute.instances.detachDisk
compute.instances.get
compute.instances.list
compute.instances.setDiskAutoDelete
compute.instances.setMachineResources
compute.instances.setMetadata
compute.instances.setMinCpuPlatform
compute.instances.setServiceAccount
compute.instances.updateNetworkInterface
compute.instances.setLabels
compute.instances.setMachineType
compute.instances.setTags
compute.instances.start
compute.instances.stop
compute.instances.use
compute.machineTypes.get
compute.machineTypes.list
compute.networks.get
compute.networks.list
compute.projects.get
compute.regionOperations.get
compute.regionOperations.list
compute.regions.get
compute.regions.list
```

```
compute.snapshots.create
compute.snapshots.delete
compute.snapshots.get
compute.snapshots.list
compute.snapshots.setLabels
compute.snapshots.useReadOnly
compute.subnetworks.get
compute.subnetworks.list
compute.subnetworks.update
compute.subnetworks.use
compute.subnetworks.useExternalIp
compute.zoneOperations.get
compute.zoneOperations.list
compute.zones.get
compute.zones.list
```

Configuring a GCP service account for CloudPoint

To protect the assets in Google Cloud Platform (GCP), CloudPoint requires permissions to be able to access and perform operations on those cloud assets. You must create a custom role and assign it with the minimum permissions that CloudPoint requires. You then associate that custom role with the service account that you created for CloudPoint.

Perform the following steps:

- 1 Create a custom IAM role in GCP. While creating the role, add all the permissions that CloudPoint requires.

See “[Google Cloud Platform permissions required by CloudPoint](#)” on page 49.

Refer to the following GCP documentation for detailed instructions:

<https://cloud.google.com/iam/docs/creating-custom-roles>

- 2 Create a service account in GCP.

Grant the following roles to the service account:

- The custom IAM role that you created in the earlier step. This is the role that has all the permissions that CloudPoint requires to access GCP resources.
- The `iam.serviceAccountUser` role. This enables the service account to connect to the GCP using the service account context.

Refer to the following GCP documentation for detailed instructions:

<https://cloud.google.com/iam/docs/creating-managing-service-accounts#iam-service-accounts-create-console>

Preparing the GCP service account for plug-in configuration

To prepare for thtte CloudPoint GCP plug-in configuration

1 Gather the GCP configuration parameters that CloudPoint requires.

See “[Google Cloud Platform plug-in configuration notes](#)” on page 49.

Do the following:

- From the Google Cloud console, navigate to **IAM & admin > Service accounts**.
- Click the assigned service account. Click the three vertical buttons on the right side and select **Create key**.
- Select **JSON** and click **CREATE**.
- In the dialog box, click to save the file. This file contains the parameters you need to configure the Google Cloud plug-in. The following is a sample JSON file showing each parameter in context. The `private-key` is truncated for readability.

```
{
  "type": "service_account",
  "project_id": "fake-product",
  "private_key_id": "somaloguid1234567890",
  "private_key": "-----BEGIN PRIVATE KEY-----\n
N11EvA18ADAN89kq4k199w08AQEF5C8KYw9951A9EAAo18AQCNvpuJ3oK974z4\n
.\n
.\n
weT9odE4ryl181tNU\nV3q1XNX4fK55QTPd6CNU+f7QjEw5x8+5ft05DU8ayQcNkX\n
4pXJoDo154N52+T4qV4WkoFD5uL4NLPz5wxfly\nnNwCnfru8K8a2q1/9o0U+99==\n
-----END PRIVATE KEY-----\n",
  "client_email": "email@xyz-product.iam.gserviceaccount.com",
  "client_id": "0000000000000001",
  "auth_uri": "https://accounts.google.com/o/oauth2/auth",
  "token_uri": "https://accounts.google.com/o/oauth2/token",
  "auth_provider_x509_cert_url": "https://www.googleapis.com \
/o/oauth2/v1/certs",
  "client_x509_cert_url": "https://www.googleapis.com/robot/v1 \\"
```

```

    /metadata/x509/ email%40xyz-product.iam.gserviceaccount.com"
}

```

- 2 Using a text editor, reformat the `private_key` so it can be entered in the CloudPoint user interface. When you look in the file you created, each line of the private key ends with `\n`. You must replace each instance of `\n` with an actual carriage return. Do one of the following:
 - If you are a UNIX administrator, enter the following command in `vi`. In the following example, the `^` indicates the `Ctrl` key. Note that only the `^M` is visible on the command line.


```
:g/\n/s//^V^M/g
```
 - If you are a Windows administrator, use WordPad or a similar editor to search on `\n` and manually replace each instance.
- 3 When you configure the plug-in from the CloudPoint user interface, copy and paste the reformatted private key into the **Private Key** field. The reformatted `private_key` should look similar to the following:

```

-----BEGIN PRIVATE KEY-----\
N11EvA18ADAN89kq4k199w08AQEFAA5C8KYw9951A9EAAo18AQcNvpvJ3oK974z4
.
.
.
weT9odE4ry181tNU\nV3q1XNX4fK55QTPd6Cnu+f7QjEw5x8+5ft05DU8ayQcNkX
4pXJoDo154N52+T4qV4WkoFD5uL4NLPz5wxf1y\nNwcnfru8K8a2q1/9o0U+99==
-----END PRIVATE KEY-----

```

HPE 3PAR plug-in configuration notes

The Hewlett-Packard Enterprise (HPE) 3PAR plug-in lets you create and delete snapshot disks on a 3PAR Array. The plug-in supports the clone and copy-on-write (COW) snapshot types.

Note: You can restore a COW snapshot, but not a clone snapshot.

Table 5-4 HPE 3PAR plug-in configuration parameters

CloudPoint configuration parameter	Description
Array IP address	The IP address of the array.

Table 5-4 HPE 3PAR plug-in configuration parameters (*continued*)

CloudPoint configuration parameter	Description
Username	The user name to access the array.
Password	The password to access the array.

Before you configure the plug-in, ensure that the specified user account has permissions to create, delete, promote, export, and deport snapshots on the array.

See [“Hewlett Packard Enterprise \(HPE\) 3PAR array”](#) on page 188.

See [“Configuring an off-host plug-in”](#) on page 40.

Microsoft Azure plug-in configuration notes

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

Before you configure the Azure plug-in, complete the following preparatory steps:

- Use the Microsoft Azure Portal to create an Azure Active Directory (AAD) application for the Azure plug-in.
- Assign the 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>

Table 5-5 Microsoft Azure plug-in configuration parameters

CloudPoint configuration parameter	Microsoft equivalent term and description
Tenant ID	The ID of the AAD directory in which you created the application.
Client ID	The application ID.
Secret Key	The secret key of the application.

Azure plug-in considerations and limitations

The Azure plug-in has the following limitations:

- The current release of the plug-in does not support snapshots of blobs.

- CloudPoint currently only supports creating and restoring snapshots of Azure-managed disks and the virtual machines that are backed up by managed disks.

See [“Configuring an off-host plug-in”](#) on page 40.

Configuring permissions on Microsoft Azure

Before CloudPoint can protect your Microsoft Azure assets, it must have access to them. You must associate a custom role that CloudPoint users can use to work with Azure assets.

The following is a custom role definition (in JSON format) that gives CloudPoint the ability to:

- Configure the Azure plug-in and discover assets.
- Create host and disk snapshots.
- Restore snapshots to the original location or to a new location.
- Delete snapshots.

```
{ "Name": "CloudPoint Admin",
  "IsCustom": true,
  "Description": "Necessary permissions for
Azure plug-in operations in CloudPoint",
  "Actions": [
    "Microsoft.Storage/*/read",
    "Microsoft.Compute/*/read",
    "Microsoft.Compute/disks/write",
    "Microsoft.Compute/disks/delete",
    "Microsoft.Compute/images/write",
    "Microsoft.Compute/images/delete",
    "Microsoft.Compute/snapshots/delete",
    "Microsoft.Compute/snapshots/write",
    "Microsoft.Compute/virtualMachines/capture/action",
    "Microsoft.Compute/virtualMachines/write",
    "Microsoft.Compute/virtualMachines/delete",
    "Microsoft.Compute/virtualMachines/generalize/action",
    "Microsoft.Compute/virtualMachines/restart/action",
    "Microsoft.Compute/virtualMachines/runCommand/action",
    "Microsoft.Compute/virtualMachines/start/action",
    "Microsoft.Compute/virtualMachines/vmSizes/read",
    "Microsoft.Network/*/read",
    "Microsoft.Network/networkInterfaces/delete",
```

```
"Microsoft.Network/networkInterfaces/effectiveNetworkSecurityGroups/action",
"Microsoft.Network/networkInterfaces/join/action",
"Microsoft.Network/networkInterfaces/write",
"Microsoft.Network/networkSecurityGroups/join/action",
"Microsoft.Network/networkSecurityGroups/securityRules/write",
"Microsoft.Network/networkSecurityGroups/write",
"Microsoft.Network/publicIPAddresses/delete",
"Microsoft.Network/publicIPAddresses/join/action",
"Microsoft.Network/publicIPAddresses/write",
"Microsoft.Network/routeTables/join/action",
"Microsoft.Network/virtualNetworks/delete",
"Microsoft.Network/virtualNetworks/subnets/delete",
"Microsoft.Network/virtualNetworks/subnets/join/action",
"Microsoft.Network/virtualNetworks/write",
"Microsoft.Resources/*/read",
"Microsoft.Resources/subscriptions/resourceGroups/write",
"Microsoft.Resources/subscriptions/resourceGroups/ \
validateMoveResources/action",
"Microsoft.Resources/subscriptions/tagNames/tagValues/write",
"Microsoft.Resources/subscriptions/tagNames/write",
"Microsoft.Subscription/*/read",
"Microsoft.Authorization/*/read" ],
"NotActions": [ ],
"AssignableScopes": [
"/subscriptions/subscription_GUID",
"/subscriptions/subscription_GUID/ \
resourceGroups/myCloudPointGroup" ] }
```

To create a custom role using powershell, follow the steps in the following Azure documentation:

<https://docs.microsoft.com/en-us/azure/role-based-access-control/tutorial-custom-role-powershell>

For example:

```
New-AzureRmRoleDefinition -InputFile "C:\CustomRoles\ReaderSupportRole.json"
```

To create a custom role using Azure CLI, follow the steps in the following Azure documentation:

<https://docs.microsoft.com/en-us/azure/role-based-access-control/tutorial-custom-role-cli>

For example:

```
az role definition create --role-definition "~/CustomRoles/
ReaderSupportRole.json"
```

Note: Before creating a role, you must copy the role definition given earlier (text in JSON format) in a .json file and then use that file as the input file. In the sample command displayed earlier, `ReaderSupportRole.json` is used as the input file that contains the role definition text.

To use this role, do the following:

- Assign the role to an application running in the Azure environment.
- In CloudPoint, configure the Azure off-host plug-in with the application's credentials.

See [“Microsoft Azure plug-in configuration notes”](#) on page 54.

Nutanix plug-in configuration notes

The Nutanix Acropolis HyperVisor (AHV) plug-in for CloudPoint lets you do the following:

- Snapshot a combination of a virtual machine and its attached disks.
- Restore a snapshot to the original virtual machine.
- Delete a snapshot.

To configure the plug-in, specify the following parameters. They are all mandatory.

- The IP address of Nutanix AHV Prism.
- The user name to access Prism.
- The password to access Prism.

See [“Configuring an off-host plug-in”](#) on page 40.

Pure Storage FlashArray plug-in configuration notes

Table 5-6 Pure Storage FlashArray configuration parameters

CloudPoint configuration parameter	Description
IP address of Pure Storage	The IP address of the array.
Username to access Pure Storage	The username to access the array.

Table 5-6 Pure Storage FlashArray configuration parameters (*continued*)

CloudPoint configuration parameter	Description
Password to access Pure Storage	The password to access the array.

Before you configure the plug-in, ensure that the user account that you provide to CloudPoint has the permissions to perform the following operations on the assets:

- Create snapshot
- Restore snapshot
- Delete snapshot

See [“Pure Storage FlashArray”](#) on page 190.

See [“Configuring an off-host plug-in”](#) on page 40.

Huawei OceanStor array plug-in configuration notes

Table 5-7 Huawei OceanStor array plug-in configuration parameters

CloudPoint configuration parameter	Description
Array IP Address	The IP address of the array.
Username	The username to access the array.
Password	The password to access the array..

Before you configure the plug-in, ensure that the specified user account has permissions to create, delete, and restore snapshots on the array.

See [“Huawei OceanStor arrays”](#) on page 192.

See [“Configuring an off-host plug-in”](#) on page 40.

Configuring the on-host agents and plug-ins

This chapter includes the following topics:

- [About agents](#)
- [Preparing to install the Linux-based on-host agent](#)
- [Preparing to install the Windows-based on-host agent](#)
- [About the installation and configuration process](#)
- [Downloading and installing the on-host agent](#)
- [Configuring the Linux-based on-host agent](#)
- [MongoDB plug-in configuration notes](#)
- [Configuring the Windows-based on-host agent](#)
- [Configuring the Windows-based agent on a host if an agent has been previously installed](#)
- [Configuring the on-host plug-in](#)
- [Configuring VSS to store shadow copies on the originating drive](#)
- [Enabling the Microsoft SQL plug-in on the Windows host](#)
- [Running the Windows agent as a service](#)

About agents

CloudPoint agents do the following:

- Translate between the message protocol and the plug-in interface.
- Ensure secure communication between the plug-ins and the 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.

CloudPoint has an off-host agent known as parent agent that manages all configurations. Each configuration has a separate agent container which manages a particular configuration and is treated as a child agent. The child agent is also an off-host type. There can be multiple child agents for each parent agent. All the operations on the plug-in, such as GET, PUT, DELETE, work on the off-host (parent) agent.

When a new configuration is added in CloudPoint, it is added to a child agent container which handles the configuration. The new configuration starts the registration with CloudPoint and it restarts automatically when the registration is finished. During this time, the child agent goes offline and comes back online after the restart of the container is completed.

See [“About plug-ins”](#) on page 37.

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

Table 6-1 Asset types and the type of plug-ins

Asset type and vendors	On-host plug-in	Off-host plug-in
Application <ul style="list-style-type: none"> ■ Amazon Relational Database Service (RDS) applications and Aurora database clusters ■ MongoDB Enterprise Edition 3.6 ■ MSSQL 2014 and 2016 ■ Oracle 12c 	✓	X

Table 6-1 Asset types and the type of plug-ins (*continued*)

Asset type and vendors	On-host plug-in	Off-host plug-in
Supported file systems on: <ul style="list-style-type: none"> ■ Linux ■ Windows 2012 and 2016 	✓	X
Public cloud (host snapshot or disk snapshot) <ul style="list-style-type: none"> ■ Amazon Web Services (AWS) EC2 instances ■ Google Cloud Platform virtual machines ■ Microsoft Azure virtual machines ■ Nutanix Acropolis Hypervisor (AHV) 	X	✓
On-premises storage array <ul style="list-style-type: none"> ■ Dell EMC Unity arrays ■ Hewlett-Packard Enterprise (HPE) 3PAR ■ Pure Storage Flash Array 	X	✓

Preparing to install the Linux-based on-host agent

Before you install the Linux-based on-host agent, make sure that you do the following:

- Install Linux networking tools using the following command:


```
# sudo yum install -y net-tools
```
- Install the extra EPEL repositories using the following command:


```
# sudo yum install https://dl.fedoraproject.org/pub/epel/epel-release-latest-7.noarch.rpm -y
```
- Install the `python2-pika` package using the following command:


```
# sudo yum install python2-pika -y
```
- Install the Open SSL version 1.0.2k or higher using the following command:


```
# sudo yum update -y openssl
```
- If you are installing the Linux-based agent to discover Oracle applications, optimize your Oracle database files and metadata files.
 See [“Optimizing your Oracle database data and metadata files”](#) on page 62.
 See [“About the installation and configuration process”](#) on page 63.

Optimizing your Oracle database data and metadata files

Veritas recommends that you do not keep the Oracle configuration files on a boot or a root disk. Use the following information to know more about how to move those files and optimize your Oracle installation.

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 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 that 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

You can configure the Oracle plug-in to discover and protect your Oracle database applications with disk-level and host-level snapshots.

Before you configure the Oracle plug-in, make sure that your environment meets the following requirements:

- A supported version of Oracle is installed in a supported Red Hat Enterprise Linux (RHEL) host environment.
See “[Meeting system requirements](#)” on page 19.
- Oracle standalone instance is discoverable.
- Oracle binary and Oracle data must be on separate volumes.
- Log archiving is enabled.
- Oracle listener is enabled.
- `tnsping` connection test to the database instance works.
- The `db_recovery_file_dest_size` parameter size is set as per Oracle recommendation.

Refer to the Oracle documentation for more information:

https://docs.oracle.com/cd/B19306_01/backup.102/b14192/setup005.htm

- The databases are running, mounted, and open.
- The databases are not in backup mode.
CloudPoint discovers databases that are in a backup mode, but snapshot operations are not supported.

Preparing to install the Windows-based on-host agent

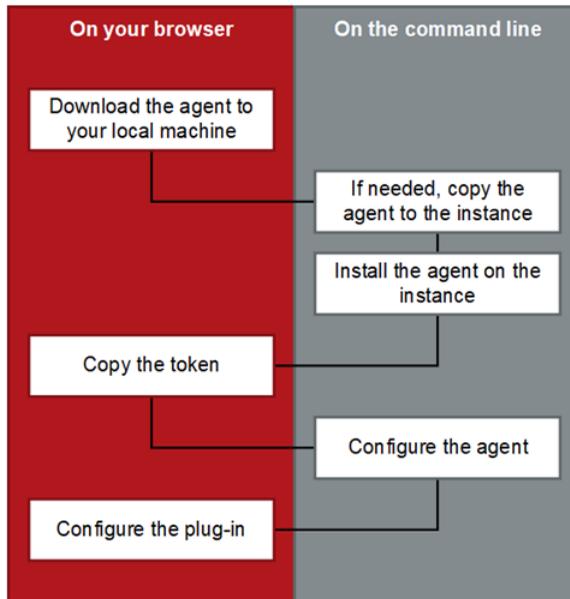
Before you install the Windows-based on-host agent, do the following on the Windows host:

- Enable port 5671 (both inbound and outbound)
- Connect to the host through Remote Desktop

See [“About the installation and configuration process”](#) on page 63.

About the installation and configuration process

To install and configure an on-host agent and plug-in, you perform tasks from the CloudPoint user interface in your browser and on the command line of your local computer or instance.

Figure 6-1 CloudPoint on-host agent installation and configuration process

See [“Preparing to install the Linux-based on-host agent”](#) on page 61.

See [“Preparing to install the Windows-based on-host agent”](#) on page 63.

See [“Downloading and installing the on-host agent”](#) on page 64.

Downloading and installing the on-host agent

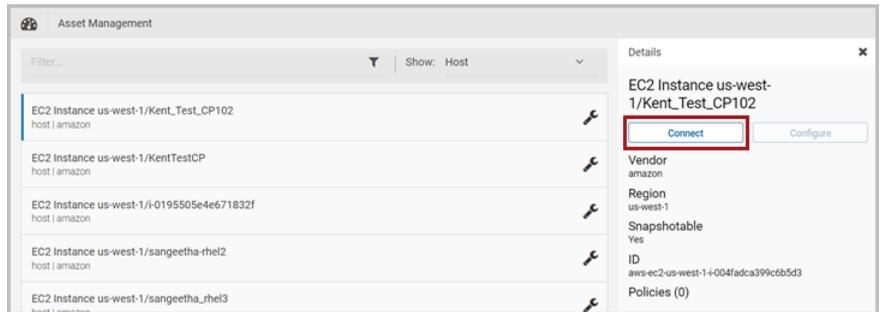
Before you complete the steps in this section, do the following:

- Make sure you have an appropriate CloudPoint license installed. The CloudPoint on-host agents are not supported with the CloudPoint Freemium license. See [“Understanding your CloudPoint license”](#) on page 11.
- Make sure you have administrative privileges on the host on which you want to install the on-host agent.
- Complete the preparatory steps and install all the dependencies for your particular agent.
See [“Preparing to install the Linux-based on-host agent”](#) on page 61.
See [“Preparing to install the Windows-based on-host agent”](#) on page 63.

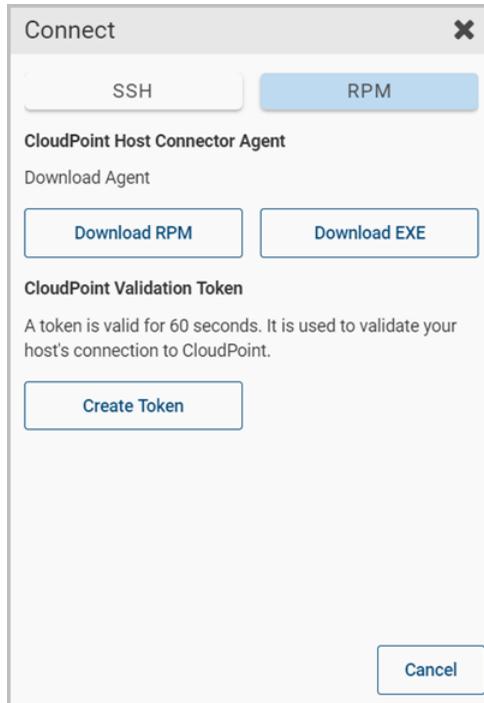
Whether you install the Linux-based on-host agent or the Windows-based on-host agent, the steps are similar.

To download and install the on-host agent

- 1 Sign in to the CloudPoint user interface (UI).
 See “[Signing in to CloudPoint](#)” on page 101.
- 2 Click **Dashboard** and from under the **Environment** section, locate the **Hosts** area, and click **Manage**.
- 3 On the **Asset Management** page, select the host on which you want to install an agent and then from the **Details** panel on the right, click **Connect**.



- 4 On the **Connect** dialog box, make sure the **RPM** tab is selected and then do the following:
 - To download the Linux-based agent, click **Download RPM**.
 - To download the Windows-based agent, click **Download EXE**.



Do not close this Connect dialog box as yet. When you configure the agent, you will return to this dialog box to get a token.

Note: The agent software download options are also available in the **Settings** (gear icon) menu from the top right corner of the user interface (UI).

- 5 If necessary, copy the downloaded agent package to the instance on which you want to run the package.
 - For the Linux-based agent, use the SCP utility to copy the package.
 - For the Windows-based agent, copy the package to `C:\Program Files\Veritas\CloudPoint` directory.
You may have to create the directory if it does not exist already.
- 6 Do the following:
 - For the Linux-based agent, type the following command:


```
# sudo rpm -ivh cloudpoint_agent_rpm_name
```

 Here, *<cloudpoint_agent_rpm_name>* is the name of the on-host agent rpm package you downloaded earlier.

For example:

```
# sudo rpm -ivh VRTScloudpoint-agent-2.1.2-RHEL7.x86_64.rpm
```

If you are re-installing the on-host agent on the instance, use the following command to upgrade the RPM:

```
# sudo rpm -Uvh CloudPoint_agent_RPM_name
```

- For the Windows-based agent, unzip the agent installer package zip file that you downloaded from the CloudPoint UI earlier.

Do not run the extracted agent package yet. You will do that during the on-host agent configuration later.

- 7 Proceed to configure the on-host agent.

See [“Configuring the Linux-based on-host agent”](#) on page 67.

See [“Configuring the Windows-based on-host agent”](#) on page 70.

Configuring the Linux-based on-host agent

Before you configure the Linux-based on-host agent, make sure you have downloaded and installed the agent on the Linux instance.

See [“Downloading and installing the on-host agent”](#) on page 64.

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

To configure the Linux-based on-host agent

- 1 If the CloudPoint Linux-based on-host agent was already configured on the host earlier, remove the `/opt/VRTScloudpoint/keys` directory.

Type the following command on that host where the agent runs:

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

- 2 Type the following command in the `/etc` directory to create a configuration file called `flexsnap.conf`.

