IBM® Storage

Achieving Hybrid Cloud Cyber Resiliency with IBM Spectrum Virtualize for Public Cloud

Contents

About this document

This document is intended to facilitate the deployment of the Cyber Resiliency solution for IBM® Spectrum Virtualize for Public Cloud. This solution is designed to protect the data on IBM Spectrum® Virtualize for Public Cloud, and the IBM FlashSystem® 9200 from external cyberattacks or insider attacks by using its feature of Transparent Cloud Tiering (TCT) to object storage, such as Amazon S3.

The information in this document is distributed on an as-is basis without any warranty that is either expressed or implied. Support assistance for the use of this material is limited to situations where IBM Spectrum Virtualize for Public Cloud is supported and entitled, and where the issues are specific to a blueprint implementation.

Executive summary

In today's data-driven world, an organization's information and data are considered the most important asset to its business, and they can serve as a key asset for the growth of an organization. As more and more data is collected by businesses, organizations, and companies, data volume is growing at a staggering pace.

With this exponential data growth, there is an increased need to protect the data from various cyberattacks in the form of malware and ransomware. These cyberattacks can have a catastrophic impact on an organization, and can result in devastating financial losses and affect an organization's reputation for many years.

The financial impact of cyberattacks is rising. According to Ponemon's *Cost of a Data Breach Report 2019*[1], the average cost of a data breach is estimated at a shocking USD 3.92 million. Moreover, that same Ponemon's report also placed the average chance of experiencing a data breach over the next two years at 29.6%. Therefore, it's a matter of when, not if.

These cyberattacks can happen in several forms. They can be in the form of malware or ransomware targeted at stealing confidential data or holding users' information for ransom. Sometimes these attacks are targeted to destroy confidential and critical data to cripple organizations. Moreover, according to Verizon[2], 34% of data breaches involved internal actors.

Per Wikipedia[3], Cyber Resiliency refers to an entity's ability to continuously deliver the intended outcome despite cyber events. Assuming that you already have an infrastructure that uses some of the current data protection techniques, such as backups, snapshots, and replication, the next step is expanding your current infrastructure to add the necessary cyber resiliency focus.

[1] https://www.ibm.com/security/data-breach?lnk=ushpv18l1
[2] https://enterprise.verizon.com/resources/reports/2019-data-breach-investigations-report.pdf
[3] https://en.m.wikipedia.org/wiki/Cyber_resilience

Support for the blueprint and its configurations

The Cyber Resiliency solution for IBM Spectrum Virtualize for Public Cloud provides an integrated support experience for clients. The information in this document (referred to throughout as *the Blueprint*) is distributed on an “as is” basis without any warranty that is either expressed or implied. Support for the underlying components that make up this solution are provided by way of the standard procedures and processes that are available for each of those components, as governed by the support entitlement that is available for those components. For more information about these components, see “Prerequisites”.

Requesting assistance

All components of the solutions are part of this unified support structure. Support assistance of the solution that is described in this Blueprint is available by requesting assistance for any of the components in the solution and is the preferred method.

Scope

This Blueprint provides the following information:

- A solutions architecture and related solution configuration workflows, with the following essential software components:
    - IBM Spectrum Virtualize for Public Cloud on AWS
- Detailed technical configuration steps for building an end-to-end solution

This technical report does not include the following:

- Provide scalability and performance analysis from a user perspective
- Provide claims of creating totally isolated air-gap infrastructure
- Replace any official manuals and documents issued by IBM

Prerequisites

This technical report assumes basic knowledge of the following prerequisites:

- IBM Spectrum Virtualize for Public Cloud on AWS installation and configuration
- IBM FlashSystem 9200
- AWS Cloud

National Institute of Standards and Technology framework

As systems became linked with external networks, organizations adopted a *defense-in-depth* security mode so that if the perimeter was breached, there were additional layers of security to protect critical information from falling into the wrong hands. The focus was on the technical aspects of recovery. However, these measures are no longer enough for protection against cyberattacks.

Organizations are beginning to understand that traditional device-centric and technology-centric security measures, such as firewalls, fail to provide security in the cyber ecosystem. Moving forward, you must take a holistic approach across your data, applications, and the entire infrastructure to not only recover, but prevent (or at the very least minimize) the attack.

Some of the following factors are considered for designing a Cyber Resiliency approach:

- Although regulations continue to play an important role, consumers decide the ultimate outcomes for a business.
- To implement an effective Cyber Resiliency approach, it must be changed from a reactive approach to a proactive approach. A repeated cycle of planning, protecting, testing, and learning must be implemented by a Cyber Resiliency team.
- Most organizations' backup and disaster recovery plans are designed around the fact that most disasters are caused by either technical failures or human errors, with secondary concern about natural disasters. Modern data protection approaches must also consider data compromise due to cyber events and be implemented accordingly.
- As attackers are getting smarter, approaches must consider continuous improvements, innovations, and reengineering to address the newer threats that are challenging organizations.
- Though effort is made to extend existing infrastructure, modern technologies help automate systems to deal more effectively with cyber threats.

In order to effectively deal with cyber events, the National Institution of Standards and Technology (NIST) provides a policy framework of computer security guidance regarding how organizations can assess and improve their ability to prevent, detect, and respond to cyberattacks. This framework is an industry-accepted methodology for building a plan to develop and implement safeguards to ensure delivery of critical business services.

As shown in Figure 1, a Cyber Resiliency plan is a continuous process that needs to be repeated in the environment to safeguard data from cyberattacks.

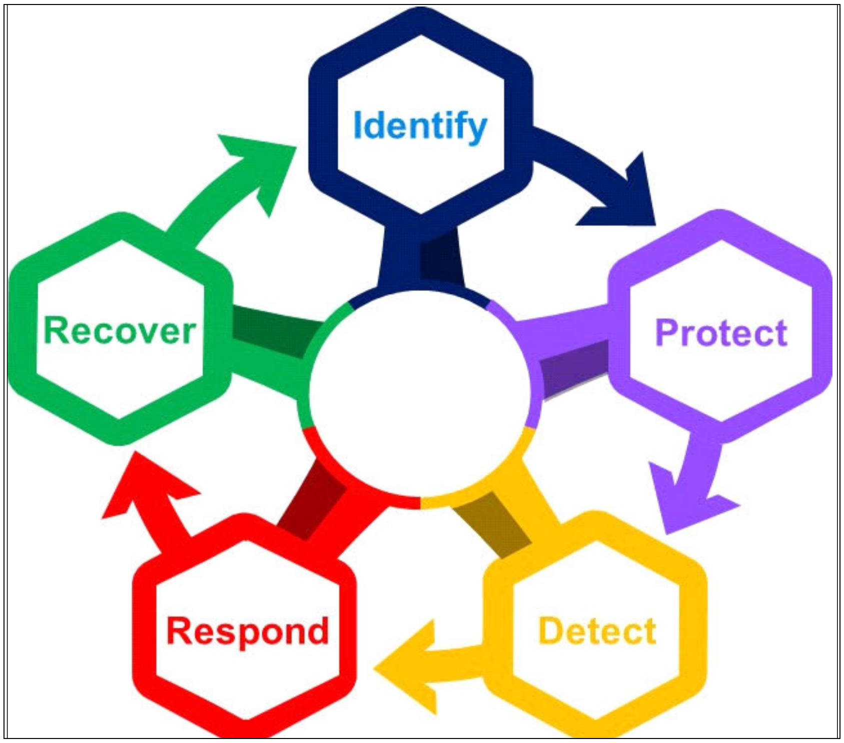

Figure 1 NIST Cybersecurity Framework

The NIST framework is a set of five Cybersecurity functions:

