IBM® Storage

Offloading Storage Volumes from Safeguarded Copy to AWS S3 Object Storage with IBM FlashSystem Transparent Cloud Tiering

Contents

About this document

The focus of this IBM® Blueprint is to showcase a method to store volumes that are created by using Safeguarded Copy off-premise to Amazon S3 object storage that uses the IBM FlashSystem Transparent cloud tiering (TCT) feature.

TCT enables volume data to be copied and transferred to object storage. The TCT feature supports creating connections to cloud service providers to store copies of volume data in private or public clouds.

This feature is useful for organizations of all sizes when planning for disaster recovery operations or storing a copy of data as extra backup. TCT provides seamless integration between the storage system and public or private clouds for Safeguarded Copy volumes and non-Safeguarded Copy volumes.

Executive summary

Increasingly, a key component to the Continuity of Operations (COOP) of any organization is cyber resiliency. Different from cybersecurity, *cyber resiliency* is about an organization's ability to continue operations despite a cyberincident. It can be combined with an organization's Disaster Recovery capabilities and their business continuity processes, but is a separate set of capabilities.

With the increased risk to data from various cyber threats (for example, ransomware), organizations are looking for new, innovated, and expanded ways to protect their business's data. Key to this protection is the ability to make immutable copies of the production data that can be quickly recovered if an attack occurs. By having these copies available, data can be quickly validated and recovered, which enables the business to return to service

IBM Spectrum Virtualize, which powers the IBM FlashSystem family and the IBM SAN Volume Controller, includes the Safeguarded Copy function that supports the ability to create cyber-resilient, point-in-time copies of volumes that cannot be changed or deleted by user errors, malicious actions, or ransomware attacks. These copies can be integrated with IBM Copy Services Manager (CSM), IBM Spectrum Copy Data Management (CDM), or the internal scheduling capability to provide automated backup copies and data recovery.

With the Safeguarded Copy providing rapid access to immutable copies for recovery on the local IBM FlashSystem, a requirement can exist to store a copy of the point-in-time copies of those volumes in another recovery zone, potentially one that is outside of the data center or to another system.

The use of public cloud by organizations is increasing every day. Organizations can exist on public cloud entirely while other enterprises use a hybrid cloud approach. The cloud is no more seen as compute entity, but it also is an excellent option from a storage perspective.

Cloud storage works by storing data on remote servers, where it can be maintained, managed, backed up, and accessed remotely. Data that is stored in the cloud is accessible by any device at any time, if permissions are in place. Despite its accessibility, data that is stored by way of the cloud is considered safe and secure.

Cloud storage generally provides high scalability and a pay-as-you-go type consumption model. The scalability aspect is further extended by long-term data retention and automatic deletion. Many cloud providers offer long-term storage in the form of object storage. Object stores are primarily used for storing unstructured data and are flexible about the size of the data.

The IBM FlashSystem family incorporates public cloud integration. With the introduction of the IBM Spectrum Virtualize 8.5.2 software release, the FlashSystem TCT function now supports normal volumes and volumes that are created by using the Safeguarded Copy function.

The cloud volumes that are moved by TCT can be restored to the original IBM FlashSystem that created the copy, or imported to an alternative IBM FlashSystem, IBM SAN Volume Controller cluster, or an instance of IBM Spectrum Virtualize for Public Cloud that is running in an available cloud provider. Therefore, TCT cloud volumes can be used for various testing, migration, and Disaster Recovery scenarios.

This combination of TCT with Safeguarded Copy provides the following benefits:

- The ability to maintain a point-in-time copy of the Safeguarded Copy volumes in an alternative failure zone
- The option to recover to a separate environment for testing and validation of the point-in-time Safeguarded Copy
- The capability to retain copies of the Safeguarded Copy volumes beyond their various policy retention times

Scope

The focus of this document is to showcase the use of the TCT feature to transfer copies of the backup volumes that are generated by using IBM FlashSystem® Safeguarded Copy capability to an Amazon S3 object storage bucket.

Although other public cloud providers, such as IBM Cloud®, OpenStack Swift, and Amazon Cloud, are supported by TCT, this document covers steps for Amazon cloud S3 object store only.

TCT supports use of encryption when transferring the data to the S3 object store when the IBM FlashSystem® has encryption that is enabled. The configuration of encryption on the storage system is *not* covered in this document.

The creation and configuration of an Amazon S3 object storage service account also is not covered in this document.

Users must have a working knowledge of the following concepts:

- Storage classes
- Account types
- Signature versions
- Access points
- Access key IDs
- Secret access keys
- Working with buckets and objects

Introduction

Transparent Cloud Tiering enables volume data to be copied and transferred to object storage in the cloud, such as Amazon S3. IBM Spectrum® Virtualize supports creating connections to cloud service providers to store copies of volume data in private or public cloud storage. These volumes now seen as cloud snapshots can be restored to the original IBM FlashSystem or to another instance of IBM Spectrum Virtualize on a different IBM FlashSystem, IBM SAN Volume Controller, or instance of IBM Spectrum Virtualize for Public Cloud.

With TCT, administrators can move older data to cloud storage to free up capacity on the system. Point-in-time snapshots of data can be created on the system and then, copied and stored on the cloud storage.

An external cloud service provider manages the cloud storage, which reduces storage costs for the system. Before data can be copied to cloud storage, a connection to the cloud service provider must be created from the system.

After the storage system is authenticated, it can access cloud storage to copy data to the cloud storage or restore data that is copied to cloud storage back to the system. The system supports one cloud account to a single cloud service provider. Migration between providers is *not* supported.

IBM Safeguarded Copy creates isolated immutable snapshots of data to help protect against cyberattacks, malware, acts of disgruntled employees, and other data corruption. Because Safeguarded Copy snapshots are on the same FlashSystem storage as operational data, recovery is faster than restoring from copies stored separately.

Safeguarded Copy provides for the capability to make a new volume in the original FlashSystem pool while not changing the original immutable source snapshot on the IBM FlashSystem. It also allows for mapping and testing through other hosts that are defined on that FlashSystem array. This capability is the safest way to preserve the volume for analysis, backup, or other uses, such as movement.

Prerequisites

This section describes the following prerequisites for the use of the TCT feature:

- Verify that your hardware model supports this function. For more information, see "Resources" on page 20.
- Ensure that a DNS server is configured on the system.

  At least one DNS server is required if you connect to cloud service providers as part of TCT support, which included establishing a cloud account and connecting to cloud-based storage.
- Determine whether encryption is required for your connection to the cloud account. Some models might require more encryption licenses. Verify these requirements before this function is used.

Note: When a connection to a cloud service provider is configured, you must decide whether to encrypt data that is at-rest in the cloud for this account. After you decide, the encryption setting for the account cannot be changed without restoring all data from the cloud, reconfiguring the account, and re-creating cloud snapshots for the data.

Configuring Transparent Cloud Tiering

The process to configure TCT includes configuring the following components:

- Domain name server (DNS)
- Cloud connection

Configuring a domain name server

Complete the following steps to configure a DNS:

1. Log in to the FlashSystem management GUI and select the **DNS configuration** option.
2. Click **Settings** → **DNS**.
3. Click **Add DNS server+**.
4. Enter a name, IP address type (IPV4 or IPV6), and a valid IP address and then, click **Save**.

The configured DNS server is shown in Figure 1.

