

IBM® Storage

# **Multicloud Solution for Business Continuity using IBM Spectrum Virtualize for Public Cloud on AWS**

**Version 1 Release 1**



**© Copyright International Business Machines Corporation 2019, 2020.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

# Contents

<b>About this document</b> .....	1
Executive summary .....	1
Scope .....	2
Prerequisites .....	2
Getting started: Multicloud Solution for Business Continuity .....	2
IBM Spectrum Virtualize for Public Cloud on AWS .....	3
IBM FlashSystem 7200 .....	3
End-to-end business continuity and data reuse solution architecture .....	4
Configuring protected on-premises site .....	5
IBM FlashSystem 7200 configuration .....	5
Configuring Disaster Recovery site at AWS Cloud .....	7
Installing IBM Spectrum Virtualize for Public Cloud .....	7
Configuring back-end storage .....	10
AWS Cloud: Configuring site-to-site VPN IPSec tunnel for hybrid cloud connectivity .....	12
Setting up Global Mirror relationship for storage .....	17
Failover of Microsoft SQL database .....	20
Summary .....	22
<b>Notices</b> .....	23
Trademarks .....	24
Terms and conditions for product documentation .....	25
Applicability .....	25
Commercial use .....	25
Rights .....	25
Privacy policy considerations .....	25





## About this document

This document is intended to facilitate the deployment of the Multicloud Solution for Business Continuity by using IBM® Spectrum Virtualize for Public Cloud on AWS. To complete the tasks it describes, you must understand the IBM FlashSystem®, and IBM Spectrum® Virtualize for Public Cloud on AWS Cloud.

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 FlashSystem or IBM FlashSystem storage devices are supported and entitled and where the issues are specific to a blueprint implementation.

## Executive summary

In today’s environment, many organizations are using some form of cloud services, whether private, public, or hybrid cloud. Storage infrastructure is an integral part of these deployments.

Virtualized storage (implemented through software in storage systems or software-defined storage) can dramatically increase operational efficiency, reduce administrative costs, improve data security, and provide cloud-based backup and disaster-recovery (DR) capabilities. It also can add these capabilities to storage you own.

Modern DR solutions can involve leveraging cloud-based resources. However, infrastructure differences between on-premises, private, or public cloud offerings for DR can complicate replicating data to secondary sites. Equivalent storage is needed at each location or significant compute resources might be required to replicate the data at an application or middleware level. Organizations often have a mix of virtualized and non-virtualized applications and associated data that must be replicated to recover from a disaster, which can complicate any implementation.

IBM Spectrum Virtualize and IBM Spectrum Virtualize for Public Cloud use advanced copy services for storage-based replication of data over IP networks in real time. With the ability to achieve a near-zero recovery point objective (RPO) and recovery time objective (RTO), IBM Spectrum Virtualize for Public Cloud can be used in a business continuity solution to address a range of recovery objectives of clients. Any of the data copying features of IBM Spectrum Virtualize on-premises and IBM Spectrum Virtualize for Public Cloud, including synchronous and asynchronous replication capabilities, can be used between the two products.

With IBM Spectrum Virtualize for Public Cloud available on IBM Cloud™ and on AWS Cloud, organizations can have multicloud environments where the data can be replicated between:

- On-premise or Private cloud to Public Cloud (IBM Cloud or AWS Cloud)
- Two Public Cloud (IBM Cloud and AWS Cloud), which makes production and DR sites in public cloud in different public cloud vendors

IBM FlashSystem is built with IBM Spectrum Virtualize, which provides a unique capability to bridge the gap in real-time replication. Consider the following points:

- IBM Spectrum Virtualize operates at the storage not server level so it can offer a consistent approach to replication across a range of virtualized, containerized, and bare-metal systems.
- IBM Spectrum Virtualize can operate within or “above” storage systems so it can offer a consistent approach to replication across a range of different storage systems (more than 440 different systems from IBM and others).

As a result, IBM Spectrum Virtualize enables the use of entirely different storage at production data center and recovery locations, which opens new opportunities for flexibility and cost savings.

## Scope

This blueprint guide provides the following information:

- A solutions architecture and related solution configuration workflow, with the following essential software and hardware components:
  - IBM FlashSystem 7200
  - IBM Spectrum Virtualize for Public Cloud on AWS
- Detailed technical configuration steps for building an end-to-end business continuity and storage as a service solution in hybrid Cloud environment

This technical report does not:

- Provide performance analysis from a user perspective
- Replace any official manuals and documents that are issued by IBM

## Prerequisites

This technical paper assumes that the user has basic knowledge of the following products and technology:

- IBM FlashSystem
- IBM Spectrum Virtualize for Public Cloud on AWS
- AWS Cloud
- IP networking

## Getting started: Multicloud Solution for Business Continuity

This section describes the essential end-to-end business continuity and storage as a service solution building material.

## IBM Spectrum Virtualize for Public Cloud on AWS

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

Designed for public cloud IaaS, IBM Spectrum Virtualize for Public Cloud represents a solution for public cloud implementations and includes technologies that 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 BYOL offering that can be purchased as 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, which offer the ability to transfer data between on-premises data centers by using any IBM Spectrum Virtualize-based appliance and AWS (see Table 1).

*Table 1 IBM Spectrum Virtualize for Public Cloud at a glance*

Storage supported	AWS EBS block storage
Licensing approach	<ul style="list-style-type: none"><li>• Simple, flat cost per managed terabyte</li><li>• Monthly or perpetual licensing</li></ul>
Platform	IBM Spectrum Virtualize for Public Cloud on AWS installed on supported EC2 instance

## IBM FlashSystem 7200

IBM FlashSystem 7200 is designed to deliver flexible, affordable scaling and performance:

- Support of NVMe over Fabrics provides for the highest end-to-end storage performance
- Utilization of IBM FlashCore®-enhanced storage media provides extraordinary flash density and storage capacity while achieving low latency in microseconds

FlashCore Modules (FCM) utilize powerful in line, hardware-accelerated compression technology that provides consistent data compression without performance impact across the full range of workloads. The FCMs are designed to support FIPS 140-2 Level 1 encryption. Built-in flexibility allows you to choose various drive types and supports all three drive types simultaneously within the array:

- FCMs in multiple capacities
- Industry-standard NVMe
- New Storage Class Memory (SCM) drives

Scaling of capacity and performance is dynamic, with always-online high-performance data compression in the FCMs, or with the Data Reduction Pool (DRP) technology using industry standard drives. Effective capacities can range up to four petabytes (PB) in a single 2U enclosure, with the ability to cluster, scale-out, or scale up capacity and performance to 32 PB and eight million input/output operations per second (IOPS).

Each controller contains a hardware compression accelerator based on Intel QuickAssist technology with an available second accelerator. Flexible host interface options include:

- 16 Gbps or 32 Gbps Fibre Channel with FC-NVMe support
- 25 Gbps Ethernet with iSCSI, iWARP, RoCE support, and 10 Gbps iSCSI

- Up to four IBM FlashSystem® 7200 arrays can be clustered and operated as a single system, with 12G, 24G, and 92G SAS expansion enclosures available. It can support up to 760 SAS drives per array controller, 96 NVMe, and 2,944 SAS drives per 4-way clustered system

For more information about IBM FlashSystem specifications, see the following link:

<https://www.ibm.com/us-en/marketplace/flashsystem-7200>

For the purposes of this paper and lab environment, the IBM FlashSystem 7200 is deployed at an on-premises environment. The IBM FlashSystem 7200 combines the performance of flash and the Non-Volatile Memory Express (NVMe) protocol with the reliability and innovation of IBM FlashCore technology and the rich feature set of IBM Spectrum® Virtualize in one powerful new storage platform for your data-driven multi-cloud enterprise.

## End-to-end business continuity and data reuse solution architecture

Figure 1 shows the end-to-end business continuity solution architecture of the IBM Spectrum Virtualize for Public Cloud running in AWS Cloud and IBM FlashSystem 7200.