- **Identify**: NIST recommends building organizational understanding during the Identify stage so that business IT systems can be confidently restored to their operational state. It is important to identify what must be protected, and then prioritize your protection plan.
- **Protect**: During the Protect stage, implement various safeguards, such as identity management, access control, awareness and training, data security, code currency procedures, and data protection technology, to ensure delivery of critical services.
- **Detect**: The best way to reduce costs during an event is to detect it early, and then rapidly recover. The point of the Detect stage is implementing activities and technologies to identify anomalies and events that are out of the ordinary. This enables you to respond quickly and limit the damage by containing the event.
- **Respond**: In the Response state, develop and implement appropriate activities to take actions regarding a detected cyber security incident.
- **Recover**: In the Recover stage, develop and implement appropriate activities to maintain plans for resilience, and to restore any capabilities or services that were impaired due to a cybersecurity incident. In this stage, the goal is to get a compromised environment back up and running quickly and efficiently.

Cyber Resiliency solution with IBM Spectrum Virtualize for Public Cloud on AWS

This section describes the components and solution building blocks used for implementing a Cyber Resiliency solution using IBM Spectrum Virtualize.

IBM Spectrum Virtualize for Public Cloud on AWS

IBM Spectrum Virtualize for Public Cloud is a version of IBM Spectrum Virtualize implemented in a cloud environment.

Designed for public cloud *infrastructure as a service* (IaaS), IBM Spectrum Virtualize for Public Cloud represents a solution for public cloud implementations, and includes technologies that both complement and enhance public cloud IaaS offering capabilities.

IBM Spectrum Virtualize for Public Cloud provides for the deployment of IBM Spectrum Virtualize-based software in public clouds, starting with IBM Cloud™, and is now available in Amazon AWS. This new offering with IBM Spectrum Virtualize for Public Cloud on AWS is a bring you own license (BYOL) offering, which can be purchased as either a perpetual license or a monthly license.

IBM Spectrum Virtualize for Public Cloud can be deployed on AWS IaaS via the AWS Marketplace to enable hybrid cloud solutions, offering the ability to transfer data between on-premises data centers using any IBM Spectrum Virtualize-based appliance and AWS. For details, see *IBM Spectrum Virtualize for Public Cloud on AWS Implementation Guide.*

IBM FlashSystem 9200

The IBM FlashSystem 9200 combines the performance of flash and a Non-Volatile Memory Express (NVMe)-optimized architecture with the reliability and innovation of IBM FlashCore® technology and the rich feature set and high availability of IBM Spectrum Virtualize. This powerful new storage platform provides the following advantages:

- The option to use large capacity IBM FlashCore modules (FCM) with inline-hardware compression, data protection, and innovative flash management features; industry standard NVMe drives; or Storage Class Memory (SCM) drives.
- The software-defined storage functionality of IBM Spectrum Virtualize with a full range of industry-leading data services such as dynamic tiering, IBM FlashCopy® management, data mobility, and high-performance data encryption, among many others.
- Innovative data reduction pool (DRP) technology that includes deduplication and hardware-accelerated compression technology, plus SCSI UNMAP support and all of the thin provisioning, copy management, and efficiency you'd expect from IBM Spectrum Virtualize-based storage.

IBM Spectrum Virtualize provides the data services foundation for every IBM FlashSystem 9200 solution. Its industry-leading capabilities include a wide range of data services that can be extended to over 450 IBM and non-IBM heterogeneous storage systems; automated data movement; synchronous and asynchronous copy services (either on-premises or to the public cloud); encryption; high-availability configurations; storage tiering; and data reduction technologies, among many others.

To further drive your IT transformation, IBM Spectrum Virtualize for Public Cloud offers multiple ways to create hybrid cloud solutions between on-premises private clouds and the public cloud. It enables real-time storage-based data replication and disaster recovery, and data migration between local storage and IBM Cloud. Furthermore thanks to its software-defined storage nature, IBM Spectrum Virtualize enables storage administration at a cloud service provider's site in the same way as on-premises, regardless of the type of storage.

Use cases

The architectural design in this Cyber Resiliency solution addresses the following use cases:

- As a storage architect and administrator, data should be safeguarded from virus attacks, ransomware encryption, or deletion by a malicious user.
- As a storage architect and administrator, data is a most-important asset, and the business of my organization relies on the data on the storage system. Business can continue even if the data on the primary system holding the data has been compromised.
- Multiple copies of data are maintained using multiple features of data protection, even if one or more copies of data are compromised.
- Copies of data are available in an immutable format to avoid overriding valid copies of data. This state provides the ability to restore valid copies of the data at a remote system to validate the authenticity of recovered data.
- Copies of data are stored in an air-gapped environment where only authorized personnel have access to the data.
- Avoid people accessing and compromising all copies of data, with a provision to store multiple copies of data at different locations, and to separate administrative access for the different copies of data.

Architectural overview and approaches

Figure 2 on page 7 shows the high-level architectural overview of a Cyber Resiliency solution to achieve protection of data on an IBM Spectrum Virtualize for Public Cloud on AWS.

The following different approaches are described in this Blueprint:

1. Back up the source volume from IBM Spectrum Virtualize for Public Cloud on AWS to Amazon S3 and restore it back on the same IBM Spectrum Virtualize for Public Cloud on AWS instance. The different scenarios are described in detail in the following sections.
2. Back up the source volume from an IBM Spectrum Virtualize for Public Cloud on AWS instance running in one AWS availability zone to Amazon S3, and restore the volume to a different IBM Spectrum Virtualize for Public Cloud on AWS instance running in a different AWS availability zone (see Figure 51 on page 18).
3. Back up the source volume from IBM FlashSystem 9200 on-premises storage to Amazon S3, and restore the volume to an IBM Spectrum Virtualize for Public Cloud on AWS instance running in AWS (see Figure 55 on page 19).