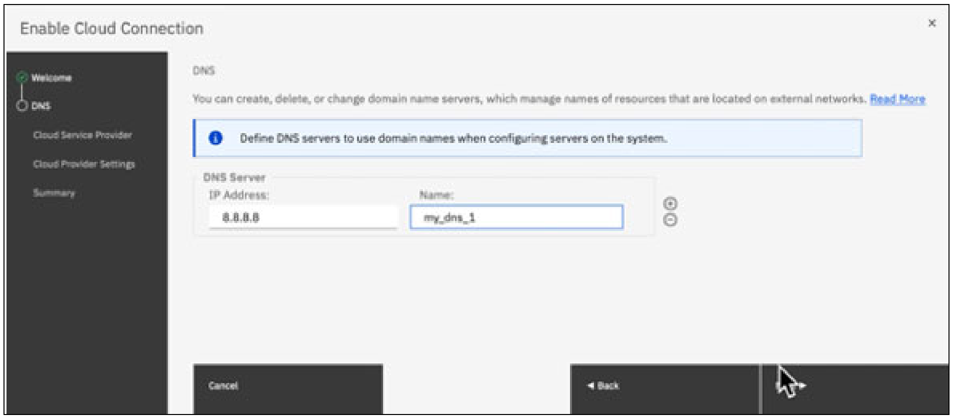

Enable Cloud Connection

DNS

You can create, delete, or change domain name servers, which manage names of resources that are located on external networks. Read More

Define DNS servers to use domain names when configuring servers on the system.

DNS Server

IP Address: 8.8.8.8

Name: my_dns_1

Figure 1 Domain Name Server configuration

> **Note:** A maximum of two DNS server entries can be configured.

Configuring a cloud connection

Complete the following steps to configure a cloud connection:

1. From FlashSystem management GUI, click **Settings** → **Transparent Cloud Tiering** → **Enable cloud connection** to start the cloud connection wizard.

   When the storage system does not use encryption, a warning message is shown. Acknowledge the message by clicking **I understand the risks and want to continue**. Then, click **Next**.

2. Click the name of the cloud provider and then, click **Next** to configure cloud credentials.
3. In the Cloud Provider Settings pane of the Enable Cloud Connection window, enter the configuration details (see Figure 2).

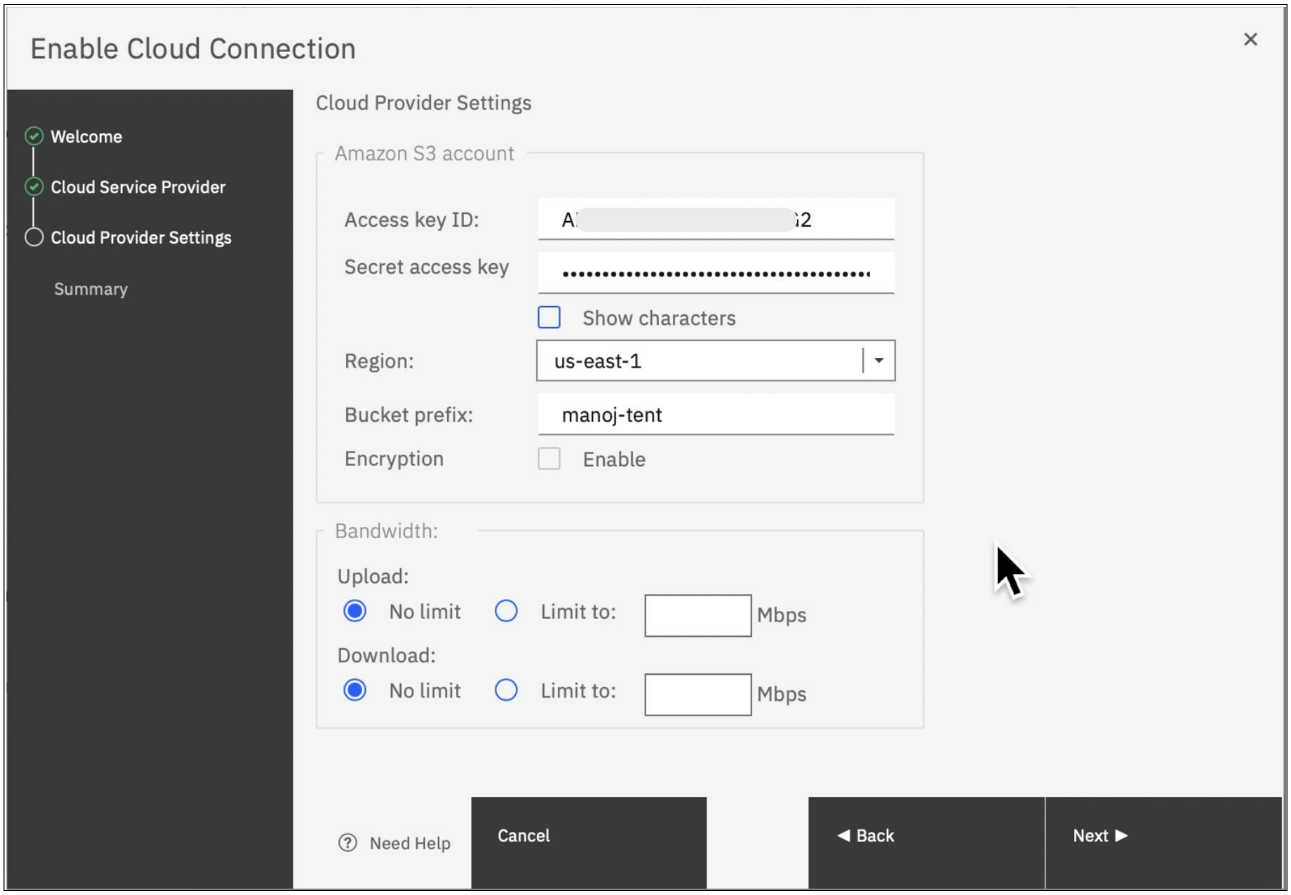

Enable Cloud Connection

Cloud Provider Settings

Amazon S3 account

Access key ID: A[redacted]2

Secret access key: ••••••••••••••••••••••••••••••••••••••

□ Show characters

Region: us-east-1

Bucket prefix: manoj-tent

Encryption □ Enable

Bandwidth:

Upload: ● No limit ○ Limit to: Mbps

Download: ● No limit ○ Limit to: Mbps

Figure 2 Configuring Amazon Cloud S3 account details

4. Click **Next** to submit all the details and configure the S3 bucket.

   Figure 3 shows the successfully configured Amazon S3 account details.