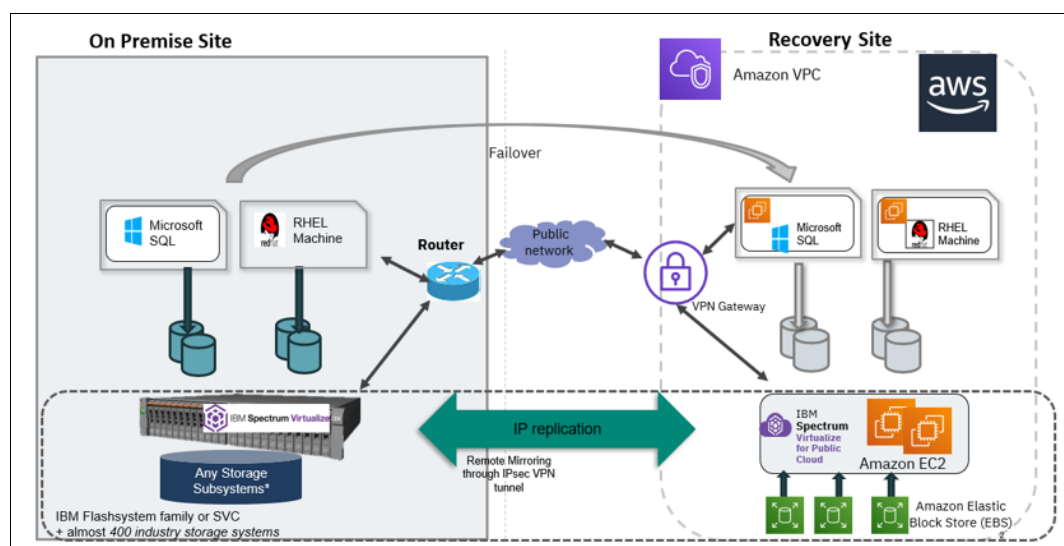


Figure 1 Multicloud business continuity architecture diagram

In this test environment, a Windows machine is installed at a protected on-premises site with SQL database running. An IBM FlashSystem 7200 system is also connected to the Windows server host.

The DR site is in an AWS Cloud and features the IBM Spectrum Virtualize for Public Cloud on AWS that is installed on EC2 servers that are deployed in AWS Cloud. The AWS Cloud site also has its own Windows EC2 instance as a DR server. The Elastic Block Storage (EBS) Volume in AWS cloud is virtualized behind the instance of IBM Spectrum Virtualize for Public Cloud running in AWS.

The production on-premises site and AWS Cloud DR site are connected to each other by using IPsec site-to-site VPN that is tunneling over the public network. IBM Global Mirror with Change Volume is used to replicate data between on-premises IBM FlashSystem FS7200 and IBM Spectrum Virtualize for Public Cloud on AWS Cloud.



This configuration provides the option to choose the flexibility of required RPO. The same solution is extended for data re-use in AWS Cloud. The snapshot copies created with IBM FlashCopy® technology in AWS Cloud can be leveraged for DevOps, analytics and reporting, and backup in AWS Cloud.

## Configuring protected on-premises site

This section describes the benefits, features, and configuration overview of the IBM® FlashSystem 7200.

### IBM FlashSystem 7200 configuration

The IBM FlashSystem 7200 system was used in the lab setup, and the drives were configured as Tier 0 with IBM Easy Tier®.

Complete the following steps to view the system overview page:

1. Log in to the IBM FlashSystem 7200 GUI. Then, log in to the cluster IP address by using a supported web browser and click **System** (see Figure 2).

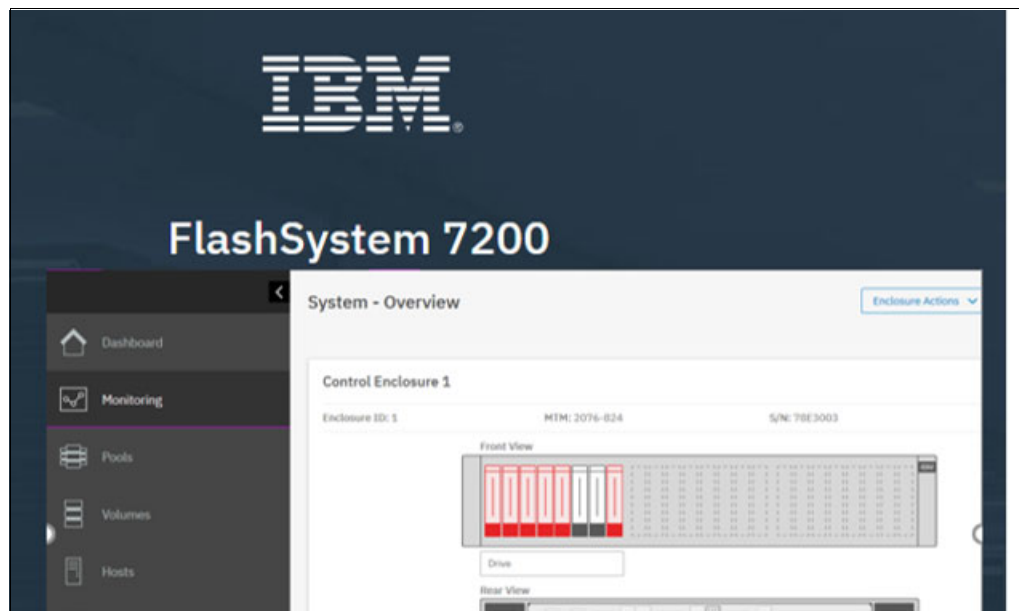


Figure 2 IBM FlashSystem 7200 Login window

2. Click **Pools** → **Internal Storage**, as shown in Figure 3.

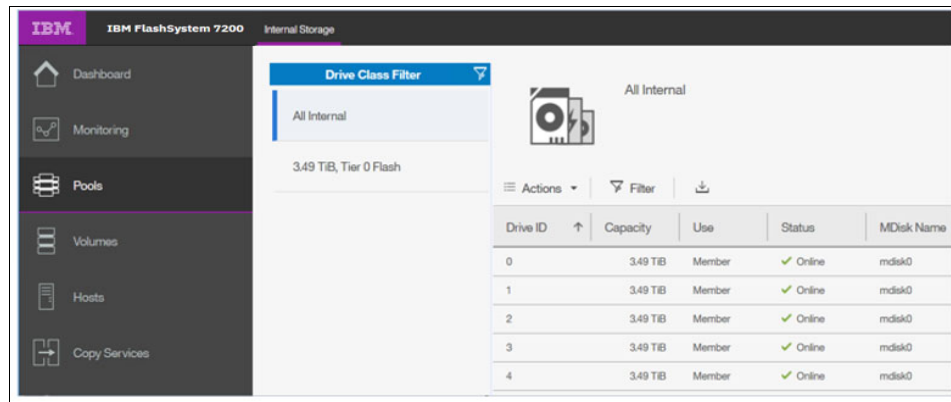


Figure 3 IBM FlashSystem 7200 disk information

3. The next step is to create a pool. Click **Pools** → **Create Pool** and follow the Create Pool wizard. Assign a managed disk (MDisk) to the pool, as shown in Figure 4.

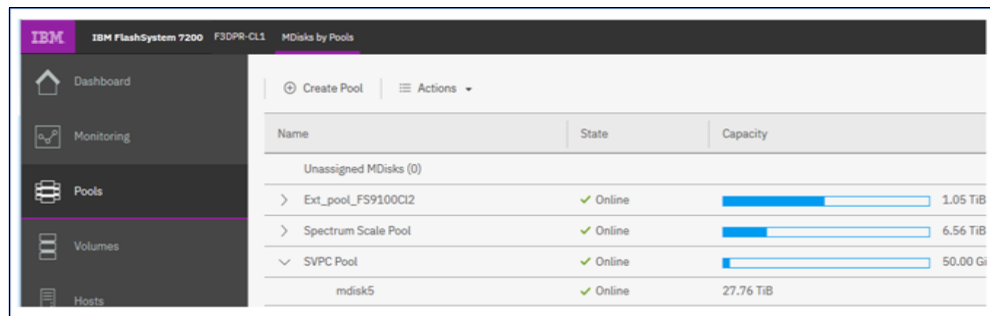


Figure 4 IBM FlashSystem 7200 pool creation

4. After the pool is created, create a VDisk and map the VDisk to the Windows host.  
To create a VDisk, click **Volumes by Pool** → **Create Volumes**, as shown in Figure 5.