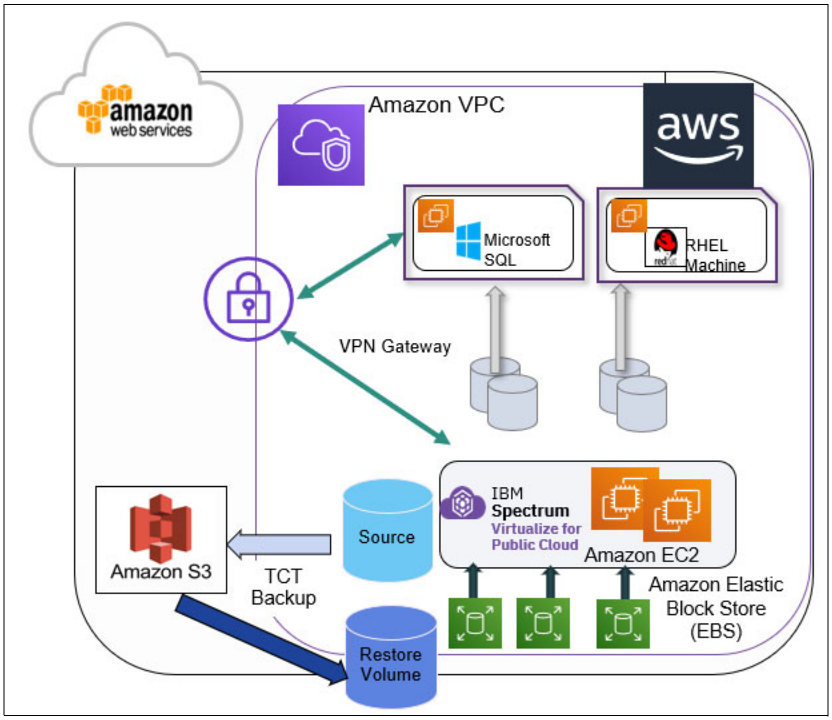

Figure 2 Architectural overview

Use-case Scenario I

In this test setup and validation, IBM Spectrum Virtualize for Public Cloud on AWS is used with the function feature called Transparent Cloud Tiering (TCT). IBM Spectrum Virtualize Transparent Cloud Tiering supports creating connections to cloud service providers to store copies of volume data on public cloud storage, such as Amazon S3, freeing up capacity on the system. The source volume copy is backed up to an Amazon S3 bucket, and can be restored back to the same original volume or a new volume.

This section covers Transparent Cloud Tiering features and functions, and how these functions help administrators create point-in-time snapshots of data on a system. Then they can copy and store the snapshots on cloud storage, enabling administrators to restore snapshots from the cloud for disaster recovery purposes.

The process for using Transparent Cloud Tiering is described in the following section. The first step is to create the cloud account on IBM Spectrum Virtualize for Public Cloud on AWS instance. For details about creating the cloud account, see Enabling a cloud connection to Amazon S3.

Cloud account

The cloud account `cloudaccount0` of type `awss3` is configured as cloud storage on IBM Spectrum Virtualize, as shown in Figure 3.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>lscloudaccount
id name          type  status mode   active_volume_count backup_volume_count
0  cloudaccount0 awss3 online normal 3                   3
```

Figure 3 Cloud account

Backup and restoration process

The following details show different options to create single point-in-time backups and different scenarios:

- Single point-in-time backup and restore to new volume
- Incremental backup and restore to the same volume
- Restore from different generations

Complete the following steps:

1. For the lab setup, `tct_win1` and `tct_win2` are the volumes from the list shown in Figure 4, which are used in the different scenarios.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>lsvdisk
id name            IO_group_id IO_group_name status  mdisk_grp_id mdisk_grp_name capacity type    FC_id FC_name RC_id RC_name vdisk_UID                        fc_map_count
py_count parent_mdisk_grp_id parent_mdisk_grp_name owner_id owner_name formatting encrypt volume_id volume_name     function protocol
0  vdisk0          0              io_grp0       online  0            mdiskgrp0      4.00MB   striped many  many                    6005076072D861A04000000000000002 2
       0                     mdiskgrp0                                     no         no      0         vdisk0
1  TCT_test        0              io_grp0       online  0            mdiskgrp0      1.00GB   striped many  many                    6005076072D861A04000000000000006 2
       0                     mdiskgrp0                                     no         no      1         TCT_test
2  vdisk1          0              io_grp0       online  0            mdiskgrp0      4.00MB   striped 0     fcmap0                  6005076072D861A04000000000000007 1
       0                     mdiskgrp0                                     no         no      2         vdisk1
3  vdisk3          0              io_grp0       online  0            mdiskgrp0      4.00MB   striped                               6005076072D861A04000000000000005 0
       0                     mdiskgrp0                                     no         no      3         vdisk3
4  vdisk2          0              io_grp0       online  0            mdiskgrp0      4.00MB   striped 1     fcmap1                  6005076072D861A04000000000000008 1
       0                     mdiskgrp0                                     no         no      4         vdisk2
5  vdisk4          0              io_grp0       offline 0            mdiskgrp0      1.00GB   striped 2     fcmap2                  6005076072D861A04000000000000009 1
       0                     mdiskgrp0                                     no         no      5         vdisk4
6  vdisk5          0              io_grp0       online  0            mdiskgrp0      1.00GB   striped 3     fcmap3                  6005076072D861A0400000000000000A 1
       0                     mdiskgrp0                                     no         no      6         vdisk5
7  tct-restore-vol 0              io_grp0       online  0            mdiskgrp0      1.00GB   striped                               6005076072D861A0400000000000000B 0
       0                     mdiskgrp0                                     no         no      7         tct-restore-vol          scsi
8  tct_win1        0              io_grp0       online  0            mdiskgrp0      30.00GB  striped                               6005076072D861A0400000000000000E 0
       0                     mdiskgrp0                                     no         no      8         tct_win1
9  tct_win2        0              io_grp0       online  0            mdiskgrp0      30.00GB  striped many  many                    6005076072D861A0400000000000000F 2
       0                     mdiskgrp0                                     no         no      9         tct_win2                 scsi
10 vdisk6          0              io_grp0       online  0            mdiskgrp0      30.00GB  striped 4     fcmap4                  6005076072D861A04000000000000010 1
       0                     mdiskgrp0                                     no         no      10        vdisk6
11 vdisk7          0              io_grp0       online  0            mdiskgrp0      30.00GB  striped 5     fcmap5                  6005076072D861A04000000000000011 1
       0                     mdiskgrp0                                     no         no      11        vdisk7
```

