Veritas NetBackup™ CloudPoint Install and Upgrade Guide

Ubuntu, RHEL, SLES

Release 10.0

February 2022

VERITAS™

Veritas NetBackup™ CloudPoint Install and Upgrade Guide

Last updated: 2022-02-28

Legal Notice

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 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:

NB.docs@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

Section 1

CloudPoint installation and configuration

Chapter 1

Preparing for CloudPoint installation

This chapter includes the following topics:

About the deployment approach

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 1-1 CloudPoint's micro-services model

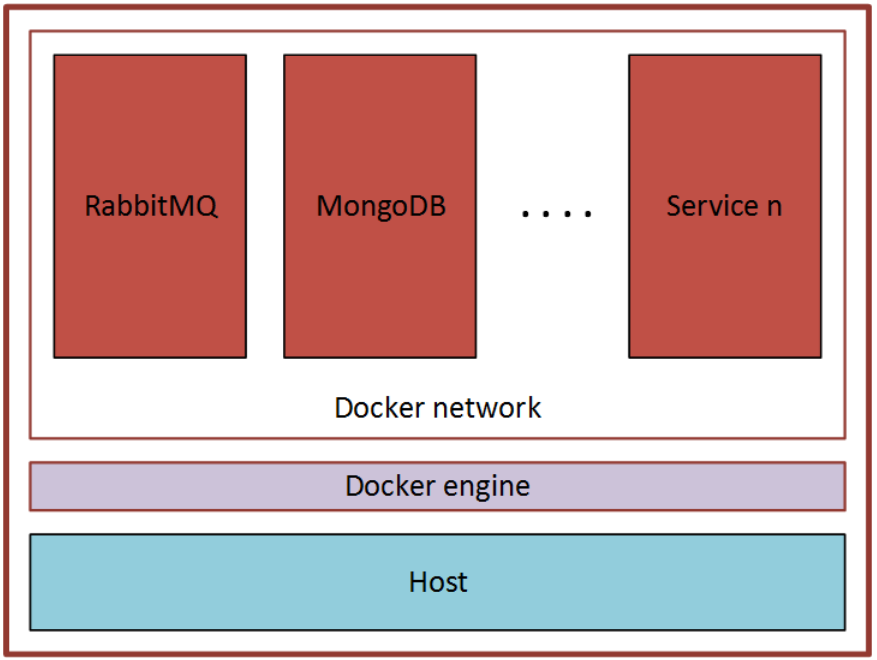

This deployment approach has the following advantages:

- 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 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 on cloud to protect your cloud assets. 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.

You can deploy CloudPoint in a NetBackup media server, but not in a NetBackup primary server.

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.

If you host the CloudPoint server and media server in the same host, do the following for proper functioning of the backup from snapshot jobs:

- Assign distinct IPs and NBU client names to the CloudPoint server and the media serve so that they can obtain different NetBackup Certificates. This is required so as have different NetBackup host ID certificates for communication. Use the following configuration:
  - Configure host with two network adapters
  - Edit the `/etc/hosts` file and make entry as mentioned in the example below:

    ```
    <IP Address MediaServer Host1>  < MediaServer Host1>
    <IP Address CloudPoint Host2>  <CloudPoint  Host2>
    ```

  - Provide the MediaServer Host1 which is mentioned in the`/etc/hosts` file during the Media server installation for Media server name.
  - Similarly select the CloudPoint Host 2 from the `/etc/hosts` file during the CloudPoint installation with non-default port other than 443.
  - Start CloudPoint and Media services and register it with NetBackup primary server.

- Once the CloudPoint server is registered, ensure that it has a different HOST DB entry.
- Before performing the backup from snapshot jobs, perform the following optimization: DISABLE SHM and NOSHM. See: https://www.veritas.com/support/en_US/article.100016170

This will ensure that NetBackup does not use shared memory for communicating between NetBackup data mover processes.

About deploying CloudPoint in the 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 cloud. You can deploy CloudPoint either manually or using the CloudPoint template available in the online marketplace.

In case of manual CloudPoint deployment, ensure the UUID of CloudPoint server boot disk is unique and does not conflict with FS UUID of any other asset node.

Refer to the following for more information on how to deploy a CloudPoint instance in the cloud:

http://veritas.com/netbackupcloud

Meeting system requirements

CloudPoint host requirements

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

See “CloudPoint host sizing recommendations” on page 22.

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

Category	Requirement
Operating system	■ Ubuntu 18.04 and 20.04 Server LTS ■ Red Hat Enterprise Linux (RHEL) 8.5, 8.4 and 7.x **Note:** CloudPoint deployment for RHEL 8.5 and 8.4 over IPV6 is not supported. ■ SUSE Linux Enterprise Server (SLES) 15 SP2
Processor architecture	x86_64 /64-bit processors

Table 1-2 System requirements for the CloudPoint host

Host on which CloudPoint is installed	Requirements
Amazon Web Services (AWS) instance	■ Elastic Compute Cloud (EC2) instance type: t3.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	■ 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. Ensure that do the following before you deploy CloudPoint on an RHEL instance in the Azure cloud: ■ Register the RHEL instance with Red Hat using Red Hat Subscription Manager ■ Extend the default LVM partitions on the RHEL instance so that they fulfil the minimum disk space requirement

Microsoft Azure Stack Hub VM	■ Virtual machine types:
	  ■ DS2_v2 Standard - CPU cores 2, RAM 7 GB
	  ■ DS3_v2 Standard - CPU cores 4, RAM 14 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.
	
	Ensure that do the following before you deploy CloudPoint on an RHEL instance in the Azure Stack Hub cloud:
	
	■ Register the RHEL instance with Red Hat using Red Hat Subscription Manager
	■ Extend the default LVM partitions on the RHEL instance so that they fulfil the minimum disk space requirement

Table 1-2 System requirements for the CloudPoint host *(continued)*

Host on which CloudPoint is installed	Requirements
Google Cloud Platform (GCP) VM	■ Virtual machine type: n2-standard-4 ■ vCPUs: 2 ■ RAM: 16 GB ■ Boot disk: 64 GB standard persistent disk ■ Data volume: 50 GB SSD persistent disk for the snapshot asset database with automatic encryption **Note:** To support LVM indexing, ensure that the Multipath service is disabled on CloudPoint host.
VMware VM	■ Virtual machine type: 64-bit with a CloudPoint supported operating system ■ vCPUs: 8 ■ RAM: 16 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)*	■ 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: 16 GB or more ■ Boot disk: 64 GB ■ Data volume: 50 GB for the snapshot asset database

Note: NetBackup CloudPoint is not fully FIPS compliant.

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 or upgrade CloudPoint has sufficient space for the following components.

Table 1-3 Space considerations for CloudPoint components

Component	Space requirements
CloudPoint containers	30 GB free space
CloudPoint agents and plug-ins	350 MB free space, for every CloudPoint plug-in and agent configured

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 1-4 Space consideration for CloudPoint data volume

Volume mount path	Size
`/cloudpoint`	50 GB or more

See "CloudPoint host sizing recommendations" on page 22.

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 plug-ins (earlier known as off-host plug-ins), or using the CloudPoint application configuration plug-ins (earlier known as on-host plug-ins), or using the CloudPoint agentless feature.

Table 1-5 Supported applications, operating systems, cloud, and storage platforms

Category	Support
Applications	■ File systems ■ Linux native file systems: ext3, ext4, and XFS ■ Microsoft Windows: NTFS ■ Microsoft SQL 2014, SQL 2016, SQL 2017, SQL 2019 See “Microsoft SQL plug-in configuration notes” on page 165. ■ Windows Server 2022 and 2019 ■ Oracle 12c, Oracle 12c R1, Oracle 18c, Oracle 19c Single node configurations are supported. See “Oracle plug-in configuration notes” on page 166. **Notes:** ■ 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.
Operating systems on supported assets	■ Red Hat Enterprise Linux (RHEL) 7.x Red Hat Enterprise Linux (RHEL) 8.2, 8.4 and 8.5 ■ Windows Server 2012 R2, 2016, 2019 and 2022 **Note:** CloudPoint agents are not supported on non-English operating systems.

Table 1-5 Supported applications, operating systems, cloud, and storage platforms *(continued)*

Category	Support
Cloud platforms	

Table 1-5 Supported applications, operating systems, cloud, and storage platforms *(continued)*

Category	Support
	■ 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. The t2 series instances are supported only if the device naming conventions recommended by AWS are followed. For more details, refer to the following links: ■ Windows: https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/device_naming.html ■ Linux: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/device_naming.html For protecting Microsoft Windows-based applications, use t2.xlarge or t3.xlarge or a higher specification 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. For protecting Microsoft Windows-based applications, use B4ms or D4s_V3 or a higher specification virtual machine. **Note:** The CloudPoint Azure plug-in supports disks of type Premium_LRS, Standard_LRS, and StandardSSD_LRS. All other disk types are defaulted to Standard_LRS during snapshot restore operations. ■ Microsoft Azure Stack Hub (2008 and later) If you wish to protect applications, the applications must be hosted on a DS2_v2 Standard or a higher specification Azure Stack Hub virtual machine type. See https://docs.microsoft.com/en-us/azure-stack/user/azure-stack-vm-sizes?view=azs-2008 **Note:** The CloudPoint Azure Stack Hub plug-in supports disks of type Premium_LRS, Standard_LRS, and StandardSSD_LRS. All other disk types are defaulted to Standard_LRS during snapshot restore operations.

Table 1-5 Supported applications, operating systems, cloud, and storage platforms *(continued)*

Category	Support
	■ Google Cloud Platform (GCP) If you wish to protect applications, the applications must be hosted on a n2-standard-4 or a higher specification GCP virtual machine type.
Storage platforms	■ NetApp storage arrays See "NetApp plug-in configuration notes" on page 105. ■ Dell EMC Unity arrays See "Dell EMC Unity array plug-in configuration parameters" on page 116. ■ HPE storage arrays See "HPE RMC plug-in configuration notes" on page 137. ■ Pure Storage FlashArray See "Pure Storage FlashArray plug-in configuration notes" on page 132. ■ Hitachi storage arrays See "Hitachi plug-in configuration notes" on page 143. ■ InfiniBox enterprise arrays See "InfiniBox plug-in configuration notes" on page 150.

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.

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.

The following table displays 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 1-6 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.

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 plug-in data to the application host, performs the discovery and configuration tasks, and then removes the plug-in package from the application host.

Veritas recommends the following configurations for the CloudPoint host:

Table 1-7 Typical CloudPoint host configuration based on the number of concurrent tasks

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 **t3.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 **t3.2xlarge** 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 plug-in configuration files.
- Depending on the number of workloads, the amount of plug-in 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.

CloudPoint sizing recommendations for cloud platforms

Note the following important points considering the standard sizing configurations:

- 20% of instances connected to CloudPoint host and performing granular restore and application consistent snapshots.
- Each protected instance has 3 disks of 100GB size attached.
- Protection cycle is twice daily with retention period of 3 months.
- /cloudpoint volume size is 50 GB or more for 400 instances and volume size is 100 GB or more for 500 instances.
- Based on cloud platform and instance types, if applicable, ensure appropriate CPU credits are available for selected instance types.

The following table provides configuration examples for the CloudPoint host:

Table 1-8 Google Cloud Platform

CloudPoint host	vCPU	Memory	Instances
■ n1-standard-2 ■ n2-standard-2	2	8	200
■ n1-standard-4 ■ n2-standard-4	4	16	400
■ n1-standard-16 ■ n2-standard-16	8	32	500

Table 1-9 Amazon Web Services

CloudPoint host	vCPU	Memory	Instances
■ t2.large ■ t3.large ■ m4.large	2	8	200
■ t2.xlarge ■ t3.xlarge ■ t3a.xlarge	4	16	400
■ m5.4xlarge ■ m4.4xlarge	8	32	500

Table 1-10 Microsoft Azure

CloudPoint host	vCPU	Memory	Instances
■ Standard_B2ms ■ Standard_D2s_v3 ■ Standard_D2_v4, standard_D2s_v4 ■ Standard_D2d_v4, Standard_D2ds_v4	2	8	200
■ Standard_B4ms ■ Standard_D4s_v3 ■ Standard_D4_v4, standard_D8s_v4 ■ Standard_D4d_v4, standard_D4ds_v4	4	16	400
■ Standard_B16ms ■ Standard_D16s_v3 ■ Standard_D16_v4, standard_D16s_v4 ■ Standard_D16d_v4, Standard_D16ds_v4	8	32	500

Table 1-11 Microsoft Azure Stack Hub

CloudPoint host	vCPU	Memory	Instances
■ Standard_DS2_v2 ■ Standard_D2_v2 ■ Standard_DS2 ■ Standard_D2	2	7	200
■ Standard_DS3_v2 ■ Standard_D3_v2 ■ Standard_DS3 ■ Standard_D3 ■ Standard_NV4as_v4	4	14	400
■ Standard_DS4_v2 ■ Standard_D4_v2 ■ Standard_DS4 ■ Standard_D4	8	28	500

CloudPoint extension sizing recommendations

The CloudPoint extension serves the purpose of scaling the capacity of the CloudPoint host to service a large number of requests concurrently running on the CloudPoint server at its peak performance capacity. You can install one or more CloudPoint extensions on-premise or in cloud, depending on your requirements to run the jobs without putting the host under additional stress. An extension can increase the processing capacity of the CloudPoint.

The CloudPoint extension can have the configuration same or higher as the CloudPoint host.

See " Meeting system requirements" on page 14.

Supported CloudPoint extension environments:

- VM based extension for on-premise
- Cloud based extension with managed Kubernetes cluster

Note: For CloudPoint 10.0, the VM based extensions are supported on Azure Stack hub and Kubernetes based extension are supported on Azure and AWS.

Veritas recommends the following configurations for the CloudPoint extensions:

Table 1-12 Typical CloudPoint extension configuration for VM based extension (Azure stack)

Workload metric	CloudPoint extension configuration
Up to 16 concurrent operational tasks	CPU: 4 CPUs Memory: 16 GB For example, in Azure stack, the CloudPoint extension should be an equivalent of a **t3.xlarge** instance in AWS.
Up to 32 concurrent operational tasks	CPU: 8 CPUs Memory: 32 GB or more For example, in Azure stack, the CloudPoint extension should be an equivalent of a **t3.2xlarge** or a higher type of instance in AWS.

Table 1-13 Typical CloudPoint extension configuration for Kubernetes based extension (Azure, AWS)

Workload metric	CloudPoint extension configuration
Up to 24 concurrent operational tasks	**For Azure** CPU: More than 2 CPU's per node Memory: 8 GB per node Maximum pods per node: 12 CP pods + 6 (system AKS pods) = 18 or more Autoscaling enabled, with minimum =1 and maximum =3
Up to 24 concurrent operational tasks	**For AWS** CPU: More than 2 CPU's per node Memory: 8 GB per node Autoscaling enabled, with minimum =1 and maximum =3

General considerations and guidelines:

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

- To achieve better performance in a high workload environment, Veritas recommends that you deploy the CloudPoint extension in the same location as that of the application hosts.
- The cloud-based extension on a managed Kubernetes cluster should be in the same VNet as that of the CloudPoint host. If it is not, then you can make use of the VNet peering mechanism available with the Azure cloud, to make sure that CloudPoint host and extension nodes can communicate with each other over the required ports
- Depending on the number of workloads, the amount of plug-in 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.
- In cases where the number of concurrent operations is higher than what the CloudPoint host and the extensions together 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 host to install CloudPoint

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

- Choose a supported Ubuntu, RHEL, or SLES instance image that meets CloudPoint installation requirements.
- Add sufficient storage to the instance to meet the installation requirements.

If you are deploying CloudPoint on an on-premise instance, do the following:

- Install a supported Ubuntu, RHEL, or SLES operating system on a physical or a virtual x86 server.
- Add sufficient storage to the server to meet the installation requirements.

Installing container platform (Docker, Podman)

Table 1-14 Installing container platform

Platform	Description
Docker on Ubuntu	Supported version: Docker 18.09 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 1-14 Installing container platform *(continued)*

Platform	Description

Docker on RHEL 7.x	Supported version: Docker 1.13.x and later
	
	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.
	
	■ (If CloudPoint is being deployed in AWS cloud) Ensure that you enable the extra repos:
	`# sudo yum-config-manager --enable rhui-REGION-rhel-server-extras`
	■ (If CloudPoint is being deployed on-premise) Enable your subscriptions:
	`# 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`
	■ Install Docker using the following command:
	`# sudo yum -y install docker`
	■ Reload the system manager configuration using the following command:
	`# sudo systemctl daemon-reload`
	■ Enable and then restart the docker service using the following commands:
	`# sudo systemctl enable docker`
	`# sudo systemctl restart docker`
	■ If SELinux is enabled, change the mode to permissive mode.
	Edit the `/etc/selinux/config` configuration file and modify the `SELINUX` parameter value to `SELINUX=permissive`.
	■ Reboot the system for the changes to take effect.
	■ Verify that the SELinux mode change is in effect using the following command:
	`# sudo sestatus`
	The `Current Mode` parameter value in the command output should appear as `permissive`.
	
	Refer to the following for detailed instructions on installing Docker on RHEL:
	
	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
	
	If the docker is using default storage driver (overlay2 or overlay) on XFS backed file system, then ensure that XFS FS has `ftype` option set to *1*. Use `xfs_info` to verify. For details, see https://docs.docker.com/storage/storagedriver/overlayfs-driver/. Otherwise, you can use different storage driver. For details, see https://docs.docker.com/storage/storagedriver/select-storage-driver/

Table 1-14 Installing container platform *(continued)*

Platform	Description

Podman on RHEL 8.3 and 8.4	Supported versions:
	For RHEL 8.5: Podman version 2.2.1 or above
	For RHEL8.4: Podman version 2.2.1 or above
	Notes:
	■ (If CloudPoint is being deployed in AWS cloud) Ensure that you enable the extra repos:
	`# sudo yum-config-manager --enable rhui-REGION-rhel-server-extras`
	■ (If CloudPoint is being deployed on-premise) Enable your subscriptions:
	`# sudo subscription-manager register --auto-attach --username=<username> --password=<password>`
	■ If SELinux is enabled, change the mode to permissive mode.
	Edit the `/etc/selinux/config` configuration file and modify the `SELINUX` parameter value to `SELINUX=permissive`.
	■ Reboot the system for the changes to take effect.
	■ Verify that the SELinux mode change is in effect using the following command:
	`# getenforce`
	The `Current Mode` parameter value in the command output should appear as `permissive`.

Creating and mounting a volume to store CloudPoint data

Before you deploy the CloudPoint or CloudPoint extension 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`.
- Ensure that UUID of the volume and the mount point (`/cloudpoint`) are mentioned in the `/etc/fstab` so that the volume is auto mounted when the host or the extension is rebooted.

Note: If you ever boot your instance without this volume attached (for example, after moving the volume to another instance), the `nofail` mount option enables the instance to boot even if there are errors mounting the volume.

Table 1-15 Volume creation steps for each supported cloud vendor

Vendor	Procedure
Amazon Web Services (AWS)	**1** On the EC2 dashboard, click **Volumes > Create Volumes**. **2** Follow the instructions on the screen and specify the following: ■ Volume type: General Purpose SSD ■ Size: 50 GB **3** Use the following instructions to create a file system and mount the device to `/cloudpoint` on the instance host. http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-using-volumes.html
Google Cloud Platform	◆ Create the disk for the virtual machine, initialize it, and mount it to `/cloudpoint`. https://cloud.google.com/compute/docs/disks/add-persistent-disk
Microsoft Azure	**1** Create a new disk and attach it to the virtual machine. https://docs.microsoft.com/en-us/azure/virtual-machines/linux/attach-disk-portal You should choose the managed disk option. https://docs.microsoft.com/en-us/azure/virtual-machines/linux/attach-disk-portal#use-azure-managed-disks **2** Initialize the disk and mount it to `/cloudpoint`. For details, see the section "Connect to the Linux VM to mount the new disk" in the following link: https://docs.microsoft.com/en-us/azure/virtual-machines/linux/add-disk
Microsoft Azure Stack Hub	**1** Create a new disk and attach it to the virtual machine. https://docs.microsoft.com/en-us/azure-stack/user/azure-stack-manage-vm-disks?adding-new-disks You should choose the managed disk option. **2** Initialize the disk and mount it to `/cloudpoint`. For details, see the section "Connect to the Linux VM to mount the new disk" in the following link: https://docs.microsoft.com/en-us/azure/virtual-machines/linux/add-disk

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 1-16 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, extensions, backup from snapshot, and restore from backup jobs.

Keep in mind the following:

- If the instance is in a cloud, configure the ports information under required inbound rules for your cloud.
- Once you configure the port when you install CloudPoint, you cannot change it when you upgrade.

Preparing CloudPoint for backup from snapshot jobs

For backup from snapshot jobs, you must have media server 9.1 or later.

Required ports:

- Port required on NetBackup primary server: 1556 and 443
- Ports required on NetBackup media server for client side deduplication: 10082 and 10102

If you use private names for installing certificates and communicating with NetBackup, which have to be resolved using `/etc/hosts` follow these steps:

- Add entries similar to `/etc/hosts` file in the */cloudpoint/openv/etc/hosts* file.
- Make sure that you use the private name during CloudPoint installation, as well as CloudPoint registration.

Chapter 2

Deploying CloudPoint using container images

This chapter includes the following topics:

Before you begin installing CloudPoint

Make sure that you complete the following before installing CloudPoint:

- Decide where to install CloudPoint.
  See “Deciding where to run CloudPoint” on page 12.

  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 14.
- Create the instance on which you install CloudPoint or prepare the physical host.
  See “Creating an instance or preparing the host to install CloudPoint” on page 28.
- Install a container platform

See Table 1-14 on page 28.

- Create and mount a volume to store CloudPoint data.
  See "Creating and mounting a volume to store CloudPoint data" on page 30.
- 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 32.

Note: RedHat 8.x has replaced the Docker ecosystem with the Podman ecosystem. Hence, for deploying CloudPoint on a RHEL8.5 or 8.4 hosts See "Installing CloudPoint in the Podman environment" on page 39.. For RHEL 7.x hosts See "Installing CloudPoint in the Docker environment" on page 34.

Installing CloudPoint in the Docker environment

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 install CloudPoint

1 Download the CloudPoint image to the system on which you want to deploy CloudPoint. Go to the Veritas support site:

https://www.veritas.com/content/support/en_US/downloads

Note: You must log on to the support site to download

From the **Products** drop-down, select **NetBackup** and select the required version from the **Version** drop-down. Click **Explore**. Click **Base and upgrade** installers.

The CloudPoint image name resembles the following format:

`VRTScloudpoint-docker-x.x.x.x.x.img.gz`

Note: The actual file name may vary depending on the release version.

---

2 Change directories to where you have downloaded the CloudPoint image.

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

```
# sudo docker load -i CloudPoint_image_name
```

For example:

```
# sudo docker load -i Veritas_CloudPoint_10.0.0.9818.img.gz
```

Messages similar to the following appear on the command line:

```
538bd068cab5: Loading layer [================>]  38.26MB/38.26MB
ed4b778f8d1d: Loading layer [================>]  1.166GB/1.166GB
c8b269899686: Loading layer [================>]  49.15kB/49.15kB
Loaded image: veritas/flexsnap-cloudpoint:10.0.0.0.9818
```

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.

4 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.

Parameter	Description
<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 `10.0.0.9818`, 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:10.0.0.9818 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:10.0.0.9818 install
```

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

The installer displays messages similar to the following:

```
Installing the services
Configuration started at time: Fri Mar 13 06:11:42 UTC 2020
WARNING: No swap limit support
Docker server version: 18.09.1
This is a fresh install of CloudPoint 10.0.0.9818
Checking if a 1.0 release container exists ...
CloudPoint currently is not configured.
Starting initial services before configuration.
Creating network: flexsnap-network ...done
Starting docker container: flexsnap-fluentd ...done
Creating docker container: flexsnap-mongodb ...done
Creating docker container: flexsnap-rabbitmq ...done
Creating docker container: flexsnap-certauth ...done
Creating docker container: flexsnap-api-gateway ...done
Creating docker container: flexsnap-coordinator ...done
Creating docker container: flexsnap-agent ...done
Creating docker container: flexsnap-onhostagent ...done
Creating docker container: flexsnap-scheduler ...done
Creating docker container: flexsnap-policy ...done
Creating docker container: flexsnap-notification ...done
Creating docker container: flexsnap-idm ...done
Starting docker container: flexsnap-config ...done
Creating self signed keys and certs for nginx ...done
Starting docker container: flexsnap-nginx ...done
```

In this step, CloudPoint does the following:

- Creates and runs the containers for each of the CloudPoint services.
- Creates self-signed keys and certificates for `nginx`.

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.

5 Provide the following details when prompted on the command prompt:

Parameter	Description
Admin username	Specify a user name for the CloudPoint administrator user account.
Admin password	Specify a password for the admin user.
Confirm Admin password	Confirm the admin user password.
Host name for TLS certificate	Specify the IP address or the Fully Qualified Domain Name (FQDN) of 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. Use commas to specify multiple entries. The names you specify here must point to the same CloudPoint host. The specified names or IP address are added to the list of host names to use for configuring CloudPoint. The installer uses these names to generate a server certificate for the CloudPoint host.
Port	Specify the port through which the CloudPoint server can communicate. Default is port 443.

The installer then displays messages similar to the following:

```
Configuring admin credentials ...done
Waiting for CloudPoint configuration to complete (22/22) ...done
Configuration complete at time Fri Mar 13 06:15:43 UTC 2020!
```

6 This concludes the CloudPoint deployment process. The next step is to register the CloudPoint server with the Veritas NetBackup primary server.

If CloudPoint is deployed in the cloud, refer to the *NetBackup Web UI Cloud Administrator's Guide* for instructions. If CloudPoint is deployed on-premise, refer to the *NetBackup Snapshot Client Administrator's Guide* for instructions.

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 41.

Installing CloudPoint in the Podman environment

CloudPoint installation prerequisites on Podman:

- Run the following commands to install the required packages (`lvm2`, `udev` and `dnsmaq`) on the hosts:

```
#yum install -y lvm2-<version>
#yum install -y lvm2-libs-<version>
#yum install -y python3-pyudev-<version>
#yum install -y systemd-udev-<version>
#yum install -y dnsmasq-<version>
```

- Run the following commands to lock the Podman and Conmon versions to the supported versions, so that they do not get updated with the *yum* update:
  For 2.2.1

```
sudo yum install -y podman-2.2.1-7.module+el8.3.1+9857+68fb1526
sudo yum install -y conmon-2:2.0.20-2.module+el8.3.0+8221+97165c3f
sudo yum install -y python3-dnf-plugin-versionlock
sudo yum versionlock podman* conmon*
```

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

See "Restarting CloudPoint" on page 41.

Verifying that CloudPoint is installed successfully

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

- Verify that a similar success message is displayed at the command prompt.

```
Configuration complete at time Fri Mar 13 06:15:43 UTC 2020!
```

- Run the following command and verify that the CloudPoint services are running and the status is displayed as `UP`:
  For Docker environment: `# sudo docker ps -a`

For Podman environment: `# podman ps -a`

The command output resembles the following:

```
CONTAINER ID   IMAGE                                              CREATED       STATUS
076d3c2252fb veritas/flexsnap-workflow:9.0.1.0.9261 system 3 days ago Up 3 days ago
flexsnap-workflow-system-0-min
07df8d5d083e veritas/flexsnap-rabbitmq:9.0.1.0.9261 3 days ago Up 3 days ago
flexsnap-rabbitmq
1d30b1922dad veritas/flexsnap-onhostagent:9.0.1.0.9261 3 days ago Up 3 days ago
flexsnap-onhostagent
4ecca5996401 veritas/flexsnap-notification:9.0.1.0.9261 3 days ago Up 3 days ago
flexsnap-notification
5c2763afe3bd veritas/flexsnap-nginx:9.0.1.0.9261 3 days ago Up 3 days ago
0.0.0.0:443->443/tcp flexsnap-nginx
5d5805787cda veritas/flexsnap-coordinator:9.0.1.0.9261 3 days ago Up 3 days ago
flexsnap-coordinator
64ebf4083dbd veritas/flexsnap-config:9.0.1.0.9261 3 days ago Exited (15) 3 days ago
flexsnap-config
6ca231fc35c2 veritas/flexsnap-certauth:9.0.1.0.9261 3 days ago Up 3 days ago
flexsnap-certauth
7356cabbb486 veritas/flexsnap-agent:9.0.1.0.9261 3 days ago Up 3 days ago
flexsnap-agent
756ba92314fb veritas/flexsnap-mongodb:9.0.1.0.9261 3 days ago Up 3 days ago
flexsnap-mongodb
79b7ad032fb7 veritas/flexsnap-workflow:9.0.1.0.9261 general 3 days ago Up 3 days ago
flexsnap-workflow-general-0-min
9018a4a7cb08 veritas/flexsnap-workflow:9.0.1.0.9261 indexing general 3 days ago Up 3 days ago
flexsnap-workflow-indexing-0-min
b9db2708f7f6 veritas/flexsnap-policy:9.0.1.0.9261 3 days ago Up 3 days ago
flexsnap-policy
cb3e69c27ab1 veritas/flexsnap-idm:9.0.1.0.9261 3 days ago Up 3 days ago
flexsnap-idm
d25d774ed2e8 veritas/flexsnap-scheduler:9.0.1.0.9261 3 days ago Up 3 days ago
flexsnap-scheduler
d58206a3c3d7 veritas/flexsnap-api-gateway:9.0.1.0.9261 3 days ago Up 3 days ago
0.0.0.0:8472->8472/tcp flexsnap-api-gateway
f522cedea280 veritas/flexsnap-listener:9.0.1.0.9261 3 days ago Up 3 days ago
flexsnap-listener
feced68604cc veritas/flexsnap-fluentd:9.0.1.0.9261 3 days ago Up 3 days ago
0.0.0.0:24224->24224/tcp flexsnap-fluentd
```

Note: The number (`9.0.1.0.9261`) displayed in the image name column represents the CloudPoint version. The version may vary depending on the actual product version being installed.

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

Restarting CloudPoint

If you need to restart CloudPoint, it's important that you restart it correctly so that your environmental data is preserved.

To restart CloudPoint in the Docker environment

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.

- 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:8.3.0.8549 restart
```

Note: Ensure that you enter the command without any line breaks.

To restart CloudPoint in the Podman environment

1. First, stop the CloudPoint by using the following command on the instance where CloudPoint is installed:

```
# podman run -it --rm --privileged
-v /cloudpoint:/cloudpoint
-v /run/podman/podman.sock:/run/podman/podman.sock
veritas/flexsnap-cloudpoint:<version> stop
```

2. Then, start it again by using the following command:

```
# podman run -it --rm --privileged
-v /cloudpoint:/cloudpoint
-v /run/podman/podman.sock:/run/podman/podman.sock
veritas/flexsnap-cloudpoint:<version> start
```

Note: Ensure that you enter the commands without any line breaks.

Troubleshooting CloudPoint restart

Refer to the following troubleshooting scenario:

Starting or restarting the CloudPoint services may fail if a stale IP address entry is retained in the Podamn layer on RHEL 8.3 environment

Sometimes the following error may be encountered when the cloudPoint service containers restart.

```
Error adding network: failed to allocate for
 range 0: 10.89.0.140 has been allocated to
 02da9e9aab2f79303c53dfb10b5ae6b6b70288d36b8fffbdfabba046da5a9afc,
 duplicate allocation is not allowed
ERRO[0000] Error while adding pod to CNI network
 "flexsnap-network": failed to allocate for
 range 0: 10.89.0.140 has been allocated to
 02da9e9aab2f79303c53dfb10b5ae6b6b70288d36b8fffbdfabba046da5a9afc,
 duplicate allocation is not allowed
Error: error configuring network namespace for container
 02da9e9aab2f79303c53dfb10b5ae6b6b70288d36b8fffbdfabba046da5a9afc:
 failed to allocate for range 0:
 10.89.0.140 has been allocated to
 02da9e9aab2f79303c53dfb10b5ae6b6b70288d36b8fffbdfabba046da5a9afc,
 duplicate allocation is not allowed"
```

The issue exists in the Podman subsystem which fails to remove the existing IP allocated for the container from `dir /var/lib/cni/networks/flexsnap-network/`, when the container is stopped.

Workaround

To remove the stale entry

1. Find the stale IP address which is retained when the containers are stopped. For example `10.89.0.140`, in the above error.
2. Run the following command to delete the stale entry from `dir`

   ```
   # rm /var/lib/cni/networks/flexsnap-network/<stale IP address>
   ```

3. Then start the service using

   ```
   # podman start <service-name>
   ```

Chapter 3

Deploying CloudPoint extensions

This chapter includes the following topics:

Before you begin installing CloudPoint extensions

The CloudPoint extensions which can be installed on a VM or on a managed Kubernetes cluster, can elastically scale out the compute infrastructure to service a large number of jobs, and then scale in as well when the jobs have completed.

Refer to the following appropriate preparatory steps for installing CloudPoint that also apply for installing CloudPoint extensions.

For a Kubernetes based extension

The CloudPoint cloud-based extension can be deployed on a managed Kubernetes cluster in Azure for scaling the capacity of the CloudPoint host to service a large number of requests concurrently. For more information on preparing the host and the managed Kubernetes cluster in Azure, See “Preparing to install the extension on a VM” on page ?.

For a VM based extension

- Decide where to install CloudPoint extension.
  See “Deciding where to run CloudPoint” on page 12.
- Ensure that your environment meets system requirements..
  See “ Meeting system requirements” on page 14.
- Create the instance or prepare the VM on which you want to install the CloudPoint extension.
  See “Creating an instance or preparing the host to install CloudPoint” on page 28.
- Install Docker on the VM or the instance on which you want to deploy the extension.
  See Table 1-14 on page 28.
- Create and mount a volume to store CloudPoint data. For a VM based extension, the volume size can be 30 GB .
  See “Creating and mounting a volume to store CloudPoint data” on page 30.
- Verify that specific ports are open on the instance or the main CloudPoint host and ensure that the hosts being protected are reachable from the extensions on required ports. Port 5671 and 443 needs to be opened for RabbitMQ communication on the CloudPoint host.

About the extension installation and configuration process

To install and configure the CloudPoint extension, perform tasks from the NetBackup user interface in your browser and on the command line of your local computer or the application host.

See “Preparing to install the extension on a VM” on page 47.

See “Installing the CloudPoint extension on a VM” on page 48.

See “Preparing to install the extension on a managed Kubernetes cluster (AKS) in Azure” on page 50.

See “Installing the CloudPoint extension on Azure (AKS)” on page 57.

See Installing the CloudPoint extension on a VM on page ?.

Downloading the CloudPoint extension

To download the extension

1. Sign in to the NetBackup Web UI.
2. From the left navigation pane, click **Cloud** and then select the **CloudPoint servers** tab.

   All the CloudPoint servers that are registered with the primary server are displayed in this pane.
3. From the desired CloudPoint server row, click the actions icon on the right and then select **Add extension**.

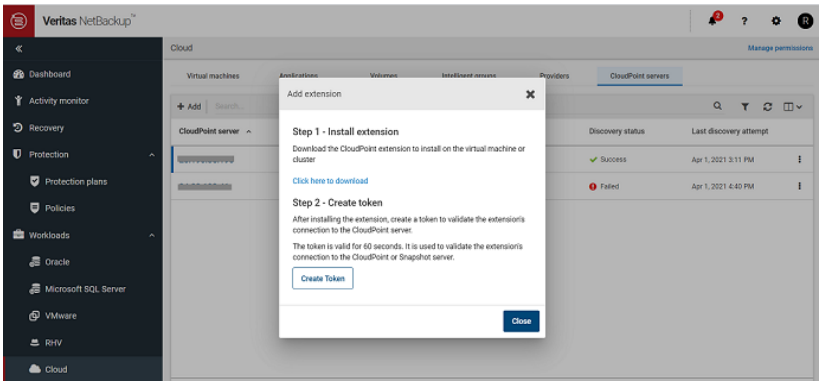

Note: For the VM-based extension you do not need to download the extension. Proceed directly to steps 7 and 8 to copy the token.

4. If you are installing the extension on a managed Kubernetes cluster, then on the Add extension dialog box, click the *download* hyperlink.

   This launches a new web browser tab.

   Do not close the Add extension dialog box yet. When you configure the extension, you will return to this dialog box to generate the validation token.
5. Switch to the new browser tab that opened and from the Add extension card, click **Download**. The extension file `nbu_cloudpoint_extension.tar` will be downloaded.

6 Copy the downloaded file to the CloudPoint host, and untar it by running the command `tar -xvf nbu_cloudpoint_extension.tar`.

See "Installing the CloudPoint extension on a VM" on page 48.

See "Installing the CloudPoint extension on a VM" on page 48.

See "Installing the CloudPoint extension on Azure (AKS)" on page 57.

7 Then to generate the validation token, on the Add extension dialog box, click **Create Token**

8 Click **Copy Token** to copy the displayed token. Then provide it on the command prompt while configuring the extension.

Note: The token is valid for 180 seconds only. If you do not use the token within that time frame, generate a new token.

Preparing to install the extension on a VM

Note: Currently, the extension is supported only on the Azure Stack Hub environment.

- Choose the CloudPoint image supported on Ubuntu or RHEL system that meets the CloudPoint installation requirements and create a host.
  See "Creating an instance or preparing the host to install CloudPoint" on page 28.
- Verify that you can connect to the host through a remote desktop.
  See "Verifying that specific ports are open on the instance or physical host" on page 32.
- Install Docker or Podman container platforms on the host.
  See Table 1-14 on page 28.
- Download the OS-specific CloudPoint image from the Veritas support site.
  - For Docker environment, load the image on the host.
    `# sudo docker load -i CloudPoint_image_name`
  - For Podman environment, un-tar the image file.
    `# gunzip VRTScloudpoint-podman-9.x.x.x.x.tar.gz`
    Run the following command to prepare the CloudPoint host for installation:
    `# ./flexsnap_preinstall.sh`

Note: The actual file name varies depending on the release version.