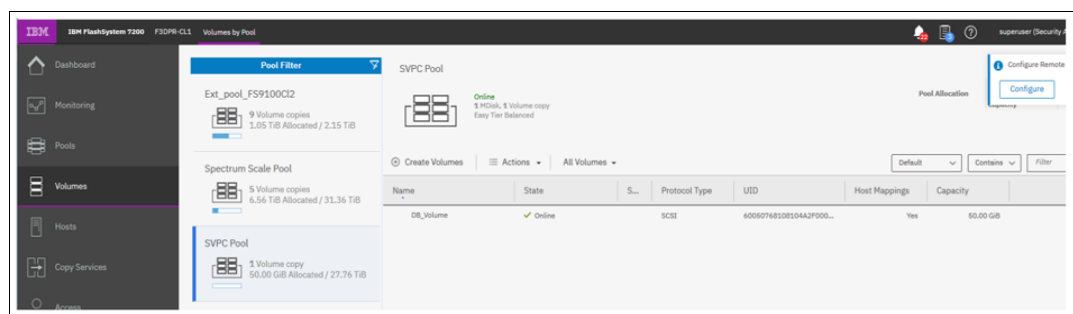


Figure 5 IBM FlashSystem 7200 volume creation

5. In the pool-creation window (see Figure 5), select the pool and provide volume details, such as the capacity and name for the VDisk. Also, select whether you want a thin provisioned volume or a thick volume, and if de-duplication must be enabled or disabled.

6. Then, click **Create and Map**, as shown in Figure 6. Follow the wizard and map the volume to the Windows host at the on-premises site.

The screenshot shows the 'Create Volumes' dialog box with the 'Basic' tab selected. The 'Pool' dropdown is set to 'SVPC Pool'. A progress bar indicates 'Total 27.76 TiB'. Under 'Volume Details', the 'Quantity' is 1, 'Capacity' is 50 GiB, and the 'Name' is 'DB\_Volume'. The 'Capacity savings' dropdown is set to 'None', and the 'Deduplicated' checkbox is unchecked. Below this, there is a link to 'Define another volume'. The 'I/O group' dropdown is set to 'Automatic'. A summary section shows '1 volume' and 'Capacity Savings: None'. At the bottom, there are buttons for 'Cancel', 'Create and Map', and 'Create'.

Figure 6 Create and map volume

## Configuring Disaster Recovery site at AWS Cloud

This section provides essential instruction for ordering, configuring, and installing IBM Spectrum Virtualize for Public Cloud.

In the proof-of-concept solution lab test environment that is described in this section, a two-node IBM Spectrum Virtualize for Public Cloud cluster is configured.

## Installing IBM Spectrum Virtualize for Public Cloud

This section describes the high-level installation steps for IBM Spectrum Virtualize for Public Cloud.

The high-level diagram for installation of IBM Spectrum Virtualize for Public Cloud in AWS is shown in Figure 7.

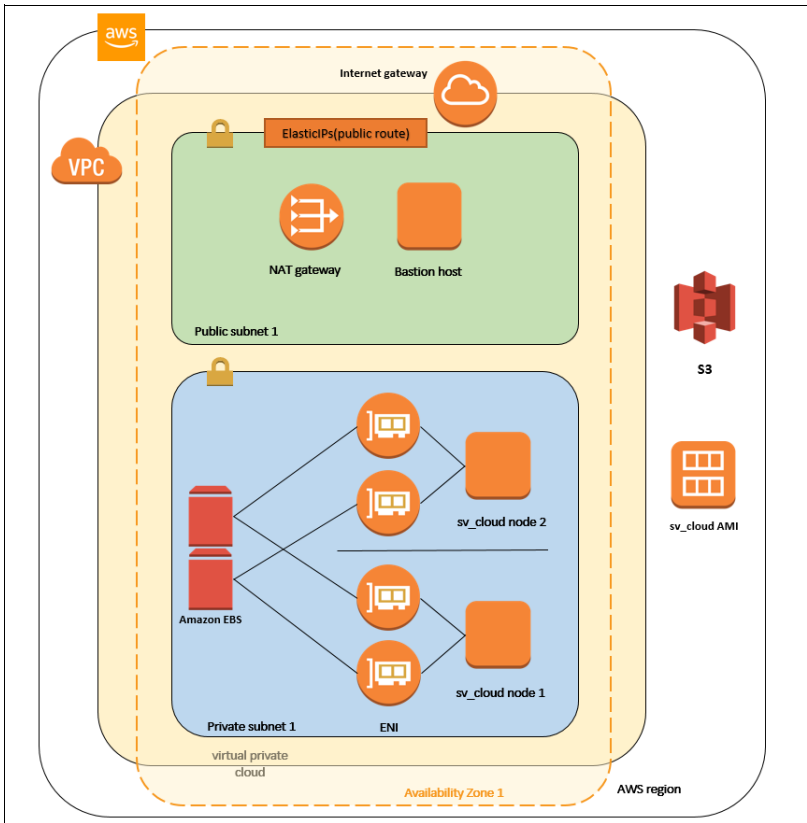


Figure 7 IBM Spectrum Virtualize for Public Cloud on AWS

As shown in Figure 7, the installation can be done in an existing VPC or a new VPC. The installation is done by using the marketplace, as shown in Figure 8.

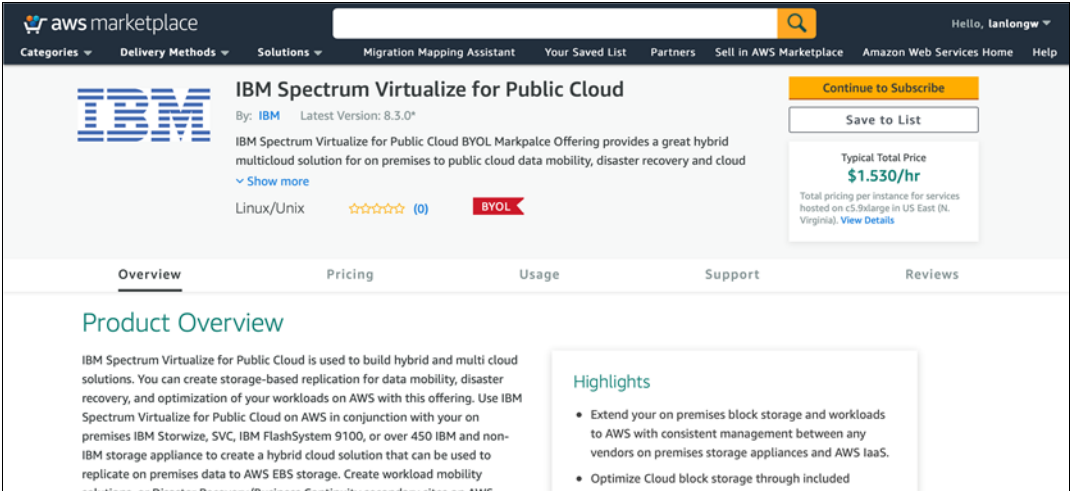


Figure 8 AWS Marketplace for IBM Spectrum Virtualize for Public Cloud

Complete the following steps to install IBM Spectrum Virtualize for Public Cloud:

1. Select the network and availability zone, as shown in Figure 9.

Network Configuration

VPC ID

ID of your existing VPC for deployment

CIDR block of VPC

The CIDR block for the VPC.

