

Veritas CloudPoint Administrator's Guide

Ubuntu, RHEL

2.2

Veritas CloudPoint Administrator's Guide

Last updated: 2019-10-01

Document version: 2.2 Rev 1

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	11
	About CloudPoint	11
	What kinds of assets can you protect?	12
	Understanding your CloudPoint license	13
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
	CloudPoint host sizing recommendations	25
	Creating an instance or preparing the physical host to install CloudPoint	27
	Installing Docker	27
	Creating and mounting a volume to store CloudPoint data	29
	Verifying that specific ports are open on the instance or physical host	30
Chapter 3	Deploying CloudPoint	31
	About deploying CloudPoint in a non-interactive mode	31
	Installing CloudPoint	32
	Configuring CloudPoint from your browser and signing in	37
	Verifying that CloudPoint installed successfully	42
	Configuring AWS KMS in CloudPoint	43
Chapter 4	Deploying CloudPoint in the AWS cloud	48
	About CloudPoint deployment in the AWS cloud	48
	About CloudPoint integration with AWS KMS	49
	About CloudPoint support for AWS IAM roles	51
	About source account and cross-account configurations	52
	How to configure CloudPoint to use IAM roles	53

	CloudPoint IAM role configuration limitations	54
	About the CloudPoint AWS CloudFormation template	54
	Resources created by the CloudPoint template	54
	CloudPoint EC2 instance configuration details	56
	Instance failures and Auto Scaling Group behavior	57
	Prerequisites for using the CloudPoint template	57
	Launching a CloudPoint CloudFormation stack	58
Chapter 5	Using plug-ins to discover assets	65
	About plug-ins	65
	Determining the types of plug-ins and agents to install	66
Chapter 6	Configuring off-host plug-ins	68
	AWS plug-in configuration notes	68
	Prerequisites for configuring the AWS plug-in	71
	Configuring AWS permissions for CloudPoint	72
	AWS permissions required by CloudPoint	73
	Before you create a cross account configuration	77
	Google Cloud Platform plug-in configuration notes	80
	Google Cloud Platform permissions required by CloudPoint	82
	Configuring a GCP service account for CloudPoint	83
	Preparing the GCP service account for plug-in configuration	84
	Microsoft Azure plug-in configuration notes	85
	Configuring permissions on Microsoft Azure	87
	Dell EMC Unity array plug-in configuration notes	89
	Pure Storage FlashArray plug-in configuration notes	89
	HPE RMC plug-in configuration notes	90
	RMC plug-in configuration parameters	90
	Supported HPE storage systems	91
	Supported CloudPoint operations on HPE storage arrays	91
	NetApp plug-in configuration notes	94
	NetApp plug-in configuration parameters	95
	Supported NetApp arrays	95
	Supported CloudPoint operations on NetApp storage	95
	Configuring an off-host plug-in	99
	About CloudPoint plug-ins and assets discovery	101
	Plug-in discovery interval requirements and limitations	102
	Viewing the assets discovery interval setting	103

Chapter 7	Configuring the on-host agents and plug-ins	104
	About agents	104
	Oracle plug-in configuration notes	106
	Optimizing your Oracle database data and metadata files	107
	MongoDB plug-in configuration notes	107
	Microsoft SQL plug-in configuration notes	108
	About the installation and configuration process	109
	Preparing to install the Linux-based on-host agent	110
	Preparing to install the Windows-based on-host agent	110
	Downloading and installing the on-host agent	111
	Configuring the Linux-based on-host agent	113
	Configuring the Windows-based on-host agent	115
	Configuring the on-host plug-in	117
	Configuring VSS to store shadow copies on the originating drive	118
Chapter 8	Protecting assets with CloudPoint's agentless feature	120
	About the agentless feature	120
	Prerequisites for the agentless configuration	121
	Granting password-less sudo access to host user account	121
	Configuring the agentless feature	122
Section 2	Configuring users	124
Chapter 9	Setting up email and adding users	125
	Configuring the CloudPoint sender email address	125
	About adding users to CloudPoint	127
	Adding AD users to CloudPoint using LDAP	128
	Adding users to CloudPoint manually	129
	Deleting a user from CloudPoint	133
Chapter 10	Assigning roles to users for greater efficiency	134
	About role-based access control	134
	Displaying role information	135
	Creating a role	135
	Editing a role	139
	Deleting a role	140

Section 3	Protecting and managing data	141
Chapter 11	User interface basics	142
	Signing in to CloudPoint	142
	Focusing on an asset type	143
	Navigating to your assets	144
	Using the action icons	146
Chapter 12	Indexing and classifying your assets	147
	About indexing and classifying snapshots	147
	Configuring classification settings using VIC	149
	Indexing and classifying snapshots	150
	Indexing and classification statuses	150
Chapter 13	Protecting your assets with policies	152
	About policies	152
	How a CloudPoint protection policy works	153
	Creating a policy	157
	Assigning a policy to an asset	160
	Listing policies and displaying policy details	163
	Editing a policy	165
	Deleting a policy	166
Chapter 14	Tag-based asset protection	168
	About tag-based asset protection	168
	How to use tag-based asset protection feature	170
	Tag-based asset protection support	173
	Tag-based asset protection considerations and limitations	173
Chapter 15	Replicating snapshots for added protection	175
	About snapshot replication	175
	About cross-account snapshot replication in the AWS cloud	176
	Requirements for replicating snapshots	177
	Cross-account snapshot replication support matrix	178
	Cross-account snapshot replication limitations	178
	Configuring replication rules	179
	Editing a replication rule	181
	Deleting a replication rule	182

Chapter 16	Managing your assets	183
	Creating a snapshot manually	183
	Displaying asset snapshots	187
	Replicating a snapshot manually	189
	About snapshot restore	193
	Restore requirements and limitations for Microsoft SQL Server	196
	Restore requirements and limitations for Oracle	197
	Restore requirements and limitations for MongoDB	197
	About single file restore (granular restore)	198
	Single file restore requirements and limitations	198
	Single file restore support on Linux	199
	Single file restore limitations on Linux	199
	Single file restore support on Windows	200
	Single file restore limitations on Windows	200
	Restoring a snapshot	201
	Additional steps required after a SQL Server snapshot restore	204
	Steps required after a SQL Server host-level restore	204
	Steps required after a SQL Server disk-level snapshot restore to new location	205
	Additional steps required after an Oracle snapshot restore	208
	Additional steps required after a MongoDB snapshot restore	209
	Additional steps required after restoring an AWS RDS database instance	210
	Restoring individual files within a snapshot	211
	Deleting a snapshot	215
Chapter 17	Monitoring activities with notifications and the job log	217
	About CloudPoint notifications	217
	Viewing notifications in the CloudPoint UI	218
	CloudPoint notification methods	222
	CloudPoint notification limitations	222
	Configuring email-based CloudPoint notifications	223
	Configuring AWS SNS-based CloudPoint notifications	224
	Using the Job Log	226
Chapter 18	Protection and disaster recovery	231
	About protection and disaster recovery	231
	Backing up CloudPoint	232
	Restoring CloudPoint	235

Section 4	Maintaining CloudPoint	238
Chapter 19	CloudPoint logging	239
	About CloudPoint logging mechanism	239
	How fluentd-based CloudPoint logging works	240
	About the CloudPoint fluentd configuration file	240
	Modifying the fluentd configuration file	242
	Fluentd-based logging requirements and considerations	242
	Viewing CloudPoint logs	242
Chapter 20	Troubleshooting CloudPoint	246
	Restarting CloudPoint	246
	Docker may fail to start due to a lack of space	247
	CloudPoint installation fails if rootfs is not mounted in a shared mode	248
	Some CloudPoint features do not appear in the user interface	249
	Off-host plug-in deletion does not automatically remove file system and application assets	250
	Disk-level snapshot restore fails if the original disk is detached from the instance	251
	Snapshot restore for encrypted AWS assets may fail	252
	Error while adding users to CloudPoint	253
	CloudPoint fails to revert restored snapshots if indexing, classification, or restore operations fail	254
	SQL snapshot or restore and SFR operations fail if the Windows instance loses connectivity with the CloudPoint host	255
	Troubleshooting CloudPoint logging	256
Chapter 21	Working with your CloudPoint license	257
	Displaying CloudPoint license and protection information	257
	Upgrading your CloudPoint license	258
Chapter 22	Upgrading CloudPoint	262
	About CloudPoint upgrades	262
	Supported upgrade path	262
	Preparing to upgrade CloudPoint	263
	Removing CloudPoint plug-in configuration	263
	Upgrading CloudPoint	264

Chapter 23	Uninstalling CloudPoint	274
	Preparing to uninstall CloudPoint	274
	Removing the CloudPoint on-host agents	274
	Removing CloudPoint from a standalone Docker host environment	275
Section 5	Reference	280
Chapter 24	Storage array support	281
	Dell EMC Unity arrays	281
	Dell EMC Unity array plug-in configuration parameters	281
	Supported Dell EMC Unity arrays	282
	Supported CloudPoint operations on Dell EMC Unity arrays	282
	Pure Storage FlashArray	284
	Pure Storage FlashArray plug-in configuration parameters	284
	Supported Pure Storage FlashArray models	284
	Supported CloudPoint operations on Pure Storage FlashArray models	285
Chapter 25	Working with CloudPoint using APIs	287
	Accessing the Swagger-based API documentation	287

Getting started with CloudPoint

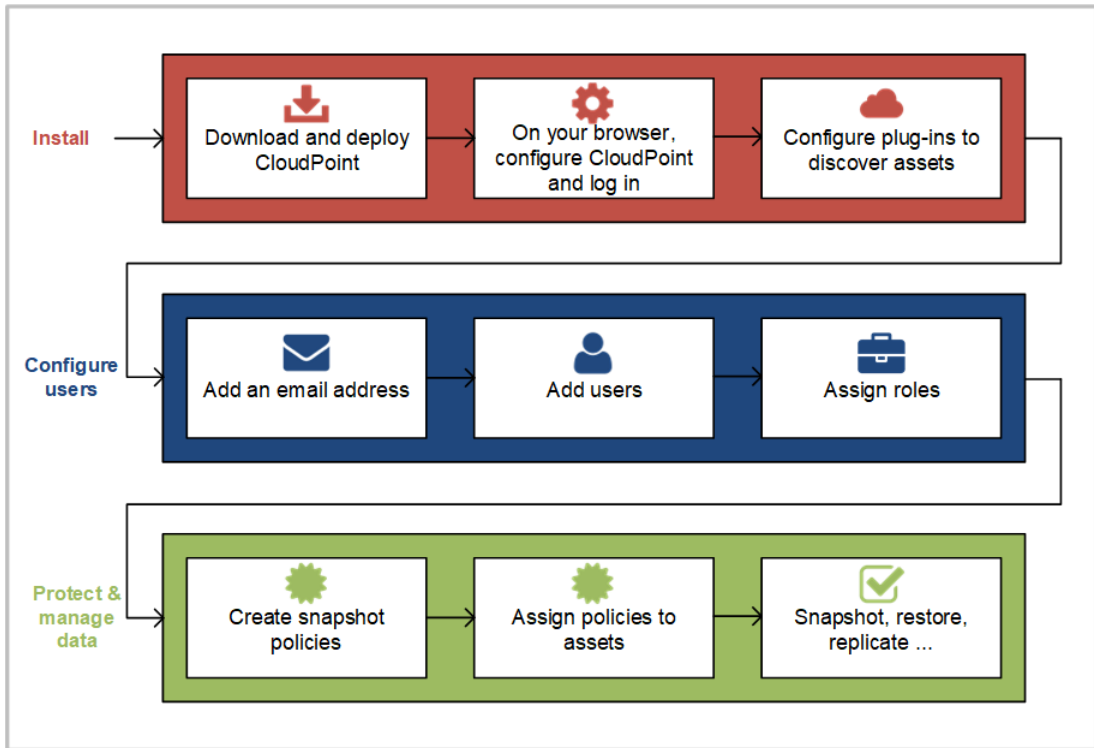
This chapter includes the following topics:

- [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, Oracle 12c R1, Oracle 18c
Disks	<ul style="list-style-type: none">■ Dell EMC Unity arrays■ Pure Storage FlashArray arrays■ HPE storage arrays■ NetApp storage 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■ 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>

Note: CloudPoint does not support AWS Nitro-based instances that use EBS volumes that are exposed as non-volatile memory express (NVMe) devices.

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 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 1024 TB of front-end terabyte data (FETB).

After installing CloudPoint when you perform 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 need more advanced features, you can upgrade the Freemium or Evaluation license to a suitable paid license and 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 license allows you to use all of the CloudPoint features in your preferred on-premise or cloud environments. This license 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 license enables you to discover the storage arrays and take crash-consistent snapshots of the on-premise assets.

This license is created specifically for CloudPoint deployments that are integrated with Veritas NetBackup. You can use this license to leverage your existing NetBackup license entitlement and perform snapshot-based data protection in your on-premise environment.

Install a perpetual or subscription-based license depending on the NetBackup license. 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 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.

Install this license if you want to use the entitled CloudPoint features in your preferred cloud environment. 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.

If CloudPoint is integrated with Veritas NetBackup

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

See [“Displaying CloudPoint license and protection information”](#) on page 257.

See [“Upgrading your CloudPoint license”](#) on page 258.

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. Deploying CloudPoint in the AWS cloud](#)
- [Chapter 5. Using plug-ins to discover assets](#)
- [Chapter 6. Configuring off-host plug-ins](#)
- [Chapter 7. Configuring the on-host agents and plug-ins](#)
- [Chapter 8. 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](#)
- [CloudPoint host sizing recommendations](#)
- [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](#)

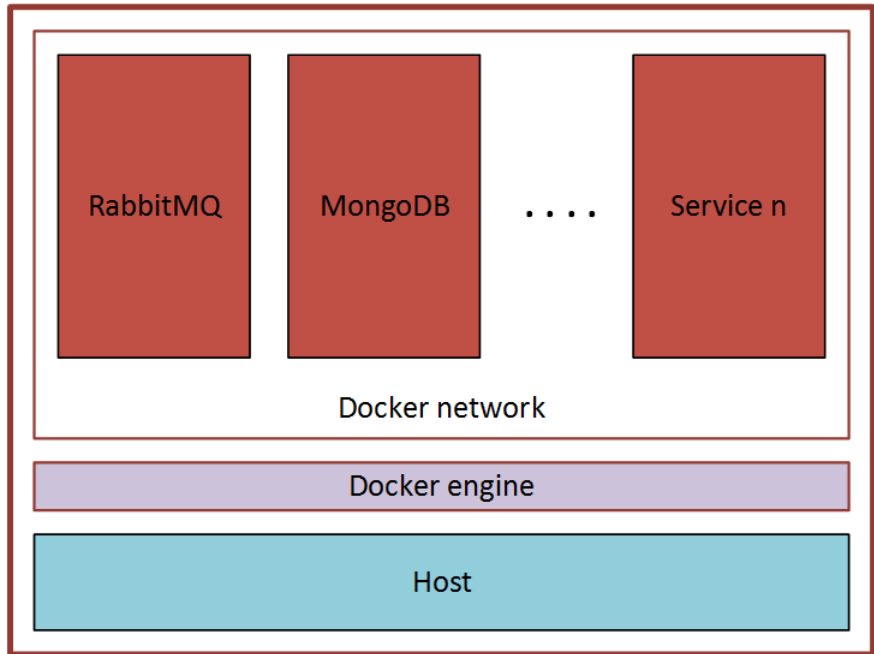
About the deployment approach

CloudPoint is distributed as an Docker image that is built on a Ubuntu 16.04 Server Long Term Support (LTS) base image or a supported 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 <p>The mount state for the root file system (<code>rootfs</code> or <code>shared subtree</code>) on the host must be set to <code>"shared"</code> mode.</p>
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
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
Physical host (x86_64 / AMD64)	<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
<code>/cloudpoint</code>	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 198.</p> Microsoft SQL 2014 and SQL 2016 See “Microsoft SQL plug-in configuration notes” on page 108. MongoDB Enterprise Edition 3.6 See “MongoDB plug-in configuration notes” on page 107. Oracle 12c, Oracle 12c R1, Oracle 18c Single node configurations are supported. See “Oracle plug-in configuration notes” on page 106. <p>Notes:</p> <ul style="list-style-type: none"> Oracle database applications are not supported in a Google Cloud Platform (GCP) cloud environment. This is a limitation imposed by the companies owning these products and services, and is currently outside the scope of CloudPoint. CloudPoint does not support application-consistent snapshots on ext2 file systems. CloudPoint does not support Microsoft SQL Server workloads in a GCP cloud environment.
Operating systems on supported assets	<ul style="list-style-type: none"> Red Hat Enterprise Linux (RHEL) 7.4, 7.5 Windows Server 2012, 2012 R2, 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.
Storage platforms	<ul style="list-style-type: none"> ■ NetApp storage arrays See “NetApp plug-in configuration notes” on page 94. ■ Dell EMC Unity arrays See “Dell EMC Unity array plug-in configuration notes” on page 89. ■ HPE storage arrays See “HPE RMC plug-in configuration notes” on page 90. ■ Pure Storage FlashArray See “Pure Storage FlashArray plug-in configuration notes” on page 89. <p>Note: Starting with CloudPoint 2.2 release, the CloudPoint plug-ins for HPE 3PAR and HDS have been deprecated. Support for these plug-in configurations is no longer available.</p>

Note: CloudPoint does not support AWS Nitro-based instances that use EBS volumes that are exposed as non-volatile memory express (NVMe) devices.

For the latest information on features, software, and hardware support, refer to the compatibility lists:

Table 2-6 CloudPoint compatibility lists

Compatibility list	Document link
Cloud Application Compatibility List (ACL)	https://www.veritas.com/support/en_US/doc/CP_22_APP_CL
Hardware Compatibility List (HCL)	https://sort.veritas.com/DocPortal/pdf/CP_22_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 later
Mozilla Firefox	52.0.0 or later

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.

Proxy server requirements

If the instance on which you are deploying CloudPoint is behind a proxy server, that is, if the CloudPoint instance connects to the internet using a proxy server, you must specify the proxy server details during the CloudPoint installation. The CloudPoint installer stores the proxy server information in a set of environment variables that are specific for the CloudPoint containers.

[Table 2-8](#) describes the environment variables and the proxy server information that you must provide to the CloudPoint installer. Make sure you keep this

information ready; you are required to provide these details during CloudPoint installation.

Table 2-8 Proxy server details required by CloudPoint

Environment variables created by CloudPoint installer	Description
VX_HTTP_PROXY	Contains the HTTP proxy value to be used for all connections. For example, "http://proxy.mycompany.com:8080/".
VX_HTTPS_PROXY	Contains the HTTPS proxy value to be used for all connections. For example, "https://proxy.mycompany.com:8080/".
VX_NO_PROXY	Contains the hosts that are allowed to bypass the proxy server. For example, "localhost,mycompany.com,192.168.0.10:80".

CloudPoint services that need to communicate externally via a proxy server use these predefined environment variables that are set during the CloudPoint installation. For example, the CloudPoint email service, the CloudPoint notifications service, and the CloudPoint plug-in agent containers.

Proxy server limitations

The following restrictions are applicable:

- If CloudPoint is deployed using proxy server settings, email configuration using SendGrid and SMTP are not supported. However, email configuration using Amazon Simple Email Service (SES) is supported.
- CloudPoint deployment using AWS CloudFormation Template (CFT) does not support proxy server configuration.

CloudPoint host sizing recommendations

The CloudPoint host configuration depends primarily on the number of workloads and also the type of workloads that you wish to protect. It is also dependent on the maximum number of simultaneous operations running on the CloudPoint server at its peak performance capacity.

Another factor that affects performance is how you use CloudPoint for protecting your assets. If you use the CloudPoint agentless option to discover and protect your assets, then the performance will differ depending on the type of workload.

With agentless, CloudPoint transfers the plugin data to the application host, performs the discovery and configuration tasks, and then removes the plugin package from the application host. Therefore, database applications such as Oracle and Microsoft SQL Server will require a higher capacity configuration, as compared to other assets.

Veritas recommends the following configurations for the CloudPoint host:

Table 2-9 Typical CloudPoint host configuration

Workload metric	CloudPoint host configuration
Up to 16 concurrent operational tasks	CPU: 2 CPUs Memory: 16 GB For example, in the AWS cloud, the CloudPoint host specifications should be an equivalent of a t2.xlarge instance.
Up to 32 concurrent operational tasks	CPU: 4 - 8 CPUs Memory: 32 GB or more For example, in the AWS cloud, the CloudPoint host specifications should be an equivalent of a t2.2xlarge or a higher type of instance.

General considerations and guidelines:

Consider the following points while choosing a configuration for the CloudPoint host:

- To achieve better performance in a high workload environment, Veritas recommends that you deploy the CloudPoint host in the same location as that of the application hosts.
- If you are using the agentless option, Veritas recommends that you allocate enough space to the `/tmp` directory on the application host. CloudPoint uses this directory for extracting the plugin configuration files.
- Depending on the number of workloads, the amount of plugin data that is transmitted from the CloudPoint host can get really large in size. The network latency also plays a key role in such a case. You might see a difference in the overall performance depending on these factors.
- If you wish to configure multiple workloads using the agentless option, then the performance will be dependent on factors such as the network bandwidth and the location of the CloudPoint host with respect to the application workload instances. You can, if desired, bump up the CloudPoint host's CPU, memory,

and network configuration to achieve a performance improvement in parallel configurations of agentless application hosts.

- In cases where the number of concurrent operations is higher than what the CloudPoint host configuration capacity can handle, CloudPoint automatically puts the operations in a job queue. The queued jobs are picked up only after the running operations are completed.

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 an on-premise instance, 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

Table 2-10 Installing Docker

Platform	Description
Docker on Ubuntu	Supported version: Docker 18.03 and later 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

Table 2-10 Installing Docker (*continued*)

Platform	Description
Docker on RHEL	<p>Supported version: Docker 1.13.x and later</p> <p>Use the following process to install Docker on RHEL. Steps may vary depending on whether CloudPoint is being deployed on-premise or in the cloud.</p> <ul style="list-style-type: none"> ■ (If CloudPoint is being deployed in AWS cloud) Ensure that you enable the extra repos: <pre># sudo yum-config-manager --enable rhui-REGION-rhel-server-extras</pre> ■ (If CloudPoint is being deployed on-premise) Enable your subscriptions: <pre># sudo subscription-manager register --auto-attach --username=<username> --password=<password> # subscription-manager repos --enable=rhel-7-server-extras-rpms # subscription-manager repos --enable=rhel-7-server-optional-rpms</pre> ■ Install Docker using the following command: <pre># sudo yum -y install docker</pre> ■ (If CloudPoint is being deployed in Azure cloud) Enable shared mounts. <ul style="list-style-type: none"> ■ Edit the <code>docker.service</code> system unit file and modify the parameter MountFlags=slave to MountFlags=shared. ■ 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 should appear as <code>MountFlags=shared</code>. ■ Reload the system manager configuration using the following command: <pre># sudo systemctl daemon-reload</pre> ■ Enable and then restart the docker service using the following commands: <pre># sudo systemctl enable docker # sudo systemctl restart docker</pre> ■ If SELinux is enabled, change the mode to permissive mode. Edit the <code>/etc/selinux/config</code> configuration file and modify the <code>SELINUX</code> parameter value to <code>SELINUX=permissive</code>. ■ Reboot the system for the changes to take effect. ■ Verify that the SELinux mode change is in effect using the following command: <pre># sudo sestatus</pre> The <code>Current Mode</code> parameter value in the command output should appear as <code>permissive</code>. <p>Refer to the following documentation for detailed instructions on installing Docker on RHEL:</p> <p>https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux_atomic_host/7/html-single/getting_started_with_containers/index#getting_docker_in_rhel_7</p>

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`.

Table 2-11 Volume creation steps for each supported cloud vendor

Vendor	Procedure
Amazon Web Services (AWS)	<ol style="list-style-type: none"> On the EC2 dashboard, click Volumes > Create Volumes. Follow the instructions on the screen and specify the following: <ul style="list-style-type: none"> Volume type: General Purpose SSD Size: 50 GB 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"> 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 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 2-12 Ports used by CloudPoint

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

- [About deploying CloudPoint in a non-interactive mode](#)
- [Installing CloudPoint](#)
- [Configuring CloudPoint from your browser and signing in](#)
- [Verifying that CloudPoint installed successfully](#)
- [Configuring AWS KMS in CloudPoint](#)

About deploying CloudPoint in a non-interactive mode

CloudPoint is distributed as a Docker image. You can use standard Docker commands to install, uninstall, or upgrade CloudPoint. During a typical installation or upgrade, the installer may display several prompts requesting for a confirmation. You have to respond to these prompts to allow the installer to proceed with the operation.

However, you can suppress these prompts and run the CloudPoint deployment in an unattended, non-interactive mode. This is particularly useful if you plan to automate the installation process using deployment scripts or a third-party application.

To run the installer in a non-interactive mode, use the following additional parameters with the installation command:

Table 3-1 Command parameters for non-interactive mode

Parameter	Description
-y	Represents a Yes. The installer considers this as an approval and proceeds with the operation. No prompts are displayed.
-n	Represents a No. The installer considers this as a disapproval and may even suspend or exit from the requested operation.

For example, the following command is used to install CloudPoint:

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


To run the same command in an unattended mode, use the following syntax:

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

The `-y` option passes an approval for all the subsequent installation prompts and allows the installer to proceed in a non-interactive mode.

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 27.
- Install Docker.
See [“Installing Docker”](#) on page 27.
- Create and mount a volume to store CloudPoint data.
See [“Creating and mounting a volume to store CloudPoint data”](#) on page 29.
- 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 30.
- If you want to install or upgrade CloudPoint in an unattended mode, ensure that you use the appropriate command syntax.
See [“About deploying CloudPoint in a non-interactive mode”](#) on page 31.
- If you want to install CloudPoint on a host where the root file system (`rootfs` or `/` shared subtree) is not mounted in a `shared` mode, ensure that you use the appropriate command syntax.
Although CloudPoint supports installation in such an environment, certain restrictions apply.
See [“CloudPoint installation fails if rootfs is not mounted in a shared mode”](#) on page 248.

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
```

If the CloudPoint host is behind a proxy server, use the following command instead:

```
# sudo docker run -it --rm
-v /<full_path_to_volume_name>:/<full_path_to_volume_name>
-e VX_HTTP_PROXY=<http_proxy_value>
-e VX_HTTPS_PROXY=<https_proxy_value>
-e VX_NO_PROXY=<no_proxy_value>
-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.
<http_proxy_value> (required only if the instance uses a proxy server)	Represents the value to be used as the HTTP proxy for all connections. For example, "http://proxy.mycompany.com:8080/".
<https_proxy_value> (required only if the instance uses a proxy server)	Represents the value to be used as the HTTPS proxy for all connections. For example, "https://proxy.mycompany.com:8080/".
<no_proxy_value> (required only if the instance uses a proxy server)	Represents the addresses that are allowed to bypass the proxy server. You can specify host names, IP addresses, and domain names in this parameter. Use commas to separate multiple entries. For example, "localhost,mycompany.com,192.168.0.10:80".

Note:

If CloudPoint is being deployed in the cloud, ensure that you set the following values in this parameter:

- For an AWS instance, add the following:
169.254.169.254
- For a GCP virtual machine, add the following:
169.254.169.254,metadata,metadata.google.internal
- For an Azure virtual machine, add the following:
169.254.169.254

CloudPoint uses these addresses to gather instance metadata from the instance metadata service.

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
```

If using a proxy server, then using the examples provided in the table earlier, the command syntax is as follows:

```
# sudo docker run -it --rm -v /cloudpoint:/cloudpoint -e
VX_HTTP_PROXY="http://proxy.mycompany.com:8080/" -e
VX_HTTPS_PROXY="https://proxy.mycompany.com:8080/" -e
VX_NO_PROXY="localhost,mycompany.com,192.168.0.10:80" -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 247.

- 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 37.


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 246.

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 32.

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.

- 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 specified email address should not include an underscore character.
CloudPoint currently does not support adding users whose email addresses contain the underscore character.

The admin password should meet the following requirements:

- At least six characters

- No spaces
 - No & (ampersand) character
- 3 Under Host information, in the **Host names or IP** field, enter any additional host names or IP address that you use to connect to this CloudPoint host and then click the **+** icon to add that entry to the list.

Repeat this for each additional host name that you wish to add.

The specified names or IP address are added to the list of host names to use for configuring CloudPoint. The names in the list are 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.

The names you specify here must point to the same CloudPoint host. The fully qualified domain name (FQDN) of the host is added by default.

Note: If you are integrating CloudPoint with Veritas NetBackup, use host names only.

- 4 Under Select Your License, select the CloudPoint trial license that you wish to activate in your CloudPoint deployment.

Pick from one of the following options depending on your requirement:

License type	Description
Freemium	A Freemium license is a permanent license that does not expire and allows you to try out a subset of the CloudPoint features. This license lets you protect up to 10 TB of front-end terabyte (FETB) data.
Evaluation	An Evaluation license is a 60-day time-bound license that allows you to try out all of the CloudPoint features. This license lets you protect up to 1000 TB of FETB data.

The selected trial license is installed on the CloudPoint instance during the initial configuration and allows you to use the features that are entitled as part of that license.


You can upgrade to a paid license any time after the initial configuration is completed.

See [“Understanding your CloudPoint license”](#) on page 13.

- 5** 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.
- 6** Read the End User License Agreement and then select the I agree to the terms and conditions option.
- 7** 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.
- 8** After the installation completes, click **Refresh browser**. If you see the login screen, it confirms that CloudPoint is installed and configured successfully.


- 9 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.

- 10 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 42.

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.

Configuring AWS KMS in CloudPoint

This is applicable only if CloudPoint instance is deployed in the AWS cloud.

Perform the following steps if you wish to configure CloudPoint to use AWS Key Management Service (KMS) for encrypting and decrypting your CloudPoint configuration information. CloudPoint provides REST APIs that you can use to configure AWS KMS in your CloudPoint environment.

These steps are required only if you have manually deployed CloudPoint using the Docker image on an AWS EC2 instance in the AWS cloud. These steps are not required if you have deployed CloudPoint using the CloudFormation Template (CFT).