- For the VM based extension installed on a RHEL OS the SElinux mode should be "*permissive*"
- Network Security Groups used by the host that is being protected should allow communication from the host where the extension is installed, on the specified ports.

See "Installing the CloudPoint extension on a VM" on page 48.

Installing the CloudPoint extension on a VM

Before you install the CloudPoint extension:

See "Preparing to install the extension on a VM" on page 47.

To install the extension

1 **For Docker environment:**

Run 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> install_extension
```

Note: This is a single command without any line breaks.

For Podman environment:

Run the following command:

```
# podman run -it --rm --privileged
-v /<full_path_to_volume_name>:/<full_path_to_volume_name>
-v /run/podman/podman.sock:/run/podman/podman.sock
veritas/flexsnap-cloudpoint:<version> install_extension
```

Note: This is a single command without any line breaks.

In this step, CloudPoint does the following:

- Creates and runs the containers for each of the CloudPoint services.

- Creates self-signed keys and certificates for `nginx`.

2 Then go to the NetBackup Web UI and follow the steps 7 and 8 described in the section *Downloading CloudPoint extension* to generate and copy the validation token.

See "Downloading the CloudPoint extension" on page 46.

Note: For the VM-based extension you do not need to download the extension. Proceed directly to steps 7 and 8 to copy the token.

3 Provide the following configuration parameters when prompted:

Parameter	Description
Cloudpoint IP address / FQDN	Provide IP address or FQDN of the main CloudPoint host.
Cloudpoint Token	Paste the token obtained in the previous step.
Extension Name Identifier	Name of the extension identifier to be visible on the NetBackup UI

The installer then displays messages similar to the following:

```
Starting docker container: flexsnap-fluentd ...done
Starting docker container: flexsnap-ipv6config ...done
Starting docker container: flexsnap-listener ...done
```

This concludes the CloudPoint extension installation on a VM.

To verify that the extension is installed successfully:

- Verify that the success message is displayed at the command prompt.
- Verify that the extension is listed on the NetBackup Web UI.
  Go to **Cloud** > **CloudPoint Servers** tab > click **Advanced settings** > go to **CloudPoint extensions** tab and verify.
- Run the following command and verify that the CloudPoint containers are running and the status is displayed as `UP`:

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

  The command output resembles the following:

```
CONTAINER ID IMAGE COMMAND CREATED STATUS PORTS NAMES
e67550304195 veritas/flexsnap-workflow:9.1.x.x.xxxx "/usr/bin/flexsnap-w…" 13 minutes ago Up 13
26472ebc6d39 veritas/flexsnap-workflow:9.1.x.x.xxxx "/usr/bin/flexsnap-w…" 13 minutes ago Up 13
```

```
4f24f6acd290 veritas/flexsnap-listener:9.1.x.x.xxxx "/usr/bin/flexsnap-l…" 13 minutes ago Up 13
4d000f2d117d veritas/flexsnap-cloudpoint:9.1.x.x.xxxx "/root/ipv6_configur…" 13 minutes ago Exit
92b5bdf3211c veritas/flexsnap-fluentd:9.1.x.x.xxxx   "/root/flexsnap-flue…" 13 minutes ago Up 13
db1f0bff1797 veritas/flexsnap-datamover:9.1.x.x.xxxx "/entrypoint.sh -c d…" 13 minutes ago Up 13
c4ae0eb61fb0 veritas/flexsnap-datamover:9.1.x.x.xxxx "/entrypoint.sh -c d…" 13 minutes ago Up 13
1bcaa2b646fb veritas/flexsnap-datamover:9.1.x.x.xxxx "/entrypoint.sh -c d…" 13 minutes ago Up 13
```

Preparing to install the extension on a managed Kubernetes cluster (AKS) in Azure

The CloudPoint cloud-based extension can be deployed on a managed Kubernetes cluster in Azure for scaling the capacity of the CloudPoint host to service a large number of requests concurrently.

Overview

- Your Azure managed Kubernetes cluster should already be deployed with appropriate network and configuration settings, and with specific roles. The cluster must be able to communicate with CloudPoint.

  The required roles are: `Azure Kubernetes Service RBAC Writer, AcrPush, Azure Kubernetes Service Cluster User Role`

  For supported Kubernetes versions, refer to the *CloudPoint Hardware Compatibility List (HCL)*.
- Use an existing Azure Container Registry or create a new one, and ensure that the managed Kubernetes cluster has access to pull images from the container registry
- A dedicated nodepool for CloudPoint workloads needs to be created with manual scaling or 'Autoscaling' enabled in the Azure managed Kubernetes cluster. The autoscaling feature allows your nodepool to scale dynamically by provisioning and de-provisioning the nodes as required automatically.
- CloudPoint extension images (`flexsnap-cloudpoint, flexsnap-listener, flexsnap-workflow, flexsnap-fluentd, flexsnap-datamover`) need to be uploaded to the Azure container registry.

Prepare the host and the managed Kubernetes cluster in Azure

- Choose the CloudPoint image supported on Ubuntu or RHEL system that meets the CloudPoint installation requirements and create a host.

  See "Creating an instance or preparing the host to install CloudPoint" on page 28.
- It is not recommended to scale the cluster up or down when a job is running. It might cause the job to fail. Set the cluster size beforehand.

- Verify that the port 5671 is open on the main CloudPoint host.
  See "Verifying that specific ports are open on the instance or physical host" on page 32.
- The public IP of the virtual machine scale set via which the node pool is configured has to be allowed to communicate through port 22, on the workloads being protected.
- Install a Docker or Podman container platform on the host and start the container service.
  See Table 1-14 on page 28.
- Prepare the CloudPoint host to access Kubernetes cluster within your Azure environment.
  - Install Azure CLI.
    https://docs.microsoft.com/
  - Install Kubernetes CLI
    https://kubernetes.io/
  - Login to the Azure environment to access the Kubernetes cluster by running this command on Azure CLI:

    ```
    # az login --identity
    # az account set --subscription <subscriptionID>
    # az aks get-credentials --resource-group <resource_group_name>
    --name <cluster_name>
    ```
- Ensure to create an Azure Container Registry or use the existing one if available, to which the CloudPoint images will be pushed (uploaded). See Azure documentation:
  https://docs.microsoft.com/
  https://docs.microsoft.com/

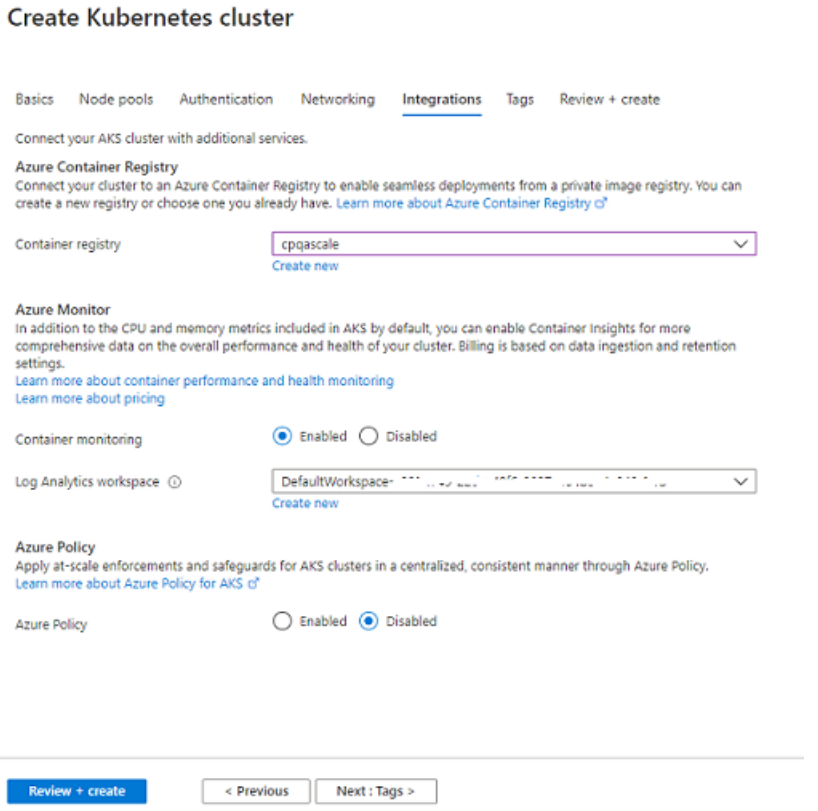

- To run the `kubectl` and container registry commands from the host system, assign the following role permissions to your VM and cluster. You can assign a 'Contributor', 'Owner', or any custom role that grants full access to manage all resources.
  - Go to your Virtual Machine > click **Identity** on the left > under **System assigned** tab, turn the **Status** to 'ON' > click **Azure role assignment** > click **Add role assignments** > select **Scope** as 'Subscription' or 'Resource Group' > select **Role** and assign the following roles : Azure Kubernetes Service RBAC Writer, AcrPush, Azure Kubernetes Service Cluster User Role, and **Save**.
  - Go to your Kubernetes cluster > click **Access Control (IAM)** on the left > click **Add role assignments** > select **Role** as 'Contributor ' > Select **Assign access to** as 'Virtual Machines' > select your VM from the drop-down and **Save**.

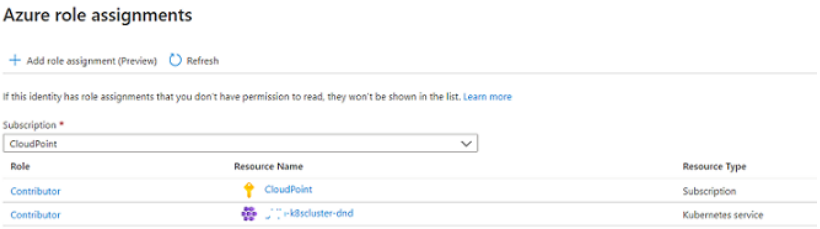

- Create a storage account in the same subscription and region your Kubernetes cluster is in, and create a file share into it. (Follow the default settings by Azure.) https://docs.microsoft.com/
- Create a namespace for CloudPoint from the command line on host system:

  ```
  # kubectl create namespace cloudpoint-system
  ```

- Then create a new or use an existing managed Kubernetes cluster in Azure, and add a new node pool dedicated for CloudPoint use. Configure Autoscaling as per your requirement.

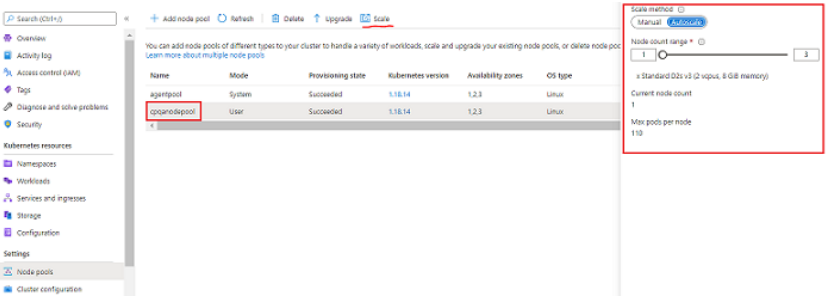

- Ensure that Azure plug-in is configured.
  See “Microsoft Azure plug-in configuration notes” on page 92.

See “Downloading the CloudPoint extension” on page 46.

See “Installing the CloudPoint extension on Azure (AKS)” on page 57.

Preparing to install the extension on a managed Kubernetes cluster (EKS) in AWS

The CloudPoint cloud-based extension can be deployed on a managed Kubernetes cluster in AWS for scaling the capacity of the CloudPoint host to service a large number of requests concurrently.

Overview

- Your AWS managed Kubernetes cluster should already be deployed with appropriate network and configuration settings, and with specific roles. The cluster must be able to communicate with CloudPoint.
  The required roles are: `AmazonEKSClusterPolicy AmazonEKSWorkerNodePolicy AmazonEC2ContainerRegistryReadOnly AmazonEKS_CNI_Policy AmazonEKSServicePolicy`
  For supported Kubernetes versions, refer to the *CloudPoint Hardware Compatibility List (HCL)*.
- Use an existing AWS Elastic Container Registry or create a new one, and ensure that the EKS has access to pull images from the elastic container registry.
- A dedicated nodepool for CloudPoint workloads needs to be created in AWS managed Kubernetes cluster. The nodepool uses AWS autoscaling group feature which allows your nodepool to scale dynamically by provisioning and de-provisioning the nodes as required automatically.
- CloudPoint extension images (`flexsnap-cloudpoint, flexsnap-listener, flexsnap-workflow, flexsnap-fluentd, flexsnap-datamover`) need to be uploaded to the AWS container registry.

Prepare the host and the managed Kubernetes cluster in AWS

- Choose the CloudPoint image supported on Ubuntu or RHEL system that meets the CloudPoint installation requirements and create a host.
  See "Creating an instance or preparing the host to install CloudPoint" on page 28.
- Verify that the port 5671 is open on the main CloudPoint host.
  See "Verifying that specific ports are open on the instance or physical host" on page 32.
- Install a Docker or Podman container platform on the host and start the container service.
  See Table 1-14 on page 28.
- Prepare the CloudPoint host to access Kubernetes cluster within your AWS environment.
  - Install AWS CLI.
    https://docs.aws.amazon.com/cli/latest/userguide/cli-chap-install.html
  - Install Kubernetes CLI
    https://docs.aws.amazon.com/eks/latest/userguide/install-kubectl.html
  - Create an AWS Container Registry or use the existing one if available, to which the CloudPoint images will be pushed (uploaded). Configure the minimum and maximum nodes as per the requirement.
    See AWS documentation https://aws.amazon.com/ecr/getting-started/.

- Create the OIDC provider for the AWS EKS cluster.
  See https://docs.aws.amazon.com/eks/latest/userguide/enable-iam-roles-for-service-accounts.html.
- Create an IAM service account for the AWS EKS cluster.
  See https://docs.aws.amazon.com/eks/latest/userguide/efs-csi.html.
- If an IAM role needs an access to the EKS cluster, run the following command from the system that already has access to the EKS cluster:
  `kubectl edit -n kube-system configmap/aws-auth`
  See https://docs.aws.amazon.com/eks/latest/userguide/add-user-role.html
- Install Amazon EFS driver .
  See https://docs.aws.amazon.com/eks/latest/userguide/efs-csi.html.
- Login to the AWS environment to access the Kubernetes cluster by running this command on AWS CLI:
  `# aws eks --region <region_name> update-kubeconfig --name <cluster_name>`

- Create a storage class.
  See https://docs.aws.amazon.com/eks/latest/userguide/efs-csi.html
- Create a namespace for CloudPoint from the command line on host system:
  `# kubectl create namespace cloudpoint-system`
- Then create a new or use an existing managed Kubernetes cluster in AWS, and add a new node pool dedicated for CloudPoint use. Configure Autoscaling as per your requirement.

See “Downloading the CloudPoint extension” on page 46.

See “Installing the CloudPoint extension on Azure (AKS)” on page 57.

See “Installing the CloudPoint extension on AWS (EKS)” on page 62.

Install extension using the Kustomize and CR YAMLs

The extension folder contains two sample YAMLs - kustomize.yaml, and cloudpoint_crd.yaml, based on which you need to create new YAMLs with the relevant values as per your environment.

kustomize.yaml

In the kustomize.yaml, update the parameters in the 'Image' section with relevant values as described in the following table.

Parameter	Description
newName	Specify the CloudPoint image name, along with the container registry path. Example: <account_id>.dkr.ecr.us-east-2.amazonaws.com/veritas/flexsnap-cloudpoint
newTag	Specify the tag of the CloudPoint image to be deployed. Example: `10.2.0.9129`

Example:

```
apiVersion: kustomize.config.k8s.io/v1beta1
kind: Kustomization
resources:
- cloudpoint_service.yaml
patchesStrategicMerge:
- node_select.yaml
namespace: demo-cloudpoint-ns
images:
- name: CLOUDPOINT_IMAGE
  newName: <account_id>.dkr.ecr.us-east-2.amazonaws.com/veritas/flexsnap-cloudpoint
  newTag: 10.2.0.9129
vars:
- name: ServiceAccount.cloudpoint-acc.metadata.namespace
  objref:
    kind: ServiceAccount
    name: cloudpoint-acc
    apiVersion: v1
  fieldref:
    fieldpath: metadata.namespace
configurations:
- cloudpoint_kustomize.yaml
```

cloudpoint_crd.yaml

In the cloudpoint_crd.yaml custom resource, update the parameters in the 'Spec' section with relevant values as described in the following table.

Parameter	Description
cloudpointHost	Specify the CloudPoint hostname or IP.

Parameter	Description
cloudpointExtensionToken	Paste the contents of the CloudPoint token that was downloaded earlier from NetBackup Web UI - Add extension dialog.
storageClassName	Kubernetes storage class that was created earlier in the preparation steps. Example: `efs-sc-new-root`
SIZE	Volume size in GB to be provisioned as per your scaling requirements.
NAMESPACE	The namespace that was created earlier in the preparation steps, in which to deploy CloudPoint.

Example:

```
apiVersion: veritas.com/v1
kind: CloudpointRule
metadata:
  name: cloudpoint-config-rule
  namespace: demo-cloudpoint-ns
spec:
  CLOUDPOINT_HOST: 3.17.**.*** .
  CLOUDPOINT_EXTENSION_TOKEN: workflow-3s3tlpwpy62dyoingxqmfeojlky7bub9rbzx8srh8kdgmsqo6f-q851fl
  RENEW: false
  LOG_STORAGE:
    STORAGE_CLASS_NAME: efs-sc-new-root
    SIZE: 100
```

Then run the following commands from the folder where the YAML files are located.

To apply the Kustomize YAML: `kubectl apply -k ./`

To apply the CloudPoint CR.`kubectl apply -f cloudpoint_crd.yaml`

Installing the CloudPoint extension on Azure (AKS)

Before you install the CloudPoint extension:

- See "Preparing to install the extension on a managed Kubernetes cluster (AKS) in Azure" on page 50.
- See "Downloading the CloudPoint extension" on page 46.

To install the extension

1 Download the extension script `nbu_cloudpoint_extension.tar`.

See "Downloading the CloudPoint extension" on page 46.

Note: Do not create the authentication token yet, as it is valid only for 180 seconds.

---

2 If the host from which you want to install the extension is not the same host where your CloudPoint is installed, load the CloudPoint container images on the extension host (`flexsnap-cloudpoint, flexsnap-listener, flexsnap-workflow, flexsnap-fluentd, flexsnap-datamover`)

The image names are in the following format:

Example: `veritas/flexsnap-cloudpoint`

3 Create image tags to map the source image with the target image, so that you can push the images to the Azure container registry.

See "Preparing to install the extension on a managed Kubernetes cluster (AKS) in Azure" on page 50.

Gather the following parameters beforehand:

Parameter	Description
container_registry_path	To obtain the container registry path, go to your container registry in Azure and from the **Overview** pane, copy the 'Login server'. Example: `mycontainer.azurecr.io`
tag	CloudPoint image version. Example: `9.0.1.0.9129`

- To tag the images, run the following command for each image, depending on the container platform running on your host:
  For Docker: `# docker tag source_image:tag target_image:tag`
  For Podman: `# podman tag source_image:tag target_image:tag`
  Where,
  - the source image tag is: `veritas/flexsnap-cloudpoint:tag>`
  - the target image tag is:
    `<container_registry_path>/<source_image_name>:<CloudPoint_version_tag>`

Example:

```
# docker tag veritas/flexsnap-cloudpoint:9.0.1.0.9129 mycontainer.azurecr.io/veritas/flexsnap-cl
# docker tag veritas/flexsnap-listener:9.0.1.0.9129 mycontainer.azurecr.io/veritas/flexsnap-list
# docker tag veritas/flexsnap-fluentd:9.0.1.0.9129 mycontainer.azurecr.io/veritas/flexsnap-fluen
# docker tag veritas/flexsnap-workflow:9.0.1.0.9129 mycontainer.azurecr.io/veritas/flexsnap-work
# docker tag veritas/flexsnap-datamover:9.0.1.0.9129 mycontainer.azurecr.io/veritas/flexsnap-dat
```

4. Then to push the images to the container registry, run the following command for each image, depending on the container platform running on your host:

   For Docker: `# docker push target_image:tag`

   For Podman: `# podman push target_image:tag`

   Example:

```
# docker push mycontainer.azurecr.io/veritas/flexsnap-cloudpoint:9.0.1.0.9129
# docker push mycontainer.azurecr.io/veritas/flexsnap-listener:9.0.1.0.9129
# docker push mycontainer.azurecr.io/veritas/flexsnap-fluentd:9.0.1.0.9129
# docker push mycontainer.azurecr.io/veritas/flexsnap-workflow:9.0.1.0.9129
# docker push mycontainer.azurecr.io/veritas/flexsnap-datamover:9.0.1.0.9129
```

5. Once the images are pushed to the container registry, execute the extension script `cp_extension.sh` that was downloaded earlier, from the host where `kubectl` is installed. The script can be executed either by providing all the required input parameters in one command, or in an interactive way where you will be prompted for input.

   Gather the following parameters before running the script:

Parameter	Description
cloudpoint_ip	Provide IP address or FQDN of the main CloudPoint host.
target_image:tag	Target image tag created for the `flexsnap-cloudpoint` image in step 3. Example: 'mycontainer.azurecr.io/veritas/flexsnap-cloudpoint:9.0.1.0.9129'
namespace	CloudPoint `namespace` that was created earlier in the preparation steps.

Parameter	Description

tag_key=tag_val	`tag_key` and `tag_val` can be retrieved by using these commands:
	1 Get the name of the node:
	`# kubectl get nodes | grep <node_name>`
	2 Get the tag key=value label:
	`# kubectl describe node <node_name> -n <namespace> | grep -i labels`
	Output example: `agentpool=cpuserpool`

storage_class	Kubernetes storage class that was created earlier in the preparation steps. Example: `cloudpoint-sc`
Size in GB	Volume size to be provisioned as per your scaling requirements.
workflow_token	Authentication token created from the NetBackup Web UI - Add extension dialog. See "Downloading the CloudPoint extension" on page 46.

Note: While deploying CloudPoint Kubernetes extension, create a storage class and provide it as an input to the CloudPoint extension installation script. By default file properties are open, hence it is recommended to create storage class by providing custom attributes in order to maintain the file/folder permissions created on extension under `/cloudpoint` directory. For more information, see Create a storage class section of the Azure product documentation.

Run the script as an executable file:

- Permit the script to run as an executable:
  ```
  # chmod +x cp_extension.sh
  ```
- Run the installation command with all the input parameters described in the above table:
  ```
  # ./cp_extension.sh install -c <cloudpoint_ip> -i
  <target_image:tag> -n <namespace> -p <tag_key=tag_val> -f
  <storage_class> -t <workflow_token>
  ```
  Example:

```
root@access-vm2-dnd:/home/cpuser/cp_ext# ./cp_extension.sh install
Veritas CloudPoint image repository path. Format=<Login-server/image:tag>: cpscale1.azurecr.io/v
CloudPoint extension namespace: ext
CloudPoint IP or fully-qualified domain name: 10.244.63.154
Node group/pool label with format key=value: agentpool=extpool1
Storage class name: nbux-sc
Size in GB : 100
CloudPoint extension token:
This is a fresh NetBackup CloudPoint Extension Installation

Starting CloudPoint service deployment
customresourcedefinition.apiextensions.k8s.io/cloudpoint-servers.veritas.com unchanged
serviceaccount/cloudpoint-acc unchanged
clusterrole.rbac.authorization.k8s.io/cloudpoint-ext unchanged
clusterrolebinding.rbac.authorization.k8s.io/cloudpoint-rolebinding-ext unchanged
deployment.apps/flexsnap-cloudpoint created
CloudPoint service deployment ...done

customresourcedefinition.apiextensions.k8s.io/cloudpoint-servers.veritas.com condition met
Generating CloudPoint Custom Resource Definition object
Waiting for deployment "flexsnap-cloudpoint" rollout to finish: 0 of 1 updated replicas are avai
deployment "flexsnap-cloudpoint" successfully rolled out
cloudpointrule.veritas.com/cloudpoint-config-rule created
CloudPoint extension installation ...done
root@access-vm2-dnd:/home/cpuser/cp_ext# kubectl get pods -n ext
NAME                                          READY   STATUS    RESTARTS   AGE
flexsnap-cloudpoint-d8fb97c49-swp7v           1/1     Running   0          5m53s
flexsnap-fluentd-b6vxz                        1/1     Running   0          5m40s
flexsnap-fluentd-collector-867c9cf776-q58bw   1/1     Running   0          5m40s
flexsnap-listener-6f9f5cf7fd-9bsm4            1/1     Running   0          5m40s
```

Run the script as an interactive file:

- Run the following command:

  ```
  # ./cp_extension.sh install
  ```

- When the script runs, provide the input parameters as described in the above table:

```
Veritas CloudPoint image repository path. Format=<Login-server/image:tag>: cpscale1.azurecr.io/v
CloudPoint extension namespace: ext
CloudPoint IP or fully-qualified domain name: 10.244.63.154
Node group/pool label with format key=value: agentpool=extpool1
```

```
Storage class name: nbux-sc
Size in GB : 100
CloudPoint extension token:
This is a fresh NetBackup CloudPoint Extension Installation

Starting CloudPoint service deployment
customresourcedefinition.apiextensions.k8s.io/cloudpoint-servers.veritas.com unchanged
serviceaccount/cloudpoint-acc unchanged
clusterrole.rbac.authorization.k8s.io/cloudpoint-ext unchanged
clusterrolebinding.rbac.authorization.k8s.io/cloudpoint-rolebinding-ext unchanged
deployment.apps/flexsnap-cloudpoint created
CloudPoint service deployment ...done

customresourcedefinition.apiextensions.k8s.io/cloudpoint-servers.veritas.com condition met
Generating CloudPoint Custom Resource Definition object
Waiting for deployment "flexsnap-cloudpoint" rollout to finish: 0 of 1 updated replicas are avai
deployment "flexsnap-cloudpoint" successfully rolled out
cloudpointrule.veritas.com/cloudpoint-config-rule created
CloudPoint extension installation ...done
```

Note: The output examples have been formatted to fit the screen.

This concludes the CloudPoint extension installation on a managed Kubernetes cluster (in Azure cloud).

To verify that the extension is installed successfully:

- Verify that the success message is displayed at the command prompt.
- Verify that the extension is listed on the NetBackup Web UI.
  Go to **Cloud** > **CloudPoint Servers** tab > click **Advanced settings** > go to **CloudPoint extensions** tab and verify.
- Run the following command and verify that there are four pods, namely, `flexsnap-cloudpoint-xxx`, `flexsnap-fluentd-xxx`, `flexsnap-listener-xxx`, `flexsnap-fluentd-collector-xxx`, `flexsnap-datamover-xxxx` in Running state:
  `# kubectl get pods -n <namespace>`
  Example: `# kubectl get pods -n cloudpoint-system`

Installing the CloudPoint extension on AWS (EKS)

Before you install the CloudPoint extension:

- See "Preparing to install the extension on a managed Kubernetes cluster (EKS) in AWS" on page 53.
- See "Downloading the CloudPoint extension" on page 46.

To install the extension

1. The extension file `nbu_cloudpoint_extension.tar` must be downloaded beforehand.

   See "Downloading the CloudPoint extension" on page 46.

   Note: Do not create the authentication token yet, as it is valid only for 180 seconds.

2. If the host from which you want to install the extension is not the same host where your CloudPoint is installed, load the CloudPoint container images on the extension host (`flexsnap-cloudpoint, flexsnap-listener, flexsnap-workflow, flexsnap-fluentd, flexsnap-datamover`)

   The image names are in the following format:

   Example: `veritas/flexsnap-cloudpoint`

3. Create image tags to map the source image with the target image, so that you can push the images to the AWS container registry.

   See "Preparing to install the extension on a managed Kubernetes cluster (EKS) in AWS" on page 53.

   Gather the following parameters beforehand:

Parameter	Description
container_registry_path	To obtain the container registry path, go to your Amazon ECR and copy the URI of each repo. Example: <accoint_id>.dkr.ecr.us-east-2.amazonaws.com/veritas/flexsnap-datamover
tag	CloudPoint image version. Example: `10.2.0.9129`

   - To tag the images, run the following command for each image, depending on the container platform running on your host:
     For Docker: `# docker tag source_image:tag target_image:tag`
     For Podman: `# podman tag source_image:tag target_image:tag`
     Where,

- the source image tag is: `veritas/flexsnap-cloudpoint:tag>`
- the target image tag is:
  `<container_registry_path>/<source_image_name>:<CloudPoint_version_tag>`

Example:

```
docker tag veritas/flexsnap-cloudpoint:10.2.0.9129
<account_id>.dkr.ecr.us-east-2.amazonaws.com/veritas/flexsnap-cloudpoint:10.2.0.9129
docker tag veritas/flexsnap-listener:10.2.0.9129
<account_id>.dkr.ecr.us-east-2.amazonaws.com/veritas/flexsnap-listener:10.2.0.9129
docker tag veritas/flexsnap-fluentd:10.2.0.9129
<account_id>.dkr.ecr.us-east-2.amazonaws.com/veritas/flexsnap-fluentd:10.2.0.9129
docker tag veritas/flexsnap-workflow:10.2.0.9129
<account_id>.dkr.ecr.us-east-2.amazonaws.com/veritas/flexsnap-workflow:10.2.0.9129
docker tag veritas/flexsnap-datamover:10.2.0.9129
<account_id>.dkr.ecr.us-east-2.amazonaws.com/veritas/flexsnap-datamover:10.2.0.9129
```

4 Then to push the images to the container registry, run the following command for each image, depending on the container platform running on your host:

For Docker: `# docker push target_image:tag`

For Podman: `# podman push target_image:tag`

Example:

```
docker push <account-id>.dkr.ecr.us-east-2.amazonaws.com/veritas/flexsnap-datamover:10.2.0.9129
docker push <account-id>.dkr.ecr.us-east-2.amazonaws.com/veritas/flexsnap-cloudpoint:10.2.0.9129
docker push <account-id>.dkr.ecr.us-east-2.amazonaws.com/veritas/flexsnap-fluentd:10.2.0.9129
docker push <account-id>.dkr.ecr.us-east-2.amazonaws.com/veritas/flexsnap-listener:10.2.0.9129
docker push <account-id>.dkr.ecr.us-east-2.amazonaws.com/veritas/flexsnap-workflow:10.2.0.9129
```

Note: The command/output examples may be formatted or truncated to fit the screen.

5 Once the images are pushed to the container registry, you can install the extension by one of the following two ways:

- Kustomization and custom resource YAML files: create and apply the `kustomize.yaml` and `cloudpoint_crd.yaml` files based on the samples provided.
  See “Install extension using the Kustomize and CR YAMLs” on page 55.
- Extension script: execute the extension script `cp_extension.sh` that is packaged within the 'tar' file that was downloaded earlier. The script can

be executed either by providing all the required input parameters in one command, or in an interactive way where you will be prompted for input.
See "Install extension using the extension script" on page 65.

After following the above instructions, you can verify if the extension was installed successfully.

To verify that the extension is installed successfully:

- Verify that the success message is displayed at the command prompt.
- Verify that the extension is listed on the NetBackup Web UI.
  Go to **Cloud** > **CloudPoint Servers** tab > click **Advanced settings** > go to **CloudPoint extensions** tab and verify.
- Run the following command and verify that there are four pods, namely, `flexsnap-cloudpoint-xxx, flexsnap-fluentd-xxx, flexsnap-listener-xxx, flexsnap-fluentd-collector-xxx, flexsnap-datamover-xxxx` in Running state:
  `# kubectl get pods -n <namespace>`
  Example: `# kubectl get pods -n cloudpoint-system`

Install extension using the extension script

Gather the following parameters before running the extension script:

Parameter	Description
cloudpoint_ip	Specify the CloudPoint hostname or IP.
target_image:tag	Target image tag created for the `flexsnap-cloudpoint` image. Example: <account_id>.dkr.ecr.us-east-2.amazonaws.com/veritas/flexsnap-cloudpoint:10.2.0.9129
namespace	The namespace that was created earlier in the preparation steps, in which to deploy CloudPoint.
tag_key=tag_val	`tag_key` and `tag_val` are the label key and value pair defined for the node on which you want to install the extension. The label key-value pair can be retrieved by using the command `kubectl describe node <node_name> -n <namespace>` Example: `eks.amazonaws.com/nodegroup=Demo-NG`

Parameter	Description
storage_class	Kubernetes storage class that was created earlier in the preparation steps. Example: `cloudpoint-sc`
Size in GB	Volume size to be provisioned as per your scaling requirements.
workflow_token	Authentication token created from the NetBackup Web UI - Add extension dialog. See "Downloading the CloudPoint extension" on page 46.

Run the script as an executable file:

- Permit the script to run as an executable:
  ```
  # chmod +x cp_extension_start.sh
  ```
- Run the installation command with all the input parameters described in the above table:
  ```
  # ./cp_extension_start.sh install -c <cloudpoint_ip> -i
  <target_image:tag> -n <namespace> -p <tag_key=tag_val> -f
  <storage_class> -t <workflow_token>
  ```
  Example:

```
root@access-vm2-dnd:/home/cpuser/cp_ext# ./cp_extension.sh install
Veritas CloudPoint image repository path. Format=<Login-server/image:tag>: cpscale1.azurecr.io/v
CloudPoint extension namespace: ext
CloudPoint IP or fully-qualified domain name: 10.244.63.154
Node group/pool label with format key=value: agentpool=extpool1
Storage class name: nbux-sc
Size in GB : 100
CloudPoint extension token:
This is a fresh NetBackup CloudPoint Extension Installation

Starting CloudPoint service deployment
customresourcedefinition.apiextensions.k8s.io/cloudpoint-servers.veritas.com unchanged
serviceaccount/cloudpoint-acc unchanged
clusterrole.rbac.authorization.k8s.io/cloudpoint-ext unchanged
clusterrolebinding.rbac.authorization.k8s.io/cloudpoint-rolebinding-ext unchanged
deployment.apps/flexsnap-cloudpoint created
CloudPoint service deployment ...done

customresourcedefinition.apiextensions.k8s.io/cloudpoint-servers.veritas.com condition met
Generating CloudPoint Custom Resource Definition object
```

```
Waiting for deployment "flexsnap-cloudpoint" rollout to finish: 0 of 1 updated replicas are avai
deployment "flexsnap-cloudpoint" successfully rolled out
cloudpointrule.veritas.com/cloudpoint-config-rule created
CloudPoint extension installation ...done
root@access-vm2-dnd:/home/cpuser/cp_ext# kubectl get pods -n ext
NAME                                        READY   STATUS    RESTARTS   AGE
flexsnap-cloudpoint-d8fb97c49-swp7v         1/1     Running   0          5m53s
flexsnap-fluentd-b6vxz                      1/1     Running   0          5m40s
flexsnap-fluentd-collector-867c9cf776-q58bw  1/1     Running   0          5m40s
flexsnap-listener-6f9f5cf7fd-9bsm4          1/1     Running   0          5m40s
```

Run the script as an interactive file:

- Run the following command:
  ```
  # ./cp_extension_start.sh install
  ```
- When the script runs, provide the input parameters as described in the above table:

Example:

```
Veritas CloudPoint image repository path. Format=<Login-server/image:tag>:
<account-id>.dkr.ecr.us-east-2.amazonaws.com/veritas/flexsnap-cloudpoint:10.2.0.9129
CloudPoint extension namespace: cloudpoint-system
CloudPoint IP or fully-qualified domain name: 18.117.***.***
Node pool with format key=value: eks.amazonaws.com/nodegroup=td-nodepool-dnd
AWS PVC Name: efs-claim
CloudPoint extension token:
This is a fresh NetBackup CloudPoint Extension Installation


Getting CloudPoint service file ...done
Getting CloudPoint CRD file ...done

Starting CloudPoint service deployment

namespace/cloudpoint-system configured

deployment.apps/flexsnap-cloudpoint created

serviceaccount/cloudpoint-acc created

clusterrole.rbac.authorization.k8s.io/cloudpoint-cloudpoint-system unchanged

clusterrolebinding.rbac.authorization.k8s.io/cloudpoint-rolebinding-cloudpoint-system unchanged
```

```
customresourcedefinition.apiextensions.k8s.io/cloudpoint-servers.veritas.com created

CloudPoint service deployment ...done

customresourcedefinition.apiextensions.k8s.io/cloudpoint-servers.veritas.com condition met