Public Subnet 1 ID

ID of public subnet 1 in Availability Zone 1 for the quorum node (e.g., subnet-a0246dcd)

Private Subnet 1 ID

ID of private subnet 1 in Availability Zone 1 for the Workload (e.g., subnet-a0246dcd)

The IP address range that can be used to visit sv\_cloud

The IP address range that can be used to visit sv\_cloud (example for full access: 0.0.0.0/0)

Figure 9 AWS Marketplace for IBM Spectrum Virtualize for Public Cloud

2. Select the EC2 configuration. The details of the three supported types of EC2 configurations are shown in Figure 10.

Model	vCPU	Memory (GiB)	Instance Storage (GiB)	Network Bandwidth (Gbps)	EBS Bandwidth (Mbps)
c5.4xlarge	16	32	EBS-Only	Up to 10	3,500
c5.9xlarge	36	72	EBS-Only	10	7,000
c5.18xlarge	72	144	EBS-Only	25	14,000

Figure 10 EC2 instance types

3. After the cloudformation script is completed, the output is displayed (see Figure 11).

Outputs ▾		
Key	Value	Description
QuorumInstancePublicIP	3.82.248.2	run "ssh -i ssh_key centos@this_I...
Cluster IP	10.10.0.25	Cluster management IP
WebGUI	<a href="https://3.82.248.2:8443">https://3.82.248.2:8443</a>	URL for web GUI
Node 1 Service IP	10.10.0.26	Node management IP
Node 2 Service IP	10.10.0.27	Node management IP

Figure 11 Cloudformation template output

- Log in to the IBM Spectrum Virtualize for Public Cloud cluster to the Web GUI.

Logging in to an IBM Spectrum Virtualize for Public Cloud cluster is almost the same as logging in to a node. Replace the service IP with the cluster IP. Log in to the cluster with a GUI by using your browser, as shown in Figure 12.

With the GUI, you are guided through the steps that help you to complete your cluster installation.

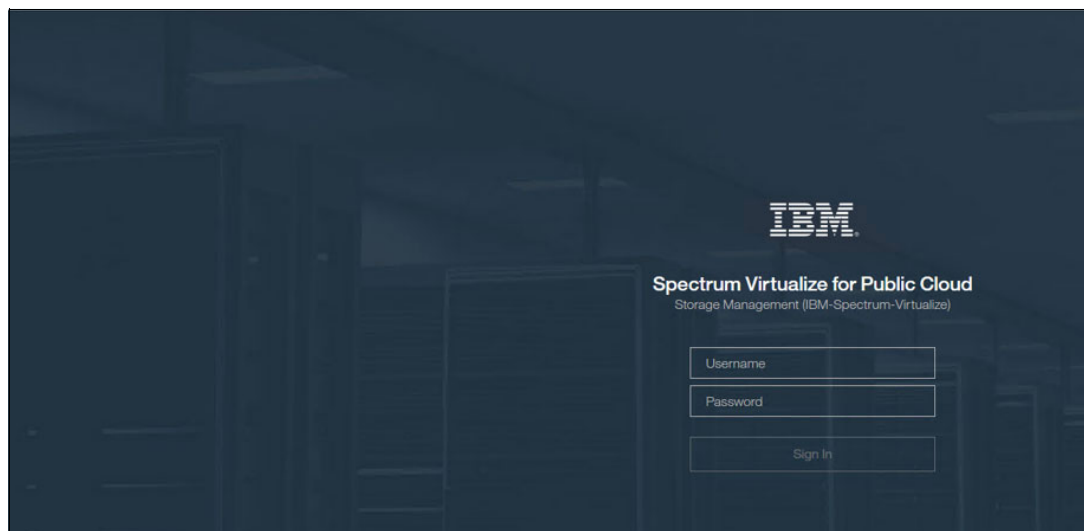


Figure 12 IBM Spectrum Virtualize for Public Cloud on AWS login window

## Configuring back-end storage

IBM Spectrum Virtualize for Public Cloud uses the back-end storage that is provided by AWS Cloud EBS Volume as external MDisk.

To order back-end storage, log in to the [AWS console](#).

Complete the following steps:

- Click **Elastic Block Storage** → **Volumes** → **Create Volumes**, as shown in Figure 13.

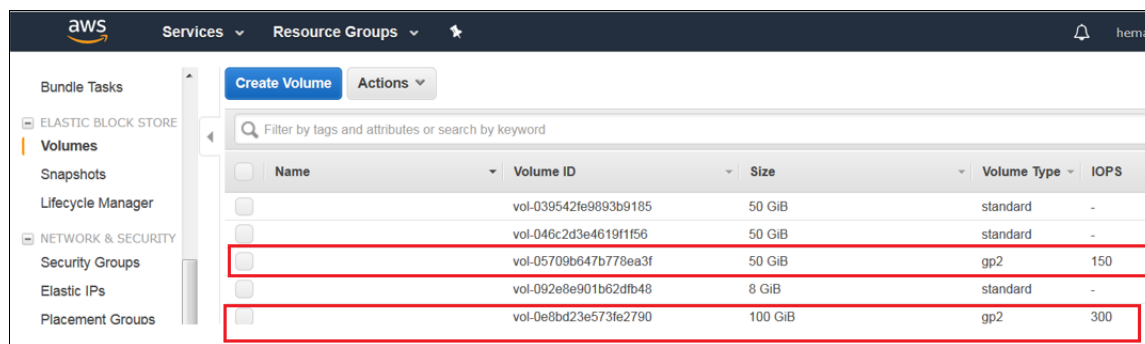


Figure 13 AWS Console with EBS volume details

2. Select the volume type and size of the volume required, as shown in Figure 14.

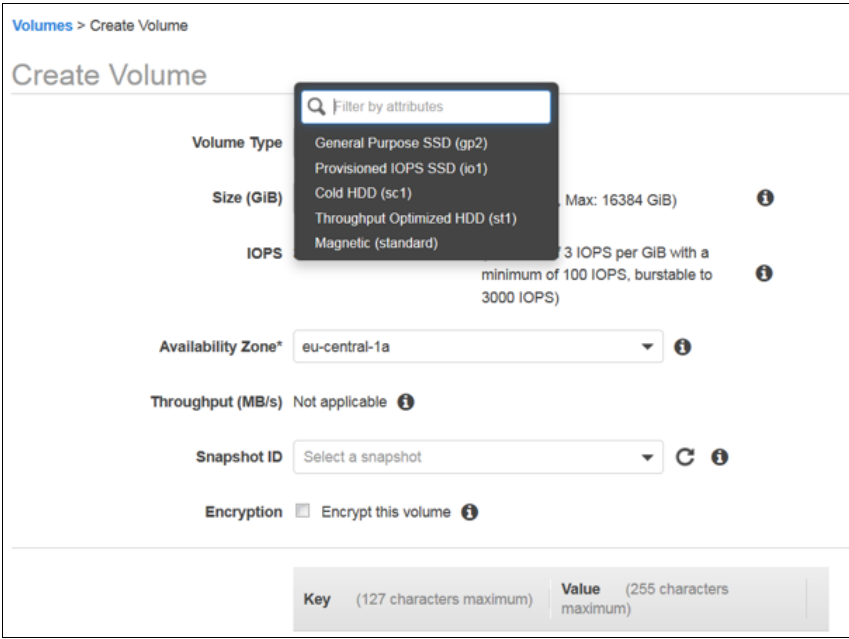


Figure 14 EBS Create volume on AWS Console

The volumes that are marked in RED are the two volumes that were created that are virtualized behind the IBM Spectrum Virtualize for Public Cloud on AWS (see Figure 13 on page 10).

As shown in Figure 15, two Pools are created on IBM Spectrum Virtualize for Public Cloud on AWS. Each Pool as one MDisk that is assigned, which is the EBS external storage that was purchased on AWS Cloud.