Figure 4 The lsvdisk output

2. Use the `tct_win2` volume as a source volume and mapped to a Microsoft Windows host, as shown in Figure 5.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>lsvdiskhostmap 9
id name     SCSI_id host_id host_name vdisk_UID                        IO_group_id IO_group_name mapping_type host_cluster_id host_cluster_name protoco
9  tct_win2 0       1       windows   6005076072D861A0400000000000000F 0           io_grp0       private                                        scsi
```

Figure 5 Example of vdisk mapping

3. On the Windows host, format `tct_win2` to `D:\`, as shown in Figure 6.

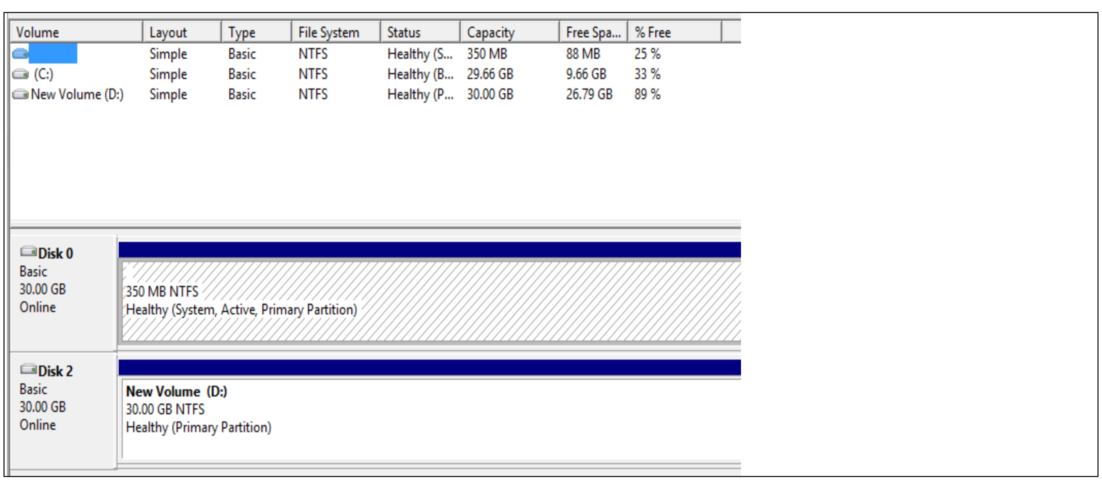

Figure 6 Disk management

4. Copy files to the `D:` drive, as shown in Figure 7.

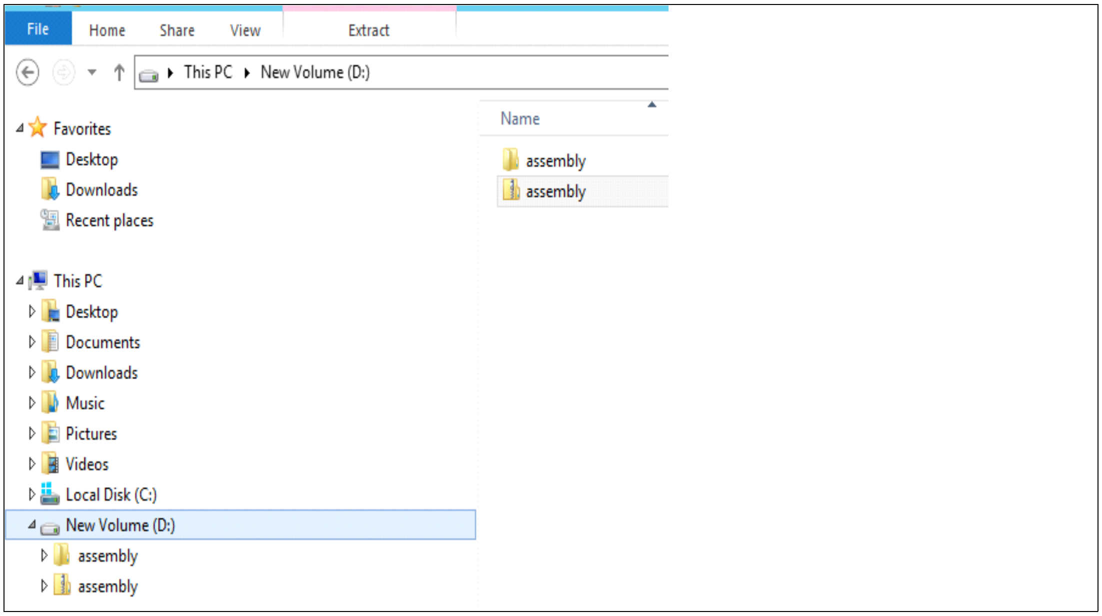

Figure 7 D: drive content

5. Next, take a backup of volume `tct_win2` using the **`backupvolume`** command, as shown in Figure 8.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>backupvolume 9
```

Figure 8 Backup volume

6. Verify the backup job using the **`lsvolumebackup`** command, as shown in Figure 9.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>lsvolumebackup
volume_UID                       volume_id volume_name volume_group_id volume_group_name cloud_account_id cloud_account_name last_backup_time generation_count backup_size
6005076072D861A04000000000000002 0         vdisk0                                        0                cloudaccount0      190606093941     1                0.01MB
6005076072D861A04000000000000006 1         TCT_test                                      0                cloudaccount0      190607072230     2                35.96MB
6005076072D861A0400000000000000F 9         tct_win2                                      0                cloudaccount0      190610091132     1                2.03GB
```

Figure 9 List volume backup

7. Verify the generation of the backup using the **`lsvolumebackupgeneration`** command, as shown in Figure 10. Note the type of `full` and the value of `backup_time`.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>lsvolumebackupgeneration -volume 9
generation_id backup_time  volume_group_name volume_size type state    cloud_upload_size
1             190610091132                   30.00GB     full complete 2.03GB
```

Figure 10 List generation of backed up volumes

Restore to a new volume

In this section, the backed-up volume is restored to a new volume and validated by mapping it to the same Windows host:

1. First, unmap the source volume from the Windows host, as shown in Figure 11.

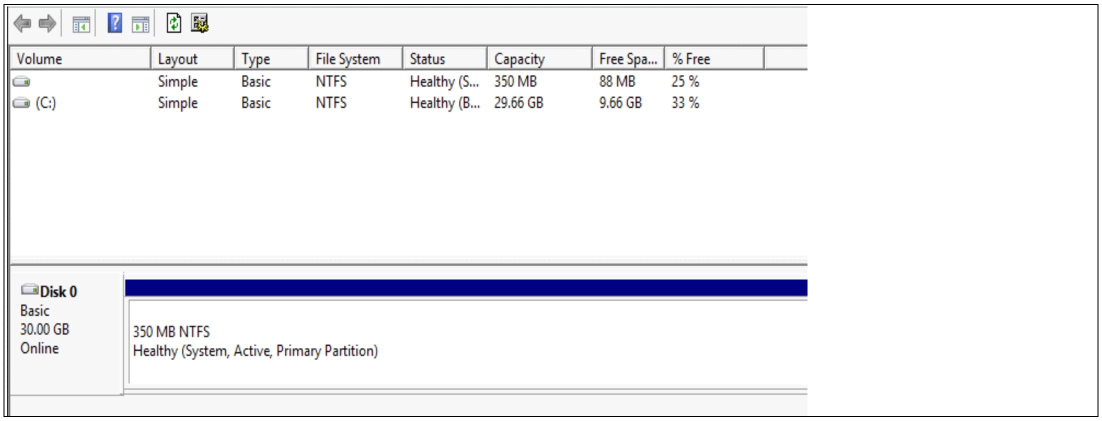

Figure 11 Unmap the source drive from the Windows host

2. Next, restore the volume from the backed up volume to a new volume `tct_win1`, using the **`restorevolume`** command, as shown in Figure 12.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>restorevolume -fromuid 6005076072D861A04000000000000000F -generation 1 tct_win1
```

Figure 12 Restore to new volume

3. Check the status of the restoration using the **`lsvdisk <volume>`** command, shown in Figure 13.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>lsvdisk 9 | grep rest
restore_status restoring
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>
```

Figure 13 Restore status

4. Using the **`datapath`** query command, verify that the new volume is mapped, as shown in Figure 14.

```
PS C:\Program Files\ibm\SDDDSM> .\datapath.exe query device

Total Devices : 1


DEV#:   0  DEVICE NAME: Disk1 Part0            TYPE: 2145       POLICY: LEAST I/O AND WEIGHT
SERIAL: 6005076072D861A04000000000000000E      Reserved: No     LUN SIZE: 30.0GB
HOST INTERFACE: ISCSI
PREFERRED PATH SET: None
===========================================================================================
Path#      Adapter/Hard Disk                   State        Mode        Select     Errors
    0      Scsi Port3 Bus0/Disk1 Part0          OPEN        NORMAL         223          0
PS C:\Program Files\ibm\SDDDSM> _
```