Generating CloudPoint Custom Resource Definition object

cloudpointrule.veritas.com/cloudpoint-config-rule created

CloudPoint extension installation ...done
```

Note: The output examples may be formatted or truncated to fit the screen.

Managing the extensions

After you have installed the VM-based or the managed Kubernetes cluster-based extensions, you may need to disable or enable them, stop, start, or restart them, or renew their certificates.

Refer to the following table that describes how to use these options to manage the extensions.

Table 3-1 Post-installation options for the extensions

Option	Procedure
Disable or enable the extension: ■ VM-based extension ■ Managed Kubernetes cluster-based extension	You can disable or enable the extensions from the NetBackup Web UI Go to **Cloud** > **CloudPoint Servers** tab > click **Advanced settings** > go to **CloudPoint extensions** tab > then disable or enable the extension as required, and click **Save.** No jobs will be scheduled on the extension that is disabled. **Note:** When CloudPoint is upgraded, all the extensions are automatically disabled. Then you need to upgrade the extensions with the same CloudPoint version and enable them manually from the NetBackup Web UI.

Table 3-1 Post-installation options for the extensions *(continued)*

Option	Procedure

Stop, start, or restart the VM-based extension
 	
To stop the extension:
	Execute the following commands on the extension host VM to stop/start/restart the extension:
	
	For Docker:
	
	`# 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`
	
	For Podman
	
	`# podman run -it --rm --privileged -v /<full_path_to_volume_name>:/<full_path_to_volume_name> -v /run/podman/podman.sock:/run/podman/podman.sock veritas/flexsnap-cloudpoint:<version> stop`

To start the extension:	For Docker:
	
	`# 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`
	
	For Podman
	
	`# podman run -it --rm --privileged -v /<full_path_to_volume_name>:/<full_path_to_volume_name> -v /run/podman/podman.sock:/run/podman/podman.sock veritas/flexsnap-cloudpoint:<version> start`

To restart the extension:	For Docker:
	
	`# 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> restart`
	
	For Podman
	
	`# podman run -it --rm --privileged -v /<full_path_to_volume_name>:/<full_path_to_volume_name> -v /run/podman/podman.sock:/run/podman/podman.sock veritas/flexsnap-cloudpoint:<version> restart`

Table 3-1 Post-installation options for the extensions *(continued)*

Option	Procedure
Renew certificate for a VM-based extension	**1** Run the following command on the extension host: `# sudo docker run -it --rm -v /<full_path_to_volume_name>:/<full_path_to_volume_nam -v /var/run/docker.sock:/var/run/docker.sock veritas/flexsnap-cloudpoint:<version> renew_extension` **2** Then provide the CloudPoint IP/FQDN, and the extension token which can be generated from NetBackup Web UI to begin renewal of the certificates. See "Installing the CloudPoint extension on a VM" on page 48.

Renew certificate for a managed Kubernetes cluster-based extension	**1** Download the extension installation script `cp_extension.sh` from the NetBackup Web UI .
	2 Execute the script from the host where `kubectl` is installed. Run the following commands:
	`# chmod +x cp_extension.sh`
	`# ./cp_extension.sh renew`
	3 Then provide the CloudPoint IP/FQDN, extension token (which can be generated from NetBackup Web UI), and the extension namespace to begin renewal of the certificates.
	See "Installing the CloudPoint extension on Azure (AKS)" on page 57.

Chapter 4

CloudPoint cloud plug-ins

This chapter includes the following topics:

How to configure the CloudPoint cloud plug-ins?

CloudPoint plug-ins are software modules that enable the discovery of your assets in the cloud or in an on-premise environment. After registering the CloudPoint server with the NetBackup primary server, you must configure the CloudPoint plug-ins to be able to protect your workloads using NetBackup.

How you configure the plug-ins depends on the asset type and how CloudPoint is deployed. If the CloudPoint server is deployed in the cloud and you want to protect workloads in the cloud, you must use the NetBackup Web UI to register the CloudPoint server and configure the CloudPoint cloud and application plug-ins. The overall steps to configure the plug-ins are similar, regardless of the asset type. Only the configuration parameters vary.

Refer to the *NetBackup Web UI Cloud Administrator's Guide* for information on how to configure cloud plug-ins.

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.

CloudPoint supports the following AWS regions:

Table 4-1 AWS regions supported by CloudPoint

AWS commercial regions	AWS GovCloud (US) regions
■ us-east-1, us-east-2, us-west-1, us-west-2 ■ ap-east-1, ap-south-1, ap-northeast-1, ap-northeast-2, ap-southeast-1, ap-southeast-2, ap-southeast-3 ■ eu-central-1, eu-west-1, eu-west-2, eu-west-3, eu-north-1, eu-south-1 Milan, eu-south-1 Cape Town ■ cn-north-1, cn-northwest-1 ■ ca-central-1 ■ me-south-1 ■ sa-east-1	■ us-gov-east-1 ■ us-gov-west-1

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 4-2 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 4-3 AWS plug-in configuration parameters: cloud deployment

CloudPoint configuration parameter	Description
For Source Account configuration	
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.
For Cross Account configuration	
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).
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 allowed list on your firewall.

- ec2.*.amazonaws.com
- sts.amazonaws.com
- rds.*.amazonaws.com
- kms. *.amazonaws.com
- ebs.*.amazonaws.com
- iam.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:

- CloudPoint does not support AWS Nitro-based instances that use EBS volumes that are exposed as non-volatile memory express (NVMe) devices.
  To allow CloudPoint to discover and protect AWS Nitro-based Windows instances that use NVMe EBS volumes, ensure that the AWS NVMe tool executable file, `ebsnvme-id.exe`, is present in any of the following locations on the AWS Windows instance:
  - `%PROGRAMDATA%\Amazon\Tools`
    This is the default location for most AWS instances.
  - `%PROGRAMFILES%\Veritas\Cloudpoint`
    Manually download and copy the executable file to this location.
  - System PATH environment variable
    Add or update the executable file path in the system's PATH environment variable.
    If the NVMe tool is not present in one of the mentioned locations, CloudPoint may fail to discover the file systems on such instances.
    You may see the following error in the logs:

    ```
    "ebsnvme-id.exe" not found in expected paths!"
    ```

- To allow CloudPoint to discover and protect Windows instances created from custom/community AMI.
  - AWS NVMe drivers must be installed on custom or community AMIs. See this link.
  - Install the `ebsnvme-id.exe` either in `%PROGRAMDATA%\Amazon\Tools` or `%PROGRAMFILES%\Veritas\Cloudpoint`
  - Friendly device name must contain the substring "NVMe", or update in Windows registry for all NVMe backed devices.
    Registry path:
    `Computer\HKEY_LOCAL_MACHINE\SYSTEM\ControlSet001\Enum\SCSI\Disk&Ven_NVMe&Prod_Amazon_Elastic_B\`
    Property Name: `FriendlyName`
    Value: `NVMe Amazon Elastic B SCSI Disk Drive`

- 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 CloudPoint 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:`.
- If you are configuring the plug-in to discover and protect AWS Nitro-based Windows instances that use NVMe EBS volumes, you must ensure that the AWS NVMe tool executable file, `ebsnvme-id.exe`, is present in any of the following locations on the AWS instance:
  - `%PROGRAMDATA%\Amazon\Tools`
    This is the default location for most AWS instances.
  - `%PROGRAMFILES%\Veritas\Cloudpoint`
    Manually download and copy the executable file to this location.
  - System PATH environment variable
    Add or update the executable file path in the system's PATH environment variable.

  If the NVMe tool is not present in one of the mentioned locations, CloudPoint may fail to discover the file systems on such instances. You may see the following error in the logs:

  `"ebsnvme-id.exe" not found in expected paths!"`

  This is required for AWS Nitro-based Windows instances only. Also, if the instance is launched using the community AMI or custom AMI, you might need to install the tool manually.
- 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. However, you cannot perform restore 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 supports commercial as well as GovCloud (US) regions. During AWS plug-in configuration, even though you can select a combination of AWS commercial and GovCloud (US) regions, the configuration will eventually fail.
- CloudPoint does not support IPv6 addresses for AWS RDS instances. This is a limitation of Amazon RDS itself and is not related to CloudPoint.
  Refer to the AWS documentation for more information:
  https://aws.amazon.com/premiumsupport/knowledge-center/rds-ipv6/
- CloudPoint does not support application consistent snapshots and granular file restores for Windows systems with virtual disks or storage spaces that are created from a storage pool. If a Microsoft SQL server snapshot job uses disks from a storage pool, the job fails with an error. But if a snapshot job for virtual machine which is in a connected state is triggered, the job might be successful. In this case, the file system quiescing and indexing is skipped. The restore job for such an individual disk to original location also fails. In this condition, the host might move to an unrecoverable state and requires a manual recovery.

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 77.
  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 84.
- 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 78.

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 71.

AWS permissions required by CloudPoint

The following is a IAM role definition (in JSON format) that gives CloudPoint the ability to configure AWS plugin and discover assets, manage the snapshots etc.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "EC2AutoScaling",
            "Effect": "Allow",
            "Action": [
                "autoscaling:UpdateAutoScalingGroup",
                "autoscaling:AttachInstances",
                "autoscaling:DescribeScalingActivities",
                "autoscaling:TerminateInstanceInAutoScalingGroup"
            ],
            "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",
            "rds:DescribeDBClusterParameterGroups"
        ],
        "Resource": [
            "*"
        ]
    },
    {
        "Sid": "EC2Backup",
        "Effect": "Allow",
        "Action": [
            "sts:GetCallerIdentity",
            "ec2:CreateSnapshot",
```

```
            "ec2:CreateSnapshots",
            "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" ,
            "ec2:DescribeNetworkInterfaces"
        ],
        "Resource": [
            "*"
        ]
    },
    {
        "Sid": "EC2Recovery",
        "Effect": "Allow",
        "Action": [
            "ec2:RunInstances",
            "ec2:AttachNetworkInterface",
            "ec2:DetachVolume",
            "ec2:AttachVolume",
            "ec2:DeleteTags",
            "ec2:CreateTags",
            "ec2:StartInstances",
            "ec2:StopInstances",
            "ec2:TerminateInstances",
```

```
            "ec2:CreateVolume",
            "ec2:DeleteVolume",
            "ec2:DescribeIamInstanceProfileAssociations",
            "ec2:AssociateIamInstanceProfile",
            "ec2:AssociateAddress",
            "ec2:DescribeKeyPairs",
            "ec2:AuthorizeSecurityGroupEgress",
            "ec2:AuthorizeSecurityGroupIngress",
            "ec2:DescribeInstanceTypeOfferings",
            "ec2:GetEbsEncryptionByDefault"
        ],
        "Resource": [
            "*"
        ]
    },
    {
        "Sid": "EBS",
        "Effect": "Allow",
        "Action": [
            "ebs:ListSnapshotBlocks",
            "ebs:GetSnapshotBlock",
            "ebs:CompleteSnapshot",
            "ebs:PutSnapshotBlock",
            "ebs:ListChangedBlocks"
        ],
        "Resource": [
            "*"
        ]
    },
    {
        "Sid": "EKS",
        "Effect": "Allow",
        "Action": [
            "eks:DescribeNodegroup",
            "eks:DescribeUpdate",
            "eks:UpdateNodegroupConfig",
            "eks:ListClusters"
        ],
        "Resource": [
            "*"
        ]
    },
    {
```

```
            "Sid": "IAM",
            "Effect": "Allow",
            "Action": [
                "iam:ListAccountAliases",
                "iam:SimulatePrincipalPolicy"
            ],
            "Resource": [
                "*"
            ]
        }
    ]
}
```

If a CloudPoint extension is installed on a managed Kubernetes cluster in AWS, then enable the following polices for a user account or a role before configuring the plugin:

```
AmazonEKSClusterPolicy
AmazonEKSWorkerNodePolicy
AmazonEC2ContainerRegistryReadOnly
AmazonEKS_CNI_Policy
AmazonEKSServicePolicy
```

Additional IAM permissions required for marketplace deployment

```
{
  "Sid": "AWSMarketplacePermissions",
  "Effect": "Allow",
  "Action": [
     "autoscaling:UpdateAutoScalingGroup",
     "autoscaling:AttachInstances",
     "sns:Publish",
     "sns:GetTopicAttributes",
     "secretsmanager:GetResourcePolicy",
     "secretsmanager:GetSecretValue",
     "secretsmanager:DescribeSecret",
     "secretsmanager:RestoreSecret",
     "secretsmanager:PutSecretValue",
     "secretsmanager:DeleteSecret",
     "secretsmanager:UpdateSecret"
  ],
  "Resource": [
    "*"
```

```
    ]
}
```

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**.

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.

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.

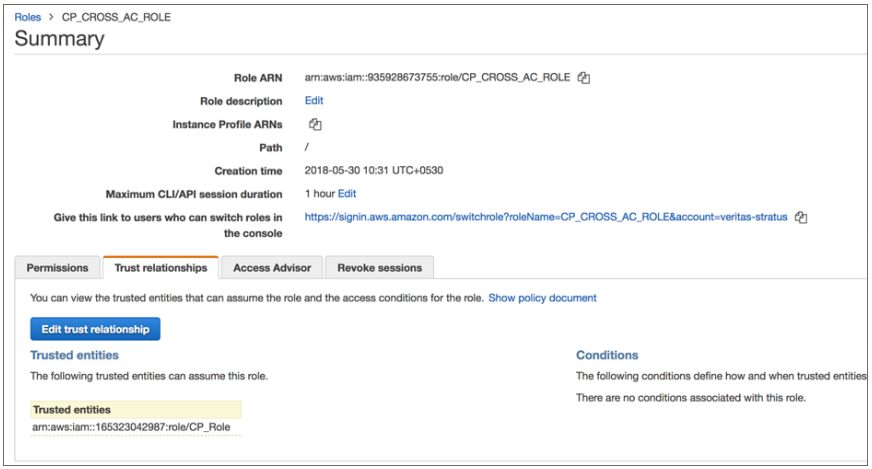

This action allows only the CloudPoint instance hosted in source AWS account to assume the target role using the credentials associated with source account's IAM role. No other entities can assume this role.

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"`).

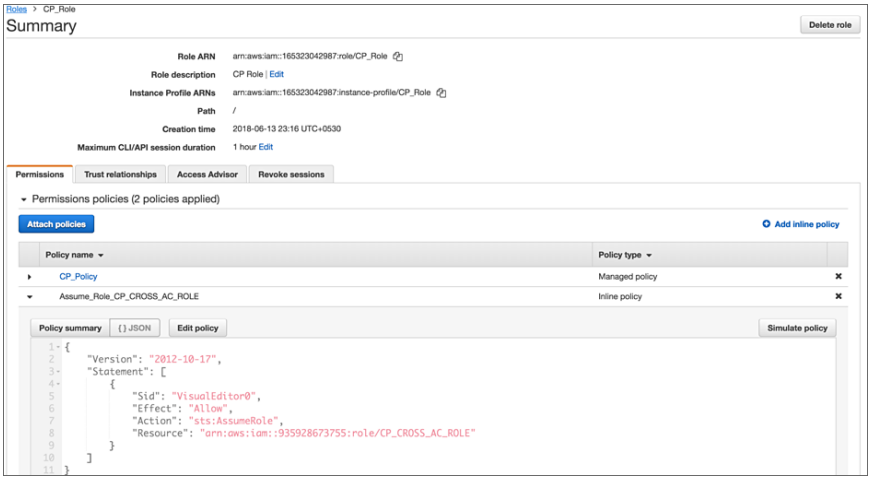

5 From the target account's Summary page, edit the **Maximum CLI/API session duration** field and set the duration to **1 hour**, at a minimum.

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.

Google Cloud Platform plug-in configuration prerequisites

Before you configure the Google Cloud Platform plug-in, enable the following APIs under **APIs & Services** from Google Cloud console:

- Cloud Resource Manager API
- Compute Engine API

Google Cloud Platform plug-in configuration parameters

The following parameters are required for configuring the Google Cloud Platform plug-in:

Table 4-4 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 `project_id` in the JSON file.
Client Email	The email address of the Client ID. Listed as `client_email` in the JSON file.
Private Key	The private key. Listed as `private_key` 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.

CloudPoint supports the following GCP zones:

Table 4-5 GCP zones supported by CloudPoint

GCP zones
■ asia-east1-a, asia-east1-b, asia-east1-c ■ asia-east2-a, asia-east2-b, asia-east2-c ■ asia-northeast1-a, asia-northeast1-b, asia-northeast1-c ■ asia-northeast2-a, asia-northeast2-b, asia-northeast2-c ■ asia-south1-a, asia-south1-b, asia-south1-c ■ asia-southeast1-a, asia-southeast1-b, asia-southeast1-c
■ australia-southeast1-a, australia-southeast1-b, australia-southeast1-c
■ europe-north1-a, europe-north1-b, europe-north1-c ■ europe-west1-b, europe-west1-c, europe-west1-d ■ europe-west2-a, europe-west2-b, europe-west2-c ■ europe-west3-a, europe-west3-b, europe-west3-c ■ europe-west4-a, europe-west4-b, europe-west4-c ■ europe-west6-a, europe-west6-b, europe-west6-c
■ northamerica-northeast1-a, northamerica-northeast1-b, northamerica-northeast1-c ■ southamerica-east1-a, southamerica-east1-b, southamerica-east1-c

Table 4-5 GCP zones supported by CloudPoint *(continued)*

GCP zones
■ us-central1-a, us-central1-b, us-central1-c, us-central1-f ■ us-east1-b, us-east1-c, us-east1-d ■ us-east4-a, us-east4-b, us-east4-c ■ us-west1-a, us-west1-b, us-west1-c ■ us-west2-a, us-west2-b, us-west2-c ■ us-west3-a Utah, us-west3-b Utah, us-west3-c Utah ■ us-west4-a Nevada, us-west4-b Nevada, us-west4-c Nevada

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. However, you cannot perform any restore 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.

See "Google Cloud Platform permissions required by CloudPoint" on page 88.

See "Configuring a GCP service account for CloudPoint" on page 90.

See "Preparing the GCP service account for plug-in configuration" on page 90.

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.setLabels
compute.disks.use
compute.globalOperations.get
compute.globalOperations.list
compute.images.get
compute.images.list
compute.instances.attachDisk
compute.instances.create
compute.instances.delete
compute.instances.detachDisk
compute.instances.get
compute.instances.list
compute.instances.setMetadata
compute.instances.setServiceAccount
compute.instances.setLabels
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.use
compute.subnetworks.useExternalIp
compute.zoneOperations.get
compute.zoneOperations.list
compute.zones.get
compute.zones.list
```

```
iam.serviceAccounts.actAs
resourcemanager.projects.get
```

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 88.

   For more information on creating and managing the custom roles, see Creating and managing custom roles section of Google documentation.

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.

   For more information on creating and managing service accounts, see Creating and managing service accounts section of Google documentation.

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 86.

   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": "some-product",
  "private_key": "-----BEGIN PRIVATE KEY-----\n
N11EvA18ADAN89kq4k199w08AQEFAA5C8KYw9951A9EAAo18AQCnvpuJ3oK974z4\n
.
.
.
weT9odE4ryl81tNU\nV3q1XNX4fK55QTpd6CNu+f7QjEw5x8+5ft05DU8ayQcNkX\n
4pXJoDol54N52+T4qV4WkoFD5uL4NLPz5wxf1y\nNWcNfru8K8a2q1/9o0U+99==\n
-----END PRIVATE KEY-----\n",
  "client_email": "email@xyz-product.iam.gserviceaccount.com",
  "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 NetBackup 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-----\
N11EvA18ADAN89kq4k199w08AQEFAA5C8KYw9951A9EAAo18AQCnvpuJ3oK974z4
.
.
.
weT9odE4ryl81tNU\nV3q1XNX4fK55QTpd6CNu+f7QjEw5x8+5ft05DU8ayQcNkX
4pXJoDol54N52+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 4-6 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.

Table 4-6 Microsoft Azure plug-in configuration parameters *(continued)*

CloudPoint configuration parameter	Microsoft equivalent term and description
Regions	One or more regions in which to discover cloud assets. **Note:** If you configure a government cloud, select US Gov Arizona, US Gov Texas US, or Gov Virginia.
Resource Group prefix	The string with which you want to append all the resources in a resource group.
Protect assets even if prefixed Resource Groups are not found	The check box determines whether the assets are protected if they are not associated to any resource groups. The prefixed Resource Group must exist in the same region as the source asset's Resource Group.

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 snapshot operations for Ultra SSD disk types in an Azure environment. Even though CloudPoint discovers the ultra disks successfully, any snapshot operation that is triggered on such disk assets fails with the following error:

  ```
  Snapshots of UltraSSD_LRS disks are not supported.
  ```

- 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.
- When you create snapshots, the Azure plug-in creates an Azure-specific lock object on each of the snapshots. The snapshots are locked to prevent unintended deletion either from the Azure console or from an Azure CLI or API call. The lock object has the same name as that of the snapshot. The lock object also includes a field named "`notes`" that contains the ID of the corresponding VM or asset that the snapshot belongs to.

  You must ensure that the "`notes`" field in the snapshot lock objects is not modified or deleted. Doing so will disassociate the snapshot from its corresponding original asset.

The Azure plug-in uses the ID from the "`notes`" fields of the lock objects to associate the snapshots with the instances whose source disks are either replaced or deleted, for example, as part of the 'Original location' restore operation.

- Azure plug-in supports the following GovCloud (US) regions:
  - US Gov Arizona
  - US Gov Texas
  - US Gov Virginia
- CloudPoint Azure plug-in does not support the following Azure regions:

Location	Region
US	■ US DoD Central ■ US DoD East ■ US Sec West
China CloudPoint does not support any regions in China.	■ China East ■ China East 2 ■ China North ■ China North 2
Germany	■ Germany Central (Sovereign) ■ Germany Northeast (Sovereign)

- Microsoft Azure gen2 type of virtual machines are not supported. Ensure that you use a gen1 type image to create a VM.
- CloudPoint does not support application consistent snapshots and granular file restores for Windows systems with virtual disks or storage spaces that are created from a storage pool. If a Microsoft SQL server snapshot job uses disks from a storage pool, the job fails with an error. But if a snapshot job for virtual machine which is in a connected state is triggered, the job might be successful. In this case, the file system quiescing and indexing is skipped. The restore job for such an individual disk to original location also fails. In this condition, the host might move to an unrecoverable state and requires a manual recovery.

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.Sql/*/read",
"Microsoft.Compute/disks/write",
"Microsoft.Compute/disks/delete",
"Microsoft.Compute/disks/beginGetAccess/action",
"Microsoft.Compute/disks/endGetAccess/action",
"Microsoft.Compute/images/write",
"Microsoft.Compute/images/delete",
"Microsoft.Compute/snapshots/delete",
"Microsoft.Compute/snapshots/write",
"Microsoft.Compute/snapshots/beginGetAccess/action",
"Microsoft.Compute/snapshots/endGetAccess/action",
"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.Compute/virtualMachines/powerOff/action",
"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/locks/*",
"Microsoft.Authorization/*/read" ],
"NotActions": [ ],
"AssignableScopes": [
"/subscriptions/subscription_GUID",
"/subscriptions/subscription_GUID/ \
resourceGroups/myCloudPointGroup" ] }
```

If CloudPoint extension is installed on a managed Kubernetes cluster in Azure, then the following permissions can also be added before configuring the plugin:

```
"Microsoft.ContainerService/managedClusters/agentPools/read",
"Microsoft.ContainerService/managedClusters/read",
"Microsoft.Compute/virtualMachineScaleSets/write",
"Microsoft.Compute/virtualMachineScaleSets/delete/action
```

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 92.

About Azure snapshots

NetBackup 9.0 introduces incremental snapshots in Azure. NetBackup creates the incremental snapshots for new changes to the disks, since the previous snapshot. The snapshots are independent of each other, for example, deletion of one snapshot, does not affect the subsequent snapshot that NetBackup creates. The incremental snapshots significantly reduce the cost of backup by reducing the required disk space, and using the Azure Standard HDD as storage, instead of Premium HDD.

Microsoft Azure Stack Hub plug-in configuration notes

The Microsoft Azure Stack Hub plug-in lets you create, delete, and restore snapshots at the virtual machine level and the managed disk level. You can configure the Azure Stack Hub plugin using AAD or ADFS authentication methods.

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

- Use the Microsoft Azure Stack Portal to create an application in the Azure Active Directory (AAD) if using AAD as the identify provider for the Azure Stack Hub plug-in.
  For more information on your identity provider options, refer to the following Azure Stack documentation:
  https://docs.microsoft.com/en-us/azure-stack/operator/azure-stack-identity-overview?view=azs-2008

- Assign the service principal to a role that has access to the resources.

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

https://docs.microsoft.com/en-us/azure-stack/operator/azure-stack-create-service-principals

Table 4-7 Azure Stack Hub plug-in configuration parameters using AAD

CloudPoint configuration parameter	Microsoft equivalent term and description
Azure Stack Hub Resource Manager endpoint URL	The endpoint URL in the following format, that allows CloudPoint to connect with your Azure resources. `https://management.<location>.<FQDN>`
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.
Authentication Resource URL (optional)	The URL where the authentication token is sent to.

Table 4-8 Azure Stack Hub plug-in configuration parameters using AD FS

CloudPoint configuration parameter	Microsoft equivalent term and description
Azure Stack Hub Resource Manager endpoint URL	The endpoint URL in the following format, that allows CloudPoint to connect with your Azure resources. `https://management.<location>.<FQDN>`
Tenant ID	The ID of the AD FS directory in which you created the application.
Client ID	The application ID.
Secret Key	The secret key of the application.
Authentication Resource URL (optional)	The URL where the authentication token is sent to.

Azure Stack Hub plug-in limitations

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

- CloudPoint currently only supports creating and restoring snapshots of Azure Stack managed disks and the virtual machines that are backed up by managed disks.
- CloudPoint currently only supports creating and restoring snapshots of Azure Stack managed disks and the virtual machines that are deployed using Azure Stack Resource Manager deployment model.
- Rollback restore operation is not supported for Azure Stack VM, because the OS disk swap not supported.
- Disk encryption is not possible with the CloudPoint Azure Stack Hub plug-in, because Azure Stack Hub 2008 does not support disk encryption.
- 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 operations for Ultra SSD disk types in an Azure Stack environment.

Azure Stack Hub plug-in considerations

- 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.
- When you create snapshots, the Azure Stack Hub plug-in creates an Azure Stack-specific lock object on each of the snapshots. The snapshots are locked to prevent unintended deletion either from the Azure console or from an Azure CLI or API call. The lock object has the same name as that of the snapshot. The lock object also includes a field named "`notes`" that contains the ID of the corresponding VM or asset that the snapshot belongs to.
  You must ensure that the "`notes`" field in the snapshot lock objects is not modified or deleted. Doing so will disassociate the snapshot from its corresponding original asset.
  The Azure Stack Hub plug-in uses the ID from the "`notes`" fields of the lock objects to associate the snapshots with the instances whose source disks are either replaced or deleted, for example, as part of the 'Original location' restore operation.

Configuring permissions on Microsoft Azure Stack Hub

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

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

- Configure Azure Stack Hub 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 Stack Hub plug-in operations in CloudPoint",
"Actions": [
"Microsoft.Storage/*/read",
"Microsoft.Storage/storageAccounts/listKeys/action",
"Microsoft.Storage/storageAccounts/ListAccountSas/action",
"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/snapshots/beginGetAccess/action",
"Microsoft.Compute/snapshots/endGetAccess/action",
"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.Compute/virtualMachines/powerOff/action",
"Microsoft.Authorization/locks/*",
"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 Stack documentation:

https://docs.microsoft.com/en-us/azure-stack/operator/azure-stack-registration-role?view=azs-2008

For example:

```
New-AzRoleDefinition -InputFile "C:\CustomRoles\registrationrole.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/
registrationrole.json"
```

Note: Before creating a role, you must copy the role definition (text in JSON format) in a .json file and then use that file as the input file. In the sample command displayed earlier, `registrationrole.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 Stack environment.
- In CloudPoint, configure the Azure Stack off-host plug-in with the application's credentials.

See "Microsoft Azure Stack Hub plug-in configuration notes" on page 97.

Configuring staging location for Azure Stack Hub VMs to restore from backup

The Azure Stack Hub requires you to create a container, inside your storage account, and use it as a staging location when you restore from backup images. The staging location is used to stage the unmanaged disks in the container during restores. Once the data is written to the disk, the disks are converted to managed disks. This is a requirement from the Azure Stack Hub platform. This is a mandatory configuration, before you can use Azure Stack Hub with NetBackup.

The `azurestack.conf` file should contain staging location details of the subscription ID, where the VMs are restored. If you plan to restore to any target subscription ID, other than the source subscription ID, then details of the target subscription ID must be present in the `azurestack.conf` file.

If you are using snapshot images for restore, you do not need to create this staging location.

Note: The staging location is specific to the subscription ID, you must create one staging location for each subscription that you are using to restore VMs.

To configure a staging location for a subscription ID:

1. In the CloudPoint server, navigate to:

   `/cloudpoint/azurestack.conf`, and open the file in a text editor. This file is created, only after you have added Azure Stack Hub as a cloud service provider in NetBackup.

2. Add the following details in the file:

   [subscription/`<subscription ID>`]

   storage_container = `<name of the storage container>`

   storage_account = `/resourceGroup/<name of the resource group where the storage account exists>/storageaccount/<name of storage account>`

   For example:
   `/resourceGroup/Harsha_RG/storageaccount/harshastorageacc`

3. Repeat step 2, for each subscription ID that you are using. Save and close the file.

Chapter 5

CloudPoint storage array plug-ins

This chapter includes the following topics:

- Hitachi (HDS VSP 5000) plug-in configuration notes
- InfiniBox plug-in configuration notes
- Dell EMC PowerScale (Isilon) plug-in configuration notes
- Dell EMC PowerMax and VMax plug-in configuration notes
- Qumulo plug-in configuration notes

How to configure the CloudPoint storage array plug-ins?

CloudPoint plug-ins are software modules that enable the discovery of your assets in the cloud or in an on-premise environment. After registering the CloudPoint server with the NetBackup primary server, you must configure the CloudPoint plug-ins to be able to protect your workloads using NetBackup.

How you configure the plug-ins depends on the asset type and how CloudPoint is deployed. If the CloudPoint server is deployed on-premise and you want to protect storage arrays, you must use the NetBackup Administration Console (Java UI) to register the CloudPoint server and configure the storage array plug-ins. The overall steps to configure the plug-ins are similar, regardless of the asset type. Only the configuration parameters vary.

Refer to the *NetBackup Snapshot Client Administrator's Guide* for information on how to configure storage plug-ins.

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.
- NetApp NFS volumes in a NAS environment.
- NetApp Storage Virtual Machines (SVM) that allow NAS clients to access storage using NFS protocols.

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.

- Ensure that a supported ONTAP version is installed on the NetApp arrays. CloudPoint supports the following:
  - ONTAP version 8.3 and later
- 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 will use to configure the plug-in has the privileges to perform the following operations on the NetApp array:
  - create snapshot
  - delete snapshot
  - restore snapshot
- Ensure that the NetApp user account that you will use to configure the plug-in is configured with `http` and `ontapi` access methods.
- Ensure that the NetApp user account that you will use to configure the plug-in has the following roles assigned:
  - Default: readonly
  - lun: all
  - volume snapshot: all
  - vserver export-policy: all

  Refer to the NetApp documentation for instructions on how to create users and roles, and assign permissions.

See "NetApp plug-in configuration parameters" on page 106.

See "Supported CloudPoint operations on NetApp storage" on page 107.

NetApp plug-in configuration parameters

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

Table 5-1 NetApp plug-in configuration parameters

CloudPoint configuration parameter	Description
Array IP address or FQDN	The cluster management IP address or the Fully Qualified Domain Name (FQDN) of the NetApp storage array or filer.

Table 5-1 NetApp plug-in configuration parameters *(continued)*

CloudPoint configuration parameter	Description
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.

Configuring a dedicated LIF for NetBackup access

NetApp NAS-based volume snapshots are exposed to NetBackup over NAS protocols. NetBackup reads these snapshots using any available Data LIF on the respective Storage Virtual Machines (SVM). If required, you can configure a Data LIF that is dedicated for NetBackup access.

While configuring a Data LIF, use the prefix "**nbu_nas_**" in the interface name for the SVM. If such a Data LIF exists, NetBackup automatically uses only that LIF for accessing the snapshots.

Note: This is an optional step. If configured, the backup reads are restricted via the dedicated LIF. If not configured, volume snapshots are accessed via any available DATA LIF of the corresponding SVM.

Supported CloudPoint operations on NetApp storage

CloudPoint performs the following management operations on the NetApp storage arrays:

Table 5-2 CloudPoint operations on NetApp storage

CloudPoint operation	Description
Discover assets	■ 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. `["state":"online","vol_type":"rw","is_snapshot_auto_delete_enabled": "false"]` **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 and volumes with security style UNIX and mixed mode on the NetApp storage. The shares must have an active `junction_path` configured so that CloudPoint can discover them.
Create snapshot	■ In a SAN deployment, CloudPoint takes a snapshot of the NetApp LUNs. When CloudPoint triggers a LUN snapshot on the NetApp storage, it internally triggers a redirect-on-write (ROW) 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. A typical snapshot created by CloudPoint has the following naming convention: `NB<unique_21digit_number>` ■ In a NAS deployment, CloudPoint takes a snapshot of the NetApp NFS shares.
Delete snapshot	■ 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 5-2 CloudPoint operations on NetApp storage *(continued)*

CloudPoint operation	Description
Restore snapshot	■ In a SAN deployment, when you restore a LUN from a snapshot, CloudPoint only restores the particular LUN on which the restore is triggered. The LUN snapshot is a ROW 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 volume using the specified snapshot.
Export snapshot	■ 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: ■ 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:** CloudPoint does not modify the SVM's "**default**" export policy. The export operation will fail if the volume is attached only to the "**default**" export policy on NetApp. You must assign the NAS volume to a non-default export policy.
Deport snapshot	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. In a NAS deployment, when a snapshot deport operation is triggered, NetBackup deletes the new rule that was created in the export policy when the snapshot was exported.

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.

ACL configuration on NetApp array

To configure ACL on NetApp array:

1. Logon to the OnCommand System Manager console.
2. Navigate to the respective SVM where you are creating the SMB volumes or shares.
3. Click the SVM setting in right pane.
4. Click the **Windows** under **Host Users and Groups** in left navigation pane.
5. **Groups** and **Users** tabs opens in the right pane.
6. In the **Groups** tab click the **BUILTIN\Backup Operators** and select **Edit** option at top
7. In the **Modify** dialog, in the **Members** frame, add your domain user and select the following Privileges: SetBackupPrivilege,SetRestorePrivilege, and SetSecurityPrivilege.

Nutanix Files plug-in configuration notes

Veritas NetBackup provides a robust data protection solution for shares that are set up on a Network Attached Storage (NAS) storage host. NetBackup extends this NAS support and allows you to protect file services that are hosted in a Nutanix Files environment. You can configure CloudPoint to discover and then perform backup and restore operations on Nutanix Files shares that are exposed as Network File System (NFS) exports.

The CloudPoint plug-in for Nutanix Files contains the necessary functional logic that enables NetBackup to discover the shares on the Nutanix Files server and then trigger snapshot create, export, deport, and delete operations for those shares. You must configure this plug-in on the NetBackup primary server.

CloudPoint uses the Nutanix REST APIs to communicate with the Nutanix Files File Server. CloudPoint establishes a connection with Nutanix Files File Server by registering itself as a backup application and then uses the API endpoints to discover the shares and their snapshots that need to be backed up.

Nutanix Files plug-in configuration prerequisites

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

- Ensure that a supported version of Nutanix Files is installed on the Nutanix arrays.
  CloudPoint supports the following:
  Nutanix Files version 3.6.1.3 and later
- Gather the following information about the Nutanix Files cluster. You will use these details while configuring the Nutanix Files plug-in:

Parameter	Description
Nutanix Files File Server FQDN	The Fully Qualified Domain Name (FQDN) of the Nutanix Files File Server.
REST API username	The user account that has the permissions to invoke the Nutanix Files REST APIs on the File Server.
REST API password	The password of the Nutanix REST API user account specified earlier.

The following screen is displayed when you configure the plug-in using the NetBackup administration console:

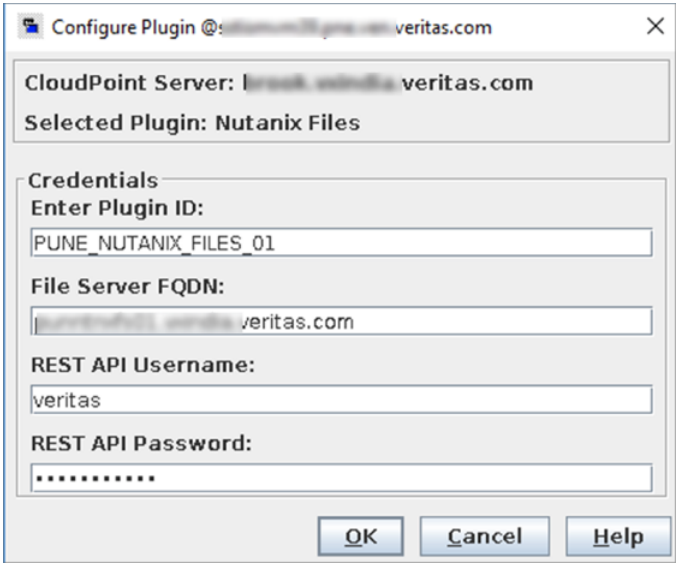

Nutanix Files plug-in considerations and limitations

The following considerations and limitations are applicable:

- Snapshot operations are not supported for nested shares on Nutanix Files File Server.
  A nested share is a share that is itself a sub-directory in an existing file share. NetBackup does not support snapshot creation for such nested shares.
- Nutanix Files File Server does not support point-in-time (PIT) rollback restore of shares using snapshots. You can use NetBackup assisted restore of shares' data.
- The maximum snapshot limit for a Nutanix Files shares is 20.
  The maximum snapshot limit defines the maximum number of policy-triggered snapshots that are retained for the specified share. When the maximum count is reached, the next snapshot that is created by the policy results in the deletion of the oldest snapshot.
  You may want to consider the policy schedule and retention for NetBackup's policy protecting Nutanix File shares.

Supported CloudPoint operations on Nutanix Files File Server

CloudPoint performs the following management operations on the Nutanix Files File Server:

Table 5-3 CloudPoint operations on Nutanix Files File Server

CloudPoint operation	Description
Discover assets	CloudPoint discovers all the shares and their snapshots along with some of their metadata. Shares that have CFT_BACKUP capabilities are eligible for snapshot diff based incremental backups. **Note:** Snapshot operations are not supported for nested shares on Nutanix Files File Server.
Create snapshot	To create a snapshot, CloudPoint triggers a POST REST API call on the `/mount_targets` API with the required share information and snapshot name. The API returns the details of the snapshot (also referred as the mount target snapshot). CloudPoint keeps polling the snapshot details until the snapshot state changes to successful (or error in case failure).
Delete snapshot	To delete a snapshot, CloudPoint triggers a DELETE REST API call with the required snapshot details in the following format: `/mount_target_snapshot/:snapshot_uuid` CloudPoint keeps polling the snapshot UUID until it returns a 404 Not Found error code. This code confirms that the snapshot has been deleted successfully.
Restore snapshot	CloudPoint does not support this operation.
Export snapshot	When a snapshot export operation is triggered, the backups host is added to the partner server that is registered during the plug-in configuration. A PUT REST API call is made to the partner server with the required mount target details. CloudPoint keeps polling the partner server to confirm the success of the operation.
Deport snapshot	When a snapshot deport operation is triggered, CloudPoint makes a PUT REST API call to the partner server to remove the mount target entry that was added during the export operation. CloudPoint keeps polling the partner server to confirm the success of the operation.