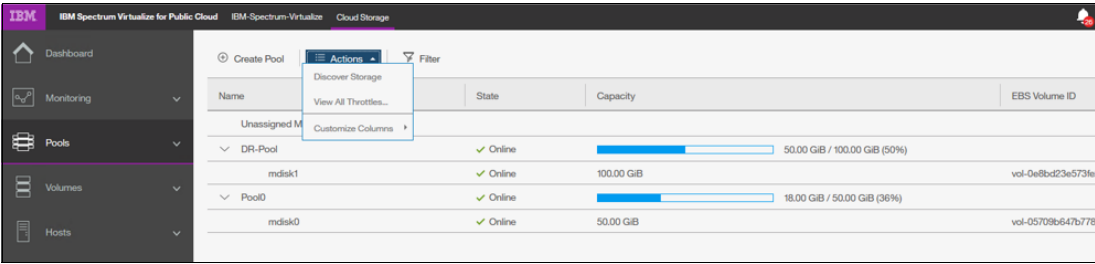


Figure 15 Pool Creation

3. A pool must be created on IBM Spectrum Virtualize for Public Cloud on AWS. Log in to IBM Spectrum Virtualize for Public Cloud by clicking **AWS GUI** → **Pools** → **Create Pool**.
4. After the Pool is created, click **Action** → **Discover storage**, as shown in Figure 15.  
The EBS volume that is purchased on AWS Cloud and is free and not used is shown under the unassigned MDisk. To cross verify that the correct volume is added to the Pool, check the EBS Volume ID (it is the same volume ID as shown on the AWS Cloud console).
5. Add storage in the form of an MDisk to the Pool. Then, create a VDisk and assign the volume for host access by using iSCSI.

# AWS Cloud: Configuring site-to-site VPN IPsec tunnel for hybrid cloud connectivity

This section describes how to configure hybrid cloud connectivity between the AWS Cloud and the on-premises environment. This section also describes the lab setup and the steps to configure the site-to-site IPsec tunnel for communication between AWS Cloud and the on-premises site.

**Note:** Although this section describes the steps for the use case that is shown, the on-premises network configuration, infrastructure, and security policy can vary on a case-by-case basis. This section is intended to give a high-level example only.

The high-level architecture diagram for hybrid cloud connectivity between on-premises and AWS cloud is shown in Figure 16.

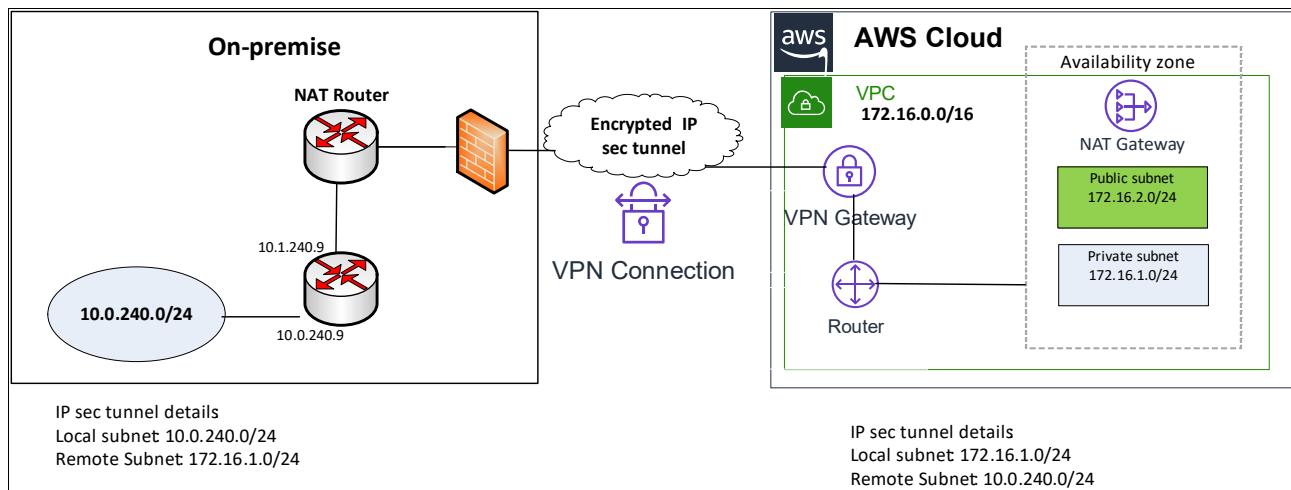


Figure 16 Hybrid cloud network connectivity topology between AWS cloud and on-premises

As shown in Figure 3 on page 6, Virtual Private Cloud (VPC) in AWS is configured with a VPN gateway and router for the CIDR block 172.16.0.0/24. VPN gateway is required for establishing the tunnel between AWS cloud and on-premises infrastructure. It acts as the default router for communication between AWS and on-premises systems. In AWS, all of the compute hosts and IBM Software defined storage systems are configured with IP addresses in the private IP subnet 172.16.1.0/24.

At the on-premises site, a network address translation (NAT) router is used (which is the core router) with a public IP address. That public IP address is NAT'd to a private IP 10.1.210.9. The second router used for lab purposes is a VyOS software gateway at the on-premises site that acts as a default gateway for a private subnet.

The VPN IPsec site-to-site tunnel creates a secure communication network between AWS Cloud infrastructure and on-premises infrastructure. Network communication between the private subnets is controlled by the access control list that is populated at the creation of the VPN IPsec site-to-site tunnel.



## AWS configuration for VPN IPsec tunnel

This section describes the various required steps at the VPC level in AWS cloud for establishing the IPsec tunnel:

1. Create customer gateway.

Log in to the AWS console with the resource provisioning privileges and scroll down to the Virtual Private Network (VPN) section in the pane. Click the customer gateways and enter the required information, as shown in Figure 17.

Customer Gateways > Create Customer Gateway

### Create Customer Gateway

Specify the Internet-routable IP address for your gateway's external interface; the address must be static and may be behind a device performing network address translation (NAT). For dynamic routing, also specify your gateway's Border Gateway Protocol (BGP) Autonomous System Number (ASN); this can be either a public or private ASN (such as those in the 64512-65534 range).

Name

Routing ☐ Dynamic ☒ Static

IP Address\*

\* Required

Cancel Create Customer Gateway

Figure 17 Customer gateway configuration in AWS

2. Create Virtual Private Gateways. Click the Virtual private gateways section in the VPC and configure the required details, as shown in Figure 18.

Virtual Private Gateways > Create Virtual Private Gateway

### Create Virtual Private Gateway

A virtual private gateway is the router on the Amazon side of the VPN tunnel.

Name tag

ASN ☒ Amazon default ASN ☐ Custom ASN

\* Required

Cancel Create Virtual Private Gateway

Figure 18 Virtual private gateway configuration in AWS

3. Attach Virtual private gateway to the VPC, as shown in Figure 19.

Virtual Private Gateways > Attach to VPC

### Attach to VPC

Select the VPC to attach to the virtual private gateway.

Virtual Private Gateway Id

VPC\*

\* Required

Cancel Yes, Attach

Figure 19 Attaching Virtual private gateway to VPC in AWS

4. Create Site-to-Site VPN connection in AWS console, as show in Figure 20. Select the virtual private gateway and customer gateway parameters as created the previous in steps (see Figure 20).

VPN Connections > Create VPN Connection

## Create VPN Connection

Select the virtual private gateway and customer gateway that you would like to connect via a VPN connection. You must have entered the vi

**Name tag**  ⓘ

**Virtual Private Gateway**  ▼ ↺

**Customer Gateway** ☒ Existing ☐ New

**Customer Gateway ID**  ▼ ↺

**Routing Options** ☒ Dynamic (requires BGP) ☐ Static

Figure 20 Creating VPN connection in AWS

Two tunnels are created in the VPC. These tunnels are used for configuration purposes at the other end of the tunnel.