Note: Veritas recommends that you use the CloudPoint CloudFormation Template to deploy CloudPoint in the AWS cloud. KMS is automatically configured as part of the template-based deployment workflow.

See [“About the CloudPoint AWS CloudFormation template”](#) on page 54.

CloudPoint AWS KMS configuration prerequisites

- Ensure that you have successfully installed and configured CloudPoint on the EC2 instance in the AWS cloud.
See [“Installing CloudPoint”](#) on page 32.
- Read about how CP integrates with AWS KMS and understand the limitations.
See [“About CloudPoint integration with AWS KMS”](#) on page 49.
- Ensure that you have created an AWS IAM role and attached it to the CloudPoint EC2 instance.

The IAM role must have the following permissions at a minimum:

```
kms:DescribeKey
kms:GenerateDataKey
kms:Decrypt
```

Refer to the following AWS KMS documentation for detailed instructions:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html#working-with-iam-roles>

- Ensure that you have created a customer managed Customer Master Key (CMK). The Key ID of the CMK is required for configuring AWS KMS in CloudPoint.

Refer to the following AWS KMS documentation for detailed instructions:

<https://docs.aws.amazon.com/kms/latest/developerguide/create-keys.html>

To configure AWS KMS in CloudPoint

- 1 Generate an authentication token for the CloudPoint administrator user account by using the following CloudPoint Identity Management API:

```
POST /v2/idm/login
```

On any system that can connect to the CloudPoint instance, type the following cURL command:

```
# curl -k https://<cloudpointhostFQDN>/cloudpoint/api/v2/idm/login  
-X 'POST' -H "Content-Type: application/json"  
-d '{"email":"<username>", "password":"<password>"}'
```

Replace the following parameters as per your environment:

Parameter	Description
<i><cloudpointhostFQDN></i>	Represents the Fully Qualified Domain Name (FQDN) that was specified while performing the initial CloudPoint configuration on the host.
<i><username></i>	Represents the user name that was specified as the CloudPoint administrator user during initial configuration.
<i><password></i>	Represents the password of the CloudPoint administrator user account.

- 2 Observe the API output on the command prompt. You will see an output that resembles the following:

```
{  
  "accessToken": "eyJhbGciOiJIUzI1NiJ9.eyJpc3MiOiJWZXJpdGFzIiwidXN",  
  "applicationId": "",  
  "applicationPath": "",  
  "errorCode": ""  
}
```

The alphanumeric entry that appears as the `accessToken` represents the token that is used to authorize all CloudPoint API requests on the host. Copy the token, it is required in the subsequent steps.

Note: The alphanumeric authentication token displayed here is for representation purpose only. Use the actual token that is generated when you run this command in your environment.

- 3 Create a new AWS KMS configuration using the CloudPoint Key Management Service API `POST /v3/kms`.

Enter the following cURL command on the same command prompt:

```
# curl -k "https://<cloudpointhostFQDN>/cloudpoint/api/v3/kms"  
-X 'POST' -H "Content-Type: application/json"  
-H "Authorization: Bearer <authtoken>"  
-d '{"platform":"aws", "masterKeyId":"<cmk_keyid>",  
"credentials":{"type":"iamrole", "regionname":"<cmk_regionname>"}}'
```

Replace the following parameters as per your environment:

Parameter	Description
<cloudpointhostFQDN>	Represents the Fully Qualified Domain Name (FQDN) that was specified while performing the initial CloudPoint configuration on the host.
<authtoken>	Represents the alpha numeric authentication token that you generated in the earlier step.
<cmk_keyid>	Represents the AWS customer managed Customer Master Keys (CMK) key ID that you created for CloudPoint.
<cmk_regionname>	Represents the CMK region where the CloudPoint instance is deployed.

- 4 Observe the API output on the command prompt and wait for the task to complete.

- 5** You can quickly verify if the AWS KMS is configured successfully by using the CloudPoint API `GET /v3/kms`.

Run the following cURL command:

```
# curl -k -X GET "https://<cloudpointhostFQDN>/cloudpoint/api/v3/kms"  
-H "accept: application/json"  
-H "Authorization: Bearer <authtoken>"
```

Replace the following parameters as per your environment:

Parameter	Description
<code><cloudpointhostFQDN></code>	Represents the Fully Qualified Domain Name (FQDN) that was specified while performing the initial CloudPoint configuration on the host.
<code><authtoken></code>	Represents the alpha numeric authentication token that you generated in step 2 earlier.

An HTTP 200 status indicates that the configuration was performed successfully.

- 6** You must now configure the CloudPoint plug-ins or the agentless feature.
- See [“Configuring an off-host plug-in”](#) on page 99.
- See [“About the agentless feature”](#) on page 120.

Deploying CloudPoint in the AWS cloud

This chapter includes the following topics:

- [About CloudPoint deployment in the AWS cloud](#)
- [About CloudPoint integration with AWS KMS](#)
- [About CloudPoint support for AWS IAM roles](#)
- [About the CloudPoint AWS CloudFormation template](#)
- [Prerequisites for using the CloudPoint template](#)
- [Launching a CloudPoint CloudFormation stack](#)

About CloudPoint deployment in the AWS cloud

A common deployment approach for CloudPoint is to set up a CloudPoint instance in the cloud and then configure it to protect and manage all the assets in the AWS cloud. You can set up the same CloudPoint instance to manage assets spread across multiple AWS accounts and regions.

Here's how you can deploy CloudPoint in the AWS cloud:

Use the CloudPoint Docker image

CloudPoint is distributed as a Docker image that you load and run and then use Docker commands to install the CloudPoint services as individual containers in the Docker network.

With this method, you first configure an EC2 instance that meets the CloudPoint requirements, download and load the CloudPoint Docker image on the instance, and then install and configure CloudPoint on that instance. This method provides

a manual deployment workflow and is the same as when you are deploying CloudPoint on an on-premise physical or a virtual host.

See “[Installing CloudPoint](#)” on page 32.

Use the CloudPoint CloudFormation Template

CloudPoint is also available on the AWS Marketplace online store in the form of an AWS CloudFormation Template (CFT).

With this method, you simply sign in to the AWS Console and use the CloudPoint template to launch a CloudPoint CloudFormation stack. This method provides a much more faster and automated deployment workflow that is fully integrated with the AWS tools and services. Veritas recommends that you use the template to deploy CloudPoint in your AWS cloud environment.

See “[About the CloudPoint AWS CloudFormation template](#)” on page 54.

About CloudPoint integration with AWS KMS

CloudPoint uses the AWS account credentials (Secret Key and Access Key pair) to connect to the AWS cloud, discover all the assets, and perform operations on those assets. The AWS account details are stored in the CloudPoint configuration in an encrypted format. CloudPoint uses the 256-bit Advanced Encryption Standard (AES) specification to encrypt and decrypt all the configuration information. The encryption keys, also referred to as the coordinator keys, are stored internally in the CloudPoint MongoDB database. This encryption mechanism is used as a default for all CloudPoint deployments, whether on-premise or in the cloud.

Starting with CloudPoint 2.2 release, CloudPoint also provides support for AWS Key Management Service (KMS) for deployments in the cloud. AWS KMS is a managed service that allows you to create and manage encryption keys that are used to encrypt your data. With AWS KMS integration, you can now utilize the AWS KMS service to encrypt and decrypt the CloudPoint configuration information rather than storing the encryption keys internally in the CloudPoint database.

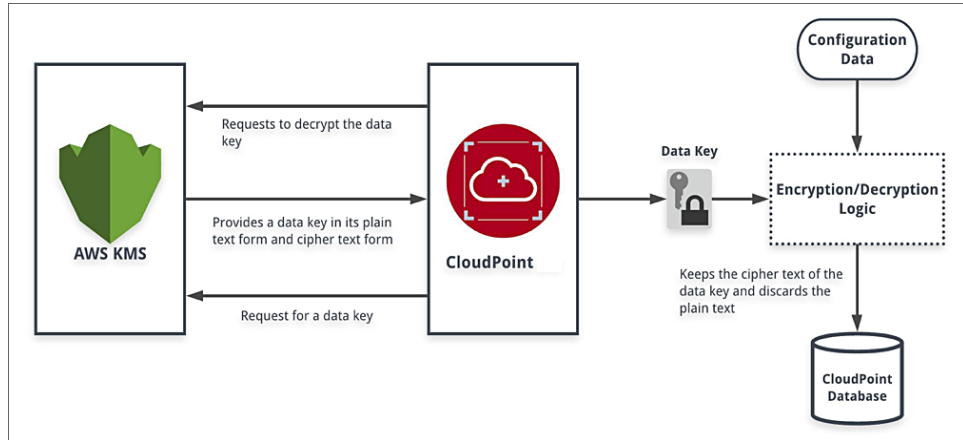
Refer to the AWS documentation for more details on KMS:

<https://docs.aws.amazon.com/kms/latest/developerguide/overview.html>

How CloudPoint uses AWS KMS

When you configure CloudPoint to use KMS, CloudPoint sends a request to KMS to generate a data key that is used to encrypt and decrypt all configuration information such as CloudPoint plug-in configuration, agentless configuration, and application configuration data. KMS provides an encryption key in two forms--a cryptic cipher form and a plain text form. CloudPoint uses the plain text form of the key to encrypt the configuration and then stores the cipher form of the key in the

CloudPoint database. The plain text form of the key is completely discarded after the encryption is completed.



To decrypt the configuration information, CloudPoint sends a request to KMS to decrypt the cipher form of the key that resides in the CloudPoint database. KMS decrypts the cipher key and generates the actual key that CloudPoint then uses to decrypt the configuration information. The cipher form of the key that resides in the database can be decrypted only by KMS.

How to configure CloudPoint to use AWS KMS

How you configure CloudPoint to use AWS KMS depends on how you deploy CloudPoint in the cloud:

- If you deploy CloudPoint manually using the Docker image on an EC2 instance, then you can configure AWS KMS using CloudPoint APIs.
 This is a manual procedure that must be performed after successfully installing and configuring CloudPoint on the EC2 instance.
 See [“Configuring AWS KMS in CloudPoint”](#) on page 43.
- If you deploy CloudPoint using the CloudFormation Template (CFT), then there are no additional steps required.
 You specify the AWS KMS Customer Master Keys (CMK) as one of the parameters in the template form. AWS KMS is automatically configured on the CloudPoint instance when you launch the CloudFormation stack.
 See [“About the CloudPoint AWS CloudFormation template”](#) on page 54.
- If you are upgrading from an older release of CloudPoint that did not support KMS, then the default encryption mechanism that was used earlier continues to work as is even after a successful upgrade. However, if you wish to use AWS KMS post upgrade, you can configure AWS KMS using CloudPoint APIs.

See “[Configuring AWS KMS in CloudPoint](#)” on page 43.

Note: Once you configure KMS, you cannot disable it or go back to using the default encryption method, in the same CloudPoint deployment.

CloudPoint and AWS KMS configuration limitations

The following conditions are applicable to CloudPoint integration with AWS KMS:

- CloudPoint supports AWS KMS customer managed Customer Master Keys (CMK) only. AWS owned CMK and AWS managed CMK keys are not supported.
- AWS KMS integration is available for CloudPoint deployments in AWS cloud only. This is not supported for on-premise deployments as well as deployments in other cloud environments.
- You must not delete the Customer Master Keys (CMK) that are used for KMS configuration in CloudPoint. If the keys are lost, your CloudPoint configuration may be irrecoverable and you may have to redeploy your entire CloudPoint environment.

About CloudPoint support for AWS IAM roles

After you deploy CloudPoint, you use the AWS Identity and Access Management (IAM) user credentials (Secret Key and Access Key pair) and configure the CloudPoint plug-in for AWS to discover the AWS assets that you wish to protect using CloudPoint. The key pair is used to get access to the AWS resources and then perform operations on the discovered assets. The AWS account credentials are permanently stored in the CloudPoint configuration database in an encrypted format. CloudPoint uses the key pair authentication mechanism for all deployments, be it on-premise or in the cloud.

Starting with release 2.2, CloudPoint provides support for using AWS IAM roles for CloudPoint deployments in the AWS cloud. IAM is an AWS service that allows you to manage access to AWS services and resources in a secure manner. You can create an IAM role, assign it with the permissions that CloudPoint requires, and then attach the role to the CloudPoint instance. CloudPoint then uses the security credentials provided by the IAM role to discover and perform snapshot operations on the assets in the cloud. You can now use IAM user or IAM roles to configure CloudPoint to protect assets that belong to multiple AWS accounts in the cloud.

Refer to the AWS documentation for more information on IAM roles:

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html

See “[About source account and cross-account configurations](#)” on page 52.

See [“How to configure CloudPoint to use IAM roles”](#) on page 53.

See [“CloudPoint IAM role configuration limitations”](#) on page 54.

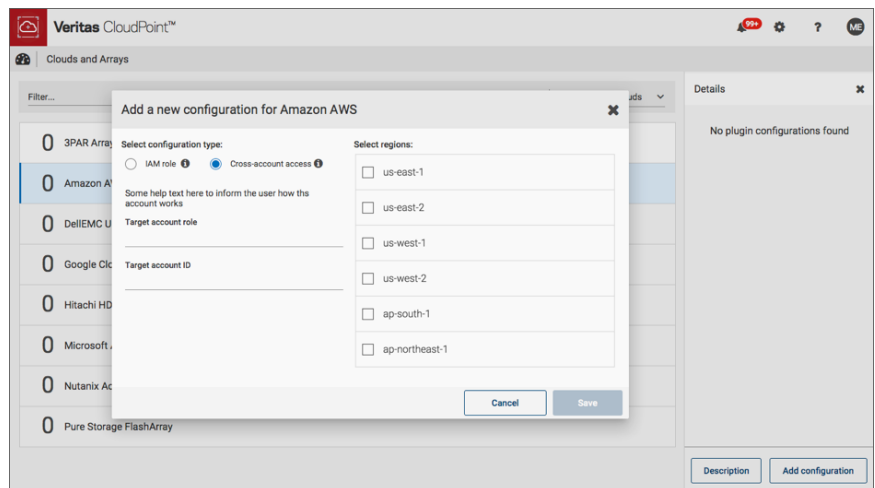
About source account and cross-account configurations

While configuring the CloudPoint plug-in for AWS, you are presented with the following configuration options:

- Source Account

Source Account configuration represents the primary AWS account whose assets you wish to protect using CloudPoint. This is the AWS account in which the CloudPoint instance resides. This is a minimal required configuration if you want to use CloudPoint to protect AWS assets in the cloud.

For Source Account configuration, the CloudPoint retrieves the security credentials from the IAM role that is attached to the CloudPoint instance.



- Cross Account

Cross Account configuration represents the additional AWS accounts whose assets you wish to protect using the same CloudPoint instance that resides in the source account.

For Cross Account configuration, CloudPoint uses the credentials associated with IAM role attached to CloudPoint instance to assume the IAM role defined in another account (target account). The target AWS account responds with a set of temporary credentials and these credentials are cached in memory and then used to access, discover, and perform operations on the assets in that AWS account.

For Cross Account configuration to work, a trust relationship is required between the IAM role of the Source Account and the IAM role of the other AWS account. The IAM role policy of the Source Account must allow to assume the IAM role attached to the other AWS account. Similarly, the IAM role of the other AWS account must trust the IAM role of the Source Account.

See [“AWS plug-in configuration notes”](#) on page 68.

See [“Prerequisites for configuring the AWS plug-in”](#) on page 71.

Note: To create a Cross Account configuration, you must first configure a Source Account.

How to configure CloudPoint to use IAM roles

How you configure CloudPoint to use AWS IAM roles depends on how you deploy CloudPoint in the AWS cloud:

- If you deploy CloudPoint manually using the Docker image on an EC2 instance, then you must create an IAM role, assign the role with the permissions required by CloudPoint, and then manually attach that role to the CloudPoint instance. After attaching the role, you then configure the AWS plug-in.
See [“AWS plug-in configuration notes”](#) on page 68.
- If you deploy CloudPoint using the CloudFormation Template (CFT), then you specify the IAM role details as an input parameter in the CFT form. The CFT attaches the IAM role to the CloudPoint instance and configures the Source Account automatically as part of the template-based deployment workflow. You can either specify an existing IAM role or have the CFT template create a new IAM role for you. If you specify an existing IAM role, you must ensure that the IAM role has all the permissions that are required by CloudPoint. If you do not specify an IAM role in the CFT form, the CFT creates a new IAM role, assigns it with all the permissions that are required by CloudPoint, and then attaches that role to the CloudPoint instance.
See [“About the CloudPoint AWS CloudFormation template”](#) on page 54.
- If you are upgrading CloudPoint from an older release that did not support IAM roles, the existing key pair-based configuration will continue to work as is even after the upgrade. However, you cannot update the secret key access key pair after the upgrade. If you wish to make any changes to the plug-in configuration, you first create an IAM role that maps to the same AWS user account that was used to create the secret key and access key pair, and then use that IAM role to update the configuration.

CloudPoint IAM role configuration limitations

The following limitations are applicable to CloudPoint support for IAM roles:

- CloudPoint supports using AWS IAM roles for CloudPoint deployments in the AWS cloud only. When you deploy CloudPoint using the CloudFormation Template or using the CloudPoint Docker image on an EC2 instance, CloudPoint requires that you use AWS IAM role for authenticating CloudPoint operations on the assets in the AWS cloud.
- CloudPoint supports using the Secret Key and Access Key pair configuration for CloudPoint deployments on-premise and in other cloud environments. The key-pair method is no longer supported for new CloudPoint deployments in the AWS cloud.

About the CloudPoint AWS CloudFormation template

Veritas provides a template that you can use to provision CloudPoint as a CloudFormation stack in your AWS cloud environment. The template contains a description of all the AWS resources, their properties, and all the dependencies that CloudPoint needs. Launch a AWS CloudFormation stack using this template to get a CloudPoint configuration up and running.

See [“Resources created by the CloudPoint template”](#) on page 54.

See [“CloudPoint EC2 instance configuration details”](#) on page 56.

See [“Instance failures and Auto Scaling Group behavior”](#) on page 57.

See [“Prerequisites for using the CloudPoint template”](#) on page 57.

See [“Launching a CloudPoint CloudFormation stack”](#) on page 58.

Resources created by the CloudPoint template

The following resources are created when you launch a CloudPoint stack using the CloudPoint CloudFormation template:

Table 4-1 CloudPoint CloudFormation template resources

Resource	Description
EBS Volume (AWS::EC2::Volume)	The volume size and availability zone are specified during the stack creation process.

Table 4-1 CloudPoint CloudFormation template resources (*continued*)

Resource	Description
EC2 Instance (AWS::EC2::Instance)	The instance type along with the required network and security configuration settings are specified during the stack creation process.
VolumeAttachment (AWS::EC2::VolumeAttachment)	An EBS volume gets attached to the CloudPoint EC2 instance that is created.
Instance Profile (AWS::IAM::InstanceProfile)	A new profile is created for the CloudPoint EC2 instance. This profile then assigns the specified IAM role to the EC2 instance.
IAM Role (AWS::IAM::Role)	<p>A new IAM role is created and attached to the CloudPoint instance during the stack creation process. The role is assigned all the AWS permissions that CloudPoint requires.</p> <p>This new role is created only if you have not specified any existing IAM role name in the CFT form (the <i>CloudPoint System Configuration</i> > <i>IAM Role</i> field is empty).</p>
Security Group (AWS::EC2::SecurityGroup)	<p>An AWS security group is created internally for the CloudPoint deployment.</p> <p>The security group contains rules that allow inbound and outbound traffic for the following:</p> <ul style="list-style-type: none"> ■ SSH on port 22 ■ RabbitMQ on port 5671 ■ HTTPS on port 443 <p>Refer to the AWS documentation for more information on security groups:</p> <p>https://docs.aws.amazon.com/vpc/latest/userguide/VPC_SecurityGroups.html</p>
Launch Configuration (AWS::AutoScaling::LaunchConfiguration)	A new launch configuration is created and is then used by the Auto Scaling Group (ASG) to scale the instance if the original CloudPoint EC2 instance status is marked as unhealthy.

Table 4-1 CloudPoint CloudFormation template resources (continued)

Resource	Description
Auto Scaling Group (AWS::AutoScaling::AutoScalingGroup)	A new Auto Scaling Group (ASG) is created and the CloudPoint instance is attached to it. If the original CloudPoint EC2 instance becomes unhealthy, this ASG automatically creates a new CloudPoint instance and attaches the existing CloudPoint metadata volume to the new instance.

CloudPoint EC2 instance configuration details

When you deploy a CloudPoint stack using the CloudFormation template, the following configuration is created on the EC2 instance where CloudPoint is deployed:

- A disk is attached to the instance and a file system of type `ext4` is created on the disk.
- The file system is mounted as a folder mount at `/cloudpoint`.
- CloudPoint is installed and the specified user account is configured as the CloudPoint administrator.
- The CloudPoint AWS plug-in is configured with the Source Account configuration. The IAM role that is attached to the CloudPoint instance is used for the plug-in configuration.
See “[About CloudPoint support for AWS IAM roles](#)” on page 51.
- A CloudPoint snapshot policy by the name **BackupCP** is created.
The policy is assigned to the file system that is mounted at `/cloudpoint`.
This is a built-in protection policy that is automatically assigned to the CloudPoint instance in the cloud. This policy takes periodic snapshots of the CloudPoint metadata. The policy schedule and retention options are fully configurable. The policy protects the CloudPoint instance assets in the background, even though the protected assets do not appear in the CloudPoint UI. If the original CloudPoint instance fails to respond, the latest snapshot created by this policy is used to create a new CloudPoint instance in the cloud.

Note: You must add an appropriate CloudPoint license to the CloudPoint configuration for the policy to take effect. The policy does not trigger file system snapshots until a valid license is installed.

Instance failures and Auto Scaling Group behavior

The Amazon EC2 Auto Scaling Group (ASG) monitors the CloudPoint EC2 instance periodically. The ASG determines the status of the instance using the default status checks or via custom health checks. After the instance passes the status checks, the instance is marked as healthy.

If the state of the instance changes due to an external event, for example--if a disaster causes a loss of the instance, or if you manually stop the instance, the ASG immediately marks the instance as unhealthy and schedules it for a replacement.

The ASG creates a new instance from the same Amazon Machine Instance (AMI) and uses the same configuration as that of the original instance that was configured earlier. The ASG creates a new EBS volume using the snapshot, attaches that volume to the new EC2 instance, and then brings CloudPoint up on that instance.

You can suspend the health check process if you do not want ASG to replace the instance, for example in cases where you want to stop the instance for maintenance purposes.

For more information on how the Amazon EC2 Auto Scaling works, refer to the following Amazon AWS documentation:

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/healthcheck.html>

Prerequisites for using the CloudPoint template

Ensure that you configure the following before you launch a CloudPoint CloudFormation stack:

- Set up AWS SNS notifications by creating an SNS topic for the CloudPoint stack. This allows you to receive notification emails each time the Auto Scaling Group (ASG) is updated.
The SNS topic must be configured in the same AWS region where the CloudPoint instance is being deployed.
<https://docs.aws.amazon.com/gettingstarted/latest/deploy/creating-an-sns-topic.html>
- Create a key pair in the region where you want to launch the CloudPoint stack.
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-key-pairs.html>
- If desired, set up a AWS customer master key (CMK) if you want to use AWS KMS with CloudPoint. This is not a mandatory requirement.
<https://docs.aws.amazon.com/kms/latest/developerguide/create-keys.html>
- Create an AWS IAM role and assign permissions that are required by CloudPoint. See [“Configuring AWS permissions for CloudPoint”](#) on page 72.

CloudPoint requires that you use AWS IAM for authenticating CloudPoint operations on the assets in the AWS cloud.

Refer to the AWS documentation for more information on IAM roles.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

Launching a CloudPoint CloudFormation stack

Perform the following steps to deploy CloudPoint in a new AWS CloudFormation stack.

To deploy CloudPoint as a CloudFormation stack

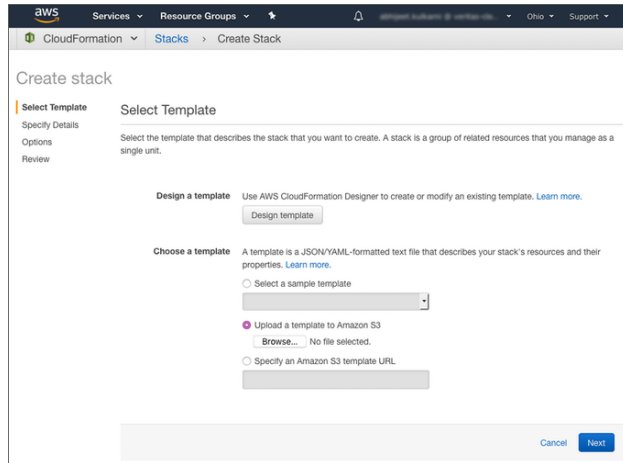
1. From the AWS Marketplace online store, download the CloudPoint CloudFormation template to a temporary location.

Alternatively, you can also make a note of the template download URL.
2. Log on to the AWS Management Console and from the top right corner select the region in which you want to run the CloudPoint instance.
3. From the Services menu, under Management & Governance, select **CloudFormation**.
4. To begin creating a new stack, click **Create Stack**.


Note: AWS provides different options to create a stack depending on whether you have an existing stack running. Refer to the following for the exact steps:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-console-create-stack.html>

5. On the Select Template page, in the Choose a template section, specify the CloudPoint template using any of the following methods:
 - Click **Upload a template to Amazon S3** and then click **Browse** and select the CloudPoint template file that you downloaded earlier.



- Alternatively, click **Specify an Amazon S3 template URL** and then enter the CloudPoint template download URL.
6. After specifying the template file, click **Next**.
 7. On the Specify Details page, in the **Stack name** field, type a name for the new stack.



Use a descriptive name that helps you identify this stack from a list of stacks later.

8. On the Specify Details page, in the Parameters section, specify the required parameter values.

These parameters allow you to customize the stack at creation time.

■ **CloudPoint System Configuration**

Parameter	Description
EC2 Instance Type	<p>From the drop-down list, select the instance type that you want to use for the CloudPoint instance.</p> <p>Specify t2.large or a higher configuration.</p>
Volume Size	<p>Specify a size for the EBS volume that is attached to the new instance. This volume is used for storing CloudPoint metadata.</p> <p>Enter a value of 60 GB or more.</p>
EBS Volume ID <i>(optional)</i>	<p>This is applicable only in case of an upgrade scenario.</p> <p>Specify the ID of the EBS volume that contains the CloudPoint metadata of an existing CloudPoint deployment.</p> <p>This parameter is not required if you are creating a fresh CloudPoint deployment on a new instance.</p>
Volume Snapshot ID <i>(optional)</i>	<p>This is applicable only in case of an upgrade scenario.</p> <p>Specify the snapshot ID of the disk that contains the CloudPoint metadata of an existing CloudPoint deployment.</p> <p>A new EBS volume is created from this snapshot and is attached to the new instance.</p> <p>This parameter is not required if you are creating a fresh CloudPoint deployment on a new instance.</p>
Availability Zone	<p>From the drop-down list, select the availability zone where you want to deploy the CloudPoint instance.</p>

Parameter	Description
IAM Role	<p>Specify the IAM role that you want to attach to the CloudPoint instance.</p> <p>Ensure that the IAM role is assigned with the permissions that CloudPoint requires.</p> <p>See “Configuring AWS permissions for CloudPoint” on page 72.</p> <p>If you do not specify any value, the CFT creates a new IAM role with requisite permissions and attaches that role to the CloudPoint instance.</p> <p>In case of an upgrade scenario, Veritas recommends that you use the same IAM role that was attached to the existing CloudPoint instance. This is the same role with which the CloudPoint plug-in for AWS was configured.</p>

■ Network Configuration

Parameter	Description
CloudPoint VPC	Specify the ID of the Virtual Private Cloud (VPC) where you want to deploy the CloudPoint instance.
CloudPoint Subnet	<p>From the drop-down list, select the subnet ID of an existing subnet in the VPC where you want to deploy the CloudPoint instance.</p> <p>The drop-down list displays all the existing subnet IDs in the region where you are deploying CloudPoint.</p>
Inbound Access CIDR	Specify the CIDR to allow inbound access to the CloudPoint instance.

■ CloudPoint Configuration

Parameter	Description
CloudPoint User Name	<p>Specify a name for the CloudPoint administrator user account that is configured on the instance.</p> <p>The user name must be a valid email address.</p>

Parameter	Description
CloudPoint Password	<p>Specify the password for the administrator user account.</p> <p>The password must include a minimum of six characters and must not contain a space or an ampersand (&) character.</p>
Confirm CloudPoint Password	<p>Re-enter the password for the administrator user account.</p>
License Type	<p>Select the CloudPoint trial license that you wish to activate on the CloudPoint instance.</p> <p>Pick from one of the following options:</p> <ul style="list-style-type: none"> ■ Freemium <p>A Freemium license is a permanent license that does not expire and allows you to try out a subset of the CloudPoint features. This license lets you protect up to 10 TB of front-end terabyte (FETB) data.</p> ■ Evaluation <p>An Evaluation license is a 60-day time-bound license that allows you to try out all of the CloudPoint features. This license lets you protect up to 1000 TB of FETB data.</p> <p>See “Understanding your CloudPoint license” on page 13.</p>
Enable Telemetry <i>(optional)</i>	<p>Specify whether you want to enable or disable the telemetry service. When enabled, your CloudPoint usage information is shared with Veritas anonymously.</p>

■ CloudPoint Recovery Notification Configuration

Parameter	Description
SNS Topic ARN <i>(optional)</i>	<p>Specify the ARN of the SNS topic that you created for the CloudPoint stack.</p> <p>The SNS topic allows you to receive notifications whenever there is a change to the Auto Scaling Group (ASG).</p> <p>Veritas recommends that you configure an SNS Topic for the CloudPoint instance. The change notifications help you keep a track of the health of the CloudPoint instance.</p> <p>See “Instance failures and Auto Scaling Group behavior” on page 57.</p>

■ CloudPoint KMS Configuration

Parameter	Description
CMK ID <i>(optional)</i>	<p>Specify the ID of the AWS KMS customer master key (CMK) that you want to use to configure AWS KMS with CloudPoint.</p> <p>This parameter is not required if you do not want to use KMS with CloudPoint. If you do not specify this parameter, CloudPoint uses the default 256-bit AES specification to encrypt and decrypt all the configuration information.</p> <p>See “About CloudPoint integration with AWS KMS” on page 49.</p>
CMK Region <i>(optional)</i>	<p>Specify the region of the CMK whose ID is specified in the CMK ID field earlier.</p> <p>This parameter is not required if the CMK region is the same as where CloudPoint is being deployed.</p>

■ Security Configuration

Parameter	Description
Key Pair Name	<p>From the drop-down list, select the EC2 Key Pair that you want to use to enable SSH access to the CloudPoint instance.</p> <p>The drop-down list displays all the Key Pair names in the region where you want to deploy CloudPoint.</p>

9. Verify the parameter values and then click **Next**.
10. On the Options page, set any additional options (such as Tags, Permissions, Rollback Triggers) for the stack and then click **Next**.
11. On the Review page, review all the details that you have provided for the stack.