Table 5-3 CloudPoint operations on Nutanix Files File Server *(continued)*

CloudPoint operation	Description
Create snapshot diff	Nutanix Files provides an API that allows to create a diff between two snapshots of a share. This process is called as Changed File Tracking (CFT). When a request to create a snapshot diff is made, CloudPoint makes a REST API call to generate the CFT between two snapshots, and then retrieves and stores the CFT data on the CloudPoint server. CFT based backups are supported only for top-level shares. Nested shares are not supported.

Troubleshooting NetBackup issues for Nutanix Files

Refer to the following:

Backup jobs for Nutanix Files fail due to snapshot import and export operations failures

Backup jobs that are scheduled for file shares on Nutanix Files may fail due to a conflict error in the snapshot import and export operations.

The job log contains the following errors:

```
Snapshot import failed (4213)
Backup from Snapshot job failed with error 4213
Snapshot import failed
(errMsg": "Failed to export Error: Edit conflict: please retry change)

WARNING: Snapshot export failed.
Failed to export. Error: Edit conflict: please retry change.
Error vfms Snapshot export API failed for snapshot ID[snapID].
```

Recommended action:

This issue occurs if the same Nutanix Files file system is configured with more than one CloudPoint server instances simultaneously.

NetBackup is registered as a partner server on the Nutanix Files platform. A one to one mapping exists between the NetBackup CloudPoint server and the Nutanix Files. If the same Nutanix Files file system is configured with multiple CloudPoint instances, it creates a resource conflict. Each CloudPoint server attempts to update the configuration with the backup job information. This concurrent configuration update on the single partner server registration fails and causes a conflict error.

NetBackup does not support such a mixed configuration. Ensure that you configure Nutanix Files with a single instance of the CloudPoint server in the NetBackup domain.

Plug-in configuration may fail if the Nutanix Files version is unsupported

The Nutanix Files plug-in configuration may fail with a http 500 status code and the following error message is displayed:

```
Minimum supported AFS version 3.6.1.3
```

This issue occurs if the Nutanix Files version in use is not supported by CloudPoint. Ensure that a supported version of Nutanix Files is installed before you configure the plug-in.

See "Nutanix Files plug-in configuration prerequisites" on page 111.

Configuring ACL for Nutanix array

To configure ACL for Nutanix array:

1. Logon to the prism console.
2. Open the file servers list, and click the file server where you want to create your SMB shares.
3. Select **User Mapping** in the **Protocol Management** link in the right corner.
4. Click **Next** multiple times, till the **Explicit Mapping** dialog appears.
5. Click **Add One to One Mapping** and add your domain user and add NFS ID, save and click **Next**.
6. You must add one domain user to the default mapping. Save the details.
7. Click **Manage Roles** in right pane for selected file server.
8. Add your domain user in the **Add Admins** section and select **Role** as **Backup admin: Backup access only**
9. Save and close the dialog.

Dell EMC Unity array plug-in configuration notes

The CloudPoint plug-in for Dell EMC Unity array plug-in also supports the Network Attached Storage (NAS) storage host and allows you to protect Network File System (NFS) and Server Message Block (SMB) exports that are hosted in a EMC Unity

array environment. You can configure CloudPoint to discover and then perform backup and restore operations on NFS and SMB exports.

The plug-in enables NetBackup to discover the NFS exports on the EMC Unity array and then trigger snapshot create, export, deport, and delete operations for those exports. You must configure this plug-in on the NetBackup primary server.

CloudPoint uses the REST API SDK of Storops Version 1.2.8 to communicate with the EMC Unity assets. CloudPoint establishes a connection with EMC Unity Array by using RestClient library exposed by SDK and then uses the SDK methods to discover the NFS exports and their snapshots that need to be backed up.

Dell EMC Unity array plug-in configuration parameters

The following parameters are required when you configure the Dell EMC Unity array plug-in:

Table 5-4 Dell EMC Unity array plug-in configuration parameters

NetBackup configuration parameter	Description
Array IP address	Array IP address that you want to be protected. Both, IPV6 and IPV4 settings are supported.
Username	A user account name that has permissions to perform snapshot operations on the EMC Unity Array Ensure that the specified user account has permissions to create, delete, and restore snapshots on the array
Password	The password of the EMC unity Array user account specified earlier.

The following screen is displayed when you configure the plug-in using the NetBackup administration console:

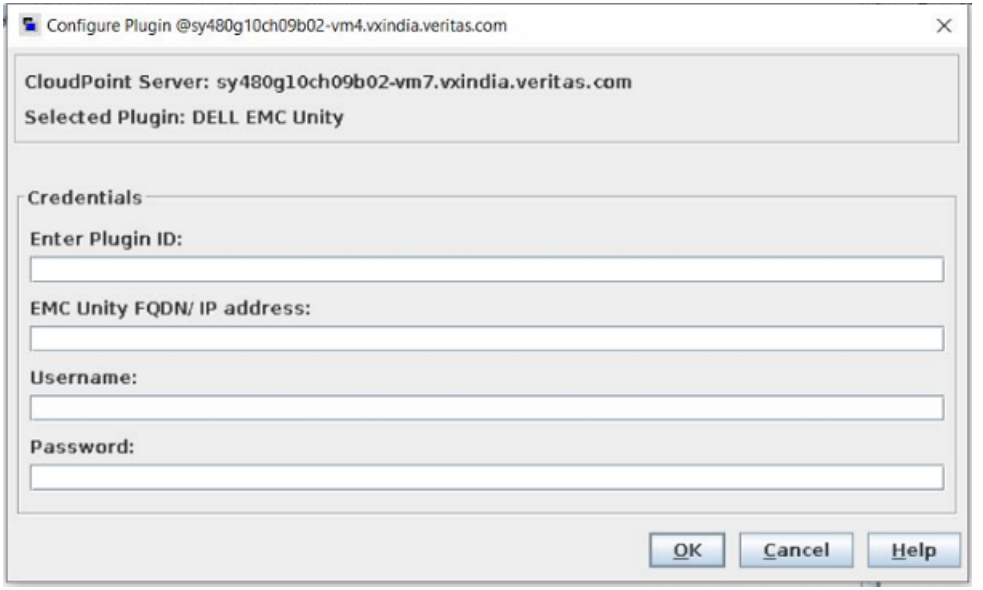

Dell EMC Unity plug-in considerations and limitations

- The array does not support share level snapshots. So the snapshots are created at a file level and they have a 1:1 mapping with each other.
- For every backup operation triggered on NetBackup a new share of the type snapshot will be created on the array and this will be automatically cleared once the Deport operation (Expire) is triggered.

Supported Dell EMC Unity arrays

You can use CloudPoint to discover and protect the following Dell EMC Unity array models.

Table 5-5 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: ■ 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

CloudPoint performs the following management operations on the Dell EMC Unity arrays.

Table 5-6 CloudPoint operations on Dell EMC Unity plug-in

CloudPoint operation	Description
Discover assets	CloudPoint discovers all the volumes and their snapshots along with their storage group. **Note:** CloudPoint only discovers assets with depth as 2.

Table 5-6 CloudPoint operations on Dell EMC Unity plug-in *(continued)*

CloudPoint operation	Description
Create snapshot	To create a snapshot, CloudPoint triggers an SDK method on the storage group within which the volumes reside, with the required information and snapshot name. A typical snapshot created by CloudPoint has the following naming convention: NB<unique_21digit_number>
Delete snapshot	To delete a snapshot, CloudPoint triggers an SDK method with the required snapshot details and confirms that the snapshot has been deleted successfully on the array.
Restore snapshot	CloudPoint offers the ability to restore with the help of SDK methods with different restore paths.
Export snapshot	When a snapshot export operation is triggered, a new NFS export is created over the same filesystem path, on which the backups host is added as a client with read-only permissions
Deport snapshot	When a snapshot deport operation is triggered, CloudPoint deletes the exported storage group created over the snapshot path and the volume inside it, and the temporary storage group that is used as a source. It essentially reverts the snapshot export operation.

You can also 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.

- 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.

Dell EMC PowerStore plug-in configuration notes

Veritas NetBackup provides a robust data protection solution for volumes that are set up on the storage array. NetBackup extends the SAN support and allows you to protect mounted iSCSI/FC volumes that are hosted on the Dell EMC PowerStore array environment. You can configure CloudPoint to discover data, perform backups, and restore operations.

Dell EMC PowerStore contains the functional logic which enables NetBackup to discover the SAN volumes on the Dell EMC PowerStore array. Then the plugin triggers the snapshot to create, export, deport, and delete operations for the exports. You must configure this plug-in on the NetBackup primary server.

CloudPoint uses the REST API supported by Dell EMC PowerStore family to communicate with the Dell EMC PowerStore assets. CloudPoint establishes a connection with Dell EMC PowerStore array by using RestClient and then uses the SDK methods to discover the SAN volumes and their snapshots that needs to be backed up.

Dell EMC PowerStore plug-in configuration parameters

Specify the following parameters when you configure the Dell EMC PowerStore plug-in:

Table 5-7 Dell EMC PowerStore plug-in configuration parameters

CloudPoint configuration parameter	Description
Plugin ID	Provide a name for the plugin.
FQDN/ IP Address	The array's IP address, in IPV4 / FQDN format.
Username	A user account that has permissions to perform snapshot operations on the Dell EMC PowerStore array.
Password	Provide a password to the user account.

The following screen is displayed when you configure the plug-in using the NetBackup administration console:

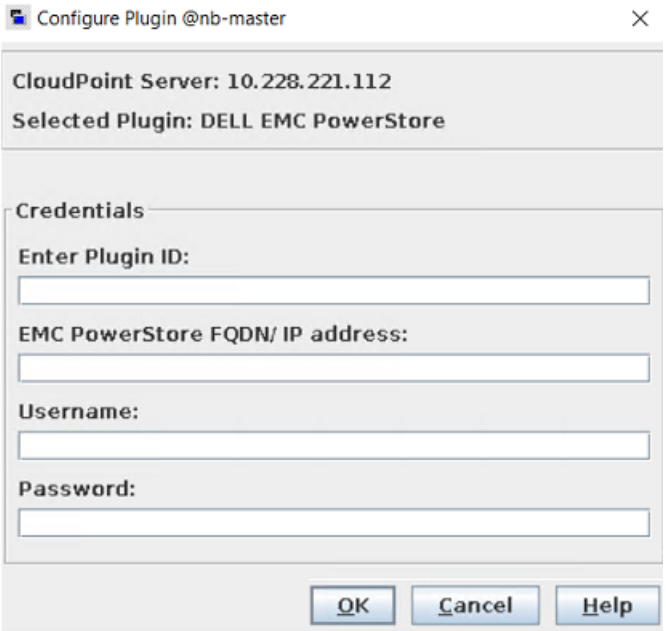

Dell EMC PowerStore plug-in considerations and limitations

The following considerations and limitations are applicable:

1. CloupPoint does not discover the clone type volumes in the discovery process.
2. When a snapshot is triggered, a volume is cloned from that snapshot. Then, the volume is cleaned up after the image is expiration happens from Netback.
3. There is no retention expiry time for the cloned volumes. The cloned volumes can only be manually deleted from NetBackup in deport and delete operations.

Supported CloudPoint operations on Dell EMC PowerStore models

You can perform the following CloudPoint operations supported on the Dell EMC PowerStore models:

Table 5-8 CloudPoint operations on the Dell EMC PowerStore array

CloudPoint operations	Description
Discover assets	CloudPoint discovers all the array volumes and snapshots inside the snapshot group flexsnap_snap_group with some metadata. The volumes which are clone types in the respective attributes and without mapping are not discovered.
Create snapshot	To create a snapshot, CloudPoint triggers a SDK method with the required snapshot details. The API returns the details of the snapshot. A typical snapshot created by CloudPoint has the following naming convention: NB<unique_21digit_number>.
Delete snapshot	To delete a snapshot, CloudPoint triggers a SDK method call with the required snapshot details. Then confirms that the snapshot is deleted successfully on the array.
Restore snapshot	CloudPoint offers the ability to restore the snapshots using the SDK method with the different restore paths.
Export snapshot	CloudPoint supports export snapshot over the iSCSI and FC. It uses the SDK to set the LUN path for snapshot.
Deport snapshot	When a snapshot deport operation is triggered, CloudPoint deletes the export mapping created between the host and the volume.

Dell EMC PowerStore NAS plug-in configuration notes

Veritas NetBackup provides a robust data protection solution for filesystems set on the storage host. NetBackup extends the NAS support to allow you to protect

filesystems based on the protocols (NFS and SMB) that are hosted in a Dell EMC PowerStore array environment.

You can configure CloudPoint to discover data, perform backup, and then restore operations on NFS and SMB filesystems.

Dell EMC PowerStore contains functional logic which enables NetBackup to discover the filesystems on the Dell EMC PowerStore array. Then triggers snapshot create, export, deport, and delete operations for filesystems.

You must configure the plug-in on the NetBackup primary server.

1. CloudPoint uses the SDK exposed by the Dell EMC called as python-powerstore (1.4.0).
2. Python-PowerStore calls the Dell EMC PowerStore family APIs to communicate and get Dell EMC PowerStore assets.
3. A connection is established to Dell EMC PowerStore array through REST and then the SDK methods are used to discover the filesystems and the snapshots to be backed up.

Dell EMC PowerStore NAS plug-in configuration parameters

Specify the following parameters when you configure the Dell EMC PowerStoreNAS plug-in:

Table 5-9 Dell EMC PowerStore NAS plug-in configuration parameters

CloudPoint configuration parameter	Description
Plugin ID	Provide a name for the plugin.
FQDN/ IP Address	The array's IP address, in IPV4 / FQDN format.
Username	A user account that has permissions to perform snapshot operations on the Dell EMC PowerStore NAS array.
Password	Provide a password to the user account.

The following screen is displayed when you configure the plug-in using the NetBackup administration console:

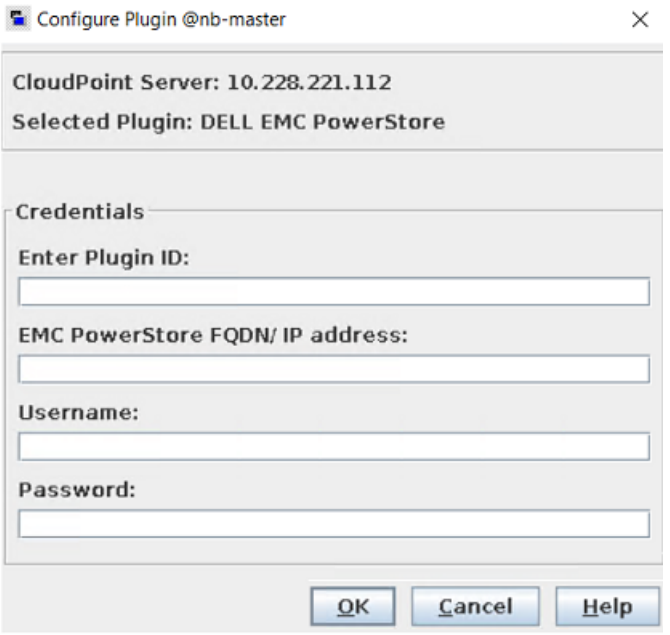

Dell EMC PowerStore NAS plug-in considerations and limitations

The following considerations and limitations are applicable:

1. All snapshots captured for filesystems are read-only mode. The host is appended as per the existing rules for a particular share.
2. The limit for any filesystem name is 128 characters on the array. In case of a snapshot copy, the maximum length you can have the volume name is 128 - 23(NB<unique_21digit_number>) = 103. A volume name must be limited to 94 characters for a successful snapshot capture.
3. You are expected to create at least one share before the backup of a filesystem. If not, the request will fail to map the resources for backup job.

Supported CloudPoint operations on Dell EMC PowerStore NAS models

You can perform the following CloudPoint operations supported on the Dell EMC PowerStore NAS models:

Table 5-10 CloudPoint operations on the Dell EMC PowerStore NAS array

CloudPoint operations	Description
Discover assets	CloudPoint discovers all the Dell EMC PowerStore NAS servers, filesystems, NFS exports, SMB shares as an asset in CP. CloudPoint triggers an SDK method which internally calls the array's API to get the assets mentioned in the list. For NAS discovery, CP doesn't skip any assets. If the current filesystem count is 50 and snapshot count is 21 then the user finds 50 directories and 21 Filesystems in the NetBackup.
Create snapshot	To create a snapshot, CloudPoint triggers an SDK method with the required information and snapshot name. The API returns with the snapshot details. The array supports two types of snapshots - protocol and snapshot type. CP triggers the snapshot type; snapshot name and retention period is not set on the array for these snapshots. All these snapshots are at a filesystem level. A typical snapshot created by CloudPoint has the following naming convention: NB<unique_21digit_number> **Note:** No other entity apart from this snapshot is created on the array as a snapshot activity.
Delete snapshot	To delete a snapshot, CloudPoint triggers an SDK method call with the required snapshot details. A confirmation of snapshot delete on the array happens through another call with the same snapshot ID and 'not found' error is reported in the CloudPoint logs. Items deleted in snapshot delete operation - NB<unique_21digit_number> snapshot
Restore snapshot	CloudPoint offers the capability for the PIT rollback and a user can use the snapshot created to restore the filesystem. **Note:** Latest snapshot is not required for the PIT. A user can perform the restore operation with old snapshots related to the filesystem.

Table 5-10 CloudPoint operations on the Dell EMC PowerStore NAS array *(continued)*

CloudPoint operations	Description
Export snapshot	CloudPoint offers the export for SMB and NFS based exports. When a user triggers an export call from NetBackup, user can specify the protocol to export and then the details are sent to CloudPoint. Example, For NFS the export path is created with - <NAS-server-ip>:<share_name>/snapshot/<UTC_for_snapshot> and the rules for the host are added as a read only on a particular share. For SMB, the shares are created using path \\<NAS-server-ip>\<share_name>\@UTC_for_snapshot and backup is performed. **Note:** Export rules are added for a particular host in read only mode.
Deport snapshot	Deport snapshot operation will remove the export rules added for host.

Dell EMC PowerFlex plug-in configuration notes

Veritas NetBackup provides a robust data protection solution for Volumes that are set up on the storage Array. NetBackup extends the SDS support and allows you to protect mounted volumes that are hosted on an Dell EMC PowerFlex array environment. You can configure CloudPoint to discover data, perform backups, and restore operations.

Dell EMC PowerFlex contains the functional logic which enables NetBackup to discover the volumes on the Dell EMC PowerFlex array. Then triggers the snapshot to create, export, deport, and delete operations for the exports. You must configure this plug-in on the NetBackup primary server.

CloudPoint uses the SDK supported by Dell EMC PowerFlex family to communicate with the Dell EMC PowerFlex assets. CloudPoint establishes a connection with Dell EMC PowerFlex array by using RestClient and then uses the SDK methods to discover the volumes and their snapshots that needs to be backed up.

Dell EMC PowerFlex plug-in configuration parameters

Specify the following parameters when you configure the Dell EMC PowerFlex plug-in:

Table 5-11 Dell EMC PowerFlex plug-in configuration parameters

CloudPoint configuration parameter	Description
Plugin ID	Provide a name for the plugin.
FQDN/ IP address	The array's IP address, in IPV4 / FQDN format.
Username	A user account that has permissions to perform snapshot operations on the Dell EMC PowerFlex array.
Password	Provide a password to the user account.

The following screen is displayed when you configure the plug-in using the NetBackup administration console:

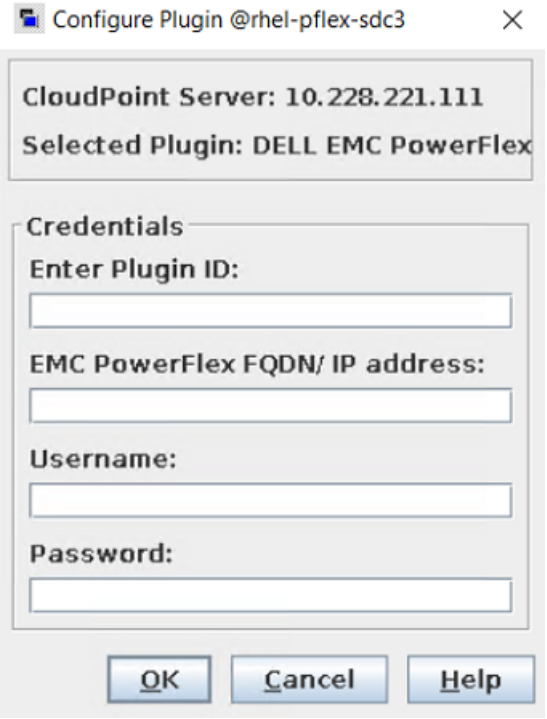

Dell EMC PowerFlex plug-in considerations and limitations

The following considerations and limitations are applicable:

1. This is a software defined storage, which requires to install the SDC (Storage Data Client) on the host where NetBackup will be configured.

2. The mapping between the volumes and SDC is completed with the help of SDC ID in CloudPoint.
3. The WWN (World Wide Name) is considered for mapping. It is developed using the $system_id$volume_id manner because it's not available directly on the array.

Supported CloudPoint operations on Dell EMC PowerFlex models

You can perform the following CloudPoint operations supported on the Dell EMC PowerFlex models:

Table 5-12 CloudPoint operations on the Dell EMC PowerFlex array

CloudPoint operations	Description
Discover assets	CloudPoint discovers all the array volumes and snapshots inside the snapshot group flexsnap_snap_group with some metadata. The volumes which have 'CMD' in the attributes and without mapping are not discovered.
Create snapshot	To create a snapshot, CloudPoint triggers an SDK method with the required snapshot details. The API returns the details of the snapshot. A typical snapshot created by CloudPoint has the following naming convention: NB<unique_21digit_number>.
Delete snapshot	To delete a snapshot, CloudPoint triggers an SDk method call with the required snapshot details. Then confirms that the snapshot is deleted successfully on the array.
Restore snapshot	CloudPoint offers the ability to restore the snapshots using a SDK method with the different restore paths.
Export snapshot	CloudPoint supports export snapshot over the SDC that is mapped on the parent volume.
Deport snapshot	When a snapshot deport operation is triggered, CloudPoint deletes the SDC mapping created between the host and the volume.

Dell EMC XtremIO SAN plug-in configuration notes

Veritas NetBackup provides a robust data protection solution for Volumes that are set up on the storage array. NetBackup extends SAN support and allows you to protect mounted iSCSI/FC volumes that are hosted on a Dell EMC XtremIO SAN

array environment. You can configure CloudPoint to discover the volumes and perform backup and restore operations.

The CloudPoint plug-in for Dell EMC XtremIO SAN contains the functional logic that enables NetBackup to discover the SAN volumes on the Dell EMC XtremIO SAN array. Then trigger snapshot create, export, deport, and delete operations for the exports. You must configure this plug-in on the NetBackup primary server.

CloudPoint uses the Rest API's exposed by Dell EMC XtremIO SAN family to communicate with SAN assets. It utilizes the latest API version V3 which is supported by the array after XMS version 6.0.1. Any firmware below XMS 6.0.1 will not be able to protect the volumes from NetBackup. The connection is a basic auth from the user used to configure the array in NetBackup.

Dell EMC XtremIO SAN plug-in configuration parameters

Specify the following parameters when you configure the Dell EMC XtremIO SAN plug-in:

Table 5-13 Dell EMC XtremIO SAN plug-in configuration parameters

CloudPoint configuration parameter	Description
Plugin ID	Provide a name for the plugin.
FQDN/ IP Address	The array's IP address, in IPV4 / FQDN format.
User name	A user account that has permissions to perform snapshot operations on the Dell EMC XtremIO SAN array.
Password	Provide a password to the user account.

The following screen is displayed when you configure the plug-in using the NetBackup Administration Console:

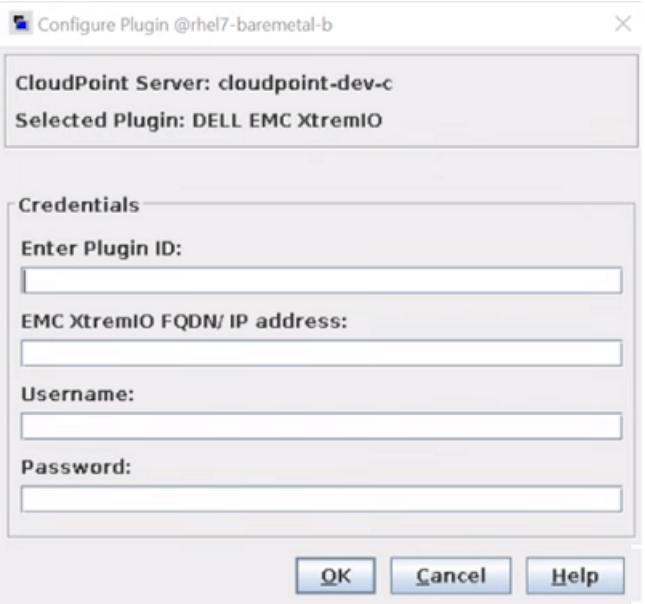

Dell EMC XtremIO SAN plug-in considerations and limitations

The following considerations and limitations are applicable:

1. All the snapshot which are taken for any volume are read-only mode. There is a new re-purpose copy cloned from the volume which is used for data backup.
2. The re-purpose clone copy is only done in the export operation not in the create snapshot directly.

   The repurpose clone copy is deleted based on the scheduled deport from NetBackup or manually triggering the expire operation. This is a thin provisioned type copy.
3. The limit for any volume name is 128 characters on the array. So, a user can have the maximum length of volume name is 128 - (23 (NB<unique_21digit_number>) - 9 (Repurpose) - 2 (Dot notations)) = 94. It is a strict requirement to limit the volume name to 94 characters for a successful snapshot creation.
4. It is strictly advised, not to write the data to repurpose copy created by NetBackup through manual mapping it on to the host server. User created re-purpose copies are considered as an individual volume in NetBackup. But changing the copies starting with

volume_name.NB<unique_21digit_number>.repurpose is strictly not recommended.

5. Don't refresh the repurpose copy, as it changes the data from the parent volume and backup, and restore operations.

   For example: If you have a volume V1 and a snapshot is triggered for the same then the protection copy occurs with the V1.NB<unique_21digit_number> format and in the export operation, a repurpose copy occurs with the name V1.NB<unique_21digit_number>.repurpose.

Supported CloudPoint operations on Dell EMC XtremIO SAN models

You can perform the following CloudPoint operations supported on the Dell EMC XtremIO SAN models:

Table 5-14 CloudPoint operations on the Dell EMC XtremIO SAN array

CloudPoint operations	Description
Discover assets	CloudPoint discovers all the array volumes and snapshots inside the snapshot group flexsnap_snap_group with some metadata. The volumes which have 'CMD' in the attributes and without mapping are not discovered.
Create snapshot	To create a snapshot, CloudPoint triggers a Post Rest API method with the required snapshot details. The API returns the details of the snapshot. A typical snapshot created by CloudPoint has the following naming convention: NB<unique_21digit_number>.
Delete snapshot	To delete a snapshot, CloudPoint triggers a Post Rest API method call with the required snapshot details. Then confirms that the snapshot is deleted successfully on the array.
Restore snapshot	CloudPoint offers the ability to restore the snapshots using a Post REST API method with the different restore paths.
Export snapshot	CloudPoint supports export snapshot over the iSCSI and FC. It uses a REST API to set the LUN path for snapshot.
Deport snapshot	When a snapshot deport operation is triggered, CloudPoint deletes the export mapping created between the host and the volume.

Pure Storage FlashArray plug-in configuration notes

Specify the following parameters when you configure the Pure Storage FlashArray plug-in:

Table 5-15 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 5-16 Supported Pure Storage FlashArray models

Category	Supported
Array model	FA-405
Firmware version	■ Software: Purity OS ■ Purity OS version: 5.1.4 ■ Rest Version: 1.11 Refer to the array-specific documentation for more information on firmware versions and how to check the current firmware on your array.

Supported CloudPoint operations on Pure Storage FlashArray models

You can perform the following CloudPoint operations on supported Pure Storage FlashArray models:

- Discover and list all volumes.
- Create a clone snapshot of a volume.

Note: A snapshot name comprises of "Diskname+ snapshotname". Snapshot suffix must be between 1 through 63 characters in length and can be alphanumeric. The snapshot name must begin and end with a letter or number. The suffix must include at least one letter or '-'.

- Delete a clone snapshot.
- Restore the original volume from a snapshot. The snapshot overwrites the original volume.
- Export a snapshot.
  When a snapshot export operation is triggered, CloudPoint creates a new volume from the snapshot and attaches it to the target host using the Fibre Channel (FC) protocol. The target host is assigned read-write privileges on the exported snapshot volume.
- Deport a snapshot.
  When a snapshot deport operation is triggered, CloudPoint detaches the exported snapshot volume from the target host and then deletes the volume.

Snapshot export related requirements and limitations

The following requirements and limitations are applicable for snapshot export and deport operations in a Pure Storage array environment:

- A snapshot cannot be exported multiple times.
- An exported snapshot cannot be deleted.

Pure Storage FlashBlade plug-in configuration notes

NetBackup provides a robust data protection solution for the file systems that are set up on the storage host. NetBackup extends the NAS support to protect the file systems based on the following protocols:

Network File System (NFS) and Server Message Block (SMB) hosted on the Pure Storage FlashBlade array environment.

You can configure CloudPoint to discover data, perform backup, and then restore operations on NFS and SMB file systems.

Pure Storage FlashBlade plug-in contains a functional logic which enables NetBackup to discover the file systems on the Pure FlashBlade array. Then the plug-in triggers the snapshot to create, export, deport, and delete operations for file systems.

Note: You must configure the Pure Storage FlashBlade plug-in on the NetBackup primary server.

1. CloudPoint uses the REST API SDK of purity-fb version 1.12.2 to communicate with the Pure Storage FlashBalade assets.
2. Using SDK exposed RestClient library CloudPoint establishes a connection with Pure Storage FlashBlade array.
3. Then uses the SDK methods to discover the file systems and its' snapshots which need to be backed up.

Pure Storage FlashBlade plug-in configuration parameters

Specify the following parameters when you configure the Pure Storage FlashBlade plug-in:

Table 5-17 Pure Storage FlashBlade plug-in configuration parameters

CloudPoint configuration parameter	Description
Plugin ID	Provide a name for the plugin.
IP Address	The array's IP address, in IPV4 format.
API Token	Provide the unique API token generated for a user.

Pure Storage FlashBlade plug-in considerations and limitations

You cannot create shares within the Pure Storage FlashBlade array, but you can create the file systems.

Supported CloudPoint operations on Pure Storage FlashBlade models

You can perform the following CloudPoint operations supported on the Pure Storage FlashBlade models:

Table 5-18 CloudPoint operations on the Pure Storage FlashBlade array

CloudPoint operations	Description
Discover assets	CloudPoint discovers all the Pure Storage FlashBlade share paths and its snapshots.

Table 5-18 CloudPoint operations on the Pure Storage FlashBlade array *(continued)*

CloudPoint operations	Description
Create snapshot	To create a snapshot, CloudPoint triggers the sdk method with the required snapshot name and details. The API returns with the snapshot details. A snapshot is created with the following naming convention: NB<unique_21digit_number>
Delete snapshot	To delete a snapshot, CloudPoint triggers the SDK method with the required snapshot details. Then confirms that the snapshot is deleted on the array.
Restore snapshot	CloudPoint restores the snapshot using the SDK methods with different restore paths.
Export snapshot	When an export snapshot is triggered, a new read-only rule is added to the host in the file system for which the snapshot is created. Before a snapshot is created an export path is generated using the vLan interface available on the array. Then the path is shared with Netbackup for mounting.
Deport snapshot	When a snapshot deport operation is triggered, CloudPoint deletes the NFS export rule created for the host.

IBM Storwize plug-in configuration notes

NetBackup provides a robust data protection solution for the volumes that are set up on the storage array.

NetBackup extends SAN support to protect the mounted iSCSI/FC volumes that are hosted on an IBM Storwize array environment.

You can configure CloudPoint to discover data, perform backup, and restore operations.

The CloudPoint plug-in for IBM Storwize contains the functional logic that enables NetBackup :

1. To discover the SAN volumes on the IBM Storwize array
2. Then triggers the snapshot to create, export, deport, and delete the operations for the exports.

Note: You must configure this plug-in on the NetBackup primary server.

1. CloudPoint uses the REST API supported by IBM Storwize family to communicate with the IBM Storwize assets.
2. Using SDK exposed RestClient library CloudPoint establishes a connection with IBM Storwize
3. Then uses the SDK methods to discover the SAN volumes and its' snapshots that needs to be backed up.

IBM Storwize plug-in configuration parameters

Specify the following parameters when you configure the IBM Storwize SAN plug-in:

Table 5-19 IBM Storwize plug-in configuration parameters

CloudPoint configuration parameter	Description
Plugin ID	Provide a name for the plugin.
FQDN/ IP Address	The array's IP address, in IPV4 / FQDN format.
Port	Port on which IBM Storwize is configured.
Username	A user account that has permissions to perform snapshot operations on the IBM Storwize array.
Password	Provide a password to the user account.

IBM Storwize plug-in considerations and limitation

The following considerations and limitations are applicable:

1. During the delete and deport operations the array is not allowed to perform the operations until the vDisk protection is disabled. So, NetBackup disables the vDisk protection when these operations are triggered and then return to the original state. These operations doesn't have any impact on the existing mapping or I/O operations performed on the array.
2. During the restore operation a standard warning about the FlashCopy mapping is displayed on the array. Then NetBackup restores the volume irrespective of the same behavior observed on the array side.
3. The array doesn't support ipV6 configuration from NetBackup and can only use ipV4 for all its' operations.

Supported CloudPoint operations on IBM Storwize models

You can perform the following CloudPoint operations supported on the IBM Storwize SAN models:

Table 5-20 CloudPoint operations on the IBM Storwize array

CloudPoint operations	Description
Discover assets	CloudPoint discovers all the array volumes and snapshots inside the snapshot group flexsnap_snap_group with some metadata. Volumes with the CMD attributes and without mapping are not discovered
Create snapshot	To create a snapshot, CloudPoint triggers a Post Rest API method with the required snapshot details. The API returns with the snapshot details. A snapshot is created with the following naming convention: NB<unique_21digit_number>
Delete snapshot	To delete a snapshot, CloudPoint triggers a Post Rest API method with the required snapshot details. Then confirms that the snapshot is deleted on the array.
Restore snapshot	CloudPoint restores the snapshot using a Post Rest API method with different restore paths.
Export snapshot	CloudPoint supports export snapshot over the iSCSI and FC. It uses a REST API to set the LUN path for snapshot.
Deport snapshot	When a snapshot deport operation is triggered, CloudPoint deletes the export mapping created between the host and the volume.

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 138.

See “Supported HPE storage systems” on page 138.

See “Supported CloudPoint operations on HPE storage arrays” on page 138.

RMC plug-in configuration parameters

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

Table 5-21 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 5-22 Supported RMC version

Category	Supported
RMC software version	■ 6.0 or later ■ 6.2 or later (for HPE Nimble)

Table 5-23 Supported RMC-managed storage systems

Category	Supported
Arrays	■ HPE 3PAR StoreServ ■ HPE Nimble Storage

Supported CloudPoint operations on HPE storage arrays

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

Table 5-24 CloudPoint operations on assets managed by HPE RMC

CloudPoint operation	Description
Discover assets	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. For snapshots, CloudPoint scans all the snapshot sets and links each snapshot to its originating parent volume.
Create snapshot	CloudPoint takes snapshots of all the volumes on the array. 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. For example, consider a volume set that contains three volumes, `vol-1`, `vol-2`, and `vol-3`. If you use CloudPoint to create a snapshot of `vol-1`, 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 `vol-1` (the selected volume) even though the snapshot set includes additional snapshots belonging to the other volumes, `vol-2,` and `vol-3`.
Delete snapshot	CloudPoint deletes the snapshot or the snapshot set (if parent volume is part of a volume set). 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.

Table 5-24 CloudPoint operations on assets managed by HPE RMC *(continued)*

CloudPoint operation	Description
Restore snapshot	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. Ensure that the parent volume is unmounted from the target host before initiating a snapshot restore.
Export snapshot	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. 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. The export operation is supported using the following protocols: ■ Fibre Channel (FC) ■ Internet Small Computer Systems Interface (iSCSI)
Deport snapshot	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.

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 NetBackup to delete snapshots that are not created using CloudPoint.
- NetBackup 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.

HPE XP plug-in configuration notes

The CloudPoint plug-in for HPE XP (XP7 and XP8) enables NetBackup to discover SAN volumes on the HPE XP array and then trigger snapshot create, export, deport, delete, and restore operations for them. You must configure this plug-in on the NetBackup primary server.

CloudPoint uses the REST API hosted on HPE XP Configuration Manager to communicate with the HPE XP storage arrays. It establishes a connection with HPE XP storage array by creating sessions in HPE XP Configuration Manager and uses the REST APIs to discover the SAN volumes and their snapshots that need to be backed up.

See “RMC plug-in configuration parameters” on page 138.

See “Supported HPE storage systems” on page 138.

See “Supported CloudPoint operations on HPE storage arrays” on page 138.

HPE XP plug-in configuration parameters

Before configuring the plug-in, make sure to create a pool with the name prefixed with`flexsnap_pool` on the storage array to store snapshots.