```
# sudo vi /etc/flexsnap.conf
```

- 3 Add the following lines to the file and save it.

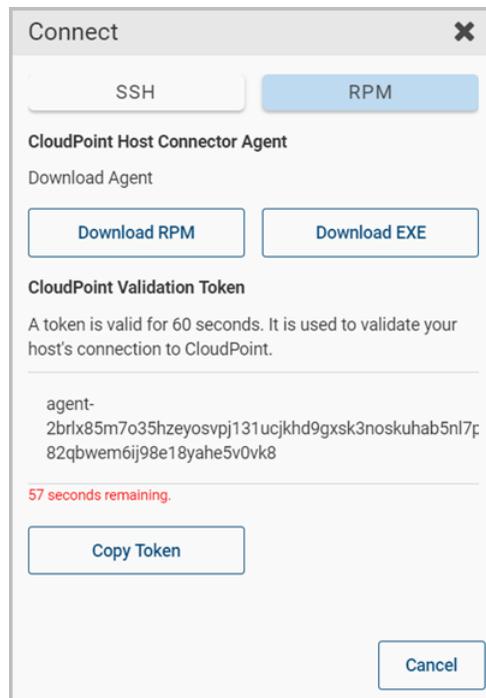
```
[global]  
target = CloudPoint_host_FQDN_or_IPaddress
```

Note: The public IP might change whenever an instance is stopped and started again. If you have added the IP address as a target (this step), then ensure that every time the IP address changes, you replace the IP address entry in the `flexsnap.conf` file and then start `flexsnap-agent` again.

- 4 On the CloudPoint dashboard, return to the **Connect** dialog box, or if you closed the dialog box, do the following:
 - On the dashboard, under the **Environment** section, locate the **Hosts** area, and then click **Manage**.
 - On the **Asset Management** page, select the host and from the **Details** panel on the right, click **Connect**.
- 5 On the **Connect** dialog box, on the **RPM** tab, click **Create Token**.

CloudPoint generates a unique sequence of alpha-numeric characters that are used as an authentication token to authorize the host connection with CloudPoint.

- 6 Click **Copy Token**.



Note: The token is valid for 60 seconds only. If you do not copy the token within that time frame, generate a new token again.

- 7 Type the following command to copy the token and start the `flexsnap-agent`:

```
# sudo flexsnap-agent copied_token
```

Note: If you encounter an error, check the logs at `/var/log/flexsnap/flexsnap-agent-onhost.log` to troubleshoot the issue.

- 8 Type the following command to enable the agent service:

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

- 9 Type the following command to start the agent service:

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

- 10 Proceed to configure the on-host plug-in.

See “[Configuring the on-host plug-in](#)” on page 73.

MongoDB plug-in configuration notes

Beginning with CloudPoint release 2.0.1, you can configure a MongoDB plug-in to discover and protect your MongoDB database applications with disk-level and host-level snapshots.

Before you configure the MongoDB plug-in, make sure that your environment meets the following requirements:

- The Linux on-host must be installed and running in a Red Hat Enterprise Linux (RHEL) 7.4 environment.
- You must be running MongoDB enterprise 3.6.
- Discovery of a MongoDB standalone instance is supported.
- Databases and journals must be stored on the same volume.
- If you want to create application-consistent snapshots, journaling must be turned on.

Have the following information ready when you configure the plug-in:

Table 6-2 Configuration parameters for MongoDB plug-in

CloudPoint configuration parameter	Description
MongoDB configuration file path	The location of the MongoDB <code>conf</code> file.
MongoDB admin user name	A MongoDB user name with administrator privileges.

Table 6-2 Configuration parameters for MongoDB plugin (*continued*)

CloudPoint configuration parameter	Description
MongoDB admin user password	The password of the MongoDB admin user account.

Note: `PyMongo` is a Python distribution that is used to work with MongoDB. During configuration, when the plug-in tries to load `pymongo` for the first time, the Linux on-host agent crashes. Restart the on-host agent. You can then configure the MongoDB plug-in successfully and begin to take snapshots.

Configuring the Windows-based on-host agent

This section describes how to configure the Windows-based agent on a host for the first time. If the host you are using has had an agent installed on it before, the configuration steps are slightly different.

See [“Configuring the Windows-based agent on a host if an agent has been previously installed”](#) on page 72.

Before you complete the steps in this section, make sure you have downloaded and installed the agent.

See [“Downloading and installing the on-host agent”](#) on page 64.

To complete the steps in this section, you need administrative privileges on the Windows instance.

To configure the Windows-based on-host agent

- 1 On the host that runs the agent, create a configuration file, `flexsnap.conf`. Navigate to `C:\ProgramData\Veritas\Cloudpoint\etc` and enter the following:

```
copy nul flexsnap.conf
```

- 2 Using Notepad, open `flexsnap.conf`, add the following lines, and save the file:

```
[global]
target = CloudPoint_HOST_FQDN_or_IPaddress
```

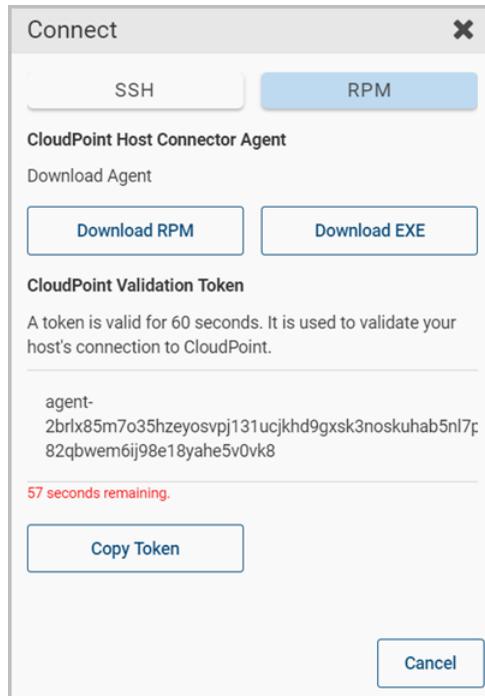
Note: The public IP might change whenever an instance is stopped and started again. If you have added the IP address as a target (this step), then ensure that every time the IP address changes, you replace the IP address entry in the `flexsnap.conf` file and then start `flexsnap-agent` again.

- 3 On the CloudPoint dashboard, return to the **Connect** dialog box or if you closed the dialog box, do the following:
 - On the dashboard, under the Environment section, locate the **Hosts** area, and then click **Manage**.
 - On the **Asset Management** page, select the host and from the Details panel on the right, click **Connect**.

- 4 On the **Connect** dialog box, on the **RPM** tab click **Create Token**.

CloudPoint generates a unique sequence of alpha-numeric characters that are used as an authentication token to authorize the host connection with CloudPoint.

- 5 Click **Copy Token**.



Note: The token is valid for 60 seconds only. If you do not copy the token within that time frame, generate a new token again.

- 6 Copy the token and start the `flexsnap-agent`.

On the Windows instance, on the command prompt, navigate to the directory where you extracted the on-host agent package .zip file, typically `C:\Program Files\Veritas\CloudPoint\`, and then type the following command:

```
flexsnap-agent_name.exe jointoken
```

- 7 Run the same .exe file without any arguments.

```
flexsnap-agent_name.exe
```

- 8 Proceed to configure the on-host plug-in.

See [“Configuring the on-host plug-in”](#) on page 73.

Configuring the Windows-based agent on a host if an agent has been previously installed

If the Windows-based on-host agent has been configured on a host before, the configuration steps are slightly different from a fresh configuration.

To configure the Windows-based agent on a host if an agent has been previously installed

- 1 Navigate to `C:\Program Files\Veritas\CloudPoint` and delete the unzipped exe folder.

Even if you do not remove the folder, remember to execute the `flexsnap-agent_name.exe` command from the latest .exe file.

- 2 If not done already, download the agent package file and unzip this file to `C:\Program Files\Veritas\CloudPoint`.

See [“Downloading and installing the on-host agent”](#) on page 64.

- 3 On the Windows instance, edit the configuration file located at
`C:\ProgramData\Veritas\CloudPoint\etc\flexsnap.conf`.

The previous installation of the on-host agent added extra lines to this file. Remove those lines and add or edit following.

```
[global]
target = CloudPoint_HOST_FQDN_or_IPaddress
```

Note: The public IP might change whenever an instance is stopped and started again. If you have added the IP address as a target (this step), then ensure that every time the IP address changes, you replace the IP address entry in the `flexsnap.conf` file and then start `flexsnap-agent` again.

- 4 From the **Connect** dialog box on the CloudPoint user interface copy the token.
- 5 Copy the token and start the `flexsnap-agent`.
On the Windows instance, from the command prompt, navigate to the directory where you extracted the .zip file and then enter the following:

```
flexsnap-agent_name.exe jointoken
```
- 6 Run the same .exe file without any arguments.

```
flexsnap-agent_name.exe
```
- 7 Proceed to configure the on-host plug-in.
See [“Configuring the on-host plug-in”](#) on page 73.

Configuring the on-host plug-in

After installing and registering the on-host agent on the host, the next step is to configure the on-host plug-in on the host.

Before you proceed, ensure you have configured the on-host agent.

See [“Configuring the Linux-based on-host agent”](#) on page 67.

See [“Configuring the Windows-based on-host agent”](#) on page 70.

To configure an on-host plug-in

- 1 Review the configuration requirements for the on-host plug-in you want to configure.

See “[Oracle plug-in configuration notes](#)” on page 62.
See “[MongoDB plug-in configuration notes](#)” on page 69.
- 2 After you configure the on-host agent, return to the CloudPoint user interface and select the asset on which you installed and configured the on-host agent.

On the **Details** page, observe that the **Configuration** button is enabled.
- 3 Click **Configuration**.
- 4 From the drop-down list, select the on-host plug-in that you want to configure.

For example, if you want to configure the CloudPoint plug-in for Microsoft SQL, choose **MSSQL Database**.
- 5 Click **Configure**.

After a few minutes, the statistics on the CloudPoint dashboard are automatically updated to indicate all the new assets that are discovered. You can view these assets by clicking the **Manage** link from under the **Applications** or the **File Systems** widgets.

For example, if you have configured the SQL plug-in, the Asset Management page displays the SQL Server database instances running on the hosts where you configured the plug-in. You can select these assets and perform snapshot operations on them.

Configuring VSS to store shadow copies on the originating drive

If you want to take disk-level, application-consistent Windows snapshots of a Windows file system or SQL application, you must configure Microsoft Volume Shadow Copy Service (VSS). VSS lets you take volume snapshots while applications continue to write to the volume.

When you configure VSS, keep in mind the following;

- CloudPoint currently has a limitation that you must manually configure the shadow copy creation location to the same drive or volume as the originating drive. This approach ensures that an application-consistent snapshot is created.
- If shadow storage already exists on an alternate drive or a dedicated drive, you must disable that storage and replace it with the configuration in the following procedure.

To configure VSS to store shadow copies on the originating drive

1. On the Windows host, open the command prompt. If User Account Control (UAC) setting is enabled on the server, launch the command prompt in the **Run as administrator** mode.
2. For each drive letter on which you want to take disk-level, application-consistent snapshots using CloudPoint, enter a command similar to the following:

```
vssadmin add shadowstorage /for=<drive being backed up> ^  
/on=<drive to store the shadow copy> ^  
/maxsize=<percentage of disk space allowed to be used>
```

Here, `maxsize` represents the maximum free space usage allowed on the shadow storage drive. The caret (^) character in the command represents the Windows command line continuation character.

For example, if the VSS shadow copies of the `D:` drive are to be stored on the `D:` drive and allowed to use up to 80% of the free disk space on `D:`, the command syntax is as follows:

```
vssadmin add shadowstorage /for=d: /on=d: /maxsize=80%
```

The command prompt displays a message similar to the following:

```
Successfully added the shadow copy storage association
```

3. Verify your changes using the following command:

```
vssadmin list shadowstorage
```

Enabling the Microsoft SQL plug-in on the Windows host

The Microsoft SQL (MS SQL) on-host plug-in lets you create disk-level and host-level snapshots of your Microsoft SQL application. When you use this plug-in, keep in mind the following:

- This plug-in is supported in Azure and AWS environments, but not in Google Cloud Platform or VMware environments.
- If you want to discover SQL applications, you cannot run the Windows agent as a service.
- SQL Server restore is not supported in this release. However, you can manually restore the SQL Server shadow copy.

To enable the SQL plug-in on the Windows host

- 1 On the CloudPoint dashboard, in the **Hosts** widget, click **Manage**.
- 2 On the **Asset Management** page, find and select the Windows host.
- 3 Click **Configure** and select the MS SQL plug-in from the drop-down list.
- 4 Return to the dashboard.
- 5 In the **Applications** widget, click **Manage**.

The **Asset Management** page lists the Microsoft SQL databases on the Windows host. If the databases are not displayed, wait for a minute and refresh your browser.

Running the Windows agent as a service

Note: If you want to discover SQL applications, you cannot run the Windows agent as a service. If you want to discover SQL applications, you must run the `flexsnap-agent.exe` executable from a command prompt that is running with `run as administrator` rights.

To run the Windows agent as a service

- 1 Make sure that the `flexsnap-agent.exe` process is not running. If it is, press `CTRL+C` in the command prompt to stop it.
- 2 Verify that the `flexsnap-agent.exe` is not running in memory. Open the **Task Manager** check the **Processes** tab.

- 3 Open a command prompt. If User Account Control is enabled, enter the following command with `run as administrator` rights. The caret (^) is in the Windows command line continuation character.

```
cd C:\Program Files\Veritas\CloudPoint\ ^  
flexsnap-windows-svc.exe --startup=delayed install
```

If you want to run the service under a domain or other (non-system) account, use the following command instead:

```
cd C:\Program Files\Veritas\CloudPoint\ ^  
flexsnap-windows-svc.exe --username=DOMAIN\username ^  
--password=password --startup=delayed install
```

- 4 Start the service. Enter the following:

```
sc start CloudPointService
```

If the operation succeeds, the Windows Task Manager displays the following processes:

```
flexsnap-agent.exe  
flexsnap-windows-svc.exe (x2)
```

Protecting assets with CloudPoint's agentless feature

This chapter includes the following topics:

- [About the agentless feature](#)
- [Prerequisites for the agentless configuration](#)
- [Configuring the agentless feature](#)

About the agentless feature

If you want CloudPoint to discover and protect on-host assets, but you want to minimize the vendor software footprint on your hosts, consider CloudPoint's agentless feature. Typically, when you use an agent, the software remains on the host at all times. In contrast, the agentless feature works as follows:

- The CloudPoint software accesses the host through SSH.
- CloudPoint performs the specified task, such as creating a snapshot.
- When the task completes, CloudPoint software deletes itself from the host.

The CloudPoint agentless feature currently discovers and operates on Linux file system assets, Oracle database, and MongoDB database assets.

See [“Prerequisites for the agentless configuration”](#) on page 79.

See [“Configuring the agentless feature”](#) on page 80.

Prerequisites for the agentless configuration

Verify the following before you configure the agentless feature:

- The agentless feature is available only with the CloudPoint Enterprise or an equivalent license. Ensure that you have installed the appropriate license. See [“Upgrading your CloudPoint license”](#) on page 180.
- Have the following information with you:
 - Host user name
 - Host password or SSH keyCloudPoint requires these details to gain access to the host and perform requested operations.
- On hosts where you wish to configure this feature, grant password-less sudo access to the host user account that you provide to CloudPoint. See [“Granting password-less sudo access to host user account”](#) on page 79.

Granting password-less sudo access to host user account

CloudPoint requires a host user account to connect and perform operations on the host. You must grant password-less sudo access to the user account that you provide to CloudPoint. This is required for all the hosts where you wish to configure the agentless feature.

Note: The following steps are provided as a general guideline. Refer to the operating system or the distribution-specific documentation for detailed instructions on how to grant password-less sudo access to a user account.

Perform the following steps on a host where you want to configure the agentless feature

1. Verify that the host user name that you provide to CloudPoint is part of the `wheel` group.

Log on as a root user and run the following command:

```
# usermod -aG wheel hostuserID
```

Here, *hostuserID* is the host user name that you provide to CloudPoint.

2. Log out and log in again for the changes to take effect.
3. Edit the `/etc/sudoers` file using the `visudo` command:

```
# sudo visudo
```

4. Add the following entry to the `/etc/sudoers` file:

```
hostuserID ALL=(ALL) NOPASSWD: ALL
```

5. In the `/etc/sudoers` file, edit the entries for the `wheel` group as follows:

- Comment out (add a `#` character at the start of the line) the following line entry:

```
# %wheel ALL=(ALL) ALL
```

- Uncomment (remove the `#` character at the start of the line) the following line entry:

```
%wheel ALL=(ALL) NOPASSWD: ALL
```

The changes should appear as follows:

```
## Allows people in group wheel to run all commands
```

```
# %wheel ALL=(ALL) ALL
```

```
## Same thing without a password
```

```
%wheel ALL=(ALL) NOPASSWD: ALL
```

6. Save the changes to the `/etc/sudoers` file.
7. Log out and log on to the host again using the user account that you provide to CloudPoint.
8. Run the following command to confirm that the changes are in effect:

```
# sudo su
```

If you do not see any prompt requesting for a password, then the user account has been granted password-less sudo access.

You can now proceed to configure the CloudPoint agentless feature.

Configuring the agentless feature

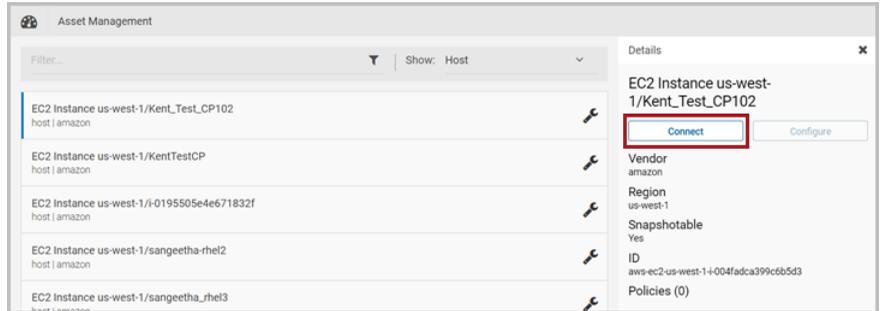
Verify all the prerequisites before you configure the CloudPoint agentless feature.

See [“Prerequisites for the agentless configuration”](#) on page 79.

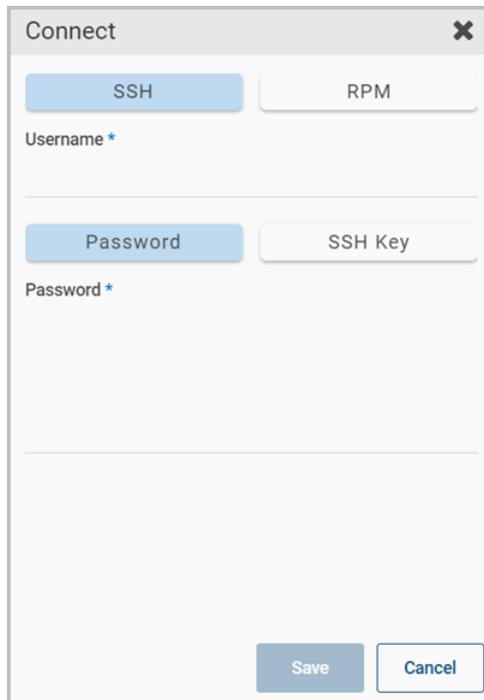
To configure the agentless feature

- 1 On the CloudPoint dashboard, in the **Environment** card, locate the **Hosts** area, and click **Manage**.
- 2 On the **Asset Management** page, select the host on which you want to use the agentless feature.

- 3 On the **Details** page, click **Connect**



- 4 On the **Connect** dialog box, select the **SSH** chip.



- 5 Enter the SSH user name, and either the SSH password or SSH key.
- 6 Click **Save**.

Configuring users

- [Chapter 8. Setting up email and adding users](#)
- [Chapter 9. Assigning roles to users for greater efficiency](#)

Setting up email and adding users

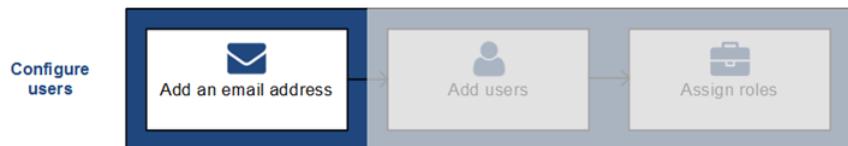
This chapter includes the following topics:

- [Configuring the CloudPoint sender email address](#)
- [About adding users to CloudPoint](#)
- [Adding AD users to CloudPoint using LDAP](#)
- [Adding users to CloudPoint manually](#)
- [Deleting a user from CloudPoint](#)

Configuring the CloudPoint sender email address

The following figure shows where you are in the CloudPoint user configuration process.

Figure 8-1 You are here in the user configuration process



To add users to the CloudPoint configuration, you must first configure a sender email address. The sender email is used as a source address for sending all CloudPoint communications.

CloudPoint sends emails for the following events:

- Whenever a new user is added to the CloudPoint configuration, CloudPoint sends an email that contains the user name and a temporary password. This email is sent from the configured sender email address to the email address that is associated with the new user account.
- If CloudPoint users forget their CloudPoint sign-in password, users can request for a password reset using the **Forgot password?** link on the CloudPoint UI sign-in page. CloudPoint then sends an email containing a new temporary password. This email is sent from the configured sender email address to the email address that is associated with that user account.

You can configure the CloudPoint sender email address using any of the following email services:

- Amazon Simple Email Service (SES)
- SendGrid email delivery service
- Simple Mail Transfer Protocol (SMTP)

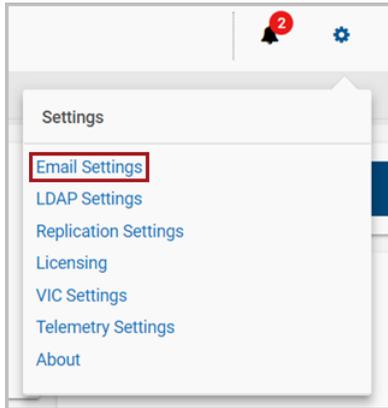
Before you configure the sender email address, gather the following information based on the email service you wish to use. You will specify this information during the actual configuration process.

Table 8-1 Email configuration parameters

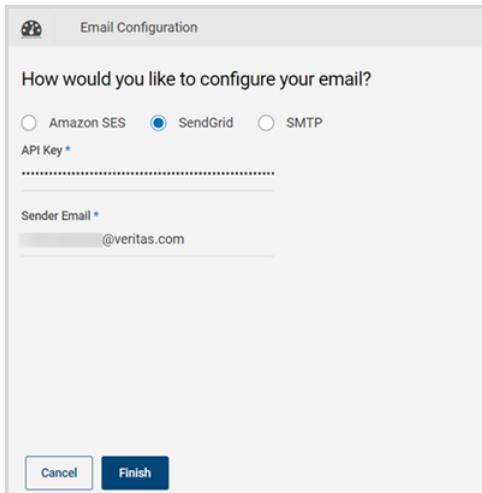
Email service	Required parameters
Amazon SES	<ul style="list-style-type: none"> ■ Region ■ Access Key ■ Secret Key ■ Sender Email
SendGrid	<ul style="list-style-type: none"> ■ API Key ■ Sender Email
SMTP	<ul style="list-style-type: none"> ■ SMTP Host ■ SMTP Port ■ User name ■ Password ■ Sender Email <p>Note: CloudPoint also supports anonymous authentication using SMTP.</p>

To configure the CloudPoint sender email address

- 1 Sign in to the CloudPoint UI and from the top right corner, click **Settings** (gear icon) and then click **Email Settings**.



- 2 On the **Email Configuration** page, select the email service to use.

A screenshot of the 'Email Configuration' page in the CloudPoint UI. The page title is 'Email Configuration'. Below the title, it asks 'How would you like to configure your email?'. There are three radio button options: 'Amazon SES', 'SendGrid' (which is selected), and 'SMTP'. Below these options are two input fields: 'API Key *' and 'Sender Email *'. The 'Sender Email *' field contains the text '@veritas.com'. At the bottom of the page, there are two buttons: 'Cancel' and 'Finish'.

- 3 Complete the form by filling in the email service-specific parameters you gathered earlier.

If you use SMTP, you can specify whether you wish to send emails anonymously. The **Authentication Required** checkbox controls whether or not you wish show the sender email address. When the checkbox is selected (default value), the sender email address is displayed in all outgoing emails.

- 4 Click **Finish**.

If you use the Amazon SES or the SendGrid service, you may have to verify the email address that you specified in the form. SES or SendGrid sends a verification email to the email address associated with the CloudPoint administrator. Click the link specified in that email address to confirm the user. Upon confirmation, the specified email address is automatically configured as the CloudPoint sender email address.

About adding users to CloudPoint

The following options are available for adding users to the CloudPoint configuration:

- Add users from an Active Directory (AD) data store using the Lightweight Directory Access Protocol (LDAP)
This method allows you to add AD users to the CloudPoint configuration using LDAP. This is a two-step process wherein you first add the LDAP configuration details to CloudPoint and then manually add the AD users.
The following conditions apply:
 - You cannot import AD users using LDAP over Secure Sockets Layer (SSL). CloudPoint does not support LDAP over SSL.
 - You cannot auto-import LDAP users in to CloudPoint.
- Add users in CloudPoint manually
This method is used to add local as well as AD user accounts individually. For local (non-AD) users, CloudPoint sends a temporary password to each user account. Users can use that password to sign-in to the CloudPoint UI.
You use the same process to add AD users. The only difference is that CloudPoint does not send a separate password to AD users. Users can use their AD credentials to sign-in to CloudPoint.

Before you begin adding users to CloudPoint, ensure that you gather the following information:

Table 8-2 User addition methods and required information

Configuration method	Information to gather
Import from LDAP	<ul style="list-style-type: none">■ The name and password of the LDAP administrator account■ The LDAP base domain■ The LDAP URL■ The network port used by the LDAP server■ The search base that is used for LDAP searches■ The LDAP email domain
Create Local Users	For each user you want to add, obtain the following information: <ul style="list-style-type: none">■ Email address■ First and last name

See [“Adding AD users to CloudPoint using LDAP”](#) on page 87.

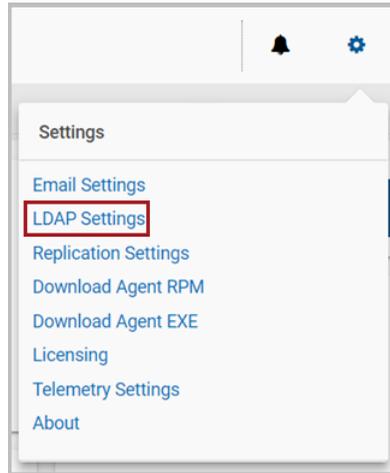
See [“Adding users to CloudPoint manually”](#) on page 89.

Adding AD users to CloudPoint using LDAP

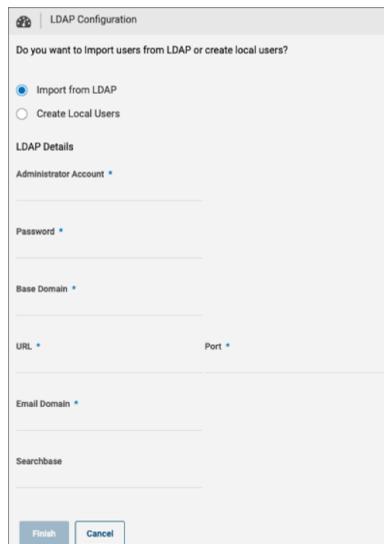
Use the following procedure to first import LDAP configuration details into CloudPoint and then proceed to manually adding AD users.

To add AD users using LDAP

- 1 Sign in to the CloudPoint UI and from the top right corner, click **Settings** (gear icon) and then click **LDAP settings**.



- 2 On the **LDAP Configuration** page, select **Import from LDAP**.

A screenshot of the LDAP Configuration page. The page title is "LDAP Configuration". Below the title, there is a question: "Do you want to Import users from LDAP or create local users?". There are two radio button options: "Import from LDAP" (which is selected) and "Create Local Users". Below this, there is a section titled "LDAP Details" with several input fields: "Administrator Account", "Password", "Base Domain", "URL", "Port", "Email Domain", and "Searchbase". At the bottom of the form, there are two buttons: "Finish" and "Cancel".

- 3 Complete the page by filling in the information that you gathered earlier.
- 4 Click **Finish**.

- 5 On the **Changing LDAP Setting** dialog box, click **Proceed**.
- 6 Proceed to adding the AD users manually.
See [“Adding users to CloudPoint manually”](#) on page 89.

Adding users to CloudPoint manually

The following figure shows where you are in the CloudPoint user configuration process.

Figure 8-2 You are here in the user configuration process

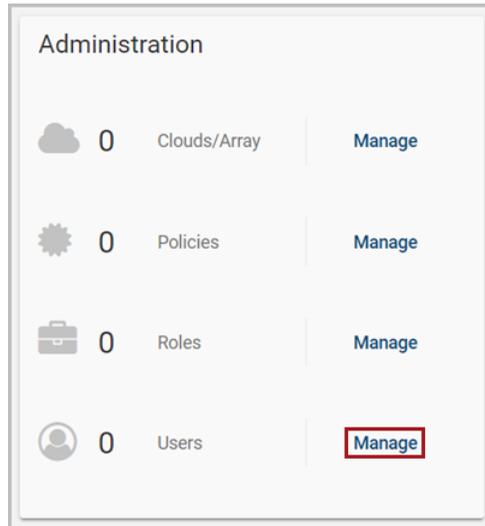


Use this procedure to add local users as well as AD users to the CloudPoint configuration. Before you proceed, ensure that you have configured a sender email address. This is the address that is used to send all CloudPoint related emails.

See [“Configuring the CloudPoint sender email address”](#) on page 83.

To add a CloudPoint user manually

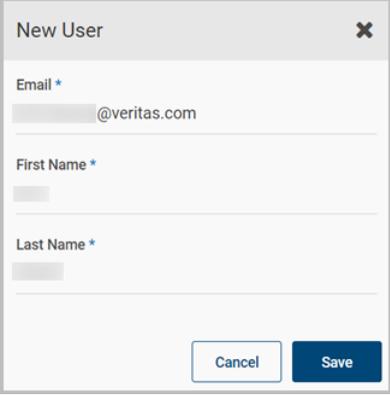
- 1 On the CloudPoint dashboard, in the **Administration** card, locate **Users**, and click **Manage**.



The User Management page displays all the users that exist in the CloudPoint configuration.

- 2 On the User Management page, click **New User**.

- 3 On the New User dialog box, specify all the requested details and then click **Save**.



The image shows a 'New User' dialog box with the following fields and buttons:

- Email ***: Input field containing '@veritas.com'
- First Name ***: Empty input field
- Last Name ***: Empty input field
- Buttons**: 'Cancel' and 'Save' buttons at the bottom right.

Go to the User Management page and verify that the user has been added successfully.

The user receives an email that they have been added to CloudPoint. The email also includes a temporary password they can use to sign-in to the CloudPoint UI.

In case of AD users, the email does not include a separate temporary password; users can use their AD password for authentication.

Note the following:

The user addition email is sent from the CloudPoint sender email address that you configured earlier. If the sender email address is configured using Amazon SES service, then you may have to verify the user email address that you just added. The user is added to CloudPoint only after a successful verification.

This is required if the Amazon SES account is placed in a sandbox environment. Refer to the following for more information:

<https://docs.aws.amazon.com/ses/latest/DeveloperGuide/request-production-access.html>

Deleting a user from CloudPoint

To delete a CloudPoint user

- 1 On the dashboard, in the **Administration** widget, locate **Users**, and click **Manage**.
- 2 On the **User Details** page, click **Delete**.

- 3** On the **Please confirm ...** dialog box, click **Delete**.
CloudPoint displays a message that the user has been removed.
- 4** On the **LDAP Users** page, verify that the user is no longer displayed.

Assigning roles to users for greater efficiency

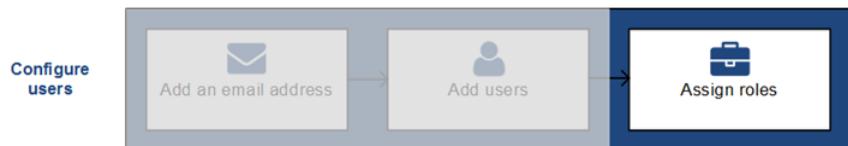
This chapter includes the following topics:

- [About role-based access control](#)
- [Displaying role information](#)
- [Creating a role](#)
- [Editing a role](#)
- [Deleting a role](#)

About role-based access control

The following figure shows where you are in the CloudPoint user configuration process.

Figure 9-1 You are here in the user configuration process



If your organization uses CloudPoint to manage a large number of assets or asset types, it may not be practical to have one CloudPoint admin account.

CloudPoint offers role-based access control which lets the administrator assign a user certain assets and privileges. With this feature, you can do the following:

- Delegate certain tasks to the people with the most expertise.

- Have multiple people in a role so there is no single point of failure.
- Control access for multiple users simultaneously.
- Clearly define ownership of assets for users.

See [“What kinds of assets can you protect?”](#) on page 10.

Displaying role information

To display role information

- 1 On the dashboard, in the **Administration** widget, locate **Roles**, and click **Manage**.
- 2 On the **Roles** page, select the check box for the role you want to view.
You can also use the **Roles** page to create a new role.
See [“Creating a role”](#) on page 94.
- 3 Review the **Role Details** page. It includes the following tabs:

Tab	Description
Users	The users who can perform this role.
Permissions	One or more sets of permissions that define the tasks users can perform.
Assets	The assets that are associated with the role.

You can also use the **Role Details** page to edit or delete the role.

See [“Editing a role”](#) on page 98.

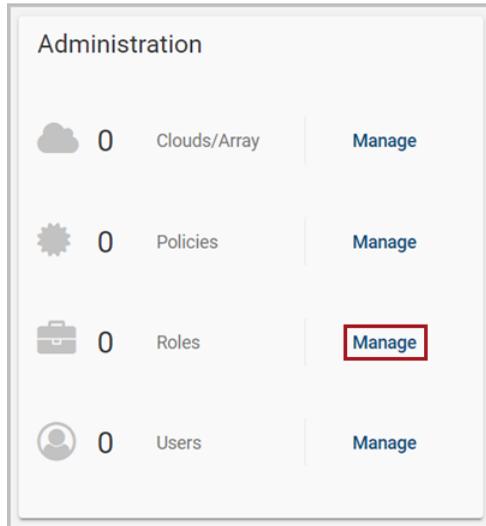
See [“Deleting a role”](#) on page 99.

Creating a role

Only the CloudPoint admin or a user with **Role management** permission can create a role.

To create a role

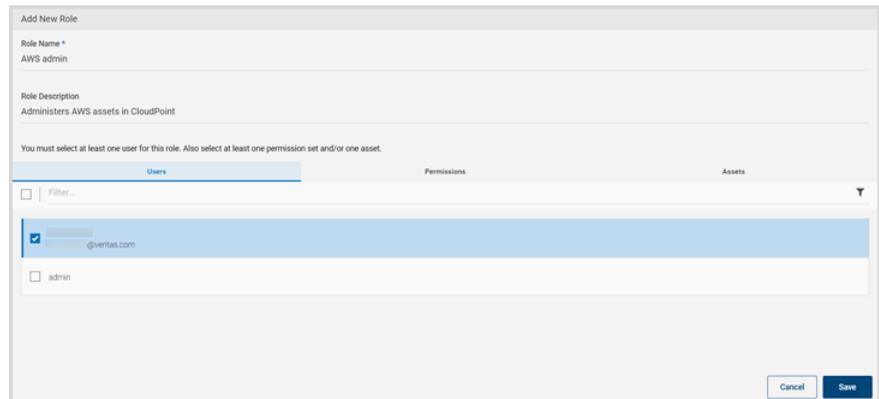
- 1 On the dashboard, in the **Administration** card, locate **Roles**, and click **Manage**.



- 2 On the **Role Management** page, click **New Role**.
- 3 On the **New Role** page, specify the name of the new role, and optionally give it a description.
- 4 Select information from the following tabs:
 - **Users**

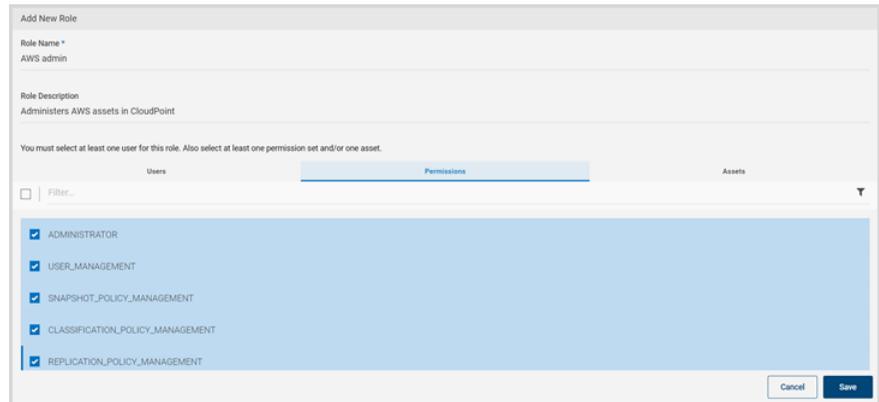
▪ **Users**

This tab displays a list of CloudPoint users and their email addresses. To assign a user to the role, select the corresponding check box. Select one or more users.



▪ **Permissions**

This tab displays a list of preconfigured permissions. Select one or more permissions.



The screenshot shows the 'Add New Role' configuration page. At the top, the role name is 'AWS admin' and the description is 'Administers AWS assets in CloudPoint'. Below this, there are three tabs: 'Users', 'Permissions', and 'Assets'. The 'Permissions' tab is selected, showing a list of five preconfigured permissions, each with a checked checkbox: ADMINISTRATOR, USER_MANAGEMENT, SNAPSHOT_POLICY_MANAGEMENT, CLASSIFICATION_POLICY_MANAGEMENT, and REPLICATION_POLICY_MANAGEMENT. At the bottom right of the permissions list, there are 'Cancel' and 'Save' buttons.

■ **Assets**

The left side of this tab displays a list of all available CloudPoint assets. The right side displays the assets that are assigned to the role. When you first assign assets to a role, the right side of the tab is blank.

Note: As the CloudPoint admin, you see all assets, regardless of whether they are appropriate for the permissions you set. The asset list is not automatically filtered based on the permission you select. If you are a non-admin user with **Role management** permission, you only see the assets assigned to you.

In the available list, select assets you want to add to the role, and click **Assigned Selected**. You can also use the buttons **Assign Selected**, **Assign All**, **Remove All**, and **Remove Selected** to create your assigned asset list.

Add New Role

Role Name *
AWS admin

Role Description
Administers AWS assets in CloudPoint

You must select at least one user for this role. Also select at least one permission set and/or one asset.

Available Assets

Filter...

- EBS Snapshot snap-000008d7349d9/936
- EBS Snapshot snap-00005302e8a42eb67
- EBS Snapshot snap-0000f5bdc6aa43653
- EBS Snapshot snap-000174b6c68db4733

Permissions

Assigned Assets

Filter...

Assign Selected
Assign All
Remove All
Remove Selected

Cancel Save

At a minimum, you must specify the following:

- One user and one permission
- One user and one asset
- One user, one permission, and one asset

5 Click **Save**.

CloudPoint displays a message that the role is added.

6 Note the new entry on the **Role Management** page.

Role Management

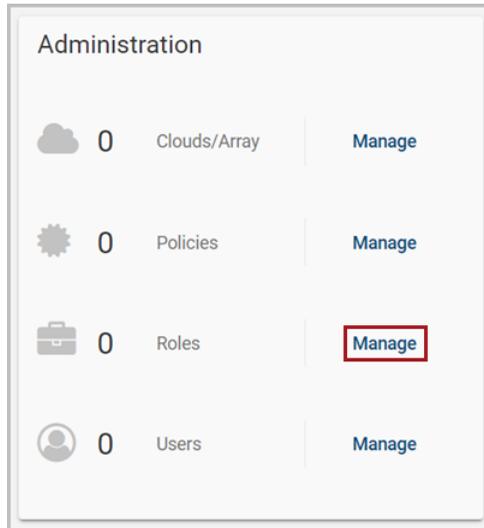
Filter...

<input type="checkbox"/>	AWS admin Administers AWS assets in CloudPoint
--------------------------	---

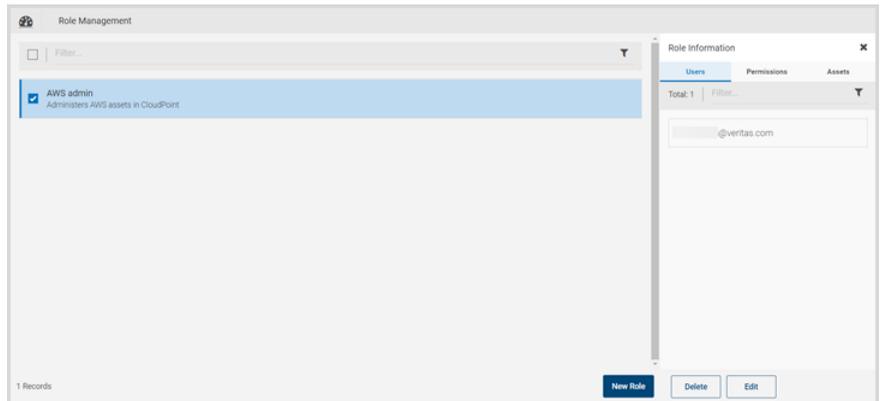
Editing a role

To edit a role

- 1 On the dashboard, in the **Administration** card, locate **Roles**, and click **Manage**.

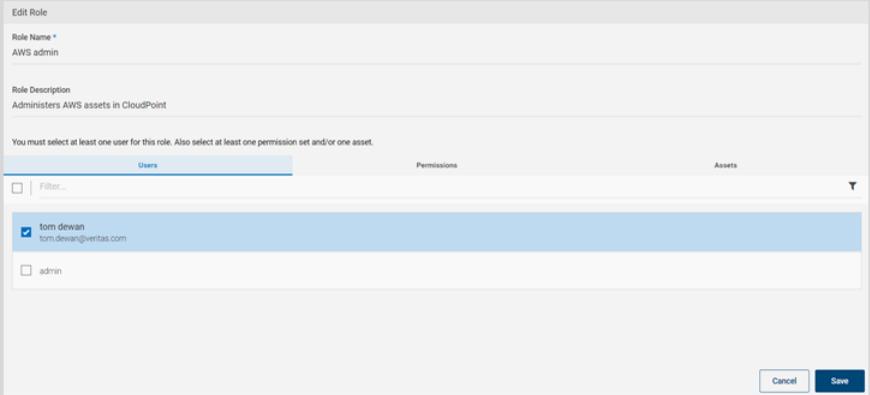


- 2 On the **Roles** page, select the check box for the role you want to view.



3 Click Edit.

The **Edit Role** page displays with the **Users** tab shown by default.



The screenshot shows the 'Edit Role' interface. At the top, the role name is 'AWS admin' and the description is 'Administers AWS assets in CloudPoint'. Below the description, there is a message: 'You must select at least one user for this role. Also select at least one permission set and/or one asset.' There are three tabs: 'Users', 'Permissions', and 'Assets'. The 'Users' tab is selected, showing a list of users with checkboxes. The user 'tom.dewan@veritas.com' is selected. At the bottom right, there are 'Cancel' and 'Save' buttons.

4 Modify the role values.

The remaining steps this procedure are the same as creating a new role.

See [“Creating a role”](#) on page 94.

5 After you edit the role, click Save.

CloudPoint displays a message that the changes have been applied.

Deleting a role

You can delete one or more CloudPoint roles in a single operation.

To delete a role

1 On the dashboard, in the **Administration** widget, locate **Roles**, and click **Manage**.