Figure 14 Datapath output

5. Import the volume and assign a drive letter, as shown in Figure 15.

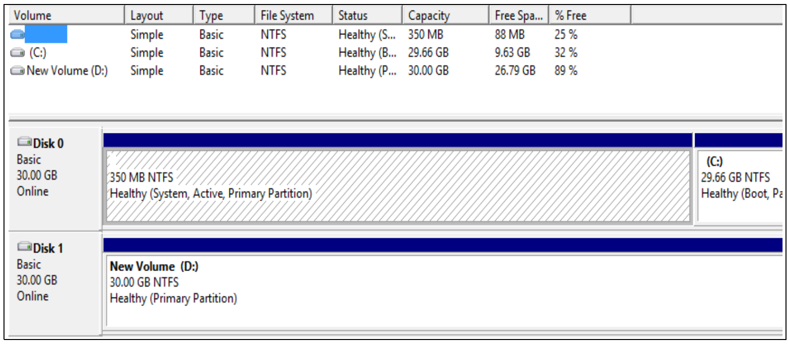

Figure 15   Import volume on Windows host

6. Next verify that the files on drive `D:` are the same as the files that were backed up, as shown in Figure 16.

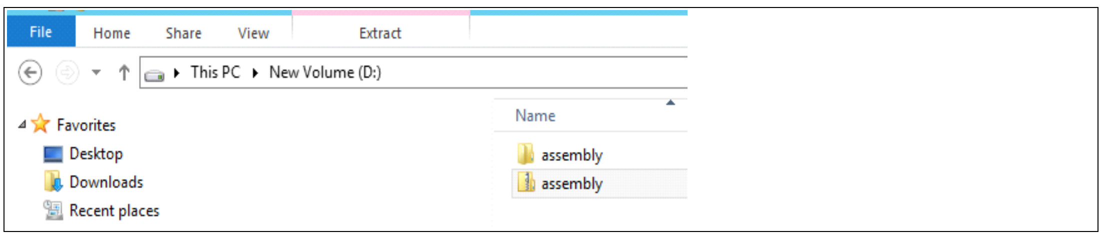

Figure 16   List content of drive D

Incremental backup and restoration to the same volume

This section covers incremental backups and restoration of a volume from different points in time:

1. Add a new directory to drive `D:` to add incremental data to, as shown in Figure 17.

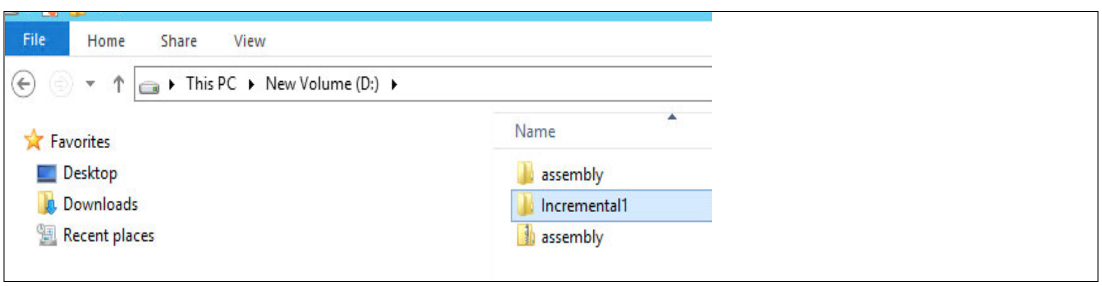

Figure 17   New directory

2. Take a backup of the volume using the **`backupvolume`** command, as shown in Figure 18.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>backupvolume 9
```

Figure 18   Backup volume

3. Run the **lsvolumebackup** command to list the latest backup job. Notice that the generation count has increased to 2 for this volume, as shown in Figure 19.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>lsvolumebackup
volume_UID                       volume_id volume_name volume_group_id volume_group_name cloud_account_id cloud_account_name last_backup_time generation_count backup_size
6005076072D861A04000000000000002 0         vdisk0                                        0                cloudaccount0      190606093941     1                0.01MB
6005076072D861A04000000000000006 1         TCT_test                                      0                cloudaccount0      190607072230     2                35.96MB
6005076072D861A0400000000000000F 9         tct_win2                                      0                cloudaccount0      190613093728     2                2.64GB
```

Figure 19 List volume backup

4. Run the **lsvolumebackupgeneration** command to list all of the generations of backups for that volume, as shown in Figure 20.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>lsvolumebackupgeneration -volume 9
generation_id backup_time  volume_group_name volume_size type        state    cloud_upload_size
1             190610091132                   30.00GB     full        complete 2.03GB
2             190613093728                   30.00GB     incremental copying  0.00MB
```

Figure 20 Generated list of backed up volumes

5. Check until the status of the backup of incremental copy is `complete`, as shown in Figure 21.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>lsvolumebackupgeneration -volume 9
generation_id backup_time  volume_group_name volume_size type        state    cloud_upload_size
1             190610091132                   30.00GB     full        complete 2.03GB
2             190613093728                   30.00GB     incremental complete 520.09MB
```

Figure 21 Status of the incremental backup

6. Add more incremental data to the same drive, as shown in Figure 22.

This PC ▸ New Volume (D:) ▸

Name	Date modified	Type	Size
assembly	6/10/2019 9:00 AM	File folder	
Incremental1	6/13/2019 8:42 AM	File folder	
Incremental2	6/13/2019 10:56 AM	File folder	
assembly	6/10/2019 9:04 AM	Compressed (zipp...	939,804 KB

Figure 22 Adding incremental data

7. Run the **backupvolume** command to create another incremental backup on the same volume, as shown in Figure 23.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>backupvolume 9
```

Figure 23 Backup volume

8. List the backup jobs. Note that the generation count has increased to 3, as shown in Figure 24.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>lsvolumebackup
volume_UID                       volume_id volume_name volume_group_id volume_group_name cloud_account_id cloud_account_name last_backup_time generation_count backup_size
6005076072D861A04000000000000002 0         vdisk0                                        0                cloudaccount0      190606093941     1                0.01MB
6005076072D861A04000000000000006 1         TCT_test                                      0                cloudaccount0      190607072230     2                35.96MB
6005076072D861A0400000000000000F 9         tct_win2                                      0                cloudaccount0      190613120558     3                3.21GB
```

Figure 24 List volume backup

9. List the generations of volumes backed up, and verify that the state is complete, as shown in Figure 25.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>lsvolumebackupgeneration -volume 9
generation_id backup_time  volume_group_name volume_size type        state    cloud_upload_siz
1             190610091132                   30.00GB     full        complete 2.03GB
2             190613093728                   30.00GB     incremental complete 620.09MB
3             190613120558                   30.00GB     incremental complete 585.59MB
```

Figure 25 Backup generations and status

10. Edit the data and take another backup for another point-in-time copy. List the generations of backup, as shown in Figure 26.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>lsvolumebackupgeneration -volume 9
generation_id backup_time  volume_group_name volume_size type        state    cloud_upload_size
1             190610091132                   30.00GB     full        complete 2.03GB
2             190613093728                   30.00GB     incremental complete 620.09MB
3             190613120558                   30.00GB     incremental complete 585.59MB
4             190614070417                   30.00GB     incremental copying  0.00MB
```