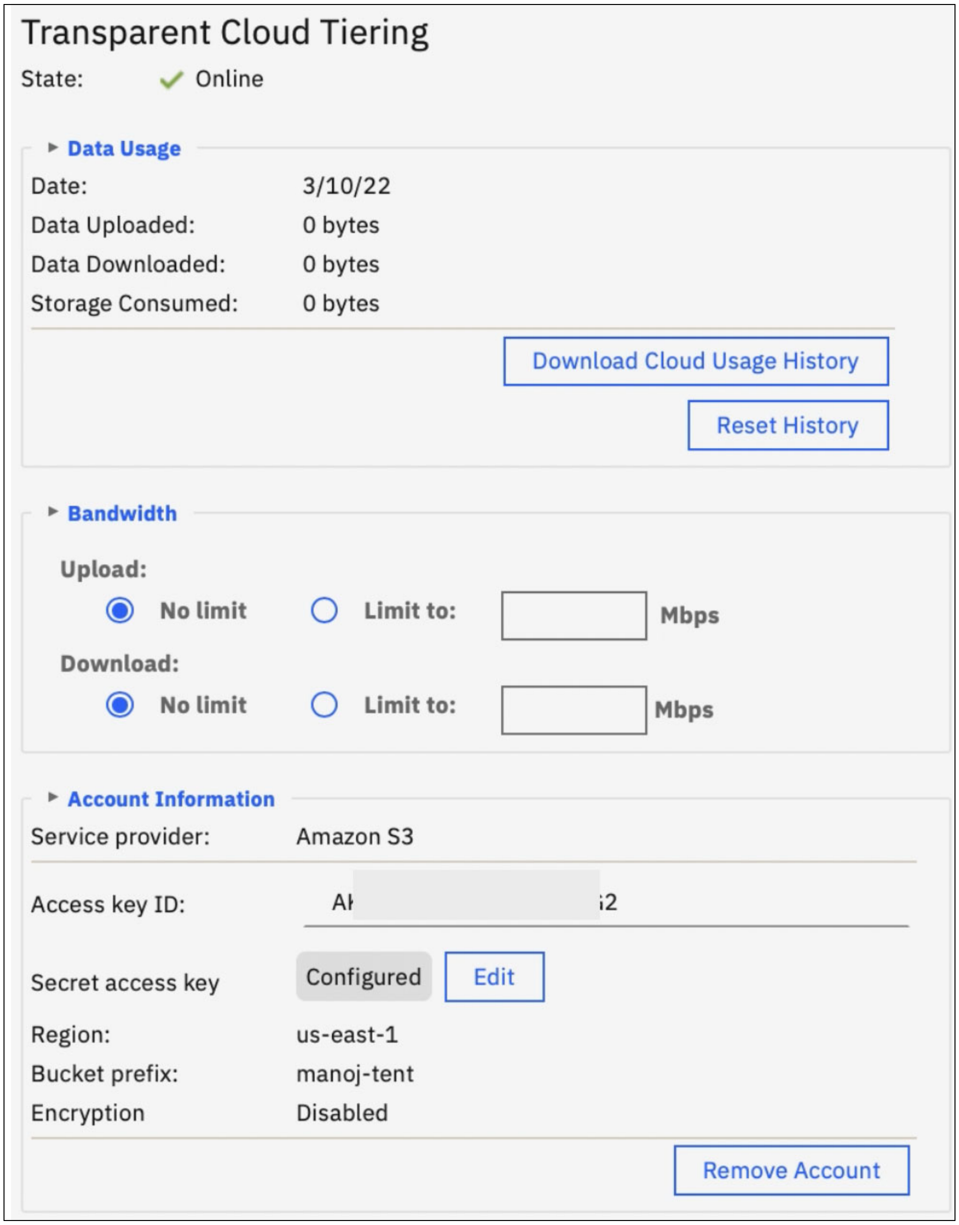

Transparent Cloud Tiering

State: ✓ Online

▸ Data Usage

Date:	3/10/22
Data Uploaded:	0 bytes
Data Downloaded:	0 bytes
Storage Consumed:	0 bytes

▸ Bandwidth

Upload:

● No limit ○ Limit to: Mbps

Download:

● No limit ○ Limit to: Mbps

▸ Account Information

Service provider:	Amazon S3
Access key ID:	[redacted]
Secret access key	Configured
Region:	us-east-1
Bucket prefix:	manoj-tent
Encryption	Disabled

Figure 3 Configured Amazon S3 account

In the following sections, we discuss creating and assigning new snapshot policies, configuring volume groups, and adding the Safeguarded attribute to volume groups.

Configuring snapshot policies for IBM Safeguarded Copy

Complete the following steps to configure snapshot policies on the IBM FlashSystem storage device:

1. From management GUI console, click **Policies** and then, select the **Snapshot Policies** option from the submenu. The system includes three predefined policies.
2. To create a policy, click **Create Snapshot Policy**. A configuration window opens, in which you enter information about the policy name, policy execution frequency, hourly interval, and retention interval fields (see Figure 4).

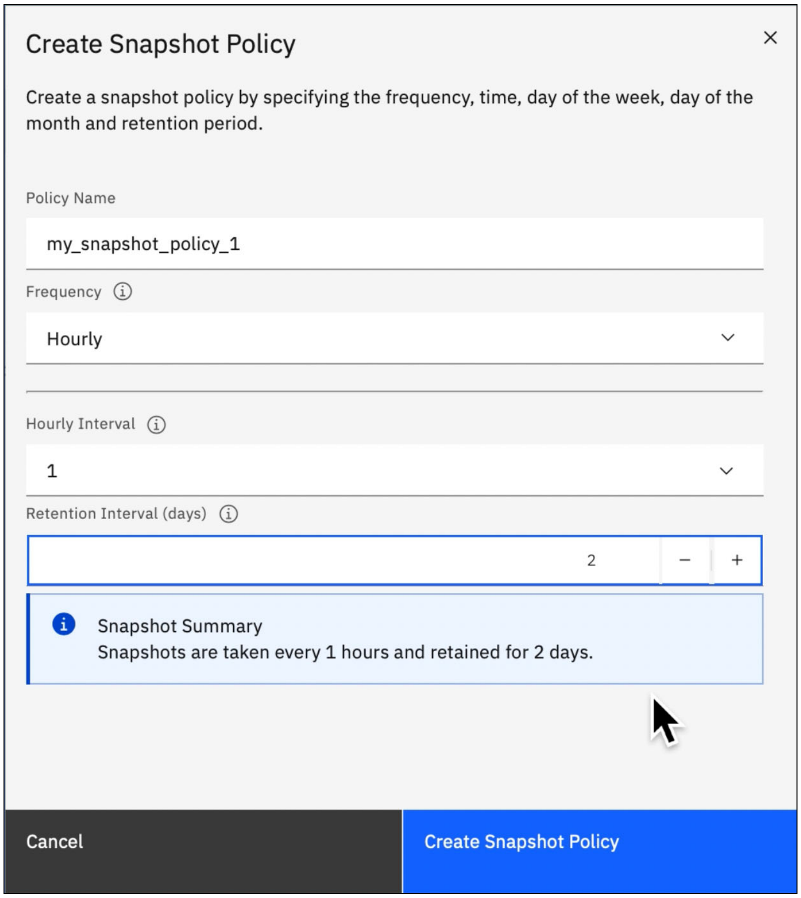

Create Snapshot Policy

Create a snapshot policy by specifying the frequency, time, day of the week, day of the month and retention period.

Policy Name: my_snapshot_policy_1

Frequency: Hourly

Hourly Interval: 1

Retention Interval (days): 2

Snapshot Summary
Snapshots are taken every 1 hours and retained for 2 days.

Figure 4 Creating a snapshot policy

Configuring volume groups

This section describes configuring volume groups on an IBM FlashSystem storage. As the name suggests, *volume groups* allow a configuration or an action (such as a point-in-time snapshot or Safeguarded copy) to be run on a set of volumes. The actions are based on policies that are assigned to a specific volume group.

Complete the following steps:

1. From management GUI console, click **Volumes** and then, select the **Volume Groups** option from the submenu.

   Create the volume group by clicking **Create Volume Group**. A configuration opens in which you enter the Volume group name. (This action always creates an empty volume group.)

2. Click the newly created volume group. From the Actions menu, choose the volumes to add to the newly created volume group.

Figure 5 shows a volume group, with a single volume.