Under the Capabilities section, you may see an information box that displays a message informing you that this template may create additional IAM resources.

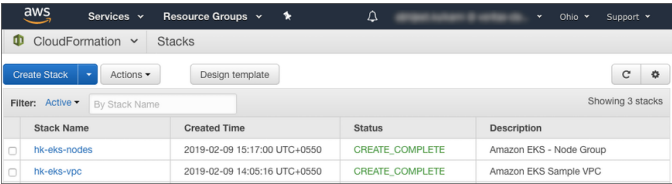
Select **I acknowledge that AWS CloudFormation might create IAM resources**. to acknowledge and confirm.

12. Verify all the details and then click **Create** to launch the stack.

Your stack now appears in the list of AWS CloudFormation stacks and the status appears as `CREATE_IN_PROGRESS`.

Select the stack and then click the **Events** tab to see the sequence of events that occur during the creation of the stack.

Click the **Resources** tab to see all the resources that are created for the stack.
13. After the stack is created successfully, the status of the stack changes to `CREATE_COMPLETE`.



Stack Name	Created Time	Status	Description
<input type="checkbox"/> hk-eks-nodes	2019-02-09 15:17:00 UTC+0550	CREATE_COMPLETE	Amazon EKS - Node Group
<input type="checkbox"/> hk-eks-vpc	2019-02-09 14:05:16 UTC+0550	CREATE_COMPLETE	Amazon EKS Sample VPC

This completes the process of setting up a CloudPoint stack using the CloudFormation template.

You can now connect to the CloudPoint instance, install required licenses, and then configure CloudPoint agents and plug-ins.

See [“Understanding your CloudPoint license”](#) on page 13.

See [“About plug-ins”](#) on page 65.

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 5-1 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 32.

See [“Configuring CloudPoint from your browser and signing in”](#) on page 37.

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 104.

See [“Determining the types of plug-ins and agents to install”](#) on page 66.

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 provides the following off-host plug-ins:
 - Amazon AWS
 - Google Cloud Platform
 - Microsoft Azure
 - Dell EMC Unity Array
 - Hewlett-Packard Enterprise Recovery Manager Central (RMC)
 - Pure Storage FlashArray
 - NetApp storage arrays
- CloudPoint provides 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:

- [AWS plug-in configuration notes](#)
- [Google Cloud Platform plug-in configuration notes](#)
- [Microsoft Azure plug-in configuration notes](#)
- [Dell EMC Unity array plug-in configuration notes](#)
- [Pure Storage FlashArray plug-in configuration notes](#)
- [HPE RMC plug-in configuration notes](#)
- [NetApp plug-in configuration notes](#)
- [Configuring an off-host plug-in](#)
- [About CloudPoint plug-ins and assets discovery](#)

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:

If CloudPoint is deployed on a on-premise host or a virtual machine:

Table 6-1 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.

If CloudPoint is deployed in the AWS cloud:

Table 6-2 AWS plug-in configuration parameters: cloud deployment

CloudPoint configuration parameter	Description
<i>For Source Account configuration</i>	
Regions	One or more AWS regions associated with the AWS source account in which to discover cloud assets. Note: If you deploy CloudPoint using the CloudFormation template (CFT), then the source account is automatically configured as part of the template-based deployment workflow.
<i>For Cross Account configuration</i>	
Account ID	The account ID of the other AWS account (cross account) whose assets you wish to protect using the CloudPoint instance configured in the Source Account.
Role Name	The IAM role that is attached to the other AWS account (cross account).

Table 6-2 AWS plug-in configuration parameters: cloud deployment
(continued)

CloudPoint configuration parameter	Description
Regions	One or more AWS regions associated with the AWS cross account in which to discover cloud assets.

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

Before you configure the plug-in, consider the following:

- You cannot delete automated snapshots of RDS instances and Aurora clusters through CloudPoint.
- You cannot take application-consistent snapshots of AWS RDS instances. Even though the CloudPoint UI allows you to create an application-consistent snapshot for such an instance, the actual snapshot that gets created is not application-consistent.
This is a limitation from AWS and is currently outside the scope of CloudPoint.
- All automated snapshot names start with the pattern `rds:.`
- CloudPoint does not support AWS Nitro-based instances that use EBS volumes that are exposed as non-volatile memory express (NVMe) devices. You will not be able to use CloudPoint to connect to and configure the AWS plug-in on such instances.

- CloudPoint does not support cross-account replication for AWS RDS instances or clusters, if the snapshots are encrypted using the default RDS encryption key (aws/rds). You cannot share such encrypted snapshots between AWS accounts. If you try to replicate such snapshots between AWS accounts, the operation fails with the following error:

```
Replication failed The source snapshot KMS key [<key>] does not exist,  
is not enabled or you do not have permissions to access it.
```

This is a limitation from AWS and is currently outside the scope of CloudPoint.

- If a region is removed from the AWS plug-in configuration, then all the discovered assets from that region are also removed from the CloudPoint assets database. If there are any active snapshots that are associated with the assets that get removed, then you may not be able perform any operations on those snapshots. Once you add that region back into the plug-in configuration, CloudPoint discovers all the assets again and you can resume operations on the associated snapshots.
- If you are creating multiple configurations for the same plug-in, ensure that they manage different regions. Two or more plug-in configurations should not manage the same set of cloud assets simultaneously.
CloudPoint currently does not block you from creating such a configuration. If there is an overlap of cloud assets between plug-in configurations, you may have to resolve the configuration issue by deleting the plug-in configurations and adding them again, ensuring that there are no overlapping assets.
However, CloudPoint does not allow you to delete a plug-in configuration if there are any snapshots associated with the assets in that configuration.

See “[Configuring an off-host plug-in](#)” on page 99.

Prerequisites for configuring the AWS plug-in

If the CloudPoint instance is deployed in the AWS cloud, do the following before you configure the plug-in:

- Create an AWS IAM role and assign permissions that are required by CloudPoint. See “[Configuring AWS permissions for CloudPoint](#)” on page 72.
Refer to the AWS documentation for instructions on how to create an IAM role: <https://docs.aws.amazon.com/IAM/latest/UserGuide/iam-roles-for-amazon-ec2.html#create-iam-role>
- Attach the IAM role to the CloudPoint instance.
Refer to the AWS documentation for instructions on how to attach an IAM role: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html#attach-iam-role>

Note: If you have deployed CloudPoint using the CloudFormation Template (CFT), then the IAM role is automatically assigned to the instance when the CloudPoint stack is launched.

- For cross account configuration, from the AWS IAM console (IAM Console > Roles), edit the IAM roles such that:
 - A new IAM role is created and assigned to the other AWS account (target account). Also, assign that role a policy that has the required permissions to access the assets in the target AWS account.
 - The IAM role of the other AWS account should trust the Source Account IAM role (**Roles > Trust relationships** tab).
 - The Source Account IAM role is assigned an inline policy (**Roles > Permissions** tab) that allows the source role to assume the role ("`sts:AssumeRole`") of the other AWS account.
 - The validity of the temporary security credentials that the Source Account IAM role gets when it assumes the Cross Account IAM role is set to 1 hour, at a minimum (**Maximum CLI/API session duration** field).

See "[Before you create a cross account configuration](#)" on page 77.

- If the assets in the AWS cloud are encrypted using AWS KMS Customer Managed Keys (CMK), then you must ensure the following:
 - If using an IAM user for CloudPoint plug-in configuration, ensure that the IAM user is added as a key user of the CMK.
 - For source account configuration, ensure that the IAM role that is attached to the CloudPoint instance is added as a key user of the CMK.
 - For cross account configuration, ensure that the IAM role that is assigned to the other AWS account (cross account) is added as a key user of the CMK.

Adding these IAM roles and users as the CMK key users allows them to use the AWS KMS CMK key directly for cryptographic operations on the assets.

Refer to the AWS documentation for more details:

<https://docs.aws.amazon.com/kms/latest/developerguide/key-policies.html#key-policy-default-allow-users>

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 73.

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 68.

AWS permissions required by CloudPoint

{

```
"Version": "2012-10-17",
"Statement": [
  {
    "Sid": "EC2AutoScaling",
    "Effect": "Allow",
    "Action": [
      "autoscaling:UpdateAutoScalingGroup",
      "autoscaling:AttachInstances"
    ],
    "Resource": [
      "*"
    ]
  },
  {
    "Sid": "KMS",
    "Effect": "Allow",
    "Action": [
      "kms:ListKeys",
      "kms:Encrypt",
      "kms:Decrypt",
      "kms:ReEncryptTo",
      "kms:DescribeKey",
      "kms:ListAliases",
      "kms:GenerateDataKey",
      "kms:GenerateDataKeyWithoutPlaintext",
      "kms:ReEncryptFrom",
      "kms:CreateGrant"
    ],
    "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",
        "rds:AddTagsToResource"
    ],
    "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": "EC2Backup",
    "Effect": "Allow",
    "Action": [
        "sts:GetCallerIdentity",
        "ec2:CreateSnapshot",
        "ec2:DescribeInstances",
        "ec2:DescribeInstanceStatus",
        "ec2:ModifySnapshotAttribute",
        "ec2:CreateImage",
        "ec2:CopyImage",

```

```

        "ec2:CopySnapshot",
        "ec2:DescribeSnapshots",
        "ec2:DescribeVolumeStatus",
        "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",
        "ec2:DescribeSecurityGroups"
    ],
    "Resource": [
        "*"
    ]
},
{
    "Sid": "EC2Recovery",
    "Effect": "Allow",
    "Action": [
        "ec2:RunInstances",
        "ec2:AttachNetworkInterface",
        "ec2:DetachVolume",
        "ec2:AttachVolume",
        "ec2>DeleteTags",
        "ec2:CreateTags",
        "ec2:StartInstances",
        "ec2:StopInstances",
        "ec2:CreateVolume",
        "ec2>DeleteVolume",
        "ec2:DescribeIamInstanceProfileAssociations",
        "ec2:AssociateIamInstanceProfile",
        "ec2:AssociateAddress"
    ],
    "Resource": [
        "*"
    ]
}

```

```
    ]  
  },  
  {  
    "Sid": "SNS",  
    "Effect": "Allow",  
    "Action": [  
      "sns:Publish",  
      "sns:GetTopicAttributes"  
    ],  
    "Resource": [  
      "arn:aws:sns:*:*:*"  
    ]  
  }  
]  
}
```

Before you create a cross account configuration

For CloudPoint cross account configuration, you need to perform the following additional tasks before you can create the configuration:

- Create a new IAM role in the other AWS account (target account)
- Create a new policy for the IAM role and ensure that it has required permissions to access the assets in that target AWS account
- Establish a trust relationship between the source and the target AWS accounts
- In the source AWS account, create a policy that allows the IAM role in the source AWS account to assume the IAM role in the target AWS account
- In the target AWS account, set the maximum CLI/API session duration to 1 hour, at a minimum

Perform the following steps:

- 1 Using the AWS Management Console, create an IAM role in the additional AWS account (the target account) whose assets you want to protect using CloudPoint.

While creating the IAM role, select the role type as **Another AWS account**.

Create role

Select type of trusted entity

AWS service
EC2, Lambda and others

Another AWS account
Belonging to you or 3rd party

Web identity
Cognito or any OpenID provider

SAML 2.0 federation
Your corporate directory

Allows entities in other accounts to perform actions in this account. [Learn more](#)

Specify accounts that can use this role

Account ID* 165323042987

Options ☒ Require external ID (Best practice when a third party will assume this role) ☒ Require MFA

- 2 Define a policy for the IAM role that you created in the earlier step.

Ensure that the policy has the required permissions that allow the IAM role to access all the assets (EC2, RDS, and so on) in the target AWS account.

Policy summary

Policy ARN: arn:aws:iam::165323042987:policy/cp-pun-test-policy

Description: To test EC2 permission template required for CR

Permissions | Policy usage | Policy versions | Access Advisor

Policy summary | {} JSON | Edit policy

Filter

Service	Access level	Resource	Request condition
Allow (3 of 146 services) Show remaining 143			
EC2	Limited: List, Read, Write	Multiple	None
RDS	Limited: List, Read, Write	All resources	None
STS	Limited: Read	All resources	None

3 Set up a trust relationship between the source and target AWS accounts.

In the target AWS account, edit the trust relationship and specify source account number and source account role.

The screenshot shows the AWS IAM console interface for the role **CP_CROSS_AC_ROLE**. The **Summary** tab is active, displaying the following details:

- Role ARN:** `arn:aws:iam::935928673755:role/CP_CROSS_AC_ROLE`
- Role description:** Edit
- Instance Profile ARNs:** Edit
- Path:** /
- Creation time:** 2018-05-30 10:31 UTC+0530
- Maximum CLI/API session duration:** 1 hour Edit
- Give this link to users who can switch roles in the console:** https://signin.aws.amazon.com/switchrole?roleName=CP_CROSS_AC_ROLE&account=veritas-status

Below the summary, the **Trust relationships** tab is selected. It shows a list of trusted entities:


- Trusted entities:** `arn:aws:iam::165323042987:role/CP_Role`

The **Conditions** section indicates that there are no conditions associated with this role.

This action allows only the CloudPoint instance hosted in source AWS account to assume the target role using the credentials associated with source account's IAM role. No other entities can assume this role.

4 Grant the source AWS account access to the target role.

In the source AWS account, from the account Summary page, create an inline policy and allow the source AWS account to assume the target role ("sts:AssumeRole").



5 From the target account's Summary page, edit the **Maximum CLI/API session duration** field and set the duration to **1 hour**, at a minimum.

This setting determines the amount of time for which the temporary security credentials that the source account IAM role gets when it assumes target account IAM role remain valid.

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 6-3 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.

Table 6-3 Google Cloud Platform plug-in configuration parameters
(continued)

CloudPoint configuration parameter	Google equivalent term and description
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.

GCP plug-in considerations and limitations

Consider the following before you configure this plug-in:

- If a zone is removed from the GCP plug-in configuration, then all the discovered assets from that zone are also removed from the CloudPoint assets database. If there are any active snapshots that are associated with the assets that get removed, then you may not be able perform any operations on those snapshots. Once you add that zone back into the plug-in configuration, CloudPoint discovers all the assets again and you can resume operations on the associated snapshots.
- If you are creating multiple configurations for the same plug-in, ensure that they manage different zones. Two or more plug-in configurations should not manage the same set of cloud assets simultaneously.
 CloudPoint currently does not block you from creating such a configuration. If there is an overlap of cloud assets between plug-in configurations, you may have to resolve the configuration issue by deleting the plug-in configurations and adding them again, ensuring that there are no overlapping assets.
 However, CloudPoint does not allow you to delete a plug-in configuration if there are any snapshots associated with the assets in that configuration.

See [“Google Cloud Platform permissions required by CloudPoint”](#) on page 82.

See [“Configuring a GCP service account for CloudPoint”](#) on page 83.

See [“Preparing the GCP service account for plug-in configuration”](#) on page 84.

See [“Configuring an off-host plug-in”](#) on page 99.

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 82.

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 the CloudPoint GCP plug-in configuration

- 1 Gather the GCP configuration parameters that CloudPoint requires.

See “Google Cloud Platform plug-in configuration notes” on page 80.

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": "someslogguid1234567890",
  "private_key": "-----BEGIN PRIVATE KEY-----\n
N11EvA18ADAN89kq4k199w08AQEFAA5C8KYw9951A9EAAo18AQCNvpuJ3oK974z4\n
.
.
.
weT9odE4ryl81tNU\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 \\"
```

```

    /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-----\
N11EvA18ADAN89kq4k199w08AQEFAA5C8KYw9951A9EAAo18AQcnpvpuJ3oK974z4
.
.
.
weT9odE4ryl81tNU\nV3q1XNX4fK55QTpd6CNu+f7QjEw5x8+5ft05DU8ayQcNkX
4pXJoDo154N52+T4qV4WkoFD5uL4NLpz5wxf1y\nNWcNfru8K8a2q1/9o0U+99==
-----END PRIVATE KEY-----

```

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 6-4 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

Consider the following before you configure the Azure plug-in:

- 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.
- CloudPoint does not support restore for Azure-managed Standard Solid State Drive (SSD) disks that are configured with the Locally redundant storage (LRS) redundancy option.
- CloudPoint does not support disk-based protection for applications that store data on virtual disks or storage spaces that are created from a storage pool. While taking snapshots of such applications, the disk-based option is not available.
- CloudPoint does not support snapshot tagging for assets in the Azure cloud environment. Even though Azure supports a maximum of up to 15 tags per snapshot, you cannot assign tags to snapshots, either manually using the APIs or via a protection policy, using CloudPoint.
- If you are creating multiple configurations for the same plug-in, ensure that they manage assets from different Tenant IDs. Two or more plug-in configurations should not manage the same set of cloud assets simultaneously.
 CloudPoint currently does not block you from creating such a configuration. If there is an overlap of cloud assets between plug-in configurations, you may have to resolve the configuration issue by deleting such plug-in configurations and adding them again, ensuring that there are no overlapping assets. However, CloudPoint does not allow you to delete a plug-in configuration if there are any snapshots associated with the assets in that configuration.

See [“Configuring an off-host plug-in”](#) on page 99.

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 85.

Dell EMC Unity array plug-in configuration notes

Table 6-5 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 281.

See [“Configuring an off-host plug-in”](#) on page 99.

Pure Storage FlashArray plug-in configuration notes

Table 6-6 Pure Storage FlashArray configuration parameters

CloudPoint configuration parameter	Description
IP address of Pure Storage	The IP address of the array.

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

CloudPoint configuration parameter	Description
Username to access Pure Storage	The username to access the array.
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 284.

See [“Configuring an off-host plug-in”](#) on page 99.

HPE RMC plug-in configuration notes

The CloudPoint plug-in for Hewlett Packard Enterprise (HPE) Recovery Manager Central (RMC) lets you create, delete, and restore snapshots of disks on all HPE storage systems that are supported by RMC. The plug-in supports clone and copy-on-write (COW) snapshot types.

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

See [“RMC plug-in configuration parameters”](#) on page 90.

See [“Supported HPE storage systems”](#) on page 91.

See [“Supported CloudPoint operations on HPE storage arrays”](#) on page 91.

See [“Configuring an off-host plug-in”](#) on page 99.

RMC plug-in configuration parameters

The following parameters are required for configuring the CloudPoint plug-in:

Table 6-7 RMC plug-in configuration parameters

CloudPoint configuration parameter	Description
IP address	The IP address of the RMC server
Username	The RMC administrator user account
Password	The password for the RMC admin user account

Before configuring the plug-in, ensure that the user account that you provide to CloudPoint has an admin role assigned on the RMC server.

Supported HPE storage systems

Table 6-8 Supported RMC version

Category	Supported
RMC software version	6.0 or later

Table 6-9 Supported RMC-managed storage systems

Category	Supported
Arrays	■ HPE 3PAR StoreServ

Supported CloudPoint operations on HPE storage arrays

CloudPoint supports the following operations on assets managed by HPE RMC:

Table 6-10 CloudPoint operations on assets managed by HPE RMC

CloudPoint operation	Description
Discover assets	<p>CloudPoint discovers all the volumes that are created on the array. If a volume is part of a multi-volume volume set, CloudPoint scans the volume set and extracts the individual volume information and then creates a list of all the unique volumes that are part of the volume set.</p> <p>For snapshots, CloudPoint scans all the snapshot sets and links each snapshot to its originating parent volume.</p>

Table 6-10 CloudPoint operations on assets managed by HPE RMC
(continued)

CloudPoint operation	Description
Create snapshot	<p>CloudPoint takes snapshots of all the volumes on the array.</p> <p>When CloudPoint takes a snapshot, it internally triggers a copy-on-write (COW) snapshot of the entire volume. If a volume is part of a multi-volume volume set, CloudPoint takes a snapshot of the entire volume set and creates a snapshot set. The snapshot set contains snapshots of all the volumes that are part of that volume set. However, CloudPoint associates that snapshot set only with the volume that was selected for the snapshot operation. Even if the volume set contains additional volumes, the snapshot set is associated only with the volume that was selected.</p> <p>For example, consider a volume set that contains three volumes, <code>vol-1</code>, <code>vol-2</code>, and <code>vol-3</code>. If you use CloudPoint to create a snapshot of <code>vol-1</code>, CloudPoint creates a snapshot set that includes snapshots of all the volumes in that volume set. But the snapshot set is marked as a snapshot of <code>vol-1</code> (the selected volume) even though the snapshot set includes additional snapshots belonging to the other volumes, <code>vol-2</code>, and <code>vol-3</code>.</p>
Delete snapshot	<p>CloudPoint deletes the snapshot or the snapshot set (if parent volume is part of a volume set).</p> <p>You can use CloudPoint to delete only those snapshots that are created using CloudPoint. If your RMC environment includes other snapshots, then CloudPoint can discover those snapshots, but the delete operation is not allowed for those snapshots.</p>
Restore snapshot	<p>When you restore a snapshot, CloudPoint only restores the particular snapshot corresponding to the selected volume. The snapshot set is a COW snapshot that can contain other snapshots belonging to the additional volumes in the volume set. However, CloudPoint only restores the snapshot for the selected volume. The other snapshots are not used during the restore operation.</p> <p>Ensure that the parent volume is unmounted from the target host before initiating a snapshot restore.</p>

Table 6-10 CloudPoint operations on assets managed by HPE RMC
(continued)

CloudPoint operation	Description
Export snapshot	<p>When a snapshot export operation is triggered, CloudPoint creates a new volume from the snapshot and then attaches the new volume to the target host.</p> <p>If the selected snapshot is a snapshot set, then while creating a new volume, CloudPoint creates a new volume set from the snapshot set. Even if the new volume set contains multiple volumes, CloudPoint attaches only the volume that corresponds to the snapshot that was selected for the export. The other volumes are not used in the export operation.</p> <p>The export operation is supported using the following protocols:</p> <ul style="list-style-type: none"> ■ Fibre Channel (FC) ■ Internet Small Computer Systems Interface (iSCSI)
Deport snapshot	<p>When a snapshot deport operation is triggered, CloudPoint detaches the volume from the target host and then deletes that volume. If the volume is part of a multi-volume volume set, then the entire volume set is detached and deleted from the host.</p>

Note: 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.

HPE RMC plug-in considerations and limitations

Consider the following when you configure the HPE EMC plug-in:

- When you delete snapshots using CloudPoint, only the snapshots that are managed by CloudPoint are available for deletion. You cannot use CloudPoint to delete snapshots that are not created using CloudPoint.
- CloudPoint operations are supported only on disks and volumes. Even if the volumes are grouped as a volume set, CloudPoint discovers and presents the volume set in the form of the individual volumes that are part of the volume set. If you create a snapshot of a volume that belongs to a multi-volume volume set, CloudPoint creates a snapshot set that includes snapshots of all the volumes in that volume set. The snapshot operation therefore results in the creation of additional snapshots and those are not tracked by CloudPoint.

If you want to use CloudPoint to protect volume sets, Veritas recommends that you configure a single volume in the volume set.

NetApp plug-in configuration notes

The CloudPoint plug-in for NetApp NAS and SAN lets you create, delete, restore, export, and deport snapshots of the following assets on the NetApp storage arrays:

- NetApp Logical Unit Number (LUNs) storage units in a SAN environment. LUNs appear under the **Disks** pane in the CloudPoint user interface (UI) dashboard.
- NetApp NFS shares in a NAS environment. Shares appear under the **Shares** pane in the CloudPoint UI dashboard.

NetApp plug-in configuration prerequisites

Before you configure the NetApp plug-in, verify the following:

- Ensure that the NetApp storage arrays have the necessary NetApp licenses that are required to perform snapshot operations.
- For NAS-based storage deployments, ensure that the NetApp shares are configured using an active `junction_path`.
- Ensure that the NetApp user account that you provide to CloudPoint has privileges to perform the following operations on the NetApp array:
 - create snapshot
 - delete snapshot
 - restore snapshot
- Ensure that the NetApp user account that you provide to CloudPoint is configured with `http` and `ontapi` access methods.
- Ensure that the NetApp user account that you provide to CloudPoint has the following roles assigned:
 - Default: readonly
 - lun: all
 - volume snapshot: all
 - vservers export-policy: all

See [“NetApp plug-in configuration parameters”](#) on page 95.

See [“Supported NetApp arrays”](#) on page 95.

See [“Supported CloudPoint operations on NetApp storage”](#) on page 95.

NetApp plug-in configuration parameters

The following parameters are required for configuring the NetApp NAS and SAN plug-in:

Table 6-11 NetApp plug-in configuration parameters

CloudPoint configuration parameter	Description
Array IP address	The cluster management IP address of the NetApp storage array or filer
Username	A NetApp user account that has permissions to perform snapshot operations on the NetApp storage array or filer
Password	The password of the NetApp user account

Supported NetApp arrays

You can use CloudPoint to discover and protect the following NetApp storage array models:

Table 6-12 Supported NetApp arrays

Category	Supported
Array model	<ul style="list-style-type: none">■ FAS/NAS 2552■ FAS/NAS 3240
Firmware version	<ul style="list-style-type: none">■ 2.3.2 (FAS/NAS 2552)■ 1.5.2 (FAS/NAS 3240)
ONTAP version	8.3.2 and later

Supported CloudPoint operations on NetApp storage

You can perform the following operations on supported NetApp storage arrays:

Table 6-13 CloudPoint operations on NetApp storage

CloudPoint operation	Description
Discover assets	<ul style="list-style-type: none"> In a SAN deployment, CloudPoint discovers the LUNs that are created from storage volumes. Only LUNs whose status is online, read-write operations are enabled, and the Snapshot auto delete parameter is set to false, are discoverable. <pre>["state": "online", "vol_type": "rw", "is_snapshot_auto_delete_enabled": "false"]</pre> Note: In a SAN deployment, CloudPoint can discover only the snapshots that are created using CloudPoint. In a NAS deployment, CloudPoint discovers all the NFS shares on the NetApp storage. The shares must have an active <code>junction_path</code> configured so that CloudPoint can discover them.
Create snapshot	<ul style="list-style-type: none"> In a SAN deployment, CloudPoint takes a snapshot of the NetApp LUNs. When CloudPoint takes a LUN snapshot, it internally triggers a copy-on-write (COW) snapshot of the entire volume to which the LUN belongs. If the volume contains multiple LUNs, the snapshot includes data from all the LUNs that reside on that volume. Additionally, CloudPoint also adds the LUN's serial number as a prefix in the snapshot name. All snapshot names use the following convention: <pre><lun_serial_number>.<specified_snapshot_name></pre> For example, if 5431fd754 is the LUN serial number and <code>test_snapshot</code> is the snapshot name provided, CloudPoint assigns the following name to the snapshot: <pre>5431fd754.test_snapshot</pre> In a NAS deployment, CloudPoint takes a snapshot of the NetApp NFS shares. Note: Snapshot names can contain integers, letters, parenthesis ('()'), hyphen ('-'), underscore ('_'), plus ('+'), and dot('.') characters.
Delete snapshot	<ul style="list-style-type: none"> In a SAN deployment, when you delete a LUN snapshot, CloudPoint internally deletes the snapshot of one or more volumes to which the LUN belongs. In a NAS deployment, CloudPoint deletes the snapshot of the share.

Table 6-13 CloudPoint operations on NetApp storage (*continued*)

CloudPoint operation	Description
Restore snapshot	<ul style="list-style-type: none">■ In a SAN deployment, when you restore a snapshot, CloudPoint only restores the particular LUN on which the restore is triggered. The LUN snapshot is a COW snapshot of the underlying volume and that volume can contain multiple additional LUNs. Even if the snapshot contains data from multiple LUNs, the restore is performed only for the selected LUN. Data on the other LUNs remains unchanged.■ In a NAS deployment, CloudPoint restores the snapshot of the share.
Export snapshot	<ul style="list-style-type: none">■ In a SAN deployment, when a snapshot export operation is triggered, CloudPoint creates a LUN from the snapshot and attaches it to target host. The target host is assigned read-write privileges on the exported LUN. The export operation is supported using the following protocols:<ul style="list-style-type: none">■ Fibre Channel (FC)■ Internet Small Computer Systems Interface (iSCSI)■ In a NAS deployment, when a snapshot export operation is triggered, a new rule is created in the export policy and is assigned to the exported snapshot that is available as a network share. The target host is assigned read-only privileges on the exported snapshot share. The export operation is supported using the NFS protocol. Note: The export policy assigned to the exported snapshot share must be other than the default policy assigned by the array. If the share is assigned the default policy, then the snapshot export operation will fail.
Deport snapshot	<p>In a SAN deployment, when a snapshot deport operation is triggered, CloudPoint removes the LUN mapping from the target host and then deletes the LUN.</p> <p>In a NAS deployment, when a snapshot deport operation is triggered, CloudPoint deletes the new rule that was created in the export policy when the snapshot was exported.</p>

Snapshot export related requirements and limitations

The following requirements and limitations are applicable in a NetApp environment:

- The host on which the snapshot is to be exported must be zoned and added to the Storage Virtual Machine (SVM) where you wish to attach or export that snapshot.
- The CloudPoint snapshot export operation fails for shares that are assigned the default array export policy. Ensure that you assign a different export policy (other than the default) to the share before you run the export operation.
- 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 SAN Export:

```
curl -H "Content-Type: application/json" -H "Authorization: Bearer  

<token>" -X POST -d '{"host-name":"offhost_server",  

"protocol":"fc|iscsi", "port":"<wwn given to host>"}' -k  

https://localhost/cloudpoint/api/v2/assets/<disk-ID>/snapshots/<snap-id>/exports/
```

For SAN 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>
```

For NAS Export:

```
curl -H "Content-Type: application/json" -H "Authorization: Bearer  

<token>" -X POST -d '{"host-name":"client ip to add rule",  

"protocol":"nfs", "port":"-"}' -k  

https://localhost/cloudpoint/api/v2/assets/<disk-ID>/snapshots/<snap-id>/exports/
```

For NAS 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 287.