Figure 26 New backup generation

11. List the backups and note the generation count, as shown in Figure 27.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>lsvolumebackup
volume_UID                       volume_id volume_name volume_group_id volume_group_name cloud_account_id cloud_account_name last_backup_time generation_count backup_size
6005076072D861A04000000000000002 0         vdisk0                                        0                cloudaccount0      190606093941     1                0.01MB
6005076072D861A04000000000000006 1         TCT_test                                      0                cloudaccount0      190607072230     2                35.96MB
6005076072D861A0400000000000000F 9         tct_win2                                      0                cloudaccount0      190614070417     4                4.16GB
```

Figure 27 New generation count

12. Confirm the backup is complete, as shown in Figure 28.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>lsvolumebackupgeneration -volume 9
generation_id backup_time  volume_group_name volume_size type        state    cloud_upload_size
1             190610091132                   30.00GB     full        complete 2.03GB
2             190613093728                   30.00GB     incremental complete 620.09MB
3             190613120558                   30.00GB     incremental complete 585.59MB
4             190614070417                   30.00GB     incrementa  complete 975.46MB
```

Figure 28 Backup status

Restoration from different generations

In this section, the restoration of a source volume from different generations or point-in-time copies is performed:

1. Restore the 3rd generation backup to a new volume `tct_win1` using the command **`restorevolume`**, as shown in Figure 29.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>restorevolume -fromuid 6005076072D861A0400000000000000F -generation 3 tct_win1
```

Figure 29 Restore to a new volume

2. Monitor the progress of the restoration, as shown in Figure 30.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>lsvolumerestoreprogress
volume_id volume_name task    status    generation_id backup_time  progress error_sequence_number
9         tct_win2    restore restoring 3             190613120558 0
```

Figure 30 Restore progress

The progress of the restoration can be seen from the status of the disk, as shown in Figure 31.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>lsvdisk 9| grep restore
restore_status restoring
```

Figure 31 Disk status

3. Confirm the restored volume is available to mount, as shown in Figure 32.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>lsvdisk 9| grep restore
restore status available
```

Figure 32 Restore status is available

4. Map the restored volume to the same host after removing the mapping of the source volume from the host, as shown in Figure 33.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>rmvdiskhostmap -host 1 tct_win2
```

Figure 33 Map restored volume

5. Confirm the previous volume is not mounted on the Windows host, as shown in Figure 34.

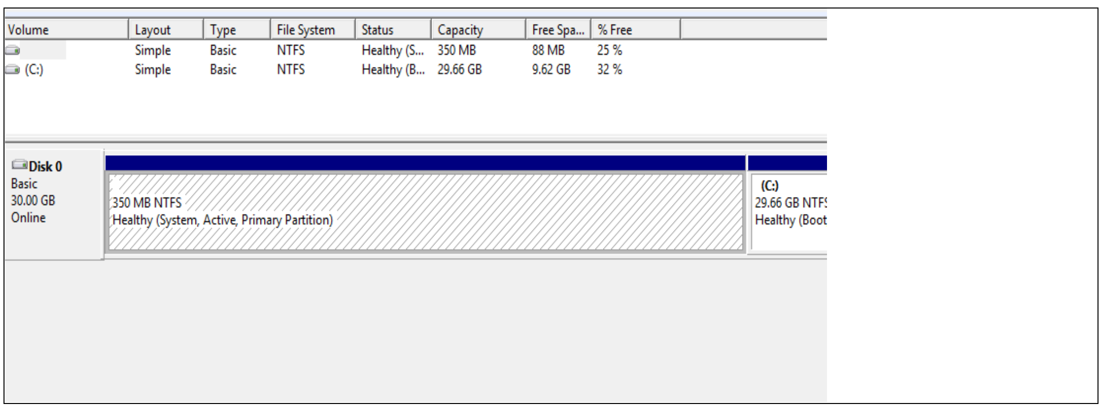

Figure 34 Disk management

6. Map the restored volume to the Windows host, as shown in Figure 35.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>lsvdiskhostmap tct_win1
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>mkvdiskhostmap -host 1 tct_win1
Virtual Disk to Host map, id [0], successfully created
```

Figure 35 Volume mapping

7. Verify the host and LUN mapping, as shown in Figure 36.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>lsvdiskhostmap 8
id name     SCSI_id host_id host_name vdisk_UID                        IO_group_id IO_group_name mapping_type host_cluster_id host_cluster_name protocol
8  tct_win1 0       1       windows   6005076072D861A0400000000000000E 0           io_grp0       private                                        scsi
```

Figure 36 List volume mapping

8. Scan for the disk on the host and assign a drive letter, as shown in Figure 37.

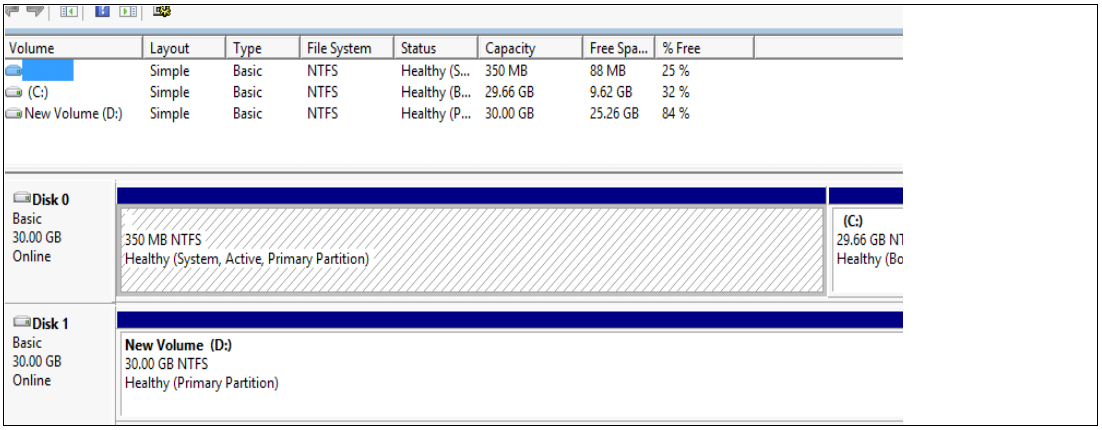

Figure 37 Import volume on Windows host

9. Verify the serial number of the restored volume on the host, as shown in Figure 38.

```
S C:\Program Files\ibm\SDDDSM> .\datapath.exe query device

otal Devices : 2

EV#:   0  DEVICE NAME: Disk1 Part0          TYPE: 2145      POLICY: LEAST I/O AND WEIGHT
ERIAL: 6005076072D861A0400000000000000E      Reserved: No    LUN SIZE: 30.0GB
OST INTERFACE: ISCSI
REFERRED PATH SET: None
=========================================================================================
ath#      Adapter/Hard Disk                   State        Mode        Select     Errors
   0      Scsi Port3 Bus0/Disk1 Part0         OPEN         NORMAL        615          5

