# Veritas CloudPoint Release Notes

Ubuntu, RHEL

2.2.1

**VERITAS**™

# Veritas CloudPoint Release Notes

Last updated: 2019-11-04

Document version: 2.2.1 Rev 0

## Legal Notice

## Technical Support

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

https://www.veritas.com/support

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

https://my.veritas.com

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

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

## Documentation

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

https://sort.veritas.com/documents

## Documentation feedback

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

cloudpointdocs@veritas.com

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

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

## Veritas Services and Operations Readiness Tools (SORT)

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

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

# Contents

# Getting help

This chapter includes the following topics:

- About this document
- CloudPoint resources

## About this document

This document provides information specific to the Veritas CloudPoint 2.2.1 release. Review this document before using the product.

The information in this document supersedes all the information provided in other product-specific documents.

For information about the operating system, hardware, and other general requirements, refer to the *Veritas CloudPoint Administrator's Guide*.

You can download the latest version of this document from the Veritas Service and Operations Readiness Tools (SORT) web site at:

https://sort.veritas.com/documents

## CloudPoint resources

For information about CloudPoint features, use cases, data sheets, white papers, and videos, refer to the following product page:

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

### Software download

You can obtain the latest release of CloudPoint from the following web site:

https://www.veritas.com/form/trialware/cloudpoint-download

The software provided is a free edition that offers limited functionality and supports up to 10 Front End Terabyte (FETB) of data usage. The free version is subject to a license agreement. Please contact your Veritas sales representative for additional details.

## User documentation

For information on supported platforms, software and hardware requirements, and installation and administration instructions, refer to the CloudPoint documentation here:

- Veritas Support

  https://www.veritas.com/support/en_US.html

  Click the documentation link and then choose CloudPoint from under the Product filter to display the latest documentation.

- Veritas Services and Operations Readiness Tools (SORT)

  https://sort.veritas.com/documents

  Select the product and the platform and apply other filters to display the desired documentation.

- Late Breaking News (LBN)

  https://www.veritas.com/support/en_US.html

  View the latest information about updates, patches, and software issues for this release.

## VOX community forum

You can use the Veritas Open eXchange (VOX) community forum to connect directly with the CloudPoint product development team:

https://vox.veritas.com/t5/CloudPoint/bd-p/CloudPoint

# New features, enhancements, and changes

This chapter includes the following topics:

- Tag-based asset protection for production environments

- Support for AWS nitro-based instances that use NVMe devices

- Single File Restore (SFR) supported on Linux root file system

- Support for Hitachi storage arrays

- Support for InfiniBox storage arrays

- CloudPoint REST API version updated to /v3

- Updates to the CloudPoint plug-in for Microsoft Azure

- Unconfigure plug-in and disconnect host option now available from the CloudPoint user interface (UI)

- Support for AWS GovCloud (US)

## Tag-based asset protection for production environments

CloudPoint introduces an enhanced version of the Tag-based asset protection feature that allows you to automate the entire process of data protection using CloudPoint protection policies. You can assign descriptive text labels, known as

tags, to the assets and configure this feature to have CloudPoint automatically protect those assets using matching protection policies.

This feature was introduced in CloudPoint 2.2 release earlier. The following changes have been made to the feature in this release:

■ CloudPoint no longer statically associates policies with assets, post discovery. Instead, CloudPoint dynamically determines whether or not to protect an asset based on tags, during every policy run.

■ In 2.2 release, this feature was offered as a technical preview. With the enhancements in this release, this feature can now be used in production environments.

Refer to the chapter on Tag-based asset protection in the *CloudPoint Administrator's Guide* for more information.

# Support for AWS nitro-based instances that use NVMe devices

CloudPoint now supports AWS nitro-based instances that use EBS volumes that are exposed as non-volatile memory express (NVMe) devices. You can now use the CloudPoint AWS plug-in to discover and protect such instances.

Refer to the AWS plug-in configuration notes in the *CloudPoint Administrator's Guide* for more information.

# Single File Restore (SFR) supported on Linux root file system

CloudPoint allows you to restore individual files from within a snapshot and this process is known as granular restore or single file restore (SFR). In this release, SFR support has been added for the Linux root file system (/).

Refer to the *CloudPoint Administrator's Guide* for more information.