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 68.

See [“Dell EMC Unity array plug-in configuration notes”](#) on page 89.

See [“Google Cloud Platform plug-in configuration notes”](#) on page 80.

See [“Microsoft Azure plug-in configuration notes”](#) on page 85.


See [“Pure Storage FlashArray plug-in configuration notes”](#) on page 89.

See [“HPE RMC plug-in configuration notes”](#) on page 90.


See [“NetApp plug-in configuration notes”](#) on page 94.

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**.

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 80.

- 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.

About CloudPoint plug-ins and assets discovery


After you configure the CloudPoint plug-ins, CloudPoint automatically starts discovering all the assets that are managed by the plug-in. Each plug-in handles its own discovery process independently. For example, when you configure the Dell EMC array plug-in, CloudPoint discovers all the assets on the configured Unity hardware array.

The discovery process is cyclic and is triggered at periodic intervals. Each plug-in has a predefined discovery interval value that is automatically set when you configure the plug-in. The discovery interval is defined in the plug-in configuration and is stored in a MongoDB database within the CloudPoint configuration.

Table 6-14 CloudPoint assets discovery interval default value

CloudPoint plug-in	Default assets discovery interval
AWS plug-in	1 hour
All other supported plug-ins	10 minutes

When the plug-ins begin their discovery cycle, the assets start appearing in the CloudPoint UI. After a discovery process is complete, the next discovery cycle is triggered only after the set interval has elapsed. The UI indicates when the discovery was last triggered and the next scheduled discovery. You can also manually trigger an asset discovery using the **Discover** button from the CloudPoint UI.



The discovery interval value and the time the last discovery was triggered are also logged in the `flexsnap-agent` logs.

See [“Viewing the assets discovery interval setting”](#) on page 103.

Plug-in discovery interval requirements and limitations

The following conditions apply to the assets discovery interval setting:

- This is not supported for the CloudPoint agentless option.
- This is not supported for the CloudPoint on-host plug-ins, such as the SQL plug-in, MongoDB plug-in, and the Oracle plug-in.
- A plug-in must first be configured before modifying the discovery interval setting.

Viewing the assets discovery interval setting

You can use the following CloudPoint REST API to view the discovery interval value that is set for a plug-in:

```
# curl -k -X GET "https://<cloudpointhostFQDN>/cloudpoint/api/v2/agents/{agentID}/plugins/{pluginName}/configs/{configID}"  
-H "accept: application/json"  
-H "Authorization: Bearer <authtoken>"
```

Parameter	Description
<cloudpointhostFQDN>	Represents the Fully Qualified Domain Name (FQDN) that was specified while performing the initial CloudPoint configuration on the host.
<authtoken>	Represents the alpha numeric authentication token that you generated in the earlier step.
{agentID}	Represents the ID of the CloudPoint agent.
{pluginName}	Represents the name of the plug-in. For example, for AWS plug-in, enter <code>aws</code> .
{configID}	Represents the ID of the plug-in configuration.

The command output displays all the details about the plug-in. The discovery interval details appear as follows:

```
"pollInterval": {  
  "value": 1,  
  "unit": "hours"  
}
```

In this sample, we can see that the interval is set to 1 hour.

Configuring the on-host agents and plug-ins

This chapter includes the following topics:

- [About agents](#)
- [Oracle plug-in configuration notes](#)
- [MongoDB plug-in configuration notes](#)
- [Microsoft SQL plug-in configuration notes](#)
- [About the installation and configuration process](#)
- [Preparing to install the Linux-based on-host agent](#)
- [Preparing to install the Windows-based on-host agent](#)
- [Downloading and installing the on-host agent](#)
- [Configuring the Linux-based on-host agent](#)
- [Configuring the Windows-based on-host agent](#)
- [Configuring the on-host plug-in](#)
- [Configuring VSS to store shadow copies on the originating drive](#)

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 65.

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

Table 7-1 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■ Oracle 18c	✓	X

Table 7-1 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	X	✓
On-premises storage array <ul style="list-style-type: none">■ Dell EMC Unity arrays■ Hewlett-Packard Enterprise (HPE) storage arrays supported by HPE RMC■ Pure Storage Flash Array■ NetApp storage arrays	X	✓

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.
- 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.

- CloudPoint supports discovery and snapshot operations on databases that are in a backup mode. After taking snapshots, the state of the databases is retained as is; CloudPoint does not change the status of such databases. However, in-place restore for such databases is not supported.

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.

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, then journaling must be turned on.

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

Table 7-2 Configuration parameters for MongoDB plugin

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.
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.

Microsoft SQL plug-in configuration notes

You can configure the CloudPoint plug-in for Microsoft SQL to discover and protect SQL Server application databases using disk-level and host-level snapshots.

After you configure the plug-in, CloudPoint automatically discovers all the file system assets and the SQL instances running on the Windows host. The discovered SQL assets appear as `MsSqlDbInstance <instancename>` on the Asset Management page in the CloudPoint user interface (UI). When you select the instance and take a snapshot, CloudPoint includes all the databases in that snapshot.

Before you configure the plug-in, ensure that your environment meets the following requirements:

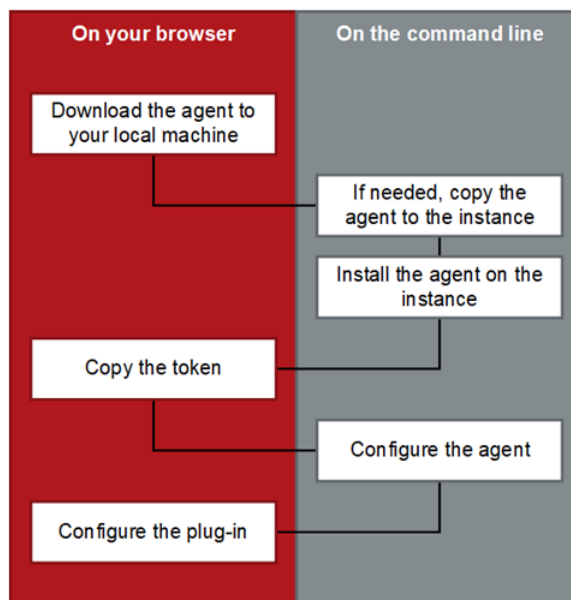
- This plug-in is supported in Microsoft Azure and Amazon AWS environments only.
- A supported version of Microsoft SQL Server is installed on the Windows instance.
See [“Meeting system requirements”](#) on page 19.
- Only standalone SQL deployments are supported. CloudPoint supports default as well as named instances.
Ensure that the instances are discoverable.

- The SQL Server instances that you want to protect must be running on a non-system drive.
 CloudPoint also does not support SQL Server instances that are installed on a mount point.
 Even though CloudPoint discovers SQL instances that are installed on a mount point, the snapshot operations will fail.
- CloudPoint uses the Microsoft Volume Shadow Copy Service (VSS) to create a single shadow copy per drive.
 Ensure that you configure VSS to store shadow copies on the same drive (the originating drive) where the database resides.
 See [“Configuring VSS to store shadow copies on the originating drive”](#) on page 118.

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 7-1 CloudPoint on-host agent installation and configuration process



See [“Preparing to install the Linux-based on-host agent”](#) on page 110.

See [“Preparing to install the Windows-based on-host agent”](#) on page 110.

See [“Downloading and installing the on-host agent”](#) on page 111.

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 the extra EPEL repositories using the following command:

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

During the on-host agent installation, the following packages and other required dependencies are automatically installed:

- Networking tools (`net-tools`)
- Python2 Pika package (`python2-pika`)
- Open SSL version 1.0.2k or higher (`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 107.
See [“About the installation and configuration process”](#) on page 109.

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
- Verify that the `pagefile.sys` is not present on the drive or volume that you wish to protect using CloudPoint. If the file exists on such drives, move it to an alternate location.

Restore of the snapshot will fail to revert the shadow copy if the `pagefile.sys` resides on the same drive or volume on which the operations are being performed.

See [“About the installation and configuration process”](#) on page 109.

Downloading and installing the on-host agent

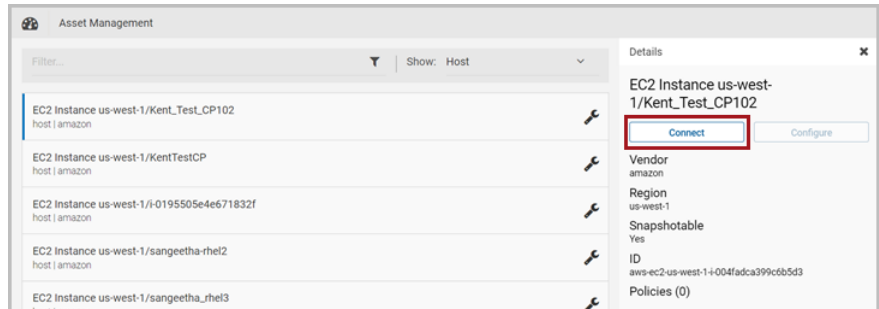
Before you complete the steps in this section, do the following:

- 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 13.
- 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 110. See [“Preparing to install the Windows-based on-host agent”](#) on page 110.

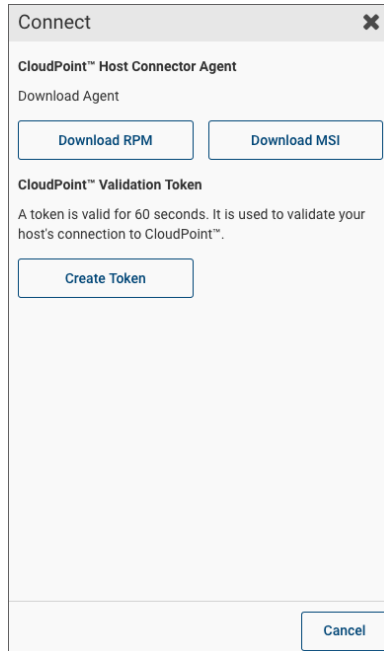
Whether you install the Linux-based on-host agent or the Windows-based on-host agent, the steps are similar.

To download and install the on-host agent

- 1 Sign in to the CloudPoint user interface (UI). See [“Signing in to CloudPoint”](#) on page 142.
- 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, do the following:
 - To download the Linux-based agent, click **Download RPM**.
 - To download the Windows-based agent, click **Download MSI**.



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.
- 6 Install the on-host agent.
 - For the Linux-based agent, type the following command on the Linux instance:

```
# sudo yum -y install 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 yum -y install VRTScloudpoint-agent-2.2-RHEL7.x86_64.rpm
```
 - For the Windows-based agent, run the agent package file and follow the installation wizard workflow to install the on-host agent on the Windows instance.

Alternatively, you can also install the Windows-based agent in a silent mode by running the following command on the Windows host:

```
msiexec /i <installpackagefilepath> /qn
```

Here, *<installpackagefilepath>* is the absolute path of the installation package. For example, if the installer is kept at `C:\temp`, then the command syntax is as follows:

```
msiexec /i C:\temp\cpwin_agent.msi /qn
```

In this mode, the installation package does not display any UI and also does not require any user intervention. The agent is installed at `C:\Program Files\Veritas\CloudPoint` by default and the path cannot be modified.

The silent mode of installation is particularly useful if you want to automate the agent installation using a third-party deployment tool.

7 Proceed to configure the on-host agent.

See [“Configuring the Linux-based on-host agent”](#) on page 113.

See [“Configuring the Windows-based on-host agent”](#) on page 115.

Configuring the Linux-based on-host agent

Verify the following before you configure the Linux-based on-host agent:

- Ensure that you have downloaded and installed the agent on the Linux instance. See [“Downloading and installing the on-host agent”](#) on page 111.
- To complete the steps in this section, you need root privileges on the Linux instance.
- If the CloudPoint Linux-based on-host agent was already configured on the host earlier, and you wish to re-register the agent with the same CloudPoint instance, then do the following on the Linux host:
 - Remove the `/opt/VRTScloudpoint/keys` directory from the Linux host.

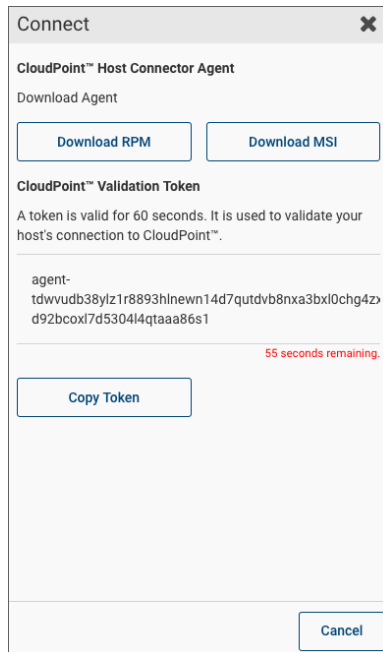
Type the following command on the host where the agent is running:

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

To configure the Linux-based on-host agent

- 1** 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**.
- 2 On the **Connect** dialog box, 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.
 - 3 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.

- 4 On the Linux host, register the on-host agent using the following command:

```
# sudo flexsnap-agent --ip <cloudpoint_host_FQDN_or_IP> --token <authtoken>
```


Here, *<cloudpoint_host_FQDN_or_IP>* is the CloudPoint host's Fully Qualified Domain Name (FQDN) or IP address that was used during the CloudPoint initial configuration.

<authtoken> is the authentication token that you copied in the earlier step.

Note: You can use `flexsnap-agent --help` to see the command help.

CloudPoint performs the following actions when you run this command:

- registers the Linux-based on-host agent
- creates a `/etc/flexsnap.conf` configuration file on the Linux instance and updates the file with CloudPoint host information
- enables and then starts the on-host agent service on the Linux host

Note: If you encounter an error, check the `flexsnap-agent` logs to troubleshoot the issue.

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

See [“Configuring the on-host plug-in”](#) on page 117.

Configuring the Windows-based on-host agent

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

See [“Downloading and installing the on-host agent”](#) on page 111.


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 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**.
- 2 On the Connect dialog box, 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.

3 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.

4 On the Windows instance, register the on-host agent.

From the command prompt, navigate to the on-host agent installation directory and type the following command:

```
flexsnap-agent.exe --ip <cloudpoint_host_FQDN_or_IP> --token  
<authtoken>
```

The agent installation directory is the path you specified while installing the Windows on-host agent using the installation wizard earlier. The default path is C:\Program Files\Veritas\CloudPoint\.

Here, <cloudpoint_host_FQDN_or_IP> is the CloudPoint host's Fully Qualified Domain Name (FQDN) or IP address that was used during the CloudPoint initial configuration.

<authtoken> is the authentication token that you copied in the earlier step.

Note: You can use `flexsnap-agent.exe --help` to see the command help.

CloudPoint performs the following actions when you run this command:

- registers the Windows-based on-host agent
- creates a `C:\ProgramData\Veritas\CloudPoint\etc\flexsnap.conf` configuration file on the Windows instance and updates the file with CloudPoint host information
- enables and then starts the on-host agent service on the Windows instance

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

See [“Configuring the on-host plug-in”](#) on page 117.

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 113.

See [“Configuring the Windows-based on-host agent”](#) on page 115.

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 106.

See [“MongoDB plug-in configuration notes”](#) on page 107.

See [“Microsoft SQL plug-in configuration notes”](#) on page 108.

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 **Configure** button is enabled.

3 Click **Configure**.

- 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 **Save**.