evice #1 is removed.
```

Figure 38 Datapath output

10. Verify that the content on drive `D:` is the same as when the generation 3 backup was taken, as shown in Figure 39.

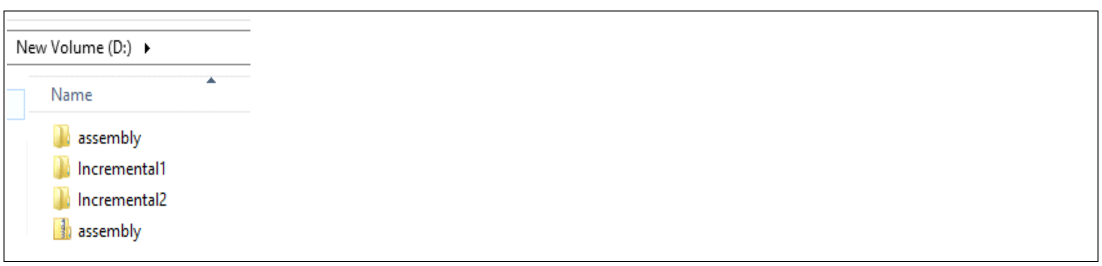

Figure 39 Restored content on volume D

Restoration from different generation to the source volume

This section covers restoring a volume from the 2nd generation backup to the same source volume.

If there are more generations after the 2nd generation, the restore fails. Therefore, the **`deletelatergenerations`** parameter must be passed at the time of restoring a volume from the previous generation. It deletes all later generations, which cannot be restored again.

It is necessary to be diligent when using this parameter, so that the multiple point-in-time copies are not lost:

1. List the generations that are available to be restored, as shown in Figure 40.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>lsvolumebackupgeneration -volume 9
generation_id backup_time  volume_group_name volume_size type        state    cloud_upload_size
1             190610091132                   30.00GB     full        complete 2.03GB
2             190613093728                   30.00GB     incremental complete 620.09MB
3             190613120558                   30.00GB     incremental complete 585.59MB
4             190614070417                   30.00GB     incremental complete 975.46MB
```

Figure 40   List generations of backed up volumes

2. Restore the volume from the 2nd generation of backup while deleting the later generations using the **restorevolume** command, as shown in Figure 41.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>restorevolume -generation 2 -deletelatergenerations tct_win2
```

Figure 41   Restore volume

3. Verify the restore status, as shown in Figure 42.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>lsvdisk 9| grep restore
restore_status restoring
```

Figure 42   Restore status

4. Wait until the restore status is available, as shown in Figure 43.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>lsvdisk 9| grep restore
restore_status available
```

Figure 43   Restore status available

5. Using the **lsvolumebackup** command, verify that the later generations of backup have been deleted, as shown in Figure 44.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>lsvolumebackup
volume_UID                       volume_id volume_name volume_group_id volume_group_name cloud_account_id cloud_account_name last_backup_time generation_count backup_size
6005076072D861A04000000000000002 0         vdisk0                                        0                cloudaccount0      190606093941     1                0.01MB
6005076072D861A04000000000000006 1         TCT_test                                      0                cloudaccount0      190607072230     2                35.96MB
6005076072D861A0400000000000000F 9         tct_win2                                      0                cloudaccount0      190614070417     2                3.21GB
```

Figure 44   List volume backup

6. Verify that the changes are reflected in the output of the **lsvolumebackupgeneration** command, as shown in Figure 45.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>lsvolumebackupgeneration -volume 9
generation_id backup_time  volume_group_name volume_size type        state    cloud_upload_size
1             190610091132                   30.00GB     full        complete 2.03GB
2             190613093728                   30.00GB     incremental complete 620.09MB
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>
```

Figure 45   List generations of backed up volumes

7. Validate the data of the 2nd generation backup by mapping the volume to the same host:

   a. First, remove the mapping of volume `tct_win1`, as shown in Figure 46.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>rmvdiskhostmap -host 1 tct_win1
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>lshostvdiskmap 1
```

Figure 46   Delete volume mapping

b. Map the volume `tct_win2` to the same host, as shown in Figure 47.

```
IBM_Spectrum_Virtualize:sv-cloud-stack-02:superuser>mkvdiskhostmap -host 1 9
Virtual Disk to Host map, id [0], successfully created
```

Figure 47 Volume mapping

c. Verify the serial number of the volume `tct_win2` on the host, as shown in Figure 48.

```
PS C:\Program Files\ibm\SDDDSM> .\datapath.exe query device

Total Devices : 2

Device #0 is removed.

DEV#:   1  DEVICE NAME: Disk2 Part0          TYPE: 2145      POLICY: LEAST I/O AND WEIGHT
SERIAL: 6005076072D861A0400000000000000F     Reserved: No    LUN SIZE: 30.0GB
HOST INTERFACE: ISCSI
PREFERRED PATH SET: None
==========================================================================================
Path#      Adapter/Hard Disk                 State     Mode      Select     Errors
    0      Scsi Port3 Bus0/Disk2 Part0        OPEN     NORMAL     22211          0
PS C:\Program Files\ibm\SDDDSM> _
```

Figure 48 Datapath output

8. Scan the volume on the host and assign a drive letter, as shown in Figure 49.

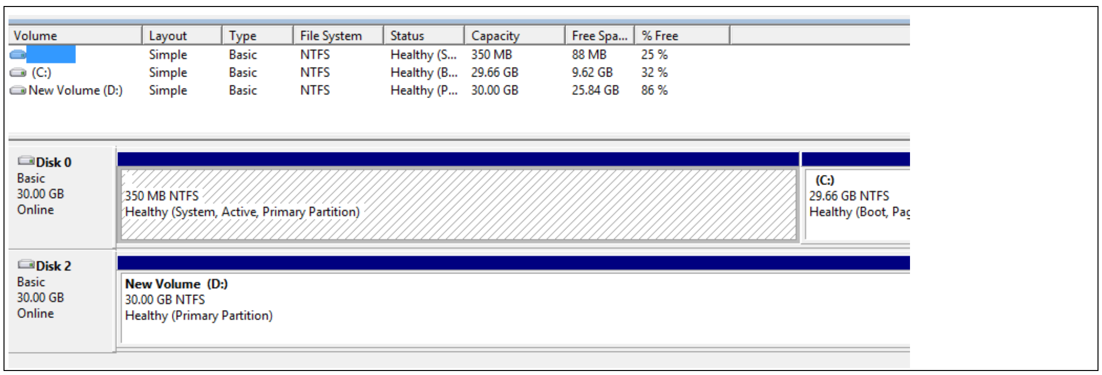

Figure 49 Import volume

9. Verify that the content of 2nd generation backup generation is restored, as shown in Figure 50.

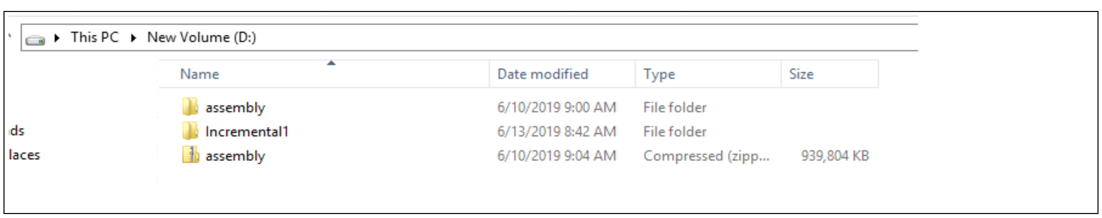

Figure 50 Restored content

Use-case Scenario II

Complete the following steps:

1. Back up the source volume from one instance in IBM Spectrum Virtualize for Public Cloud on AWS running in one availability zone in AWS to Amazon S3.

2. Then, restore the volume to a different IBM Spectrum Virtualize for Public Cloud on AWS running in a different availability zone in AWS, as shown in Figure 51.