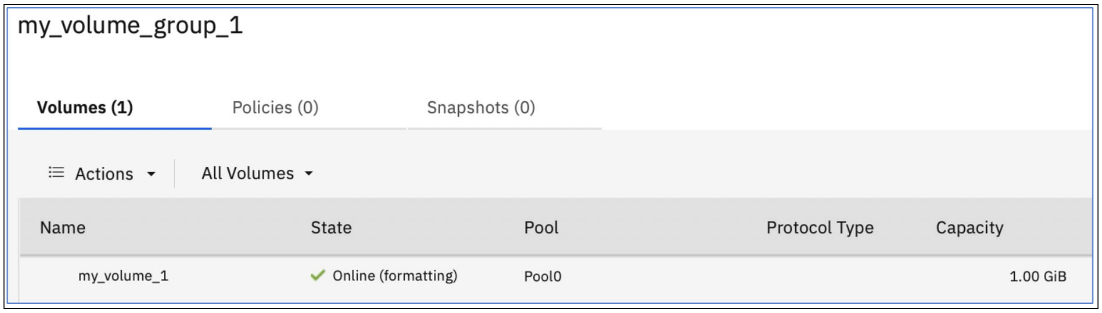

Name	State	Pool	Protocol Type	Capacity
my_volume_1	✓ Online (formatting)	Pool0		1.00 GiB

Figure 5 Volume group with single volume

3. To assign an internal snapshot policy to the volume group, click **Policies** and then, **click Assign internal snapshot policy**.

4. From the Assign Internal Snapshot Policy window, select attributes, such as execution date and time. To make the snapshot immutable, click the **Safeguarded** option and then, click **Assign Policy** (see Figure 6).

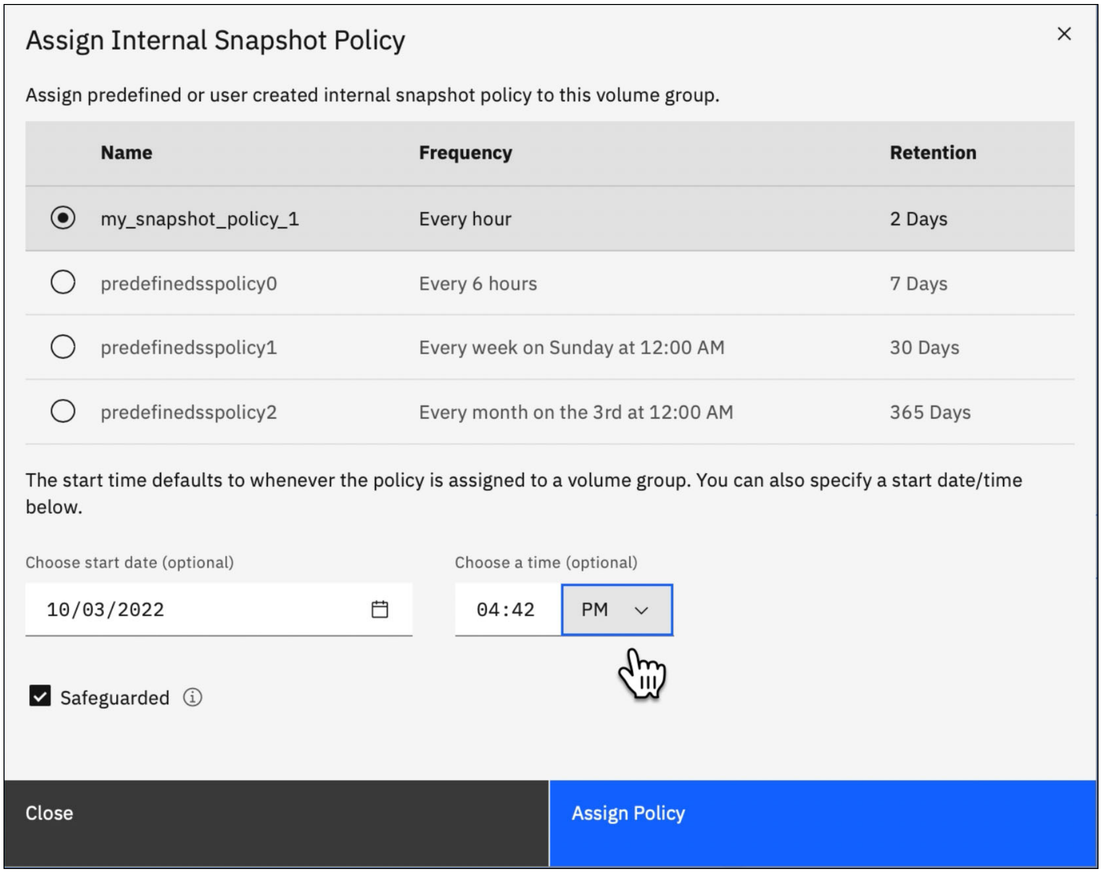

Assign Internal Snapshot Policy

Assign predefined or user created internal snapshot policy to this volume group.

	Name	Frequency	Retention
●	my_snapshot_policy_1	Every hour	2 Days
○	predefinedsspolicy0	Every 6 hours	7 Days
○	predefinedsspolicy1	Every week on Sunday at 12:00 AM	30 Days
○	predefinedsspolicy2	Every month on the 3rd at 12:00 AM	365 Days

The start time defaults to whenever the policy is assigned to a volume group. You can also specify a start date/time below.

Choose start date (optional): 10/03/2022

Choose a time (optional): 04:42 PM

☑ Safeguarded

Figure 6 Assign Internal Snapshot Policy to volume group

Note: The time field denotes the next execution time that is based on the frequency for the policy. Referring to Figure 6, when you choose 04:42 PM as the time, the frequency is added to the time and a next execution time 05:42 PM is assigned. If you want immediately run the policy action, you must back calculate time value that is relative to current time.

A Safeguarded snapshot policy that is assigned to a volume group is shown in Figure 7.

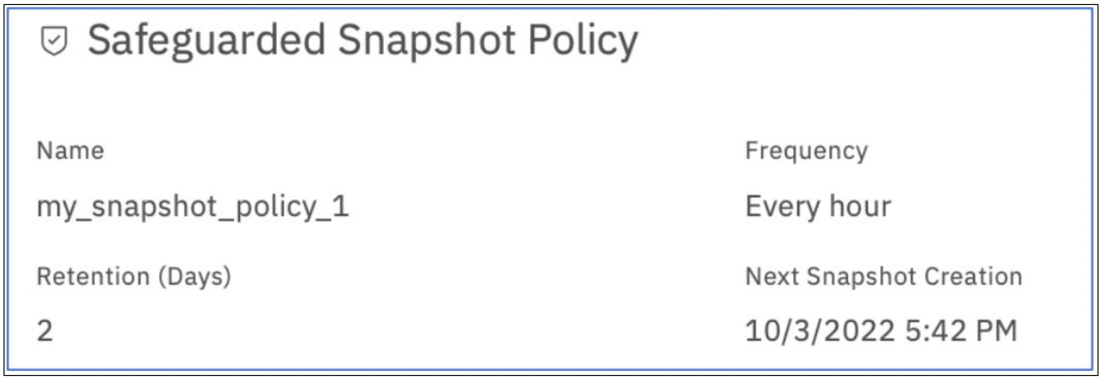

Safeguarded Snapshot Policy

Name: my_snapshot_policy_1

Frequency: Every hour

Retention (Days): 2

Next Snapshot Creation: 10/3/2022 5:42 PM

Figure 7 Safeguarded snapshot policy assigned to volume group

An internal scheduler takes the snapshot at the time interval that is defined by the policy. All the available snapshots for a specific volume group are listed under the Snapshot tab (see Figure 8).