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
```

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 121.

See [“Configuring the agentless feature”](#) on page 122.

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 258.
- 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 121.

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

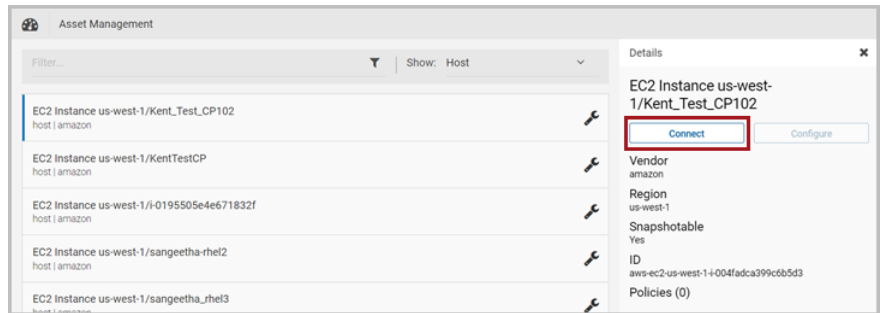
Verify all the prerequisites before you configure the CloudPoint agentless feature.

See [“Prerequisites for the agentless configuration”](#) on page 121.

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.

The screenshot shows the 'Connect' dialog box. At the top, there are two tabs: 'SSH' (which is selected and highlighted in blue) and 'RPM'. Below the tabs, there are two main sections. The first section is for 'Username' with a required field indicator (*). The second section has two tabs: 'Password' (selected and highlighted in blue) and 'SSH Key'. Below these tabs is a 'Password' field with a required field indicator (*). At the bottom right of the dialog, there are 'Save' and 'Cancel' buttons.

5 Enter the SSH user name, and either the SSH password or SSH key.

6 Click **Save**.

Configuring users

- [Chapter 9. Setting up email and adding users](#)
- [Chapter 10. 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 9-1 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 9-1 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.

- 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 9-2 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 128.

See [“Adding users to CloudPoint manually”](#) on page 129.

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**.

- 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 129.

Adding users to CloudPoint manually

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

Figure 9-2 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 125.

To add a CloudPoint user manually

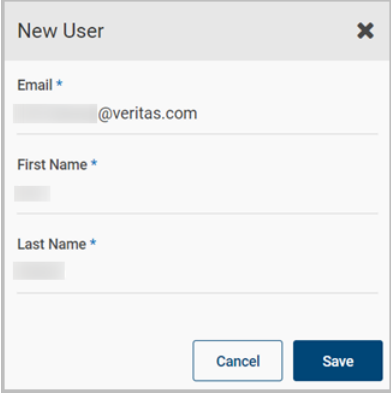
- 1 On the CloudPoint dashboard, in the **Administration** card, locate **Users**, and click **Manage**.



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**.

A screenshot of a 'New User' dialog box. The dialog has a title bar with 'New User' and a close button (X). It contains three text input fields: 'Email *' with a placeholder '@veritas.com', 'First Name *', and 'Last Name *'. At the bottom right are two buttons: 'Cancel' and 'Save'.

Note: Ensure that the specified email address does not include an underscore character. CloudPoint currently does not support adding users whose email addresses contain the underscore character.

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 10-1 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 12.

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 135.
- 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 139.


See [“Deleting a role”](#) on page 140.

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

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.

The screenshot shows the 'Add New Role' form. The 'Role Name' field contains 'AWS admin'. The 'Role Description' field contains 'Administers AWS assets in CloudPoint'. Below the description, a message states: '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 and highlighted in blue. It shows a list of users with a search filter 'Filter...' and a dropdown arrow. The first user is '@veritas.com' with a checked checkbox. The second user is 'admin' with an unchecked checkbox. At the bottom right, there are 'Cancel' and 'Save' buttons.

■ Permissions

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

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.

Users Permissions Assets

☐ Filter...

- ☒ ADMINISTRATOR
- ☒ USER_MANAGEMENT
- ☒ SNAPSHOT_POLICY_MANAGEMENT
- ☒ CLASSIFICATION_POLICY_MANAGEMENT
- ☒ REPLICATION_POLICY_MANAGEMENT

Cancel Save

■ 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.

Users Permissions **Assets**

Available Assets

Filter...

- ☐ EBS Snapshot snap-000008d7349d52936
- ☐ EBS Snapshot snap-00005302e8a42eb67
- ☐ EBS Snapshot snap-0000f5bdc6aa43653
- ☐ EBS Snapshot snap-000174b6c68db4733

Assign Selected
Assign All
Remove All
Remove Selected

Assigned Assets

Filter...

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...

- ☐ **AWS admin**
Administers AWS assets in CloudPoint

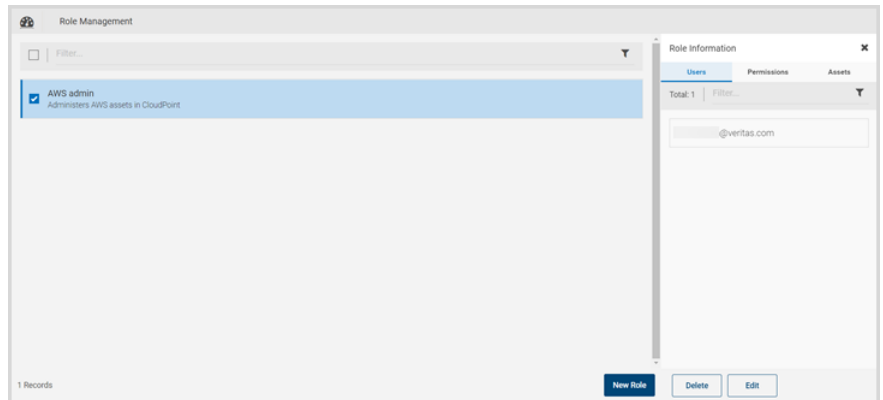
Editing a role

To edit a role

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

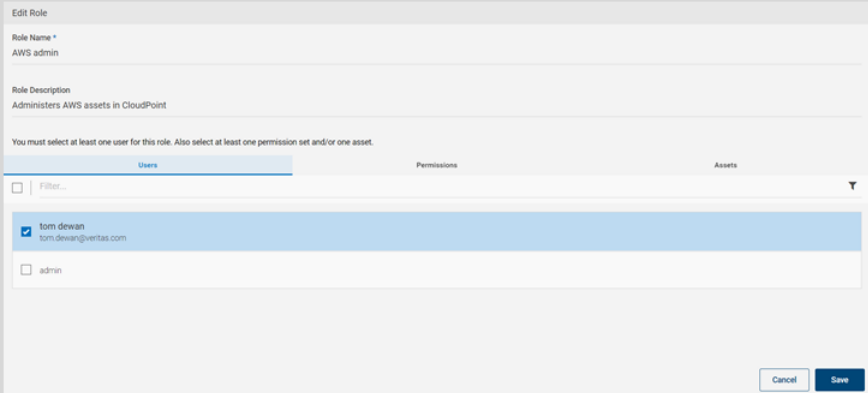


- 2 On the **Roles** page, select the check box for the role you want to view.



3 Click Edit.

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

**4 Modify the role values.**

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

See [“Creating a role”](#) on page 135.

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 11. User interface basics](#)
- [Chapter 12. Indexing and classifying your assets](#)
- [Chapter 13. Protecting your assets with policies](#)
- [Chapter 14. Tag-based asset protection](#)
- [Chapter 15. Replicating snapshots for added protection](#)
- [Chapter 16. Managing your assets](#)
- [Chapter 17. Monitoring activities with notifications and the job log](#)
- [Chapter 18. Protection and disaster recovery](#)

User interface basics

This chapter includes the following topics:

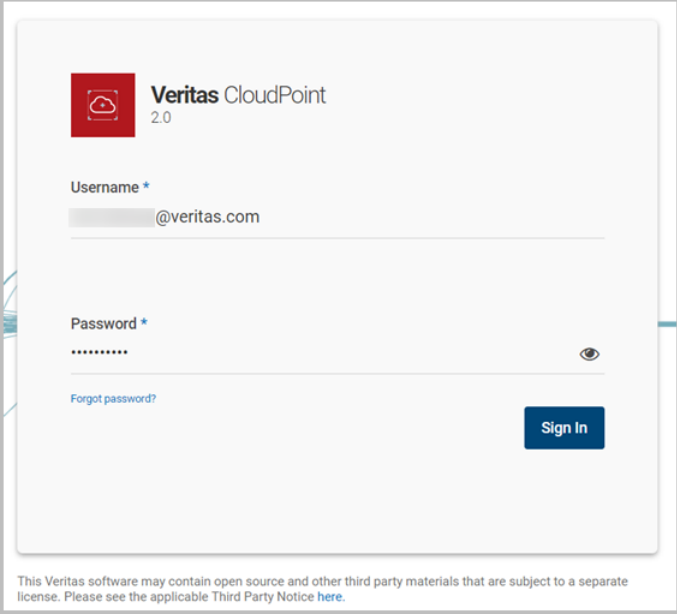
- [Signing in to CloudPoint](#)
- [Focusing on an asset type](#)
- [Navigating to your assets](#)
- [Using the action icons](#)

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.

The image shows the Veritas CloudPoint 2.0 sign-in interface. At the top left is the Veritas logo (a red square with a white cloud icon) followed by the text "Veritas CloudPoint 2.0". Below this are two input fields: "Username *" and "Password *". The username field contains the text "@veritas.com". The password field is filled with dots. To the right of the password field is an eye icon for toggling password visibility. Below the password field is a link that says "Forgot password?". At the bottom right is a blue button labeled "Sign In". At the very bottom of the screen, there is a small disclaimer: "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**.

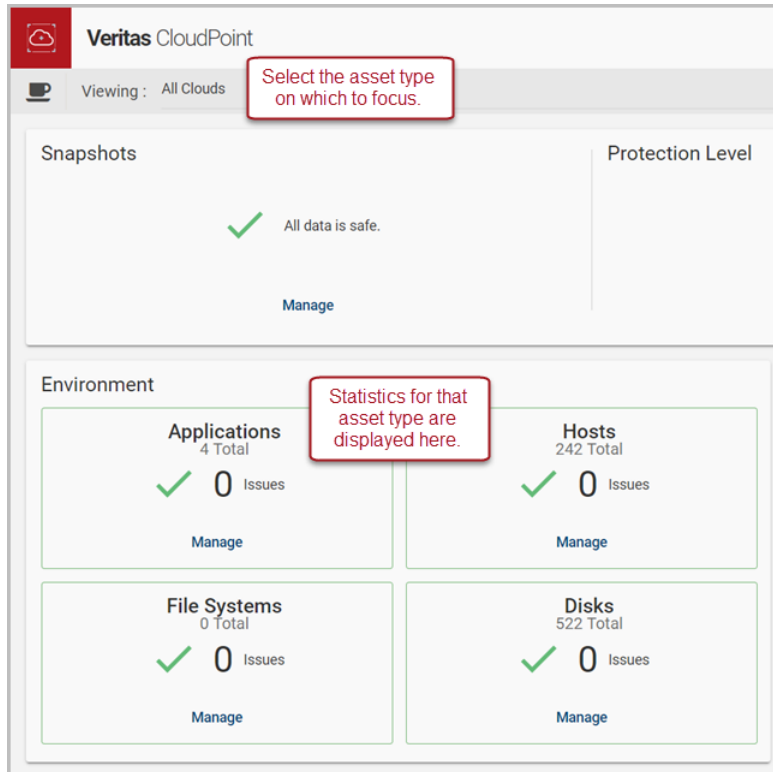
If this is the first time you have signed in to CloudPoint, verify that CloudPoint was installed successfully.

See [“Verifying that CloudPoint installed successfully”](#) on page 42.

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 11-1 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 11-1 Navigating to your assets *(continued)*





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

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 11-2 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).

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](#)
- [Indexing and classification statuses](#)

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.

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:

- Classification and indexing are licensed features and are not available with the CloudPoint Freemium license. Install or upgrade to a CloudPoint Enterprise or an equivalent license to enable and use these features in your CloudPoint deployment.
See [“Understanding your CloudPoint license”](#) on page 13.
See [“Upgrading your CloudPoint license”](#) on page 258.
- 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 to be successful, the snapshot that is being indexed must be local to the CloudPoint host so that it can be restored on the CloudPoint host in order to mount and scan the file system contents.
Indexing fails if the snapshot that is being indexed is not in the same region, cloud account, availability zone, or project as where the CloudPoint host resides. The job fails irrespective of whether the operation is triggered manually from the CloudPoint UI or using CloudPoint REST APIs.
In such cases, the CloudPoint UI does not display any errors indicating the failure. The only way to find out the status of an indexing job is to look through the coordinator logs and check for error messages.
- 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 149.

See [“Indexing and classifying snapshots”](#) on page 150.

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 147.
- enable at least one classification policy in Veritas Information Classifier (VIC)
See [“Configuring classification settings using VIC”](#) on page 149.

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 [“Indexing and classification statuses”](#) on page 150.

Indexing and classification statuses

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 12-1 Statuses for indexing and classification

Status	Description
Classified	The classification process is complete. No tags were found.

Table 12-1 Statuses for indexing and classification (*continued*)

Status	Description
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.
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.
Retrieving	CloudPoint is retrieving the snapshot to perform indexing and classification operations on the snapshot files.

See [“Indexing and classifying snapshots”](#) on page 150.

Protecting your assets with policies

This chapter includes the following topics:

- [About policies](#)
- [How a CloudPoint protection policy works](#)
- [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.
- 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 [“How a CloudPoint protection policy works”](#) on page 153.

See [“Creating a policy”](#) on page 157.

See [“Assigning a policy to an asset”](#) on page 160.

See [“Listing policies and displaying policy details”](#) on page 163.

How a CloudPoint protection policy works

After you create a CloudPoint protection policy with the desired parameters, you have to assign the policy to one or more assets. CloudPoint then triggers the policy runs as per the defined policy schedule. During each policy cycle, CloudPoint scans the assigned assets and performs the following actions:

- Creates snapshots of the assets to which the policy is assigned
- Replicates the snapshots if the replication option is enabled
- Deletes the asset snapshots, beginning with the oldest snapshot copy first, if the policy-triggered snapshot count is more than the defined retention value

During a policy run, when CloudPoint takes asset snapshots, it holds them together in a virtual object called as a snapshot group. A snapshot group represents a unique set of snapshots taken at a particular point in time, where each snapshot belongs to a particular asset that is included in the policy. A snapshot group contains only one snapshot per asset, it never includes two snapshots that belong to the same asset. CloudPoint creates a new and unique snapshot group in each policy run.

CloudPoint uses a snapshot group as the unit of reference for running policy-driven snapshot deletion operations. If the snapshot count exceeds the retention value defined in a policy, CloudPoint deletes the oldest policy-triggered snapshot copy during the policy run. CloudPoint uses the age (time when a snapshot was taken) and number (snapshot group count) of the snapshot group to determine which

snapshots are to be deleted. CloudPoint does not use the age or the number of the individual asset snapshots within a group.


For example, consider a policy **P1** that is assigned to assets **A1**, **A2**, and **A3** and is scheduled to run on the Monday, Tuesday, and Wednesday of a week. The retention count is set to 1, which means CloudPoint should maintain only one copy of the snapshots at any given time.

Table 13-1 CloudPoint policy behavior example


Policy name	Policy schedule	Retention value	Assigned to assets
P1	Run once a day on Mondays, Tuesdays, and Wednesdays	1	A1, A2, A3

Now let us look at the CloudPoint policy behavior, depending on different scenarios:

- During the first policy run on a Monday, CloudPoint creates three snapshots, **A1-Snap-Mon**, **A2-Snap-Mon**, **A3-Snap-Mon**, and puts them in a snapshot group called **G1**. The three snapshots represent one for each asset that is included in the policy.



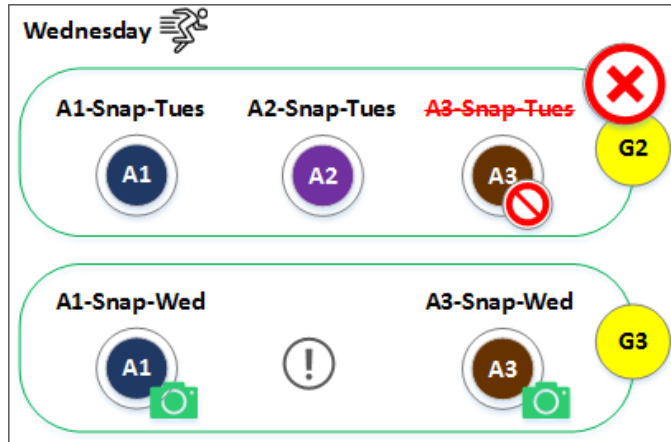
- In the Tuesday policy run, consider a case where the snapshot creation for asset **A3** fails. CloudPoint creates two snapshots, **A1-Snap-Tue** and **A2-Snap-Tues**, and puts them in a new snapshot group called **G2**.
With **G1** and the newly created snapshot group **G2**, the snapshot group count exceeds the policy retention value of 1. CloudPoint immediately triggers a snapshot delete operation and deletes all the snapshots included in group **G1**, as **G1** is the older of the two snapshot groups.



As $G1$ is deleted, there is only one snapshot group ($G2$) that remains at the end of the Tuesday policy run. $G2$ includes only two snapshots, one each for assets $A1$ and $A2$. Asset $A3$ does not have any snapshot as snapshot creation had failed in the Tuesday policy run and CloudPoint has already deleted snapshot group $G1$ that included an older snapshot of asset $A3$.

- Now consider a case where asset $A2$ is removed from the policy. Policy $P1$ is no longer associated with the asset $A2$. In the Wednesday policy run, CloudPoint creates two snapshots $A1$ -Snap-Wed and $A3$ -Snap-Wed and puts them in snapshot group $G3$. Snapshot for asset $A2$ is not created as $A2$ is no longer included in the policy.

With $G2$ and the newly created snapshot group $G3$, the snapshot count exceeds the policy retention value. CloudPoint triggers a snapshot delete operation once again. $G2$ is the older between $G2$ and $G3$, so CloudPoint deletes the snapshots in group $G2$.



As *G2* is deleted, there is only one snapshot group (*G3*) that remains at the end of the Wednesday policy run. *G3* includes two snapshots, one each for asset *A1* and asset *A3*. Asset *A2* does not have any snapshots.


Note the following:

- CloudPoint does not consider the individual asset snapshot count when performing the delete operation. Also, if the snapshot creation fails in a subsequent policy run, CloudPoint does not retain the asset snapshot that was created during the earlier policy cycle.
In this case, asset *A3* had only one snapshot in group *G1*, as the snapshot operation had failed in the Tuesday policy run. But with creation of *G2*, the snapshot group count exceeded the policy retention value and *G1* was deleted. As a result, asset *A3* does not have any snapshots remaining at the end of the Tuesday policy run, even if the policy retention value is 1.
- CloudPoint does not retain policy-created snapshots if an asset is removed from a policy. CloudPoint does not consider that an asset is no longer associated with the policy and proceeds with delete operation. An asset snapshot is deleted even if the policy is no longer associated with the asset.
In this case, during the Wednesday policy run, CloudPoint does not create a new snapshot for asset *A2* as it was excluded from the policy. However, with the creation of *G3*, the snapshot group count exceeded the policy retention value and *G2* was deleted. *G2* included a snapshot of asset *A2*, and even though *A2* is no longer associated with the policy, deletion of *G2* resulted in the deletion of the snapshot that belonged to *A2*.

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.

The 'New Policy' form is divided into two main sections: 'Policy Information' and 'Retention *'.
Policy Information: Includes fields for 'Policy Name *' (with a 12-character limit), 'Description' (with a 500-character limit), and 'Storage Level *' (a dropdown menu with the text 'Please select a storage level'). At the bottom, there are two checkboxes: 'Application Consistent' (checked) and 'Enable Replication' (unchecked).
Retention *: Features a numeric input field set to '0', a '+' icon, and buttons for 'Copies', 'Days', 'Weeks', 'Months', and 'Years'.
Scheduling *: Includes buttons for 'Hourly', 'Daily', 'Weekly', and 'Monthly'.
 At the bottom right, there are 'Save' and 'Cancel' buttons.

Enter the following:

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

Field	Description
Policy Name	<p>Enter a name for the policy.</p> <p>The name can contain 2 to 12 characters. The name can only contain lower case letters, numbers, and hyphens. The name should begin 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 2 to 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, Application, or Share (applicable to NetApp).</p>
Application Consistent	<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>

- **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.

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	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.

New Policy

Policy Information

Policy Name *
test

Description
test

Storage Level *

Host

☒ Application Consistent
☐ Enable Replication

Retention *

4 Copies Days Weeks Months Years

Scheduling *

Hourly Daily Weekly Monthly

Run at 12:00 AM on...

S M T W T F S

Save Cancel

- 4 Click **Save**.

CloudPoint displays a message that the new policy is created.

- 5 The Policies page displays all the policies that are created.

Policies

Filter...

☐ weeklyhost
Host-level application consistent snapshot every week. Keep 2 copies.
Host

☐ hourlydisk
Disk snapshots every 6 hours. Keep 1 copy.
Disk

☐ dailydisk
Disk snapshots daily. Keep 3 copies.
Disk

DISABLED

ENABLED

ENABLED

3 Records

New Policy

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.

Before you assign policies to assets, 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 steps to change the policy that is associated with an asset or to un-assign a policy from an asset.
- Do not assign a host-level protection policy to a disk asset. Doing so may lead to the creation of several unnamed snapshots that include snapshot of the disk itself, snapshot of the host to which that disk is attached, and snapshot of all the other disks that are attached to the host.

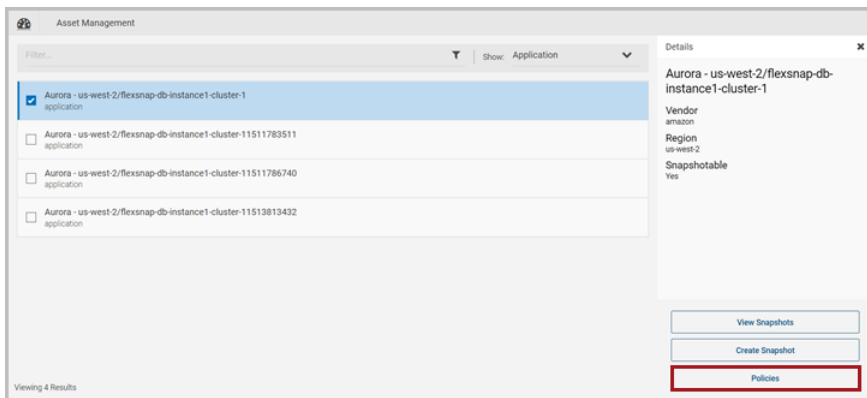
CloudPoint currently does not block you from assigning a host-level policy to a non-host entity such as a disk. But doing so can lead to unintended snapshot creation.

- 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**.

Policies for Aurora - us-west-2/flexsnap-db-instance1-cluster-1

Available Policies

Filter...

- ☐ daily disk level
Protection Level: disk
- ☒ weekly app consistent
Protection Level: application
- ☐ weekly disk level
Protection Level: disk

Applied Policies

Filter...

Assign Selected

Assign All

Remove All

Remove Selected

Cancel Save

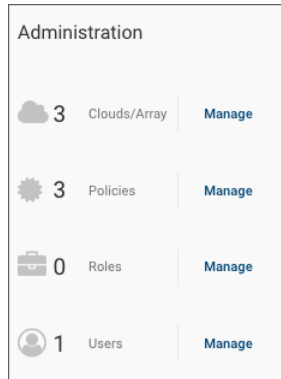
You can also assign or remove multiple policies at the same time.

- 4 Click **Save**.

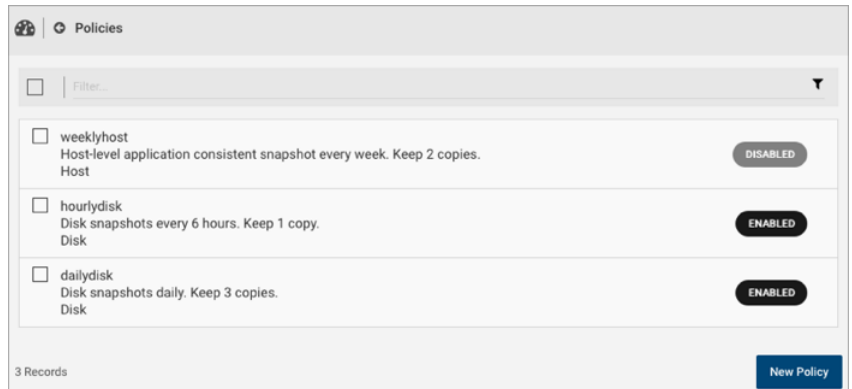
Listing policies and displaying policy details

To list policies and display policy details

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

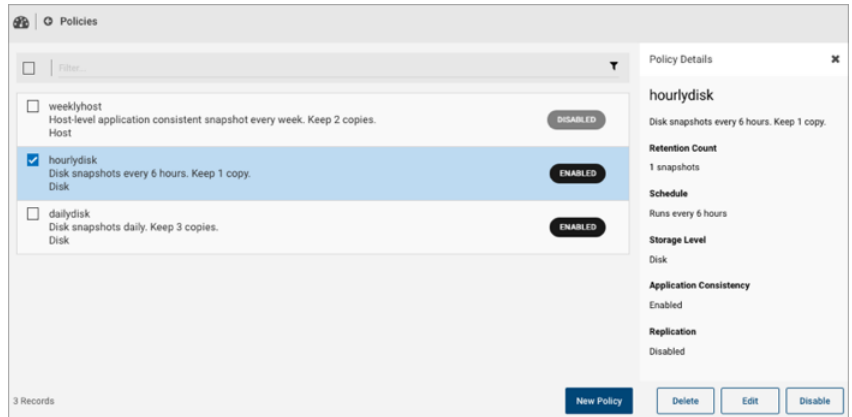


The **Policies** page displays a list of the policies that are created.



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
- The storage level that displays the kind of snapshot that the policy takes; whether disk, host, application, or share.

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

- Edit a policy.
- Delete a policy.
- Enable or disable a policy.

See [“About policies”](#) on page 152.

See [“Creating a policy”](#) on page 157.

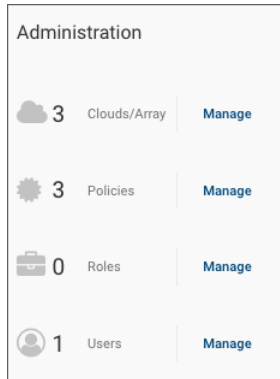
See [“Deleting a policy”](#) on page 166.

See [“Editing a policy”](#) on page 165.


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 On the **Edit Policy** page, modify the policy parameters as per your requirement.

The screenshot shows the 'Edit Policy' interface. On the left, under 'Policy Information', the 'Policy Name' is 'hourlydisk', the 'Description' is 'Disk snapshots every 6 hours. Keep 1 copy.', the 'Storage Level' is 'Disk', and 'Application Consistent' is checked. On the right, the 'Retention' section shows '1' copies. The 'Scheduling' section shows 'Hourly' scheduling with a frequency of '6' hour(s) OR '0' minute(s). 'Save' and 'Cancel' buttons are at the bottom right.

The remaining steps in this procedure are the same as those performed while creating a new policy.

See [“Creating a policy”](#) on page 157.

- 5 After you have finished editing the policy, click **Save**.

CloudPoint displays a message that the policy is updated.


See [“About policies”](#) on page 152.

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**.
- CloudPoint displays a message that the policy has been deleted.
- The deleted policy no longer appears on the **Policies** page.

See [“About policies”](#) on page 152.

See [“Creating a policy”](#) on page 157.

Tag-based asset protection

This chapter includes the following topics:

- [About tag-based asset protection](#)
- [How to use tag-based asset protection feature](#)
- [Tag-based asset protection support](#)
- [Tag-based asset protection considerations and limitations](#)

About tag-based asset protection

Note: Tag-based asset protection is introduced in CloudPoint 2.2 release and is offered only as a Technical Preview feature. It is included in this release to introduce you to the concept and the functionality. The feature should not be used in production environments. Veritas is under no obligation to develop, modify, improve, maintain, or market Technical Preview features or release production builds or general availability (GA) versions of the software.

The process of deploying a workload in the cloud is getting easier than in the past. Whether it is a development environment, a simple application instance, or a complex production deployment, virtual instances can be up and running within a few clicks. This ease of provisioning has led to a proliferation in the kinds of workloads and also the number of instances that are getting deployed in the cloud. The challenge is limited not only to the management of such a diverse environment, but also in the implementation of data protection policies in an ever expanding cloud footprint.

CloudPoint makes it easier by automatically discovering all the assets in your cloud environment. The periodic discovery ensures that any addition or deletion of workloads does not go unnoticed and all asset changes remain accounted for. You can easily configure a CloudPoint protection policy and assign it to the desired workloads.

The responsibility of data protection is entirely on the backup and data protection administrators. They need to monitor the number of workloads, determine the kind of data protection that is needed, and ensure that the required workloads are protected by assigning the correct protection policy. Also, if any of the protected workloads no longer exist or are no longer required to be protected, the task of removing that asset from the assigned protection policies also remains with the administrators.

With release 2.2, CloudPoint introduces a feature called as Tag-based asset protection that is designed to automate the data protection policy assignment process. Tag-based asset protection provides an intelligent and automated mechanism to protect assets based on user-defined tags. Tagging is a method where you can use descriptive text labels and assign them to the assets, either during the asset creation or at any time during the active life of the asset. When CloudPoint discovers the assets, it also scans the tags associated with those assets and then automatically assigns a matching pre-defined protection policy to the asset.

During each discovery cycle, CloudPoint checks the asset tags and automatically assigns a matching protection policy to the asset. Similarly, if a tag is removed from an asset, and if that asset is associated with a matching policy, then CloudPoint removes that asset from the policy. Then, the asset is no longer protected by that policy. Tag-based asset protection allows you to assign or unassign protection policies to a large number of assets simultaneously. This eliminates the manual task of navigating through the CloudPoint UI and choosing a protection policy for an asset individually.

How tag-based asset protection works

CloudPoint stores all the information about configured policies, discovered assets and their associated tags in a MongoDB database. The CloudPoint coordinator service and the policy engine service together make use of this information to control and manage data protection using the tag-based asset protection.

The CloudPoint coordinator keeps a track of all the changes to the assets and their tags. During the CloudPoint discovery cycle, the CloudPoint plug-in sends all the asset data to the coordinator. The coordinator compares this information with the records in the MongoDB database and identifies all the assets and asset tags that are added, modified, or removed since the last discovery and then sends that information to the policy engine. The policy engine uses this information and

determines whether to assign a policy to an asset or to remove an existing policy from an asset.

The CloudPoint policy engine keeps a track of all the changes that are made to CloudPoint policies. Whenever a new policy is created, the policy engine requests the coordinator for a list of tagged assets. The coordinator polls the MongoDB database and sends that information to the policy engine. If any of the asset tag matches the new policy name, the policy engine requests the coordinator to attach that policy to the asset.

Depending on the change, CloudPoint performs the following actions during each discovery cycle:

- If a new tagged asset is added, CloudPoint associates a matching policy to that asset.
- If a tagged asset is deleted, CloudPoint disassociates the asset from the policy.
- If an existing tag value is modified such that an existing policy name is replaced by a new one, CloudPoint disassociates the existing policy and associates a new matching policy to that asset.
- If a new policy name is appended to the existing asset tag, CloudPoint associates that new policy to the asset. The same asset is now protected by the existing as well as the new protection policy.
- If an existing tag is removed from an asset or if the tag does not include a policy name, CloudPoint disassociates the policy from that asset. The asset is then no longer being protected and is excluded from future policy runs.
- If an asset tag contains a policy name that does not exist, CloudPoint generates a notification alert. As and when a matching policy is created, CloudPoint automatically assigns it to that asset.

How to use tag-based asset protection feature

Using the tag-based asset protection feature involves the following steps:

Table 14-1 Using tag-based asset protection

Step #	Action
Step 1 - Create protection policies	<p>Create one or more CloudPoint protection policies with the desired snapshot and replication settings and a policy schedule. Ensure that you create the appropriate host or disk-level policy depending on the type of asset that you wish to protect.</p> <p>Make a note of the policy names. You are going to use the policy names in the tags that you create and assign to the assets.</p> <p>See "About policies" on page 152.</p>
Step 2 - Add tags to the assets	<p>Using the cloud vendor management tools, add tags to the assets that you wish to protect using the configured CloudPoint policies.</p> <p>The tags must be defined as a "key": "value" pair and must use the following convention:</p> <pre>"veritas-protection-policy": "<polycynname>"</pre> <p>Here, veritas-protection-policy is the name of the key.</p> <p>The key name is case sensitive; use this name as is, do not modify the key name.</p> <p><polycynname> is the value of the key and represents the CloudPoint protection policy that you wish to associate with the asset.</p> <p>For example, if the name of the policy is Policy1, the asset tag will be as follows:</p> <pre>"veritas-protection-policy": "Policy1"</pre> <p>Note: The policy name that you specify here is not case sensitive.</p> <p>If you want to associate multiple protection policies to the same asset, you can specify all the policies within a single tag.</p> <p>This is the most critical step in the process. If the assets are not tagged correctly, the feature will not work.</p>

Table 14-1 Using tag-based asset protection (*continued*)

Step #	Action
	<ul style="list-style-type: none"> In case of AWS and Azure cloud assets, specify multiple policy names separated using commas. For example, if want to associate policies <code>Policy1</code>, <code>Policy2</code>, and <code>Policy3</code>, the asset tag will be as follows: <code>"veritas-protection-policy":"Policy1, Policy2, Policy3"</code> In case of GCP cloud assets, specify multiple policy names separated using the underscore character. For example, using the same policy names listed earlier, the asset tag will be as follows: <code>"veritas-protection-policy":"Policy1_Policy2_Policy3"</code>
Step 3 - Configure the CloudPoint plug-in	<p>Configure the CloudPoint plug-in for the desired cloud environment (AWS, GCP, or Azure).</p> <p>See “About plug-ins” on page 65.</p> <p>When you configure the plug-in, CloudPoint immediately begins discovering the assets in that cloud and starts scanning the tags associated with the discovered assets. If an asset tag matches an existing CloudPoint policy name, CloudPoint automatically assigns that protection policy to the asset.</p> <p>If the asset tag does not match with any existing policy name, an informational alert notification is generated informing that the policy does not exist. However, if a policy by that name is created later, CloudPoint automatically assigns that policy to the asset in the next discovery cycle.</p>

The order of the tasks mentioned in the table does not matter. You do not need to perform these tasks in a sequential manner. All that is required is that you create the CloudPoint policy and ensure that the corresponding tag is associated with the asset that you wish to protect.

For example, you can add tags to the assets even before you install and configure CloudPoint. Once you have installed CloudPoint and configured the plug-in, you can create the requisite protection policies. During the discovery cycle, CloudPoint automatically checks the asset tags and assigns a matching policy.

Similarly, if you have already deployed CloudPoint, you can add appropriate tags to new assets that are not yet discovered by CloudPoint. During the discovery cycle, CloudPoint scans all the assets and their tags and automatically assigns a matching policy.

Tag-based asset protection support

Tag-based asset protection is offered as a Technical Preview feature. It should not be used in a production environment.

The following cloud environments are supported:

- Amazon Web Services (AWS)
- Microsoft Azure
- Google Cloud Platform (GCP)

The following asset types are supported:

- Hosts, instances, or virtual machines
- Disks
- RDS database instances (AWS)

Tag-based asset protection considerations and limitations

The following conditions are applicable to the tag-based asset protection feature:

- Tag-based asset protection is offered as a Technical Preview feature. It should not be used in a production environment.
- Tag-based asset protection works only for assets that support associating tags. If the assets cannot be assigned a tag in the form of a `key:value` pair, then you cannot use the CloudPoint tag-based asset protection for those assets. You will have to manually assign CloudPoint policies to protect such assets.
- The CloudPoint user interface (UI) does not support configuring tag-based asset protection. You will not be able to see the tags that are assigned to the assets or administer this feature from the UI. Tags must be assigned to the assets using the cloud vendor management tools.
- Newly created tagged assets are not assigned to a CloudPoint protection policy until the assets are discovered by CloudPoint.
- You cannot modify a policy name if the policy is associated with a tagged asset using tag-based asset protection.
- You cannot delete a policy if the policy is associated with a tagged asset using tag-based asset protection. To delete a policy, you must remove the policy name from the `"veritas-protection-policy"` key value (in case multiple policies

are specified in the tag), or delete the key if there are no other policies specified in the asset tag.

- You cannot manually disassociate a policy from a tagged asset if the policy is assigned using tag-based asset protection.
 You must first remove the policy name from the asset tag or delete the asset tag altogether. After the next discovery cycle, CloudPoint automatically removes the asset from that policy in a subsequent policy run.
- Tag-based asset protection feature will not take effect if the same policy is manually assigned to the tagged asset earlier. For example, consider a policy `Policy1` that is manually assigned to an asset `VM1`. Now, even if you add a tag `"veritas-protection-policy": "Policy1"` to `VM1`, CloudPoint ignores that tag and does not reassign the policy to that asset again. The policy association is considered as manual and tag-based asset protection functionality does not take effect on that asset or policy.
- All tag-based asset protection operational logs are stored in the `flexsnap-policy` and `flexsnap-coordinator` logs.

Replicating snapshots for added protection

This chapter includes the following topics:

- [About snapshot replication](#)
- [About cross-account snapshot replication in the AWS cloud](#)
- [Requirements for replicating snapshots](#)
- [Cross-account snapshot replication support matrix](#)
- [Cross-account snapshot replication limitations](#)
- [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.

See [“About cross-account snapshot replication in the AWS cloud”](#) on page 176.

About cross-account snapshot replication in the AWS cloud

You can use CloudPoint to replicate snapshots across regions that are associated with the same AWS account. Beginning with release 2.2, CloudPoint extends this feature support and also allows you to replicate snapshots across regions that are associated with different AWS accounts. You can now select an asset snapshot from a specific region associated with a particular AWS account and replicate that snapshot to a region that is associated with a different AWS account.

In case of cross-account replication, CloudPoint initiates a 3-step process where it first shares the snapshot with the selected region that belongs to a different AWS account, copies that snapshot to that region, and then unshares the snapshot from that region. The replication mechanism is handled internally and the entire process is completely transparent to the user.

There is no change to the overall snapshot replication workflow. You can perform the replication operation from the CloudPoint user interface (UI), where you select the replication target from the list of available regions and CloudPoint then replicates the snapshot to the selected target. You must first configure the cross account to be able to select the regions belonging to that account. There are no additional steps required for cross-account replication.

You can also use the replicated snapshot to perform a restore operation and launch a new instance. This allows you to easily bring up another instance of your application workload in the desired region, in case the original workload becomes unavailable due to an unplanned event.

Note: Cross-account snapshot replication support is available for AWS cloud only.

See [“Requirements for replicating snapshots”](#) on page 177.

See [“Cross-account snapshot replication support matrix”](#) on page 178.

See [“Cross-account snapshot replication limitations”](#) on page 178.

Requirements for replicating snapshots

For replicating unencrypted snapshots

Ensure that you add the AWS source and cross account (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>
```
- For cross-account replication, the Customer Managed Key (CMK) that is used for encrypting the snapshot in the source region must be shared with the AWS account associated with the target region.
- For cross-account replication, the AWS IAM user or role associated with the snapshot source region's AWS account (source AWS account) must have the following permissions:
 - `ModifySnapshotAttribute` and `CopySnapshot` on the EC2 instance
 - `DescribeKey` and `ReEncrypt` on the KMS key that is used to encrypt the original snapshot
- For cross-account replication, the AWS IAM user or role associated with the snapshot replication target region's AWS account (target AWS account) must have the following permissions:
 - `CreateGrant`, `DescribeKey`, and `Decrypt` on the KMS key that is used to encrypt the original snapshot
 - `CreateGrant`, `Encrypt`, `Decrypt`, `DescribeKey`, and `GenerateDataKeyWithoutPlainText` on the KMS encryption key used while performing the `CopySnapshot` operation on the original snapshot

See [“AWS permissions required by CloudPoint”](#) on page 73.

Cross-account snapshot replication support matrix

Table 15-1 displays the cross-account replication support matrix for assets in the AWS cloud environment.

Table 15-1 AWS cross-account support matrix

Asset type	Cross-account replication: Unencrypted snapshots	Cross-account replication: Encrypted snapshots
Host-level snapshot (AMI)	Not supported	Not supported
EBS snapshot	Supported	Supported
RDS automated snapshot	Not supported	Not supported
RDS manual snapshot	Supported	Replicate to source region only
Aurora automated snapshot	Not supported	Not supported
Aurora manual snapshot	Replicate to source region only	Replicate to source region only

Cross-account snapshot replication limitations

Make a note of the following limitations:

- Cross-account replication support is not available for host-level snapshots.
- Cross-account replication support is not available for Amazon Machine Instances (AMI).
- The time required for the snapshot replication to complete is entirely dependent on the size and the target region's location. If the status of the replication job in the CloudPoint dashboard remains static, Veritas recommends that you verify the replication status in the target region using the AWS Management Console.
- AWS has default limits defined for AWS services, accounts, and regions. The number of simultaneous operations that you can perform on a region are therefore restricted by these default settings. If you trigger multiple snapshot replication operations on a region, those jobs may appear queued in the CloudPoint dashboard, if the number of jobs cross the predefined limits for that region.

You must also note that snapshot replication utilizes AWS resources and that may involve an additional cost. Ensure that you consider the total cost involved and plan snapshot replication tasks accordingly.

https://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html

- After a snapshot is replicated between AWS accounts, you cannot remove the source account configuration using the CloudPoint UI.

The only way to remove the source account configuration is by using the following methods:

- First, delete all the snapshots that are replicated to another AWS account region and then remove the source account configuration.
- Use the following CloudPoint REST API to remove the source account configuration:

```
/v2/agents/{agentId}/plugins/{pluginName}/configs/{configId}?force=True
```

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

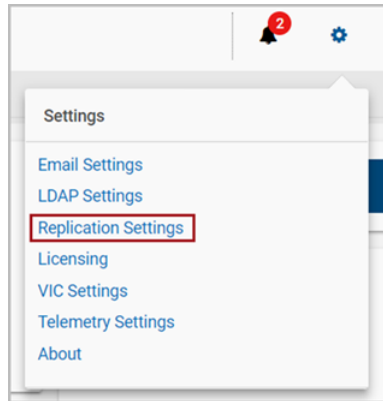
You can configure up to three replication destination for each source.

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 **Add New Replication Rule** page, specify the required parameters to configure a new rule.

Drop-down list	Description
Platform	Displays 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 platform type, such as Amazon Web Services (AWS)
- The destination regions to which the snapshots are replicated

- 3 Select the source location whose replication rules you want to edit and then click **Edit**.

- 4 The **Edit Replication Rule** page displays the source account ID and the account region.

Use the drop-down lists to change the replication locations or the order of the locations.


- 5 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](#)
- [Additional steps required after a SQL Server snapshot restore](#)
- [Additional steps required after an Oracle snapshot restore](#)
- [Additional steps required after a MongoDB snapshot restore](#)
- [Additional steps required after restoring an AWS RDS database instance](#)
- [Restoring individual files within 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 16-1 Assets and supported snapshot types

Asset	Supported snapshot types
Dell EMC Unity array	Copy-on-write (COW) snapshots on LUNs
HPE storage arrays	<p>COW and clone snapshot types</p> <p>Note the following:</p> <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.
Pure Storage FlashArray	Clone snapshots of volumes
NetApp storage arrays	COW snapshots of LUNs (SAN deployment) or NetApp NFS shares (NAS deployment)


Before you proceed, keep in mind the following:

- 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.
- 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.