# Support for Hitachi storage arrays

CloudPoint now adds support for Hitachi storage arrays. You can configure the new CloudPoint plug-in for Hitachi to discover and protect disks on Hitachi storage arrays that are managed using Hitachi Configuration Manager (HCM).

For more details on the Hitachi plug-in, refer to the *CloudPoint Administrator's Guide*.

# Support for InfiniBox storage arrays

CloudPoint now adds support for INFINIDAT InifiniBox enterprise storage arrays. You can configure the new CloudPoint plug-in for InfiniBox to discover and protect disk pools created on InfiniBox enterprise storage arrays.

For more details on the InfiniBox plug-in, refer to the *CloudPoint Administrator's Guide*.

# CloudPoint REST API version updated to `/v3`

The CloudPoint REST APIs have been updated and the version has been modified from `/v2` to `/v3`. All the CloudPoint APIs are now accessible at `/v3`.

For example, to get the CloudPoint version details from a CloudPoint host, the path of the API endpoint earlier was as follows:

```
/cloudpoint/api/v2/version
```

The same API is now available at the following path:

```
/cloudpoint/api/v3/version
```

The updated APIs will work with CloudPoint 2.2 and later releases. For working with older releases of CloudPoint, for example CloudPoint 2.1, you can use the APIs available at `/v2`.

You can access all the APIs using the following URL:

```
https://<cloudpoint_hostFQDN>/cloudpoint/docs
```

# Updates to the CloudPoint plug-in for Microsoft Azure

The CloudPoint Azure plug-in has been updated in this release. As a result, the following changes and enhancements have been introduced:

- CloudPoint now supports restore for Azure-managed Standard Solid State Drive (SSD) disks that are configured with the Locally redundant storage (LRS) redundancy option.

- CloudPoint now supports the **Overwrite existing** option for restoring asset snapshots in an Azure cloud environment. When you use this restore option, the existing asset is replaced with the selected snapshot.
  This restore option is enabled only for snapshots that are created using CloudPoint 2.2.1 or later.

- The Azure plug-in uses lock objects on asset snapshots to protect the snapshots from unintentional deletion. The lock object also includes a "`notes`" field that contains the ID of the asset that the snapshot belongs to. The plug-in uses this ID to associate the snapshots with the instances.

# Unconfigure plug-in and disconnect host option now available from the CloudPoint user interface (UI)

The CloudPoint UI has been enhanced in this release. You can now use the UI to unconfigure the CloudPoint plug-ins from the hosts and also disconnect the hosts to clean up their configuration entries from the CloudPoint database. These operations were earlier available only via CloudPoint REST APIs.

Refer to the *CloudPoint Administrator's Guide* for more information.

# Support for AWS GovCloud (US)

CloudPoint provides support for AWS GovCloud (US) in this release. You can now configure the CloudPoint AWS plug-in to discover and protect assets in the following AWS GovCloud (US) regions:

- us-gov-east-1
- us-gov-west-1

For more details, refer to the AWS plug-in configuration notes in the *Veritas CloudPoint Administrator's Guide*.

# Limitations

This chapter includes the following topics:

- CloudPoint support for AWS regions and GCP zones (14446)
- Limitations on replicating and restoring assets
- If two snapshot operations are performed on an instance at the same time, the second one fails
- You cannot delete snapshots created by an Amazon snapshot policy
- CloudPoint cannot snapshot LUNs which are under a consistency group for Dell EMC Unity arrays (3977)

## CloudPoint support for AWS regions and GCP zones (14446)

CloudPoint supports AWS regions and GCP zones that exist at the time CloudPoint software is released as a generally available (GA) release. You can use CloudPoint to protect cloud assets that reside in the regions and zones that already exist at GA.

CloudPoint currently does not support newer regions or zones that are created after a CloudPoint version is released.

## Limitations on replicating and restoring assets

When you work with CloudPoint's replication feature, keep in mind the following;

- You can restore an encrypted snapshot. To enable the restoring of encrypted snapshots, add a Key Management Service (KMS) policy, and grant the

CloudPoint user access to KMS keys so that they can restore encrypted snapshots.

# If two snapshot operations are performed on an instance at the same time, the second one fails

If a snapshot operation is in progress and a second operation is triggered on the same application or cluster, then the second snapshot operation fails with a `operation already in progress` error.

The failure occurs because the instance or cluster must be in the available state for the operation to be performed on it. If the asset is not available, the snapshot operation fails.