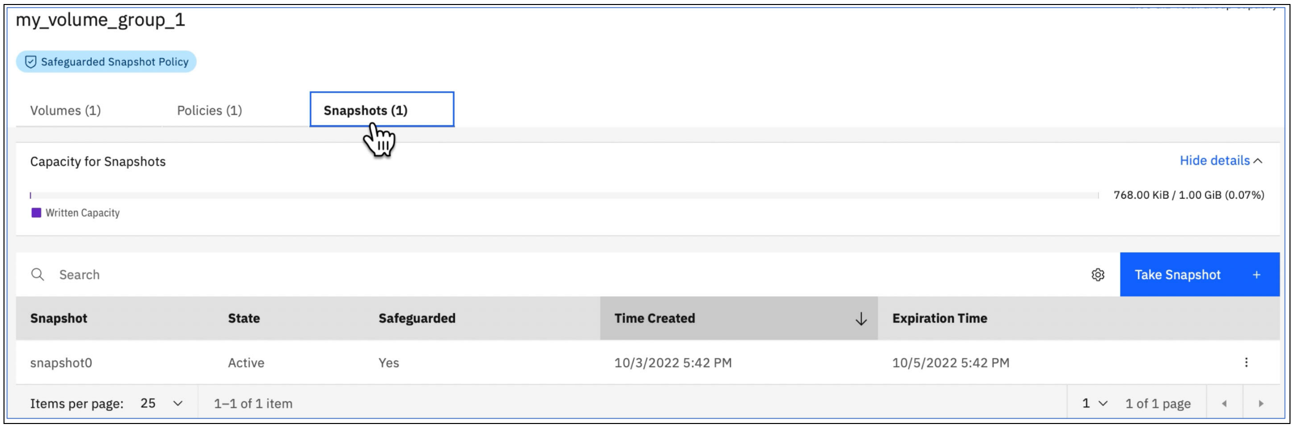

Snapshot	State	Safeguarded	Time Created	Expiration Time
snapshot0	Active	Yes	10/3/2022 5:42 PM	10/5/2022 5:42 PM

Figure 8   Available snapshots for volume group

5. Ad hoc point in time snapshots can be created by clicking **Take Snapshot**. Such snapshots are mutable or not Safeguarded, even if the Safeguarded attribute is selected when the policy is created (see Figure 9).

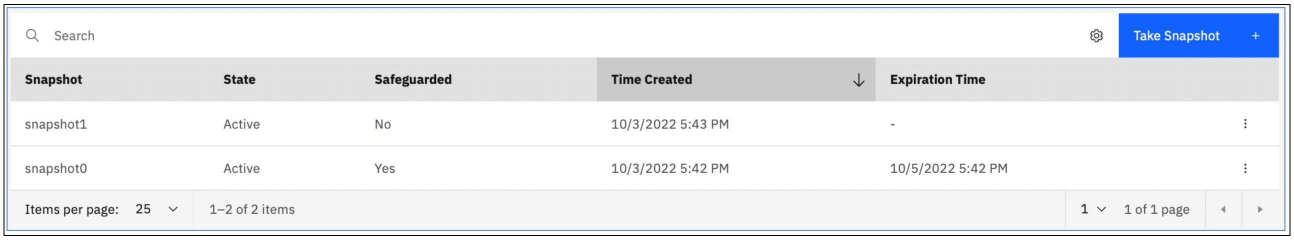

Snapshot	State	Safeguarded	Time Created	Expiration Time
snapshot1	Active	No	10/3/2022 5:43 PM	-
snapshot0	Active	Yes	10/3/2022 5:42 PM	10/5/2022 5:42 PM

Figure 9   Safeguarded and non-safeguarded snapshots

> **Note:** Safeguarded or immutable snapshots are created by the internal scheduler that is based on the Safeguarded attribute of the policy. Alternatively, IBM Safeguarded Copies can be managed by using IBM Copy Services Manager (CSM) or IBM Spectrum Copy Data Management (CDM).

Creating cloud volumes

This section describes how to create cloud volumes.

The cloud volume can be created from Safeguarded copies that were restored or traditional volumes. To create cloud volume from Safeguarded volume, a clone volume must be created.

Complete the following steps:

1. From the list of Safeguarded volumes that are listed under Volume Groups, click the three vertical dots and then, click **Create Clone**.
2. A Create Volume Group window opens. Select the type of clone to create, the target pool, and the I/O group along with and optional name for the clone volume group (see Figure 10).

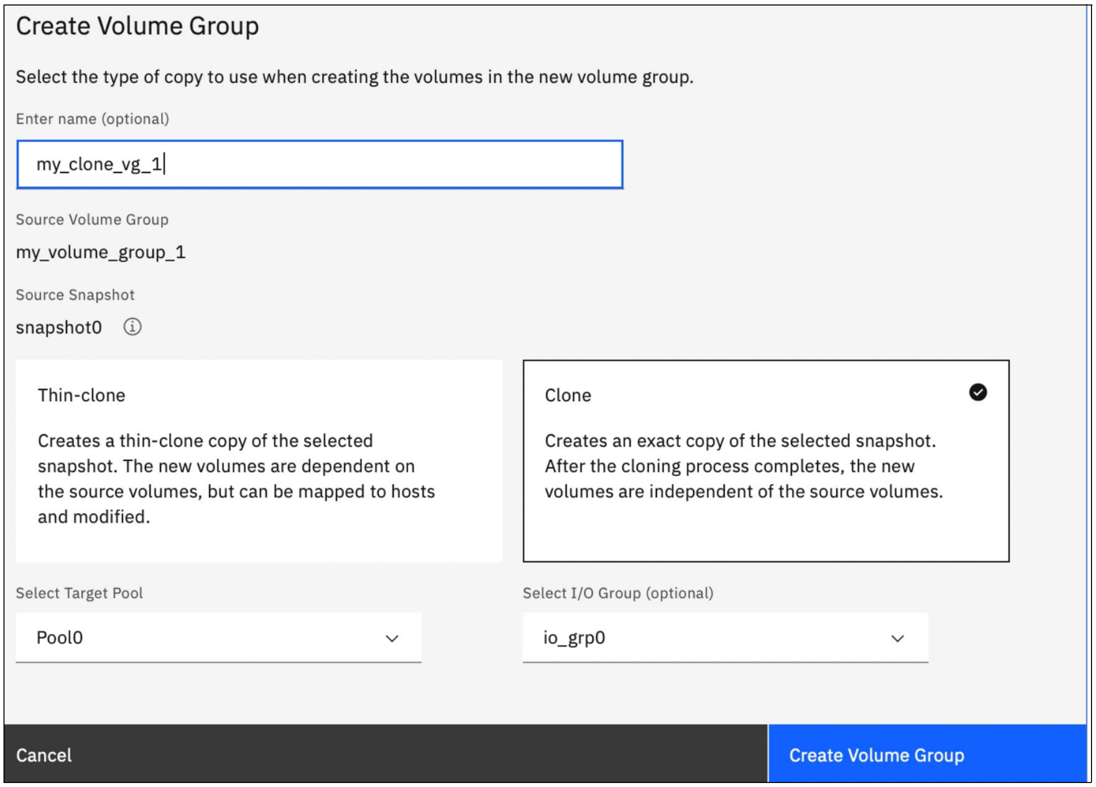

Create Volume Group

Select the type of copy to use when creating the volumes in the new volume group.

Enter name (optional): my_clone_vg_1

Source Volume Group: my_volume_group_1

Source Snapshot: snapshot0

Thin-clone

Creates a thin-clone copy of the selected snapshot. The new volumes are dependent on the source volumes, but can be mapped to hosts and modified.