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**

Asset Name	Asset Type	Vendor	Snap Count	Snapshotable
Aurora - us-west-2/rdsclus-harbhajanrai-606175228-06061246	application	amazon	0	Yes
Aurora - us-west-2/rdsclus-sayalpawar-606133038-06060828	application	amazon	0	No
RDS mysql Instance - us-west-2/rdsinst-divtas-eng-sdio-flexsnap-dev-60606060	application	amazon	1	Yes
RDS mysql Instance - us-west-2/rdsinst-harbhajanrai-606153425-06061020	application	amazon	1	Yes
RDS mysql Instance - us-west-2/rdsinst-harbhajanrai-606153425-06061020	application	amazon	0	No
RDS mysql Instance - us-west-2/rdsinst-harbhajanrai-606175228-06061241	application	amazon	1	Yes
RDS mysql Instance - us-west-2/rdsinst-nilehsawant-05301818	application	amazon	8	Yes

Showing 1 of 7 (1 selected)

Asset Details

RDS mysql Instance - us-west-2/rdsinst-nilehsawant-05301818

Vendor
amazon

Region
us-west-2

Snapshotable
Yes

ID
aws-rds-us-west-2-165323042987-rdsinst-nilehsawant-05301818

Policies (0)

[View Snapshots \(8\)](#)

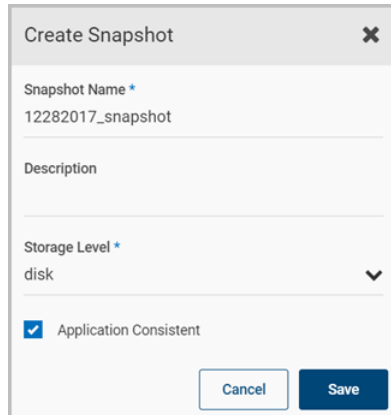
[Create Snapshot](#)

[Policies](#)

- 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.

A screenshot of a 'Create Snapshot' dialog box. The dialog has a title bar with a close button (X). Inside, there are three input fields: 'Snapshot Name' with the value '12282017_snapshot', 'Description' (empty), and 'Storage Level' with a dropdown menu showing 'disk'. Below these is a checkbox labeled 'Application Consistent' which is checked. At the bottom right are 'Cancel' and 'Save' buttons.

Create Snapshot

Snapshot Name *
12282017_snapshot

Description

Storage Level *
disk

☒ Application Consistent

Cancel 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.



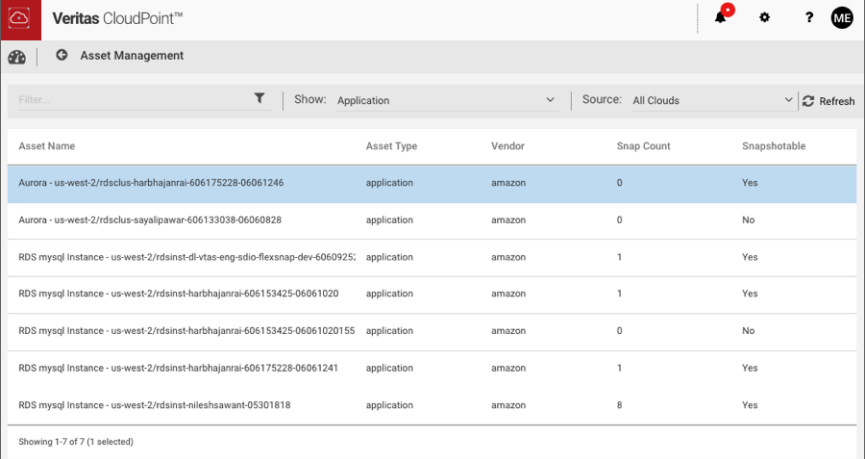
2 On the **Asset Management** page, select the application whose snapshots you want to view and then click **View Snapshots** from the Details pane on the right.

Asset Name	Asset Type	Vendor	Snap Count	Snapshotable
Aurora - us-west-2/rdscus-harbhajanrai-606175228-06061246	application	amazon	0	Yes
Aurora - us-west-2/rdscus-sayalpawar-606133038-06060828	application	amazon	0	No
RDS mysql Instance - us-west-2/rdscus-dl-vtas-eng-sdio-flexnap-dev-60606	application	amazon	1	Yes
RDS mysql Instance - us-west-2/rdscus-harbhajanrai-606153425-06061020	application	amazon	1	Yes
RDS mysql Instance - us-west-2/rdscus-harbhajanrai-606153425-06061020	application	amazon	0	No
RDS mysql Instance - us-west-2/rdscus-harbhajanrai-606175228-06061241	application	amazon	1	Yes
RDS mysql Instance - us-west-2/rdscus-nilehsawant-05301818	application	amazon	8	Yes

Showing 1 of 7 (1 selected)

Asset Details
RDS mysql Instance - us-west-2/rdscus-nilehsawant-05301818
Vendor: amazon
Region: us-west-2
Snapshotable: Yes
ID: aws-rds-us-west-2-165323042987-rdscus-nilehsawant-05301818
Policies (0)
[View Snapshots \(8\)](#)
[Create Snapshot](#)
[Policies](#)

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



The screenshot shows the Veritas CloudPoint™ Asset Management interface. At the top, there's a header with the Veritas logo and 'Veritas CloudPoint™'. Below it, a navigation bar shows 'Asset Management'. A filter bar includes a 'Filter...' dropdown, a 'Show: Application' dropdown, a 'Source: All Clouds' dropdown, and a 'Refresh' button. The main content is a table with columns: Asset Name, Asset Type, Vendor, Snap Count, and Snapshotable. The table lists several assets, including Aurora instances and RDS MySQL instances, all of which are 'application' type and 'amazon' vendor. The first row is highlighted in blue. At the bottom, it says 'Showing 1-7 of 7 (1 selected)'.

Asset Name	Asset Type	Vendor	Snap Count	Snapshotable
Aurora - us-west-2/rdsclus-harbhajanrai-606175228-06061246	application	amazon	0	Yes
Aurora - us-west-2/rdsclus-sayalipawar-606133038-06060828	application	amazon	0	No
RDS mysql Instance - us-west-2/rdsinst-dl-vtas-eng-sdio-flexsnap-dev-6060925	application	amazon	1	Yes
RDS mysql Instance - us-west-2/rdsinst-harbhajanrai-606153425-06061020	application	amazon	1	Yes
RDS mysql Instance - us-west-2/rdsinst-harbhajanrai-606153425-06061020155	application	amazon	0	No
RDS mysql Instance - us-west-2/rdsinst-harbhajanrai-606175228-06061241	application	amazon	1	Yes
RDS mysql Instance - us-west-2/rdsinst-nilehsaawant-05301818	application	amazon	8	Yes

Showing 1-7 of 7 (1 selected)

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

- Restore a snapshot
See [“Restoring a snapshot”](#) on page 201.
- Replicate a snapshot
See [“Replicating a snapshot manually”](#) on page 189.
- Classify a snapshot
- Delete a snapshot
See [“Deleting a snapshot”](#) on page 215.

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 175.

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**

A screenshot of the Veritas CloudPoint 'Asset Management' page. The main area contains a table with columns: Asset Name, Asset Type, Vendor, Snap Count, and Snapshotable. One row is highlighted in blue. To the right, the 'Asset Details' sidebar is open, showing information for the selected asset, including Vendor (amazon), Region (us-west-2), and a 'View Snapshots (2)' button.

Asset Name	Asset Type	Vendor	Snap Count	Snapshotable
EBS Volume us-west-2/vol-0008d868b28edfa6	disk	amazon	0	Yes
EBS Volume us-west-2/Reena_MSSQL_DND	disk	amazon	2	Yes
EBS Volume us-west-2/vol-000ae29500ad4fc9f	disk	amazon	0	Yes
EBS Volume us-west-2/vol-0013539709457e7e5	disk	amazon	0	Yes
EBS Volume us-west-2/vol-001b94c42b8442605	disk	amazon	0	Yes
EBS Volume us-west-2/vol-001c9be8402598fd1	disk	amazon	2	Yes
EBS Volume us-west-2/grijja_multidb_DND	disk	amazon	0	Yes
EBS Volume us-west-2/vol-002d80b46d9c099fe	disk	amazon	0	Yes
EBS Volume us-west-2/vol-0032781e79e292b13	disk	amazon	1	Yes
EBS Volume us-west-2/vol-003971008dc02da9e	disk	amazon	0	Yes

Showing 1-10 of 3356 (1 selected)

1 2 3 4 5

Asset Details

EBS Volume us-west-2/vol-001c9be8402598fd1

Vendor
amazon

Region
us-west-2

Snapshotable
Yes

ID
aws-efs-us-west-2-vol-001c9be8402598fd1

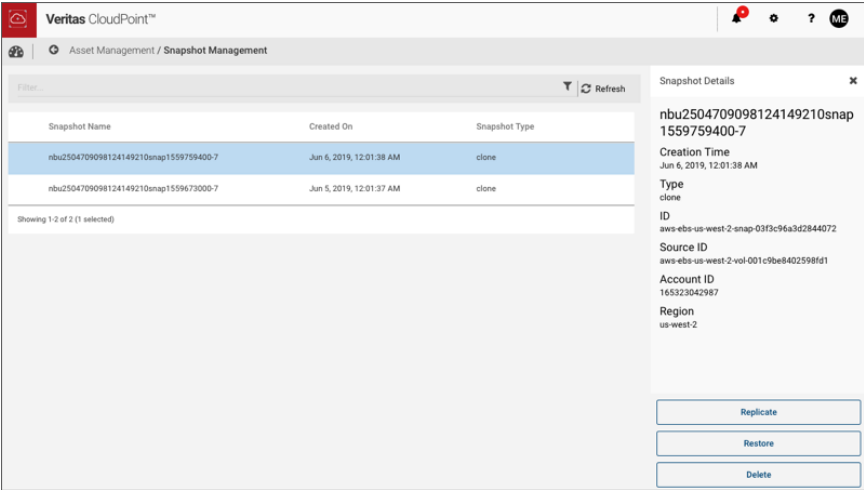
Policies (0)

[View Snapshots \(2\)](#)

[Create Snapshot](#)

[Policies](#)

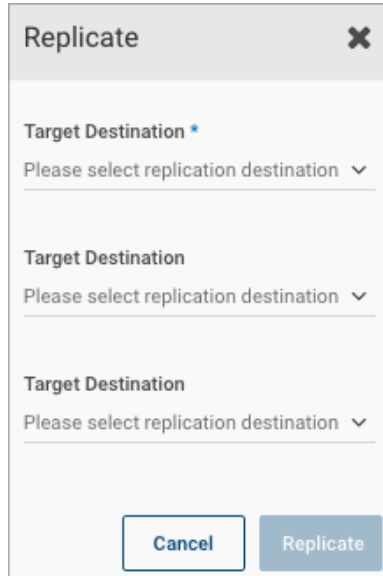
- 4
- On the **Snapshot Management** page, select the snapshot you want to replicate. You can only select one.



- 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**.

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

For cross-account replication, ensure that you select a different AWS account where you want to replicate the snapshot to.



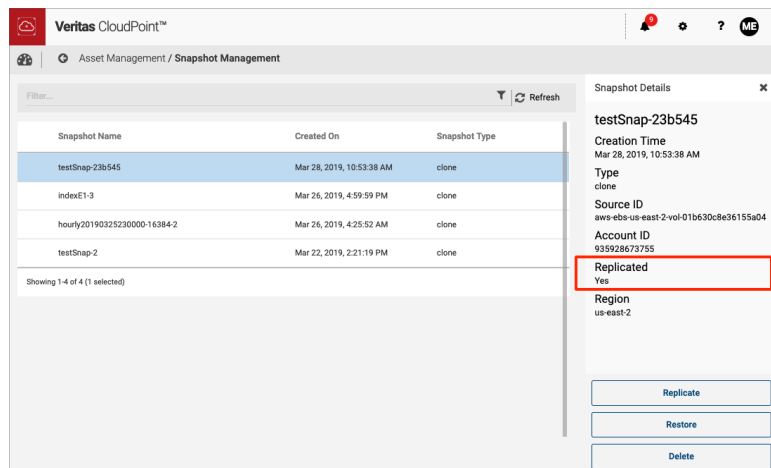
The image shows a 'Replicate' dialog box with a close button (X) in the top right corner. Inside the dialog, there are three identical sections, each labeled 'Target Destination' with a blue asterisk. Below each label is a dropdown menu with the text 'Please select replication destination' and a downward arrow. At the bottom of the dialog, there are two buttons: 'Cancel' and 'Replicate'.

7 Click **Replicate**.

8 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.



About snapshot restore

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

Table 16-2 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.

Table 16-2 Assets and supported restore options (*continued*)

Asset	Supported restore options
HPE storage arrays	<p>Restore a COW volume snapshot to the same volume with the Overwrite Existing option.</p> <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	<p>Restore a clone volume snapshot to the same volume with the Overwrite Existing option.</p>
NetApp storage arrays	<p>Restore the LUN snapshot to the same LUN (SAN deployment) or restore the NetApp NFS shares (NAS deployment).</p>

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 16-3 Default volume types

Region	Default volume type
us-east-1, eu-west-1, eu-central-1, us-west-1, us-west-2 ap-northeast-1, ap-northeast-2, ap-southeast-1, ap-southeast-2, ap-south-1 sa-east-1, us-gov-west-1, cn-north-1	standard
All other regions	gp2

- If you are performing a disk-level snapshot restore to the same location, then verify that the original disk is attached to the instance, before you trigger a restore.

If the existing original disk is detached from the instance, then the restore operation might fail.

See [“Disk-level snapshot restore fails if the original disk is detached from the instance”](#) on page 251.

- 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.
- If you intend to restore multiple file systems or databases on the same instance, then Veritas recommends that you perform these operations one after the other, in a sequential manner.
Running multiple restore operations in parallel can lead to an inconsistency at the instance level and the operations might fail eventually.
- If a region or zone is removed from the AWS or GCP plug-in configuration, then all the discovered assets from that region or zone are also removed from the CloudPoint assets database. If there are any active snapshots that are associated with the assets that get removed, then you may not be able perform any restore operations on those snapshots.
Once you add that zone back into the plug-in configuration, CloudPoint discovers all the assets again and you can resume the restore operations on the associated snapshots.

See [“Restore requirements and limitations for Microsoft SQL Server”](#) on page 196.

See [“Restore requirements and limitations for Oracle”](#) on page 197.

See [“Restore requirements and limitations for MongoDB”](#) on page 197.

Restore requirements and limitations for Microsoft SQL Server

Consider the following before you restore a SQL Server snapshot:

- Ensure that you close SQL Management Studio before you restore a SQL Server snapshot.
This is applicable only if you are restoring the snapshot to replace the current asset (Overwrite existing option) or restoring the snapshot to the same location as the original asset (Original Location option).
- Disk-level restore to a new location fails if the target host is connected or configured.
In such a case, to complete the SQL Server snapshot restore to a new location successfully, you must perform the restore in the following order:
 - First, perform a SQL Server disk-level snapshot restore.

Ensure that you restore the disk snapshots of all the disks that are used by SQL Server. These are the disks on which SQL Server data is stored.

See [“Restoring a snapshot”](#) on page 201.

- Then, after the disk-level restore is successful, perform the additional manual steps.

See [“Additional steps required after a SQL Server snapshot restore”](#) on page 204.

Restore requirements and limitations for Oracle

Consider the following before you restore an Oracle snapshot:

- 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 new 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.
- Disk-level restore to a new location fails if the CloudPoint plug-in for Oracle is not configured on the target host.

In such a case, to complete the Oracle snapshot restore to a new location successfully, you must perform the restore in the following order:

- First, perform a Oracle disk-level snapshot restore.
Ensure that you restore the disk snapshots of all the disks that are used by Oracle. These are the disks on which Oracle data is stored.
See [“Restoring a snapshot”](#) on page 201.
- Then, after the disk-level restore is successful, perform the additional manual steps.
See [“Additional steps required after an Oracle snapshot restore”](#) on page 208.

Restore requirements and limitations for MongoDB

Consider the following before you restore a MongoDB snapshot:

- Disk-level restore to a new location fails if the target host is connected or configured.

In such a case, to complete the MongoDB snapshot restore to a new location successfully, you must perform the restore in the following order:

- First, perform a MongoDB disk-level snapshot restore.

Ensure that you restore the disk snapshots of all the disks that are used by MongoDB. These are the disks on which MongoDB data is stored.

See [“Restoring a snapshot”](#) on page 201.

- Then, after the disk-level restore is successful, perform the additional manual steps.

See [“Additional steps required after a MongoDB snapshot restore”](#) on page 209.

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 198.

See [“Configuring classification settings using VIC”](#) on page 149.

See [“About indexing and classifying snapshots”](#) on page 147.

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 150.

- 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 13.
See [“About indexing and classifying snapshots”](#) on page 147.
- SFR is supported only for disk-level file system snapshots.
- 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 199.
See [“Single file restore limitations on Linux”](#) on page 199.
See [“Single file restore support on Windows”](#) on page 200.
See [“Single file restore limitations on Windows”](#) on page 200.

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 ext3, ext4, and XFS file systems.
- SFR is supported for multi-partition disks and whole disks.
- SFR is supported for Amazon Web Services (AWS), Google Cloud Platform (GCP), and Microsoft Azure cloud assets.

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**.

The screenshot shows the Veritas CloudPoint Asset Management interface. The main table lists assets with columns: Asset Name, Asset Type, Vendor, Snap Count, and Snapshotable. One asset is selected, and its details are shown on the right.

Asset Name	Asset Type	Vendor	Snap Count	Snapshotable
EBS Volume us-west-2/vol-008d868b28ed6a6	disk	amazon	0	Yes
EBS Volume us-west-2/reena_MSSQL_DND	disk	amazon	2	Yes
EBS Volume us-west-2/vol-009ae2960da4fc9f	disk	amazon	0	Yes
EBS Volume us-west-2/vol-0013539709457e7e5	disk	amazon	0	Yes
EBS Volume us-west-2/vol-001b94c42b8442605	disk	amazon	0	Yes
EBS Volume us-west-2/vol-001c9be8402598fd1	disk	amazon	2	Yes
EBS Volume us-west-2/gfija_multidb_DND	disk	amazon	0	Yes
EBS Volume us-west-2/002d80646d9c099fe	disk	amazon	0	Yes
EBS Volume us-west-2/vol-0032781e79e292b13	disk	amazon	1	Yes
EBS Volume us-west-2/vol-003971008dc02da9e	disk	amazon	0	Yes


Showing 1-10 of 3356 (1 selected)

Asset Details for EBS Volume us-west-2/vol-001c9be8402598fd1:

- Vendor: amazon
- Region: us-west-2
- Snapshotable: Yes
- ID: aws-efs-us-west-2-vol-001c9be8402598fd1
- Policies (0)

Buttons: View Snapshots (2), Create Snapshot, Policies

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



- 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	Restore option	Description
Cloud snapshot <ul style="list-style-type: none"> ■ Host / instance ■ Disk ■ Application (host, disk) 	Overwrite existing	<p>Replaces the current asset with the snapshot.</p> <p>Following is the behavior for this option:</p> <ul style="list-style-type: none"> ■ 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. <p>Notes:</p> <ul style="list-style-type: none"> ■ 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.
	Original Location	Restores the snapshot to the same location as the asset, without overwriting the existing asset.
	New location	<p>Restores the snapshot to a completely different location in the cloud.</p> <p>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.</p>
Array snapshot <ul style="list-style-type: none"> ■ Disk 	Overwrite existing	<p>Replaces the current asset with the snapshot.</p> <p>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.</p>

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 snapshot restore

The following steps are required after you restore a SQL Server snapshot from the CloudPoint user interface (UI). Even though the restore operation is successful, these steps are required for the application database to be available for normal use again.

The post-restore steps might vary depending on whether it is a host-level restore or a disk-level restore and whether restore is to the same location as the original database or to a new location. Ensure that you perform the steps that are relevant to your restore scenario.

Table 16-4 SQL post-restore steps

SQL snapshot restore scenario	Steps to perform post restore
Host-level restore (original or new location)	See “Steps required after a SQL Server host-level restore” on page 204.
Disk-level restore to new location	See “Steps required after a SQL Server disk-level snapshot restore to new location” on page 205.

Steps required after a SQL Server host-level restore

Perform these steps after you have restored a host-level SQL Server snapshot from the CloudPoint UI. These steps are required irrespective of whether you are restoring the snapshot to the original location or to a new location.

Before you proceed, verify the following:

- Ensure that the SQL Server user account on the Windows host where you intend to revert the shadow copy, has full access to the restore data.
- Ensure that the `pagefile.sys` is not present on the drive that is selected for the snapshot creation or snapshot restore.

The snapshot creation and snapshot restore operations will fail if the file is present on the selected drives.

Perform the following steps to revert the shadow copy

- 1 Connect to the Windows host where the SQL Server instance is running.
Ensure that you use an account that has administrator privileges on the host.
- 2 Stop the SQL Server service on the Windows host.
- 3 Open a command prompt window. If Windows UAC is enabled on the host, open the command prompt in the **Run as administrator** mode.
- 4 Navigate to
`%programdata%\Veritas\CloudPoint\tmp\tools\windows\tools\` directory,
 and then run the following command from there:


```
vss_snapshot.exe --revertSnapshot
```

The command displays a json output with Status = 0 that confirms that the operation is successful.

This command reverts the shadow copies for all the drives, except the system drive. The SQL Server service is stopped before the snapshot is reverted and automatically started after the revert operation is successful.
- 5 Start the SQL Server service on the Windows host.

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

Perform these steps after you have restored a disk-level SQL Server snapshot from the CloudPoint UI. These steps are required only if the snapshot is restored to a new location. New location refers to a new host that is different from the one where the SQL instance is running.

Clear the read-only mode of the new disk attached to the host

Perform the following steps

- 1 Connect to the new Windows host where the SQL Server instance is running.
Ensure that you use an account that has administrator privileges on the host.
- 2 Open a command prompt window. If Windows UAC is enabled on the host, open the command prompt in the **Run as administrator** mode.
- 3 Start the diskpart utility using the following command:

```
diskpart
```

- 4 View the list of disks on the new host using the following command:

```
list disk
```

Identify the new disk that is attached due to the snapshot restore operation and make a note of the disk number. You will use it in the next step.

- 5 Select the desired disk using the following command:

```
select disk <disknumber>
```

Here, <disknumber> represents the disk that you noted in the earlier step.

- 6 View the attributes of the selected disk using the following command:

```
attributes disk
```

The output displays a list of attributes for the disk. One of the attributes is `read-only`, which we will modify in the next step.

- 7 Modify the read-only attribute for the selected disk using the following command:

```
attributes disk clear readonly
```

This command changes the disk to read-write mode.

- 8 Bring the disk online.

From the Windows Server Manager console, navigate to **Files and Storage Devices > Disks** and then right click on the newly attached disk and select **Bring online**.

- 9 Quit the diskpart utility using the following command:

```
exit
```

Do not close the command prompt yet; you can use the same window to perform the remaining steps described in the next section.

Revert shadow copy using the Microsoft DiskShadow utility

Perform the following steps

- 1 From the same command window used earlier, start the diskshadow command interpreter in the interactive mode using the following command:

```
diskshadow
```

- 2 View the list of all the shadow copies that exist on the new host. Type the following command:

```
list shadows all
```

Identify the shadow copy that you want to use for the revert operation and make a note of the shadow copy ID. You will use the shadow ID in the next step.

- 3 Revert the volume to the desired shadow copy using the following command:

```
revert <shadowcopyID>
```

Here, <shadowcopyID> is the shadow copy ID that you noted in the earlier step.

- 4 Exit the DiskShadow utility using the following command:

```
exit
```

Attach .mdf and .ldf files to the instance database

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 and .ldf 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.

Additional steps required after an Oracle snapshot restore

The following steps are required after you restore an Oracle snapshot. Even though the restore operation itself is successful, these steps are required for the application database to be available for normal use again.

These manual steps are not required in case of a disk-level restore in the following scenario:

- You are performing the disk-level restore to a different host
- The target host is connected to the CloudPoint host
- The CloudPoint Oracle plug-in is configured on the target host

Perform the following steps:

- 1 Ensure that the snapshot restore operation has completed successfully and a new disk is created and mounted on the application host (in case of a disk-level restore) or the application host is up and running (in case of a host-level restore).
- 2 Connect to the virtual machine and then log on to the Oracle database as a database administrator (sysdba).
- 3 Start the Oracle database in mount mode using the following command:

```
# STARTUP MOUNT
```

Verify that the database is mounted successfully.

- 4 Remove the Oracle database from the backup mode using the following command:

```
# ALTER DATABASE END BACKUP
```

- 5 Open the Oracle database for normal usage using the following command:

```
# ALTER DATABASE OPEN
```


- 6 Add an entry of the newly created database in the Oracle `listener.ora` and `tnsnames.ora` files.

- 7 Restart the Oracle listener using the following command:

```
# lsnrctl start
```

Additional steps required after a MongoDB snapshot restore

The following steps are required after you restore a MongoDB snapshot. Even though the restore operation itself is successful, these steps are required for the application database to be available for normal use again.

Note: These manual steps are not required in case of a disk-level restore to the same location.

Perform the following steps

- 1 Ensure that the snapshot restore operation has completed successfully and a new disk is created and attached to the application host (in case of a disk-level restore) or the application host is up and running (in case of a host-level restore).

- 2 Connect to the application host.

- 3 Mount the attached disk on the application host using the following command:

```
# sudo mount /dev/<diskname> /<mountdir>
```

Here, `<diskname>` is the name of the new disk that was created after restore, and `<mountdir>` is the path where you want to mount the disk.

- 4 Edit the MongoDB config file `/etc/mongod.conf` and set the **dbPath** parameter value to the `<mountdir>` path that you specified in the earlier step.

- 5 Start the MongoDB service on the application host and verify that the service is running.

Use the following commands:

```
# sudo systemctl start mongod.service
# sudo systemctl status mongod.service
```

Note: In case of a disk-level restore to a new host, ensure that `mongo` is installed on that host.

- 6 Log on to the MongoDB server using the MongoDB client and verify that the database is running.

Additional steps required after restoring an AWS RDS database instance

The following steps are required after you restore an AWS RDS database instance snapshot. Even though the restore operation is successful, these manual steps are required so that the instance is available for normal use.

After restoring an AWS RDS database instance successfully, you have to manually check and reassign certain properties of the restored instance. This is required because even though the restore operation itself is successful, one or more instance properties are not restored completely. In some cases, CloudPoint resets the property values to their default settings.

The following RDS database instance or cluster properties are not restored completely and will need modification:

- **VPC security groups** value (*AWS Management Console > RDS Database instance > Connectivity & security tab*)
- **Deletion protection** setting (*AWS Management Console > RDS Database instance > Configuration tab*)
- **Copy tags to snapshots** setting (*AWS Management Console > RDS Database instance > Maintenance & backups tab*)

Perform the following steps:

- 1 Verify that the RDS database instance snapshot restore is successful.
- 2 Sign in to the AWS Management Console and from the top right corner, select the region in which you have restored the RDS instance.
- 3 From the Services menu, under Database, click **RDS**.

- 4 From the Dashboard menu on the left, click **Databases**.
- 5 In the Databases panel, select the restored RDS database instance and then click **Modify** from the menu bar on the top right.
- 6 On the Modify DB panel, check for the following properties and ensure that the attribute values match with those of the original instance:
 - Under Network & Security, verify that the **Security group** attribute has the correct security group name assigned.
 - Under Backup, verify that the **Copy tags to snapshots** option is set as per the original instance.
 - Under Deletion protection, verify that the **Enable deletion protection** option is set as per the original instance.
 - If required, verify all the other parameter values and set them as per your preference.
- 7 Once you have modified the desired RDS instance properties, click **Continue**.
- 8 Under Scheduling of modifications, choose an appropriate option depending on when you wish to apply the modifications to the instance and then click **Modify DB instance**.
- 9 Verify the RDS instance properties and ensure that the changes have taken effect.

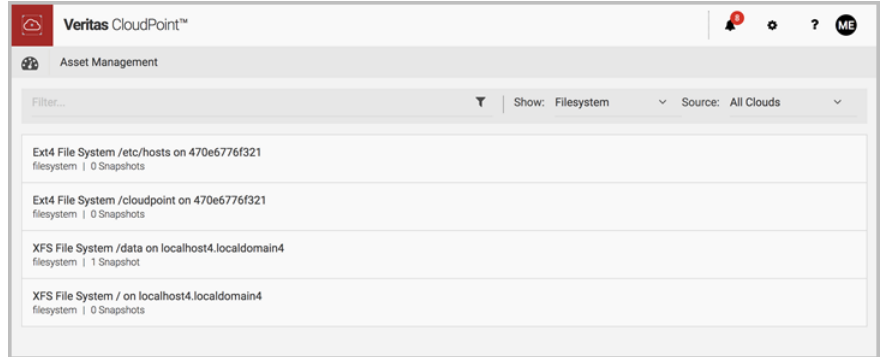
Restoring individual files within a snapshot

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 198.

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**.




- 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.

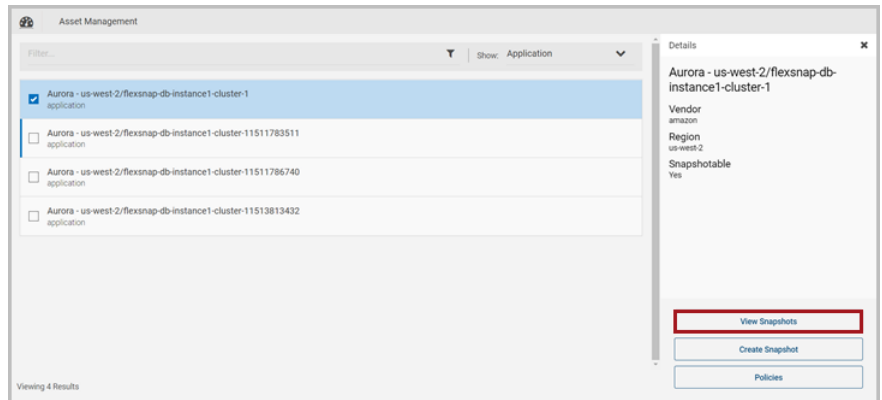
To delete a snapshot

- 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:

- [About CloudPoint notifications](#)
- [Viewing notifications in the CloudPoint UI](#)
- [CloudPoint notification methods](#)
- [CloudPoint notification limitations](#)
- [Configuring email-based CloudPoint notifications](#)
- [Configuring AWS SNS-based CloudPoint notifications](#)
- [Using the Job Log](#)

About CloudPoint notifications

CloudPoint notifications allow you to keep a track of all the critical events and failures that are happening in your CloudPoint environment. CloudPoint provides different kinds of notifications depending on the types of events that occur. For example, you will see a notification message whenever a snapshot or a replication job has failed, a restore operation has been completed successfully, and if a classification job has completed (failed or successful) for a snapshot.