2 On the **Roles** page, select the check boxes for the roles you want to delete.

The **Role Details** page is displayed. If you select one role to delete, it displays the **Users** tab, **Permissions** tab, and **Assets** tab. If you select multiple roles to delete, the page displays the number of roles you selected.

3 On the **Role Details** page, click **Delete**.

4 On the **Please confirm ...** dialog box, click **Delete**.

CloudPoint displays a message that the role has been deleted.

5 Note that the role is no longer on the **Roles** page.

Protecting and managing data

- [Chapter 10. User interface basics](#)
- [Chapter 11. Protecting your assets with policies](#)
- [Chapter 12. Replicating snapshots for added protection](#)
- [Chapter 13. Managing your assets](#)
- [Chapter 14. Monitoring activities with notifications and the job log](#)
- [Chapter 15. Indexing and classifying your assets](#)
- [Chapter 16. Protection and disaster recovery](#)

User interface basics

This chapter includes the following topics:

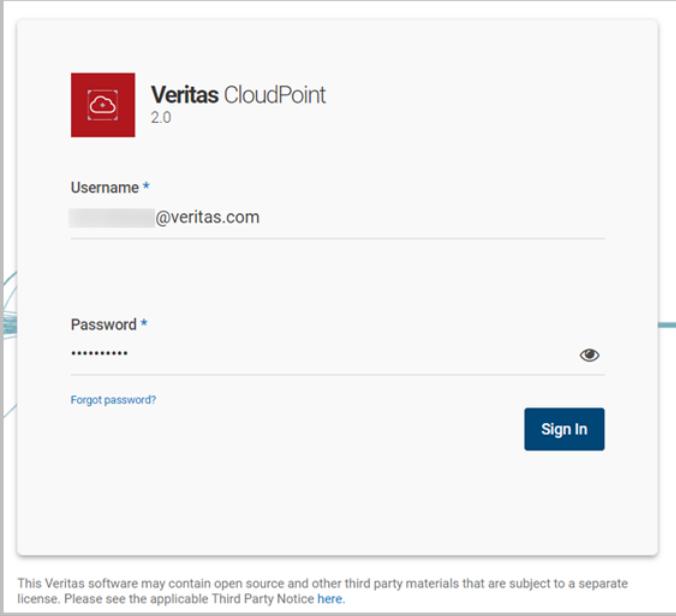
- [Signing in to CloudPoint](#)
- [Focusing on an asset type](#)
- [Navigating to your assets](#)
- [Using the action icons](#)

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.

To sign in to CloudPoint

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



Veritas CloudPoint
2.0

Username *
@veritas.com

Password *
.....

[Forgot password?](#)

Sign In

This Veritas software may contain open source and other third party materials that are subject to a separate license. Please see the applicable Third Party Notice [here](#).

- 2 Click **Sign In**.

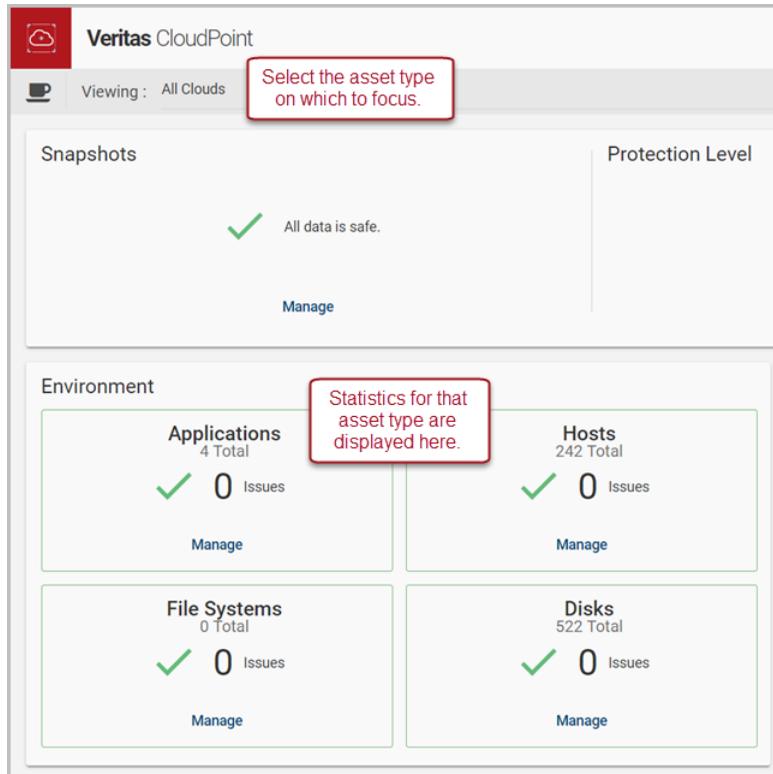
If this is the first time you have signed in to CloudPoint, verify that CloudPoint was installed successfully.

See “[Verifying that CloudPoint installed successfully](#)” on page 36.

Focusing on an asset type

By default, the dashboard displays statistics on all the clouds in your environment.

You can use the **Viewing** drop-down list to select a particular asset type. Then, the dashboard only displays statistics on that type.



The **Viewing** drop-down list has the following options:

- All clouds
- Amazon Web Services (AWS)
- Microsoft Azure
- Google Cloud
- OnPrem

Navigating to your assets

Many CloudPoint tasks consist of navigating to an asset and performing an action. Actions can include taking a snapshot, viewing a snapshot, or associating an asset with a policy.

The **Asset Management** page is the starting point for all these activities. You can filter the information on the Asset Management page to display the following:

- Everything (all asset types)
- Disks
- Hosts
- Applications
- File systems

The following example shows the **Asset Management** page listing only applications.



Type a search string in the **Filter** field and then press **Enter** to filter your search results further.

Note: If the search string you specify includes a hyphen, enclose the string in double quotes. For example, to show only the assets that include the string `prod-pipeline`, type `"prod-pipeline"`.

From here, you can select an application and perform a number of tasks.

The following table lists the ways you can navigate to the **Asset Management** page and what is displayed.

Table 10-1 Navigating to your assets

When you click here ...	The Asset Management page displays ...
Snapshots > Manage Classification > Manage Protection Summary > Manage	Everything (default) or the last asset type displayed
Protect Assets	Everything

Table 10-1 Navigating to your assets (*continued*)

When you click here ...	The Asset Management page displays ...
Applications > Manage Hosts > Manage File Systems > Manage Disks > Manage	The specified asset type

Using the action 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 10-2 CloudPoint icons

Click this icon ...	To display ...
	Notifications Recent CloudPoint activity, including creating, restoring, and deleting snapshots.
	Settings
	The CloudPoint online Help. The online Help displays information on CloudPoint deployment and administration.
	The logged on CloudPoint user name. You can perform the following actions from this screen: <ul style="list-style-type: none"> ■ Change the logged on CloudPoint user account password. ■ Display the installed CloudPoint version. ■ Sign out from the CloudPoint user interface (UI).

Protecting your assets with policies

This chapter includes the following topics:

- [About policies](#)
- [Creating a policy](#)
- [Assigning a policy to an asset](#)
- [Listing policies and displaying policy details](#)
- [Editing a policy](#)
- [Deleting a policy](#)

About policies

A policy lets you automate your asset protection. When you create a policy, you define the following:

- The type of snapshot to take, either a crash-consistent snapshot (the default) or an application-consistent snapshot.
- Whether or not to replicate the snapshot. For added protection, you can specify that CloudPoint stores a copy of the snapshot at another physical location.
- Whether or not to analyze snapshots using CloudPoint's classification feature. If you enable classification, CloudPoint analyzes your snapshots and displays an alert if they contain sensitive data such as personally identifiable information (PII).
- The number of snapshots to retain and how long to retain them before the snapshots and their replicated copies are deleted.

- The frequency with which the policy runs.

You can then assign the policy to your assets to ensure regular, consistent protection. You can assign more than one policy to an asset. For example, you can create a policy that takes asset snapshots on a weekly basis, and another that takes asset snapshots daily. You can then associate both the policies to the same asset.

Note: If you have an asset in multiple policies and the policy run times overlap, one of the policies may fail. For example, suppose an asset is in both Policy 1 and Policy 2. If Policy 1 is running when Policy 2 starts, Policy 2 may fail. It takes an average of 10 minutes to create an Oracle snapshot. Allow at least a 10 minute gap between two policies that are assigned to the same asset.

See [“Creating a policy”](#) on page 107.

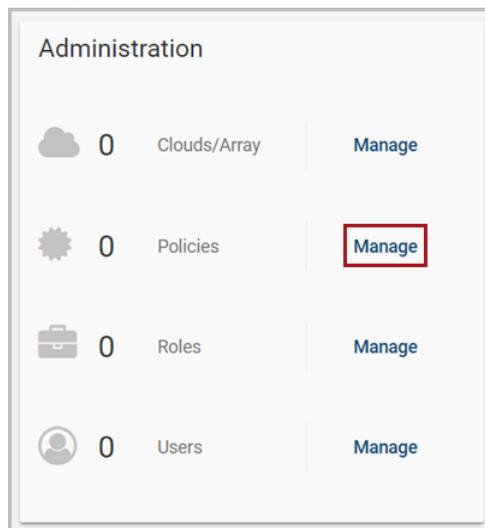
See [“Assigning a policy to an asset”](#) on page 111.

See [“Listing policies and displaying policy details”](#) on page 113.

Creating a policy

To create a policy

- 1 On the dashboard, in the **Administration** widget, locate **Policies**, and click **Manage**.



- 2 On the **Policies** page, click **New Policy**.

3 Complete the **New Policy** page.

Enter the following:

- **Policy Information**
Name and describe the policy, and enable features.

Field	Description
Policy Name	<p>A 2- to 13-character string.</p> <p>The name can only contain lower case letters, numbers, and hyphens. The name should begin and end with a letter.</p> <p>Notes:</p> <ul style="list-style-type: none"> ■ In Google Cloud, a policy name cannot contain an underscore. ■ If policy contains any on-premise array disk, then the policy name must be a 1-12 character string. ■ In case of a Pure Storage array, the policy name must also not contain an underscore.
Description <i>(optional)</i>	<p>A short description to remind you about what the policy does.</p>
Storage level	<p>The level at which the snapshot is taken: Disk, Host, or Application</p>

Field	Description
Application Consistent Snapshot	<p>Click the check box to enable application-consistent snapshots.</p> <p>In an application consistent snapshot, CloudPoint notifies the application that it is about to take a snapshot. The application completes its transactions and writes data to memory. It is then briefly frozen and CloudPoint takes the snapshot. The application resumes activity after the snapshot is taken.</p> <p>The default is to create a crash-consistent snapshot. This snapshot type does not capture data in memory or pending operations.</p> <p>An application-consistent snapshot is recommended for database applications. A crash-consistent snapshot is acceptable for other types of assets.</p> <p>This option is not available with the CloudPoint Freemium license.</p> <p>Note: CloudPoint does not support application-consistent snapshots on ext2 file systems.</p>
Enable Replication	<p>Click the check box to enable replication.</p>
Enable Classification	<p>Click the check box to enable classification.</p> <p>Note: This option is available only with the CloudPoint Enterprise or an equivalent license.</p>

■ Retention

Use the retention parameter to define how many copies of the asset snapshots to create and for how long you wish to keep them. You can choose to retain the snapshot copies for days, weeks, months, or years. After the retention period expires, CloudPoint automatically deletes all the snapshots and their replicated copies.

Note: Use careful planning and consideration when using this parameter. The retention period applies to the policy-created snapshot copies as well as the replicated snapshot copies that are stored at a different location than the source. All the snapshots are completely lost and you will not have access to them after they are deleted.

Use the up and down arrows and the retention tabs to specify how many snapshots of the asset you want to retain or for how long.

The following table shows some sample settings.

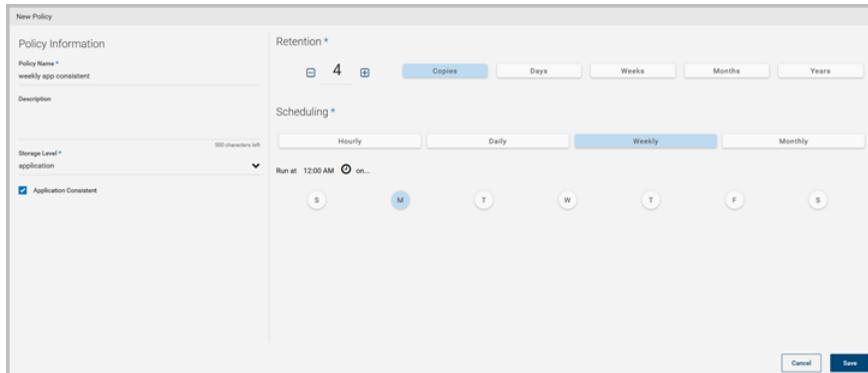
Number	Tab	Description
5	Copies	Retains the last five snapshots. Note: An asset may have more total snapshots than the number specified here. If an asset is associated with multiple policies, it has snapshots with each policy. Also, the snapshots you create manually do not count toward the retention total. Manual snapshots are not automatically deleted.
7	Days	Retains all snapshots for a week.
3	Months	Retains all snapshots for 3 months.

■ Scheduling

Use this part of the page to determine how often the policy runs.

Tab	Description
Hourly	Use the up and down arrows to specify the hour or minute interval at which the policy runs.
Daily	Click the clock icon to specify the time the policy runs each day.
Weekly	Use the clock icon and day buttons to specify the day of the week and the time the policy runs.
Monthly	Use the clock icon and calendar to specify the time and the date each month on which the policy runs.

The following example takes application consistent snapshots each Monday at 12:00 AM. CloudPoint retains four snapshots before it discards the oldest one.



- 4 Click **Save**.
 CloudPoint displays a message that the new policy is created.
- 5 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:

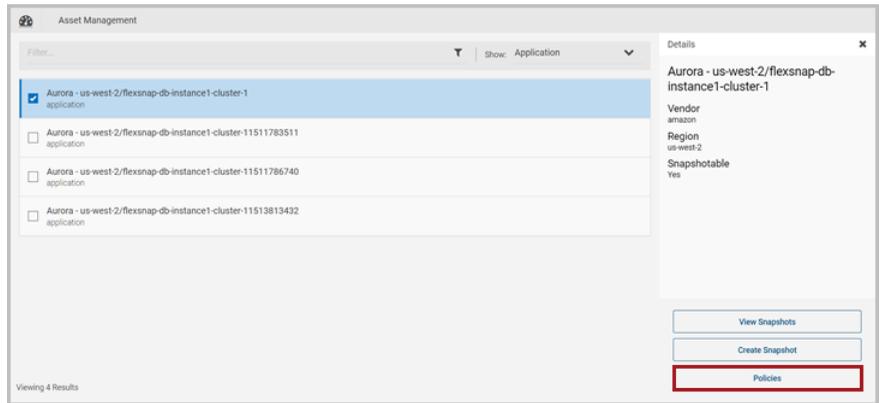
- The steps for assigning a policy are the same regardless of the type of asset you assign it to.
- You can follow the same step to change the policy that is associated with an asset or to un-assign a policy from an asset.
- CloudPoint does not support running multiple operations on the same asset simultaneously. If you have an asset in multiple policies and the policy run times overlap, one of the policies may fail.

For example, suppose an asset is in both Policy 1 and Policy 2. If Policy 1 is running when Policy 2 starts, Policy 2 may fail. It takes an average of 10 minutes

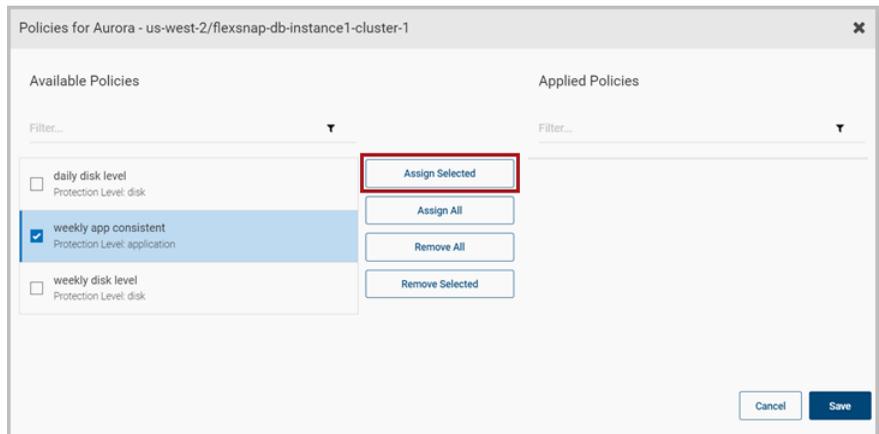
to create an Oracle snapshot. Allow at least a 10 minute gap between two policies that have the same asset.

To assign a policy to an asset

- 1 On the CloudPoint dashboard, in the **Environment** area, find the asset type you want to protect, and click **Manage**. This example protects an application.
- 2 On the **Asset Management** page, select the application you want to protect. On the **Details** page, click **Policies**.



- 3 On the **Policies for asset name** screen assign one or more policies to the asset. In the **Available Policies** column, select the policy you want to assign and click **Assign Selected**.



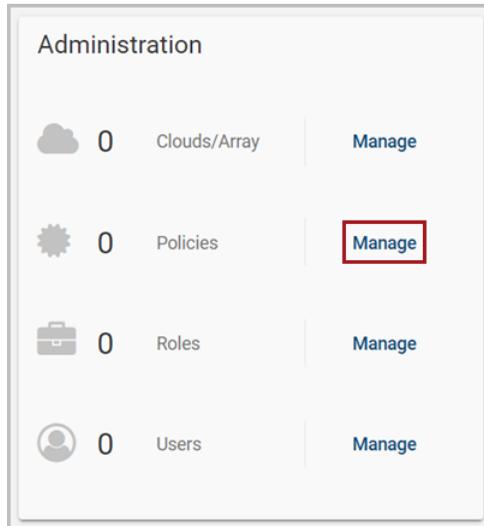
You can also assign or remove multiple policies at the same time.

- 4 Click **Save**.

Listing policies and displaying policy details

To list policies and display policy details

- 1 On the dashboard, in the **Administration** card, locate **Policies**, and click **Manage**.

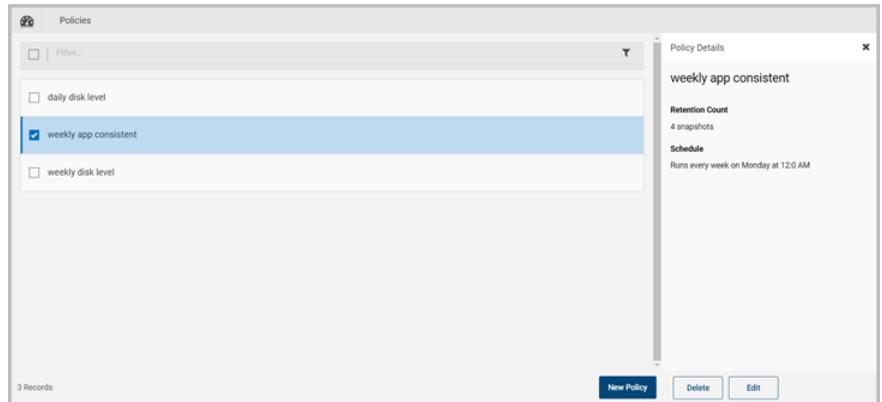


The **Policies** page displays with a list of policies.



From the **Policies** page, you can create a new policy.

- 2 To display a policy's details, select it from the list.



The **Policy Details** page displays the following information:

- The policy name
- The description (if available)
- The retention count; that is, number of snapshots that are kept for each asset before the oldest one is removed
- When the policy is scheduled to run

From **Policy Details** page, you can do the following:

- Edit a policy.
- Delete a policy.

See [“About policies”](#) on page 106.

See [“Creating a policy”](#) on page 107.

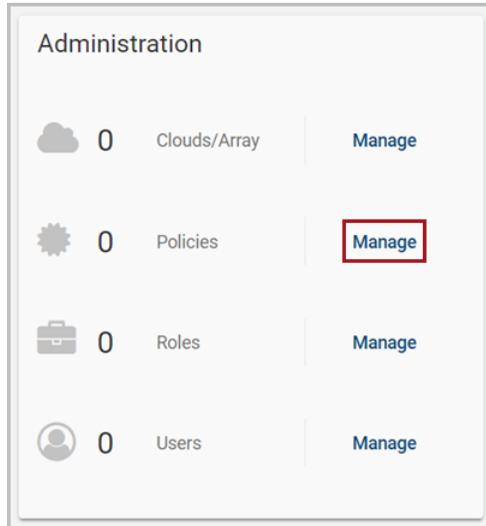
See [“Deleting a policy”](#) on page 116.

See [“Editing a policy”](#) on page 115.

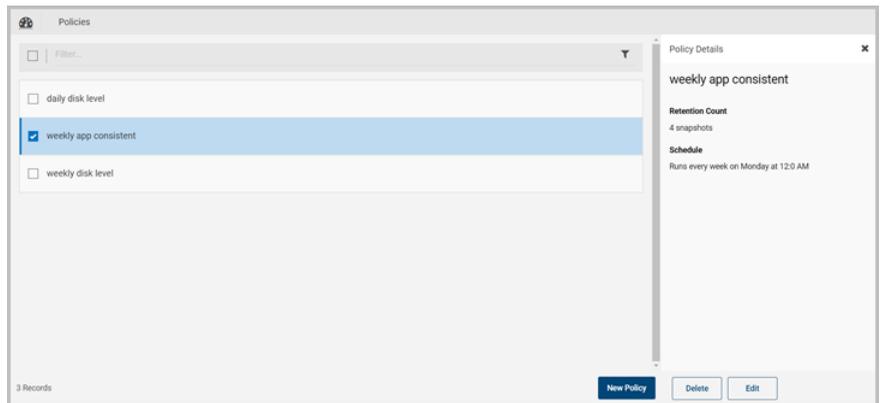
Editing a policy

To edit a policy

- 1 On the dashboard, in the **Administration** widget, locate **Policies**, and click **Manage**.



- 2 On the **Policies** page, select the check box for the policy you want to modify.



- 3 On the **Policy Details** page, click **Edit**.

4 Modify the policy values.

The screenshot shows the 'Edit Policy' interface. On the left, under 'Policy Information', the 'Policy Name' is 'weekly app consistent' and the 'Storage Level' is 'application'. The 'Application Consistent' checkbox is checked. On the right, under 'Retention *', the value is '4' and the unit is 'Copies'. Under 'Scheduling *', the frequency is 'Weekly' and the time is '12:00 AM on...'. The day of the week is set to 'M' (Monday). At the bottom right, there are 'Cancel' and 'Save' buttons.

The remaining steps this procedure are the same as creating a new policy.

See [“Creating a policy”](#) on page 107.

5 After you edit the policy, click **Save**.

CloudPoint displays a message that the policy is updated.

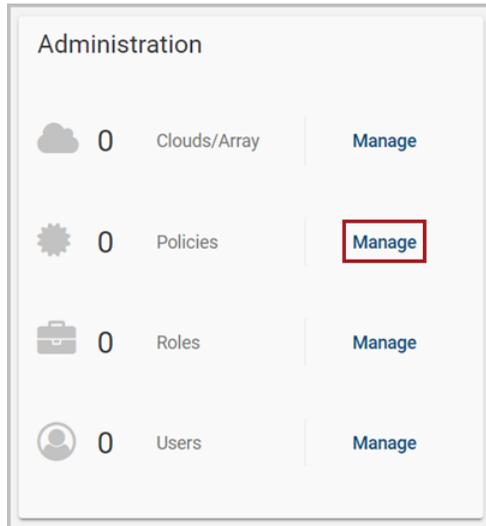
See [“About policies”](#) on page 106.

Deleting a policy

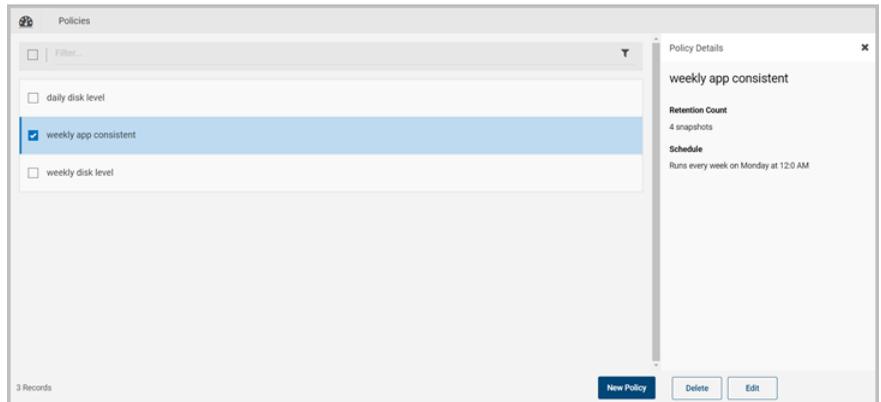
Policy deletion fails if there are assets assigned to the policy. You must unassign all assets that are associated with a policy before attempting to delete that policy.

To delete a policy

- 1 On the dashboard, in the **Administration** card, locate **Policies**, and click **Manage**.



- 2 On the **Policies** page, select the check box for the policy you want to delete. You can select multiple policies.



- 3 On the **Policy Details** page, click **Delete**.
- 4 On the **Please confirm ...** dialog box, click **Delete**.

5 CloudPoint displays a message that the policy has been deleted.

6 Note that the policy is no longer on the **Policies** page.



See [“About policies”](#) on page 106.

See [“Creating a policy”](#) on page 107.

Replicating snapshots for added protection

This chapter includes the following topics:

- [About snapshot replication](#)
- [Requirements for replicating snapshots](#)
- [Configuring replication rules](#)
- [Editing a replication rule](#)
- [Deleting a replication rule](#)

About snapshot replication

When you replicate a snapshot, you save a copy of the snapshot to another physical location. For example, suppose that you administer an Amazon Web Services (AWS) cloud and your assets are in the region `us-east-1`. Your asset snapshots will also be stored in `us-east-1` region. However, you can also replicate the snapshots to the region `us-west-1` for an added level of protection. In CloudPoint terminology, the original location (`us-east-1`) is the replication source, and the location where snapshots are replicated (`us-west-1`) is the replication destination.

As an administrator, you can configure up to three replication targets for each source region. You can replicate a snapshot manually or using a policy. When you create a policy, replication is one of the policy parameters that you can enable. Note that replication via policy works only when there is a replication rule available for the particular region.

Requirements for replicating snapshots

For replicating unencrypted snapshots

Ensure that you add the AWS accounts (using the CloudPoint AWS plug-in) configuration in CloudPoint. These are the AWS accounts between which you want to replicate snapshots.

There are no additional requirements for replicating unencrypted snapshots.

For replicating encrypted snapshots

Prerequisites for replicating encrypted snapshots:

- Encryption key (KMS key) used for encryption must have the same name in both regions; that is, they should have the same key alias (in terms of AWS). If encryption key with the same name is not present, then the replication fails with the following error:

```
KMS key <encryption_key_arn> not present in target region:  
<target_region>
```

See [“AWS permissions required by CloudPoint”](#) on page 46.

Configuring replication rules

A replication rule consists of the following:

- The original location of your assets and snapshots
- One or more alternate physical locations where snapshots are replicated

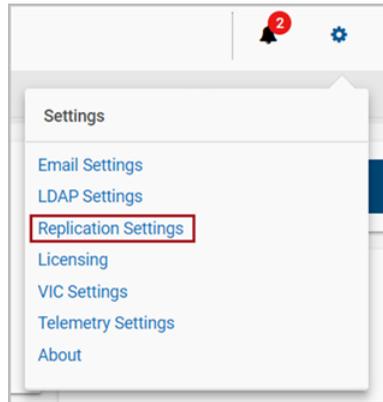
You can configure up to three replication destination for each source.

You can use a replication rule in the following ways:

- You can automate replication. On a snapshot policy, select **Enable Replication**. When the policy runs, snapshots are automatically replicated to the targets that are configured in the rule.
- You can replicate a snapshot manually. On the **Snapshot Details** page, select **Replicate**.

To create a replication rule

- 1 On the CloudPoint dashboard, click the **Settings** (gear) icon, and select **Replication Settings** from the drop-down list.



- 2 On the **Replication Settings** page, click **New Rule**.
- 3 On the **New Replication Rule** page, use the drop-down lists to configure your rule.

Drop-down list	Description
Platform	Specify the asset vendor. Currently, CloudPoint supports Amazon Web Services (AWS).
Location/Region	The choices here are based on what you select on the Platform list. The location you select becomes the Source Name on the Replication Settings page.
Destination 1, Destination 2, Destination 3	Use these drop-down lists to select one or more alternate physical locations where replicated snapshots are stored. Note: For AWS, you cannot replicate snapshots between two accounts. You can only replicate snapshots between locations in the same account.

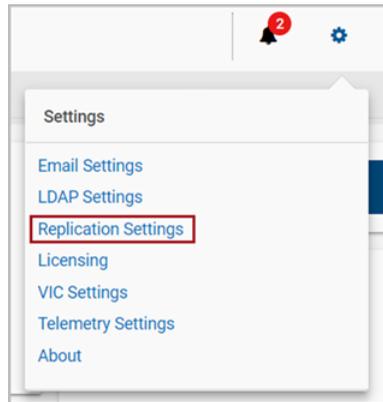
- 4 Click **Save**.
CloudPoint displays a message that a new rule has been created.
- 5 Note that the **Replication Settings** screen displays the new rule.

Editing a replication rule

You can edit a replication rule to change the location where snapshots are replicated or the order of the locations. You cannot edit the vendor platform or source location.

To edit a replication rule

- 1 On the CloudPoint dashboard, click the **Settings** (gear) icon, and select **Replication Settings** from the drop-down list.



- 2 Review the **Replication Setting** page.

This page lists each replication source in your environment. It includes the following information for each source:

- The source name
- The source server
- The source platform type, such as Amazon Web Services (AWS)
- The regions to which the snapshots are replicated

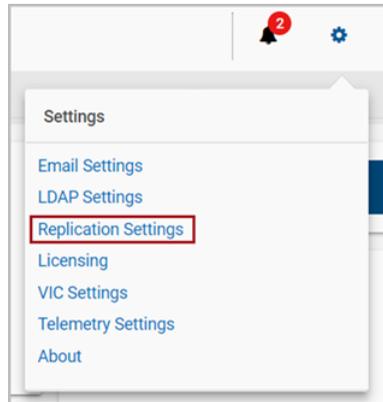
- 3 Select the source location whose replication rules you want to edit.
- 4 Click **Edit**.
- 5 Use the drop-down lists to change the replication locations or the order of the locations.
- 6 Click **Save**.

CloudPoint displays a message that a new rule has been updated.

Deleting a replication rule

To delete a replication rule

- 1 On the CloudPoint dashboard, click the **Settings** (gear) icon, and select **Replication Settings** from the drop-down list.



- 2 Select the replication rules you want to delete. You can select more than one rule.
- 3 Click **Delete**.
- 4 On the **Please confirm ...** dialog box, click **Delete**.
CloudPoint displays a message that the rule has been deleted

Managing your assets

This chapter includes the following topics:

- [Creating a snapshot manually](#)
- [Displaying asset snapshots](#)
- [Replicating a snapshot manually](#)
- [About snapshot restore](#)
- [About single file restore \(granular restore\)](#)
- [Single file restore requirements and limitations](#)
- [Restoring a snapshot](#)
- [Deleting a snapshot](#)

Creating a snapshot manually

One of CloudPoint's most important features is the ability to create snapshot policies. These policies let you take snapshots of specific assets on a regular schedule.

However, you can also take a snapshot of an asset manually. That is, you can navigate to a particular asset at any time and create a snapshot.

The types of snapshots you can create vary depending on the asset type. Review the following table:

Table 13-1 Assets and supported snapshot types

Asset	Supported snapshot types
Dell EMC Unity array	Copy-on-write (COW) snapshots on LUNs

Table 13-1 Assets and supported snapshot types (*continued*)

Asset	Supported snapshot types
HPE 3PAR array	COW and clone snapshot types Note the following: <ul style="list-style-type: none"> ■ HPE 3PAR Virtual Copy Software is responsible for the snapshot operation. ■ You can have 500 snapshots per volume. 256 can be read/write. ■ When a volume is involved in a Remote Copy with a secondary array, the operation fails. ■ You can take a clone snapshot, however you cannot restore it.
Hitachi HDS array	COW snapshots; Hitachi Thin Image (HTI) volumes P-VOL or S-VOL The following are not supported: <ul style="list-style-type: none"> ■ Clone snapshots; Multi Raid Coupling Facility (MRCF): ShadowImage volume P-VOL or S-VOL ■ The VVol volume type
Pure Storage FlashArray	Clone snapshots of volumes
Huawei OceanStor arrays	COW snapshots on LUNs

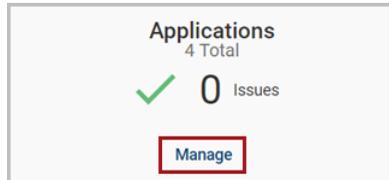
Note: CloudPoint does not support running multiple operations on the same asset simultaneously. You can perform only one operation at any given time. If multiple operations are submitted for the same asset, then only the first operation is triggered and the remaining operations will fail.

Regardless of the asset type you work with, the steps for creating a snapshot are the same. Depending on the asset, some parameters you enter may be slightly different. They are explained in the procedure.

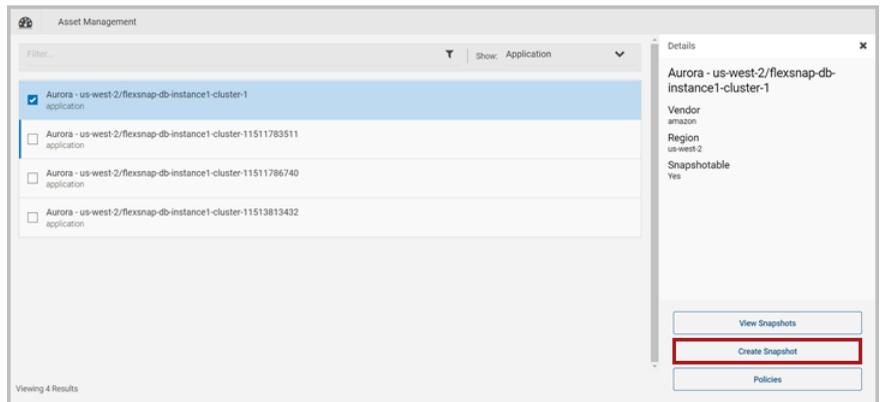
To create a snapshot manually

- 1 Navigate to your list of assets.

On the CloudPoint dashboard, in the **Environment** card, select the asset type you want to work with, and click **Manage**. This example creates an application snapshot.



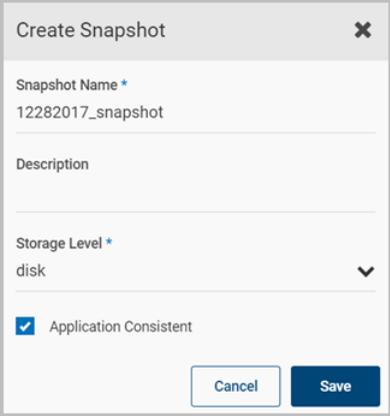
- 2 On the **Asset Management** page, select the application for which you want to create a snapshot.
- 3 On the asset's **Details** page, click **Create Snapshot**



- 4 On the **Create Snapshot** page, complete the following fields.

Field	Description
Snapshot name	<p>A 2- to 32-character string.</p> <p>Cloud vendors have additional restrictions on the snapshot name.</p> <ul style="list-style-type: none"> ■ In Amazon Web Services, an RDS snapshot or Aurora cluster snapshot name has the following restrictions: <ul style="list-style-type: none"> ■ The name cannot be null, empty, or blank. ■ The first character must be a letter. ■ The name cannot end with a hyphen or contain two consecutive hyphens. ■ In Google Cloud, an application snapshot name has the following restrictions: <ul style="list-style-type: none"> ■ The name can only contain lower case letters, numbers, and hyphens. You cannot use an underscore. ■ The name should begin and end with a letter.
Description	<p>This field is optional. You can create a summary to remind you of the snapshot content.</p>
Storage level	<p>This option only displayed for application snapshots.</p> <p>host takes a snapshot of all the disks that are associated with the instance. You cannot restore an application snapshot that has the host protection level.</p> <p>disk takes a snapshot of the disks the application uses.</p>
Applicaiton Consistent	<p>Click this option to enable an application-consistent snapshot.</p> <p>In an application consistent snapshot, CloudPoint notifies the application that it is about to take a snapshot. The application completes its transactions and writes data to memory. It is then briefly frozen and CloudPoint takes the snapshot. The application resumes activity after the snapshot is taken.</p> <p>The default is to create a crash-consistent snapshot. This snapshot type does not capture data in memory or pending operations. An application-consistent snapshot is recommended for database applications. A crash-consistent snapshot is acceptable for other types of assets</p> <p>Note: CloudPoint does not support application-consistent snapshots on ext2 file systems.</p>

The following example creates a disk level snapshot with application consistency.



The screenshot shows a 'Create Snapshot' dialog box. It has a title bar with the text 'Create Snapshot' and a close button (X). Below the title bar, there are four main sections: 'Snapshot Name' with a red asterisk and the value '12282017_snapshot'; 'Description' with an empty text area; 'Storage Level' with a red asterisk and a dropdown menu showing 'disk'; and 'Application Consistent' with a checked checkbox. At the bottom, there are two buttons: 'Cancel' and 'Save'.

5 Click **Save**.

CloudPoint displays a message that the snapshot is created.

About resource limits for Amazon RDS

By default, AWS allows up to a 100 RDS manual snapshots per region. You may get an error if you try to take more than a 100 snapshots.

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Limits.html#RDS_Limits.Limits

You can work around this issue using any of the following options:

- Contact AWS support and request them for an increase in the number of snapshots allowed. Once they do that, you will not get an error until you reach the new limit.
- Reduce the retention in your policies so as to keep the snapshots count within the maximum limit.

Displaying asset snapshots

You can display all the snapshots for an asset, when they were created, and the region they are located in.

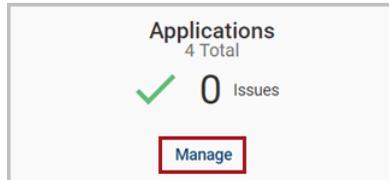
In addition, displaying an asset's snapshots is your gateway to other activities, including the following:

- Restoring a snapshot
- Replicating a snapshot manually
- Deleting a snapshot

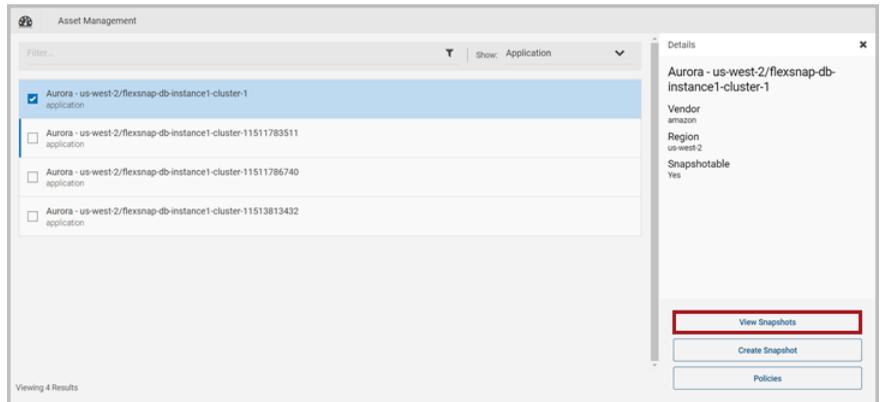
To display an asset's snapshots

- 1 Navigate to your list of assets.

On the CloudPoint dashboard, in the **Environment** card, select the asset type you want to work with, and click **Manage**. This example displays the snapshots for an application.



- 2 On the **Asset Management** page, select the application whose snapshots you want to view and then click **View Snapshots** from the Details pane on the right.



- 3 The **Snapshot Management** page lists all the snapshots. You can filter and sort the list to find the snapshot you are interested in.



From this page, you can select a snapshot and perform the following actions:

- Restore a snapshot
See [“Restoring a snapshot”](#) on page 138.
- Replicate a snapshot
See [“Replicating a snapshot manually”](#) on page 130.
- Classify a snapshot
- Delete a snapshot
See [“Deleting a snapshot”](#) on page 145.

Replicating a snapshot manually

When you replicate a snapshot, you save a copy of the snapshot to another physical location. Replication gives your data extra protection in case of a disaster at the original site.

The most efficient way to use replication is to define replication rules and then apply the rules to your snapshot policies. That way, replication takes on a regular schedule. Setting up replication rules is described in the chapter titled *“Replicating snapshots for added protection.”*

See [“About snapshot replication”](#) on page 119.

However, you can also replicate a snapshot manually. That is, you can navigate to a particular snapshot at any time, specify an alternate location, and replicate it.

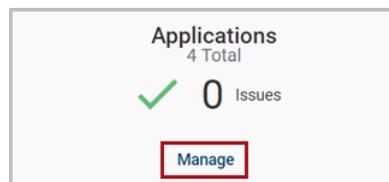
Regardless of the asset type you work with, the steps for replicating a snapshot are the same.

To replicate a snapshot manually

- 1 Navigate to your list of assets.

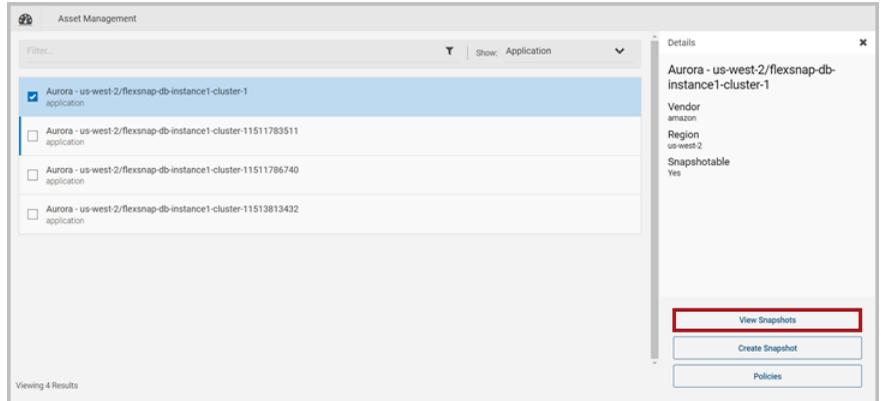
On the CloudPoint dashboard, in the **Environment** card, select the asset type you want to work with, and click **Manage**.

This procedure uses an application snapshot for replication as an example.

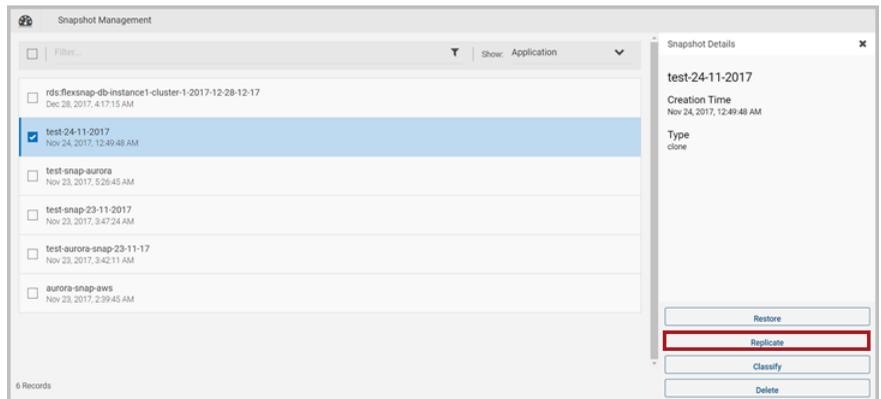


- 2 On the **Asset Management** page, select the application whose snapshot you want to replicate.

3 On the **Details** page click **View Snapshots**



4 On the **Snapshot Management** page, select the snapshot you want to replicate. You can only select one.



5 Depending on the structure for the snapshot, do one of the following:

- If the snapshot does not have any sub-assets, click **Replicate**.
- If the snapshot has sub-assets, a **Snapshot Assets** page is displayed. By default, all sub-assets are checked. Select the sub-assets you want to replicate and click **Replicate**.

- On the **Replicate** page, use the **Target Destination** drop-down list to select an alternate physical location.

- Click **Replicate**.
- On the **Please Confirm ...** dialog box, click **Replicate**.

A message on the CloudPoint UI confirms that the replication job has been triggered. You can view the details of the replication job in the Job Log panel.

After the replication is completed, open the Asset Management pane to view the replicated snapshot. The Snapshot Details panel displays additional information such as the creation time, the type, and the source AWS account. The Replicated field value is displayed as Yes, indicating that it is a replicated snapshot.

Snapshot Name	Created On	Snapshot Type
testSnap-23b545	Mar 28, 2019, 10:53:38 AM	clone
indexE1-3	Mar 26, 2019, 4:59:59 PM	clone
hourly20190325230000-16384-2	Mar 26, 2019, 4:25:52 AM	clone
testSnap-2	Mar 22, 2019, 2:21:19 PM	clone

Snapshot Details	
testSnap-23b545	
Creation Time	Mar 28, 2019, 10:53:38 AM
Type	clone
Source ID	aws-ebb-us-east-2-vol-01b630c8e36155a04
Account ID	935928673735
Replicated	Yes
Region	us-east-2

About snapshot restore

The types of snapshots you can restore and where you can restore them varies depending on the asset type.

Table 13-2 Assets and supported restore options

Asset	Supported restore options
Dell EMC Unity array	Restore a copy-on-write (COW) LUN snapshot to the same LUN with the Overwrite Existing option.
HPE 3PAR array	Restore a COW volume snapshot to the same volume with the Overwrite Existing option. Note the following: <ul style="list-style-type: none"> ■ Although you can take a clone snapshot, you cannot restore it. ■ When a volume has both COW and clone snapshot type, restore operations fail on that volume. ■ When a volume is involved in a Remote Copy with a secondary array, the operation fails. ■ When the array operation begins, the array creates a backup point for the volume.
Pure Storage FlashArray	Restore a clone volume snapshot to the same volume with the Overwrite Existing option.

When you restore a snapshot, keep in mind the following:

- You can restore an encrypted snapshot. To enable the restoring of encrypted snapshots, add a Key Management Service (KMS) policy, and grant the CloudPoint user access to KMS keys so that they can restore encrypted snapshots.
- If you are restoring a replicated host snapshot to a location that is different from the source region, then the restore might fail as the key is not available at the target location.
 As a prerequisite, create a key-pair with the same name as the source of the snapshot, or import the key-pair from the source to the target region.
 Then, after the restore is successful, change the security groups of the instance from the network settings for the instance.
- When you have created a snapshot of a disk of supported storage arrays from 'Disk' section in CloudPoint dashboard, which has a file system created and mounted on it, you must first stop any application that is using the file system and then unmount the file system and perform restore.
 For AWS/Azure/GCP cloud disk/volume snapshots, you must first detach the disk from the instance and then restore the snapshot to original location.

- (Applicable to AWS only) When you restore a host-level application snapshot, the name of the new virtual machine that is created is the same as the name of the host-level snapshot that corresponds to the application snapshot.
For example, when you create an application snapshot named `OracleAppSnap`, CloudPoint automatically creates a corresponding host-level snapshot for it named `OracleAppSnap-<number>`. For example, the snapshot name may resemble `OracleAppSnap-15`.
Now, when you restore the application snapshot (`OracleAppSnap`), the name of the new VM is `OracleAppSnap-<number> (timestamp)`.
Using the example cited earlier, the new VM name may resemble `OracleAppSnap-15 (restored Nov 20 2018 09:24)`.
Note that the VM name includes "*Oracle-AppSnap-15*" which is the name of the host-level snapshot.
- (Applicable to AWS only) When you restore a disk-level application snapshot or a disk snapshot, the new disk that is created does not bear any name. The disk name appears blank.
You have to manually assign a name to the disk to be able to identify and use it after the restore.
- When you restore a snapshot of a Windows instance, you can log in to the newly restored instance using original instance's username/password/pem file.
By default, AWS disables generating a random encrypted password after launching the instance from AMI. You must set `Ec2SetPassword` to `Enabled` in `config.xml` to generate new password every time. For more information on how to set the password, see the following link.
https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/ec2config-service.html#UsingConfigXML_WinAMI
- The volume type of newly created volumes for replicated snapshots is according to the region's default volume type.
If volume type is not specified, the following default values are used:

Table 13-3 Default volume types

Region	Default volume type
<ul style="list-style-type: none"> ■ us-east-1 ■ eu-west-1 ■ eu-central-1 ■ us-west-2 ■ us-west-1 ■ sa-east-1 ■ ap-northeast-1 ■ ap-northeast-2 ■ ap-southeast-1 ■ ap-southeast-2 ■ ap-south-1 ■ us-gov-west-1 ■ cn-north-1 	standard
All other regions	gp2

- You can perform only one restore operation on a snapshot at any given time. If multiple operations are submitted on the same asset, then only the first operation is triggered and the remaining operations will fail.
 This is applicable for all CloudPoint operations in general. CloudPoint does not support running multiple jobs on the same asset simultaneously.
- The destination host where you wish to restore the snapshot must have the same Oracle version installed as that at the source.
- If you are restoring the snapshot to a different location, verify the following:
 - Ensure that there is no database with the same instance name running on the target host.
 - The directories that are required to mount the application files are not already in use on the target host.

About single file restore (granular restore)

You can use CloudPoint to restore individual files within a snapshot. This process is known as "granular restore" (each single file in the snapshot is considered as a granule) or more commonly referred to as "single file restore" (SFR). CloudPoint makes an inventory of all the files within a snapshot using a simple indexing process. You can restore specific files from a snapshot only if that snapshot has been indexed by CloudPoint.

CloudPoint also supports a deeper and more intelligent scan of the snapshot files using a process known as Classification. This process goes a little further into the data than indexing. During classification, CloudPoint first indexes a snapshot and then identifies items that contain tags that describe the type of the data in the snapshot files. Tags indicate the type of information in a file, but not the actual data. For example, if a snapshot file contains credit card data, the tag indicates that the file includes information about a credit card, but does not identify the actual credit card number in the file. To classify individual files within a snapshot, CloudPoint uses a built-in set of data tags that are predefined in Veritas Information Classifier (VIC).

Both indexing and classification are two independent processes. You can choose to index a snapshot without classifying or to index and classify a snapshot.

See [“Single file restore requirements and limitations”](#) on page 136.

See [“Configuring classification settings using VIC”](#) on page 154.

See [“About indexing and classifying snapshots”](#) on page 152.

Single file restore requirements and limitations

If you wish to use single file restore (SFR) feature, make a note of the following:

- To restore individual files within a snapshot, the snapshot must be indexed or classified first.
See [“Indexing and classifying snapshots”](#) on page 154.
- Indexing is a licensed feature and is not available with the CloudPoint basic freemium license. Install a CloudPoint Enterprise or an equivalent license to enable and use the feature in your CloudPoint deployment.
See [“Understanding your CloudPoint license”](#) on page 11.
See [“About indexing and classifying snapshots”](#) on page 152.
- SFR is supported only for disk-level file system snapshots.
- SFR is supported only for Amazon Web Services (AWS) and Microsoft Azure assets.
- For indexing and classification to work, the CloudPoint host and the Windows or Linux instances must belong to the same region.

See [“Single file restore support on Linux”](#) on page 137.

See [“Single file restore limitations on Linux”](#) on page 137.

Single file restore support on Linux

If you wish to use single file restore (SFR) on Linux instances, make a note of the following:

- SFR is supported on ext4 and XFS file systems (*starting with CloudPoint 2.0.1 release*).

Single file restore limitations on Linux

The following limitations are applicable to CloudPoint single file restore (SFR) on Linux:

- SFR is not supported for file systems managed using Logical Volume Manager (LVM).
- SFR is not supported for the / (root file system).
- Restoring a file to a custom user-defined location on the original instance is not supported. You can restore a file only to its existing location on the original instance.
- Restoring a file to an altogether different instance is not supported. You can restore files only on the original instance.

Single file restore support on Windows

If you wish to use single file restore (SFR) on Windows instances, make a note of the following:

- SFR is supported on NTFS file system on Windows (*starting with CloudPoint release 2.2*).
- SFR is supported only for NTFS file system on disks that are formatted as MBR (Primary, Extended or Logical) and GPT partitions (Basic disks).
- SFR is supported for NTFS junction points (NTFS volume mounted as a directory path).
- SFR is supported on multi-partition disks.
- SFR is supported for Amazon Web Services (AWS) and Microsoft Azure cloud assets.

Single file restore limitations on Windows

The following limitations are applicable to CloudPoint single file restore (SFR) support on Windows:

- SFR is not supported for Google Cloud Platform (GCP) cloud assets.

- SFR is not supported for FAT, FAT32, and ReFS file systems.
- SFR is not supported for disks and volumes managed using Windows Logical Disk Manager (LDM). The Windows on-host plug-in does not discover them.
- SFR is not supported for encrypted or compressed files.
- When you perform a restore, directory Access Control Lists (ACLs), permissions, attributes, and time stamps are not restored.