---

**Note:** The CloudPoint user interface does not display whether an instance or application is available.

---

# You cannot delete snapshots created by an Amazon snapshot policy

CloudPoint not only displays the snapshots you create, but also the snapshots created outside CloudPoint by an Amazon snapshot policy. You cannot delete Amazon-created snapshots using the CloudPoint **Delete Snapshot** operation. You can only delete a snapshot created from within CloudPoint.

# CloudPoint cannot snapshot LUNs which are under a consistency group for Dell EMC Unity arrays (3977)

If you use CloudPoint with Dell EMC Unity arrays, be aware that you cannot snapshot LUNs which are under a consistency group. The reason for this limitation is that to restore a single LUN snapshot restores the entire consistency group.

# Known issues

This chapter includes the following topics:

- Azure instance image cleanup failed (7253)

- Agent services restarting continuously (8030)

- CloudPoint does not support 'Overwrite existing' restore option for file system and application's host-level snapshot (8924)

- If a snapshot creation policy is run in parallel with in-place restore, the policy may fail. When the policy is run again, it succeeds (8142)

- When AWS is configured with different regions in parallel, stacks are added in the logs (7481)

- Some errors are logged in the nginx_error.log file (6593)

- CloudPoint is not able to find Google cloud network configuration (3254)

- Signing out from a non-admin account and signing in as an admin gives limited access (2862)

- Configuring multiple plug-ins may cause an error (6562)

- Assets (instances) are not getting discovered for Azure (6953)

- CloudPoint incorrectly allows snapshot operations on operating system native file systems (12285)

- Application-consistent snapshot operations for ext2 file system assets fail (12948)

- Snapshot operations might hang if CloudPoint host restarts (14757, 9039)

- Indexing may fail if snapshot is not local to the CloudPoint host (14127)

- Disk-level snapshot restore to the same location fails if an application was previously added and removed on the same disk (13196)

- Disk-level snapshot restore fails if MongoDB configuration file is at a non-default location (13130)

- Multiple simultaneous restore operations on the same AWS instance may fail (14286)

- During a policy run, a snapshot failure of one of the assets causes the other successful snapshots to be orphaned (14760)

- Classification fails if the snapshot contains encrypted files (15858)

- CloudPoint upgrade may stall if the Docker version is not 18.03 (15846)

- Classification job may fail with a partition detection error (15554)

- CloudPoint UI incorrectly displays internal storage objects (15611)

# Azure instance image cleanup failed (7253)

Azure operations (snapshot creation, snapshot deletion ) fail intermittently with http status codes 429, 502, or 503.

The http status codes are returned from the Azure portal and then you need to retry the operation again.

# Agent services restarting continuously (8030)

The agent services keep restarting continuously due to insufficient memory for CloudPoint processes.

Sometimes agent services restart continuously due to high memory pressure generated by large numbers of workload. In this situation some of services fail to work due to memory crunch and are not able to recover. This leads to agent continuously restarting for a healthy communication. As some of services do not respond, agent continuously retries to establish a connection. CloudPoint services need to be restarted to recover to normal state.

# CloudPoint does not support 'Overwrite existing' restore option for file system and application's host-level snapshot (8924)

CloudPoint does not support the **Overwrite existing** restore option for host-level snapshot assets of file systems and applications. If you trigger such restore operations using the CloudPoint REST API, the operation fails with an error.

# If a snapshot creation policy is run in parallel with in-place restore, the policy may fail. When the policy is run again, it succeeds (8142)

This occurs when in-place restore and the policy were ran at the same time. In-place restore deletes currently attached volume to the instance at the same time when the policy to take snapshot is triggered.

The attached volumes of the previously started policy are not detached from the instance. Hence, those volumes are counted to take snapshot are deleted. By the time create_snapshot called, source volumes got deleted. Hence snapshot creation fails.

# When AWS is configured with different regions in parallel, stacks are added in the logs (7481)

This is a known issue. When an agent is created dynamically, stack are added in the logs but it does not impact any functionality.

# Some errors are logged in the nginx_error.log file (6593)

Some error related to variables are logged in the `nginx_error.log` file. The following errors are logged in the `nginx_error.log` file.

```
using uninitialized "csrfcookie" variable

using uninitialized "authcookieexists" variable

using uninitialized "csrfmismatch" variable
```

# CloudPoint is not able to find Google cloud network configuration (3254)

If CloudPoint host networking is configured using network manager before installing the docker then CloudPoint is unable to find Google cloud network configuration.