Clone ✓

Creates an exact copy of the selected snapshot. After the cloning process completes, the new volumes are independent of the source volumes.

Select Target Pool: Pool0

Select I/O Group (optional): io_grp0

Figure 10 Create Volume Group window

Note: Two types of clones are possible: thin-clone and thick-clone. Only thick clones are considered for cloud tiering. Thick clones are indicated as clones.

When the clone operation completes, a clone volume group is created along with clone copies of volumes available in the source volume group.

3. Click the cloned volume group and then, click the list of volumes to select the volume. Click **Actions** and then, from the Cloud Volumes menu option, select **Create Cloud Snapshot**.

4. From the Create Cloud Snapshots for Volumes window, confirm the default settings and click **Create** to start the volume tiering to the preconfigured cloud provider (see Figure 11).

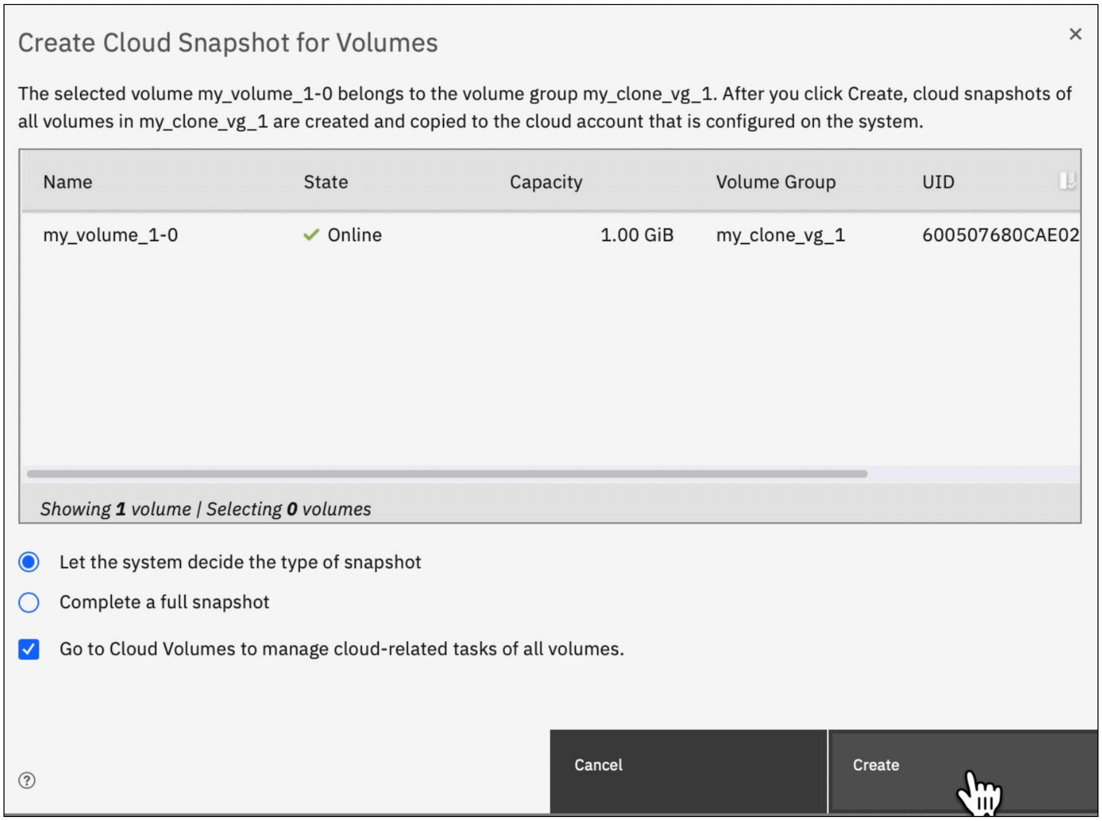

Create Cloud Snapshot for Volumes

The selected volume my_volume_1-0 belongs to the volume group my_clone_vg_1. After you click Create, cloud snapshots of all volumes in my_clone_vg_1 are created and copied to the cloud account that is configured on the system.

Name	State	Capacity	Volume Group	UID
my_volume_1-0	✓ Online	1.00 GiB	my_clone_vg_1	600507680CAE02

*Showing **1** volume | Selecting **0** volumes*

● Let the system decide the type of snapshot

○ Complete a full snapshot

☑ Go to Cloud Volumes to manage cloud-related tasks of all volumes.

Figure 11 Create Cloud Snapshot for Volumes window

5. To view the list cloud volumes and their status, click **Volumes** and select **Cloud Volumes** from the submenu (see Figure 12).

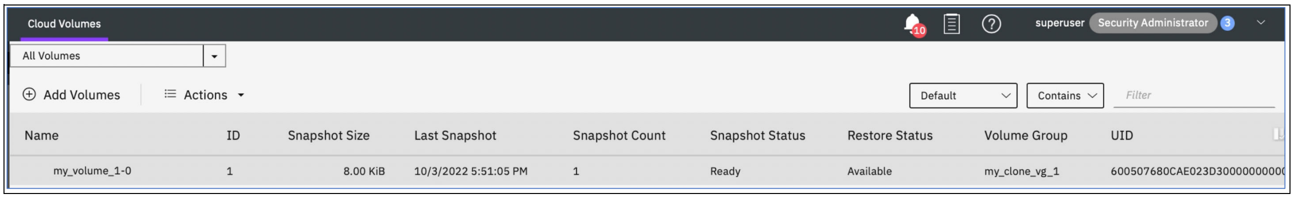

Name	ID	Snapshot Size	Last Snapshot	Snapshot Count	Snapshot Status	Restore Status	Volume Group	UID
my_volume_1-0	1	8.00 KiB	10/3/2022 5:51:05 PM	1	Ready	Available	my_clone_vg_1	600507680CAE023D3000000000(

Figure 12 Available cloud volumes

6. Clean-up any specific version or to remove all versions of a volume, click **Actions** → **Manage Cloud Volumes** (see Figure 13).

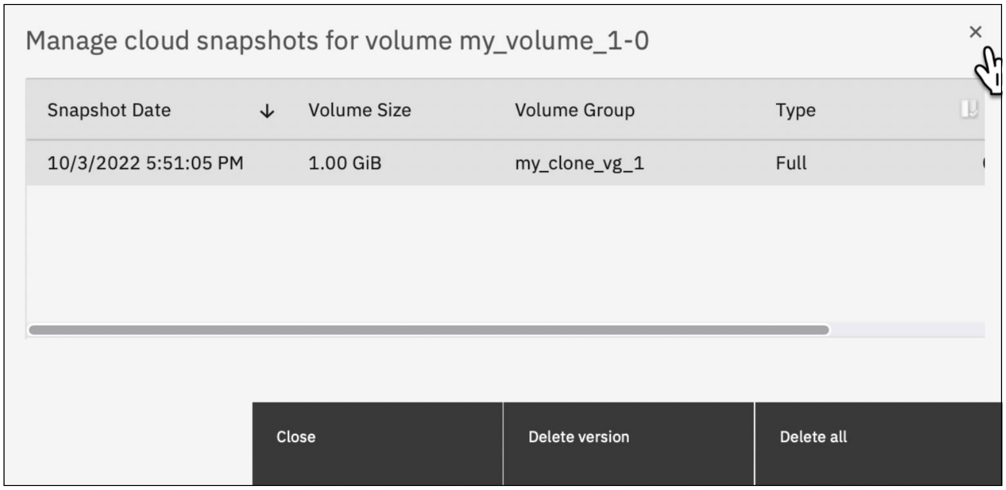

Manage cloud snapshots for volume my_volume_1-0

Snapshot Date	Volume Size	Volume Group	Type
10/3/2022 5:51:05 PM	1.00 GiB	my_clone_vg_1	Full

Figure 13 Manage cloud volumes

Here, the delete operation is possible on the cloud volume because it is a mutable clone that was created from Safeguarded volume. The original Safeguarded volume remains immutable. It cannot be removed from storage until its policy for retention expires nor can it be mapped to host.