- Restoring a file to a custom user-defined location on the original instance is not supported. You can restore a file only to its existing location on the original instance.
- Restoring a file to an altogether different instance is not supported. You can restore files only on the original instance.
- Restore is not supported for file names or paths that contain Unicode characters.

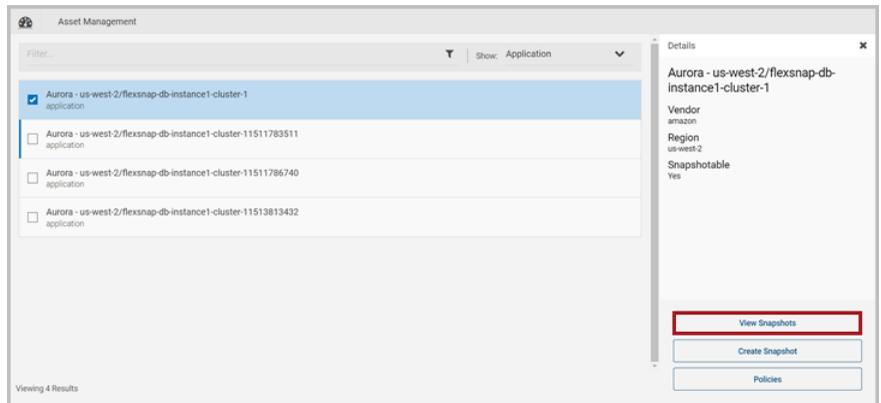
Restoring a snapshot

To restore a snapshot

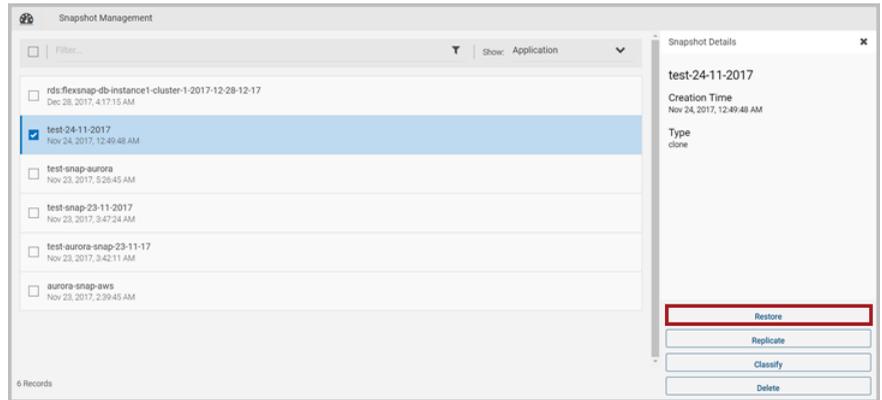
- 1 Navigate to your list of assets.

On the CloudPoint dashboard, in the **Environment** card, select the asset type you want to work with, and click **Manage**. This example restores an application snapshot.

- 2 On the **Asset Management** page, select the application whose snapshot you want to restore.
- 3 On the **Details** page click **View Snapshots**.



- 4 On the **Snapshot Management** page, select the snapshot you want to restore and then click **Restore**.



- 5 On the **Restore** page, complete the following.
 - Specify a **Restore Job Name** and **Description**.
 - Select one of the following restore options, depending on the snapshot type:

Snapshot type

- Cloud snapshot
- Host / instance
 - Disk
 - Application (host, disk)

Restore option

Overwrite existing

Description

Replaces the current asset with the snapshot.

Following is the behavior for this option:

- CloudPoint creates new EBS volumes from the VM (disk) snapshots and stops the original instance. It detaches the existing volumes and attaches them to the stopped instance to start the instance.
- VM or instance ID remain the same, but as new disks are created from the snapshots, the disk IDs are different.
- Instance and volume tags are copied properly.
- Policies assigned to hosts are preserved.

Notes:

- This restore option is supported on AWS only.
- CloudPoint does not remove or delete the older volumes in any restore scenario.
- If an instance is corrupted, Veritas recommends that you revert that instance to a previous working state instead of spinning up a new instance by restoring an existing snapshot. This helps in avoiding any orchestration tasks that may be needed to integrate the instance into other workflows.

Snapshot type	Restore option	Description
	Original Location	Restores the snapshot to the same location as the asset, without overwriting the existing asset.
	New location	Restores the snapshot to a completely different location in the cloud. You can select a target destination from the list of available options displayed in the drop-down list. For example, in case of AWS cloud, the list displays all the subnets in the AWS region where the asset resides. Note: Currently, you cannot restore an Oracle snapshot to a new location.
Array snapshot	Overwrite existing	Replaces the current asset with the snapshot. CloudPoint sends a snapshot restore request to the underlying storage array and presents it with the selected snapshot. The storage array then performs the actual snapshot restore operation.

■ Disk

6 Click Restore.

(Applicable for AWS only) If you are performing a disk-level snapshot restore, then the new disk that is created after the restore does not bear a name. The disk name appears blank. In such a case, you have to manually assign a name to the disk to be able to identify and use the disk after the restore is complete.

Note: Starting with release 2.0.2, you can restore an Azure instance snapshot to a private network. The instance does not require a public IP address.

Additional steps required after a SQL Server disk-level snapshot restore to a new location

The following steps are required after you restore a Microsoft SQL Server disk-level snapshot to a new location.

Even though the restore operation is successful, these steps are required for the application database to be available for normal use again.

Perform the following steps:

- 1 Ensure that the disk-level snapshot restore operation has completed successfully and a new disk is created and mounted on the application host.
- 2 Log on to Microsoft SQL Server Management Studio as a database administrator.
- 3 From the Object Explorer, connect to an instance of the SQL Server Database Engine and then click to expand the instance view.
- 4 In the expanded instance view, right-click **Databases** and then click **Attach**.
- 5 In the Attach Databases dialog box, click **Add** and then in the Locate Database Files dialog box, select the disk drive that contains the database and then find and select all the .mdf files associated with that database. Then click **OK**.

 The disk drive you selected should be the drive that was newly created by the disk-level snapshot restore operation.
- 6 Wait for the requested operations to complete and then verify that the database is available and is successfully discovered by CloudPoint.

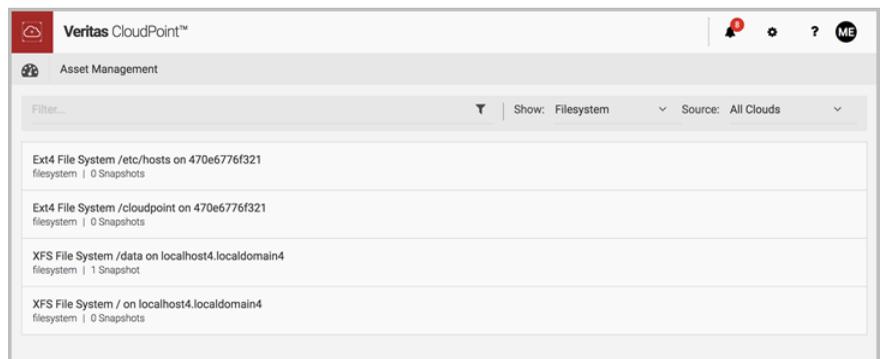
Restoring individual files within a snapshot

CloudPoint provides a way for you to restore specific files that are part of a snapshot. This process is known as granular file restore and is also commonly referred to as single file restore (SFR). Before restoring individual files, ensure that you are aware of the feature support, requirements, and its limitations.

See [“About single file restore \(granular restore\)”](#) on page 135.

To restore individual files within a snapshot

- 1 On the CloudPoint dashboard, in the **File Systems** area, click **Manage**.
- 2 On the **Asset Management** page, select the file system whose snapshots you want to view.



3 On the **Details** page, click **View Snapshots**.

The screenshot displays the Veritas CloudPoint Asset Management interface. The main area shows a list of file systems with the following details:

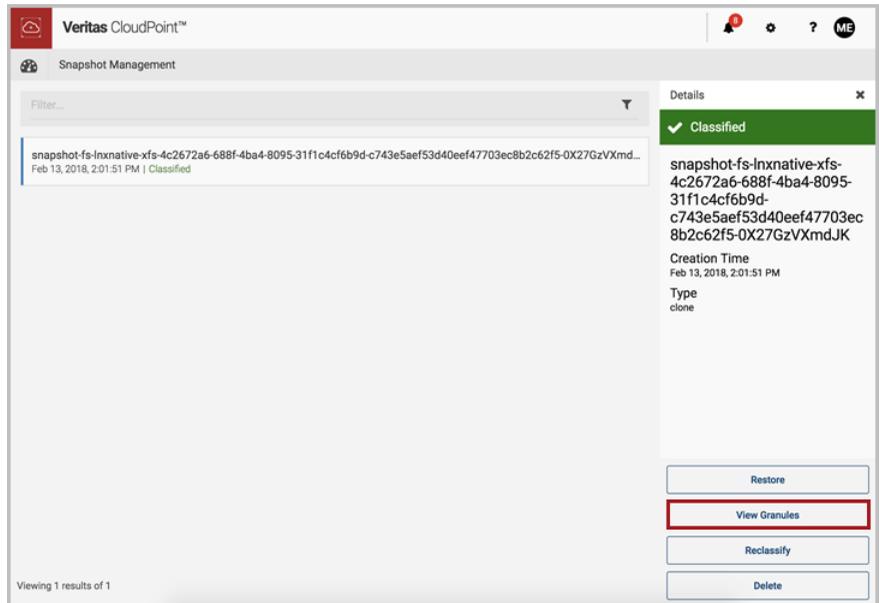
File System	Filesystem	Snapshots
Ext4 File System /etc/hosts on 470e6776f321	filesystem	0 Snapshots
Ext4 File System /cloudpoint on 470e6776f321	filesystem	0 Snapshots
XFS File System /data on localhost4.localdomain4	filesystem	1 Snapshot
XFS File System / on localhost4.localdomain4	filesystem	0 Snapshots

The details panel on the right for the selected XFS File System shows the following information:

- XFS File System /data on localhost4.localdomain4**
- Vendor**: -
- Region**: -
- Snapshotable**: Yes
- ID**: fs-Inactive-xfs-4c2672a6-688f-4ba4-8095-31f1c4cf6b9d-c743e3aeaf53d40ee147703ec8b2c62f5
- Policies (0)**

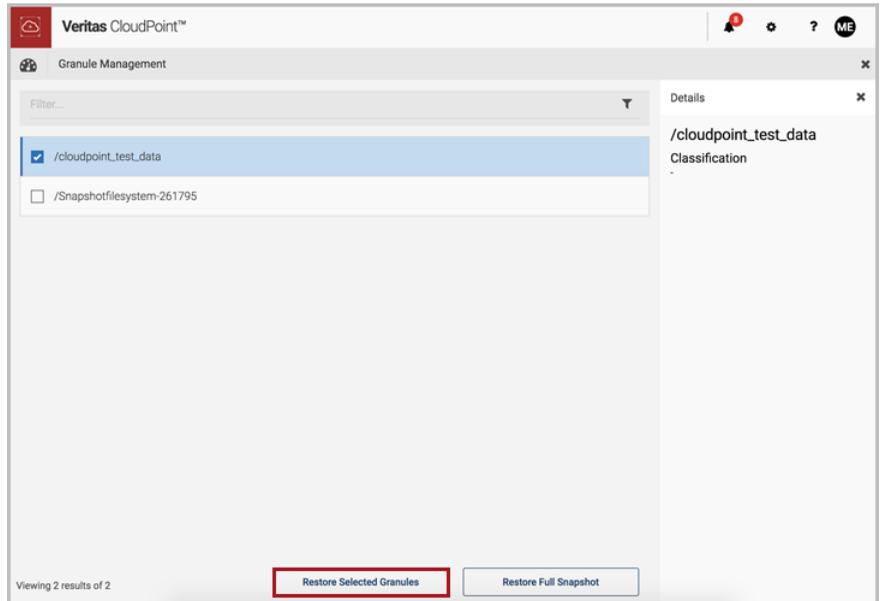
At the bottom of the details panel, there are three buttons: **View Snapshots (1)** (highlighted with a red box), **Create Snapshot**, and **Policies**. The main interface also includes a filter bar, a 'Show: Filesystem' dropdown, and a 'Source: All Clouds' dropdown. The bottom left corner indicates 'Viewing 4 results of 4'.

4 On the **Snapshot Management** page, click **View Granules**.



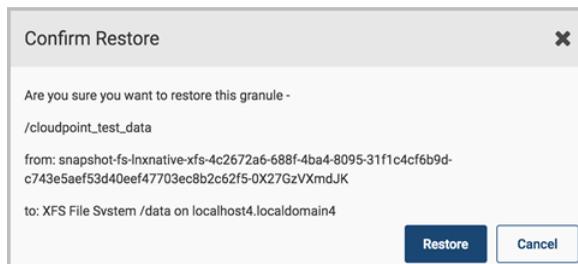
Note: The **View Granules** option is available only after indexing and classification is complete.

- 5 On the **Granule Management** page, select one or more files to restore and then click **Restore Selected Granules**.



The Granule Management page displays a list of all the files that are included in the indexed snapshot. If there are a large number of indexed granules, you can sort the files using the **Filter** field at the top of the page. CloudPoint currently does not provide any other methods (for example, alphabetical or time-based list) to sort the granules.

- 6 On the **Confirm Restore** page, select **Restore**.



Note: When you restore a granule, the existing copy is overwritten.

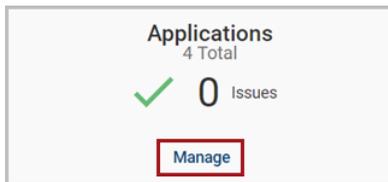
Deleting a snapshot

Regardless of the asset type you work with, the steps for deleting a snapshot are the same.

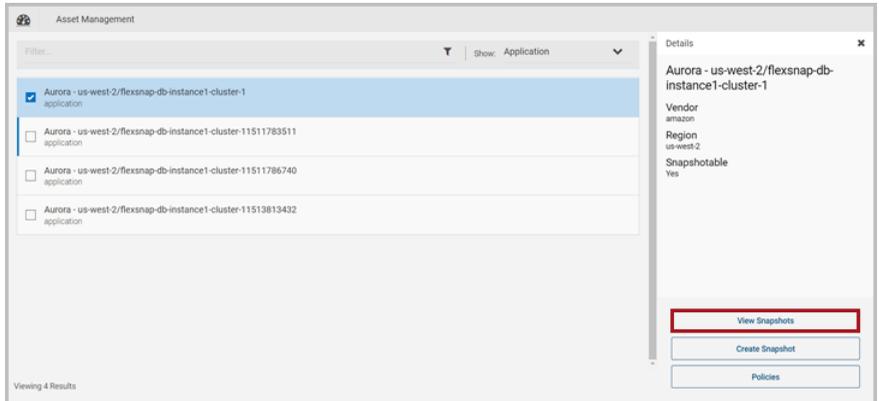
To delete a snapshot

- 1 Navigate to your list of assets.

On the CloudPoint dashboard, in the **Environment** card, locate the asset type you want to work with and click its **Manage** link. This example deletes an application snapshot.

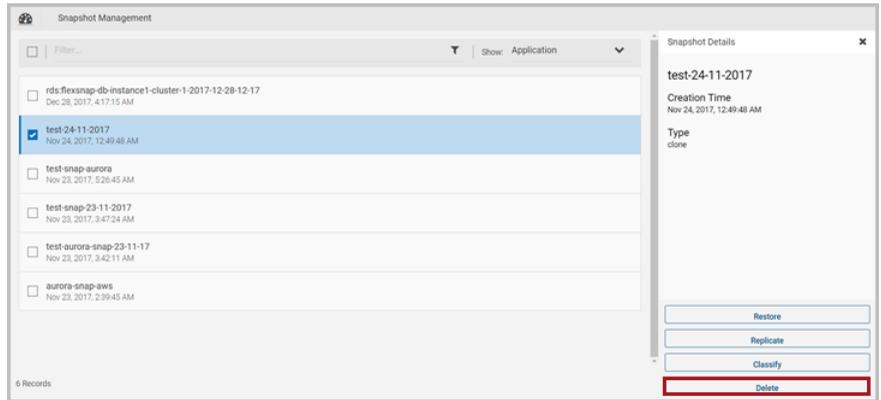


- 2 On the **Asset Management** page, select the application whose snapshot you want to restore. You can select multiple applications.



- 3 On the **Details** page click **View Snapshots**.

- 4 On the **Snapshot Management** page, select the snapshot (or snapshots) you want to delete. You can select multiple snapshots.



- 5 Depending on the structure of the snapshot, do one of the following:
 - If the snapshot does not have any sub-assets, click **Delete**.
 - If the snapshot has sub-assets, a **Snapshot Assets** page is displayed. By default, all sub-assets are checked. Select the sub-assets you want to delete and click **Delete**.
- 6 On the **Please Confirm ...** dialog box, click **Delete**.
 CloudPoint displays a message that the snapshot has been deleted.
 The snapshot is removed from the **Snapshot Management** page.

Monitoring activities with notifications and the job log

This chapter includes the following topics:

- [Working with notifications](#)
- [Using the Job Log](#)

Working with notifications

CloudPoint notifies you if any of the following occur:

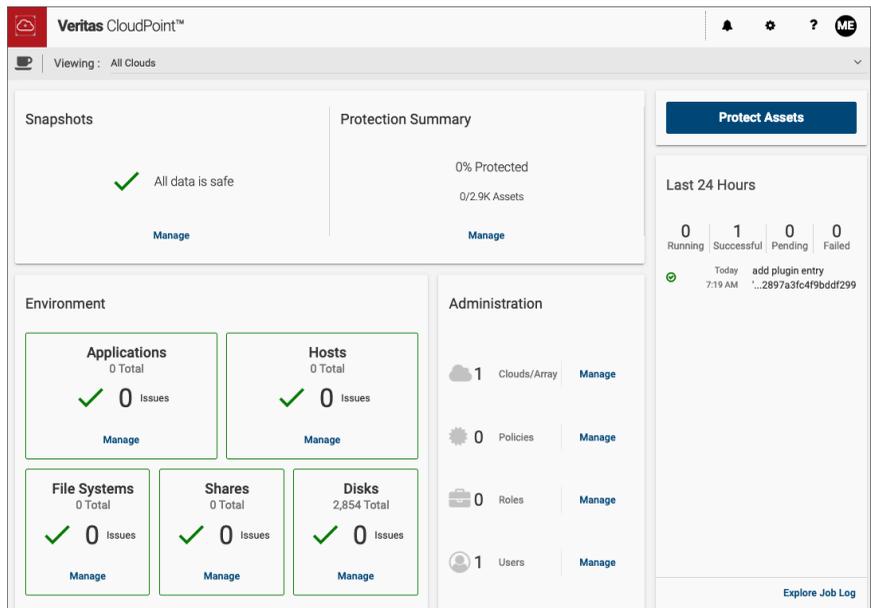
- Your role changes; for example if any of your permissions are added, deleted, or changed.
- The assets that are assigned to you change; for example, if assets are added, removed, or their policies change.
- If an operation on one of your assigned assets fails; for example, if a snapshot, restore, or replication fails.

CloudPoint writes notifications to the **Notifications** panel. To access the **Notifications** panel, click the bell icon at the top of the CloudPoint dashboard. The bell icon also indicates how many messages are in the log.

CloudPoint displays notifications for the last 7 days.

Using the Job Log

The CloudPoint Job Log allows you to keep a track of all the activities that are happening in your CloudPoint environment. These include manual and policy-driven tasks such as snapshot creation and deletion, as well as plug-in configuration and removal operations. On the right-hand side of the CloudPoint dashboard, you will see the Last 24 Hours pane that shows a preview of the most recent activities that have occurred.



The Last 24 Hours pane also displays the following information:

- Number of running tasks
- Number of successfully completed tasks
- Number of tasks that are pending
- Number of tasks that have failed

Towards the bottom of the pane is a link that takes you to the Job Log page where you can see more detailed information about all the tasks.

To use the Job Log

- 1 On the CloudPoint dashboard, in the **Last 24 Hours** panel, click **Explore Job Log**.
- 2 Review the **Job Log** page.

Job Type	Job Name	Started on
Export Snapshot Job	export snapshot NetApp-sharesnap-94709da3-77a4-43d3-8342-a3cf4f4d8f99	Apr 11, 2019, 4:39:46 PM
Deport Snapshot Job	deport snapshot NetApp-sharesnap-94709da3-77a4-43d3-8342-a3cf4f4d8f99	Apr 11, 2019, 4:39:28 PM
Export Snapshot Job	export snapshot NetApp-sharesnap-94709da3-77a4-43d3-8342-a3cf4f4d8f99	Apr 11, 2019, 4:36:10 PM
Export Snapshot Job	export snapshot NetApp-sharesnap-aae23852-39c1-4343-bace-0b0a89c42ce5	Apr 11, 2019, 4:30:17 PM
Create Snapshot Job	create snapshot (snap_Task_1) of share TASKV0276049_001	Apr 11, 2019, 4:29:31 PM
Delete Snapshot Job	delete snapshot (54315756693f492d47377647.Cp_scale62873a3d1) of disk-rcsi-600a098054315756693f492...	Apr 11, 2019, 4:26:28 PM
Delete Snapshot Job	delete snapshot (54315756693f492d47377649.Cp_scale62873a3d3) of disk-rcsi-600a098054315756693f492...	Apr 11, 2019, 4:26:28 PM
Delete Snapshot Job	delete snapshot (54315756693f492d47377648.Cp_scale62873a3d2) of disk-rcsi-600a098054315756693f492...	Apr 11, 2019, 4:26:28 PM
Delete Snapshot Job	delete snapshot (54315756693f492d47377646.Cp_scale62873a3d0) of disk-rcsi-600a098054315756693f492...	Apr 11, 2019, 4:26:28 PM
Delete Snapshot Job	delete snapshot (54315756693f492d4737764a.Cp_scale62873a3d4) of disk-rcsi-600a098054315756693f492...	Apr 11, 2019, 4:26:28 PM

The Job Log page displays a tabular list of all the tasks. Each log entry includes the following:

- An icon that indicates the job status--whether completed successfully, completed with errors, failed, or in progress
 - The job name, includes details about the assets that are involved in the task
 - The job type
 - The job start time and the end time (if applicable)
- 3 Use the filter and sorting tools as needed to locate the job you are interested in. You can filter the jobs based on the job type, job name, or job status.
 - 4 Click anywhere in a job row to see detailed information about that job. The Job Details pane on the right displays additional details.

A successful job appears as follows:

Job Details ✕

✓ Job Completed Successfully

Create Snapshot Job

create snapshot (CPAUTOSNAPf4bbc) of host DND_AUTOMATION_1

Task ID

79d5edf5-2586-447e-9554-9496ad4345ae

Started on Dec 26, 2018, 5:05:13 PM

Ended on Dec 26, 2018, 5:05:34 PM

Summary

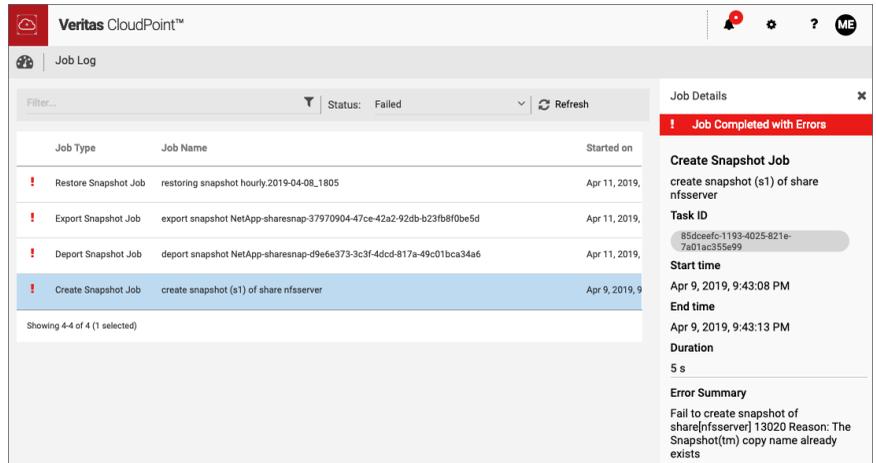
EC2 Snapshot of i-04bbdf126091076e8 (ami-08bf1bd07c553f973)

Snapshot type: clone

Depending on the type of job, the Job Details pane displays the following information:

- The job type
- A description of the job
- A task ID that uniquely identifies the job
- Job start time
- Job end time (applicable only if the job has completed)
- The time it took for the job to complete (applicable only if the job has completed)
- A summary of the underlying tasks
- The type of snapshot that was involved in the job (applicable only for create snapshot operations)

A failed job appears as follows:



For jobs that have failed, the Job Details pane also displays an error summary that indicates why a particular job has failed.

Note: You may notice that the Job Log page displays several "delete snapshots of asset" messages on a regular basis. The frequency of such messages is higher especially during a policy run. However, do not be alarmed.

During every policy run, CloudPoint performs the delete snapshots task to verify and ensure that the retention count specified in the policy configuration is maintained. These messages do not indicate a user-initiated snapshot deletion case.

Indexing and classifying your assets

This chapter includes the following topics:

- [About indexing and classifying snapshots](#)
- [Configuring classification settings using VIC](#)
- [Indexing and classifying snapshots](#)
- [Statuses for indexing and classification](#)

About indexing and classifying snapshots

Taking a snapshot protects your asset data, but does not give you insight into the data itself. You know the time that you created the snapshot and the asset that was protected, but little else. Knowing the content of the snapshot can be crucial. A snapshot may contain personally identifiable information (PII) and other sensitive data. If a snapshot contains sensitive data, you might treat it differently, or even delete it.

The classification feature lets you analyze your snapshot content, flag sensitive data, and take further actions as necessary.

Note: Classification is supported in CloudPoint Release 2.0.1 and later, and it is only available through the CloudPoint Enterprise license.

Indexing creates an index of the files in a snapshot. Having an index of the files enables you to restore a single file from a snapshot. Classification goes deeper into the data than indexing. During classification, indexing is performed automatically before the classification process identifies items that contain tags from the Veritas

Information Classifier. Tags indicate the type of data that is in a file, such as a credit card number, but not the actual data. For any snapshot, you can choose to index without classifying or to index and classify.

After a snapshot has been classified, you can always reclassify it. Reclassifying is useful if you have changed the settings in the Veritas Information Classifier since the last classification of a snapshot. During reclassification, CloudPoint can reclassify the snapshot contents based on the newly enabled or disabled classification policies in VIC and then display all the tags that are assigned to the files.

Considerations for indexing and classifying snapshots

Consider the following when you work with indexing and classification:

- Indexing and classification options are not available for file systems that are discovered from a Windows system.
- Indexing and classification are supported on Amazon Web Services (AWS) cloud, Microsoft Azure, and Google Cloud Platform (GCP), and in the same region and the same cloud account or project as the CloudPoint server. Each account or project will need its own CloudPoint configuration.
- For indexing and classification operations to run successfully, the target instance must be running.
- Indexing and classification are supported only for file system snapshots that you take at the disk level.
- Indexing and classification operations are not performed on symbolic links (symlink or soft link) that are references to another file or directory in the form of an absolute or a relative path.
- Only one classification or indexing job can run at a time. Additional snapshots are put in a queue until the previous classification or indexing job completes.
- A snapshot that is in the process of being indexed cannot be classified. The indexing process must complete before classification can start.
- You cannot delete a snapshot, either manually or using a policy, if indexing or classification is in progress, or if a granular restore operation (SFR) is being performed on the snapshot.
Similarly, if a snapshot is being deleted, no other simultaneous operations can be triggered on that snapshot.
- Sometimes, a classification job might fail even if the indexing job on the asset has completed successfully. You might not see any snapshot granules in the CloudPoint UI.
In such cases, you may have to reinitiate the indexing and classification jobs on the same asset again.

- Classification is supported for a maximum file size of 128 MB.
- All Veritas Information Classifier (VIC) policies are disabled by default. Before you trigger a classification job from CloudPoint, ensure that at least one classification policy is enabled in VIC.

See “[Configuring classification settings using VIC](#)” on page 154.

See “[Indexing and classifying snapshots](#)” on page 154.

Configuring classification settings using VIC

Veritas Information Classifier (VIC) lets you classify items based on their content and metadata. The classification tags that are configured in VIC are used when you select the **Classify** option or the **Index and Classify** option in CloudPoint.

To configure classification settings using VIC

- 1 On the CloudPoint dashboard, click the **Settings** (gear) icon, and then select **VIC Settings**.
- 2 You may be prompted to confirm that you want to leave CloudPoint and go to the Veritas Information Classifier. Click **Leave** to launch Veritas Information Classifier in a separate browser window.
- 3 In the Veritas Information Classifier UI, from the left-hand side menu, click **Tags**.

The UI displays all the built-in tags that are included in VIC.

- 4 Use the built-in tags or set up custom tags as required.

All policies in VIC are disabled by default. You must enable a policy if you want VIC to check for and tag the items that match the policy.

Refer to the VIC documentation for more details:

https://veritashelpsupport.com/Welcome?locale=EN_US&context=VIC2.1.3

Indexing and classifying snapshots

This section describes how you can index and classify snapshots manually.

Before you attempt to index or classify a snapshot, ensure that you:

- review the considerations for using these features
See “[About indexing and classifying snapshots](#)” on page 152.
- enable at least one classification policy in Veritas Information Classifier (VIC)
See “[Configuring classification settings using VIC](#)” on page 154.

Note: Classification is supported in CloudPoint Release 2.0.1 and later, and it is only available through the CloudPoint Enterprise or an equivalent license.

To index and classify a snapshot

- 1 Navigate to the asset that contains the snapshots you want to index or classify.
 - In the CloudPoint UI, click Dashboard and from under the Environment section, locate the **File Systems** area, and click **Manage**.
 - Select a file system asset and then go to the disk level snapshots of the file system asset.
- 2 On the snapshots page, select the snapshot, and then do one of the following:
 - To index the snapshot without classifying it, click **Index Only**.
After the snapshot is indexed, you can select the option to classify it.
 - To index and classify the snapshot in one step, click **Index and Classify**.
- 3 (Optional) If you selected the **Index Only** option in step 2, click **Classify** if you want to classify this snapshot.
- 4 (Optional) If you want to reclassify this snapshot, click **Reclassify**.
Reclassifying is useful if you have changed the settings in the Veritas Information Classifier since the last classification of a snapshot.

See “[Statuses for indexing and classification](#)” on page 155.

Statuses for indexing and classification

When an indexing or a classification operation is being performed on a snapshot, the following states indicate the status of the operation. In the CloudPoint UI, click on an asset to see the status in the upper corner of right hand side panel.

Table 15-1 Statuses for indexing and classification

Status	Description
Classified	The classification process is complete. No tags were found.
Classified - Tags Found	The classification process is complete. Tags that are configured in the Veritas Information Classifier were found in the selected snapshot. These tags may require your attention or additional action.
Classifying	The classification process is in progress.
Classifying Failed	The classification process cannot be completed.

Table 15-1 Statuses for indexing and classification (*continued*)

Status	Description
Indexed	The indexing process is complete.
Indexing	The indexing process is in progress.
Indexing - Classification Queued	The indexing process is in progress. The classification process begins when the indexing progress is complete. This status appears only if you selected Index and Classify .
Indexing Failed	The indexing process failed.
Unindexed	The selected snapshot has not been indexed yet. Click Index or Index and Classify to index the snapshot.

Note: Classification is supported in CloudPoint Release 2.0.1 and later, and it is only available through the CloudPoint Enterprise or an equivalent license.

See [“Indexing and classifying snapshots”](#) on page 154.

Protection and disaster recovery

This chapter includes the following topics:

- [About protection and disaster recovery](#)
- [Backing up CloudPoint](#)
- [Restoring CloudPoint](#)

About protection and disaster recovery

As of CloudPoint 2.0.x, CloudPoint cannot protect itself from disaster scenarios. This section describes how to backup and recover CloudPoint in case of a disaster.

Backing up CloudPoint

CloudPoint deployed in a cloud

To back up CloudPoint when it is deployed in a cloud

- 1 Sign out of the CloudPoint user interface (UI).
- 2 Stop CloudPoint services.

Use the following command:

```
# sudo docker run -it --rm
-v /full_path_to_volume_name:
/full_path_to_volume_name
-v /var/run/docker.sock:/var/run/docker.sock
veritas/flexsnap-cloudpoint:version stop
```

Here, *version* represents the currently installed CloudPoint product version.

For example:

```
# sudo docker run -it --rm -v /cloudpoint:/cloudpoint
-v /var/run/docker.sock:/var/run/docker.sock
veritas/flexsnap-cloudpoint:2.0.1.5300 stop
```

Note: This is a single command. Ensure that you enter the command without any line breaks.

Use the following API to determine CloudPoint version installed and configured on user setup:

```
# curl -H "Content-Type: application/json" -H
"Authorization: Bearer $token" -X GET -k
https://localhost:443/cloudpoint/api/v2/version
{ "Version": "2.1.0.7425",
"Commit": "b3916cd6fd62039f8f6dbf0dc1afd625f2066431" }
```

Use the following API to get the CloudPoint authentication token

```
# curl -k -X POST -H 'Content-type: application/json'
-d '{"email": "<email>", "password": "<pass>" }'
-k https://localhost/cloudpoint/api/v2/idm/login/
```

- 3 Make sure that all CloudPoint containers are stopped. This step is important because all activity and connections to and from CloudPoint must be stopped to get a consistent CloudPoint backup.

Enter the following:

```
# docker ps | grep veritas
```

This command should not return any actively running CloudPoint containers.

- 4 (Optional) If you still see any active containers, repeat step 3. If that does not work, run the following command on each active container:

```
# docker kill container_name
```

For example:

```
# docker kill flexsnap-api
```

- 5 After all the containers are stopped, take a snapshot of the volume on which you installed CloudPoint. Use the cloud provider's snapshot tools.
- 6 After the snapshot completes, restart CloudPoint services.

Use the following command:

```
# sudo docker run -it --rm -v /full_path_to_volume_name:  
/full_path_to_volume_name  
-v /var/run/docker.sock:/var/run/docker.sock  
veritas/flexsnap-cloudpoint:version start
```

Here, *version* represents the currently installed CloudPoint product version.

For example:

```
# sudo docker run -it --rm -v /cloudpoint:/cloudpoint  
-v /var/run/docker.sock:/var/run/docker.sock  
veritas/flexsnap-cloudpoint:2.0.1.5300 start
```

Note: This is a single command. Ensure that you enter the command without any line breaks.

CloudPoint deployed on-premise

To backup CloudPoint when it is deployed on-premise

- 1 Sign out of the CloudPoint user interface (UI).
- 2 Stop CloudPoint services.

Use the following command:

```
# sudo docker run -it --rm
-v /full_path_to_volume_name:/full_path_to_volume_name
-v /var/run/docker.sock:/var/run/docker.sock
veritas/flexsnap-cloudpoint:version stop
```

```
# sudo docker run -it --rm
```

Here, *version* represents the currently installed CloudPoint product version.

For example:

```
# sudo docker run -it --rm -v /cloudpoint:/cloudpoint
-v /var/run/docker.sock:/var/run/docker.sock
veritas/flexsnap-cloudpoint:2.0.1.5300 stop
```

Note: This is a single command. Ensure that you enter the command without any line breaks.

- 3 Make sure that all CloudPoint containers are stopped. This step is important because all activity and connections to and from CloudPoint must be stopped to get a consistent CloudPoint backup.

Enter the following:

```
# docker ps | grep veritas
```

This command should not return any actively running CloudPoint containers.

- 4 (Optional) If you still see any active containers, repeat step 3. If that does not work, run the following command on each active container:

```
# docker kill container_name
```

For example:

```
# docker kill flexsnap-api
```

- 5 Back up the folder `/cloudpoint`. Use any backup method you prefer.

For example:

```
# tar -czvf cloudpoint_dr.tar.gz /cloudpoint
```

This command creates a compressed archive file named `cloudpoint_dr.tar.gz` that contains the data in the `/cloudpoint` directory.

Restoring CloudPoint

You can restore CloudPoint using any of the following methods:

- Recover CloudPoint using a snapshot you have in the cloud
- Recover CloudPoint using a backup located on-premises

Using CloudPoint snapshot located in the cloud

To recover CloudPoint using a snapshot you have in the cloud

- 1 Using your cloud provider's dashboard or console, create a volume from the existing snapshot.
- 2 Create a new virtual machine with specifics equal to or better than your previous CloudPoint server.
- 3 Install docker on the new server.
See ["Installing CloudPoint"](#) on page 28.
- 4 Attach the newly-created volume to this CloudPoint server instance.
- 5 Create the CloudPoint installation directory on this server.

Use the following command:

```
# mkdir /full_path_to_cloudpoint_installation_directory
```

For example:

```
# mkdir /cloudpoint
```

- 6 Mount the attached volume to the installation directory you just created.

Use the following command:

```
# mount /dev/device-name  
/full_path_to_cloudpoint_installation_directory
```

For example:

```
# mount /dev/xvdb /cloudpoint
```

- 7 Verify that all CloudPoint related configuration data and files are in the directory.

Enter the following command:

```
# ls -l /cloudpoint
```

- 8 Download or copy the CloudPoint installer binary to the new server.

- 9 Install CloudPoint.

Use the following command:

```
# sudo docker run -it --rm  
-v /cloudpoint:/cloudpoint  
-v /var/run/docker.sock:/var/run/docker.sock  
veritas/flexsnap-cloudpoint:2.0.1.5300 install
```

Here, 2.0.1.5300 represents the CloudPoint version. Replace it as per your currently installed product version.

Note: This is a single command. Ensure that you enter the command without any line breaks.

The installation program detects an existing version of CloudPoint and re-installs all CloudPoint services without overwriting existing content.

Messages similar to the following are displayed on the command prompt:

```
Configuration started at time Fri May 4 22:20:47 UTC 2018  
This is a re-install.  
Checking if a 1.0 release container exists ...
```

Note the line that indicates that the operation is a re-install.

- 10 When the installation completes, you can resume working with CloudPoint using your existing credentials.

Using CloudPoint backup located on-premise

To recover CloudPoint using a backup located on-premise

- 1 Copy the existing CloudPoint backup to the new CloudPoint server and extract it to the CloudPoint installation directory.

In the following example, because `/cloudpoint` was backed up, the command creates a new `/cloudpoint` directory.

```
# tar -zxvf cloudpoint_dr.tar.gz -C /cloudpoint/
```

- 2 Download or copy the CloudPoint installer binary to the new server.
- 3 Install CloudPoint.

Use the following command:

```
# sudo docker run -it --rm  
-v /cloudpoint:/cloudpoint  
-v /var/run/docker.sock:/var/run/docker.sock  
veritas/flexsnap-cloudpoint:2.0.1.5300 install
```

Here, `2.0.1.5300` represents the CloudPoint version. Replace it as per your currently installed product version.

Note: This is a single command. Ensure that you enter the command without any line breaks.

The installation program detects an existing version of CloudPoint and re-installs all CloudPoint services without overwriting existing content.

Messages similar to the following are displayed on the command prompt:

```
Configuration started at time Fri May 4 22:20:47 UTC 2018  
This is a re-install.  
Checking if a 1.0 release container exists ...
```

Note the line that indicates that the operation is a re-install.

Note: When CloudPoint recovers, no licenses are installed. Hence, you must install the CloudPoint licenses manually. This is applicable if you are using CloudPoint version 2.1.x.

Maintaining CloudPoint

- [Chapter 17. CloudPoint logs](#)
- [Chapter 18. Troubleshooting CloudPoint](#)
- [Chapter 19. Upgrading CloudPoint](#)
- [Chapter 20. Working with your CloudPoint license](#)

CloudPoint logs

This chapter includes the following topics:

- [CloudPoint logs](#)

CloudPoint logs

CloudPoint maintains the following logs to monitor activity and troubleshoot issues. The logs are stored on the path `installation_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 17-1 CloudPoint logs

Log	Description
<code>flexsnap-agent-<agnet-id>.log</code>	<p>The log file for storing specific child agent configuration.</p> <p>There can be multiple child agent log files. The log file is generated after a new configuration is validated and the child agent is spawned to handle that configuration.</p> <p>The log file for the agents that stores all the error logs related to agent and the plugins that the agent is managing. The offhost-agent only deals with offhost plugins like AWS, Azure, GCP, or array plugins. All the tasks like discovering the assets, creating, restoring, and deleting snapshots which are done by the agent and the plugin are stored in this log file. The flexsnap-coordinator requests the agent services based on the asset type to create, restore, delete, or find asset in the cloud.</p>

Table 17-1 CloudPoint logs (*continued*)

Log	Description
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 API service to the appropriate agents.
flexsnap-telemetry.log	The log file for the telemetry service which contains information about service life cycle including successful telemetry operations as well as any errors related to that service.
init.log	The log for recording the installation activities.
flexsnap-classifier.log	<p>The log file for storing the error logs related to the classification and indexing activity performed on the snapshot. As the flexsnap-classifier interfaces with VIC and MongoDB you can also find logs related to connection to these containers.</p> <p>Note: This log file is available in CloudPoint Release 2.0.1 and later.</p>
flexsnap-agent-offhost.log	<p>The log file for the parent offhost agent that stores the error logs related to the new plugins configuration addition. This log file is generated by parent agent. Parent agent is a stand alone agent which validate to new plugin configuration which is not owned any configuration. It does not contain any specific plugin discovery log. This file contains initial configuration validation log before spawning child agent. Once child agent is spawned to handle plugin configuration, the configuration log is redirected to new log file with the name flexsnap-agent-<agentid>.log</p>

Table 17-1 CloudPoint logs (*continued*)

Log	Description
<code>flexsnap-agent-onhost.log</code>	<p>The log file for the agents that stores all the error logs related to agent and the plugins that the agent is managing. The onhost-agent deals with plugins that can run inside a host like the application plugins like Oracle, Linux, Mongo, and so on.</p> <p>All the tasks like discovering the assets, creating, restoring, deleting snapshots which are done by the agent and the plugin are stored in this log file. The flexsnap-coordinator requests the agent services based on the asset type to create, restore, delete, or find asset in the cloud.</p>
<code>email_service.log</code>	<p>The log file for storing the logs related to the email service. The log file stores the start up information of the service, the RabbitMQ calls made to the service, connection issues while setting up RabbitMQ. and any errors during an internal call.</p>
<code>identity_manager_service.log</code>	<p>The log file for the Identity Management Service (IDM). It stores the logs for any REST call received by IDM (create user, modify user, login), any requests over RabbitMQ received by IDM (create prebake user, validate token), errors on IDM (unauthorized), step by step information of certain operations done by IDM.</p>
<code>flexsnap-indexingsupervisor.log</code>	<p>The log file for storing the logs related to coordinating the workflow for indexing and classification. The indexing supervisor service cooperates with the flexsnap-coordinator and flexsnap-classifier service to index and/or classify snapshots. The indexing supervisor is responsible for queuing and subsequently running indexing and classification jobs.</p>
<code>nginx_access.log</code>	<p>The log file is generated by nginx web-server.</p>
<code>nginx_error.log</code>	<p>The log file is generated by the nginx web-server.</p>
<code>api-gateway.log</code>	<p>The log file for storing the details of the proxy that routes requests/responses between the application's web console and back-end services. This log file is configured by the API and not from the flexsnap.conf file.</p>

Troubleshooting CloudPoint

This chapter includes the following topics:

- [Restarting CloudPoint](#)
- [Docker may fail to start due to a lack of space](#)
- [Some CloudPoint features do not appear in the user interface](#)

Restarting CloudPoint

If you need to restart CloudPoint after an error, it's important that you restart it correctly so that your environmental data is preserved.

Warning: Do not use commands such as `docker restart` or `docker stop` and `docker start` to restart CloudPoint. Use the `docker run` command described below.

To restart CloudPoint

- ◆ On the instance where CloudPoint is installed, enter the following command:

```
# sudo docker run -it --rm -v
/cloudpoint:/cloudpoint -v /var/run/docker.sock:/var/run/docker.sock
veritas/flexsnap-cloudpoint:version restart
```

Here, *version* represents the currently installed CloudPoint product version.

For example:

```
# sudo docker run -it --rm -v
/cloudpoint:/cloudpoint -v /var/run/docker.sock:/var/run/docker.sock
veritas/flexsnap-cloudpoint:2.0.2.4815 restart
```

Note: This is a single command. Ensure that you enter the command without any line breaks.

Docker may fail to start due to a lack of space

During CloudPoint deployment, the Docker image may fail to start if there is not enough space for the MongoDB database. The failure occurs after you enter the `docker run` command.

Workaround:

The following procedure shows the steps to take if the image fails to start.

- 1 Check the log file `/mount-point-from-host/logs/init.log`.

Note that MongoDB starts, then immediately stops. (See the information messages in bold.)

```
# sudo cat /mount-point-from-host/logs/init.log
Oct 03 11:24:45 init:INFO - Veritas CloudPoint init process starting up.
Oct 03 11:24:45 init:INFO - Veritas CloudPoint init process starting up.
Oct 03 11:24:45 init:INFO - Started mongod[9]
Oct 03 11:24:45 init:INFO - Started mongod[9]
Oct 03 11:24:45 init:INFO - mongod already stopped, 100
Oct 03 11:24:45 init:INFO - mongod already stopped, 100
```

- 2 Verify the amount of available space on the host boot disk. MongoDB needs about 4 GB of space.

In the following example, only 1.6 GB is available.

```
# sudo df -kh /
Filesystem      Size  Used Avail Use% Mounted on
/dev/xvda1      7.7G  6.2G  1.6G  80% /
```

- 3 Free up space on the boot disk.
- 4 After the boot disk has more than 4.0 GB of available space, restart the container.

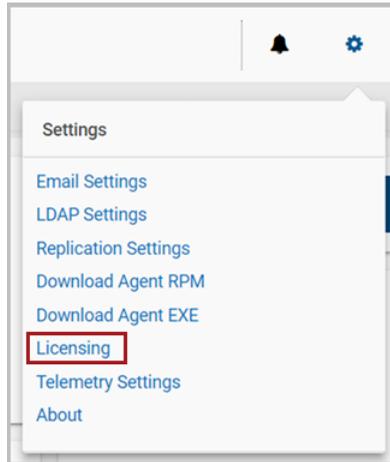
```
# sudo docker restart container-id
```

Some CloudPoint features do not appear in the user interface

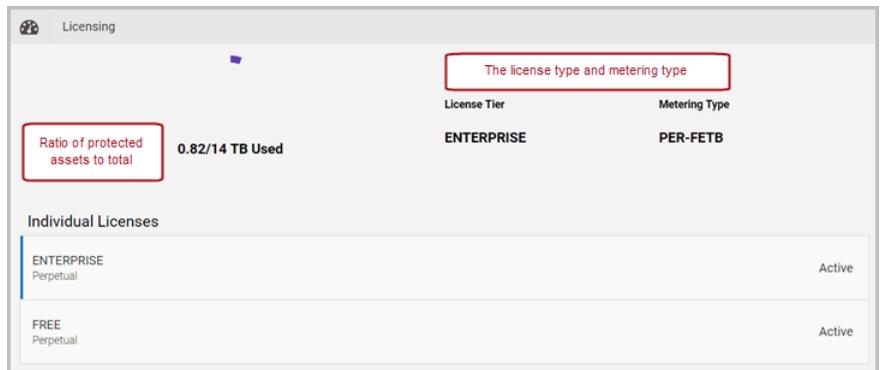
If certain CloudPoint features do not appear in the user interface, the first step is to verify which CloudPoint license you have. The license type determines which features you can access.

To display your CloudPoint license type

- 1 From the top of any CloudPoint page click the **Settings** icon (gear) and select **Licensing**.



- 2 On the Licensing page, note the type of license you have.



- 3 Review the features supported by your license.
See [“Understanding your CloudPoint license”](#) on page 11.
- 4 If your license does not support the feature you want, consider upgrading your license.
See [“Upgrading your CloudPoint license”](#) on page 180.

Upgrading CloudPoint

This chapter includes the following topics:

- [About CloudPoint upgrades](#)
- [Preparing to upgrade CloudPoint](#)
- [Upgrading CloudPoint](#)

About CloudPoint upgrades

Two versions of CloudPoint on two different hosts should not manage the same assets.

When you upgrade CloudPoint, all the snapshot data and configuration data from your previous version is maintained in the external `/cloudpoint` data volume. We strongly recommend that you upgrade CloudPoint on the same host or on a different host to which the CloudPoint data volume of the previous version is attached.

Supported upgrade path

The following table displays the supported upgrade paths for CloudPoint.

Table 19-1 CloudPoint upgrade path

Upgrade from version	Upgrade to version
<ul style="list-style-type: none">▪ 2.0.2▪ 2.1	2.1.2

Preparing to upgrade CloudPoint

Note the following before you upgrade CloudPoint:

- Ensure that the virtual machine or physical host meets the requirements of the CloudPoint version that you wish to upgrade to.
See “[Meeting system requirements](#)” on page 19.
- When you upgrade CloudPoint, all the snapshot data and configuration data from your previous version is maintained in the external `/cloudpoint` data volume. This information is external to the CloudPoint container and the image and is preserved during the upgrade.
However, you can take a backup of all the data in the `/cloudpoint` volume, if desired.
See “[Backing up CloudPoint](#)” on page 158.

Upgrading CloudPoint

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

To upgrade CloudPoint

- 1 Make sure that the CloudPoint host (physical host, virtual machine or a cloud instance) meets the requirements of the new CloudPoint version.
See “[Meeting system requirements](#)” on page 19.
- 2 Open the Veritas CloudPoint trial page.
In your browser's address bar, type the following URL:
<https://www.veritas.com/trial/en/us/cloud-point.html>
- 3 On the trial page, provide the requested details and then click **Submit** to register.
- 4 On the CloudPoint download page, click **Download Now** to download the CloudPoint installer.

The CloudPoint software components are available in the form of Docker images and these images are packaged in a compressed file. The file name has the following format:

```
Veritas_CloudPoint_2.x.x_IE.img.gz
```

The numerical sequence in the file name represents the CloudPoint product version.

Note: The actual compressed image file name may vary depending on the product release version.

- 5 Copy the downloaded compressed image file to the computer on which you want to deploy CloudPoint.
- 6 Load the image file using the following command:

```
# docker load -i <imagefilename>
```

For example, if the CloudPoint version is 2.1.2, the command syntax is as follows:

```
# docker load -i Veritas_CloudPoint_2.1.2_IE.img.gz
```

Messages similar to the following appear on the command line:

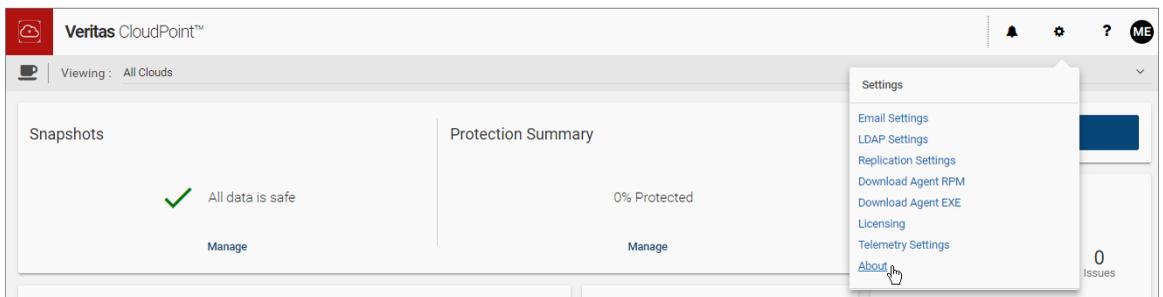
```
644879075e24: Loading layer [=====>] 117.9MB/117.9MB
d7ff1dc646ba: Loading layer [=====>] 15.87MB/15.87MB
d73dd9qwer58: Loading layer [=====>] 1.812GB/1.812GB
3167ba895aec: Loading layer [=====>] 352.9MB/352.9MB
fd22ad285778: Loading layer [=====>] 41.98kB/41.98kB
Loaded image: veritas/flexsnap-cloudpoint:2.1.2.7542
```

Make a note of the loaded image name and version that appears on the last line. This represents the new CloudPoint version that you wish to upgrade to. You will need this information in the subsequent steps.

Note: The version displayed here is used for representation only. The actual version will vary depending on the product release you are installing.

- 7 Make a note of the current CloudPoint version that is installed. You will use the version number in the next step.

Log on to the CloudPoint user interface (UI) and from the top right corner, click **Settings** and then click **About**.



The Current Version field in the About dialog box displays the installed version.

- 8 From the Job Log page, verify that there are no protection policy snapshots or other operations in progress and then stop CloudPoint by running the following command:

```
# sudo docker run --rm -it -v /cloudpoint:/cloudpoint
-v /var/run/docker.sock:/var/run/docker.sock
veritas/flexsnap-cloudpoint:current_version stop
```

Here, *current_version* represents the currently installed CloudPoint version. Use the version number you noted in step 7 earlier.

For example, if the installed CloudPoint version is 2.0.2.4722, the command will be as follows:

```
# sudo docker run --rm -it -v /cloudpoint:/cloudpoint
-v /var/run/docker.sock:/var/run/docker.sock
veritas/flexsnap-cloudpoint:2.0.2.4722 stop
```

Note: This is a single command. Ensure that you enter the command without any line breaks.

The CloudPoint containers are stopped one by one. Messages similar to the following appear on the command line:

```
Stopping the services
Trying to stop container: flexsnap-mongodb
flexsnap-mongodb
Stopped container: flexsnap-mongodb
Trying to stop container: flexsnap-rabbitmq
flexsnap-rabbitmq
Stopped container: flexsnap-rabbitmq
Trying to stop container: flexsnap-auth
flexsnap-auth
Stopped container: flexsnap-auth
Trying to stop container: flexsnap-coordinator
flexsnap-coordinator
Stopped container: flexsnap-coordinator
...
```

Wait for all the CloudPoint containers to be stopped and then proceed to the next step.

9 Upgrade CloudPoint by running the following command:

```
# sudo docker run --rm -it -v /cloudpoint:/cloudpoint
-v /var/run/docker.sock:/var/run/docker.sock
veritas/flexsnap-cloudpoint:new_version install
```

Here, *new_version* represents the CloudPoint version you are upgrading to.

For example, using the version number specified in step 6 earlier, the command will be as follows:

```
# sudo docker run --rm -it -v /cloudpoint:/cloudpoint
-v /var/run/docker.sock:/var/run/docker.sock
veritas/flexsnap-cloudpoint:2.1.2.7542 install -y
```

Note: This is a single command. Ensure that you enter the command without any line breaks.

10 The new CloudPoint installer detects the existing CloudPoint containers that are running and asks for a confirmation for removing them.

Press **Y** to confirm the removal of the old CloudPoint containers.

The installer first loads the individual service images and then launches them in their respective containers.

Wait for the installer to display messages similar to the following and then proceed to the next step:

```
Trying to run docker container: flexsnap-cloudpointconsole
7cb1b17688a88098679de8a69fbec5d10fcf4b7035c0be2f4
Successfully ran docker container: flexsnap-cloudpointconsole
```

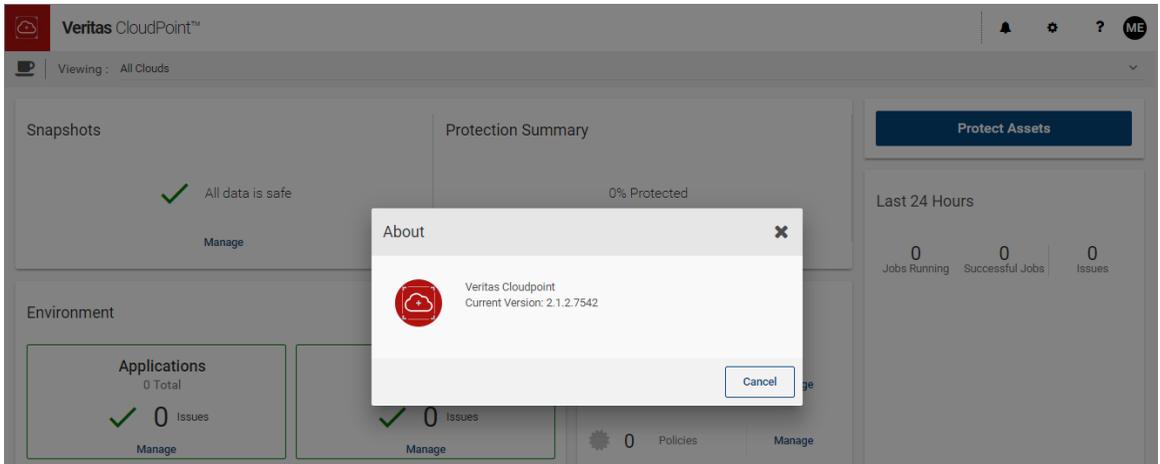
```
Please go the UI and configure CloudPoint now.
Waiting for CloudPoint configuration to complete ...
```

11 Refresh your web browser and log in to the CloudPoint user interface.

12 Verify the CloudPoint version.

From the UI, click on **Settings** from the top right corner and select **About**.

The Current Version field in the About dialog box should now indicate the new version you just installed.



13 This concludes the upgrade process on the host. Verify that your CloudPoint configuration settings and data are preserved as is.

14 After upgrading CloudPoint containers on the CloudPoint host, the next step is to upgrade the on-host agents on the Linux and Windows hosts.

Perform the following steps to upgrade the agent on Linux hosts:

- Download the newer version of the agent installation package by logging in to the CloudPoint UI.
- Stop the flexsnap agent service on the host where you want to upgrade the agent.