This is a known issue and the workaround can be found at
https://gist.github.com/JPvRiel/dcb9e2866a9d0aa19042028cca3306c7

# Signing out from a non-admin account and signing in as an admin gives limited access (2862)

If you sign out of a non-admin CloudPoint account and then sign in as an admin, the user interface does not give you access to admin functions. For example, on the Administration tab, the links for Clouds/Arrays, Policies, and Users say Need access?. They should say Manage.

This occurs sometimes. As a workaround, log out and refresh the page and then log on again as admin user.

# Configuring multiple plug-ins may cause an error (6562)

This issue may occur intermittently. If multiple CloudPoint plug-ins are configured simultaneously, then the operations might fail with a `dictionary changed size during iteration` error.

**Workaround:**

Try re-configuring the plug-in again.

# Assets (instances) are not getting discovered for Azure (6953)

CloudPoint currently does not support the following types of Azure instances:

- Virtual machine Classic
- Virtual machine Scale Sets

These Azure instances are not discovered.

# CloudPoint incorrectly allows snapshot operations on operating system native file systems (12285)

CloudPoint discovers operating system native file systems as assets and the user interface displays the `Snapshotable` parameter as "`Yes`" for those assets, which means that you can perform snapshot operations on those native file system assets.

For example,

- XFS file system `/` (root)
- XFS file system configuration file `/etc/hosts`

Do not perform any snapshot or other operations on the native file system objects that appear as assets in CloudPoint.

# Application-consistent snapshot operations for ext2 file system assets fail (12948)

If you try to create application-consistent snapshots of `ext2` file system assets, either from the CloudPoint user interface (UI) or using the CloudPoint API, the snapshot operation fails.

The following errors appear in the CloudPoint log file:

```
ERROR - Failed to freeze filesystem: ext2 on device <device>.
ERROR: fsfreeze: /ext2mnt: freeze failed: Operation not supported
```

While creating a snapshot, CloudPoint attempts to freeze the file system. But the `fsfreeze` command itself is not supported on `ext2` file systems and therefore the operation fails.

**Workaround:**

There is no known workaround at the moment. Veritas recommends that if you wish to create application-consistent snapshots for file system assets, you use only the supported file systems such as ext3, ext4, or XFS.

# Snapshot operations might hang if CloudPoint host restarts (14757, 9039)

If the CloudPoint host is restarted while the snapshot operations are in progress, then those operations may hang. Even after the host restarts, these operations continue to remain in a hung state.

**Workaround:**

After the CloudPoint host restarts successfully and all the containers are up and running, restart the coordinator service on the host:

```
# sudo docker restart flexsnap-coordinator
```

Snapshot operations that were stuck are automatically rerun after the service is restarted.

# Indexing may fail if snapshot is not local to the CloudPoint host (14127)

Indexing operations triggered manually from the CloudPoint UI or using CloudPoint REST APIs may fail if the snapshot being indexed is not local to the CloudPoint host. The CloudPoint UI does not display any errors indicating the failure.

The `flexsnap-coordinator.log` file may contain the following errors:

```
flexsnap-indexingsupervisor: flexsnap-indexingsupervisor[1]
Thread-2 flexsnap.indexingsupervisor:
ERROR - Coordinator indexSnapshot failed: Method 'restore to <ID>'
not supported on asset <ID>

flexsnap-indexingsupervisor: raise flexsnap.GenericError.generate(msg)
flexsnap-indexingsupervisor: MethodNotSupported: Method 'restore to <ID>'
not supported on asset <ID>
```

**Workaround:**

There is no workaround to resolve this issue.

CloudPoint does not support indexing for snapshots that are not local to the CloudPoint host. For indexing to be successful, ensure that the snapshots being indexed belong to the same region, cloud account, availability zone, or project as that of the CloudPoint host.

# Disk-level snapshot restore to the same location fails if an application was previously added and removed on the same disk (13196)

If you create and remove an application from a disk and then use the same disk to create a new application, then when you try to restore disk-level snapshots of the new application to the same location, the restore operation fails.

For example, consider a disk `Disk1` on which you created an application `MyApp1`. Now remove `MyApp1` and create another application `MyApp2` on the same disk (`Disk1`) and then take disk-level snapshots of that disk. When you try to restore the new application (`MyApp2`) disk-level snapshot to the same location, the restore operation fails.