Gather the following information about the HPE XP array. You will need to use these details while configuring the plug-in.

Table 5-25 HPE XP plug-in configuration parameters

CloudPoint configuration parameter	Description
HPE XP Configuration Manager Server IP	IP of the HPE XP Configuration Manager REST server which is configured with the storage array to be used.
HPE XP Configuration Manager Server Port	Port on which the HPE XP Configuration Manager REST server is hosted.
Array Username	HPE XP Storage Array user account which have permissions for snapshot operations.
Array Password	The password associated with the array username.
Array Storage Device ID	Storage device ID of the array that is already registered with the HPE XP Configuration Manager.

The following screen is displayed when you configure the plug-in using the NetBackup administration console:

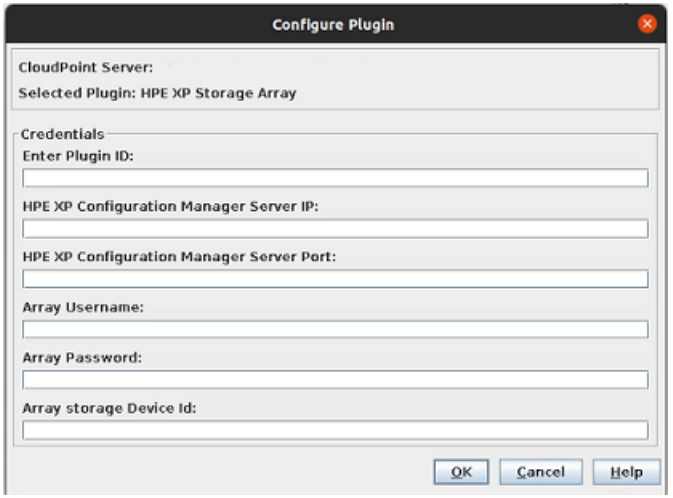

HPE XP plug-in considerations and limitations

- CloudPoint uses a snapshot group while creating snapshots, So maximum number of snapshots in CloudPoint for an array is 8192.
- The pool created must be large enough to accommodate all the snapshot needs.

Supported CloudPoint operations on HPE XP storage arrays

CloudPoint performs the following management operations on the HPE XP Storage Array:

Table 5-26 CloudPoint operations on assets managed by HPE XP

CloudPoint operation	Description
Discover assets	CloudPoint discovers the Logical Devices (LDEV) created on the storage array and snapshots inside the snapshot group named similar to `flexsnap_snap_group` along with some of their metadata. The LDEVs that have "CMD" in their attributes and those without any logical unit number (LUN) mapped are not discovered.
Create snapshot	For snapshots, CloudPoint uses HPE XP Fast Snap Pairs and triggers a sequence of REST API requests with the required information and snapshot name. The API returns the details of the snapshot. A typical snapshot created by CloudPoint has the following naming convention: `NB<unique_21digit_number>`
Delete snapshot	To delete a snapshot, CloudPoint triggers a sequence of REST API requests with the required snapshot details and confirms that the snapshot has been deleted successfully on the cluster.
Restore snapshot	To restore a snapshot, CloudPoint triggers a REST API request where the fast snap is restored to the parent volume.
Export snapshot	Export over iSCSI and FC is supported by CloudPoint. Cloudpoint uses REST API to set LUN path of snapshot.
Deport snapshot	When a snapshot deport operation is triggered, CloudPoint deletes the export created over the snapshot path at the time of Export operations. It essentially reverts the Export operation.

Hitachi plug-in configuration notes

The CloudPoint plug-in for Hitachi lets you create, delete, export, deport, and restore storage snapshots of a supported Hitachi storage array that is registered with Hitachi Configuration Manager (HCM). The plug-in supports the copy-on-write (COW) snapshot type.

Hitachi plug-in configuration prerequisites

Before you configure the Hitachi plug-in, perform the following steps on the storage system:

- Ensure that you create a pool named `flexsnap_pool` on the Hitachi storage array. This is required for the CloudPoint plug-in to work.
- Create a snapshot group named `flexsnap_default_group` on the storage array.

  Note: This is not a prerequisite. If you do not create this snapshot group, the plug-in automatically creates it during the configuration.

- Ensure that the Hitachi storage arrays are registered with Hitachi Configuration Manager (HCM). CloudPoint uses the HCM REST APIs to communicate with the storage arrays.
- Ensure that the Hitachi storage arrays have the necessary licenses that are required to perform snapshot operations.
- Ensure that the user account that you provide to CloudPoint has general read permissions as well as the permissions to create, delete, export, deport, and restore snapshots on the storage array.

See "Hitachi plug-in configuration parameters" on page 144.

See "Supported Hitachi storage arrays" on page 145.

See "Supported CloudPoint operations on Hitachi arrays" on page 145.

Hitachi plug-in configuration parameters

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

Table 5-27 Hitachi plug-in configuration parameters

CloudPoint configuration parameter	Description
Hitachi Configuration Manager Server URL	The base URL for accessing the Hitachi Configuration Manager (HCM) server. The URL has the following format: *protocol*://*host-name*:*port-number*/ConfigurationManager

Table 5-27 Hitachi plug-in configuration parameters *(continued)*

CloudPoint configuration parameter	Description
Array IP address	The IP address of the Hitachi storage array.
Array Username	The name of the user account that has access to the Hitachi storage array. In addition to general read permissions, the user account must have the permissions to create, delete, export, deport, and restore snapshots on the storage array.
Array Password	The password of the user account that is used to access the Hitachi storage array.

Supported Hitachi storage arrays

You can use CloudPoint to discover and protect the following Hitachi G Series array models:

Table 5-28 Supported Hitachi arrays

Category	Supported
Array model	VSP G1000 VSP G1500
Firmware version	80-01-21-XX/XX or later
Software development kit (SDK) required	Hitachi Configuration Manager (HCM)

For the latest information on hardware support, refer to the *CloudPoint Hardware Compatibility List (HCL)*.

See “ Meeting system requirements” on page 14.

Supported CloudPoint operations on Hitachi arrays

You can perform the following CloudPoint operations on the supported Hitachi storage arrays that are registered with Hitachi Configuration Manager (HCM):

Table 5-29 Supported CloudPoint operations on Hitachi arrays

CloudPoint operation	Description
Discover assets	CloudPoint discovers all the Logical Devices (LDEV) created on the storage array. The primary LDEV objects appear as disk assets. The secondary LDEV objects that are part of a Thin Image (TI) pair appear under snapshots. One or more LDEV objects are grouped in a logical entity called as a pool. For the CloudPoint Hitachi plug-in to work, you must create a pool named `flexsnap_pool` on the storage array.
Create snapshot	NetBackup takes a snapshot of all the LDEV objects that are attached to a hostgroup. When CloudPoint takes a snapshot, it performs the following actions: ■ Creates a new LDEV object that is of the same size as the original (base) LDEV. ■ Puts the base LDEV and the new LDEV into a Thin Image (TI) pair. The base LDEV is the primary LDEV and the new LDEV is the secondary LDEV. ■ Splits the TI pair to create a point-in-time snapshot of the base LDEV and then updates the snapshot LUN path to point to the secondary LDEV. ■ Attaches the snapshot to the same hostgroup where the base LDEV is attached.
Delete snapshot	When CloudPoint deletes a snapshot, it performs the following actions: ■ Deletes the snapshot. ■ Removes the LUN path to the secondary LDEV associated with the snapshot. ■ Deletes the secondary thin LDEV.
Restore snapshot	CloudPoint performs a restore operation on a thin image snapshot of an LDEV. All the data in the primary LDEV is overwritten by the data from the secondary LDEV.

Table 5-29 Supported CloudPoint operations on Hitachi arrays *(continued)*

CloudPoint operation	Description
Export snapshot	When a snapshot export operation is triggered, CloudPoint searches for the target host based on the world wide name (WWN) or the iSCSI Qualified Name (IQN) specified in the export request. After the host is identified on the storage array, CloudPoint updates the path attribute of the secondary LDEV with the target host where the snapshot is to be exported. Once the target host is added to the secondary LDEV host ports, the exported snapshot is immediately visible on the target host.
Deport snapshot	When a snapshot deport operation is triggered, CloudPoint removes the target host from the secondary LDEV path attribute. Once the target host entry is removed from the secondary LDEV host ports, the exported snapshot is no longer visible on the target host and the deport operation is complete.

Snapshot related requirements and limitations

Consider the following when you configure the Hitachi 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.
- The export operation is supported using the following protocols:
  - Fibre Channel (FC)
  - Internet Small Computer Systems Interface (iSCSI)

Hitachi (HDS VSP 5000) plug-in configuration notes

The CloudPoint plug-in for Hitachi (HDS VSP 5000) enables NetBackup to discover the SAN volumes on the Hitachi HDS VSP 5000 array and then trigger snapshot create, export, deport, delete, and restore operations for those exports. You must configure this plug-in on the NetBackup primary server.

CloudPoint uses the REST API SDK hosted on Hitachi Configuration Manager to communicate with the Hitachi storage arrays. CloudPoint establishes a connection with Hitachi storage array by creating sessions in Hitachi Configuration Manager

and uses the REST APIs to discover the SAN volumes and their snapshots that need to be backed up.

Hitachi (HDS VSP 5000) plug-in configuration parameters

- Create one Hitachi Thin Image (HTI) Pool with the name prefixed with `flexsnap_pool` to store snapshots.
- Gather the following information about the Hitachi (HDS VSP 5000). You will use these details while configuring the plug-in:

Table 5-30 Hitachi (HDS VSP 5000) plug-in configuration parameters

CloudPoint configuration parameter	Description
Hitachi Configuration Manager Server IP	IP of the Hitachi Configuration Manager REST server which is configured with the storage array to be used.
Hitachi Configuration Manager Server port	Port on which Hitachi Configuration Manager REST server is hosted.
Array Username	The name of the user account that has access to the Hitachi storage array. In addition to general read permissions, the user account must have the permissions to create, delete, export, deport, and restore snapshots on the storage array.
Array Password	The password of the user account that is used to access the Hitachi storage array.
Array Storage Device ID	ID of the storage array device that is already registered with the Hitachi Configuration Manager.

The following screen is displayed when you configure the plug-in using the NetBackup administration console:

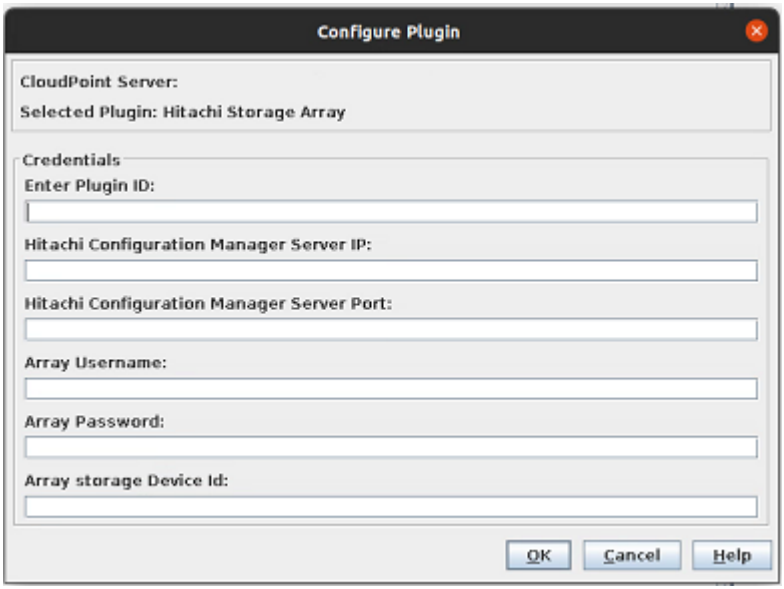

Hitachi (HDS VSP 5000) plug-in considerations and limitations

- CloudPoint uses a snapshot group while creating snapshots, so maximum number of snapshots in CloudPoint for an array is 8192 per pool.
- The Thin Image Pool must be large enough to accommodate all snapshot needs.

Supported CloudPoint operations on Hitachi (HDS VSP 5000) array

CloudPoint performs the following management operations on the Hitachi (HDS VSP 5000) storage array.

Table 5-31 Supported CloudPoint operations on Hitachi (HDS VSP 5000) arrays

CloudPoint operation	Description
Discover assets	CloudPoint discovers the Logical Devices (LDEV) created on the storage array and snapshots inside the snapshot group named similar to `flexsnap_snap_group` along with some of their metadata. The LDEVs that have "CMD" in their attributes and those without any logical unit number (LUN) mapped are not discovered

Table 5-31 Supported CloudPoint operations on Hitachi (HDS VSP 5000) arrays *(continued)*

CloudPoint operation	Description
Create snapshot	For snapshots, CloudPoint uses Hitachi Thin Image Pairs and triggers a sequence of REST API requests with the required information and snapshot name. The API returns the details of the snapshot. A typical snapshot created by CloudPoint has the following naming convention: `NB<unique_21digit_number>`
Delete snapshot	To delete a snapshot, CloudPoint triggers a sequence of REST API requests with the required snapshot details and confirms that the snapshot has been deleted successfully on the cluster.
Restore snapshot	To restore a snapshot, CloudPoint triggers a REST API request where the Thin Image is restored to the parent volume.
Export snapshot	Export over iSCSI and FC is supported by CloudPoint. Cloudpoint uses REST API to set LUN path of snapshot.
Deport snapshot	When a snapshot deport operation is triggered, CloudPoint deletes the export created over the snapshot path at the time of Export operations. It essentially reverts the Export Operation.

InfiniBox plug-in configuration notes

The CloudPoint plug-in for InfiniBox lets you create, delete, restore, export, and deport snapshots of the SAN volumes (virtual disks) that are part of storage pools on the INFINIDAT InfiniBox storage arrays.

CloudPoint supports all the InfiniBox storage arrays that are compatible with InfiniSDK.

InfiniBox plug-in configuration prerequisites

Before you configure the InfiniBox plug-in, perform the following steps on the storage system:

- Ensure that the InfiniBox storage arrays have the necessary licenses that are required to perform snapshot operations.

- Ensure that the user account that you provide to CloudPoint has administrative privileges to all the storage pools that you wish to protect using CloudPoint.

See "InifiniBox plug-in configuration parameters" on page 151.

See "Supported CloudPoint operations on InfiniBox arrays" on page 151.

InifiniBox plug-in configuration parameters

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

Table 5-32 InfiniBox plug-in configuration parameters

CloudPoint configuration parameter	Description
InfiniBox System IP Address	The IP address of the InfiniBox storage array.
Username	The name of the user account that has access to the InfiniBox storage array. The user account must have administrative privileges (`POOL_ADMIN` role) to the storage pools on the array.
Password	The password of the user account that is used to access the InfiniBox storage array.

Supported CloudPoint operations on InfiniBox arrays

CloudPoint supports the following operations on the InfiniBox storage arrays:

Table 5-33 Supported CloudPoint operations on InfiniBox arrays

CloudPoint operation	Description
Discover assets	CloudPoint discovers all the SAN volumes (virtual disks) that are part of storage pools that are created on the InfiniBox storage array. The plug-in sends a request to the array to return a list of all the volumes that have the type set as `PRIMARY`. Such volumes are considered as base volumes and appear as disk assets. To discover snapshot objects, the plug-in sends a request to the array to return a list of all the volumes that have the type set as `SNAPSHOT` and the depth attribute set as 1. Such volumes are considered as snapshots. InfiniBox arrays support creating a snapshot of a snapshot. The depth attribute identifies the snapshot type. A snapshot depth value greater than 1 indicates that it is a snapshot of an existing snapshot. CloudPoint does not support discovery and operations on snapshot volumes that have a depth value other than 1.
Create snapshot	CloudPoint takes a snapshot of all the SAN volumes that are part of a storage pool. When a snapshot is created, CloudPoint plug-in uses InfiniSDK to send a `create_snapshot` method request on the selected volume and passes a snapshot name as an argument in that request. The InifiniBox array creates a snapshot volume, sets the type as `SNAPSHOT` and the depth attribute value as 1, and returns that information to CloudPoint.
Delete snapshot	When a snapshot is deleted, CloudPoint plug-in sends a `delete_snapshot` method request on the parent volume that is associated with the snapshot and passes the snapshot volume name as an argument in that request. The InfiniBox array deletes the specified snapshot associated with the parent volume.
Restore snapshot	When a snapshot restore operation is triggered, CloudPoint first gets details about the parent volume that is associated with the snapshot that is being restored. CloudPoint plug-in then sends the `restore_snapshot` method request on the parent volume and passes the selected snapshot as an argument in that request. The array uses the selected snapshot to perform the restore on the parent volume. All the data in the parent volume is overwritten by the data in the snapshot volume.

Table 5-33 Supported CloudPoint operations on InfiniBox arrays *(continued)*

CloudPoint operation	Description
Export snapshot	When a snapshot export operation is triggered, CloudPoint searches for the target host based on the world wide name (WWN) or the iSCSI Qualified Name (IQN) specified in the export request. After the host is identified, CloudPoint plug-in sends a `map_volume` method request on the target host and passes the selected snapshot ID as an argument in that request. The InfiniBox array returns a LUN ID as a response to the restore request. CloudPoint stores the LUN ID and the target host ID mapping information internally in the CloudPoint database. The export operation also creates a new virtual asset of type `disk:snapshot:export` and that is saved in the CloudPoint database.
Deport snapshot	When a snapshot deport operation is triggered, CloudPoint first gets the target host ID from the database. The CloudPoint plug-in then sends a `unmap_volume` method request on the target host and passes the selected snapshot ID as an argument in that request. The InfiniBox array removes the snapshot volume mapping from the specified target host.

InfiniBox plug-in and snapshot related requirements and limitations

Consider the following when you configure the InfiniBox plug-in:

- The InfiniBox plug-in supports discovery and snapshot operations only on volume snapshots that have the depth attribute value set to 1. Volume snapshots that have the depth attribute value other than 1 are not supported.
- All parent volume objects and snapshot objects on an InfiniBox array are unique. While creating a snapshot of a volume, if an object with the same name already exists on the array, the create operation fails. You must ensure that the snapshot names are unique.
- 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.
- The snapshot export operation is supported using the following protocols:
  - Fibre Channel (FC)
  - Internet Small Computer Systems Interface (iSCSI)

Dell EMC PowerScale (Isilon) plug-in configuration notes

Veritas NetBackup provides a robust data protection solution for shares that are set up on a Network Attached Storage (NAS) storage host. NetBackup extends the NAS support and allows you to protect NFS exports that are hosted in a Dell EMC PowerScale (Isilon) environment. You can configure CloudPoint to discover and perform backup and restore operations on Network File System (NFS) exports.

The CloudPoint plug-in for Dell EMC PowerScale contains the necessary functional logic that enables NetBackup to discover the NFS exports on the PowerScale (Isilon) and trigger snapshot create, export, deport, snapshot diff (changelist), and delete operations for the exports. You must configure this plug-in on the NetBackup primary server.

CloudPoint uses the REST API SDK that PowerScale (Isilon) (isilon_sdk_python) provides to communicate with the PowerScale (Isilon) NFS exports and snapshots. CloudPoint establishes a connection with PowerScale (Isilon) by registering itself as a backup application and then uses the API endpoints to discover the NFS exports and their snapshots that need to be backed up.

Dell EMC PowerScale (Isilon) plug-in configuration prerequisites

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

- Ensure that the OneFS version of Dell EMC PowerScale (Isilon) is supported. CloudPoint supports the following:
  - OneFS version 8.0 and later
  - For vendor change tracking OneFS version 8.2.1 and later
- Gather the following information about the Dell EMC PowerScale (Isilon). You will use these details while configuring the PowerScale plug-in:

Parameter	Description
Cluster Address	An Isilon cluster consists of three or more hardware nodes. You can add any management IP address or the Fully Qualified Domain Name (FQDN) of the Node.
Username	A user account that has permissions to perform the snapshot operations on the PowerScale cluster.

Parameter	Description
Password	The password of the PowerScale (Isilon) user account specified earlier.

The following screen is displayed when you configure the plug-in using the NetBackup administration console:

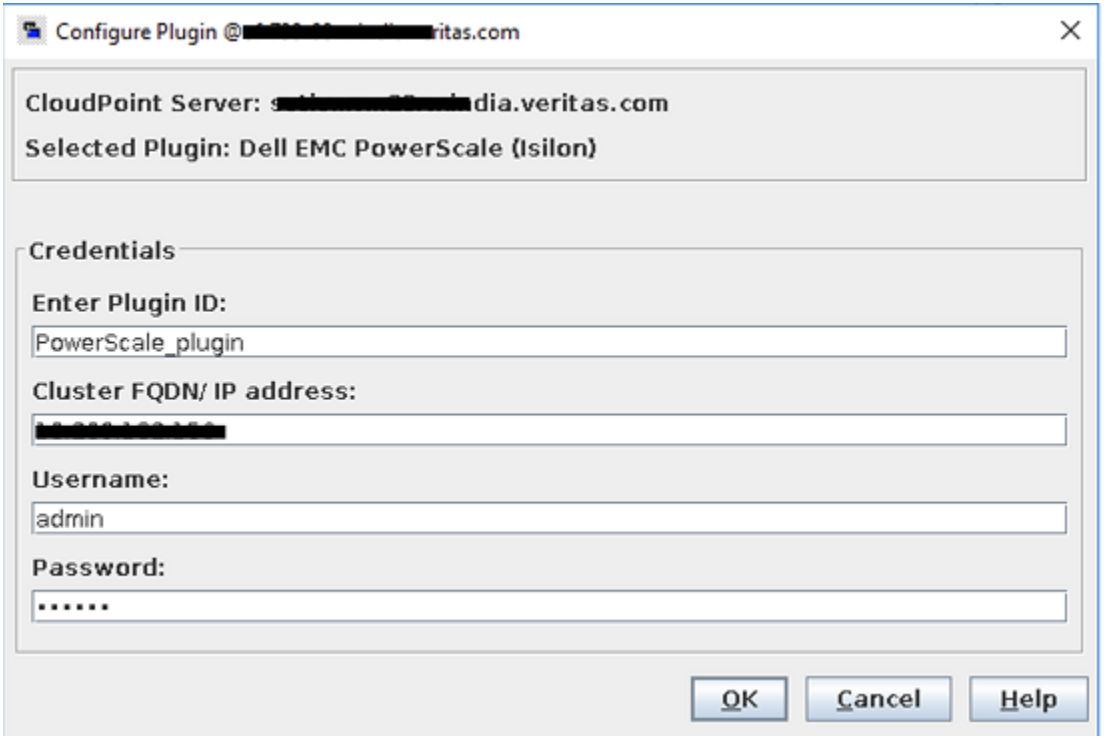

Supported CloudPoint operations on Dell EMC PowerScale (Isilon) plug-in

CloudPoint performs the following management operations on the Dell EMC PowerScale (Isilon):

Table 5-34 CloudPoint operations on Dell EMC PowerScale (Isilon) plug-in

CloudPoint operation	Description
Discover assets	CloudPoint discovers all the NFS exports and their snapshots along with some of their metadata. **Note:** CloudPoint only discovers assets with depth as 2. For example, if on NFS exports you have: ["/ifs", "/ifs/test_fs1", "/ifs/test_fs2", "/ifs/test_fs1/test_data", "/ifs/smb_03/test_data/dir01"] so NFS exports discovered in cloudpoint are ["/ifs/test_fs1", "/ifs/test_fs2"].
Create snapshot	To create a snapshot, CloudPoint triggers a POST REST API call on the nfs_export with the required information and the snapshot name. The API returns the details of the snapshot. A typical snapshot created by CloudPoint has the following naming convention: NB<unique_21digit_number>
Delete snapshot	To delete a snapshot, CloudPoint triggers a DELETE REST API call with the required snapshot details and confirms that the snapshot has been deleted successfully on the Cluster.
Restore snapshot	CloudPoint is uses the JobAPI to revert a snapshot. To revert a snapshot that contains a directory, it is recommended that you create a SnapRevert domain for a directory. To revert a snapshot, perform the following steps: **1** Create a SnapRevert domain for the directory. **2** Create a snapshot revert job.
Export snapshot	When a snapshot export operation is triggered, a new NFS export is created over the snapshot path ("/ifs/test_fs/.snapshot/NB15985918570166499611/") and the backup host is added as a Root Client with the read-only permission.
Deport snapshot	When a snapshot deport operation is triggered, CloudPoint deletes the NFS export created over the snapshot path at the time of the export operation.

Table 5-34 CloudPoint operations on Dell EMC PowerScale (Isilon) plug-in *(continued)*

CloudPoint operation	Description
Create snapshot diff	CloudPoint use the JobAPI to create a changelist between snapshots. To create a changelist, perform the following steps: **1** Use the JobAPI to create job for creating ChangeList between snapshots. **2** Use the get_changelist_entries API to fetch the changelist entries between snapshots. **Note:** The following important points: ■ The get_changelist_entries API is available for OneFS version 8.2.1 and above only. ■ For creating a changelist, use the JobAPI. The job engine allows only 3 different types of jobs to run simultaneously. To allow multiple instances of the changelist run the following CLI: ■ isi_gconfig -t job-config jobs.types.changelistcreate.allow_multiple_instances=true' (the default is false) ■ isi_gconfig -t job-config jobs.types.changelistcreate.allow_multiple_instances'

Dell EMC PowerMax and VMax plug-in configuration notes

The CloudPoint plug-in for Dell EMC PowerMax and VMax enables NetBackup to discover the SAN Volumes mounted on PowerMax/VMax and then trigger snapshot create, export, deport, restore and delete operations for those volumes. You must configure this plug-in on the NetBackup primary server.

CloudPoint uses the REST API SDK provided by PowerMax/VMax (PyU4V) to communicate with the PowerMax/ VMax assets. CloudPoint establishes a connection with PowerMax/VMax array by registering itself as a backup application and then uses the API endpoints to discover the SAN volumes and their snapshots that needs to be backed up.

Dell EMC PowerMax and VMax plug-in configuration prerequisites

Before you configure the plug-in:

- Note the following:
  - The minimum Unisphere version required is 9.2.0.1 of Unisphere Management console
  - Supported array models are PowerMax, VMAX-3, VMAX-AFA.
  - Array uCode should be above 5978.669.669 for PowerMax OS, HyperMax OS to support 'SnapSet Id'.
- Gather the following information about the Dell EMC PowerMax/VMax. You will use these details while configuring the plug-in:

Parameter	Description
Unisphere Address	The Unisphere Management console through which all the arrays are managed. You can add any management IP Address or the FULL Qualified Domain Name (FQDN) of the Unisphere Management console.
Unisphere Port	The Unisphere Management port through which the console is accessed (Dell EMC suggests 8443). You can configure any port through which you can access the Unisphere console.
Array ID	A 12 digit unique Array ID which you want to be protected.
Username	Unisphere console user account that has permissions to perform snapshot operations on the PowerMax/VMax array.
Password	The password of the Unisphere user account specified earlier.

The following screen is displayed when you configure the plug-in using the NetBackup administration console:

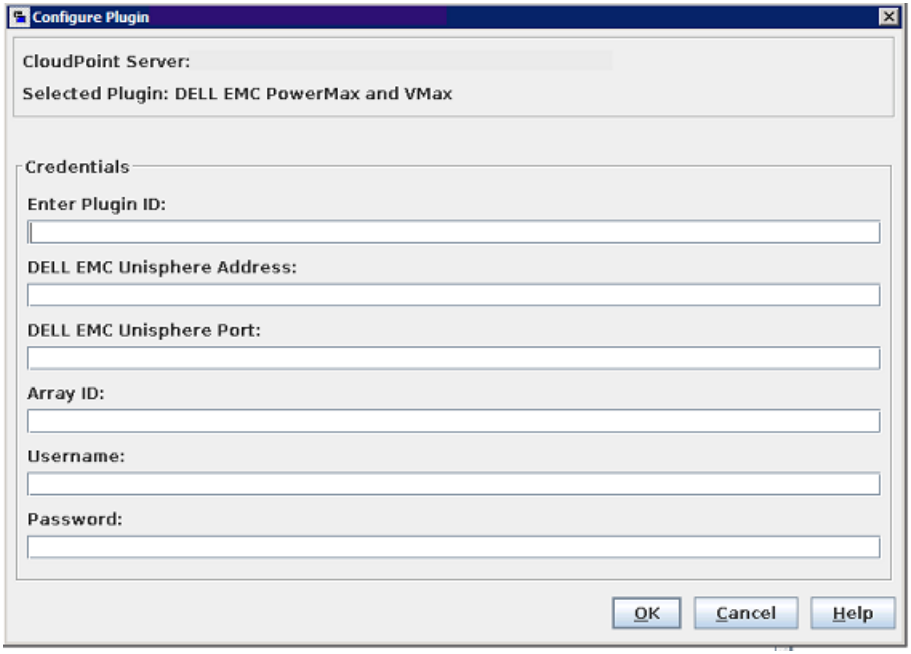

Supported CloudPoint operations on Dell EMC PowerMax and VMax

CloudPoint performs the following management operations on the Dell EMC PowerMax and VMax:

Table 5-35 CloudPoint operations on Dell EMC PowerMax/ VMax plug-in

CloudPoint operation	Description
Discover assets	CloudPoint discovers all the volumes and their snapshots along with their storage group. **Note:** CloudPoint only discovers assets with depth as 2.
Create snapshot	To create a snapshot, CloudPoint triggers a POST API call on the storage group within which the volumes resides, with the required information and snapshot name. A typical snapshot created by CloudPoint has the following naming convention: NB<unique_21digit_number>

Table 5-35 CloudPoint operations on Dell EMC PowerMax/ VMax plug-in *(continued)*

CloudPoint operation	Description
Delete snapshot	To delete a snapshot, CloudPoint triggers a DELETE REST API call with the required snapshot details and confirms that the snapshot has been deleted successfully on the array.
Restore snapshot	CloudPoint uses storage group snapshot restore API from Unisphere. To restore a snapshot to the point in time image on the volume. **1** Create an empty temporary storage group. **2** Add a volume which is to be restored in the storage group. **3** Restore the temporary storage group. **4** Delete the temporary storage group.
Export snapshot	When a snapshot export operation is triggered, a volume is carved out of the snapshot and attached to the host on which it is to be exported. Perform the following steps: **1** Fetch the initiators on which you want to perform the export. **2** Create an empty temporary storage group. **3** Add a source volume whose snapshot is to be exported in the storage group. **4** Now, considering temporary storage as source storage group, create an export storage group from snapshot and link the snapshot to the exported storage group. **5** Fetch the Host ID and Port group ID. **6** Using the export storage group, Host ID and Port group ID, create a masking view group which would attach the exported storage group to the host.
Deport snapshot	When a snapshot deport operation is triggered, CloudPoint deletes the exported storage group and the volume inside it, and the temporary storage group that is used as a source. It essentially reverts the snapshot export operation.

Qumulo plug-in configuration notes

NetBackup provides a robust data protection solution for shares that are set up on a Network Attached Storage (NAS) storage host. NetBackup extends this NAS

support and allows you to protect NFS exports that are hosted in a Qumulo environment. You can configure CloudPoint to discover and then perform backup and restore operations on Network File System (NFS) exports.

The CloudPoint plug-in for Qumulo contains the necessary functional logic that enables NetBackup to discover the NFS exports on the Qumulo cluster and then trigger snapshot create, export, deport, and delete operations for those exports. You must configure this plug-in on the NetBackup primary server.

CloudPoint uses the REST API SDK Qumulo (qumulo-api) provides to communicate with the Qumulo assets. CloudPoint establishes a connection with Qumulo by using the RestClient library exposed by SDK and then uses the SDK methods to discover the NFS exports and their snapshots that need to be backed up.

Qumulo plug-in configuration prerequisites

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

- Ensure that the Qumulo Core version is supported.
  CloudPoint supports version 3.0.5 and later.
- Gather the following information about the Qumulo cluster. You will use these details while configuring the plug-in:

Parameter	Description
Cluster Address	You can add any management IP address or the Fully Qualified Domain Name (FQDN) of the Node. You can also use Qumulo DNS Roundrobin FQDN here.
Username	A user account that has permissions to perform snapshot operations on the Qumulo cluster.
Password	The password of the Qumulo user account specified earlier.

The following screen is displayed when you configure the plug-in using the NetBackup administration console:

Qumulo plug-in considerations and limitations

The following considerations and limitations are applicable:

- Snapshot operations are not supported for nested shares on Qumulo file server. A nested share is a share that is itself a sub-directory in an existing file share. NetBackup does not support snapshot creation for such nested shares.
- Qumulo File Server does not support point-in-time (PIT) rollback restore of shares using snapshots. You can use NetBackup assisted restore of share's data.
- NFSv4 is not supported by the Qumulo plug-in. NetBackup provides an explicit option in NAS policy to configure NFS mount version NFSv3 and NFSv4 for backup jobs but by default the NFSv3 is configured for NAS Policy.

Supported CloudPoint operations on Qumulo plug-in

CloudPoint performs the following management operations on the Qumulo plug-in:

Table 5-36 CloudPoint operations on Qumulo plug-in

CloudPoint operation	Description
Discover assets	CloudPoint discovers all the Qumulo file system paths and their snapshots along with some of their metadata. Single depth discovery is supported.. For example, if there filesystem directories are [/home, /home/user1, /home/user2, /user1], the discovered filesystem are [/home, /user1].
Create snapshot	To create a snapshot, CloudPoint triggers an SDK method with the required information and snapshot name. The API returns the details of the snapshot. A typical snapshot created by CloudPoint has the following naming convention: NB<unique_21digit_number>
Delete snapshot	To delete a snapshot, CloudPoint triggers a SDK method call with the required snapshot details. Then CloudPoint confirms that the snapshot has been deleted successfully on the cluster.
Restore snapshot	CloudPoint does not support this operation.
Export snapshot	When a snapshot export operation is triggered, a new NFS export is created over the same filesystem path on which the backup hosts is added as a client with the read-only permission.
Deport snapshot	When a snapshot deport operation is triggered, CloudPoint deletes the NFS export created over the snapshot path at the time of the export operation.
Create snapshot diff	CloudPoint does not support this operation.

Chapter 6

CloudPoint application agents and plug-ins

This chapter includes the following topics:

Microsoft SQL plug-in configuration notes

You can configure the CloudPoint plug-in for Microsoft SQL to discover SQL application instances and databases and protect them using disk-level snapshots. After you configure the plug-in, CloudPoint automatically discovers all the file system assets, SQL instances and databases that are configured on the SQL server host. The discovered SQL assets then appear in the NetBackup user interface (UI) from where you can protect the assets by subscribing them to a protection plan or by taking snapshots manually.

Microsoft SQL plug-in configuration requirements

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

- This plug-in is supported in Microsoft Azure, Google Cloud Platform and Amazon AWS environments.
- A supported version of Microsoft SQL server is installed on the Windows instance.
  See “ Meeting system requirements” on page 14.
- 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.
- CloudPoint uses the Microsoft Volume Shadow Copy Service (VSS).
  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 179.

Note: CloudPoint does not support discovery, snapshot, and restore operations for SQL databases that contain leading or trailing spaces or non-printable characters. This is because the VSS writer goes into an error state for such databases. Refer to the following for more details:

https://support.microsoft.com/en-sg/help/2014054/backing-up-a-sql-server-database-using-a-vss-backup-application-may-fa

Oracle plug-in configuration notes

You can configure the Oracle plug-in to discover and protect your Oracle database applications with disk-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 14.
- Oracle standalone instance is discoverable.
- Oracle binary and Oracle data must be on separate volumes.
- Log archiving 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.

Veritas 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.

About the installation and configuration process

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

You can also establish the agent connection using agentless connection mechanism, See “About the agentless feature” on page 198.

Figure 6-1 CloudPoint agent installation and configuration process

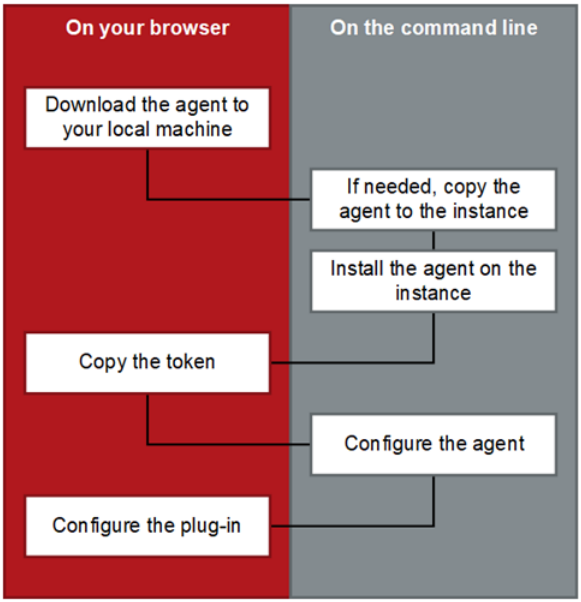

See "Preparing to install the Linux-based agent" on page 168.

See "Preparing to install the Windows-based agent" on page 168.

See "Downloading and installing the CloudPoint agent" on page 169.

Preparing to install the Linux-based agent

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

- 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 166.
  See "About the installation and configuration process" on page 167.

Preparing to install the Windows-based agent

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

- Verify that the required ports are enabled on the CloudPoint host.

See "Verifying that specific ports are open on the instance or physical host" on page 32.

- Verify that you can 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.

Downloading and installing the CloudPoint agent

Download and install the appropriate CloudPoint agent depending on the application that you wish to protect. Whether you install the Linux-based agent or the Windows-based agent, the steps are similar.

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

- Make sure that you have administrative privileges on the application host on which you want to install the agent.
  If a non-admin user attempts the installation, the installer displays the Windows UAC prompt where the user must specify the credentials of an admin user.
- Complete the preparatory steps and install all the dependencies for the respective agent.
  See "Preparing to install the Linux-based agent" on page 168.
  See "Preparing to install the Windows-based agent" on page 168.

To download and install the agent

1. Sign in to the NetBackup Web UI.
2. From the left navigation pane, click **Cloud** and then select the **CloudPoint servers** tab.