```
# sudo systemctl stop flexsnap-agent.service
```

- Upgrade the agent on the Linux host.

```
# sudo rpm -Uvh cloudpoint_agent_rpm_name
```

Here, *cloudpoint_agent_rpm_name* is the name of the on-host agent rpm package you downloaded earlier.

- Start the flexsnap agent service on the host.

```
# sudo systemctl start flexsnap-agent.service
```

- Reload the daemon, if prompted.

```
# sudo systemctl daemon-reload
```

- Repeat these steps on all the Linux hosts where you wish to upgrade the Linux-based on-host agent.

Perform the following steps to upgrade the agent on Windows hosts:

- Download the newer version of the agent installation package by logging in to the CloudPoint UI.
- Upgrade the agent on the Windows host.
Unzip the agent installation package on the host and then perform the steps mentioned in the following topic:
See [“Configuring the Windows-based on-host agent”](#) on page 70.
- Repeat these steps on all the Windows hosts where you wish to upgrade the Windows-based on-host agent.

For details on how to download the agent installation package from the CloudPoint UI, refer to the following:

See [“Downloading and installing the on-host agent”](#) on page 64.

- 15** If your CloudPoint deployment is integrated in to a Veritas NetBackup environment, the next step is to refresh the NetBackup configuration so that the upgraded CloudPoint configuration details are available with NetBackup.
Refer to the Veritas NetBackup documentation for instructions on how to refresh the configuration.

Working with your CloudPoint license

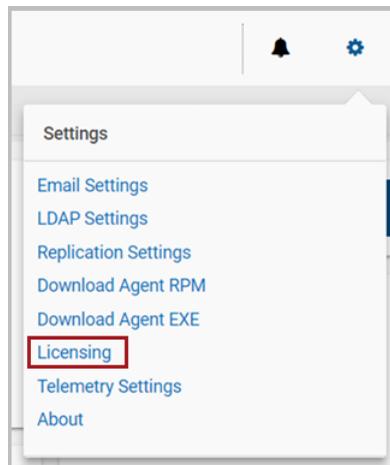
This chapter includes the following topics:

- [Displaying CloudPoint license and protection information](#)
- [Upgrading your CloudPoint license](#)

Displaying CloudPoint license and protection information

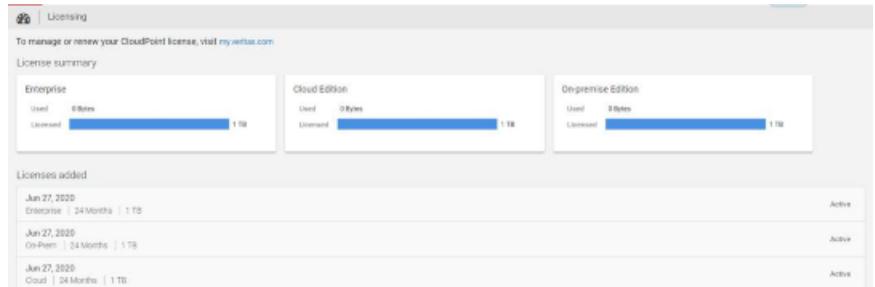
To display CloudPoint license and protection information

- 1 From the **Settings** drop-down list, select **Licensing**.



- 2 Review the **Licensing** page. Note the following:

- Under the **License summary** you can view the type of license in effect and the amount of license used.
- Under **License summary**, you can, view the license metering type; Instance or FETB, current license in effect, current consumption, number of remaining months in case of subscription based licensing, and the last date.
When you upgrade from free license to paid license, your free license consumption is transferred to the paid license.



See [“Understanding your CloudPoint license”](#) on page 11.

See [“Upgrading your CloudPoint license”](#) on page 180.

Upgrading your CloudPoint license

CloudPoint is distributed with a free license. It does not expire, and it gives you a chance to try out a subset of features in your preferred cloud. This license lets you protect up to 10 TB of front-end terra byte data (FETB).

CloudPoint also offers three paid subscription licenses. If you need more advanced features, you can upgrade your license and unlock the bundle that is right for you. CloudPoint's paid licenses are the following:

- Enterprise - This license lets you take application-consistent snapshots of your workloads, such as Oracle, SQL, and Amazon Web Services (AWS). This license also gives you advanced features such as snapshot replication.
- Cloud - This license supports only cloud plug-ins. It lets you take application-consistent snapshots of your workloads, such as AWS, GCP, and Azure.
- On-prem - This license supports only on-prem plug-ins. It lets you take application-consistent snapshots of your workloads, such as array plug-ins, hypervisor, and so on.

Your Veritas representative can help you decide which paid license is right for you.

A CloudPoint license is an XML file with a `.slf` file extension.

See [“Understanding your CloudPoint license”](#) on page 11.

To upgrade your CloudPoint license

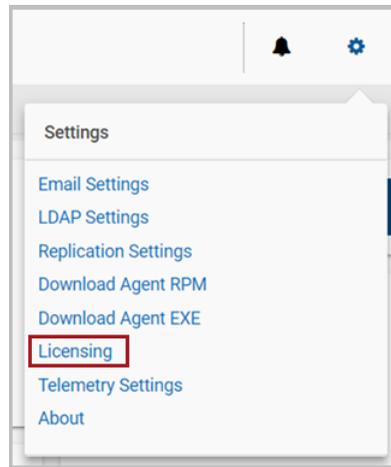
- 1 Use the download link that is provided by your Veritas representative to download the license file to your local machine. If necessary, copy the license file to the machine from where you will access the CloudPoint user interface.

The following example upgrades the CloudPoint Basic license to an Enterprise license.

- 2 Sign in to the CloudPoint user interface.

See [“Signing in to CloudPoint”](#) on page 101.

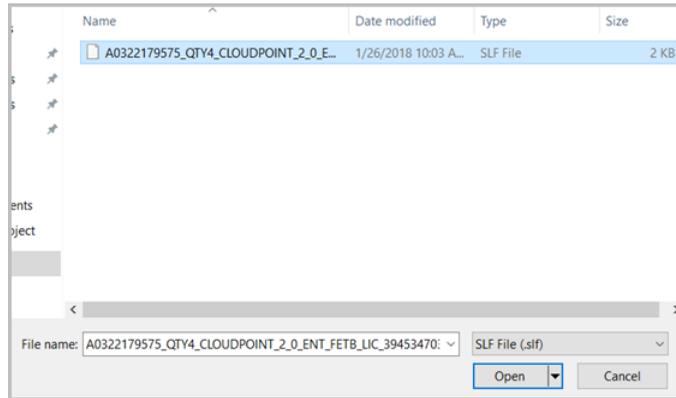
- 3 From the **Settings** drop-down list, select **Licensing**.



- 4 On the **Licensing** page, click **Upload License**.

- 5 On the **Upload License** dialog box, click **Select File**.

- 6 Navigate to the location where you copied the license file earlier, select the file, and then click **Open**.



- 7 On the **Upload License** dialog box, click **Upload File**.



- 8 The **License** page lists the new license. The following example shows that the Enterprise license is active and in effect. The license is measured in terms of front-end terabyte (FETB) data. You can also purchase an Enterprise license based on the number of instances to protect.

Licensing

To manage or renew your CloudPoint license, visit my.veritas.com

License summary

Enterprise	Cloud Edition	On-premise Edition
Used 0 Bytes Licensed 1 TB	Used 0 Bytes Licensed 1 TB	Used 0 Bytes Licensed 1 TB

Licenses added

Jun 27, 2020 Enterprise 24 Months 1 TB	Active
Jun 27, 2020 On-Prem 24 Months 1 TB	Active
Jun 27, 2020 Cloud 24 Months 1 TB	Active

See [“Understanding your CloudPoint license”](#) on page 11.

See [“Displaying CloudPoint license and protection information”](#) on page 179.

Reference

- [Chapter 21. Storage array support](#)
- [Chapter 22. Working with CloudPoint using APIs](#)

Storage array support

This chapter includes the following topics:

- [Dell EMC Unity arrays](#)
- [Hewlett Packard Enterprise \(HPE\) 3PAR array](#)
- [Pure Storage FlashArray](#)
- [Huawei OceanStor arrays](#)

Dell EMC Unity arrays

This section describes the following:

- The parameters you must supply to configure the Dell EMC Unity array plug-in
- The Dell EMC Unity arrays that CloudPoint supports
- The CloudPoint operations you can perform on Dell EMC Unity array assets

Dell EMC Unity array plug-in configuration parameters

When you configure the Dell EMC Unity array plug-in, specify the parameters shown in the following table.

Table 21-1 Dell EMC Unity array plug-in configuration parameters

CloudPoint configuration parameter	Description
Array IP Address	The array's IP address
Username	The user name used to access the array
Password	The password used to access the array

Before you configure the plug-in, ensure that the specified user account has permissions to create, delete, and restore snapshots on the array.

Supported Dell EMC Unity arrays

You can use CloudPoint to discover and protect the following Dell EMC Unity array models.

Table 21-2 Supported EMC arrays

Category	Supported
Array model	Unity 600 Theoretically, other models will work also because CloudPoint does not include any model-specific coding. Other models include the following: <ul style="list-style-type: none">■ Unity 300 and Unity 300F ("F" indicates that it is a flash array)■ Unity 400 and Unity 400F■ Unity 500 and Unity 500F■ Unity 600F
Software	UnityOS
Firmware version	4.2.1.9535982 or later Refer to the array-specific documentation for more information on firmware versions and how to check the current firmware on your array.
Library	<code>storops</code> Note: CloudPoint automatically installs all the required libraries during installation.

Supported CloudPoint operations on Dell EMC Unity arrays

You can perform the following CloudPoint operations on supported Dell EMC Unity arrays:

- List all the disks.
- Create a copy-on-write (COW) snapshot of a LUN.

Note: Snapshot name can be lowercase or uppercase, can contain any ASCII character, and can include special characters.

- Export snapshot
When a snapshot is exported, CloudPoint attaches the snapshot to the target host and keeps a track of it using the export ID.
- Deport snapshot
When a snapshot is deported, CloudPoint detaches the exported snapshot from the target host and removes the export ID.
- Delete a COW snapshot of a LUN.
- Restore a LUN using a COW snapshot. The snapshot overwrites the original object.

Note: You cannot snapshot LUNs which are under a consistency group. The reason for this limitation is that to restore a single LUN snapshot would restore the entire consistency group.

Snapshot export related requirements and limitations

The following requirements and limitations are applicable in a Dell EMC Unity array environment:

- The host on which the snapshot is to be exported must be attached to the array.

Note: The exported snapshot is attached to the host and is accessible using a world wide name (WWN) that is assigned by the array.

- Snapshot export is supported using the following protocols:
 - Fibre Channel (FC)
 - Internet Small Computer Systems Interface (iSCSI)
- A snapshot cannot be exported multiple times.
- An exported snapshot cannot be deleted.
- The CloudPoint user interface (UI) does not support running the snapshot export and deport operations.

Use the following CloudPoint API to perform these operations:

```
(POST) /v2/assets/<disk-ID>/snapshots/<snap-id>/exports/  
(DELETE) /v2/assets/<disk-ID>/snapshots/<snap-id>/exports/<export-ID>
```

Here are some sample cURL commands:

For Export:

```
curl -H "Content-Type: application/json" -H "Authorization: Bearer  
<token>" -X POST -d '{"host-name":"offhost_server", "protocol":"fc",  
"port":"<uuid given to host>"}' -k  
https://localhost/cloudpoint/api/v2/assets/<disk-ID>/snapshots/<snap-id>/exports/
```

For Deport:

```
curl -H "Content-Type: application/json" -H "Authorization: Bearer  
<token>" -X DELETE -k  
https://localhost/cloudpoint/api/v2/assets/<disk-id>/snapshots/<snap-id>/exports/<export-id>
```

You can access the CloudPoint REST APIs using Swagger.

See ["Accessing the Swagger-based API documentation"](#) on page 197.

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

3PAR array plug-in configuration parameters

Plug-in configuration prerequisites

Ensure that the following prerequisites are fulfilled before you configure the CloudPoint plug-in:

- The HPE 3PAR Web Services API (WSAPI) is installed and enabled on the array.
CloudPoint plug-in uses the `python-3parclient` Python library package for all its communications with the 3PAR storage arrays. This Python library uses WSAPI.
<https://github.com/hpe-storage/python-3parclient>
- The specified user account has read privileges on the array and is able to perform the following actions:
 - create snapshot
 - delete snapshot
 - promote snapshot
 - export snapshot
 - deport snapshot

When you configure the 3PAR array plug-in, specify the following parameters:

Table 21-3 HPE 3PAR array configuration parameters

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

Supported 3PAR arrays

Table 21-4 Supported HPE 3PAR arrays

Category	Supported
Array model	HP_3PAR 8200
Firmware version	3.1.3 or later Refer to the array-specific documentation for more information on firmware versions and how to check the current firmware on your array.
Required software development kit	HP 3PAR Management Console 4.5.0
Library	<code>hpe3parclient</code>

Note: CloudPoint plugin uses the `python-3parclient` Python library package for all its communications with the 3PAR storage arrays. CloudPoint supports any firmware and 3PAR array that works with this Python library.

<https://github.com/hpe-storage/python-3parclient>

Supported CloudPoint operations on 3PAR array assets

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

- List all the disks.
- Create a copy-on-write (COW) virtual copy or clone (physical copy) snapshots of a volume.

Note: Snapshot name must be between 1 through 31 characters in length. For a snapshot of a volume set, use name patterns that are used to form the snapshot volume name. Refer to VV Name Patterns in the HPE 3PAR Command Line Interface Reference available from the HPE Storage Information Library.

- Delete a COW virtual copy or clone physical snapshots of a volume.
- Restore COW virtual copy snapshots of a volume, overwriting the original object.

Pure Storage FlashArray

This section describes the following:

- The parameters you must supply to configure the Pure Storage FlashArray plug-in
- The FlashArray models that CloudPoint supports
- The CloudPoint operations you can perform on FlashArray assets

Pure Storage FlashArray plug-in configuration parameters

When you configure the Pure Storage FlashArray plug-in, specify the parameters shown in the following table.

Table 21-5 Pure Storage FlashArray plug-in configuration parameters

CloudPoint configuration parameter	Description
IP Address	The array's IP address
Username	The user name used to access the array
Password	The password used to access the array

Before you configure the plug-in, ensure that the specified user account has permissions to create, delete, and restore snapshots on the array.

Supported Pure Storage FlashArray models

You can use CloudPoint to discover and protect the following Pure Storage FlashArray models.

Table 21-6 Supported Pure Storage FlashArray models

Category	Supported
Array model	FA-405
Firmware version	4.10.6 Refer to the array-specific documentation for more information on firmware versions and how to check the current firmware on your array.

Supported CloudPoint operations on Pure Storage FlashArray models

You can perform the following CloudPoint operations on supported Pure Storage FlashArray models:

- Discover and list all volumes.
- Create a clone snapshot of a volume.

Note: A snapshot name comprises of "Diskname+ snapshotname". Snapshot suffix must be between 1 through 63 characters in length and can be alphanumeric. The snapshot name must begin and end with a letter or number. The suffix must include at least one letter or '-'.

- Delete a clone snapshot.
- Restore the original volume from a snapshot. The snapshot overwrites the original volume.
- Export a snapshot.
When a snapshot export operation is triggered, CloudPoint creates a new volume from the snapshot and attaches it to the target host using the Fibre Channel (FC) protocol. The target host is assigned read-write privileges on the exported snapshot volume.
- Deport a snapshot.
When a snapshot deport operation is triggered, CloudPoint detaches the exported snapshot volume from the target host and then deletes the volume.

Snapshot export related requirements and limitations

The following requirements and limitations are applicable for snapshot export and deport operations in a Pure Storage array environment:

- A snapshot cannot be exported multiple times.

- An exported snapshot cannot be deleted.
- The CloudPoint user interface (UI) does not support running the snapshot export and deport operations.

Use the following CloudPoint API to perform these operations:

```
(POST) /v2/assets/<disk-ID>/snapshots/<snap-id>/exports/  
(DELETE) /v2/assets/<disk-ID>/snapshots/<snap-id>/exports/<export-ID>
```

Here are some sample cURL commands:

For Export:

```
curl -H "Content-Type: application/json" -H "Authorization: Bearer  
<token>" -X POST -d '{"host-name":"offhost_server", "protocol":"fc",  
"port": "<uuid given to host>"}' -k  
https://localhost/cloudpoint/api/v2/assets/<disk-ID>/snapshots/<snap-id>/exports/
```

For Deport:

```
curl -H "Content-Type: application/json" -H "Authorization: Bearer  
<token>" -X DELETE -k  
https://localhost/cloudpoint/api/v2/assets/<disk-id>/snapshots/<snap-id>/exports/<export-id>
```

You can access the CloudPoint REST APIs using Swagger.

See [“Accessing the Swagger-based API documentation”](#) on page 197.

Huawei OceanStor arrays

This section describes the following:

- The parameters you must supply to configure the Huawei OceanStor Storage Array plug-in
- The Huawei OceanStor models that CloudPoint supports
- The CloudPoint operations you can perform on assets Huawei OceanStor plug-in configuration parameters

Huawei OceanStor array plug-in configuration parameters

When you configure the Huawei OceanStor plug-in, specify the parameters shown in the following table.

Table 21-7 Huawei OceanStor array plug-in configuration parameters

CloudPoint configuration parameter	Description
Array IP Address	The array's IP address
Username	The user name used to access the array
Password	The password used to access the array

Before you configure the plug-in, ensure that the specified user account has permissions to create, delete, and restore snapshots on the array.

Supported Huawei OceanStor arrays

You can use CloudPoint to discover and protect the following Huawei array models.

Table 21-8 Supported Huawei arrays

Category	Supported
Array model	OceanStor 5600 v3
Version	V300R006C10
Patch	SPC100

Table 21-9 List of supported model on Huawei array plugin by CloudPoint

Series	Model
OceanStor V3 series	<ul style="list-style-type: none"> ■ OceanStor 2200 V3 ■ OceanStor 2600 V3 ■ OceanStor 2800 V3 ■ OceanStor 5300 V3 ■ OceanStor 5500 V3 ■ OceanStor 5600 V3 ■ OceanStor 5800 V3 ■ OceanStor 6800 V3 ■ OceanStor 18500 V3 ■ OceanStor 18800 V3

Table 21-9 List of supported model on Huawei array plugin by CloudPoint
(continued)

Series	Model
OceanStor V3 Flash series	<ul style="list-style-type: none">■ OceanStor 2600F V3■ OceanStor 5500F V3■ OceanStor 5600F V3■ OceanStor 5800F V3■ OceanStor 6800F V3■ OceanStor 18500F V3■ OceanStor 18800F V3
OceanStor V5 series	<ul style="list-style-type: none">■ OceanStor 2800 V5■ OceanStor 5300 V5■ OceanStor 5500 V5■ OceanStor 5500 V5 Elite■ OceanStor 5600 V5■ OceanStor 5800 V5■ OceanStor 6800 V5■ OceanStor 18500 V5■ OceanStor 18800 V5
OceanStor V5 Flash series	<ul style="list-style-type: none">■ OceanStor 5300F V5■ OceanStor 5500F V5■ OceanStor 5600F V5■ OceanStor 5800F V5■ OceanStor 6800F V5■ OceanStor 18500F V5■ OceanStor 18800F V5
OceanStor Dorado V3 series	<ul style="list-style-type: none">■ OceanStor Dorado5000 V3■ OceanStor Dorado6000 V3■ OceanStor Dorado18000 V3

Supported CloudPoint operations on Huawei OceanStor array

You can perform the following CloudPoint operations on supported Huawei OceanStor models:

- Discover and list all luns.
- Create a snapshot of a lun.

Note: Snapshot name must be between 1 through 63 characters in length and can be alphanumeric. Snapshot name can contain integers, letters, hyphen ('-'), underscore ('_'), and dot ('.').

- Delete a snapshot.
- Restore the original lun from a snapshot. The snapshot overwrites the original lun.
- Export a snapshot.
In a SAN deployment, when a snapshot export operation is triggered, CloudPoint attaches the snapshot to the target host. The target host is assigned read-write privileges on the exported LUN.
- Deport a snapshot.
In a SAN deployment, when a snapshot deport operation is triggered, CloudPoint removes the exported snapshots mapping from the target host.

Snapshot export related requirements and limitations

The following requirements and limitations are applicable in a Huawei OceanStor array environment:

- A maximum of 255 snapshots or LUNs can be attached to a single Mapping View. The CloudPoint plug-in therefore supports export for up to 255 snapshots or LUNs per WWN/IQN on a target host.

Note: The exported snapshot is attached to the host and is accessible using a world wide name (WWN) or an iSCSI Qualified Name (IQN) that is assigned by the array.

- Snapshot export is supported using the following protocols:
 - Fibre Channel (FC)
 - Internet Small Computer Systems Interface (iSCSI)
- The CloudPoint user interface (UI) does not support running the snapshot export and deport operations.

Use the following CloudPoint API to perform these operations:

```
(POST) /v2/assets/<disk-ID>/snapshots/<snap-id>/exports/  
(DELETE) /v2/assets/<disk-ID>/snapshots/<snap-id>/exports/<export-ID>
```

Here are some sample cURL commands:

For Export:

```
curl -H "Content-Type: application/json" -H "Authorization: Bearer  
<token>" -X POST -d '{"host-name":"offhost_server", "protocol":"fc",  
"port":"<uuid given to host>"}' -k  
https://localhost/cloudpoint/api/v2/assets/<disk-ID>/snapshots/<snap-id>/exports/
```

For Deport:

```
curl -H "Content-Type: application/json" -H "Authorization: Bearer  
<token>" -X DELETE -k  
https://localhost/cloudpoint/api/v2/assets/<disk-id>/snapshots/<snap-id>/exports<export-id>
```

You can access the CloudPoint REST APIs using Swagger.

See ["Accessing the Swagger-based API documentation"](#) on page 197.

Working with CloudPoint using APIs

This chapter includes the following topics:

- [Accessing the Swagger-based API documentation](#)

Accessing the Swagger-based API documentation

You can access the CloudPoint APIs and documentation using the Swagger URL.

To access CloudPoint APIs from a browser

- ◆ Open your browser and enter the following URL in the address bar:

```
https://cloudpoint_hostFQDN/cloudpoint/docs
```

Here, *cloudpoint_hostFQDN* is the name used during the initial CloudPoint configuration. Typically, it is the Fully Qualified Domain Name (FQDN) of the host.