## Configuring the VYOS router at on-premises

Complete the following steps:

1. Enable NAT traversal (NAT-T).

The address of the external interface for your customer gateway must be a static address. In the LAB configuration, we used the VYOS gateway that is behind a device that is performing network address translation (NAT). To ensure that NAT-T can function, you must adjust your firewall rules to unblock UDP port 4500:

```
set security vpn ipsec nat-traversal enable
```

2. Complete the following steps to configure the IPsec tunnel #1:

- a. Use the following Internet Key Exchange (IKE) configuration:

```
set vpn ipsec ike-group AWS lifetime '28800'
set vpn ipsec ike-group AWS proposal 1 dh-group '2'
set vpn ipsec ike-group AWS proposal 1 encryption 'aes128'
set vpn ipsec ike-group AWS proposal 1 hash 'sha1'
set vpn ipsec site-to-site peer <Public IP of tunnel1@AWS> authentication
mode 'pre-shared-secret'
set vpn ipsec site-to-site peer <Public IP of tunnel1@AWS> authentication
pre-shared-secret 'mD2U0cZmKY23sX30u.Iox_OFj_GYcsEd'
set vpn ipsec site-to-site peer <Public IP of tunnel1@AWS> description 'VPC
tunnel 1'
```

```

set vpn ipsec site-to-site peer <Public IP of tunnel1@AWS> ike-group 'AWS'
set vpn ipsec site-to-site peer <Public IP of tunnel1@AWS> local-address
<local Public IP>
set vpn ipsec site-to-site peer <Public IP of tunnel1@AWS> vti bind 'vti0'
set vpn ipsec site-to-site peer <Public IP of tunnel1@AWS> vti esp-group
'AWS'

```

- b. Use the following Encapsulating Security Payload (ESP) configuration:

```

set vpn ipsec ipsec-interfaces interface 'eth0'
set vpn ipsec esp-group AWS compression 'disable'
set vpn ipsec esp-group AWS lifetime '3600'
set vpn ipsec esp-group AWS mode 'tunnel'
set vpn ipsec esp-group AWS pfs 'enable'
set vpn ipsec esp-group AWS proposal 1 encryption 'aes128'
set vpn ipsec esp-group AWS proposal 1 hash 'sha1'

```

- c. Configure the IPSec dead peer detection parameters:

```

set vpn ipsec ike-group AWS dead-peer-detection action 'restart'
set vpn ipsec ike-group AWS dead-peer-detection interval '15'
set vpn ipsec ike-group AWS dead-peer-detection timeout '30'

```

- d. Configure IPSec tunnel parameters:

```

set interfaces vti vti0 address '169.254.40.102/30'
set interfaces vti vti0 description 'VPC tunnel 1'
set interfaces vti vti0 mtu '1436'

```

- e. Configure the Border Gateway Protocol (BGP). BGP is used within the tunnel to exchange prefixes between the Virtual Private Gateway and your Customer Gateway. The Virtual Private Gateway announces the prefix that corresponds to your VPC:

```

set protocols bgp 65000 neighbor <Inside IP CIDR of tunnel1@AWS> remote-as
'64512'
set protocols bgp 65000 neighbor <Inside IP CIDR of tunnel1@AWS>
soft-reconfiguration 'inbound'
set protocols bgp 65000 neighbor <Inside IP CIDR of tunnel1@AWS> timers
holdtime '30'
set protocols bgp 65000 neighbor <Inside IP CIDR of tunnel1@AWS> timers
keepalive '10'

```

Your Customer Gateway can announce a default route (0.0.0.0/0), which can be done with the “network” statement.

```

set protocols bgp 65000 network 0.0.0.0/0

```

To advertise more prefixes to Amazon VPC, replace the 0.0.0.0/0 with the prefix you want to advertise. Ensure that the prefix is present in the routing table of the device with a valid next-hop.

3. Repeat steps 1 and 2 to configure IPSec tunnel #2:

- a. Use the following IKE parameters:

```

set vpn ipsec ike-group AWS lifetime '28800'
set vpn ipsec ike-group AWS proposal 1 dh-group '2'
set vpn ipsec ike-group AWS proposal 1 encryption 'aes128'
set vpn ipsec ike-group AWS proposal 1 hash 'sha1'
set vpn ipsec site-to-site peer <Public IP of tunnel2@ AWS> authentication
mode 'pre-shared-secret'
set vpn ipsec site-to-site peer <Public IP of tunnel2@ AWS> authentication
pre-shared-secret '63VRY4qF.6viqkguHB8w0wQK0FJdsvk1'

```

```

set vpn ipsec site-to-site peer <Public IP of tunnel2@ AWS> description 'VPC
tunnel 2'
set vpn ipsec site-to-site peer <Public IP of tunnel2@ AWS> ike-group 'AWS'
set vpn ipsec site-to-site peer <Public IP of tunnel2@ AWS> local-address
<Local Public IP>
set vpn ipsec site-to-site peer <Public IP of tunnel2@ AWS> vti bind 'vti1'
set vpn ipsec site-to-site peer <Public IP of tunnel2@ AWS> vti esp-group
'AWS'

```

- b. Use the following Tunnel Interface configuration:

```

set interfaces vti vti1 address '169.254.42.106/30'
set interfaces vti vti1 description 'VPC tunnel 2'
set interfaces vti vti1 mtu '1436'

```

- c. Use the following BGP configuration:

```

set protocols bgp 65000 neighbor <Inside IP CIDR of tunnel2@AWS> remote-as
'64512'
set protocols bgp 65000 neighbor <Inside IP CIDR of tunnel2@AWS>
soft-reconfiguration 'inbound'
set protocols bgp 65000 neighbor <Inside IP CIDR of tunnel2@AWS> timers
holdtime '30'
set protocols bgp 65000 neighbor <Inside IP CIDR of tunnel2@AWS> timers
keepalive '10'

```

4. Check the tunnel status. On the VyOS router, VPN IPsec tunnel status can be verified by running the **show vpn ipsec sa** command (see Figure 21).

```

vyos@VyOS-class24:~$ show vpn ipsec sa
Peer ID / IP                               Local ID / IP
-----
                               10.1.240.9

Description: VPC tunnel 3

Tunnel  State  Bytes Out/In  Encrypt  Hash  NAT-T  A-Time  L-Time  Proto
-----
vti      up        2.3K/2.8K    aes128   sha1   yes    1208    3600    all

Peer ID / IP                               Local ID / IP
-----
                               10.1.240.9

Description: VPC tunnel 4

Tunnel  State  Bytes Out/In  Encrypt  Hash  NAT-T  A-Time  L-Time  Proto
-----
vti      up      221.3M/125.8M  aes128   sha1   yes    3462    3600    all

```

Figure 21 Checking VPN IPsec status at VYOS router

5. You can also check the status from the VPC AWS console by clicking **Site-to-site connections** → **Tunnel Details** (see Figure 22).

Name	VPN ID	State	Virtual Private Gateway	Transit Gateway
Kelsterbach	vpn-00a71e802a290e3b2	available	vgw-071860227e2c36fb6   cactus private gat...	-

VPN Connection: vpn-00a71e802a290e3b2

Details Tunnel Details Tags

Outside IP Address	Inside IP CIDR	Status	Status Last Changed	Details
		UP	April 3, 2019 at 11:06:18 AM UTC+5:30	1 BGP ROUTES
		UP	April 3, 2019 at 11:06:29 AM UTC+5:30	1 BGP ROUTES

Figure 22 Checking VPN IPsec status from AWS console

## Setting up Global Mirror relationship for storage

This section describes the creating the Global Mirror IP replication relationship between on-premises IBM FlashSystem 7200 and IBM Spectrum Virtualize for Public Cloud cluster running in AWS.

After the VPN tunnel is up and running, the machines in the on-premises private subnet 10.0.240.0/26 can connect and ping any machine running in the AWS Cloud private VPC subnet 172.16.1.0/26.

Complete the following steps to create Global Mirror IP replication between on-premises IBM FlashSystem 7200 and IBM Spectrum Virtualize for Public Cloud cluster:

1. On the IBM Spectrum Virtualize for Public Cloud GUI, click **Settings** → **Network** → **Ethernet** and select the **IP port**. Then, right-click and select **Modify Remote Copy**, as shown in Figure 23.