Notifications are triggered for all operational tasks and events related to CloudPoint features. They include events related to the following:

- If a CloudPoint policy corresponding to the protection tag that is assigned to an asset does not exist.
- CloudPoint operations such as snapshot creation, replication, and restore have failed.
These operations could be triggered manually or using a policy.
- CloudPoint licensing changes, including license expiry, and license capacity usage.
- Indexing and classification of assets.
- CloudPoint assets discovery.

How CloudPoint notifications work

CloudPoint notifications are handled by the CloudPoint notification service that runs in a separate container named `flexsnap-notification`. This container is created when you install and configure CloudPoint. When an event occurs, the service or component that generates that event sends the details of that event to the notification service. The details include the context of the event along with the status, whether the event is a success or a failure type of event. The notification service then generates an appropriate notification based on the event information. The notification is then displayed on the Notifications panel in the CloudPoint user interface (UI), or sent to email recipients, or to AWS SNS topic subscribers, if configured.


See [“Viewing notifications in the CloudPoint UI”](#) on page 218.

See [“CloudPoint notification methods”](#) on page 222.

See [“Configuring AWS SNS-based CloudPoint notifications”](#) on page 224.

Viewing notifications in the CloudPoint UI

You can see CloudPoint notifications in the Notifications panel from the CloudPoint user interface (UI). To access the Notifications panel, click the bell icon in the top right corner of the CloudPoint UI.






On clicking the bell icon, a Notifications pop-up window displays a preview of a smaller list of the latest notification updates. Click **View all** to see the full list of notification messages.

Veritas CloudPoint™					
Notifications					
Status	Active	Severity	All	Category	All
Status	Severity	Message	Category	Date received	
Active	Error	Failed to restore newsnap-test for the asset named EBS Volume ap-south-1/Reena-attach-encrypt.	Task	Apr 10 2019, 07:37:35 PM	
Active	Error	Failed to restore myinstsnap9572 for the asset named EC2 instance ap-south-1/Reena_snap_host	Task	Apr 10 2019, 07:34:25 PM	
Active	Informational	Snapshot nrrefd for the asset named Ext4 File System /fstest on ip-172-31-19-103.us-west-2.compu	Task	Apr 10 2019, 07:31:38 PM	
Active	Error	Failed to create a snapshot of asset Ext4 File System /fstest on ip-172-31-19-103.us-west-2.compu	Task	Apr 10 2019, 06:33:08 PM	
Active	Informational	Successfully classified renmasnp (groupsnap-Scea8e9d0abdf252f5a698fc9079b69390371d6bdfd)	Classification	Apr 10 2019, 05:32:51 PM	
Active	Error	Unable to replicate the snapshot named reenardinst of asset RDS mysql Instance - us-east-2/reen	Task	Apr 10 2019, 05:30:19 PM	
Active	Critical	Received 3 times. One or more snapshot operations scheduled in policy rep-policy have failed.	Task	Apr 10 2019, 04:00:01 PM	

The Notifications panel displays a consolidated list of all the notification messages that have been generated in your CloudPoint environment. You can sort through the messages based on the following filter criteria:

- **Status:** Indicates whether the notification is in an active state or has been resolved.

- **Severity:** Indicates the criticality of the event for which the notification is generated. The severity is of the following types:

Severity type	Description
	<p>Informational</p> <p>Informational notification messages are generated for operational tasks that have been completed successfully.</p> <p>Example:</p> <pre>Snapshot <snapshotname> for the asset named <assetname> was restored successfully.</pre>
	<p>Warning</p> <p>Warning notification messages are generated for events such as when the license usage has reached near capacity or if a license is nearing its expiry and a renewal is required.</p> <p>While these messages do not call for an immediate action, failing to take remedial steps in the near term can potentially lead up to an operational stall.</p> <p>Example:</p> <pre>Enterprise license capacity usage has reached <util%>. Upgrade your license before it reaches the maximum entitled capacity.</pre>
	<p>Error</p> <p>Error notification messages are generated when a requested operation has failed. For example, if a manual snapshot or a replication job fails, an error is generated for the event. Errors are also generated when the license usage has exceeded 95% of the entitled capacity, or if the license has expired and the grace period has started.</p> <p>Example:</p> <pre>Unable to replicate snapshot named <snapshotName> of asset EBS volume <volumeName> to <regionName>.</pre>

Severity type **Description**



Critical

Critical notification messages are generated for failed scenarios that may indicate that CloudPoint operations may have stalled and your data protection plan may not be functioning as intended. Such notifications require immediate attention and swift remedial actions.

For example, when the license grace period has expired or the usage exceeds the total entitled capacity, a critical alert notification is generated. Similarly, if a policy-based snapshot or replication job fails, a critical alert is generated.

Example:

Received 3 times. One or more snapshot operations scheduled in policy <policyName> have failed.

Observe the text at the beginning of the sample message. If the same notification is generated multiple times, CloudPoint keeps a count of the occurrences and displays that number just before the actual message text in the Notifications panel.

In the Notifications pop-up window that appears when you click the bell icon, the count appears below the actual notification message text. Notification occurrence count is displayed for Critical and Error notification messages.

- **Category:** Indicates whether the notification is related to an operational task or an event related to a specific feature.

Task, Classification, Indexing, License, Connectivity, and Tag Based Protection are the various categories available.

- **Date range:** Indicates the time when the event has occurred. You can view messages that were generated in a specific time period such as the last 48 hours, the last 7 days, and so on. You can even specify a custom time frame.

CloudPoint displays all notification messages on the CloudPoint UI by default. Additionally, CloudPoint also supports sending notifications to email accounts or to AWS Simple Notification Service (SNS) topic. Once configured, you can receive notifications directly in those configured accounts, in addition to seeing them in the CloudPoint UI.

See [“CloudPoint notification methods”](#) on page 222.

CloudPoint notification methods

The Notifications panel in the CloudPoint UI displays real-time notifications for all the events in your CloudPoint environment. You can monitor the UI to keep a track of alerts that might need attention. However, this UI-based monitoring approach is more suited for a passive mode of monitoring, where a CloudPoint administrator performs random and infrequent checks on the notifications and the overall CloudPoint configuration status.

For a more active mode of monitoring, CloudPoint supports a push-based mechanism for notifications where, in addition to the CloudPoint UI, CloudPoint users can also receive these notifications using the following methods:

- **Email**
You can specify email addresses for receiving CloudPoint notifications in the form of emails.
- **AWS SNS**
You can specify an AWS Simple Notification Service (SNS) topic for receiving notifications. Users can subscribe to that SNS topic to receive CloudPoint notifications.

The most immediate advantage of using these methods is that you get notified of events that need urgent attention almost instantaneously, without having to be logged on to the CloudPoint UI. It allows you to take remedial actions and resolve the issue quickly. With AWS SNS, the CloudPoint notifications are delivered in a JavaScript Object Notation (JSON) format that can be used to set up an automated response mechanism to handle common failure scenarios. For example, if a hardware issue is causing a backup job failure, a support ticket can be raised with the hardware array vendor automatically.

Note that CloudPoint only sends creation and resolution alert notifications to the configured email recipients or SNS topic. No event notifications are sent in case there are any updates to an existing alert event.

See [“Configuring email-based CloudPoint notifications”](#) on page 223.

See [“Configuring AWS SNS-based CloudPoint notifications”](#) on page 224.

CloudPoint notification limitations

The following limitations apply to CloudPoint notifications:

- Notifications are generated for all types of supported events by default. You cannot choose specific events for generating notifications.
- You cannot manually set the severity for the notification alerts.

- The decision to send notification alerts to email recipients and SNS topic subscribers is made by the CloudPoint notification service. You cannot customize this behavior.
- By default, only alert creation and alert resolution event notifications are sent to email recipients and SNS topic subscribers. You cannot modify or customize this behavior.
- Only one SNS topic can be configured for SNS-based CloudPoint notifications. You cannot add additional SNS topics to the same configuration.

Configuring email-based CloudPoint notifications

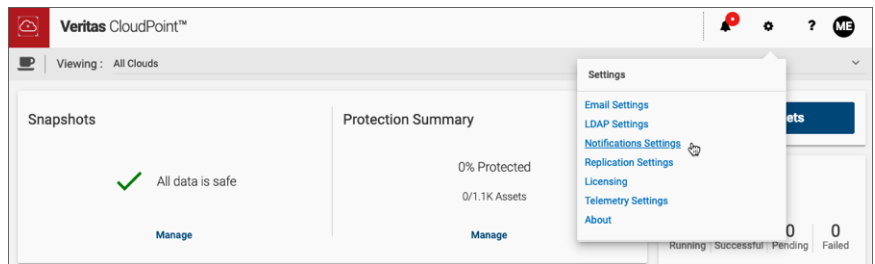
If you want to receive CloudPoint notifications as emails, add all the required recipient email addresses to the CloudPoint configuration.

Prerequisites

- Ensure that the Email Service is configured in CloudPoint.
See [“Configuring the CloudPoint sender email address”](#) on page 125.

To configure email-based notifications

- 1 Sign in to the CloudPoint UI with an administrator user account.
- 2 On the CloudPoint dashboard, click **Settings** (gear icon) from the top right corner and then click **Notification Settings** from the drop-down list.



- 3 Click to select the **Email** tab.

- 4 In the **Add email address to list of recipients** field, type a valid email address and then click **Add**.

Veritas CloudPoint™

Notifications Settings

How would you like to receive notifications?

Email Amazon Simple Notification Service (SNS)

Add email address to list of recipients

john.doe@mycompany.com

Add

Email Address	Type
No data to display	
Showing 0-0 of 0 (0 selected)	

A message on the UI confirms that the email address is added successfully.

- 5 Repeat step 4 for every recipient email address that you want to configure to receive CloudPoint notification emails.

When you add a valid email address, email-based notification mechanism is enabled automatically and CloudPoint begins sending notification emails to the configured email addresses.

The email-based notification mechanism remains active as long as there is at least one valid email address in the configuration. If you remove all the configured email addresses, the email-based notifications are disabled.

Configuring AWS SNS-based CloudPoint notifications

Add an AWS Simple Notification Service (SNS) topic to the CloudPoint configuration to send CloudPoint notifications to that SNS topic.

Prerequisites


- CloudPoint instance is deployed on an AWS EC2 instance in the AWS cloud.
- An AWS IAM role with the required minimum AWS SNS permissions is attached to the CloudPoint EC2 instance.

See [“AWS permissions required by CloudPoint”](#) on page 73.

- An AWS SNS topic is created for CloudPoint notifications and intended subscribers have subscribed to that topic.


To configure AWS SNS-based notifications

- 1 Sign in to the CloudPoint UI with an administrator user account.
- 2 On the CloudPoint dashboard, click **Settings** (gear icon) from the top right corner and then click **Notification Settings** from the drop-down list.



- 3 Click to select the **Amazon Simple Notification Service (SNS)** tab.

- 4
- Click **Add SNS Topic**.
- 5
- On the Add SNS dialog box, specify the required AWS SNS topic details and then click **Save**.



Parameter	Description
Topic Name	Specify the topic name of the AWS SNS topic that you have created for receiving CloudPoint notifications.
AWS Region	From the drop-down list, select the AWS region where the SNS topic you specified earlier is configured.

A message on the UI confirms that the SNS topic is added successfully.


When you add a valid SNS topic, SNS-based notification mechanism is enabled automatically and CloudPoint begins sending notifications to that SNS topic. To receive those notifications, users will have to subscribe to that SNS topic.

The SNS-based notification mechanism remains active as long as there is a valid SNS topic in the configuration. If you remove the configured SNS topic, the SNS-based notifications are disabled.

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 Jog Log

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

Veritas CloudPoint™			   	
Job Log				
Filter...			Status: All	Refresh
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- SCSI-600a098054315756693f492...	Apr 11, 2019, 4:26:28 PM		
✓ Delete Snapshot Job	delete snapshot (54315756693f492d47377649.Cp_scale62873a3d3) of disk- SCSI-600a098054315756693f492...	Apr 11, 2019, 4:26:28 PM		
✓ Delete Snapshot Job	delete snapshot (54315756693f492d47377648.Cp_scale62873a3d2) of disk- SCSI-600a098054315756693f492...	Apr 11, 2019, 4:26:28 PM		
✓ Delete Snapshot Job	delete snapshot (54315756693f492d47377646.Cp_scale62873a3d0) of disk- SCSI-600a098054315756693f492...	Apr 11, 2019, 4:26:28 PM		
✓ Delete Snapshot Job	delete snapshot (54315756693f492d4737764a.Cp_scale62873a3d4) of disk- SCSI-600a098054315756693f492...	Apr 11, 2019, 4:26:28 PM		
Showing 1-10 of 64 (0 selected)			1 2 3 4 5	

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:

Veritas CloudPoint™

Job Log

Filter...

Status: All

Refresh

Job Type	Job Name	Started o
Export Snapshot Job	export snapshot NetApp-sharesnap-94709da3-77a4-43d3-8342-a3cf4f4d8f99	Apr 11, 20
Deport Snapshot Job	deport snapshot NetApp-sharesnap-94709da3-77a4-43d3-8342-a3cf4f4d8f99	Apr 11, 20
Export Snapshot Job	export snapshot NetApp-sharesnap-94709da3-77a4-43d3-8342-a3cf4f4d8f99	Apr 11, 20
Export Snapshot Job	export snapshot NetApp-sharesnap-aae23852-39c1-4343-bace-0b0a89c42ce5	Apr 11, 20
Create Snapshot Job	create snapshot (snap_Task_1) of share TASKV0276049_001	Apr 11, 20
Delete Snapshot Job	delete snapshot (54315756693f492d47377647.Cp_scale62873a3d1) of disk-scsl-600a098054315756	Apr 11, 20
Delete Snapshot Job	delete snapshot (54315756693f492d47377649.Cp_scale62873a3d3) of disk-scsl-600a098054315756	Apr 11, 20
Delete Snapshot Job	delete snapshot (54315756693f492d47377648.Cp_scale62873a3d2) of disk-scsl-600a098054315756	Apr 11, 20
Delete Snapshot Job	delete snapshot (54315756693f492d47377646.Cp_scale62873a3d0) of disk-scsl-600a098054315756	Apr 11, 20
Delete Snapshot Job	delete snapshot (54315756693f492d4737764a.Cp_scale62873a3d4) of disk-scsl-600a098054315756	Apr 11, 20

Showing 1-10 of 64 (1 selected)

Job Details

Job Completed Successfully

Create Snapshot Job

create snapshot (snap_Task_1) of share TASKV0276049_001

Task ID

23e9908d-9569-49f9-ba15-039b5a09607d

Start time

Apr 11, 2019, 4:29:31 PM

End time

Apr 11, 2019, 4:29:38 PM

Duration

6 s

Summary

NetApp Snapshot snap_Task_1

Snapshot type

cow

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:

Veritas CloudPoint™

Job Log

Filter...

Status: Failed

Refresh

Job Type	Job Name	Started on
Restore Snapshot Job	restoring snapshot hourly.2019-04-08_1805	Apr 11, 2019,
Export Snapshot Job	export snapshot NetApp-sharesnap-37970904-47ce-42a2-92db-b23fb8f0be5d	Apr 11, 2019,
Deport Snapshot Job	deport snapshot NetApp-sharesnap-d9e6e373-3c3f-4dcd-817a-49c01bca34a6	Apr 11, 2019,
Create Snapshot Job	create snapshot (s1) of share nfsserver	Apr 9, 2019, 9

Showing 4-4 of 4 (1 selected)

Job Details

Job Completed with Errors

Create Snapshot Job

create snapshot (s1) of share nfsserver

Task ID

85dceefc-1193-4025-821e-7a01ac3f5ee9

Start time

Apr 9, 2019, 9:43:08 PM

End time

Apr 9, 2019, 9:43:13 PM

Duration

5 s

Error Summary

Fail to create snapshot of share[nfsserver]. 13020 Reason: The Snapshot(tm) copy name already exists

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.

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 your setup:

```
# curl -H "Content-Type: application/json" -H
"Authorization: Bearer $token" -X GET -k
https://localhost/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 32.
- 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 19. CloudPoint logging](#)
- [Chapter 20. Troubleshooting CloudPoint](#)
- [Chapter 21. Working with your CloudPoint license](#)
- [Chapter 22. Upgrading CloudPoint](#)
- [Chapter 23. Uninstalling CloudPoint](#)

CloudPoint logging

This chapter includes the following topics:

- [About CloudPoint logging mechanism](#)
- [How fluentd-based CloudPoint logging works](#)
- [Viewing CloudPoint logs](#)

About CloudPoint logging mechanism

Beginning with CloudPoint 2.2 release, CloudPoint uses the Fluentd-based logging framework for log data collection and consolidation. Fluentd is an open source data collector that provides a unified logging layer for structured log data collection and consumption.

Refer to the following for more details on Fluentd:

<https://www.fluentd.org/>

All the CloudPoint container services generate and publish service logs to the configured Docker logging driver. In CloudPoint release 2.2, the logging driver is the fluentd framework that is running as a separate `flexsnap-fluentd` container on the CloudPoint host. With the Fluentd framework, these individual service logs are now structured and routed to the Fluentd data collector from where they are sent to the configured output plugins. In CloudPoint release 2.2, MongoDB collection and the flexsnap-fluentd container logs are the two output plugins that are configured by default.

Using Fluentd-based logging provides several benefits including the following:

- A persistent structured repository that stores the logs of all the CloudPoint services
- A single stream of all CloudPoint logs (vs disparate individual log files) makes it easy to trail and monitor specific logs

- Metadata associated with the logs allow for a federated search that speeds up troubleshooting
- Ability to integrate and push CloudPoint logs to a third-party tool for analytics and automation
- Ability to view individual service logs using the mongodb helper utility (`flexsnap-log`)

How fluentd-based CloudPoint logging works

When you install CloudPoint, or upgrade to release 2.2, the following changes occur on the CloudPoint host:

- A new container service named `flexsnap-fluentd` is started on the CloudPoint host. This service is started before all the other CloudPoint container services. The `flexsnap-fluentd` service serves as the `fluentd` daemon on the host.
- All the CloudPoint container services are then started with `fluentd` as the Docker logging driver.
- A `fluentd` configuration file is created at `/cloudpoint/fluent/fluent.conf`. This file contains the output plugin definitions that are used to determine where the CloudPoint logs are redirected for consumption.

Once all the infrastructure components are ready, each of the CloudPoint services begin to send their respective log messages to the configured Docker `fluentd` logging driver. The `fluentd` daemon then redirects the structured logs to the output plugins configured in the `fluentd` configuration file. These logs are sent to the CloudPoint MongoDB collection as per the default settings in the `fluent.conf` configuration file. All the log messages are stored in a JSON format.

About the CloudPoint fluentd configuration file

Fluentd uses a configuration file that defines the source of the log messages, the set of rules and filters to use for selecting the logs, and the target destinations for delivering those log messages.

The `fluentd` daemon running on the CloudPoint host is responsible for sending the CloudPoint logs to various destinations. These target destinations, along with the other details such as input data sources and required fluentd parameters are defined in the plugin configuration file. For CloudPoint, these plugin configurations are stored in a `fluentd` configuration file that is located at `/cloudpoint/fluent/fluent.conf` on the CloudPoint host. The `fluentd` daemon reads the output plugin definition from this configuration file to determine where to send the CloudPoint log messages.

The following output plugin definitions are added to the configuration file by default:

- CloudPoint MongoDB collection (`Fluentd::logs`)

This represents the default MongoDB collection that stores all the CloudPoint service logs on the CloudPoint host.

The plugin is defined as follows:

```
# Send to a mongodb collection fluentd:logs
<store>
@type mongo
host flexsnap-mongodb
port 27017
database fluentd
collection logs

ssl true
ssl_cert /cloudpoint/keys/mongodb.pem
ssl_key /cloudpoint/keys/mongodb.pem
ssl_ca_cert /cloudpoint/keys/cacert.pem

capped
capped_size 5120m
</store>
```

- STDOUT

This is used to send the CloudPoint log messages to the `flexsnap-fluentd` service container logs. These logs can be obtained using standard Docker commands.

The plugin is defined as follows:

```
# Send to fluentd docker logs
<store>
@type stdout
</store>
```

Additionally, the CloudPoint fluentd configuration file includes plugin definitions for the following destinations:

- Splunk
- ElasticSearch

These plugin definitions are provided as a template and are commented out in the file. To configure an actual Splunk or ElasticSearch target, you can uncomment these definitions and replace the parameter values as required.

Modifying the fluentd configuration file

Modify the `fluentd.conf` configuration file if you want to modify the existing plugin definitions.

To modify the fluentd.conf file

- 1 On the CloudPoint host, open the `/cloudpoint/fluent/fluent.conf` configuration file in a text editor of your choice and then edit the contents to add or remove a plugin definition.
- 2 If desired, increase the maximum size for the MongoDB collection for CloudPoint logs.

The default size is set to 5120m, which is 5120 MB or 5 GB.

- 3 Save all the changes to the file.
- 4 Restart the `flexsnap-fluentd` container service using the following command:

```
# sudo docker restart flexsnap-fluentd
```

Note that the changes take effect immediately and are applicable only to the newer log messages that get generated after the change. The file changes do not apply to the older logs that were generated before the configuration file was updated.

Fluentd-based logging requirements and considerations

- If you are attempting a real time analysis of the logs, then you might see a noticeable delay when using the CloudPoint plugin for the MongoDB collection. This happens because the plugin performs a periodic data flush in to the MongoDB database. The default flush rate is set to 10 seconds and is defined in the `/cloudpoint/fluent/fluent.conf` configuration file on the CloudPoint host.

An alternative approach is to use the `STDOUT` plugin for such requirements. The logs appear as the logs of the `flexsnap-fluentd` container and can be obtained using Docker commands.

- To use the `flexsnap-log` utility, ensure that Python is installed on the CloudPoint host.

Viewing CloudPoint logs

CloudPoint provides a MongoDB client helper utility (`flexsnap-log`) that is located within the `flexsnap-coordinator` service. This utility allows you to access the MongoDB logs collection.

The general command syntax for using the `flexsnap-log` utility is as follows:

```
# sudo docker exec flexsnap-coordinator flexsnap-log <options>
```

Table 19-1 Flexsnap-log command options

Command option	Description
<service>	The CloudPoint service name. The command displays the logs of the specified service.
- h --help	Displays the command syntax and a description of the available options.
-n <N> --limit <N>	Displays the last "N" number of log messages. For example, to view the last 50 log messages, specify the following: -n 50
-t --tail	Use this option to follow and monitor the log messages in real time.
-F <format> --format <format>	Displays the log messages in the specified output format. For example, -F {container_name}: {log}.
-v --verbose	Displays the command output in a verbose mode.
-j --json	Displays the logs in a JavaScript Object Notation (JSON) format.
-d <days> --days <days>	Displays the logs for the last "DAYS" number of days. For example, to view the logs for the last seven days, specify the following: -d 7
-f <filename> --file <filename>	Dumps the logs to the file specified in <filename>.

You can view the CloudPoint logs using any of the following commands on the CloudPoint host:

- To obtain all the CloudPoint service logs, run the following command:

```
# sudo docker exec flexsnap-coordinator flexsnap-log
```

- To obtain logs of a specific CloudPoint container service, run the following command:

```
# sudo docker exec flexsnap-coordinator flexsnap-log  
<flexsnap-service name>
```
- To tail or follow log messages, run the following command:

```
# sudo docker exec flexsnap-coordinator flexsnap-log -t
```
- To obtain the last "N" number of log messages, run the following command:

```
# sudo docker exec flexsnap-coordinator flexsnap-log -n <N>
```
- You can also combine these options to achieve a specific output. For example, to obtain the last 10 log messages for the `flexsnap-agent` service, run the following command:

```
# sudo docker exec flexsnap-coordinator flexsnap-log -n 10  
flexsnap-agent
```

The command output displays messages similar to the following:

```
flexsnap-agent: flexsnap-agent-offhost[1] flexsnap.updates: INFO - find_files:netapp.zip  
flexsnap-agent: flexsnap-agent-offhost[1] flexsnap.updates: INFO - find_files:nutanix.zip  
flexsnap-agent: flexsnap-agent-offhost[1] flexsnap.updates: INFO - find_files:oracle.zip  
flexsnap-agent: flexsnap-agent-offhost[1] flexsnap.updates: INFO - find_files:purestg.zip  
flexsnap-agent: flexsnap-agent-offhost[1] flexsnap.updates: INFO - find_files:windows.zip  
flexsnap-agent: flexsnap-agent-offhost[1] INFO - Beginning registration with coordinator  
flexsnap-agent: flexsnap-agent-offhost[1] INFO - loaded plugin, sending configId status: {}  
flexsnap-agent: flexsnap-agent-offhost[1] INFO - Sending list of sources  
flexsnap-agent: flexsnap-agent-offhost[1] INFO Registration complete
```

The most recent CloudPoint logs are also available in the `flexsnap-fluentd` container logs. You can use standard Docker commands to obtain the logs.

Run the following command:

```
# sudo docker logs flexsnap-fluentd | grep flexsnap-agent | head -10
```

The command output displays messages similar to the following:

```
flexsnap-agent: {"container_name":"flexsnap-agent","source":"stdout","log":  
"Mar 04 09:10:20 f5dlae1c4808 flexsnap-agent-offhost[1] MainThread agent:  
INFO - Not generating certificate. Join token not passed for role agent"}  
  
flexsnap-agent: {"container_name":"flexsnap-agent","source":"stdout","log":  
"Mar 04 09:10:20 f5dlae1c4808 flexsnap-agent-offhost[1] MainThread  
flexsnap.ca: INFO - Loading /opt/VRTScloudpoint/keys/agent.6c5c9.cert.pem  
/opt/VRTScloudpoint/keys/cacert.pem"}
```

```
flexsnap-agent: {"container_name":"flexsnap-agent","source":"stdout","log":  
"Mar 04 09:10:20 f5d1ae1c4808 flexsnap-agent-offhost[1] MainThread  
flexsnap.connectors.rabbitmq: INFO - Starting service"}
```

To view the flexsnap-fluentd container logs in real time, run the following command:

```
# sudo docker logs flexsnap-fluentd -f | grep <flexsnap-service-name>
```

Troubleshooting CloudPoint

This chapter includes the following topics:

- [Restarting CloudPoint](#)
- [Docker may fail to start due to a lack of space](#)
- [CloudPoint installation fails if rootfs is not mounted in a shared mode](#)
- [Some CloudPoint features do not appear in the user interface](#)
- [Off-host plug-in deletion does not automatically remove file system and application assets](#)
- [Disk-level snapshot restore fails if the original disk is detached from the instance](#)
- [Snapshot restore for encrypted AWS assets may fail](#)
- [Error while adding users to CloudPoint](#)
- [CloudPoint fails to revert restored snapshots if indexing, classification, or restore operations fail](#)
- [SQL snapshot or restore and SFR operations fail if the Windows instance loses connectivity with the CloudPoint host](#)
- [Troubleshooting CloudPoint logging](#)

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 mongodb[9]
Oct 03 11:24:45 init:INFO - Started mongodb[9]
Oct 03 11:24:45 init:INFO - mongodb already stopped, 100
Oct 03 11:24:45 init:INFO - mongodb 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 book disk.
- 4 After the boot disk has more than 4.0 GB of available space, restart the container.