The `flexsnap-coordinator` logs contain the following:

```
flexsnap-coordinator: <ID> flexsnap-coordinator[1] Thread-26113
flexsnap.connectors.base: ERROR - Request failed unexpectedly
flexsnap-coordinator: File "/opt/VRTScloudpoint/lib/flexsnap/coordinator.py"
, line 5542, in verify_single_leaf_child
flexsnap-coordinator: therefore, cannot do original location restore"
% asset_type)
flexsnap-coordinator: GenericError: Another application would be affected
therefore, cannot do original location restore
```

This issue occurs because references of the application (`MyApp1`) that was deleted from the disk (`Disk1`) are retained in the CloudPoint assets database. A restore of the new application (`MyApp2`) disk-level snapshot fails to overwrite the old references on the same disk (`Disk1`).

**Workaround:**

There is no known workaround for this issue. As an alternative, you can try performing the disk-level restore to a new location.

# Disk-level snapshot restore fails if MongoDB configuration file is at a non-default location (13130)

If you configure a MongoDB instance from a configuration file that is located at a non-default location, then a restore operation for the MongoDB disk-level snapshot fails with the following error:

```
Failed to unmount /mongodata (umount: /mongodata: target is busy.
(In some cases useful info about processes that use the device is
found by lsof(8) or fuser(1)) )
```

**Workaround:**

There is no known workaround for this issue. As an alternative, you can launch MongoDB instances using a configuration file from the default location; snapshot restore does not fail for such instances.

# Multiple simultaneous restore operations on the same AWS instance may fail (14286)

If you trigger multiple disk-level snapshot restore operations on the same AWS instance in parallel, CloudPoint may choose conflicting device names for attaching the restored disks to the instance. This can lead to an inconsistency at the instance level and one or more restore operations might eventually fail.

**Workaround:**

There is no known workaround for this issue. CloudPoint does not support running multiple jobs on the same asset simultaneously. If you intend to perform multiple restore operations on the same instance, then Veritas recommends that you perform the operations one after the other, in a sequential manner.

# During a policy run, a snapshot failure of one of the assets causes the other successful snapshots to be orphaned (14760)

This issue may occur when a CloudPoint protection policy is assigned to two or more assets. If the CloudPoint on-host plug-in that manages the asset is not online on the host when the policy is run, then the application-consistent snapshot creation for that asset might fail. If the snapshot operation for at least one of the assets in the policy fails, the policy run is marked as failed even if the snapshots for all the other remaining assets get created successfully.

As a result, the snapshots that get created successfully are no longer associated with the policy that triggered those snapshots. These snapshots are not accounted for in the snapshot retention count defined in the policy. These snapshots do not get deleted automatically during subsequent policy runs and remain in an orphaned state.

This behavior is applicable only in case of application-consistent snapshots. For host and disk-level snapshots, even if snapshot creation fails for one of the assets, the policy run is not marked as failed and the policy retention count works as intended.

**Workaround:**

There is no known workaround for this issue. You may have to troubleshoot the snapshot failure and then manually remove such orphaned policy-based snapshots from the CloudPoint configuration.

# Classification fails if the snapshot contains encrypted files (15858)

If you trigger a classification job on a snapshot that includes one or more encrypted files, then the classification operation fails. CloudPoint is unable to discover the tags that are assigned to any of the files, encrypted as well as unencrypted, that are included in the snapshot.

This issue occurs on Microsoft Windows only.

**Workaround:**

There is no known workaround for this issue. You may have to remove the encrypted files from the asset, take a fresh snapshot, and then run the classification operation on that snapshot again.

# CloudPoint upgrade may stall if the Docker version is not 18.03 (15846)

CloudPoint upgrade process may hang and eventually fail if the Docker version is greater than 18.03.

You may see errors similar to the following:

```
dockerd[23589]: level=info msg="Container 2bbbc67f579f6877c4c1 failed to
exit within 10 seconds of signal 15 - using the fo
level=info msg="shim reaped" id=2bbbc67f579f6877c4c1
level=info msg="ignoring event" module=libcontainerd namespace=moby
topic=/tasks/delete type="*events.TaskDelete"
```

**Workaround:**

First downgrade the Docker version to 18.03 and then proceed with the CloudPoint upgrade.