Name	Port	State	IP	Speed	Host Attach	Remote Copy	Storage Port IPv4
~io_grp0							
node1	1	✓ Configured	172.16.1.220	10Gb/s	Yes		Enabled
node2	1	✓ Configured	172.16.1.63	10Gb/s	Yes		Enabled
node1		✓ Configured	172.16.1.158	10Gb/s	Yes	Copy Group 1	Enabled
node2		✓ Configured	172.16.1.141	10Gb/s	Yes	Copy Group 1	Enabled

Figure 23 Setting ports for remote replication

2. Enable IPv4 for the remote group and select the group, as shown in Figure 24.

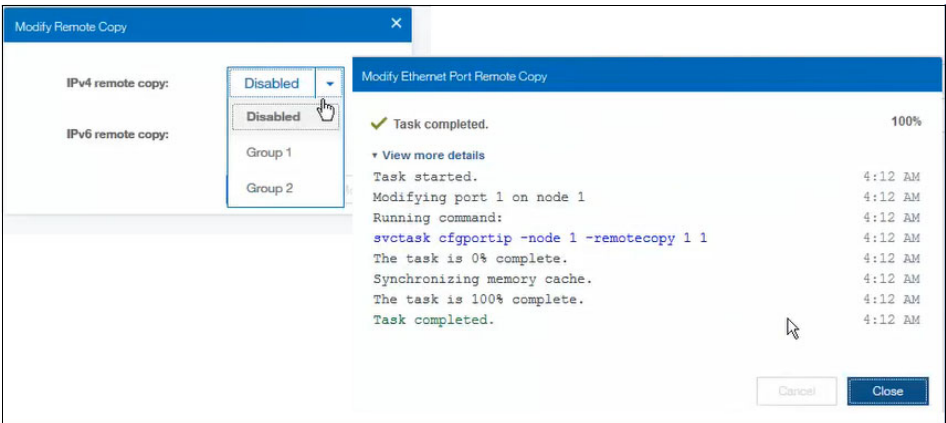


Figure 24 Setting remote copy port

3. After you enabled the port (see Figure 24), set the remote copy group, as shown in Figure 25.

Ethernet Ports

The Ethernet ports can be used for iSCSI connections, host attachment, and remote copy.

Actions Filter

Name	Port	State	IP	Speed	Host Attach	Remote Copy
io_grp0						
node1	1	✓ Configured	172.16.1.220	10Gb/s	Yes	
node2	1	✓ Configured	172.16.1.63	10Gb/s	Yes	
node1	2	✓ Configured	172.16.1.158	10Gb/s	Yes	Copy Group 1
node2	2	✓ Configured	172.16.1.141	10Gb/s	Yes	Copy Group 1

Figure 25 Setting remote copy group

4. Repeat steps 1 and 2 for the on-premises IBM FlashSystem 7200 system.
5. Select **Copy services** → **Partnership** → **Create Partnership**, as shown in Figure 26.


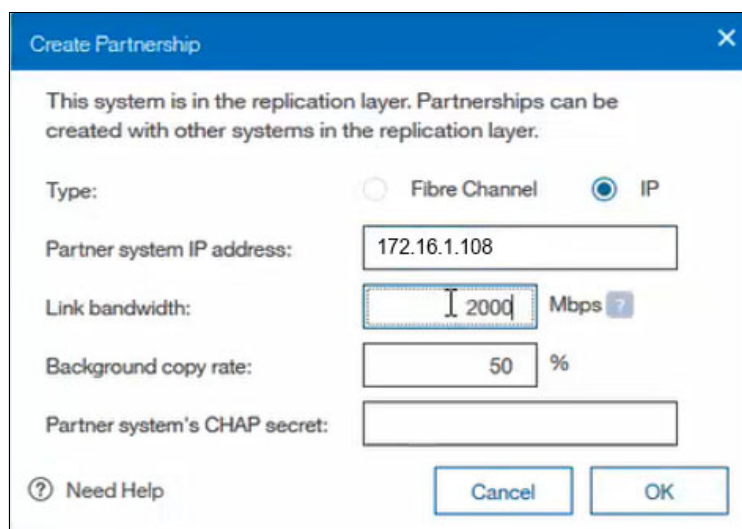


Figure 26 Create partnership

6. Enter the cluster IP address of the remote system, as shown in Figure 27.



**Create Partnership**

This system is in the replication layer. Partnerships can be created with other systems in the replication layer.

Type: ☐ Fibre Channel ☒ IP

Partner system IP address:

Link bandwidth:  Mbps

Background copy rate:  %

Partner system's CHAP secret:

[? Need Help](#)

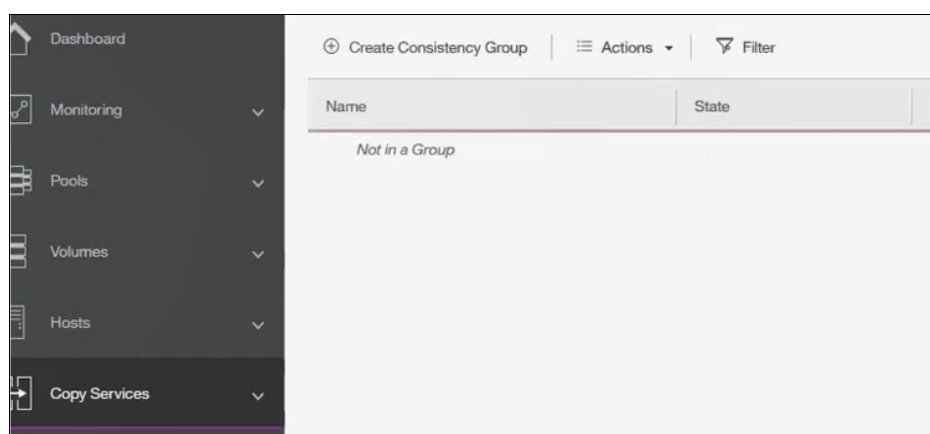
Figure 27 Setting IP address for remote replication

The Global Mirror IP replication is set up relationship between the on-premises IBM FlashSystem 7200 system and the IBM Spectrum Virtualize for Public Cloud cluster running in AWS Cloud, as shown in Figure 28.

Name	Location	State	Type	IP Address
F3DPR-CL1	Local	-	-	-
IBM-Spectrum-Virtualize	Remote	✓ Fully Configured	IPv4	172.16.1.108

Figure 28 IP replication

7. After the Global Mirror relationship is created, set the consistency group and replicate the volumes for DR between the on-premises site and AWS Cloud, as shown in Figure 29.



**Create Consistency Group**

Actions Filter

Name	State
Not in a Group	

Figure 29 Create consistency group

8. Click **Copy Services** → **Remote Copy** → **Create Consistency Group**.

9. Select the type of copy, as shown in Figure 30.

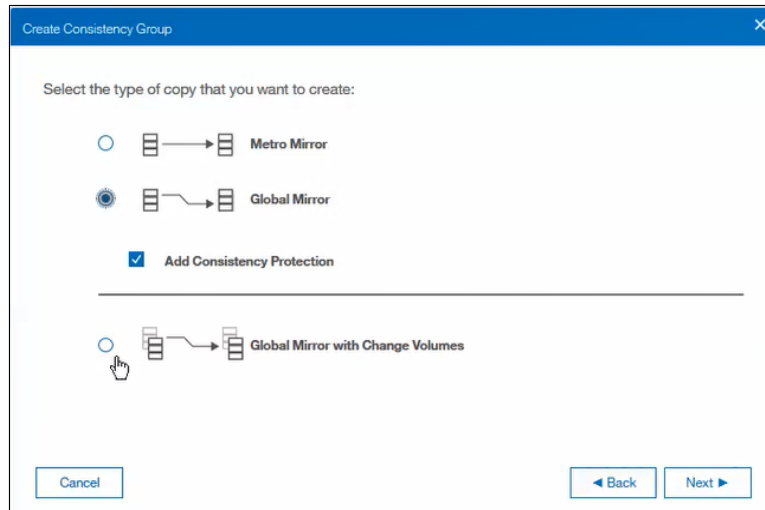


Figure 30 Select replication type

10. Add local and remote volume for the copy and begin copying, as shown in Figure 31.

Name	State	Master Volume	Auxiliary Volume
Not in a Group			
DR-SQL	State: Consistent Synchronized	Master System: F3DPR-CL1	Auxiliary System: IBM-Spectrum-Vii
rcrel0	Consistent Synchronized	DB_Volume	DR_DB_Volume

Figure 31 Local and remote copy volume

## Failover of Microsoft SQL database

In the lab environment, the Microsoft SQL Server Express database was running at on-premises where the underline storage is from IBM FlashSystem 7200.