Restoring cloud volumes

This section describes how to perform a restore operation from cloud volumes. Such volumes can be restored to the original IBM FlashSystem, or to an alternative instance of IBM Spectrum Virtualize that is running on an IBM FlashSystem, IBM SAN Volume Controller, or IBM Spectrum Virtualize for Public Cloud, which expands the potential recovery zone or analysis options.

Complete the following steps to perform the restore:

1. In the Management GUI console, click **Volumes** and then, click the **Cloud Volumes** option.
2. Right-click the volume from the list of volumes that is displayed under Cloud Volumes. Click the **Restore** option from the menu.

3. In the Restore wizard window, select the snapshot that you want to restore and then, click **Next**. In the next window of the wizard, two potions are available: Restore directly to the production volume or Restore to a new volume, as shown in Figure 14.

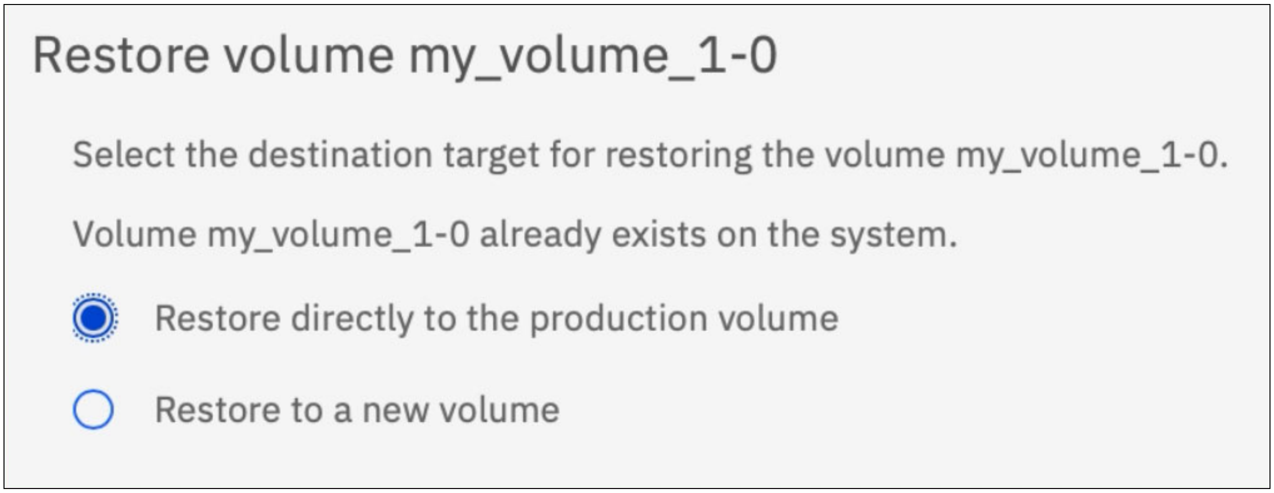

Restore volume my_volume_1-0

Select the destination target for restoring the volume my_volume_1-0.

Volume my_volume_1-0 already exists on the system.

● Restore directly to the production volume

○ Restore to a new volume

Figure 14 Cloud volume restore to an existing volume

If you select **Restore to a new volume**, enter the name for the new volume in the Restore volume window (see Figure 15). The restore action is performed on the production volume which, overwrites the data. (During the restore, the production volume goes offline.)

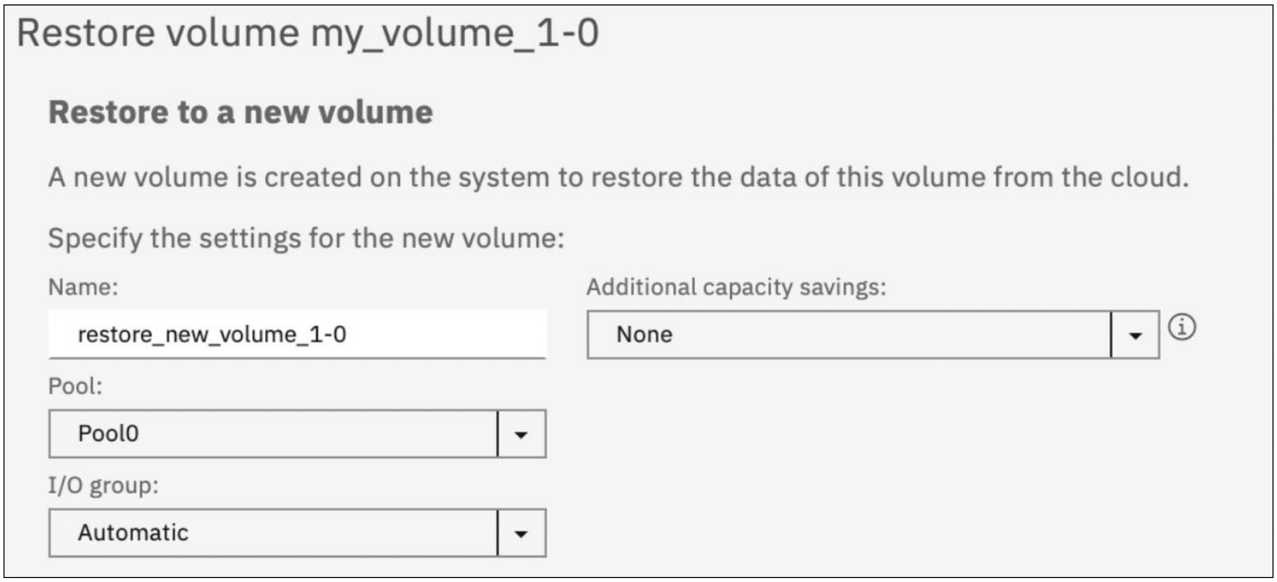

Restore volume my_volume_1-0

Restore to a new volume

A new volume is created on the system to restore the data of this volume from the cloud.

Specify the settings for the new volume:

Name: restore_new_volume_1-0

Additional capacity savings: None

Pool: Pool0

I/O group: Automatic

Figure 15 Cloud volume restore to a new volume

4. Click **Finish** in the Summary window to start the restore.

The Restore Status column under Cloud Volumes shows the restore status of each volume that is restored. The status is changed to `Available` when the restore completes.

The cloud volume creation or restore from cloud volume operations are audited. The events for these operations can be seen by clicking **Monitoring** → **Events**.

A sample event is shown in Figure 16.