   All the CloudPoint servers that are registered with the primary server are displayed in this pane.

3 From the desired CloudPoint server row, click the actions icon on the right and then select **Add agent**.

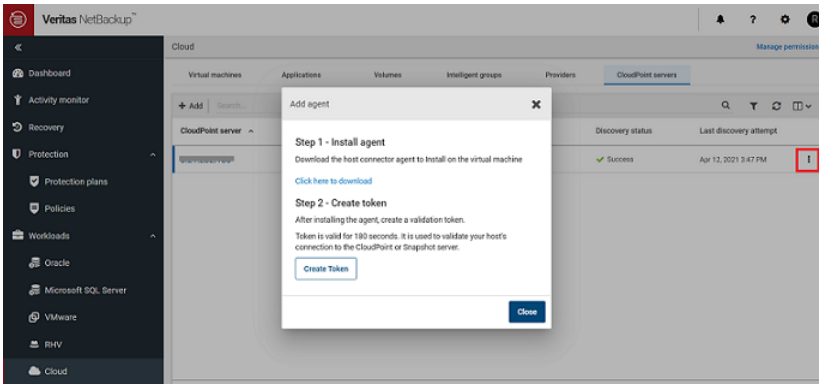

4 On the Add agent dialog box, click the 'download' link.

This launches a new browser window.

Do not close the existing Add agent dialog box on the NetBackup Web UI as yet. When you configure the agent, you will return to this dialog box to get the authentication token.

5 Switch to the new web page browser window and from the Add Agent section, click on the download link to download the desired CloudPoint agent installation package.

The web page provides separate links to download the Linux and Windows agents.

6 If necessary, copy the downloaded agent package to the application host on which you want to install the agent.

7 Install the agent.

- For the Linux-based agent, type the following command on the Linux host:

  ```
  # sudo yum -y install <cloudpoint_agent_rpm_name>
  ```

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

  For example:

  ```
  # sudo yum -y install
  VRTScloudpoint-agent-8.3.0.8549-RHEL7.x86_64.rpm
  ```

- For the Windows-based agent, run the agent package file and follow the installation wizard workflow to install the agent on the Windows application host.

Note: To allow the installation, admin users will have to click Yes on the Windows UAC prompt. Non-admin users will have to specify admin user credentials on the UAC prompt.

The installer installs the agent at `C:\Program Files\Veritas\CloudPoint` by default and the path cannot be modified.

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:

```
msiexe /i
C:\temp\VRTScloudpoint-agent-8.3.0.8549-Windows.x64.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 useful if you want to automate the agent installation using a third-party deployment tool.

8 This completes the agent installation. You can now proceed to register the agent.

See "Registering the Linux-based agent" on page 171.

See "Registering the Windows-based agent" on page 174.

Registering the Linux-based agent

Verify the following before you register the Linux-based agent:

- Ensure that you have downloaded and installed the agent on the application host.

  See "Downloading and installing the CloudPoint agent" on page 169.

- Ensure that you have root privileges on the Linux instance.
- If the CloudPoint Linux-based 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
    ```

- If the CloudPoint Linux-based agent was already registered on the host earlier, and you wish to register the agent with a different CloudPoint instance, then do the following on the Linux host:
  - Uninstall the agent from the Linux host.
    See "Removing the CloudPoint agents" on page 254.
  - Remove the `/opt/VRTScloudpoint/keys` directory from the Linux host.
    Type the following command:
    ```
    # sudo rm -rf /opt/VRTScloudpoint/keys
    ```
  - Remove the `/etc/flexsnap.conf` configuration file from the Linux host.
    Type the following command:
    ```
    sudo rm -rf /etc/flexsnap.conf
    ```
  - Re-install the agent on the Linux host.
    See "Downloading and installing the CloudPoint agent" on page 169.

  If you do not perform these steps, then the on-host agent registration may fail with the following error:

  ```
  On-host registration has failed. The agent is already registered
  with CloudPoint instance <instance>.
  ```

To register the Linux-based agent

1. Return to the NetBackup Web UI, and on the Add agent dialog box, click **Create Token**.

   If you have closed the dialog box, sign in to the NetBackup Web UI again and do the following:
   - Click **Cloud** from the left navigation menu, and select the **CloudPoint servers** tab.
   - From the desired CloudPoint server row, click the actions button on the right and then select **Add agent**.

- On the Add agent dialog box, click **Create Token**.

2 Click **Copy Token** to copy the displayed CloudPoint validation token.

The token is a unique sequence of alpha-numeric characters and is used as an authentication token to authorize the host connection with CloudPoint.

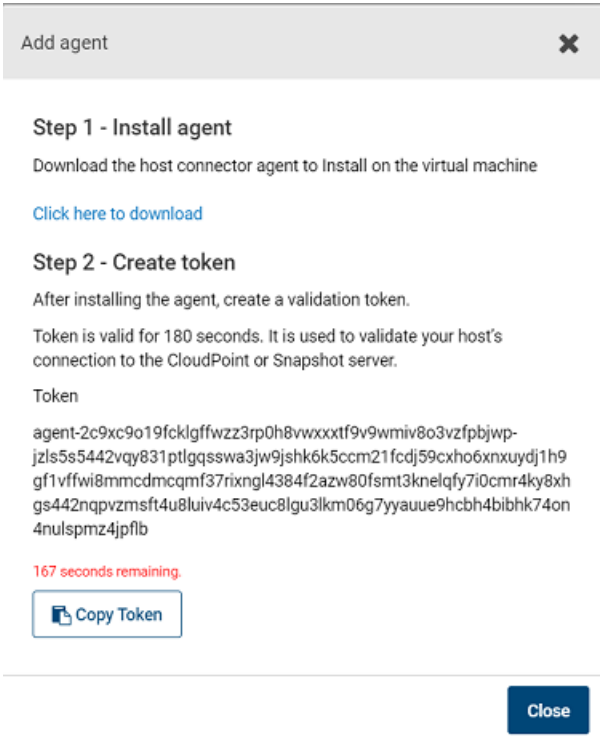

Note: The token is valid for 180 seconds only. If you do not copy the token within that time frame, generate a new token again.

3 Connect to the Linux host and register the 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 server's Fully Qualified Domain Name (FQDN) or IP address that was specified during the CloudPoint 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 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 agent service on the Linux host

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

4 Return to the NetBackup Web UI, close the Add agent dialog box, and then from the CloudPoint server row, click the actions button on the right and then click **Discover**.

This triggers a manual discovery of all the assets that are registered with the CloudPoint server.

5 Click on the **Virtual machines** tab.

The Linux host where you installed the agent should appear in the discovered assets list.

Click to select the Linux host. If the host status is displayed as **VM Connected** and a **Configure Application** button appears, it confirms that the agent registration is successful.

6 This completes the agent registration. You can now proceed to configure the application plug-in.

See "Configuring the CloudPoint application plug-in" on page 178.

Registering the Windows-based agent

Verify the following before you register the Windows-based agent:

- Ensure that you have downloaded and installed the agent on the Windows application host.
  See "Downloading and installing the CloudPoint agent" on page 169.
- Ensure that you have administrative privileges on the Windows host.

To register the Windows-based agent

1 Return to the NetBackup Web UI, and on the Add agent dialog box, click **Create Token**.

If you have closed the dialog box, sign in to the NetBackup Web UI again and do the following:

- Click **Cloud** from the left navigation menu, and select the **CloudPoint servers** tab.
  From the desired CloudPoint server row, click the actions button on the right and then select **Add agent**.

- On the Add agent dialog box, click **Create Token**.

2 Click **Copy Token** to copy the displayed CloudPoint validation token.

The token is a unique sequence of alpha-numeric characters and is used as an authentication token to authorize the host connection with CloudPoint.

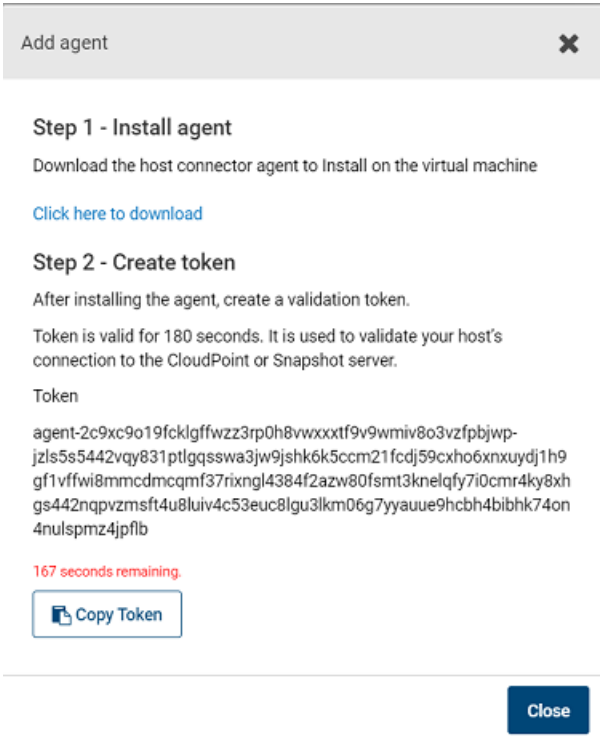

Note: The token is valid for 180 seconds only. If you do not copy the token within that time frame, generate a new token again.

3 Connect to the Windows instance and register the agent.

From the command prompt, navigate to the 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 agent using the installation wizard earlier. The default path is `C:\Program Files\Veritas\CloudPoint\`.

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

NetBackup performs the following actions when you run this command:

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

---

Note: If you intend to automate the agent registration process using a script or a 3rd-party deployment tool, then consider the following:

Even if the agent has been registered successfully, the Windows agent registration command may sometimes return error code 1 (which generally indicates a failure) instead of error code 0.

An incorrect return code might lead your automation tool to incorrectly indicate that the registration has failed. In such cases, you must verify the agent registration status either by looking in to the flexsnap-agent-onhost logs or from the NetBackup Web UI.

---

4 Return to the NetBackup Web UI, close the Add agent dialog box, and then from the CloudPoint server row, click the actions button on the right and then click **Discover**.

This triggers a manual discovery of all the assets that are registered with the CloudPoint server.

5 Click on the **Virtual machines** tab.

The Windows host where you installed the agent should appear in the discovered assets list.

Click to select the Windows host. If the host status is displayed as **VM Connected** and a **Configure Application** button appears, it confirms that the agent registration is successful.

6 This completes the agent registration. You can now proceed to configure the application plug-in.

See "Configuring the CloudPoint application plug-in" on page 178.

Configuring the CloudPoint application plug-in

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

Before you proceed, ensure that you do the following:

- Verify that you have configured the agent on the host.
  See "Registering the Linux-based agent" on page 171.
  See "Registering the Windows-based agent" on page 174.
- Review the configuration requirements for the plug-in you want to configure.
  See "Oracle plug-in configuration notes" on page 166.
  See "Microsoft SQL plug-in configuration notes" on page 165.

To configure an application plug-in

1 Sign in to the NetBackup Web UI and from the left navigation pane, click **Cloud** and then select the **Virtual machines** tab.

2 From the list of assets, search for the application host where you installed and registered the CloudPoint agent.

Click to select the application host and verify that the **Configure application** button appears in the top bar.

3 Click **Configure application** and from the drop-down list, select the application plug-in that you want to configure, and then click **Configure**.

For example, if you want to configure the CloudPoint plug-in for Microsoft SQL, choose **Microsoft SQL Server**.

4 After the plug-in is configured, trigger an assets discovery cycle.

Click the **CloudPoint servers** tab and then from the desired CloudPoint server row, click the action button from the right and then click **Discover**.

5 After the discovery is completed, click the **Virtual machines** tab and verify the state of the application host. The Application column in the assets pane displays a value as **Configured** and this confirms that the plug-in configuration is successful.

6 Click on the **Applications** tab and verify that the application assets are displayed in the assets list.

For example, if you have configured the Microsoft SQL plug-in, the Applications tab displays the SQL Server instances, databases, and SQL Availability Group (AG) databases that are running on the host where you configured the plug-in.

You can now select these assets and start protecting them using protection plans.

Configuring VSS to store shadow copies on the originating drive

If you want to take disk-level, application-consistent snapshots of a Windows file system or Microsoft 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.
- CloudPoint does not support discovery, snapshot, and restore operations for SQL databases that contain leading or trailing spaces or non-printable characters. This is because the VSS writer goes into an error state for such databases. Refer to the following for more details:
  https://support.microsoft.com/en-sg/help/2014054/backing-up-a-sql-server-database-using-a-vss-backup-application-may-fa

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
   ```

Creating a NetBackup protection plan for cloud assets

A protection plan defines when backups are performed, how long the backups are retained, and the type of storage to use. Once you have set up a protection plan, you can subscribe assets to that protection plan.

To create a protection plan

1. Sign in to the NetBackup Web UI.
2. From the left navigation pane, click **Protection plans** and then click **Add** from the right hand side.
3. On the Basic properties panel, do the following:
   - Enter a **Name** and **Description** for the plan.
   - From the **Workload** drop-down, select **Cloud**.
   - From the **Cloud Provider** drop-down, select a cloud provider. NetBackup supports homogenous cloud asset subscriptions. While subscribing an

asset to a protection plan, the cloud provider of the asset must be the same as the cloud provider defined in the protection plan.

- Click **Next**.

4 On the Schedules and retention panel, specify the desired backup schedule and then click **Next**.

5 Configure the remaining options as per your requirement and click **Finish** to create the protection plan.

The Protection plans pane displays the plan you created.

6 You can now proceed to assign assets to this protection plan.

See "Subscribing cloud assets to a NetBackup protection plan" on page 181.

For detailed information about managing protection plans, refer to the *NetBackup Web UI Backup Administrator's Guide*.

Subscribing cloud assets to a NetBackup protection plan

You can subscribe a single asset or a group of assets to a protection plan. For example, you can create a plan 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 take a snapshot once a month.

NetBackup supports homogenous cloud asset subscriptions. While subscribing an asset to a protection plan, the cloud provider of the asset must be the same as the cloud provider defined in the protection plan.

Before you proceed, ensure that you have sufficient privileges to assign assets to a protection plan from the NetBackup Web UI.

To subscribe cloud assets to a protection plan

1 Sign in to the NetBackup Web UI.

2 From the left navigation pane, click **Cloud** and then select the **Applications** tab.

The Application tab displays a list of assets that you can protect.

3 On the Applications tab, search and select the asset that you wish protect and then click **Add Protection**.

For example, to protect Microsoft SQL, you can select a SQL instance, a standalone database, or an Availability Group (AG) database.

Note: If instance level SQL server backup is selected, only the databases that are online are included in the snapshot. The snapshot does not include databases that are offline or in an erroneous state.

---

4 On the Choose a protection plan panel, search and select the appropriate protection plan and then click **Protect**.

Verify that on the Applications tab, the Protected by column for the selected asset displays the protection plan that you just assigned. This indicates that the asset is now being protected by the configured protection plan.

The backup jobs should automatically get triggered as per the schedule defined in the plan. You can monitor the backup jobs from the Activity monitor pane.

For more detailed information on how to subscribe assets to a protection plan, refer to the *NetBackup Web UI Backup Administrator's Guide*.

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).
- In case of a SQL instance 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 “Recovering a SQL database to the same location” on page 186.

  - 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 190.

- CloudPoint does not support discovery, snapshot, and restore operations for SQL databases that contain leading or trailing spaces or non-printable characters. This is because the VSS writer goes into an error state for such databases.
  Refer to the following for more details:
  https://support.microsoft.com/en-sg/help/2014054/backing-up-a-sql-server-database-using-a-vss-backup-application-may-fa
- Before you restore a SQL Availability Group (AG) database, perform the pre-restore steps manually.
  See "Steps required before restoring SQL AG databases" on page 185.
- New location restore of system database is not supported.
- If destination instance has AG configured, restore is not supported.
- If database exists on new location destination and the overwrite existing option is not selected, the restore job will fail.
- If the overwrite existing option is selected for database that is a part of an AG, the restore job will fail.
- For system database restore, the SQL Server version must be same. For user databases, restore from a higher SQL version to a lower version is not allowed.
- Default timeout of 6 hours is not allowing restore of larger database (size more than 300 GB). Configurable timeout parameter value can be set to restore larger database. See ??? on page 274. section.

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 NetBackup 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.
  - Then, after the disk-level restore is successful, perform the additional manual steps.
    See "Additional steps required after an Oracle snapshot restore" on page 184.
- In an Azure environment, it is observed that the device mappings may sometimes get modified after performing a host-level restore operation. As a result, the Oracle application may fail to come online on the new instance, after the restore.
  To resolve this issue after the restore, you have to manually unmount the file systems and then mount them again appropriately as per the mappings on the original host.
  If you are using the `/etc/fstab` file to store file systems, mount points, and mount settings, Veritas recommends that you use the disk UUID instead of device mappings. Using disk UUIDs ensures that the file systems are mounted correctly on their respective mount points.
- Snapshots of application data residing on a filesystem that is part of an LVM type of partition are not supported. If you try to take a snapshot of such a filesystem, the following error is displayed:
  ```
  *flexsnap.GenericError: Unable to protect asset *
  ```

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 a disk-level restore to the original location or an alternate location
- 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 `listerner.ora` and `tnsnames.ora` files.
7. Restart the Oracle listener using the following command:

   ```
   # lsnrctl start
   ```

Steps required before restoring SQL AG databases

You must perform the following steps before you restore a SQL Availability Group (AG) database:

Note: If you are restoring the AG database to multiple replicas, perform the entire restore process on the primary replica first, and then repeat the steps for each secondary replica.

---

1. For the database that you want to restore, suspend data movement from the replica.

   From the SQL Server Management Studio, right-click on the database and select **Suspend Data Movement**.
2. Remove the database from the AG on the replica.

From the SQL Server Management Studio, right-click on the database and select **Remove Database from Availability Group**.

Confirm that the database is no longer part of the AG. Observe that the database on the primary replica is no longer in synchronized mode, and the status of the corresponding database on the secondary replica appears as `(Restoring...)`.

3. Delete the database from the replica.

   From the SQL Server Management Studio, right-click on the database and select **Delete**.

Recovering a SQL database to the same location

Perform the following steps to restore SQL server snapshots to the same location as that of the asset. Before you proceed, note the following:

- SQL AG databases do not support recovering to the same location.
- The RECOVERY and NORECOVERY restore options are applicable to SQL databases only.

To restore a SQL snapshot to the same location

1. Sign in to the NetBackup Web UI.
2. From the left navigation pane, click **Workloads > Cloud** and then select the **Applications** tab.
3. Select the SQL asset that you want to recover, then click **View details**, and then select the **Recovery points** tab.

   The pane displays all the recovery point snapshots that are available for restore.
4. Click to select a recovery point snapshot that you want to use for the restore.
5. From the right side, click **Recover** and then select **Original location** from the drop-down menu.

6 On the Recover to original location dialog box, choose the database recovery options and then click **Start recovery** to trigger the recovery job.

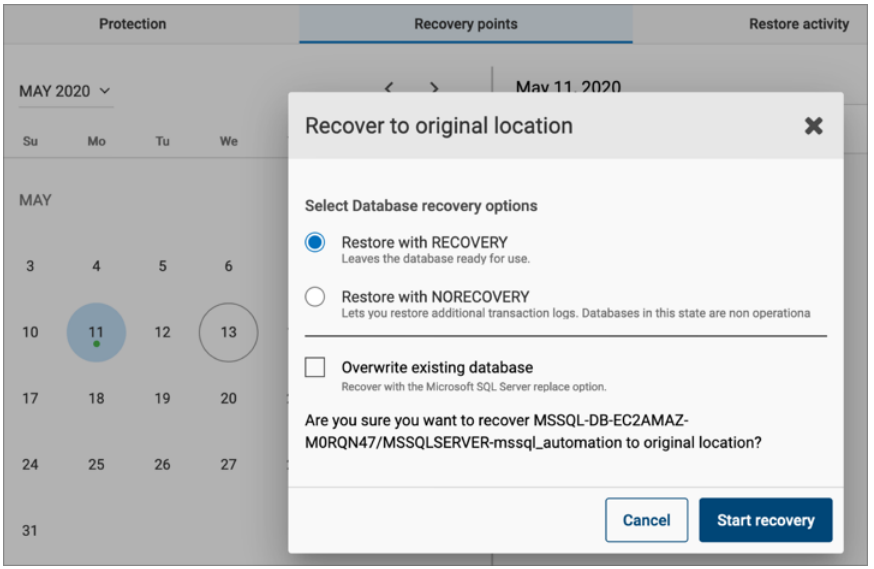

The following options are available:

Recovery option	Description
Restore with RECOVERY	Select this option if you want to perform a single restore on the database and bring it back to a consistent and operational state. The database becomes accessible immediately after the restore is complete.
Restore with NORECOVERY	Select this option if you intend to perform multiple database restores from a group of backups. For example, if you want to perform a restore using a full backup snapshot and then restore transaction logs. The database remains in the restoring state and remains inaccessible. You can work with the database only after the transaction logs are restored with the RECOVERY option.
Overwrite existing database	Select this option if you want the restore operation to replace the original database.

7 You can monitor the recovery job from the Activity monitor pane.

A status code 0 indicates that the recovery job is successful. You can now verify that the SQL database is recovered.

Recovering a SQL database to an alternate location

Perform the following steps to restore SQL databases to a new location. Before you proceed, note the following:

- SQL AG databases support recovering to an alternate location only.
- The RECOVERY and NORECOVERY restore options are applicable to SQL databases only.
- For AG databases, if you are recovering to a primary replica you must select the RECOVERY option during the restore. If you are recovering to a secondary replica, select the NORECOVERY option during the restore.
- The same steps are applicable for restoring a same name database to a new location.

  If a database with the same name already exists at the new location, you must select the overwrite existing option to perform the restore successfully.

To restore a SQL database to an alternate location

1 Sign in to the NetBackup Web UI.

2 From the left navigation pane, click **Workloads > Cloud** and then select the **Applications** tab.

3 Select the SQL asset that you want to recover, then click **View details**, and then select the **Recovery points** tab.

The pane displays all the recovery points snapshots that are available for restore.

4 Click to select a recovery point snapshot that you want to use for the restore.

5 From the right side, click **Recover** and then select **Alternate location** from the drop-down menu.

6 On the Recover to alternate location dialog box, choose the database recovery options and then click **Start recovery** to trigger the recovery job.

The following options are available:

Recovery option	Description
Restore with RECOVERY	Select this option if you want to perform a single restore on the database and bring it back to a consistent and operational state. The database becomes accessible immediately after the restore is complete. **Note:** Select this option if you are recovering an AG database to a primary replica.
Restore with NORECOVERY	Select this option if you intend to perform multiple database restores from a group of backups. For example, if you want to perform a restore using a full backup snapshot and then restore transaction logs. The database remains in the restoring state and remains inaccessible. You can work with the database only after the transaction logs are restored with the RECOVERY option. **Note:** Select this option if you are recovering an AG database to a secondary replica.
Overwrite existing database	If a database with the same name exists at the target location, select this option if you want the restore operation to replace that database.

7 You can monitor the recovery job from the Activity monitor pane.

A status code 0 indicates that the recovery job is successful. You can now verify that the SQL database is recovered.

8 If recovering SQL database in restoring mode, then after the recovery operation is complete, verify that the state of the database on the SQL host appears as `(Restoring...)`.

9 If applicable, you can now manually restore any transaction logs on the recovered database.

Additional steps required after a SQL Server snapshot restore

The following steps are required after you restore a SQL Server snapshot from the NetBackup 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.

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 NetBackup 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.

Note: These steps are applicable only in case of a SQL Server instance snapshot restore to a new location. These are not applicable for a SQL Server database snapshot restore.

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 Assign drive letters to the volumes on the disk that you brought online in the earlier step. Drive letters are required to view the shadow copies associated with each volume on the disk.

Go back to the command prompt window and perform the following steps:

- View the list of volumes on the new host using the following command:
  `list volume`
  From the list of volumes displayed, identify the volume for which you want to assign, modify, or remove a drive letter.
- Select the desired volume using the following command:
  `select volume <volnumber>`
  Here, <*volnumber*> represents the volume that you noted in the earlier step.
- Assign a drive letter to the selected volume using the following command:
  `assign letter=<driveletter>`
  Here, <*driveletter*> is the drive letter that you wish to assign to the volume. Ensure that the specified drive letter is not already in use by another volume.
- Repeat these steps to assign a drive letter to all the SQL Server volumes on the disk.

10 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 NetBackup.

Additional steps required after restoring SQL AG databases

You must perform the following steps after restoring a SQL Availability Group (AG) database:

Note: If you are restoring the AG database to multiple replicas, perform the entire restore process on the primary replica first, and then repeat the steps for each secondary replica.

- Add the restored database to the AG on the primary replica.
  From the SQL Server Management Studio, right-click on the AG entry and select **Add Database**. In the wizard workflow, select the database, and on the Initial Data Synchronisation page, select the **Skip Initial Data Synchronization** option. You can select the other options depending on the requirement.

If you restoring the same database to a secondary replica, perform the following steps:

1. Restore database to the secondary SQL instance in "Not recovered" state. Restore with no recovery should be successful.
2. Join the database to the AG on the secondary replica.

   From the SQL Server Management Studio, connect to the secondary replica node, then right-click on the database and select **Join Availability Group**.

   Observe that the database status on the secondary replica change from `(Restoring...)` to `(Synchronized)`, indicating that AG database snapshot restore is successful.

You must repeat these steps for each replica where you wish to restore an AG database.

SQL snapshot or restore and granular restore operations fail if the Windows instance loses connectivity with the CloudPoint host

This issue occurs if the CloudPoint 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 granular restore 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.

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, NetBackup 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 NetBackup UI:

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

The NetBackup 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 NetBackup 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 NetBackup UI.

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, NetBackup 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.

Chapter 7

Protecting assets with CloudPoint's agentless feature

This chapter includes the following topics:

About the agentless feature

If you want NetBackup to discover and protect assets on a host, but you want to minimize the vendor software footprint on the 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 on Linux and WMI and SMB in case of Windows.
- CloudPoint performs the specified task, such as creating a snapshot.
- When the task completes, CloudPoint software stops the process.

The CloudPoint agentless feature currently discovers and operates on Windows or Linux file system assets, Oracle database and Ms SQL database assets.

See “Prerequisites for the agentless configuration” on page 199.

See “Configuring the agentless feature” on page 201.

Prerequisites for the agentless configuration

Prerequisites for using the agentless feature in Linux

- Have the following information with you:
  - Host user name
  - Host password or SSH key

  CloudPoint 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.

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.

---

1. Perform the following steps on a host where you want to configure the agentless feature
2. 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.
3. Log out and log in again for the changes to take effect.
4. Edit the `/etc/sudoers` file using the `visudo` command:

   ```
   # sudo visudo
   ```
5. Add the following entry to the `/etc/sudoers` file:

   ```
   hostuserID ALL=(ALL) NOPASSWD: ALL
   ```
6. 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
```

7. Save the changes to the `/etc/sudoers` file.
8. Log out and log on to the host again using the user account that you provide to CloudPoint.
9. 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.

Prerequisites for using the agentless feature in Windows

- The user account used to connect to remote instance should be able to:
  - Access remote admin share (ADMIN$). Enabled by default.
  - Access to `root\cimv2`
- Configure the following ports:
  - Modify the security group to allow inbound traffic on the ports 135, 445 and dynamic port or fixed port for WMI .
  - Enable inbound rules in the firewall for the ports 135, 445 and the dynamic or fixed WMI-IN ports on Windows hosts.

Note: The dynamic range for the ports is 49152-65535.

  - You can use fixed or dynamic WMI-IN ports. If you want to configure a fixed WMI-IN port, see https://docs.microsoft.com/en-us/windows/win32/wmisdk/setting-up-a-fixed-port-for-wmi
- Disable User Account Control for the users groups accessing the agentless feature.
- For protecting SQL applications, the user account used for connecting to the cloud host, must have the required admin privileges to access the SQL server.

Configuring SMB for Windows (Optional)

Perform the following Server Message Block (SMB) configurations before configuring the agentless feature on Windows.

- Restrict unencrypted access to SMB share by setting the value to `True`.

  ```
  RejectUnencryptedAccess: True
  ```

- Disable SMB 1.0 by running the following command on Windows powershell:

  ```
  Set-SmbServerConfiguration –EnableSMB1Protocol $false
  ```

  For more details, see https://docs.microsoft.com/en-us/windows-server/storage/file-server/smb-security#disabling-smb-10

For more details on SMB security, see:
https://docs.microsoft.com/en-us/windows-server/storage/file-server/smb-security

Configuring WMI security for Windows (optional)

Windows Management Instrumentation (WMI) security protects access to the namespace data. CloudPoint uses the `root\cimv2` namespace. This name space must be accessible to only those users that are configured using the connect option. For details, see https://docs.microsoft.com/en-us/windows/win32/wmisdk/maintaining-wmi-security?redirectedfrom=MSDN

Configuring the agentless feature

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

See “Prerequisites for the agentless configuration” on page 199.

To configure the agentless feature

1. Sign in to the NetBackup Web UI and from the left navigation pane, click **Cloud** and then select the **Virtual machines** tab.
2. From the list of assets, search for the host on which you want to use the agentless feature.

   Note: The CloudPoint agentless feature currently discovers and operates on Windows or Linux file system assets, Oracle database and Ms SQL database assets.

3. Click to select the host and then click **Connect** in the top bar.

Note: If you have not assigned any credential to the VM, a message prompts you to assign the credentials before you can connect the VM. See the *Managing Credentials* section, in the *Web UI Administrator's Guide.*

Configuring the agentless feature after upgrading CloudPoint

After upgrade the cloud assets which were already in connected state, continues to work. If you want to change the asset's credentials for Linux agentless instance(s), which are already in connected state, the credentials must be associated and updated for the asset(s) from credential management.

Chapter 8

Volume Encryption in NetBackup CloudPoint

This chapter includes the following topics:

About volume encryption support in CloudPoint

NetBackup CloudPoint supports disk volume encryption for AWS, Azure, and Google Cloud Platforms. Volume encryption is provided using customer keys or system keys from the cloud provider Key Management Service (KMS).

Volume encryption for Azure

You can encrypt disks in Azure using the following methods:

- Default encryption, using Platform Managed Key (PMK)
- Customer Managed Key (CMK) using Azure Key vault

For more information on Azure encryption, see:
https://docs.microsoft.com/en-us/azure/security/fundamentals/encryption-models

Table 8-1 Encryption for creating snapshots

Disk encryption	Snapshot encryption
Platform Managed Key (PMK)	Same PMK is used as the source disk.
Customer Managed Key (CMK)	Same CMK is used as the source disk.

Table 8-2 Encryption for restoring snapshots

Snapshot encryption	Restored disk encryption
PMK	Same PMK is used as the snapshot.
CMK	Same CMK is used as the snapshot.

Volume encryption for GCP

You can encrypt disks in GCP using the following methods:

- Encryption by default (PMK or Google Managed Key)
- Customer Managed Encryption Key (CMEK) using Google Cloud KMS

For more information on GCP encryption, see:
https://cloud.google.com/security/encryption-at-rest

Table 8-3 Encryption for creating snapshots

Disk encryption	Snapshot encryption
Platform Managed Key (PMK)	Same PMK is used as the source disk.
CMEK	Same CMEK is used as the source disk.

Table 8-4 Encryption for restoring snapshots

Snapshot encryption	Restored disk encryption
PMK	Same PMK is used as the snapshot.
CMEK	Same CMEK is used as the snapshot, if the target restore location is within the scope of the key.

Note: For successful restoration, the target restore location must be inside the scope of the key during restoration.

Volume encryption for AWS

You can encrypt disks in AWS using the following methods:

- Default encryption, using Platform Managed Key (PMK).
- Customer Managed Encryption Key (CMEK), using AWS KMS.

For more information on AWS encryption, see:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html

Table 8-5 Encryption for creating snapshots

Disk encryption	Snapshot encryption
Platform Managed Key (PMK)	Same PMK is used as the source disk.
CMEK	Same CMEK is used as the source disk.

Table 8-6 Encryption for restoring snapshots

Snapshot encryption	Restored disk encryption
PMK	Same PMK is used as the snapshot.
CMEK	Same CMEK is used as the snapshot.

Chapter 9

CloudPoint security

This chapter includes the following topics:

Configuring security for Azure and Azure Stack

You can connect to an Azure or Azure Stack workload in two ways.

- The CloudPoint server can connect to the cloud workload using provider plugins.
- The data mover container present in the CloudPoint server, can connect to the workload, through the cloud connector plug-in component.

For Azure and Azure Stack workloads, these components connect using the HTTPS protocol. By default, peer and hosts validations are always enabled.

Configuring the cloud connector for Azure and Azure Stack

The cloud connector component connects to the workloads through a secure mechanism. You need to perform the following configurations.

SSL peer and host validations

By default, peer and host validations are enabled. You can disable peer and host validations only for Azure Stack.

To disable peer and host validation, set the parameter *VIRTUALIZATION_HOSTS_SECURE_CONNECT_ENABLED=NO* in the `/cloudpoint/openv/netbackup/bp.conf` file in the CloudPoint server. You must use HTTPS protocol, even after you disable peer and host validation.

For cloud workloads, the public root certificates are a part of the container image. NetBackup maintains the `cacert.pem` file which has root certificates of public cloud, at the following location:

`/usr/openv/var/global/wmc/cloud/cacert.pem`

For Azure Stack, you must specify the file path of the root certificates using the *ECA_TRUST_STORE_PATH* parameter in the `/cloudpoint/openv/netbackup/bp.conf` file in the CloudPoint server. The value of *ECA_TRUST_STORE_PATH* must be in the */cloudpoint/eca/trusted/cacerts.pem* file.

Configuring CRL validations

From release 10.0 onwards CloudPoint will be treated as NetBackup entity while communicating with NetBackup. Certificate Revocation List (CRL) check is enabled by default while communication happens between NetBackup entities.

- ECA_CRL_CHECK: This flag is used while communicating between two NetBackup entities. By default CRL check is enabled for ECA_CRL_CHECK flag. In case CloudPoint machines certificate revoked then communication between NetBackup and CloudPoint will fail with the following error:

  `"The CloudPoint server's certificate is not valid or doesn't exist.(9866)"`

- VIRTUALIZATION_CRL_CHECK: Before 10.0 CloudPoint was considered as workload while communication happens with NetBackup. Value of VIRTUALIZATION_CRL_CHECK flag used for CRL check whenever communication happens between NetBackup and workload. By default CRL check is disabled for VIRTUALIZATION_CRL_CHECK flag.

Note: If NetBackup is upgraded from version 9.1 to 10.0, then user can delete the VIRTUALIZATION_CRL_CHECK flag which was enabled for CRL check between NetBackup and CloudPoint.

Specifying the CRL path

If you enable CRL validations, you need to specify the path to the directory containing revoked certificates of the external CA.

In the *ECA_CRL_PATH* parameter in the `/cloudpoint/openv/netbackup/bp.conf` file in the CloudPoint server, specify the path to the directory where the certificate revocation lists (CRL) of the external CA are located. The path must be `/cloudpoint/eca/crl`.

If the *ECA_CRL_PATH* option is not specified, NetBackup downloads the CRLs from the URLs that are specified in the CRL Distribution Point (CDP) and uses them to verify revocation status of the peer host's certificate.

CA configuration for Azure Stack

You can sign the Azure Stack workloads with a different ECA than NetBackup. You can also configure in NBCA mode. You can have the following configurations:

1. NetBackup, CloudPoint configured with ECA-1 and Azure Stack with either ECA-1 or ECA-2.
   - You need to configure the *ECA_TRUST_STORE_PATH* parameter in the `/cloudpoint/openv/netbackup/bp.conf` file.
   - The trust store file is available in `/cloudpoint/eca/trusted/cacerts.pem`.The trust store file is in PEM format.
   - The file contains both NetBackup and Azure Stack appliance public root certificates. Manually append the NetBackup root CA certificates as well as the Azure Stack appliance root public certificates in this file.
2. NBCA, CloudPoint and AzureStack are configured with different ECAs: Only the Azure stack appliance public root certificates need to be present in the: /cloudpoint/eca/trusted/cacerts.pem file.
3. NBCA, CloudPoint is configured with CPCA and AzureStack is configured with ECA.
   - Use the `cacert.pem` file available under the data-mover container for peer and host validations.
   - Configure *ECA_TRUST_STORE_PATH* the CloudPoint server. *ECA_TRUST_STORE_PATH* should point to a file that contain the NetBackup root CA certificates, so that the `vnetd` is able to connect back to NetBackup servers.

Securing the connection to CloudPoint

In the CloudPoint server, you can upload CRLs of the external CA at `/cloudpoint/eca/crl`. The uploaded CRL does not work, if the `crl` directory is not present or empty.

For the data mover container, add this path against the *ECA_CRL_PATH* parameter in the `/cloudpoint/openv/netbackup/bp.conf` file.

Following three parameters are tuneable, you can add the entry under *eca* section in the `/cloudpoint/flexsnap.conf` file.

Table 9-1 ECA parameters

Parameter	Default	Value	Remarks
eca_crl_check	0 (Disabled)	0 (disabled) 1 (leaf) 2 (chain)	Certificate check level. Used to control the CRL/OCSP validation level for CloudPoint host connecting to On-prem/cloud workloads. ■ **0 (disabled)**: No CRL/OCSP is performed during validation ■ **1 (leaf)**: CRL/OSCP validation is performed only for leaf ■ **2 (chain)**: CRL/OSCP validation is performed for the whole chain
eca_crl_refresh_hours	24	Numerical value between 0 and 4830	Time interval in hours to update the CloudPoint CRLs cache from CA through the certificate CDP URL. Option is not applicable if `/cloudpoint/eca/crl` is present and contains CRL files. If it is set as 0, cache does not refresh.
eca_crl_path_sync_hours	1	Numerical value between 1 and 720	Time interval in hours to update the CloudPoint CRL cache from `/cloudpoint/eca/crl`. Option is not applicable if `/cloudpoint/eca/crl` is not present or empty.

Note: Cache is invalidated if any of ECA tuneable are added or modified manually inside the /cloudpoint/flexsnap.conf .

Note: The scope of CRL is check is limited to Azure and Azure Stack only.

Section 2

CloudPoint maintenance

Chapter 10

CloudPoint logging

This chapter includes the following topics:

About CloudPoint logging mechanism

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. 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 plug-ins. The MongoDB collection and the flexsnap-fluentd container logs are the two output plug-ins 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

How Fluentd-based CloudPoint logging works

When you install or upgrade CloudPoint, 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 plug-in 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 plug-ins configured in the `fluentd` configuration file. These logs are then sent to the `/cloudpoint/logs/flexsnap.log` file on the CloudPoint host.

Note that the `flexsnap.log` file gets rotated after the file size reaches a maximum of 100 MB. A total of 30 generations (rotated files) of the `flexsnap.log` file are maintained. These conditions are applicable because of the new log file rotate (`log-rotate-age`) and log size (`log-rotate-size`) command options that are introduced in the fluentd command.

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 plug-in configuration file. For CloudPoint, these plug-in 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 plug-in definition from this configuration file to determine where to send the CloudPoint log messages.

The following output plug-in definitions are added to the configuration file by default:

- `STDOUT`

  This is used to send the CloudPoint log messages to `/cloudpoint/logs/flexsnap.log`.

  The plug-in is defined as follows:

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

Additionally, the CloudPoint fluentd configuration file includes plug-in definitions for the following destinations:

- MongoDB
- Splunk
- ElasticSearch

These plug-in definitions are provided as a template and are commented out in the file. To configure an actual MongoDB, 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 plug-in 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 plug-in definition.
2. Save all the changes to the file.
3. 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.

CloudPoint logs

CloudPoint maintains the following logs that you can use to monitor CloudPoint activity and troubleshoot issues, if any. The logs are stored at `<install_path>/cloudpoint/logs` on the CloudPoint host.

Table 10-1 CloudPoint log files

Log	Description
`/cloudpoint/logs/flexsnap.log`	This log file contains all the product logs.
`/cloudpoint/logs/flexsnap-cloudpoint.log`	This log file contains all the CloudPoint installation related logs.
`/cloudpoint/logs/flexsnap-ipv6config.log`	This log file contains all the IPv6 related logs.

Logs for backup from snapshot and restore from backup jobs.

Navigate to: `/cloudpoint/openv/dm/datamover.<id>`

Here, logs can be found in the following directories: `logs`, `opt` and the `netbackup`.

- `nbpxyhelper` and `nbsubscriber` logs can be found inside the `logs` directory
- `VRTSpbx` logs can be found inside the `opt` directory
- `bpbkar, bpcd, bpclntcmd, nbcert, vnetd, vxms` and all other services logs can be found inside `netbackup` directory

To increase logging verbosity, `bp.conf` and `nblog.conf` files can be updated on CloudPoint server at `/cloudpoint/openv/netbackup`. See *NetBackup Logging Reference Guide*

Changes to the `bp.conf` and `nblog.conf` files come to effect when the next backup from snapshot or restore job runs.

Log retention

The default configuration for datamover logs is as follows:

- Log retention maximum period is 30 days. Logs older than 30 days are deleted.

- The default configuration for high and low water marks for datamover logs is 70% and 30% of the size of "/cloudpoint" mount point. For example, if the usable size of the `/cloudpoint` folder is 30 GB, then the high water mark is 21 GB (70%) and low water mark is 9GB (30%). In case, the logs directory (`/cloudpoint/openv/dm/`) size reaches to high water mark, older logs for which the datamover containers are cleaned up and no longer running are considered for deletion. The logs are deleted for such datamover containers until low water mark is reached or no logs are remaining for the datamover containers cleaned up or no longer running.

Modifying the default configuration:

You can modify the default configuration for log retention by adding such a section in the `flexsnap.conf` on the primary CloudPoint server. Open the `flexsnap.conf` file from the path `/cloudpoint/flexsnap.conf` and add the following section:

```
[datamover]
high_water_mark = 50
low_water_mark = 20
log_retention_in_days = 60
```

In case of CloudPoint extensions, the configuration from the primary server are used. Once the configuration is changed in primary CloudPoint server, the configuration is updated on each CloudPoint extension within one hour. It is not possible to have separate custom configurations for primary CloudPoint or the CloudPoint extensions and configurations should only be changed in the primary CloudPoint server. Though the configuration is same for primary as well as CloudPoint extensions, the high water mark and low water mark for log size are calculated based on the `/cloudpoint` mounted on each primary or CloudPoint extensions.

CloudPoint extension logs

Each CloudPoint extension maintains the logs under its own `/cloudpoint/logs` location.

- VM-based extension logs: Under the directory `/cloudpoint/logs`.
- Managed Kubernetes cluster-based extension logs: Under the directory `/cloudpoint/logs` which belongs to a file share.

Agentless logs

Logs for agentless connection to cloud instance(s) are present on the cloud instance at following locations based on the platform:

- **Linux**: `/tmp/ directory`
- **Windows**: `C:\\ProgramData\\Veritas\\CloudPoint\\logs\\`

Troubleshooting CloudPoint logging

You can retrieve the logs of a CloudPoint service from the `/cloudpoint/logs/flexsnap.log` file by running the following command:

For Docker environment: `# sudo cat /cloudpoint/logs/flexsnap.log | grep` ***`<flexsnap-service name>`***

For Podman environment: `# tail /cloudpoint/logs/flexsnap.log | grep` ***`<flexsnap-service name>`***

Chapter 11

Upgrading CloudPoint

This chapter includes the following topics:

About CloudPoint upgrades

You should not use two versions of CloudPoint on two different hosts to 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. Veritas recommends 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

Table 11-1 CloudPoint upgrade path

Upgrade from version	Upgrade to version
8.3	9.0, 9.0.0.1, 9.1, 9.1.0.1
9.0/9.0.0.1	9.1, 9.1.0.1
9.1/9.1.0.1	10.0
8.3/9.0/9.0.0.1	9.1/9.1.0.1 upgraded to 10.0

Notes:

- Direct upgrade from older versions to 10.0 is not supported. We need to first upgrade to 9.1 before upgrading to 10.0 for any upgrade path.
- Upgrading CloudPoint across the OS versions is not supported. If you are using CloudPoint on a RHEL7.x host, then you can only migrate it to a RHEL 8.5 or 8.4 host. Then follow the upgrade paths mentioned in the above table for upgrading CloudPoint on a RHEL 8.5 or 8.4 host.
- See “Upgrade scenarios” on page 219., for more information on upgrading NetBackup 8.3.x to NetBackup 10.0.

Upgrade scenarios

The following table lists the CloudPoint upgrade scenarios.

Note: For the NetBackup versions 9.0 and later, both NetBackup and CloudPoint versions should be at the same level. During the upgrade, first upgrade the CloudPoint server and then upgrade the NetBackup server.

Table 11-2 Upgrade scenarios

Scenario	Description	Action
Full upgrade from NetBackup 8.3 or 9.0 to NetBackup 9.1 or later	If you plan to upgrade NetBackup to 9.1 or later that includes upgrading all CloudPoint servers.	■ Disable CloudPoint servers ■ Upgrade CloudPoint servers ■ Upgrade NetBackup primary server ■ Then enable CloudPoint servers See "Upgrading CloudPoint" on page 221. **Note:** If you do not plan to upgrade one or more CloudPoint servers, then you must disable them using the NetBackup Web UI. In that case, any assets associated with the disabled CloudPoint servers cannot be protected by NetBackup.
Only CloudPoint upgrades to version 9.1 or later	If you plan to upgrade only the CloudPoint servers to 9.1 or later, but do not plan to upgrade NetBackup to 9.1 or later.	■ Please contact Veritas Technical Support to obtain an Emergency Engineering Binary (EEB) to support the incompatibility between the CloudPoint and NetBackup versions. ■ Disable CloudPoint servers ■ Apply the EEB patch on the NetBackup primary server and associated media servers. ■ Upgrade CloudPoint servers ■ Then enable CloudPoint servers See "Upgrading CloudPoint using patch or hotfix" on page 235.
Upgrading to NetBackup version 10.0	If your NetBackup 8.3.x server has CloudPoint, you must first upgrade CloudPoint to NetBackup 9.1.x before you upgrade to NetBackup 10.0. Then you can proceed to upgrade NetBackup 8.3.x to NetBackup 10.0.	The process for this upgrade is: ■ Disable the CloudPoint server for maintenance in the NetBackup web UI. ■ Upgrade the CloudPoint server from NetBackup 8.3.x to NetBackup 9.1.x. ■ Upgrade the CloudPoint server from NetBackup 9.1.x to NetBackup 10.0. ■ Enable the CloudPoint server in the NetBackup web UI. ■ Upgrade the NetBackup server from 8.3.x directly to 10.0.

Preparing to upgrade CloudPoint

Note the following before you upgrade

- Ensure that the CloudPoint instance, virtual machine, or physical host meets the requirements of the CloudPoint version you are upgrading to.
  See " Meeting system requirements" on page 14.

- Ensure that the ports required by NetBackup server meet the requirements as mentioned in the *Required Ports* section of Preparing CloudPoint for backup from snapshot jobs.
- 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 during the upgrade process when prompted or manually, if required.
  See "Backing up CloudPoint" on page 250.
- Ensure that no jobs are running on CloudPoint.
  - If you are using NetBackup Web UI, disable the CloudPoint server and wait for all the in-progress jobs to complete. Use the `nbstlutil` command to cancel all the pending SLP operations. Use one of the following commands:
    - To cancel the pending SLP operation for a specific image, use `nbstlutil cancel -backupid <value>`
    - To cancel the pending SLP operation for images that belong to specific lifecycle, use `nbstlutil cancel -lifecycle <name>`
  - If you are using NetBackup Administration console (Java UI), on the NetBackup primary server, run the following command to stop all NetBackup processes:
    - UNIX: `/usr/openv/netbackup/bin/bp.kill_all`
    - Windows: `install_path\NetBackup\bin\bpdown -f`
- After you upgrade CloudPoint, if required you can upgrade the NetBackup primary server. Also, you must enable the CloudPoint server from NetBackup Web UI.
- After upgrading, all the CloudPoint servers that you want to use for backup from snapshot or restore from backup jobs, must be re-edited by providing a token so that NetBackup certificates are generated in the CloudPoint server. See *Edit a CloudPoint server* section, in the *NetBackup Web UI Cloud Administrator's Guide*.

Upgrading CloudPoint

The following procedures describe how to upgrade your CloudPoint deployment. During the upgrade, you replace the container that runs your current version of CloudPoint with a newer container.

Upgrade in Docker environment

To upgrade CloudPoint server in Docker environment

1. Download the CloudPoint upgrade installer.

   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_10.x.x.img.gz
   ```

   The numerical sequence in the file name represents the product version.

2. Copy the downloaded compressed image file to the computer on which you want to deploy CloudPoint.

3 Load the image file using the following command:

```
# sudo docker load -i <imagefilename>
```

For example, if the version is `10.0.0.0.9800`, the command syntax is as follows:

```
# sudo docker load -i Veritas_CloudPoint_10.0.0.0.9800.img.gz
```

Messages similar to the following appear on the command line:

```
Load -i VRTScloudpoint-docker-10.0.0.0.9800.img.gz

# docker load -i VRTScloudpoint-docker-10.0.0.0.9800.img.gz
84d4c3ed8c69: Loading layer [==================================================>]
d84981a44f6f: Loading layer [==================================================>]
d16d96b1c9c2: Loading layer [==================================================>]
c3e609767136: Loading layer [==================================================>]
831c43947848: Loading layer [==================================================>]
e4e2d238ca51: Loading layer [==================================================>]
Loaded image: veritas/flexsnap-coordinator:10.0.0.0.9800
b4c897c6bcb4: Loading layer [==================================================>]
41ddf91b45b9: Loading layer [==================================================>]
8e48499e30a3: Loading layer [==================================================>]
Loaded image: veritas/flexsnap-onhostagent:10.0.0.0.9800
b0220c3f236d: Loading layer [==================================================>]
9e90f5eac12e: Loading layer [==================================================>]
98765bae5c6b: Loading layer [==================================================>]
1be3ffeabb75: Loading layer [==================================================>]
3042d315be1a: Loading layer [==================================================>]
0a344789d6a2: Loading layer [==================================================>]
a050f5e0c3b2: Loading layer [==================================================>]
Loaded image: veritas/flexsnap-config:10.0.0.0.9800
e8228e50fe18: Loading layer [==================================================>]
4716779a2c02: Loading layer [==================================================>]
c2f47959d296: Loading layer [==================================================>]
91008c4f5394: Loading layer [==================================================>]
Loaded image: veritas/flexsnap-nginx:10.0.0.0.9800
293db2c3b6b9: Loading layer [==================================================>]
dd80a649929d: Loading layer [==================================================>]
Loaded image: veritas/flexsnap-workflow:10.0.0.0.9800
3f57d7433543: Loading layer [==================================================>]
1f73b2a83aa0: Loading layer [==================================================>]
1e3d4cc3ad29: Loading layer [==================================================>]
1009be131793: Loading layer [==================================================>]
```

```
95a6aa756736: Loading layer [==================================================>
Loaded image: veritas/flexsnap-datamover:10.0.0.0.9800
2ec23a901704: Loading layer [==================================================>
493fc1cdb251: Loading layer [==================================================>
b72568ad4206: Loading layer [==================================================>
fec969040470: Loading layer [==================================================>
54b1b851f710: Loading layer [==================================================>
Loaded image: veritas/flexsnap-cloudpoint:10.0.0.0.9800
Loaded image: veritas/flexsnap-scheduler:10.0.0.0.9800
4de8cab6c086: Loading layer [==================================================>
b355e933b16a: Loading layer [==================================================>
05b7dc006c75: Loading layer [==================================================>
Loaded image: veritas/flexsnap-rabbitmq:10.0.0.0.9800
744c86b54390: Loading layer [==================================================>
1323ffbff4dd: Loading layer [==================================================>
b3c084534040: Loading layer [==================================================>
9e3242667b03: Loading layer [==================================================>
44bf86c64a25: Loading layer [==================================================>
f525c078fda1: Loading layer [==================================================>
3658a337606a: Loading layer [==================================================>
Loaded image: veritas/flexsnap-api-gateway:10.0.0.0.9800
6e61460d1644: Loading layer [==================================================>
ed2435cdcf18: Loading layer [==================================================>
Loa
```

`Loaded image: veritas/flexsnap-cloudpoint:10.0.0.0.9800`

Make a note of the loaded image name and version that appears towards the end of the status messages on the command prompt. 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.

---

4 Make a note of the current CloudPoint version that is installed. You will use the version number in the next step.

5 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 -it --rm -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 the earlier step.

For example, if the installed CloudPoint version is 9.1.0.1.9408, the command will be as follows:

```
# sudo docker run -it --rm -v /cloudpoint:/cloudpoint
-v /var/run/docker.sock:/var/run/docker.sock
veritas/flexsnap-cloudpoint:9.1.0.1.9408 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.8a51aac1848c404ab61e4625d7b88703 ...done
Stopping container: flexsnap-workflow-long-15 ...done
Stopping container: flexsnap-workflow-long-14 ...done
Stopping container: flexsnap-workflow-long-13 ...done
Stopping container: flexsnap-workflow-long-12 ...done
Stopping container: flexsnap-workflow-long-11 ...done
Stopping container: flexsnap-workflow-long-10 ...done
Stopping container: flexsnap-workflow-long-9 ...done
Stopping container: flexsnap-workflow-long-8 ...done
Stopping container: flexsnap-workflow-long-7 ...done
Stopping container: flexsnap-workflow-long-6 ...done
Stopping container: flexsnap-workflow-long-5 ...done
Stopping container: flexsnap-workflow-long-4 ...done
Stopping container: flexsnap-workflow-long-3 ...done
Stopping container: flexsnap-workflow-long-2 ...done
Stopping container: flexsnap-workflow-long-1 ...done
Stopping container: flexsnap-workflow-long-0 ...done
Stopping container: flexsnap-workflow-15 ...done
Stopping container: flexsnap-workflow-14 ...done
Stopping container: flexsnap-workflow-13 ...done
```

```
Stopping container: flexsnap-workflow-12 ...done
Stopping container: flexsnap-workflow-11 ...done
Stopping container: flexsnap-workflow-10 ...done
Stopping container: flexsnap-workflow-9 ...done
Stopping container: flexsnap-workflow-8 ...done
Stopping container: flexsnap-workflow-7 ...done
Stopping container: flexsnap-workflow-6 ...done
Stopping container: flexsnap-workflow-5 ...done
Stopping container: flexsnap-workflow-4 ...done
Stopping container: flexsnap-workflow-3 ...done
Stopping container: flexsnap-workflow-2 ...done
Stopping container: flexsnap-workflow-1 ...done
Stopping container: flexsnap-workflow-0 ...done
Stopping container: flexsnap-nginx ...done
Stopping container: flexsnap-notification ...done
Stopping container: flexsnap-policy ...done
Stopping container: flexsnap-scheduler ...done
Stopping container: flexsnap-idm ...done
Stopping container: flexsnap-onhostagent ...done
Stopping container: flexsnap-agent ...done
Stopping container: flexsnap-coordinator ...done
Stopping container: flexsnap-api-gateway ...done
Stopping container: flexsnap-certauth ...done
Stopping container: flexsnap-rabbitmq ...done
Stopping container: flexsnap-mongodb ...done
Stopping container: flexsnap-fluentd ...done
```

Wait for all the CloudPoint containers to be stopped and then proceed to the next step.

6 Upgrade CloudPoint by running the following command:

```
# sudo docker run -it --rm -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 -it --rm -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 earlier, the command will be as follows:

```
# sudo docker run -it --rm -v /cloudpoint:/cloudpoint
-v /var/run/docker.sock:/var/run/docker.sock
veritas/flexsnap-cloudpoint:10.0.0.0.9800 install -y
```

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

---

7 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:

```
[root@ip-172-31-28-254 ec2-user]# docker run -it --rm -v /cloudpoint:/cloudpoint -v /var/run/doc
Installing the services
Configuration started at time: Thu Jan 13 16:12:30 UTC 2022
WARNING: You're not using the default seccomp profile
docker server version: 1.13.1 Supported: true
WARNING: You're not using the default seccomp profile
This is an upgrade to NetBackup CloudPoint 10.0.0.0.9800
Previous CloudPoint version: 9.1.0.1.9408
Do you want to take a backup of the CloudPoint metadata prior to upgrade? (y/n): y
Removing exited container flexsnap-ipv6config ...done
The containers flexsnap-agent.8d0bd0c9929840d29ae64613b4dd9287 flexsnap-agent.97324e7ceff84b4092
Do you wish to remove them ? (y/n): y
Removing container flexsnap-agent.8d0bd0c9929840d29ae64613b4dd9287 ...done
Removing container flexsnap-agent.97324e7ceff84b4092cefdab1cfcf253 ...done
Removing container flexsnap-workflow-longrun-0-min ...done
Removing container flexsnap-workflow-system-0-min ...done
Removing container flexsnap-workflow-general-0-min ...done
Removing container flexsnap-listener ...done
Removing container flexsnap-nginx ...done
Removing container flexsnap-notification ...done
Removing container flexsnap-policy ...done
Removing container flexsnap-scheduler ...done
Removing container flexsnap-idm ...done
Removing container flexsnap-onhostagent ...done
Removing container flexsnap-agent ...done
Removing container flexsnap-coordinator ...done
Removing container flexsnap-api-gateway ...done
Removing container flexsnap-certauth ...done
Removing container flexsnap-rabbitmq ...done
Removing container flexsnap-mongodb ...done
Removing container flexsnap-fluentd ...done
```

```
Deleting network : flexsnap-network ...done
Taking backup of CloudPoint metadata...done
Backup completed successfully.
Backup file located at /cloudpoint/backup/cloudpoint_9.1.0.1.9408.tar.gz.
Creating network: flexsnap-network ...done
Starting container: flexsnap-fluentd ...done
Starting container: flexsnap-ipv6config ...done
Starting container: flexsnap-mongodb ...done
Starting container: flexsnap-rabbitmq ...done
Waiting for flexsnap-rabbitmq container to move to healthy state...done
Starting container: flexsnap-certauth ...done
Starting container: flexsnap-api-gateway ...done
Starting container: flexsnap-coordinator ...done
Starting container: flexsnap-listener ...done
Starting container: flexsnap-agent ...done
Starting container: flexsnap-onhostagent ...done
Starting container: flexsnap-scheduler ...done
Starting container: flexsnap-policy ...done
Starting container: flexsnap-notification ...done
Starting container: flexsnap-idm ...done
Starting container: flexsnap-config ...done
Starting container: flexsnap-nginx ...done
Finalizing the upgrade..done
Upgrade finished at time: Thu Jan 13 16:14:12 UTC 2022
```

8. (Optional) Run the following command to remove the previous version images.

   ```
   # docker rmi -f <imagename>:<oldimage_tagid>
   ```

9. To verify that the new CloudPoint version is installed successfully:

   See "Verifying that CloudPoint is installed successfully" on page 39.

10. This concludes the upgrade process. Verify that your CloudPoint configuration settings and data are preserved as is.

11. If CloudPoint is not registered with the NetBackup primary server, you must register it now.

    Refer to the *NetBackup Web UI Cloud Administrator's Guide* for instructions.

Upgrade in Podman environment

To upgrade CloudPoint server in Podman environment

1. Download the CloudPoint upgrade installer.

   On the CloudPoint download page, click **Download Now** to download the CloudPoint installer.

   The CloudPoint software components are available in the form of images which are packaged in a compressed file. The file name has the following format:

   ```
   Veritas_CloudPoint_9.x.x.x.x.tar.gz
   ```

   Example: `Veritas_CloudPoint_9.1.0.0.9349.tar.gz`

2. Copy the downloaded compressed image file to the computer on which you want to deploy CloudPoint.

3. Unzip and un-tar the image file and list the contents:

   ```
   # gunzip VRTScloudpoint-podman-9.1.0.0.9349.tar.gz
   ```

   ```
   # tar -xvf VRTScloudpoint-podman-9.1.0.0.9349.tar
   ```

   The output resembles the following:

   ```
   flexsnap-cloudpoint-9.x.x.x.x.img
   flexsnap-coordinator-9.x.x.x.x.img
   flexsnap-agent-9.x.x.x.x.img
   flexsnap-onhostagent-9.x.x.x.x.img
   flexsnap-policy-9.x.x.x.x.img
   flexsnap-scheduler-9.x.x.x.x.img
   flexsnap-config-9.x.x.x.x.img
   flexsnap-certauth-9.x.x.x.x.img
   flexsnap-rabbitmq-9.x.x.x.x.img
   flexsnap-api-gateway-9.x.x.x.x.img
   flexsnap-notification-9.x.x.x.x.img
   flexsnap-fluentd-9.x.x.x.x.img
   flexsnap-nginx-9.x.x.x.x.img
   flexsnap-idm-9.x.x.x.x.img
   flexsnap-workflow-9.x.x.x.x.img
   flexsnap-listener-9.x.x.x.x.img
   flexsnap-datamover-9.x.x.x.x.img
   flexsnap-mongodb-9.x.x.x.x.img
   flexsnap-podman-api.service
   flexsnap-podman-containers.service
   flexsnap_preinstall.sh
   dnsname
   ```

4 Run the following command to prepare the CloudPoint host for installation:

```
# ./flexsnap_preinstall.sh
```

The output resembles the following:

```
Executing the following changes on this node to prepare the CloudPoint
server for installation:
1) Loading CloudPoint service images.
2) Copying dnsname plugin to the /usr/libexec/cni folder.
3) Check if dnsmasq rpm is installed on the host, if not,
   installation is done automatically.
4) Creating and starting the systemd service for the Podman API server.
Do you want to continue? (Yes/No): Yes
Loaded image(s): localhost/veritas/flexsnap-agent:9.1.0.0.9349
Loaded image(s): localhost/veritas/flexsnap-api-gateway:9.1.0.0.9349
Loaded image(s): localhost/veritas/flexsnap-certauth:9.1.0.0.9349
............
............
Loaded image(s): localhost/veritas/flexsnap-workflow:9.1.0.0.9349
Copying dnsname plugin...done
Starting Podman API service...done
```

Note: The output is truncated to fit the page.

5 Verify that there are no protection policy snapshots or other operations in progress and then stop CloudPoint by running the following command:

```
# podman run -it --rm --privileged
-v /cloudpoint:/cloudpoint
-v /run/podman/podman.sock:/run/podman/podman.sock
veritas/flexsnap-cloudpoint:<current_version> stop
```

Here, *current_version* represents the currently installed CloudPoint version, for example '9.0.0.0.9234'

Note: 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-workflow-system-0-0 ...done
Stopping container: flexsnap-workflow-indexing-0-0 ...done
Stopping container: flexsnap-workflow-general-0-0 ...done
Stopping container: flexsnap-listener ...done
Stopping container: flexsnap-nginx ...done
Stopping container: flexsnap-notification ...done
Stopping container: flexsnap-policy ...done
Stopping container: flexsnap-scheduler ...done
Stopping container: flexsnap-idm ...done
Stopping container: flexsnap-onhostagent ...done
Stopping container: flexsnap-agent ...done
Stopping container: flexsnap-coordinator ...done
Stopping container: flexsnap-api-gateway ...done
Stopping container: flexsnap-certauth ...done
Stopping container: flexsnap-rabbitmq ...done
Stopping container: flexsnap-mongodb ...done
Stopping container: flexsnap-fluentd ...done
```

Wait for all the CloudPoint containers to be stopped and then proceed to the next step.

6 Upgrade CloudPoint by running the following command:

```
# podman run -it --rm --privileged
-v /cloudpoint:/cloudpoint
-v /run/podman/podman.sock:/run/podman/podman.sock
veritas/flexsnap-cloudpoint:<new_version> install
```

For an unattended installation, use the following command:

```
# podman run -it --rm --privileged
-v /cloudpoint:/cloudpoint
-v /run/podman/podman.sock:/run/podman/podman.sock
veritas/flexsnap-cloudpoint:<new_version> install -y
```

Here, *new_version* represents the CloudPoint version you are upgrading to, for example '9.1.0.0.9349'

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

Note: Ensure that you enter the command without any line breaks.

---

7 The installer first loads the individual service images and then launches them in their respective containers.

The output resembles the following:

```
Installing the services
Configuration started at time: Mon May  3 11:57:33 UTC 2021
podman server version: 2.0.5 Supported: true
This is an upgrade to NetBackup CloudPoint 9.1.0.0.9349
Previous CloudPoint version: 9.0.0.0.9234
Do you want to take a backup of the CloudPoint metadata prior to upgrade?
(y/n): y
Taking backup of CloudPoint metadata...done
Backup completed successfully.
Backup file located at /cloudpoint/backup/cloudpoint_9.0.0.0.9234.tar.gz.
[Storing /cloudpoint/keys/idm_store]
[Storing /cloudpoint/keys/flexsnap-idm_store]
Creating network: flexsnap-network ...done
Starting container: flexsnap-fluentd ...done
Starting container: flexsnap-mongodb ...done
Starting container: flexsnap-rabbitmq ...done
Starting container: flexsnap-certauth ...done
Starting container: flexsnap-api-gateway ...done
Starting container: flexsnap-coordinator ...done
Starting container: flexsnap-listener ...done
Starting container: flexsnap-agent ...done
Starting container: flexsnap-onhostagent ...done
Starting container: flexsnap-scheduler ...done
Starting container: flexsnap-policy ...done
Starting container: flexsnap-notification ...done
Starting container: flexsnap-idm ...done
Starting container: flexsnap-config ...done
Starting container: flexsnap-nginx ...done
Upgrade finished at time: Mon May  3 11:58:51 UTC 2021
Before using backups from cloud snapshots, re-register CloudPoint with the
NetBackup primary server
```

8 (Optional) Run the following command to remove the previous version images.

```
# podman rmi -f <imagename>:<oldimage_tagid>
```

9 To verify that the new CloudPoint version is installed successfully:

See "Verifying that CloudPoint is installed successfully" on page 39.

10 This concludes the upgrade process. Verify that your CloudPoint configuration settings and data are preserved as is.

11 If CloudPoint is not registered with the NetBackup primary server, you must register it now.

Refer to the *NetBackup Web UI Cloud Administrator's Guide* for instructions.

Upgrading CloudPoint using patch or hotfix

You can also upgrade your current CloudPoint server using a patch or a hotfix. All the considerations and steps that apply for a normal upgrade, also apply to the upgrade being done using a patch or a hotfix, except that instead of downloading a new CloudPoint image, you download the patch/hotfix binaries.

Contact Veritas Technical Support at https://www.veritas.com/content/support/en_US/contact-us to obtain an Emergency Engineering Binary (EEB) for patch/hotfix.

Following are the brief steps explained with an example. For the detailed upgrade procedures

See "Upgrading CloudPoint" on page 221.

Consider that the currently installed version is CloudPoint `9.1.0.0.9344` and you are upgrading to a CloudPoint patch version `9.1.0.0.9349` on a RHEL8.3 system in a Podaman environment. The same steps apply for the Docker environment with the appropriate Docker commands.

To upgrade CloudPoint using a patch or a hotfix

1 Download the CloudPoint EEB obtained from Veritas Technical Support.

Example: `Veritas_CloudPoint_9.1.0.0.9349.img.gz`

2 Unzip and un-tar the binaries and list the contents:

```
# gunzip VRTScloudpoint-podman-9.1.0.0.9349.tar.gz
```

```
# tar -xvf VRTScloudpoint-podman-9.1.0.0.9349.tar
```

3 Run the following command to prepare the CloudPoint host for installation:

```
# ./flexsnap_preinstall.sh
```

4 Verify that there are no protection policy snapshots or other operations in progress and then stop CloudPoint by running the following command:

```
# podman run -it --rm --privileged
-v /cloudpoint:/cloudpoint
-v /run/podman/podman.sock:/run/podman/podman.sock
veritas/flexsnap-cloudpoint:<current_version> stop
```

5 Upgrade CloudPoint by running the following command:

```
# podman run -it --rm --privileged
-v /cloudpoint:/cloudpoint
-v /run/podman/podman.sock:/run/podman/podman.sock
veritas/flexsnap-cloudpoint:<EEB_version>install
```

For an unattended installation, use the following command:

```
# podman run -it --rm --privileged
-v /cloudpoint:/cloudpoint
-v /run/podman/podman.sock:/run/podman/podman.sock
veritas/flexsnap-cloudpoint:<EEB_version>install -y
```

Here, *EEB_version* represents the CloudPoint patch/hotfix version you are upgrading to, for example '9.1.0.0.9349'

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

Note: Ensure that you enter the command without any line breaks.

The installer first loads the individual service images and then launches them in their respective containers.

6 (Optional) Run the following command to remove the previous version images.

```
# podman rmi -f <imagename>:<oldimage_tagid>
```

7 To verify that the new CloudPoint version is installed successfully:

See "Verifying that CloudPoint is installed successfully" on page 39.

8 This concludes the CloudPoint upgrade process using a patch or a hotfix . Verify that your CloudPoint configuration settings and data are preserved as is.

9 If CloudPoint is not registered with the NetBackup primary server, you must register it now.

Refer to the *NetBackup Web UI Cloud Administrator's Guide* for instructions.

Migrating and upgrading CloudPoint

Before you begin migrating CloudPoint

Make sure that you complete the following before installing CloudPoint:

- Ensure that your environment meets system requirements.
  See " Meeting system requirements" on page 14.
- Create the instance on which you install CloudPoint or prepare the physical host.
  See "Creating an instance or preparing the host to install CloudPoint" on page 28.
- Prepare a RHEL 8.3 or 8.4 host for installation. You can either upgrade your existing RHEL 7.x OS to RHEL 8.3 or 8.4 OS, or create a new system with RHEL 8.3 or 8.4.
  - For upgrading the system from RHEL 7.x to RHEL 8.3 or 8.4, follow the Red Hat documentation:
    https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html/upgrading_from_rhel_7_to_rhel_8/index
  - For creating a new system with RHEL 8.3 or 8.4, configure a Podman container platform
    See Table 1-14 on page 28.
    The brief steps include:
    - Setup the RHEL repos
      For AWS cloud, enable the extra repos
      ```
      # sudo yum-config-manager --enable
      rhui-REGION-rhel-server-extras
      ```
      For on-premise, enable your subscriptions:
      ```
      # sudo subscription-manager register --auto-attach
      --username=<username> --password=<password>
      ```
    - Install Podman if required:
      ```
      # sudo yum install -y podman
      ```
    - If SELinux is enabled, change the mode to permissive mode and restart the system.
      Edit the `/etc/selinux/config` configuration file and modify the `SELINUX` parameter value to `SELINUX=permissive`.
- Run the following commands to install the required packages (`lvm2`, `udev` and `dnsmaq`) on the hosts:
  ```
  #yum install -y lvm2-<version>
  #yum install -y lvm2-libs-<version>
  ```

```
#yum install -y python3-pyudev-<version>
#yum install -y systemd-udev-<version>
#yum install -y dnsmasq-<version>
```

- Run the following commands to lock the Podman and Conmon versions to the supported versions, so that they do not get updated with the *yum* update:

```
sudo yum install -y podman-2.2.1-7.module+el8.3.1+9857+68fb1526
sudo yum install -y conmon-2:2.0.20-2.module+el8.3.0+8221+97165c3f
sudo yum install -y python3-dnf-plugin-versionlock
sudo yum versionlock podman* conmon*
```

- 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 32.

Next, you migrate CloudPoint from the RHEL 7.x host to the newly prepared RHEL 8.3 or 8.4 host.

See "Migrate and upgrade CloudPoint on RHEL 8.5 or 8.4" on page 238.

Migrate and upgrade CloudPoint on RHEL 8.5 or 8.4

Perform the following steps to migrate CloudPoint 9.1 or 9.1.0.1 from your RHEL 7.x host to the new RHEL 8.5 or 8.4 host.

To upgrade CloudPoint in Podman environment

1. Download the CloudPoint upgrade installer.

   Example: `Veritas_CloudPoint_10.0.0.9818.img.gz`

2. Unzip and un-tar the image file and list the contents:

```
# gunzip VRTScloudpoint-podman-10.0.0.9818.tar.gz
# tar -xvf VRTScloudpoint-podman-10.0.0.9818.tar
```

3. Run the following command to prepare the CloudPoint host for installation:

```
# ./flexsnap_preinstall.sh
```

4 Upgrade CloudPoint by running the following command:

```
# podman run -it --rm --privileged
-v /cloudpoint:/cloudpoint
-v /run/podman/podman.sock:/run/podman/podman.sock
veritas/flexsnap-cloudpoint:<new_version> install
```

For an unattended installation, use the following command:

```
# podman run -it --rm --privileged
-v /cloudpoint:/cloudpoint
-v /run/podman/podman.sock:/run/podman/podman.sock
veritas/flexsnap-cloudpoint:<new_version> install -y
```

Here, *new_version* represents the CloudPoint version you are upgrading to, for example `10.0.0.9818`

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

Note: Ensure that you enter the command without any line breaks.

The installer first loads the individual service images and then launches them in their respective containers.

5 (Optional) Run the following command to remove the previous version images.

```
# podman rmi -f <imagename>:<oldimage_tagid>
```

6 To verify that the new CloudPoint version is installed successfully:

See "Verifying that CloudPoint is installed successfully" on page 39.

To migrate CloudPoint

1. On the RHEL 7.x host, 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 -it --rm
-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.

Example:

```
# sudo docker run -it --rm -v /cloudpoint:/cloudpoint
-v /var/run/docker.sock:/var/run/docker.sock
veritas/flexsnap-cloudpoint:9.1.0.0.9349 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.8a51aac1848c404ab61e4625d7b88703 ...done
Stopping container: flexsnap-workflow-long-15 ...done
Stopping container: flexsnap-workflow-long-14 ...done
Stopping container: flexsnap-workflow-long-13 ...done
Stopping container: flexsnap-workflow-long-12 ...done
Stopping container: flexsnap-workflow-long-11 ...done
Stopping container: flexsnap-workflow-long-10 ...done
Stopping container: flexsnap-workflow-long-9 ...done
Stopping container: flexsnap-workflow-long-8 ...done
Stopping container: flexsnap-workflow-long-7 ...done
Stopping container: flexsnap-workflow-long-6 ...done
Stopping container: flexsnap-workflow-long-5 ...done
Stopping container: flexsnap-workflow-long-4 ...done
Stopping container: flexsnap-workflow-long-3 ...done
Stopping container: flexsnap-workflow-long-2 ...done
Stopping container: flexsnap-workflow-long-1 ...done
Stopping container: flexsnap-workflow-long-0 ...done
Stopping container: flexsnap-workflow-15 ...done
Stopping container: flexsnap-workflow-14 ...done
Stopping container: flexsnap-workflow-13 ...done
```

```
Stopping container: flexsnap-workflow-12 ...done
Stopping container: flexsnap-workflow-11 ...done
Stopping container: flexsnap-workflow-10 ...done
Stopping container: flexsnap-workflow-9 ...done
Stopping container: flexsnap-workflow-8 ...done
Stopping container: flexsnap-workflow-7 ...done
Stopping container: flexsnap-workflow-6 ...done
Stopping container: flexsnap-workflow-5 ...done
Stopping container: flexsnap-workflow-4 ...done
Stopping container: flexsnap-workflow-3 ...done
Stopping container: flexsnap-workflow-2 ...done
Stopping container: flexsnap-workflow-1 ...done
Stopping container: flexsnap-workflow-0 ...done
Stopping container: flexsnap-nginx ...done
Stopping container: flexsnap-notification ...done
Stopping container: flexsnap-policy ...done
Stopping container: flexsnap-scheduler ...done
Stopping container: flexsnap-idm ...done
Stopping container: flexsnap-onhostagent ...done
Stopping container: flexsnap-agent ...done
Stopping container: flexsnap-coordinator ...done
Stopping container: flexsnap-api-gateway ...done
Stopping container: flexsnap-certauth ...done
Stopping container: flexsnap-rabbitmq ...done
Stopping container: flexsnap-mongodb ...done
Stopping container: flexsnap-fluentd ...done
```

Wait for all the CloudPoint containers to be stopped and then proceed to the next step.