The volume that consists of the database is replicated by using Global mirror/Global mirror with changed volumes to IBM Spectrum Virtualize for Public Cloud running in AWS.

The D drive is the volume that includes TESTdb running, as shown in Figure 32.

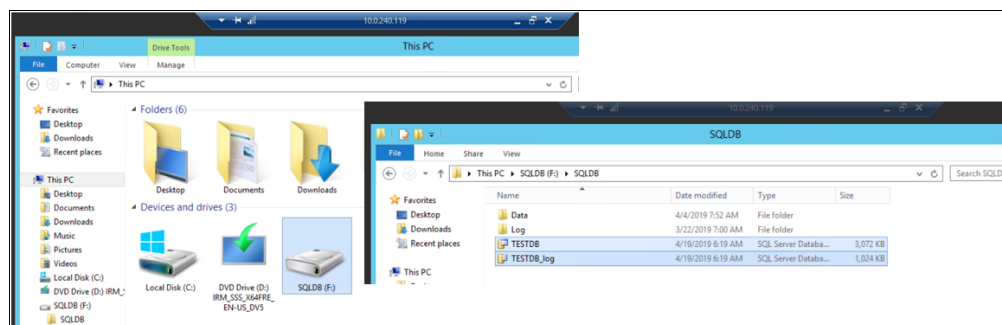


Figure 32 Primary on-premises server with TESTdb



The database has 1000 rows in the table at on-premises, as shown in Figure 33.

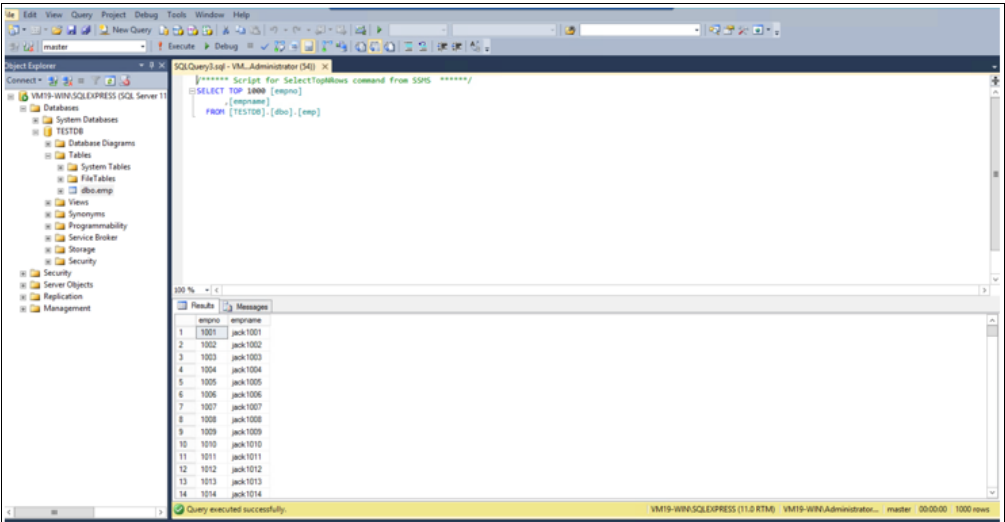


Figure 33 On-premises server with TESTdb

During failover of this volume from on-premises to IBM Spectrum Virtualize for Public Cloud on AWS, stop the Global Mirror replication and make the secondary volume read/write enabled, as shown in Figure 34.




Figure 34 Replication stopped due to disaster at on-premises site

The target volume in IBM spectrum Virtualize for Public Cloud on AWS is attached to a DR server/EC2 windows instance with SQL running in AWS Cloud, as shown in Figure 35. The volume mapping to the EC2 instances is through iSCSI.

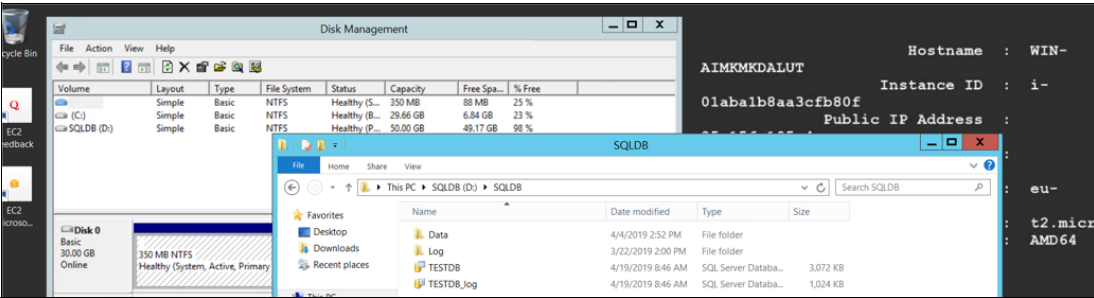


Figure 35 Replication stopped due to disaster at on-premises site

The Test DB is attached to the SQL instance running in EC2 instance in AWS cloud, as shown in Figure 36.

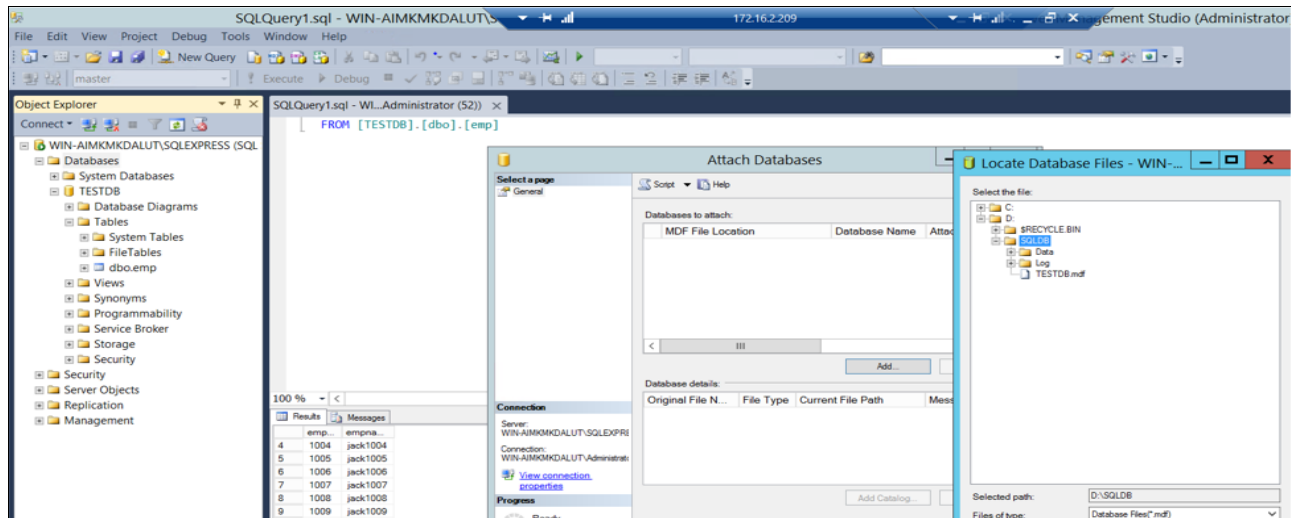


Figure 36 Attach a TESTDB at EC2 instance in AWS Cloud post DR

It is also possible to failback the SQL instance to the original on-premises site using by reversing the Global Mirror replication from AWS Cloud to On-premise and failback the data to on-premises location.

## Summary

With IBM FlashSystem and IBM Spectrum Virtualize for Public Cloud, customers can optimize their heterogeneous storage infrastructure and plan for hybrid-cloud DR between on-premises and AWS Cloud EBS storage.

# 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) ®  
Easy Tier®  
FlashCopy®

IBM®  
IBM Cloud™  
IBM FlashCore®

IBM FlashSystem®  
IBM Spectrum®

The following terms are trademarks of other companies:

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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>.







© Copyright IBM Corporation

February 2020

US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule  
Contract with IBM Corp.



Please recycle

ISBN 0738458538

REDP-5545-01