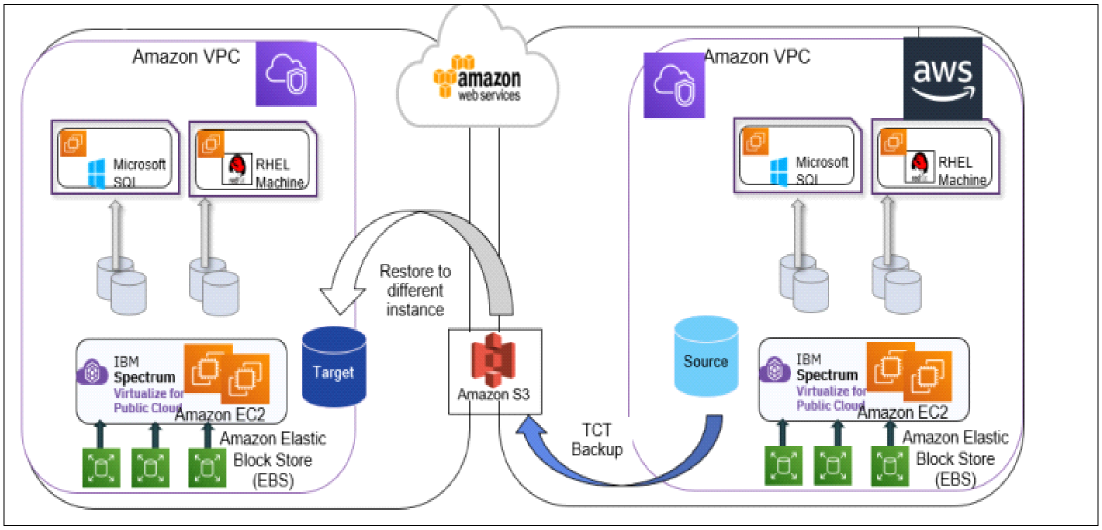

Figure 51 Multiple AWS availability zones architecture

The backup of the source volume is exactly the same as explained in the previous section. For restoring the volume on a different instance of IBM Spectrum Virtualize for Public Cloud on AWS in a different availability zone is detailed in the following steps:

1. To restore the volume on a different instance of IBM spectrum Virtualize Public Cloud on AWS, the first step is to create a cloud account and connect to the same AWS bucket prefix that has the source volume stored, as shown in Figure 52.

Figure 52 Create cloud account

2. After the second instance is connected to the same AWS S3 bucket, run the **lscloudimportcandidate** command. It shows the details of the source IBM Spectrum Virtualize for Public cloud instance that was used to copy the volume to Amazon S3, as shown in Figure 53.

```
IBM_Spectrum_Virtualize:IBM-SVC-Oregan:superuser>lscloudaccountimportcandidate
cloud_account_id cloud_account_name import_system_id import_system_name          backup_volume_count backup_size backup_timestamp
0                cloudaccount0      0000001CA7B18C24 Spectrum-Virtualize-aws-cloud 2                   1.45GB      190919152428
IBM_Spectrum_Virtualize:IBM-SVC-Oregan:superuser>
```

Figure 53 Import cloud candidate

3. Next, run the **chcloudaccounts3** command to import the volume and instance on the target IBM Spectrum Virtualize for Public Cloud on AWS instance, as shown in Figure 54.

```
IBM_Spectrum_Virtualize:IBM-SVC-Oregan:superuser>chcloudaccountawss3 -mode import -importsystem  0000001CA7B18C24 0
IBM_Spectrum_Virtualize:IBM-SVC-Oregan:superuser>lscloudaccount
id name          type  status mode   active_volume_count backup_volume_count import_system_id import_system_name            error_sequence_number
0  cloudaccount0 awss3 online import 0                   2                   0000001CA7B18C24 Spectrum-Virtualize-aws-cloud
IBM_Spectrum_Virtualize:IBM-SVC-Oregan:superuser>lsvolumebackup
volume_UID                       volume_id volume_name volume_group_id volume_group_name cloud_account_id cloud_account_name last_backup_time generation_count backup_size
60050760729EC6309000000000000001           test                                          0                cloudaccount0      190919122748     1                0.01MB
60050760729EC6309000000000000004           Volume_01                                     0                cloudaccount0      190919152428     2                1.45GB
IBM_Spectrum_Virtualize:IBM-SVC-Oregan:superuser>
```

Figure 54   Import instance on target AWS instance

4. After the instance is imported, list the details of the volumes backed up from the source IBM Spectrum Virtualize for Public Cloud on AWS instance.
5. Restore the volume on the target IBM Spectrum Virtualize for Public Cloud on AWS instance. The process is the same as explained in the section “Restore to a new volume” on page 10.

Use-case Scenario III

Complete the following steps:

1. Back up the source Volume from on-premise IBM FlashSystem 9200 to Amazon S3.
2. Restore the volume to IBM Spectrum Virtualize for Public Cloud on AWS running in AWS, as shown in Figure 55.

   The configuration, backup and restore process is same as described in the first two scenarios.

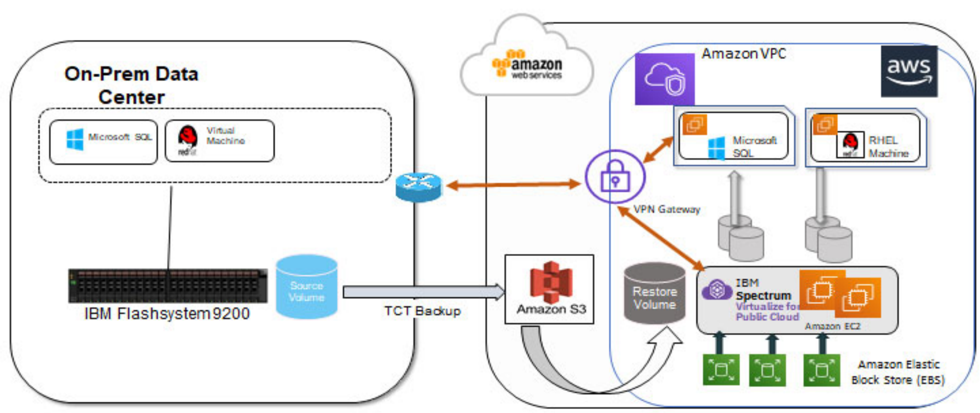

Figure 55   Architecture of backup from IBM FlashSystem 9200 to AWS

Summary

Cyberattacks are likely to remain a significant risk for the foreseeable future. Attacks on organizations can be external and internal. Investing in technology and processes to prevent these cyberattacks is the highest priority for these organizations. Organizations need well-designed procedures and processes to recover from attacks.

The NIST framework provides standards, guidelines, and best practices to manage cybersecurity-related risks. Adoption of the NIST framework, the proper discipline of risk management, and IBM Storage offerings can be used to create and implement recovery plans that ensure the safety of business-critical data.

Using the TCT feature that is available on IBM Spectrum Virtualize for Public Cloud on AWS, the volume is backed up to Amazon S3 and stored in object format, providing the logical air-gapping for the data.

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.

FlashCopy®
IBM®
IBM Cloud™
IBM FlashCore®
IBM FlashSystem®
IBM Spectrum®
Redbooks (logo) ®

The following terms are trademarks of other companies:

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

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.

February 2020

ISBN 0738458449

REDP-5585-00