2. Migrate the CloudPoint configuration data to the RHEL 8.5 or 8.4 host:
   - If you have upgraded from RHEL 7.x to RHEL 8.5 or 8.4, copy the `/cloudpoint` mountpoint data from RHEL 7.x system and move it to the RHEL8.5 or 8.4 system under `/cloudpoint` folder.
   - If you have created a new system with RHEL 8.5 or 8.4:
     - Run the following command to unmount `/cloudpoint` from the current host.

       ```
       # umount /cloudpoint
       ```

     - Detach the data disk that was mounted on `/cloudpoint` mountpoint.

Note: For detailed instructions to detach or attach the data disks, follow the documentation provided by your cloud or storage vendor.

- On the RHEL8.5 or 8.4 host, run the following commands to create and mount the disk:

  ```
  # mkdir /cloudpoint
  # mount /dev/<diskname> /cloudpoint
  ```

  For vendor-specific details
  See "Creating and mounting a volume to store CloudPoint data" on page 30.

This concludes the CloudPoint migration process.

After migration, install the new_version on the new host by following the steps mentioned in the To upgrade CloudPoint in Podman environment.

3 During migration process, if CloudPoint server is migrated to another system or IP address is changed, then regenerate the certificates as follows:

- Stop the CloudPoint services using the following command:

  ```
  [root@ip-172-31-24-178 ec2-user]# podman run -it --rm
  --privileged -v /cloudpoint:/cloudpoint -v
  /run/podman/podman.sock:/run/podman/podman.sock
  veritas/flexsnap-cloudpoint:10.0.0.9818 stop
  ```

- Regenerate the certificates using the following command:

  ```
  # podman run -it --rm -v /cloudpoint:/cloudpoint --entrypoint
  "/cloudpoint/scripts/cp_regenerate_certs.sh"
  veritas/flexsnap-cloudpoint:10.0.0.9818 -i <CP_IP_ADDRESS> -h
  <CP_HOSTNAME>

  Setting up certificate authority ...done
  Generating certificates for servers ...done
  Generating certificates for clients ...done
  Adding MongoDB and RabbitMQ certificate to the trust store ...[Storing
  [Storing /cloudpoint/keys/flexsnap-idm_store]
  done
  Creating symlinks for nginx certificates ...done
  ```

- Start the CloudPoint services using the following command:

  ```
  [root@ip-172-31-24-178 ec2-user]# podman run -it --rm
  --privileged -v /cloudpoint:/cloudpoint -v
  /run/podman/podman.sock:/run/podman/podman.sock
  veritas/flexsnap-cloudpoint:10.0.0.9818 start
  ```

4 Depending on the following appropriate scenario, update the `/cloudpoint/openv/netbackup/bp.conf` file to update the value of **CLIENT_NAME** to new CloudPoint IP/hostname.

- If IP address doe not change, then edit the CloudPoint server entry and provide reissue token.
- If IP address changes, disable previous CloudPoint host, and add CloudPoint server with new IP. And then edit the entry with reissue token.

5 After migrating CloudPoint to a RHEL 8.5 or 8.4 host, perform the following steps to upgrade CloudPoint to 10.0.

Refer to the detailed procedure *"To upgrade CloudPoint server in Podman environment"*.

See "Upgrading CloudPoint" on page 221.

6 This concludes the migration and upgrade process for CloudPoint. Verify that your CloudPoint configuration settings and data are preserved as is.

7 If CloudPoint is not registered with the NetBackup primary server, you must register it now.

Refer to the *NetBackup Web UI Cloud Administrator's Guide* for instructions.

Post-upgrade tasks

You may need to perform the following tasks after a successful upgrade of the CloudPoint server.

Post-upgrade tasks

1 Upgrade the CloudPoint agents on the Linux and Windows application hosts.

Note: If you are upgrading from CloudPoint 8.3 to 9.0 or 9.1, then you must manually upgrade the on-host agents. If you are upgrading from CloudPoint 9.0 to 9.1, upgrading the on-host agents is optional.

Perform the following steps to upgrade the agent on Linux hosts:

- Sign in to NetBackup UI and download the newer agent package. Navigate to **Cloud > CloudPoint servers > Actions > Add agent**.
- Stop the flexsnap agent service on the Linux host where you want to upgrade the agent.
  Run the following command on the Linux host:

```
# sudo systemctl stop flexsnap-agent.service
```

- Upgrade the agent on the Linux host.
  Run the following command on the Linux host:
  ```
  # sudo rpm -Uvh --force cloudpoint_agent_rpm_name
  ```
  Here, *cloudpoint_agent_rpm_name* is the name of the agent rpm package you downloaded earlier.
- Generate the token for agent configuration. Navigate to **NetBackup Web UI > Cloud > CloudPoint Servers > Actions > Add agent > Create Token**.
- Start the flexsnap agent service on the Linux host.
  Run the following command on the Linux host:
  ```
  # sudo systemctl start flexsnap-agent.service --renew --token
  <auth_token>
  ```
- Reload the daemon, if prompted.
  Run the following command on the Linux host:
  ```
  # sudo systemctl daemon-reload
  ```
- Repeat these steps on all the Linux hosts where you wish to upgrade the Linux-based agent.

Perform the following steps to upgrade the agent on Windows hosts:

- Sign in to NetBackup UI and download the newer agent package.
  Navigate to **Cloud > CloudPoint servers > Actions > Add agent**.
- Stop the Veritas CloudPoint Agent service that is running on the host.
- Run the newer version of 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.
- Generate the token for agent configuration. Navigate to **NetBackup Web UI > Cloud > CloudPoint Servers > Actions > Add agent > Create Token**.
- ([Not applicable when upgrading from 9.1.x to 10.0]) Register the agent on the host again.
  From the command prompt, navigate to the agent installation directory (C:\Program Files\Veritas\CloudPoint\) and run the following command:
  ```
  # flexsnap-agent.exe --renew --token <auth_token>
  ```
- Repeat these steps on all the Windows hosts where you wish to upgrade the Windows-based agent.

For details on how to download the agent installation package from the NetBackup UI, refer to the following:

See "Downloading and installing the CloudPoint agent" on page 169.

2 If you want to run backup from snapshot and restore from backup jobs after upgrade, you must update the NetBackup configuration so that the upgraded CloudPoint configuration details are available with NetBackup. After upgrading, all the CloudPoint servers that you want to use for backup from snapshot or restore from backup jobs, must be re-edited by providing a token so that NetBackup certificates are generated. See *Edit a CloudPoint server* section, in the *NetBackup Web UI Cloud Administrator's Guide*.

Perform one of the following actions:

- From the NetBackup Web UI, edit the CloudPoint server information.
  - In the Web UI, click **Workloads > Cloud** from the left navigation pane and then click the **CloudPoint servers** tab.
  - Select the CloudPoint server that you just upgraded, and then click **Edit** from the ellipsis action button on the right.
  - In the Edit CloudPoint server dialog, specify all the requested details.
  - Click **Validate** to validate the CloudPoint server certificate.
  - In the **Token** field enter the Standard Host Token.
  - Click **Save** to update the CloudPoint server configuration.
- Or, on the NetBackup primary server, run the following command:

```
# ./tpconfig -update -cloudpoint_server
cp-hostname-cloudpoint_server_user_id admin -manage_workload
<manage_workload>
```

Note: Additional option `-security_token` is required for updating CloudPoint which is managing cloud workloads. The token must be Standard host token. This is required for NetBackup certificates generation on CloudPoint.

On UNIX systems, the directory path to this command is `/usr/openv/volmgr/bin/`. On Windows systems, the directory path to this command is `install_path\Volmgr\bin\`. Refer to the *Veritas NetBackup Commands Reference Guide* for details.

- Or, make a PATCH API call to the NetBackup primary server using the following URL:

`https://primaryserver.domain.com/netbackup/config/servers/snapshot-mgmt-servers/cp-hostname`

For more details about the `tpconfig` command and its options, refer to the *Veritas NetBackup Commands Reference Guide*.

Chapter 12

Uninstalling CloudPoint

This chapter includes the following topics:

Preparing to uninstall CloudPoint

Note the following before you uninstall CloudPoint:

- Ensure that there are no active CloudPoint operations in progress. For example, if there are any snapshot, replication, restore or indexing 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.
- Ensure that you remove the CloudPoint agents that are installed on the application hosts. The application hosts are the systems where the applications that are being protected by CloudPoint are running.
  See “Removing the CloudPoint agents” on page 254.

- Ensure that you disable the CloudPoint server from NetBackup. Depending on how you have set up your CloudPoint server, whether on-premise or in the cloud, you can disable CloudPoint server either from the NetBackup Web UI or from the NetBackup Administration console (Java UI).
  Refer to the *NetBackup Web UI Backup Administrator's Guide* or the *NetBackup Snapshot Client Administrator's Guide* for instructions.
- 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 deleted after the uninstallation. You can take a backup of all the data in the `/cloudpoint` volume, if desired.
  See "Backing up CloudPoint" on page 250.

Backing up CloudPoint

If CloudPoint is deployed in a cloud

To back up CloudPoint when it is deployed in a cloud

1 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. You can retrieve the version using the following command:

```
# cat /cloudpoint/version
```

For example:

```
# sudo docker run -it --rm -v /cloudpoint:/cloudpoint -v
/var/run/docker.sock:/var/run/docker.sock
veritas/flexsnap-cloudpoint:8.3.0.8549 stop
```

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

2 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:

```
# sudo docker ps | grep veritas
```

This command should not return any actively running CloudPoint containers.

3 (Optional) If you still see any active containers, repeat step 2. If that does not work, run the following command on each active container:

```
# sudo docker kill container_name
```

For example:

```
# sudo docker kill flexsnap-api
```

4. After all the containers are stopped, take a snapshot of the volume on which you installed CloudPoint. Use the cloud provider's snapshot tools.

5. 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:8.3.0.8549 start
   ```

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

If CloudPoint is deployed on-premises

To backup CloudPoint when it is deployed on-premise

1. 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:8.3.0.8549 stop
   ```

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

2. 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:

   ```
   # sudo docker ps | grep veritas
   ```

   This command should not return any actively running CloudPoint containers.

3. (Optional) If you still see any active containers, repeat step 2. If that does not work, run the following command on each active container:

   ```
   # sudo docker kill container_name
   ```

   For example:

   ```
   # sudo docker kill flexsnap-api
   ```

4. 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.

Unconfiguring CloudPoint plug-ins

CloudPoint plug-ins allow CloudPoint to discover the assets on the host so that you can protect those assets by taking snapshots. If required, you can remove a CloudPoint plug-in configuration using the NetBackup UI.

Before you remove a plug-in configuration from the host, consider the following:

- You must remove all the snapshots of the assets that are related to the plug-in that you wish to unconfigure.
  Plug-in unconfiguration fails if asset snapshots exist.
- Unconfiguring a plug-in removes the plug-in from the selected host. To protect the plug-in related assets on the same host again, you will have to reconfigure that plug-in on the host.
- Once you unconfigure a plug-in, all the assets that are related to the plug-in are removed from the CloudPoint configuration and you will no longer be able to protect those assets.

To unconfigure a plug-in from a host

1. Sign in to the NetBackup UI.
2. Verify that you have removed all the plug-in related asset snapshots.
3. From the menu on the left, click **Workloads > Cloud** and then click the **Virtual machines** tab.
4. On the Virtual machines tab, select the host where you want unconfigure the agent and then from the menu bar that appears at the top, click **Unconfigure**.

   CloudPoint unconfigures the plug-in from the host. Observe that the **Unconfigure** button now changes to **Configure**. This indicates that the plug-in unconfiguration is successful on the host.

Unconfiguring CloudPoint agents

To enable CloudPoint to protect assets on a remote host, you first need to establish a connection between the CloudPoint server and the remote host. Depending on how the connection is configured (either with agents or using the agentless feature), CloudPoint uses agents that manage the plug-ins that are used to discover all the assets and perform the operations on the host.

Whenever you configure a remote host for protection, the agent registration and the plug-in configuration information is added to the CloudPoint database on the CloudPoint server. You can, if required, remove an agent entry from the CloudPoint database by performing the disconnect operation from the NetBackup UI.

Before you unconfigure an agent, consider the following:

- Once you unconfigure an agent, you cannot re-configure a CloudPoint plug-in on the same host, if you had installed the CloudPoint agent on that host. To be able to configure a plug-in on the host again, you must first uninstall the agent package from the host, connect the host and install and register the agent with the CloudPoint server again.
- You must first unconfigure the CloudPoint plug-in from the host before you proceed with the disconnect operation. The disconnect option is not enabled if a CloudPoint plug-in is configured on the host.
- Unconfiguring an agent entry from the CloudPoint server does not uninstall the agent package from the host. You have to manually remove the agent binaries from the host after completing the disconnect operation.
- Once you unconfigure an agent, all the file system assets that belong to that host are removed from the CloudPoint configuration.

To unconfigure the agent entry from the CloudPoint server

1. Sign in to the NetBackup UI.
2. Remove CloudPoint plug-in configuration from the host that you wish to disconnect.

   See "Unconfiguring CloudPoint plug-ins" on page 253.
3. From the menu on the left, click **Workloads > Cloud** and then click the **Virtual machines** tab.
4. On the Virtual machines tab, select the host where you want unconfigure the agent and then from the menu bar that appears at the top, click **Disconnect**.

   CloudPoint begins to unconfigure the agent. Observe that the Disconnect button now changes to Connect. This indicates that the disconnect operation is successful and the agent has been unconfigured successfully.

   The agent registration and all the assets information about that host is completely removed from the database.
5. The next step is to manually uninstall the agent from the host on which you performed the disconnect operation. This is required if you wish to protect this host and its assets using CloudPoint at a later time.

   See "Removing the CloudPoint agents" on page 254.

Removing the CloudPoint agents

You must first remove the CloudPoint agents before you remove CloudPoint. The agents are installed directly on the host where the applications are running.

CloudPoint agents manage the CloudPoint plug-ins that discover assets and perform snapshot operations on the host.

To uninstall the CloudPoint on-host agents

1. Connect to the host where you have installed the CloudPoint agent.

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

2. For Linux-based 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 agent rpm package, without the version number and the file extension (`.rpm`).

   For example, if the name of the 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 agent, do the following:

   From Windows Control Panel > Programs and Features, select the entry for the CloudPoint agent (**Veritas CloudPoint Agent**) and then click **Uninstall**.

   Follow the wizard workflow to uninstall the agent from the Windows instance.

   Note: To allow the uninstallation, admin users will have to click Yes on the Windows UAC prompt. Non-admin users will have to specify admin user credentials on the UAC prompt.

   ---

4. This completes the agent uninstallation.

   You can now proceed to uninstall CloudPoint.

   See "Removing CloudPoint from a standalone Docker host environment" on page 255.

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 agents from all the hosts that are part of the CloudPoint configuration.

   See "Removing the CloudPoint agents" on page 254.

2. Verify that there are no protection policy snapshots or other operations in progress, and then 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
   ```

   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 `/cloudpoint`.

   For example, if the product version is `8.3.0.8549`, 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:8.3.0.8549 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:8.3.0.8549 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
```

3. Confirm that the CloudPoint containers are removed.

   Use the following docker command:

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

4. 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>
   ```

5. This completes the CloudPoint uninstallation on the host.

   Possible next step is to re-deploy CloudPoint.

   See “Installing CloudPoint in the Docker environment” on page 34.

Removing CloudPoint extensions - VM-based or managed Kubernetes cluster-based

During uninstallation, the installer performs the following tasks on the CloudPoint extension host:

- Stops all the CloudPoint containers that are running
- Removes the CloudPoint containers

To uninstall a VM-based extension

1 **For Docker environment:**

Run 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> uninstall
```

Example:

```
# sudo docker run -it --rm
-v /cloudpoint:/cloudpoint
-v /var/run/docker.sock:/var/run/docker.sock
veritas/flexsnap-cloudpoint:<9.1.x.x.xxx> uninstall
```

Note: This is a single command without any line breaks.

For Podman environment:

Run the following command:

```
# podman run -it --rm --privileged
-v /<full_path_to_volume_name>:/<full_path_to_volume_name>
-v /run/podman/podman.sock:/run/podman/podman.sock
veritas/flexsnap-cloudpoint:<version> uninstall
```

Example:

```
# podman run -it --rm --privileged
-v /cloudpoint:/cloudpoint
-v /run/podman/podman.sock:/run/podman/podman.sock
veritas/flexsnap-cloudpoint:<9.1.x.x.xxx> uninstall
```

Note: This is a single command without any line breaks.

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 `/cloudpoint`.

2 If desired, remove the CloudPoint container images from the extension host.

Use the following docker command to view the docker images that are loaded on the host and remove the CloudPoint images based on their IDs.

```
# sudo docker images -a
```

```
# sudo docker rmi <image ID>
```

This completes the CloudPoint extension uninstallation on a VM host.

To uninstall a managed Kubernetes cluster-based extension

- Execute the extension script `cp_extension_start.sh` that was downloaded at the time of extension installation, from the host where `kubectl` is installed.

  Run the following command:

  ```
  bash cp_extension_start.sh uninstall
  ```

  Once the uninstallation is triggered, provide the namespace as an input, from which the extension services need to be uninstalled.

After the uninstallation, the provisioned cloud resources associated with the uninstalled extension can be terminated or removed.

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 container platform (Docker, Podman)” on page 28.

4 Attach the newly-created volume to this CloudPoint server instance.

5 Create the CloudPoint installation directory on this server.

Use the following command:

```
# mkdir /full_path_to_cloudpoint_installation_directory
```

For example:

```
# mkdir /cloudpoint
```

6 Mount the attached volume to the installation directory you just created.

Use the following command:

```
# mount /dev/device-name
/full_path_to_cloudpoint_installation_directory
```

For example:

```
# mount /dev/xvdb /cloudpoint
```

7 Verify that all CloudPoint related configuration data and files are in the directory.

Enter the following command:

```
# ls -l /cloudpoint
```

8 Download or copy the CloudPoint installer binary to the new server.

9 Install CloudPoint.

Use the following command:

```
# sudo docker run -it --rm
-v /cloudpoint:/cloudpoint
-v /var/run/docker.sock:/var/run/docker.sock
veritas/flexsnap-cloudpoint:8.3.1.5300 install
```

Here, `8.3.1.5300` 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 Wed May 13 22:20:47 UTC 2020
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 -zxf 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:8.3.1.5300 install
```

Here, `8.3.1.5300` 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 Wed May 13 22:20:47 UTC 2020
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.

4 When the installation completes, you can resume working with CloudPoint using your existing credentials.

Chapter 13

Troubleshooting CloudPoint

This chapter includes the following topics:

Troubleshooting CloudPoint

Refer to the following troubleshooting scenarios:

- **CloudPoint agent fails to connect to the CloudPoint server if the agent host is restarted abruptly.**
  This issue may occur if the host where the CloudPoint agent is installed is shut down abruptly. Even after the host restarts successfully, the agent fails to establish a connection with the CloudPoint server and goes into an offline state.
  The agent log file contains the following error:

  ```
  Flexsnap-agent-onhost[4972] mainthread
  flexsnap.connectors.rabbitmq: error - channel 1 closed
  unexpectedly: (405) resource_locked - cannot obtain exclusive
  access to locked queue '
  flexsnap-agent.a1f2ac945cd844e393c9876f347bd817' in vhost '/'
  ```

  This issue occurs because the RabbitMQ connection between the agent and the CloudPoint server does not close even in case of an abrupt shutdown of the agent host. The CloudPoint server cannot detect the unavailability of the agent until the agent host misses the heartbeat poll. The RabbitMQ connection remains open until the next heartbeat cycle. If the agent host reboots before the next heartbeat poll is triggered, the agent tries to establish a new connection with the CloudPoint server. However, as the earlier RabbitMQ connection already exists, the new connection attempt fails with a resource locked error.

As a result of this connection failure, the agent goes offline and leads to a failure of all snapshot and restore operations performed on the host.

Workaround:

Restart the Veritas CloudPoint Agent service on the agent host.

- On a Linux hosts, run the following command:

  ```
  # sudo systemctl restart flexsnap-agent.service
  ```

- On Windows hosts:

  Restart the `Veritas CloudPoint™ Agent` service from the Windows Services console.

- **CloudPoint agent registration on Windows hosts may time out or fail.**

  For protecting applications on Windows, you need to install and then register the CloudPoint agent on the Windows host. The agent registration may sometimes take longer than usual and may either time out or fail.

  Workaround:

  To resolve this issue, try the following steps:

  - Re-register the agent on the Windows host using a fresh token.
  - If the registration process fails again, restart the CloudPoint services on the CloudPoint server and then try registering the agent again.

  Refer to the following for more information:

  See "Registering the Windows-based agent" on page 174.
  See "Restarting CloudPoint" on page 41.

- **Disaster recovery when DR package is lost or passphrase is lost.**

  This issue may occur if the DR package is lost or the passphrase is lost.

  In case of Catalog backup, 2 backup packages are created:

  - DR package which contains all the certs
  - Catalog package which contains the data base

  The DR package contains the NetBackup UUID certs and Catalog DB also has the UUID. When you perform disaster recovery using the DR package followed by catalog recovery, both the UUID cert and the UUID are restored. This allows NetBackup to communicate with CloudPoint since the UUID is not changed.

  However if the DR package is lost or the Passphrase is lost the DR operation cannot be completed. You can only recover the catalog without DR package after you reinstall NetBackup. In this case, a new UUID is created for NetBackup which is not recognised by CloudPoint. The one-to-one mapping of NetBackup and CloudPoint is lost.

  Workaround:

  To resolve this issue, you must update the new NBU UUID and Version Number after NetBackup primary is created.

- The NetBackup administrator must be logged on to the NetBackup Web Management Service to perform this task. Use the following command to log on:
  `/usr/openv/netbackup/bin/bpnbat -login -loginType WEB`
- Execute the following command on the primary server to get the NBU UUID:
  `/usr/openv/netbackup/bin/admincmd/nbhostmgmt -list -host <primary server host name> | grep "Host ID"`
- Execute the following command to get the Version Number:
  `/usr/openv/netbackup/bin/admincmd/bpgetconfig -g <primary Ssrver host name> -L`

After you get the NBU UUID and Version number, execute the following command on the CloudPoint host to update the mapping:

`/cloudpoint/scripts/cp_update_nbuuid.sh -i <NBU UUID> -v <Version Number>`

- **The snapshot job is successful but backup job fails with error "The CloudPoint server's certificate is not valid or doesn't exist.(9866)" when ECA_CRL_CHECK disabled on master server.**
  If ECA_CRL_CHECK is configured on master server and is disabled then it must be configured in `bp.conf` on CloudPoint setup with same value.
  For example, considering a scenario of backup from snapshot where NetBackup is configured with external certificate and certificate is revoked. In this case, if ECA_CRL_CHECK is set as DISABLE on master then set the same value in `bp.conf` of CloudPoint setup, otherwise snapshot operation will be successful and backup operation will fail with the certificate error.
  See “Configuring security for Azure and Azure Stack ” on page 206.
- **CloudPoint fails to establish connection using agentless to the Windows cloud instance**
  Error 1: `<Instance_name>: network connection timed out.`
  Case 1: CloudPoint server log message:

```
WARNING - Cannot connect to the remote host. SMB Connection timeout
 <IP address> <user>

…

flexsnap.OperationFailed: Could not connect to the remote server
<IP address>
```

  Workaround
  To resolve this issue, try the following steps:

- Verify if the SMB port 445 is added in the Network security group and is accessible from the CloudPoint server.
- Verify if the SMB port 445 is allowed through cloud instance firewall.

Case 2: CloudPoint Server log message:

```
WARNING - Cannot connect to the remote host. WMI Connection
timeout <IP address> <user>

…

flexsnap.OperationFailed: Could not connect to the remote
server <IP address>
```

Workaround:
To resolve this issue, try the following steps:

- Verify and add DCOM port (135) in the Network security group and is accessible from CloudPoint server.
- Verify if the port 135 is allowed through cloud instance firewall.

Case 3: CloudPoint Server log message:

```
Exception while opening SMB connection, [Errno Connection error
(<IP address>:445)] [Errno 113] No route to host.
```

Workaround: Verify if the cloud instance is up and running or not in inconsistent state.
Case 4: CloudPoint Server log message:

```
Error when closing dcom connection: 'Thread-xxxx'"
```

Where, *xxxx* is the thread number.
Workaround:
To resolve this issue, try the following steps:

- Verify if the WMI-IN dynamic port range or the fixed port as configured is added in the Network security group.
- Verify and enable WMI-IN port from the cloud instance firewall.

Error 2: `<Instance_name>: Could not connect to the virtual machine.`
CloudPoint server log message:

```
Error: Cannot connect to the remote host. <IP address> Access denied.
```

Workaround:
To resolve this issue, try the following steps:

- Verify if the user is having administrative rights.
- Verify if the UAC is disabled for the user.

- **CloudPoint cloud operations fail on a RHEL system if a firewall is disabled**
  The CloudPoint operations fail for all the supported cloud plugins on a RHEL system, if a firewall is disabled on that system when the CloudPoint services are running. This is a network configuration issue that prevents the CloudPoint from accessing the cloud provider REST API endpoints.
  Workaround
  - Stop CloudPoint
    ```
    # docker run --rm -it
    -v /var/run/docker.sock:/var/run/docker.sock
    -v /cloudpoint:/cloudpoint veritas/flexsnap-cloudpoint:<version>
    stop
    ```
  - Restart Docker
    ```
    # systemctl restart docker
    ```
  - Restart CloudPoint
    ```
    # docker run --rm -it
    -v /var/run/docker.sock:/var/run/docker.sock
    -v /cloudpoint:/cloudpoint veritas/flexsnap-cloudpoint:<version>
    start
    ```

- **Backup from Snapshot job and Indexing job fails with the errors**

  ```
  Jun 10, 2021 2:17:48 PM - Error mqclient (pid=1054) SSL
  Connection failed with string, broker:<hostname>
  Jun 10, 2021 2:17:48 PM - Error mqclient (pid=1054) Failed SSL
  handshake, broker:<hostname>
  Jun 10, 2021 2:19:16 PM - Error nbcs (pid=29079) Invalid
  operation for asset: <asset_id>
  Jun 10, 2021 2:19:16 PM - Error nbcs (pid=29079) Acknowledgement
  not received for datamover <datamover_id>
  ```

  and/or

  ```
  Jun 10, 2021 3:06:13 PM - Critical bpbrm (pid=32373) from client
  <asset_id>: FTL - Cannot retrieve the exported snapshot details
  for the disk with UUID:<disk_asset_id>
  Jun 10, 2021 3:06:13 PM - Info bptm (pid=32582) waited for full
  buffer 1 times, delayed 220 times
  Jun 10, 2021 3:06:13 PM - Critical bpbrm (pid=32373) from client
  <asset_id>: FTL - cleanup() failed, status 6
  ```

This can happen when the inbound access to CloudPoint on port 5671 and 443 port gets blocked at the OS firewall level (firewalld). Hence, from the datamover container (used for the Backup from Snapshot and Indexing jobs), communication to CloudPoint gets blocked. This results in the datamover container not being able to start the backup or indexing.

Workaround

Modify the rules in OS firewall to allow the inbound connection from 5671 and 443 port.

- **Agentless connection fails for a VM with an error message.**

  Agentless connection fails for a VM with the following error message when user changes the authentication type from SSH Key based to password based for a VM through the portal:

```
User does not have the required privileges to establish an agentless connection
```

  This issue occurs when the sudoers file has the order messed up for the user as mentioned in the above error message.

  Workaround:

  Resolve the sudoers file issue for the user by providing the required permissions to perform the passwordless sudo operations.

- **When CloudPoint is deployed in private subnet (without internet) CloudPoint function fails**

  This issue occurs when CloudPoint is deployed in private network where firewall is enabled or public IP which is disabled. The customer's information security team would not allow full internet access to the virtual machine's.

  Workaround

  Enable the ports from the firewall command line using the following commands:

```
firewall-cmd --add-port=22/tcp
firewall-cmd --add-port=5671/tcp
firewall-cmd --add-port=443/tcp
```

- **Restoring asset from backup copy fails**

  In some of the scenarios it is observed that the connection resets intermittently in Docker container. Due to this the server sends more tcp payload than the advertised client window. Sometimes Docker container drops **SYN+ACK** packet from new TCP connection handshake. To allow these packets, use the `nf_conntrack_tcp_be_liberal` option.

  If `nf_conntrack_tcp_be_liberal = 1` then the following packets are allowed:

  - ACK is under the lower bound (possible overly delayed ACK)
  - ACK is over the upper bound (ACKed data not seen yet)
  - SEQ is under the lower bound (already ACKed data retransmitted)

- SEQ is over the upper bound (over the window of the receiver)

If `nf_conntrack_tcp_be_liberal = 0` then those are also rejected as invalid.

Workaround

To resolve the issue of restore from backup copy, use the `nf_conntrack_tcp_be_liberal = 1` option and set this value on node where datamover container is running.

Use the following command for setting the value of `nf_conntrack_tcp_be_liberal`:

```
sysctl -w net.netfilter.nf_conntrack_tcp_be_liberal=1
```

- **Some pods on Kubernetes extension progressed to completed state**

  Workaround

  Disable Kubernetes extension.

  Delete listener pod using the following command:

  ```
  #kubectl delete pod flexnsap-listener-xxxxx -n <namespace>
  ```

  Enable Kubernetes extension.

- **User is not able to customize a cloud protection plan**

  Workaround

  Create a new protection plan with the desired configuration and assign it to the asset.

- **Podman container not starting or containers are not up after reboot**

  On RHEL 8.x platform, restarting container or machine reboot, the container displays the following error message:

  ```
  # podman restart flexsnap-coordinator
  47ca97002e53de808cb8d0526ae033d4b317d5386ce085a8bce4cd434264afdf:
  "2022-02-05T04:53:42.265084989+00:00 Feb 05 04:53:42
  flexsnap-coordinator flexsnap-coordinator[7]
  agent_container_health_check flexsnap.container_manager: INFO -
  Response: b'{""cause"":""that name is already in
  use"",""message"":""error creating container storage: the container
  name \\""flexsnap-agent.15bd0aea11164f7ba29e944115001d69\\"" is
  already in use by
  \\""30f031d586b1ab524511601aad521014380752fb127a9440de86a81b327b6777\\"".
  You have to remove that container to be able to reuse that name.:
  that name is already in use"",""response"":500}\n'"
  ```

  Workaround

  Check if there is a file with IP address entry mapping to the container that could not be started at `/var/lib/cni/networks/flexsnap-network/` file system location.

```
[ec2-user@ip-172-31-44-163 ~]$ ls -latr
/var/lib/cni/networks/flexsnap-network/ total 16 -rwxr-x---. 1
root root 0 Jan 22 12:30 lock drwxr-xr-x. 4 root root 44 Jan 22
12:30 .. -rw-r--r--. 1 root root 70 Feb 4 14:47 10.89.0.150
-rw-r--r--. 1 root root 70 Feb 4 14:47 10.89.0.151 -rw-r--r--. 1
root root 70 Feb 4 14:47 10.89.0.152 -rw-r--r--. 1 root root 11
Feb 7 11:09 last_reserved_ip.0 drwxr-xr-x. 2 root root 101 Feb 7
11:13 . [ec2-user@ip-172-31-44-163 ~]$
```

From the above directory , delete the duplicate IP address file and perform the stop and start operation as follows:

Stop the container: `#podman stop <container_name>`

Start the container:`#podman start <container_name>`

- **After starting the start/stop services, CloudPoint, RabbitMQ and MongoDB containers are still in the starting state**

  It was observed that flexsnap-mongodb and flexsnap-rabbitmq containers did not go into healthy state. Following is the Below is the state of flexsnap-mongodb container:

```
[ec2-user@ip-172-31-23-60 log]$ sudo podman container inspect
--format='{{json .Config.Healthcheck}}' flexsnap-mongodb
{"Test":["CMD-SHELL","echo 'db.runCommand({ping: 1}).ok' | mongo
--ssl --sslCAFile /cloudpoint/keys/cacert.pem --sslPEMKeyFile
/cloudpoint/keys/mongodb.pem flexsnap-mongodb:27017/zenbrain
--quiet"],"Interval":60,"Timeout":30000000000,"Retries":3}
[ec2-user@ip-172-31-23-60 log]$ sudo podman container inspect
--format='{{json .State.Healthcheck}}' flexsnap-mongodb
{"Status":"starting","FailingStreak":0,"Log":null}
[ec2-user@ip-172-31-23-60 log]$
```

  Workaround

  Run the following #podman CLI(s) command:

```
[ec2-user@ip-172-31-23-60 log]$ sudo podman healthcheck run flexsnap-mongo

[ec2-user@ip-172-31-23-60 log]$ sudo podman ps -a

CONTAINER ID  IMAGE                            COMMAND  CREATED      STATUS

fe8cf001032b  localhost/veritas/

  flexsnap-fluentd:10.0.0.0.9817         2 days ago    Up 45 hours ago

2c00500c1ac6  localhost/veritas/
```

```
  flexsnap-mongodb:10.0.0.0.9817         2 days ago    Up 45 hours ago (he

7ab3e248024a  localhost/veritas/

  flexsnap-rabbitmq:10.0.0.0.9817        2 days ago    Up 45 hours ago (st

[ec2-user@ip-172-31-23-60 log]$ sudo podman healthcheck run flexsnap-rabbi

[ec2-user@ip-172-31-23-60 log]$ sudo podman ps -a

CONTAINER ID  IMAGE                              COMMAND  CREATED       STATU

fe8cf001032b  localhost/veritas/

  flexsnap-fluentd:10.0.0.0.9817        2 days ago    Up 45 hours ago

2c00500c1ac6  localhost/veritas/

  flexsnap-mongodb:10.0.0.0.9817        2 days ago    Up 45 hours ago (hea

7ab3e248024a  localhost/veritas/

  flexsnap-rabbitmq:10.0.0.0.9817       2 days ago    Up 45 hours ago (hea


[ec2-user@ip-172-31-23-60 log]$ sudo podman container inspect --format='{{

{"Status":"healthy","FailingStreak":0,"Log":

[{"Start":"2022-02-14T07:32:13.051150432Z","End":"2022-02-14T07:32:13.4446

[ec2-user@ip-172-31-23-60 log]$ sudo podman container inspect --format='{{

{"Status":"healthy","FailingStreak":0,"Log":

[{"Start":"2022-02-14T07:32:46.537804403Z","End":"2022-02-14T07:32:47.2936

[ec2-user@ip-172-31-23-60 log]$
```

- **Certificate generation would fail while registering CloudPoint with NetBackup**

  Starting CloudPoint release 9.1.2, NetBackup certificate generation will happen synchronously with registration in register API of CloudPoint. Hence, any failure in certificate generation will cause failure while registering CloudPoint with

NetBackup, that is ading or editing the CloudPoint server entry from Web UI. These certificates are used for datamover which is launched for operations like backup from snapshot, restore from backup, indexing (VxMS based), and so on. Hence, if certificate generation fails, these jobs cannot be performed. Hence CloudPoint on cloud VMs cannot connect to NetBackup on lab VMs, hence the registration will fail, and hence CloudPoint cannot be added to NetBackup.

Workaround

To add CloudPoint in such scenario requires to skip certificate generation on CloudPoint by adding the following entry in `/cloudpoint/flexsnap.conf` file:

```
[client_registration] skip_certificate_generation = yes
```

- **Default timeout of 6 hours is not allowing restore of larger database (size more than 300 GB)**

  Workaround

  Configurable timeout parameter value can be set to restore larger database. The timeout value can be specified in `/etc/flexsnap.conf` file of `flexsnap-coordinator` container. It does not require restart of the coordinator container. Timeout value would be picked up in next database restore job. User must specify the timeout value in seconds as follows:

  ```
  docker exec -it flexsnap-coordinator bash
  root@flexsnap-coordinator:/# cat /etc/flexsnap.conf [global] target
  = flexsnap-rabbitmq grt_timeout = 39600
  ```

- **Plugin information is duplicated, if CloudPoint registration has failed in previous attempts**

  This occurs only when CloudPoint has been deployed using the MarketPlace Deployment Mechanism. This issue is observed when the plugin information is added before the registration. This issue creates duplicate plugin information in the **CloudPoint_plugin.conf** file.

  Workaround

  Manually delete the duplicated plugin information from the **CloudPoint_plugin.conf** file.

  For example, consider the following example where the duplicate entry for GCP plugin config is visible (in bold) in **CloudPoint_plugin.conf** file:

```
[
  {
    "CPServer1": [
      {
        "Plugin_ID": "test",
        "Plugin_Type": "aws",
        "Config_ID": "aws.8dda1bf5-5ead-4d05-912a-71bdc13f55c4",
        "Plugin_Category": "Cloud",
```

```
      "Disabled": false
    }
  ]
},
{
  "CPServer2": [
    {
      "Plugin_ID": "gcp.2080179d-c149-498a-bf1f-4c9d9a76d4dd",
      "Plugin_Type": "gcp",
      "Config_ID": "gcp.2080179d-c149-498a-bf1f-4c9d9a76d4dd",
      "Plugin_Category": "Cloud",
      "Disabled": false
    },
    {
      "Plugin_ID": "gcp.2080179d-c149-498a-bf1f-4c9d9a76d4dd",
      "Plugin_Type": "gcp",
      "Config_ID": "gcp.2080179d-c149-498a-bf1f-4c9d9a76d4dd",
      "Plugin_Category": "Cloud",
      "Disabled": false
    }
  ]
}
```