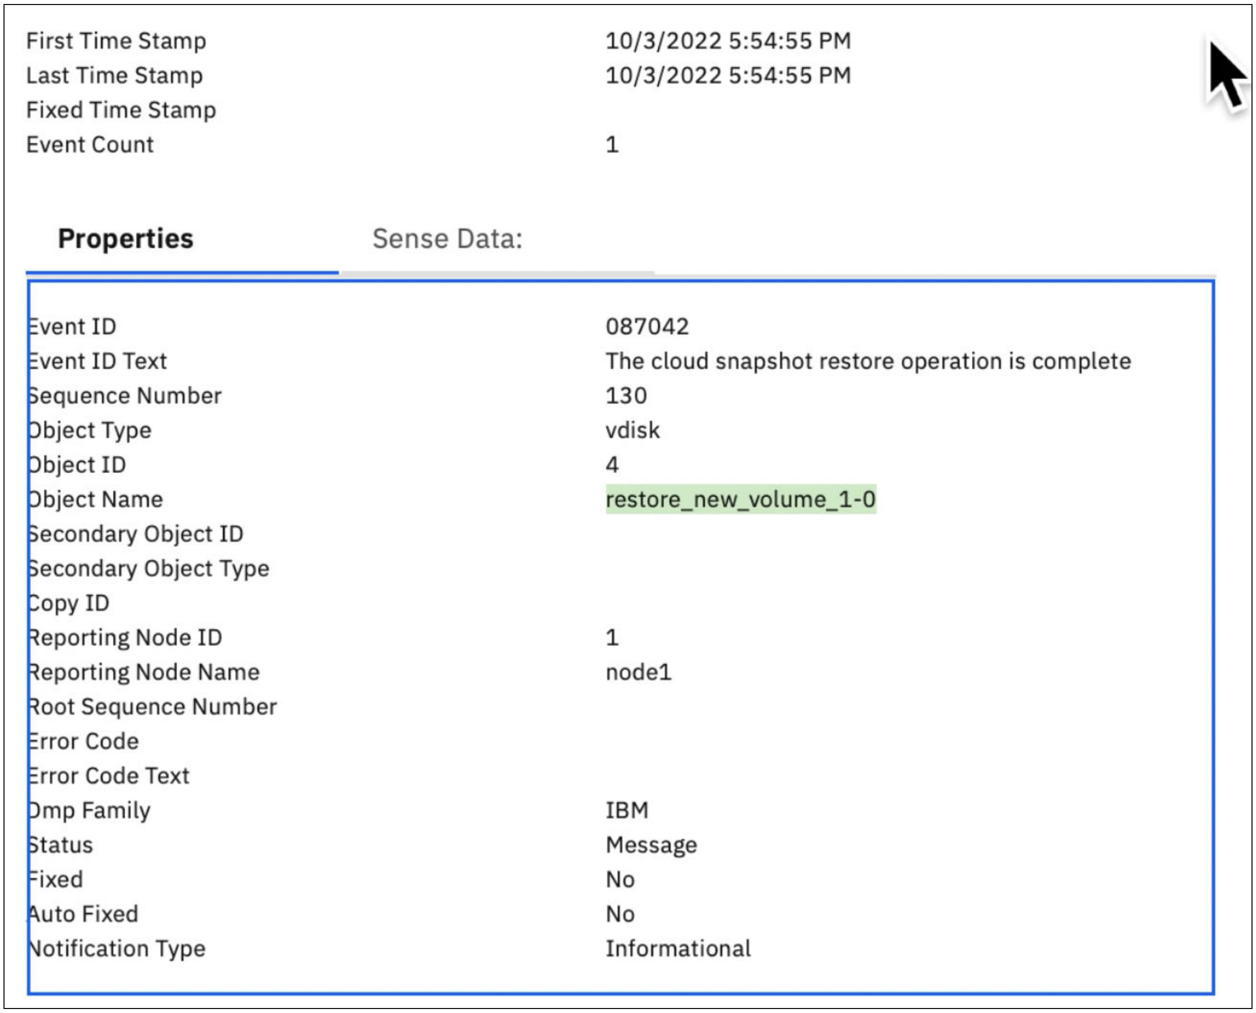

First Time Stamp	10/3/2022 5:54:55 PM
Last Time Stamp	10/3/2022 5:54:55 PM
Fixed Time Stamp	
Event Count	1

Properties Sense Data:

Event ID	087042
Event ID Text	The cloud snapshot restore operation is complete
Sequence Number	130
Object Type	vdisk
Object ID	4
Object Name	restore_new_volume_1-0
Secondary Object ID	
Secondary Object Type	
Copy ID	
Reporting Node ID	1
Reporting Node Name	node1
Root Sequence Number	
Error Code	
Error Code Text	
Dmp Family	IBM
Status	Message
Fixed	No
Auto Fixed	No
Notification Type	Informational

Figure 16 Sample event

Summary

This blueprint explains the use of IBM FlashSystem Transparent Cloud Tiering for Safeguarded and non-Safeguarded volumes. A cloud provider configuration must be created before tiering to cloud volumes is started.

For the Safeguarded volumes, a volume clone must be created before tiering can occur. Cloned volumes that are created from Safeguarded volumes and cloud volumes are mutable.

When performing a restore from cloud volumes, volumes can be overwritten or a restore to new volume is possible.

Authors

This blueprint guide was produced by a team of specialists from around the world working at IBM Redbooks, Poughkeepsie Center.

Shashank Shingornikaris a Storage Solutions Architect with IBM® Systems, ISDL Lab Pune, India, for over 12 years. He has worked extensively with IBM Storage products, such as IBM Spectrum Virtualize, IBM FlashSystems, and IBM Spectrum Scale building solutions that combine Oracle and Red Hat OpenShift features. Currently, he is working on demonstrating cyber resilience solutions with IBM QRadar® and IBM Storage Systems. Before joining IBM, Shashank worked in The Netherlands on various high availability, Disaster Recovery, cluster, and replication solutions for database technologies, such as Oracle, MSSQL, and MySQL.

Manoj Kateja is a Software Engineer with IBM Systems, ISDL Lab Pune, India with over 10 years of industry experience. He has worked extensively with IBM Storage products, such as IBM Spectrum Virtualize and IBM FlashSystem. Currently, he is leading the testing efforts for the Transparent Cloud Tiering feature of IBM FlashSystem.

Christopher Vollmar is a Cyber and Data Resiliency Architect with IBM Technology. He is focused on helping customers design solutions to support cyber resiliency on primary and backup data to complement their cybersecurity practices. He is and author of several IBM Redbooks®, an Enterprise Design Thinking® Co-Creator, and a frequent speaker at customer-focused events, such as IBM TechU.

The authors thank the following contributor for their design and architectural support:

Kosta Makropoulos, Senior Storage Partner Technical Specialist
IBM Technology, Canada

Resources

For more information about the topic that is discussed in this blueprint, see the following resources:

- IBM FlashSystem hardware overview:

  https://www.ibm.com/docs/en/flashsystem-9x00/8.5.x?topic=overview-system

- Configuring transparent cloud tiering:

  https://ibm.com/docs/en/flashsystem-9x00/8.5.x?topic=configuring-transparent-cloud-tiering

- Managing cloud volumes:

  https://ibm.com/docs/en/flashsystem-9x00/8.5.x?topic=administering-managing-cloud-volumes

- Configuring Safeguarded Copy function:

  https://ibm.com/docs/en/flashsystem-9x00/8.5.x?topic=configuring-safeguarded-copy-function

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at `http://www.ibm.com/legal/copytrade.shtml`

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

Redbooks (logo) ®
Enterprise Design Thinking®
IBM®
IBM Cloud®
IBM FlashSystem®
IBM Spectrum®
QRadar®
Redbooks®

The following terms are trademarks of other companies:

OpenShift, Red Hat, are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM’s Privacy Policy at http://www.ibm.com/privacy and IBM’s Online Privacy Statement at http://www.ibm.com/privacy/details in the section entitled “Cookies, Web Beacons and Other Technologies,” and the “IBM Software Products and Software-as-a-Service Privacy Statement” at http://www.ibm.com/software/info/product-privacy.

November 2022

ISBN 073846094x

REDP-5697-00