```
# sudo docker restart container-id
```

CloudPoint installation fails if rootfs is not mounted in a shared mode

CloudPoint installation fails if the mount state of the root file system (`rootfs` or `shared subtree`) on a host is not set to `"shared"` mode. One or more CloudPoint containers may fail to start due to insufficient privileges on the host.

Workaround:

Use the following command to install CloudPoint on such a host:

```
# sudo docker run --rm -it --name CloudPoint -e SKIP_SHARED_MNT=Yes
-v /<full_path_to_volume>:/<full_path_to_volume> -v
/var/run/docker.sock:/var/run/docker.sock
veritas/flexsnap-cloudpoint:<version> install
```


Here, `<full_path_to_volume>` is the path to the CloudPoint data volume, which typically is `/cloudpoint`.

`<version>` represents the CloudPoint product version.

The installer displays messages similar to the following:

```
Some features of CloudPoint like indexing, classification and SFR
will be disabled.
```

```
If this is not desired then change the mount propagation and
start installation again.
```

```
Do you wish to continue CloudPoint installation
with limited features? (y/n):
```

Enter **y** to proceed with the installation.

For detailed information on CloudPoint installation, refer to the following:

See [“Installing CloudPoint”](#) on page 32.

Note the following:

Even though CloudPoint supports installation on a root file system that is not mounted in a shared mode, you will not be able to use CloudPoint features such as Indexing, Classification, and Single File Restore (SFR) in such deployments. These features do not work in such an environment.

To make full use of CloudPoint and all its features, Veritas recommends that you deploy CloudPoint on a host where the mount mode of the root file system is set to `shared`.

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 13.
- 4 If your license does not support the feature you want, consider upgrading your license.
See [“Upgrading your CloudPoint license”](#) on page 258.

Off-host plug-in deletion does not automatically remove file system and application assets

When you configure an off-host plug-in, CloudPoint immediately starts discovering all the assets that are associated with that plug-in. For example, in case of the AWS plug-in, CloudPoint discovers all the EC2 instances and the corresponding EBS volumes that are being used.

However, for CloudPoint to be able to discover the file systems and applications on those instances, you are also required to configure the CloudPoint on-host plug-in or the agentless feature. From the CloudPoint UI, you have to connect to each of the instances and then click Configure.

When you remove the off-host plug-in, CloudPoint automatically removes the discovered instances and storage assets that are associated with that plug-in. For example, in case of AWS plug-in, an EC2 instance and the corresponding EBS volumes are removed from the CloudPoint configuration.

However, the off-host plug-in deletion does not automatically remove the file system and application assets that belong to the on-host or agentless host instances corresponding to the off-host plug-in. Those assets are discovered by the on-host plug-ins or the agentless feature and remain in the CloudPoint configuration even after the off-host plug-in is deleted.

On the Asset Management page in the CloudPoint user interface (UI), the **Snapshotable** column for such assets displays as **No**. You may not be able to perform any snapshot operations (for example, create snapshot, view snapshot) on such file system and application assets.

Workaround:

After deleting the off-host plug-in, you also have to manually remove the on-host agents and plug-ins and the agentless feature to cleanup all the associated assets from the CloudPoint configuration.

You can remove agents and plug-ins using the following CloudPoint APIs:

Table 20-1 CloudPoint APIs: Delete agents and plug-ins

Objective	HTTP Method	API URI
Delete an agent	DELETE	/v2/agents/{agentId}
Delete a plug-in from an agent	DELETE	/v2/agents/{agentId}/plugins/{pluginName}
Delete a configuration entry for a plug-in	DELETE	/v2/agents/{agentId}/plugins/{pluginName}/configs/{configId}

Disk-level snapshot restore fails if the original disk is detached from the instance

This issue occurs if you are performing a disk-level snapshot restore to the same location.

When you trigger a disk-level snapshot restore to the same location, CloudPoint first detaches the existing original disk from the instance, creates a new volume from the disk snapshot, and then attaches the new volume to the instance. The original disk is automatically deleted after the restore operation is successful.

However, if the original disk whose snapshot is being restored is manually detached from the instance before the restore is triggered, the restore operation fails.

You may see the following message on the CloudPoint UI:

```
Request failed unexpectedly: [Errno 17] File exists: '/<app.diskmount>'
```

The CloudPoint coordinator logs contain messages similar to the following:

```
flexsnap.coordinator: INFO - configid : <app.snapshotID> status changed to
{u'status': u'failed', u'discovered_time': <time>, u'errmsg': u'
Could not connect to <application> server localhost:27017:
[Errno 111]Connection refused'}
```

Workaround:

If the restore has already failed in the environment, you may have to manually perform a disk cleanup first and then trigger the restore job again.

Perform the following steps:

- 1** Log on to the instance for which the restore operation has failed.

Ensure that the user account that you use to connect has administrative privileges on the instance.

- 2** Run the following command to unmount the application disk cleanly:

```
# sudo umount /<application_diskmount>
```

Here, *<application_diskmount>* is the original application disk mount path on the instance.

If you see a "device is busy" message, wait for some time and then try the `umount` command again.

- 3** From the CloudPoint UI, trigger the disk-level restore operation again.

In general, if you want to detach the original application disks from the instance, use the following process for restore:

1. First take a disk-level snapshot of the instance.
2. After the snapshot is created successfully, manually detach the disk from the instance.

For example, if the instance is in the AWS cloud, use the AWS Management Console and edit the instance to detach the data disk. Ensure that you save the changes to the instance.

3. Log on to the instance using an administrative user account and then run the following command:

```
# sudo umount /<application_diskmount>
```

If you see a "device is busy" message, wait for some time and then try the `umount` command again.

4. Now trigger a disk-level restore operation from the CloudPoint UI.

Snapshot restore for encrypted AWS assets may fail

If the AWS assets that you wish to protect using CloudPoint are encrypted using AWS KMS Customer Managed Keys (CMK), then the CloudPoint restore operations may fail due to access permission errors. The AWS instances that are spawned may go in to a terminated state.

You may see the following error messages in the AWS CloudTrail logs:

```
"errorCode": "AccessDenied", with the following error message:
User: arn:aws:sts::<ID>:assumed-role/<iamrole>/<awsinstance> is not
authorized to perform: kms:ReEncryptFrom on resource: <resourcename>"
```

```
Waiter InstanceRunning failed: Waiter encountered a terminal failure state
```

This error occurs because CloudPoint is unable to perform encrypt and decrypt operations on the AWS assets. This happens because the IAM user or role that is provided to CloudPoint does not have the requisite permissions to use the CMKs.

Workaround:

To resolve this issue, do the following:

- If using an IAM user for CloudPoint plug-in configuration, ensure that the IAM user is added as a key user of the CMK.
- For source account configuration, ensure that the IAM role that is attached to the CloudPoint instance is added as a key user of the CMK.
- For cross account configuration, ensure that the IAM role that is assigned to the other AWS account (cross account) is added as a key user of the CMK.

Adding these IAM roles and users as the CMK key users allows them to use the AWS KMS CMK key directly for cryptographic operations on the assets. Refer to the AWS documentation for more details:

<https://docs.aws.amazon.com/kms/latest/developerguide/key-policies.html#key-policy-default-allow-users>

Error while adding users to CloudPoint

You might not be able to add users to the CloudPoint configuration. The following email verification error is displayed on the UI:

```
User not added. Email address is not verified.
```

If using the Amazon SES email service, the following error is displayed on the UI:

```
User not added. Email address is not verified.
The following identities failed the check in region <regionname>
(Service: AmazoneSimpleEmailService; Status Code: 400;
Error Code: Message Rejected; Request ID: <ID>)
```

Third-party email services such as Amazon SES and SendGrid require that you verify the email address accounts before you start using them. When you try to add users, CloudPoint reports this error if the specified user email address is not verified.

Workaround:

Before you add users to CloudPoint, ensure that you do the following:

- If using email services such as Amazon SES or SendGrid, then before adding email addresses to CloudPoint (either a CloudPoint sender email address, or a CloudPoint user email address), ensure that you add those email addresses to the email service accounts. This ensures that the email addresses are verified.
- In case of Amazon SES, you might also have to move the SES account out of the sandbox environment to be able to use that email address for sending emails to non-verified user email addresses.

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

CloudPoint fails to revert restored snapshots if indexing, classification, or restore operations fail

When you trigger CloudPoint operations such as indexing, classification, and restore (full snapshot restore, single file restore, or application restore), CloudPoint first attaches the snapshot to the CloudPoint host (in case of indexing and classification) or to the target host (in case of snapshot restore) and then performs the operation. The snapshot is detached from the host once the operation is completed successfully.

Sometimes, these CloudPoint operations might fail due to an internal workflow error, an unknown exception, or an unlikely error scenario. In such failure cases, it is observed that CloudPoint is unable to detach the snapshot that was exported on the host for performing the operation. The snapshot remains attached to the host, even if the operation itself has failed. As a result, you may not be able to trigger any subsequent indexing, classification, or restore operations. Even if you are able to initiate these operations, they remain in a queued state and may eventually fail.

Workaround:

If you encounter such an issue, you may have to manually detach the snapshot from the host. This will ensure that subsequent indexing, classification, or restore jobs do not fail.

Manual steps required in case of an indexing or classification failure on the CloudPoint host

Perform the following steps on the CloudPoint host:

- Check if the restored snapshot file system is mounted on the CloudPoint host. If it is mounted, unmount the file system from the host.
- Check if the restored snapshot disk is attached to the CloudPoint host. If it is attached, detach the volume from the host.

SQL snapshot or restore and SFR operations fail if the Windows instance loses connectivity with the CloudPoint host

- Check if the restored snapshot volume is visible in the cloud. If it is available, delete that volume using the cloud management console.

Manual steps required in case of an SFR failure on Linux hosts

Perform the following steps on the Linux host:

- Check if the restored snapshot file system is mounted on the Linux host. If it is mounted, unmount the file system from the host.
- Check if the restored snapshot disk is attached to the Linux host. If it is attached, detach the volume from the host using the cloud management console.
- Check if the restored snapshot volume is visible in the cloud. If it is available, delete that volume using the cloud management console.

Manual steps required in case of SFR failure on Windows hosts

Perform the following steps on the Windows host:

- Check if the restored snapshot volume is mounted on the Windows host. If it is mounted, unmount the volume from the host.
- Check if the restored snapshot disk is attached to the Windows host. If it is attached, take the disk offline using the diskpart command line utility or from the Windows Computer Management UI.
- Verify that the restored snapshot disk is offline on the Windows host, and then detach the volume from the host.
- Check if the restored snapshot volume is visible in the cloud. If it is available, delete that volume using the cloud management console.

SQL snapshot or restore and SFR operations fail if the Windows instance loses connectivity with the CloudPoint host

This issue occurs if the CloudPoint on-host agent that is configured on a Windows instance loses network connectivity with the CloudPoint host. CloudPoint operations such as snapshot creation or restore for SQL Server and single file restore (SFR) begin to fail for the Windows instance.

The connectivity failure may occur due to various reasons such as a services restart on the CloudPoint host as part of a CloudPoint software upgrade or a general network disruption.

The flexsnap-agent logs may contain messages similar to the following:

```
flexsnap-agent-onhost[2720] MainThread flexsnap.connectors.rabbitmq:
ERROR - Unexpected exception() in main loop
flexsnap-agent-onhost[2720] MainThread agent: ERROR - Agent failed
unexpectedly
```

If CloudPoint is deployed in a Veritas NetBackup environment, the NetBackup logs may contain messages similar to the following:

```
Error nbcs (pid=5997) Failed to create snapshot for asset: <sqlassetname>
Error nbcs (pid=5997) Operation failed. Agent is unavailable.
```

Workaround:

To resolve this issue, restart the Veritas CloudPoint Agent service on the Windows instance.

Troubleshooting CloudPoint logging

If the flexsnap-coordinator service is down

If the flexsnap-coordinator service itself is down, you cannot use the flexsnap-log utility to obtain the logs. In such a case, you can retrieve the logs of a CloudPoint service from the logs of the flexsnap-fluentd container itself.

Use the following command:

```
# sudo docker logs flexsnap-fluentd -f | grep <flexsnap-service name>
```

If the CloudPoint installation fails with a fluentd logging driver error

If the CloudPoint installation fails with an error that the CloudPoint container services are unable to find the fluentd logging driver, do the following:

- Verify that the flexsnap-fluentd service is running.
Run the following command on the CloudPoint host:

```
# sudo docker ps
```

The command output should display the fluentd container as running.
- If the fluentd service has not started on the CloudPoint host, look at the errors in the flexsnap-fluentd service logs.

Use the following command:

```
# sudo docker logs flexsnap-fluentd
```

In general, if the flexsnap-log utility fails to display log messages, you can retrieve the logs from the flexsnap-fluentd container log using Docker commands.

Working with your CloudPoint license

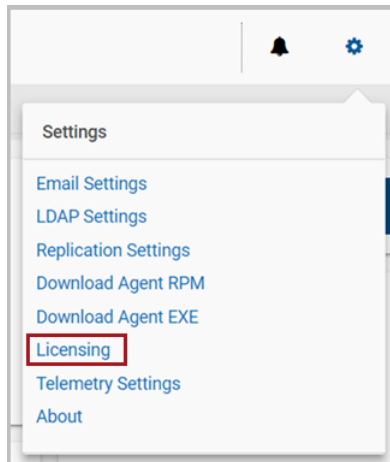
This chapter includes the following topics:

- [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 13.

See [“Upgrading your CloudPoint license”](#) on page 258.

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 13.

To upgrade your CloudPoint license

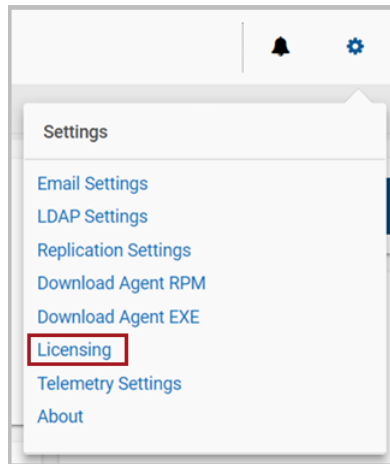
- 1 Use the download link that is provided by your Veritas representative to download the license file to your local machine. If necessary, copy the license file to the machine from where you will access the CloudPoint user interface.

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 142.

- 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

Used 0 Bytes

Licensed 1 TB

Cloud Edition

Used 0 Bytes

Licensed 1 TB

On-premise Edition

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 13.

See [“Displaying CloudPoint license and protection information”](#) on page 257.

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 22-1 CloudPoint upgrade path

Upgrade from version	Upgrade to version
<ul style="list-style-type: none">■ 2.0.2■ 2.1■ 2.1.2	2.2

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 232.
- Ensure that you remove CloudPoint plug-ins that are no longer supported by the newer CloudPoint release.
See [“Removing CloudPoint plug-in configuration”](#) on page 263.
- If you want to install or upgrade CloudPoint in an unattended mode, ensure that you use the appropriate command syntax.
See [“About deploying CloudPoint in a non-interactive mode”](#) on page 31.

Removing CloudPoint plug-in configuration

Before you upgrade CloudPoint, you must remove the CloudPoint plug-ins that are no longer supported by the newer CloudPoint release. If you upgrade CloudPoint without removing the deprecated plug-ins, then CloudPoint may hang and you may not be able to log on to the CloudPoint UI after the upgrade.

For example, CloudPoint 2.2 release no longer supports the HPE 3PAR and HDS plug-ins. If you have configured these plug-ins in your CloudPoint environment, you must remove them before you upgrade to release 2.2.

To remove a CloudPoint plug-in


- 1 Ensure that there are no active snapshots of the assets that are managed by the CloudPoint plug-in. You cannot remove a CloudPoint plug-in if there are snapshots.

For example, if you want to remove the HPE 3PAR plug-in, ensure that you delete all the CloudPoint snapshots from the HPE 3PAR array.

See [“Deleting a snapshot”](#) on page 215.

- 2 From the CloudPoint dashboard, in the **Administration** card, locate **Clouds/Arrays** and then click **Manage**.

- 3 On the Clouds and Arrays page, select the CloudPoint plug-in that you wish to remove and then click **Manage** from the Plugin Details panel that appears on the right side.
- 4 On the plug-in configurations page, select the plug-in configuration and then click **Remove Configuration** from the Configuration Details panel that appears on the right side.



- 5 Respond to the dialog that asks you to confirm whether you want to remove the plug-in configuration.
The UI displays a message that the plug-in has been removed successfully.
- 6 Repeat these steps for each deprecated CloudPoint plug-in that you wish to remove.
- 7 The next step is to upgrade CloudPoint.
See [“Upgrading CloudPoint”](#) on page 264.

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:

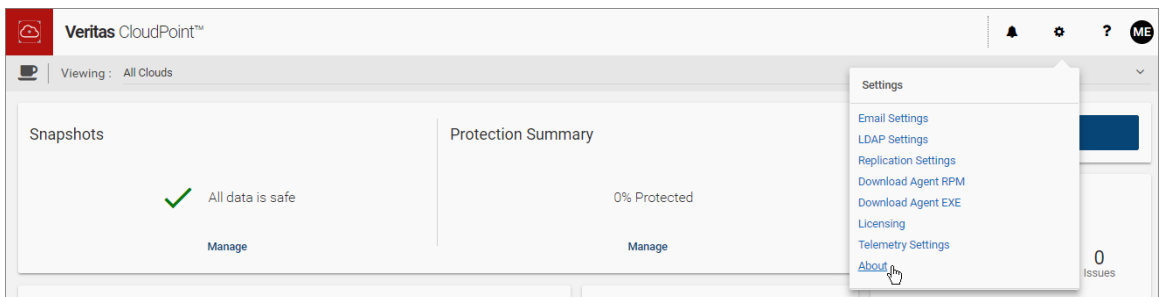
```
644879075e24: Loading layer [=====>] 117.9MB/117.9MB
d7ff1dc646ba: Loading layer [=====>] 15.87MB/15.87MB
d73dd9qwer58: Loading layer [=====>] 1.812GB/1.812GB
3167ba895aec: Loading layer [=====>] 352.9MB/352.9MB
fd22ad285778: Loading layer [=====>] 41.98kB/41.98kB
Loaded image: veritas/flexsnap-cloudpoint:2.1.2.7542
```

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
Stopping container: flexsnap-agent.e425d969dd4 ...done
Stopping container: flexsnap-agent.4704fd318322 ...done
Stopping container: flexsnap-fluentd ...done
Stopping container: flexsnap-mongodb ...done
Stopping container: flexsnap-rabbitmq ...done
Stopping container: flexsnap-authorization-service ...done
Stopping container: flexsnap-auth ...done
Stopping container: flexsnap-api-gateway ...done
Stopping container: flexsnap-api ...done
Stopping container: flexsnap-coordinator ...done
Stopping container: flexsnap-agent ...done
Stopping container: flexsnap-scheduler ...done
Stopping container: flexsnap-vic ...done
Stopping container: flexsnap-indexingsupervisor ...done
Stopping container: flexsnap-telemetry ...done
Stopping container: flexsnap-licensing ...done
Stopping container: flexsnap-policy ...done
Stopping container: flexsnap-cloudpointconsole ...done
Stopping container: flexsnap-notification ...done
```

```
Stopping container: flexsnap-identity-manager-service ...done
Stopping container: flexsnap-email-service ...done
Stopping container: flexsnap-onhostagent ...done
```

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
```

For an unattended installation, use 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 -y
```

Here, *new_version* represents the CloudPoint version you are upgrading to.

The `-y` option passes an approval for all the subsequent installation prompts and allows the installer to proceed in a non-interactive mode.

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.

Note: No inputs are required if the installer is run in a non-interactive mode.

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:

```
Installing the services
Configuration started at time: Thu Jul 11 09:58:02 UTC 2019
Do you wish to continue CloudPoint installation with older
docker version. ? (y/n): y
This is an upgrade to CloudPoint 2.1.2.7542
Previous CloudPoint version: 2.0.2.4722
Checking if a 1.0 release container exists ...
Removing exited container flexsnap-agent.e425d969d ...done
Removing exited container flexsnap-agent.47033896e ...done
Removing exited container flexsnap-cloudpointconsole ...done
Removing exited container flexsnap-authorization-service ...done
Removing exited container flexsnap-email-service ...done
Removing exited container flexsnap-identity-manager-service ...done
Removing exited container flexsnap-licensing ...done
Removing exited container flexsnap-vic ...done
Removing exited container flexsnap-telemetry ...done
Removing exited container flexsnap-notification ...done
Removing exited container flexsnap-indexingsupervisor ...done
Removing exited container flexsnap-policy ...done
Removing exited container flexsnap-scheduler ...done
Removing exited container flexsnap-onhostagent ...done
Removing exited container flexsnap-agent ...done
Removing exited container flexsnap-coordinator ...done
Removing exited container flexsnap-api ...done
Removing exited container flexsnap-api-gateway ...done
Removing exited container flexsnap-auth ...done
Removing exited container flexsnap-rabbitmq ...done
Removing exited container flexsnap-mongodb ...done
Removing exited container flexsnap-fluentd ...done
Deleting network : flexsnap-network ...done
Loading images for the CloudPoint services ...done
```


```
Creating network: flexsnap-network ...done
Generating certificates for MongoDB server ...done
Starting docker container: flexsnap-fluentd ...done
Starting docker container: flexsnap-mongodb ...done
Starting docker container: flexsnap-rabbitmq ...done
Generating certificates for API-gateway container ...done
Generating certificates for few other service container ...done
Generating certificates for OnhostAgent container ...done
Adding MongoDB certificate to the trust store ...done
Starting docker container: flexsnap-auth ...done
Starting docker container: flexsnap-api-gateway ...done
Starting docker container: flexsnap-api ...done
Starting docker container: flexsnap-coordinator ...done
Starting docker container: flexsnap-agent ...done
Starting docker container: flexsnap-onhostagent ...done
Starting docker container: flexsnap-scheduler ...done
Starting docker container: flexsnap-policy ...done
Starting docker container: flexsnap-indexingsupervisor ...done
Starting docker container: flexsnap-notification ...done
Starting docker container: flexsnap-telemetry ...done
Starting docker container: flexsnap-vic ...done
Starting docker container: flexsnap-licensing ...done
Starting docker container: flexsnap-identity-manager-service ...done
Starting docker container: flexsnap-email-service ...done
Starting docker container: flexsnap-authorization-service ...done
Starting docker container: flexsnap-cloudpointconsole ...done
```

- 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.
Run the agent package file and follow the installation wizard workflow to upgrade the on-host agent on the Windows host. The installer detects the existing installation and upgrades the package to the new version automatically.
- 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 111.

- 15** If you have deployed CloudPoint on an EC2 instance in the AWS cloud, you now have the option to configure CloudPoint to use AWS KMS service for encrypting and decrypting CloudPoint configuration.

This is an optional step and is applicable only if your CloudPoint deployment is using the default encryption mechanism and AWS KMS is not already configured in that environment.

See [“About CloudPoint integration with AWS KMS”](#) on page 49.

See [“Configuring AWS KMS in CloudPoint”](#) on page 43.

- 16** If CloudPoint is deployed in the AWS cloud and is integrated into a Veritas NetBackup environment, the next step is to update the NetBackup configuration so that the upgraded CloudPoint configuration details are available with NetBackup.

Performing this step ensures that the AWS IAM configuration settings in CloudPoint are updated in the NetBackup configuration.

A NetBackup configuration update (this step) is not required if CloudPoint is not deployed in the AWS cloud.

Run the following command on the NetBackup Master server:

```
# ./tpconfig -update -cloudpoint_server <cloudpoint_server_name>  
-cloudpoint_server_user_id <user_ID> [-requiredport  
<IP_port_number>]
```

Here,

- `<cloudpoint_server_name>` is the fully qualified domain name (FQDN) of the CloudPoint host.
- `<cloudpoint_server_user>` is the CloudPoint administrator user account that is configured on the CloudPoint server.
- `<IP_port_number>` represents the port number used by the CloudPoint server. The default port number is 443. This parameter is required only if CloudPoint host is using a different port.

For example, if the CloudPoint host name is `mycphost.mydomain.dom` and the configured CloudPoint admin user is `mycpuser@mycp.com`, then the command syntax is as follows:

```
# ./tpconfig -update -cloudpoint_server mycphost.mydomain.com  
-cloudpoint_server_user_id mycpuser@mycp.com
```

When prompted, enter the password for the CloudPoint user that you specified in the command earlier, and then enter the password again to confirm.

Once the NetBackup configuration is updated, you can start using CloudPoint with NetBackup immediately. All the existing CloudPoint configuration settings such as plug-ins, assets, snapshots, restore, and replication jobs are retained and continue to work as is. New assets discovery is performed as per the discovery schedule. If you add a new AWS region, AWS IAM is used to discover and perform operations on the assets in that region.

For more details about the `tpconfig` command and its options, refer to the *Veritas NetBackup Commands Reference Guide*. For more information about CloudPoint and Veritas NetBackup integration, refer to the *Veritas NetBackup Web UI Cloud Administrator's Guide*.

https://www.veritas.com/support/en_US/article.100040135

Note: If you have upgraded to CloudPoint 2.2 without removing deprecated plug-in configurations such as HPE 3PAR and HDS, then you may not be able to sign in to the CloudPoint UI after the upgrade. Contact Veritas Technical Support to help you clean the deprecated plug-in entries from the CloudPoint database and get your CloudPoint deployment up and running.

Uninstalling CloudPoint

This chapter includes the following topics:

- [Preparing to uninstall CloudPoint](#)
- [Removing the CloudPoint on-host agents](#)
- [Removing CloudPoint from a standalone Docker host environment](#)

Preparing to uninstall CloudPoint

Note the following before you uninstall CloudPoint:

- Ensure that there are no active CloudPoint operations in progress. If there are any snapshot, replication, or restore jobs running, wait for them to complete. If you have configured policies, ensure that you stop the scheduled policy runs. You may even want to delete those policies.
- All the snapshot data and configuration data from your existing installation is maintained in the external `/cloudpoint` data volume. This information is external to the CloudPoint containers and images and is preserved after the uninstallation. However, you can take a backup of all the data in the `/cloudpoint` volume, if desired.
See [“Backing up CloudPoint”](#) on page 232.

Removing the CloudPoint on-host agents

You must first remove the CloudPoint on-host agents before you remove CloudPoint. On-host agents are installed directly on the host where the applications are running. CloudPoint on-host agents manage on-host plug-ins that are used to discover assets and perform snapshot operations on the host.

See [“About agents”](#) on page 104.

To uninstall the CloudPoint on-host agents

- 1 Connect to the host where you have installed the CloudPoint on-host agent.

Ensure that the user account that you use to connect has administrative privileges on the host.

- 2 For Linux-based on-host agent, do the following:

Remove the .rpm package using the following command:

```
# sudo yum -y remove <cloudpoint_agent_package>
```

Here, *<cloudpoint_agent_package>* is the name of the on-host agent rpm package, without the version number and the file extension (.rpm).

For example, if the name of the on-host agent rpm package is *VRTScloudpoint-agent-2.2-RHEL7.x86_64.rpm*, the command syntax is as follows:

```
# sudo yum -y remove VRTScloudpoint-agent
```

- 3 For Windows-based on-host agent, do the following:

From Windows Control Panel > Programs and Features, select the entry for the CloudPoint on-host agent (**Veritas CloudPoint Agent**) and then click **Uninstall**.

Follow the wizard workflow to uninstall the on-host agent from the Windows instance.

- 4 This completes the on-host agent uninstallation.

Possible next steps are to either re-install the agents or uninstall CloudPoint.

See [“Downloading and installing the on-host agent”](#) on page 111.

See [“Removing CloudPoint from a standalone Docker host environment”](#) on page 275.

Removing CloudPoint from a standalone Docker host environment

The process for uninstalling CloudPoint is the same as that followed for installation. The only difference is that you specify `"uninstall"` in the command, which tells the installer to remove the components from the host.

During uninstallation, the installer performs the following tasks on the CloudPoint host:

- Stops all the CloudPoint containers that are running

- Removes the CloudPoint containers
- Unloads and removes the CloudPoint images

To uninstall CloudPoint

1. Ensure that you have uninstalled the CloudPoint on-host agents from all the hosts that are part of the CloudPoint configuration.

See [“Removing the CloudPoint on-host agents”](#) on page 274.

2. Sign in to the CloudPoint user interface (UI) and from the Job Log page, verify that there are no protection policy snapshots or other operations in progress, and then stop CloudPoint.

Run the following command on the CloudPoint host:

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

Here, *<version>* represents the currently installed CloudPoint version.

<full_path_to_volume> represents the path to the CloudPoint data volume, which typically is */cloudpoint*.

For example, if the installed CloudPoint version is 2.0.2.4722, the command syntax is 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.

3. Uninstall CloudPoint by running the following command on the host:

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

If the CloudPoint host is behind a proxy server, use the following command instead:

```
# sudo docker run -it --rm
-v /<full_path_to_volume_name>:/<full_path_to_volume_name>
-e VX_HTTP_PROXY=<http_proxy_value>
-e VX_HTTPS_PROXY=<https_proxy_value>
-e VX_NO_PROXY=<no_proxy_value>
-v /var/run/docker.sock:/var/run/docker.sock
veritas/flexsnap-cloudpoint:<version> uninstall
```

Replace the following parameters as per your environment:

Parameter	Description
<version>	Represents the CloudPoint product version that is installed on the host.
<full_path_to_volume>	Represents the path to the CloudPoint data volume, which typically is <code>/cloudpoint</code> .
<http_proxy_value> (required only if the instance uses a proxy server)	Represents the value to be used as the HTTP proxy for all connections. For example, "http://proxy.mycompany.com:8080/".

Parameter	Description
<i><https_proxy_value></i> (required only if the instance uses a proxy server)	Represents the value to be used as the HTTPS proxy for all connections. For example, "https://proxy.mycompany.com:8080/".
<i><no_proxy_value></i> (required only if the instance uses a proxy server)	Represents the hosts that are allowed to bypass the proxy server. Use commas to separate multiple host names. For example, "localhost,mycompany.com,192.168.0.10:80".

For example, if the product 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 uninstall
```

If using a proxy server, then using the examples provided in the table earlier, the command syntax is as follows:

```
# sudo docker run -it --rm -v /cloudpoint:/cloudpoint -e
VX_HTTP_PROXY="http://proxy.mycompany.com:8080/" -e
VX_HTTPS_PROXY="https://proxy.mycompany.com:8080/" -e
VX_NO_PROXY="localhost,mycompany.com,192.168.0.10:80" -v
/var/run/docker.sock:/var/run/docker.sock
veritas/flexsnap-cloudpoint:2.0.2.5300 uninstall
```

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

The installer begins to unload the relevant CloudPoint container packages from the host. Messages similar to the following indicate the progress status:

```
Uninstalling Veritas CloudPoint
-----
Stopping flexsnap-mongodb ... done
Stopping flexsnap-rabbitmq ... done
Stopping flexsnap-auth ... done
Stopping flexsnap-coordinator ... done
Removing flexsnap-mongodb ... done
Removing flexsnap-rabbitmq ... done
```

```
Removing flexsnap-auth ... done
Removing flexsnap-coordinator ... done
Unloading flexsnap-mongodb ... done
Unloading flexsnap-rabbitmq ... done
Unloading flexsnap-auth ... done
Unloading flexsnap-coordinator ... done
```

4. Confirm that the CloudPoint containers are removed.

Use the following docker command:

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

5. If desired, remove the CloudPoint container images from the host.

Use the following docker command to view the docker images that are loaded on the host:

```
# sudo docker images -a
```

Use the following docker command to remove the CloudPoint container images from the host:

```
# sudo docker rmi <image ID>
```

6. This completes the CloudPoint uninstallation on the host.

Possible next step is to re-deploy CloudPoint.

See [“Installing CloudPoint”](#) on page 32.

Reference

- [Chapter 24. Storage array support](#)
- [Chapter 25. Working with CloudPoint using APIs](#)

Storage array support

This chapter includes the following topics:

- [Dell EMC Unity arrays](#)
- [Pure Storage FlashArray](#)

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 24-1 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 24-2 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	storops 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 287.

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 24-3 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 24-4 Supported Pure Storage FlashArray models

Category	Supported
Array model	FA-405
Firmware version	<ul style="list-style-type: none">■ Software: Purity OS■ Purity OS version: 5.1.4■ Rest Version: 1.11 <p>Refer to the array-specific documentation for more information on firmware versions and how to check the current firmware on your array.</p>

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 287.

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.