# Classification job may fail with a partition detection error (15554)

While running an indexing or classification operation, CloudPoint temporarily mounts the snapshot as a mount point on the CloudPoint host. If the mount point where the snapshot is mounted is accessed by any other application or service while the operation is in progress, then the classification or indexing operation may or may not succeed.

If you trigger another indexing or classification operation on the same asset, the job fails with the following error:

```
Unable to detect partition for device /dev/xvdg
```

This issue occurs because CloudPoint fails to clear the mount point entries that are created for the indexing and classification operations. The stale mount point entries persist even after the operations are complete. As a result, subsequent operations fail with a partition detection error.

This issue occurs on Linux only.

**Workaround:**

Perform the following steps on the CloudPoint host:

1.  Identify the processes that are consuming the mount point and then close or terminate all those processes.

2.  Unmount the disk or partition.

3.  From the CloudPoint UI, detach the respective volume.

4.  If the detach operation from the UI fails, then you may have to reboot the CloudPoint server and then try the detach operation again.

# CloudPoint UI incorrectly displays internal storage objects (15611)

While performing certain operations such as agent installation, plug-in registration, or granular restore, CloudPoint creates internal storage objects and mounts them on the protected host where the operation is being performed. These storage objects remain on the target host while the operation is in progress and are automatically dismounted and removed once the operation is complete.

If a CloudPoint discovery cycle is triggered while such an operation is in progress, CloudPoint discovers these internal-only temporary storage objects as assets and the CloudPoint UI incorrectly displays them and allows you to select them for performing CloudPoint operations.

These objects appear as file system or disk assets and are typically named as follows:

```
Disk /dev/<diskmount> on <hostname>.internal
```

**Workaround:**

The appearance of these objects in the UI does not cause an issue with the actual operation itself. However, you must ensure that you do not select these temporary

storage objects for snapshot or restore operations or assign a CloudPoint protection policy to such objects.

# Fixed issues

This chapter includes the following topics:

■ Fixed issues

## Fixed issues

The following issues are fixed in this release. If you contact Veritas about any of these issues, use the incident number as a reference.

**Table 5-1**    CloudPoint fixed issues

| Incident # | Description |
|---|---|
| 3928828 | Delete snapshot operation not visible on the Recent Activity tab in the UI. |
| 3931139 | Failed to delete snapshot errors are logged incorrectly until the policy snapshot retention count is reached. |
| 12286 | MongoDB database application snapshot creation might fail. |
| 13940 | Indexing or classification jobs may fail in an Azure environment. |
| 13502 | Indexing or classification appears as completed but the actual operation may have failed. |
| 15063 | Removing a plug-in fails to delete assets if an asset discovery is already in progress. |
| 3941239 | If a CloudPoint policy protects a large number of assets, for example 100 devices or more, it takes longer to delete snapshots. If a delete operation occurs at a scheduled snapshot time, the delete may fail. This issue is now resolved. |

**Table 5-1**     CloudPoint fixed issues *(continued)*

| Incident # | Description |
| --- | --- |
| 8250 | After restoring an instance on AWS, tags are not restored. This issue is now resolved. |
| 6274 | An issue related to role API returning 500 Internal Server Error when an incorrect role ID is provided is now resolved. |
| 6098 | The support for replication and restore of encrypted snapshots when using Customer Managed Keys is now added. |
| 5286, 5611 | An issue related to snapshot replication through policy not getting triggered is now resolved. |
| 4604 | An issue related to anonymous relay for SMTP is now resolved. |
| 4427 | An error is generated when GCP host is restored from one zone to another. This issue is now resolved. |
| 5623 | An issue related to replication of replicated snapshot caused an exception "`TypeError: cannot concatenate 'str' and 'list' objects`". This issue is now resolved. |
| 4604 | An issue related to Anonymous relay for SMTP is now resolved. |
| 3210 | An issue about not being able to find the server at www.googleapis.com was resolved. |
| 9519, 5495 | An LDAP integration error in the CloudPoint UI "`Error updating LDAP`". The issue is resolved. CloudPoint now supports LDAP over Secure Sockets Layer (SSL). |
| 8291 | After restoring a host-level snapshot, the instance name appeared blank in the AWS console. This issue is now resolved. |
| 6128, 10335 | CloudPoint did not provide an option to suspend a policy. This issue is now resolved. An option to disable a policy is now available. |