IBM FlashSystem 9100 Multi-Cloud Solution for Business Continuity and Data Reuse

Version 1 Release 1

IBM

Contents

About this document

This document is intended to facilitate the deployment of the IBM® FlashSystem® 9100 Multi-Cloud Solution for Business Continuity and Data Reuse. To complete the tasks it describes, you must understand IBM FlashSystem 9100, IBM Spectrum Virtualize™ for Public Cloud, IBM Spectrum™ Copy Data Management and VMware Site Recovery Manager.

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 Storwize® and/or IBM FlashSystem storage devices are supported and entitled and where the issues are not specific to a blueprint implementation.

Report an IBM FlashSystem 9100 Multi-Cloud Solution for Business Continuity and Data Reuse software problem to IBM Support

The solution described in this Blueprint involves the combination of IBM FlashSystem 9100 system(s) and one or more of the software offerings associated with IBM FlashSystem 9100 system(s). Technical support for the solution and offering(s) referenced in this document is limited to situations where support is entitled for those offerings. Technical support is provided for:

- The IBM software offerings, provided as a base option of the IBM FlashSystem 9100 system(s), associated with this solution:
  - IBM Spectrum Copy Data Management Multi-Cloud starter for IBM FlashSystem 9100
  - IBM Spectrum Virtualize for Public Cloud Multi-Cloud starter for IBM FlashSystem 9100
- The underlying IBM software components provided when you order the IBM FlashSystem 9100 Multi-Cloud Solution for Business Continuity and Data Reuse offering/option of the IBM FlashSystem 9100
- Assistance in using this Blueprint to combine these software offerings to create the IBM FlashSystem 9100 Multi-Cloud Solution for Business Continuity and Data Reuse
- Assistance in setting up the IBM Cloud™ environment to host these cloud-based software offerings

Note: Support for the third-party software described in this Blueprint is not provided for the multi-cloud solutions discussed here. Please reference the information provided with those offerings for more information on how to obtain support.

The IBM FlashSystem 9100 Knowledge Center contains more detailed planning and troubleshooting information. Please reference the IBM FlashSystem 9100 Knowledge Center for details on planning (i.e., how to obtain the software packages) and troubleshooting (i.e., specifics on how to correctly request assistance from IBM).

How to obtain software

Client access to the software packages is available only through the Entitled System Support (ESS) system (www.ibm.com/eserver/ess). It is important to note that this is a different repository than the repositories used for generally licensed software (such as IBM Passport Advantage® or Fix Central).

- Client access to the software packages is based on the same entitlement model (system entitlement) described in the previous section.
- Access to the software packages is entitled using the same process as that described in the previous section for problem submission, which requires identifying the associated IBM FlashSystem 9100 hardware, either by system serial # or the IBM Customer Number (ICN) associated with the order.

You can download all the packages associated with an offering, or specific packages. Use the following tables to understand and select the appropriate software packages to download:

Ordered offering name (PID)	Related software packages	
	For packages beginning:	Description
IBM Spectrum Copy Data Management Multi-Cloud starter for IBM FlashSystem 9100	“IBM Spectrum CDM”	Packages containing the software related to IBM Spectrum Copy Data Management
IBM Spectrum Virtualize for Public Cloud Multi-Cloud starter for IBM FlashSystem 9100	“IBM SV Cloud”	Packages containing the software related to IBM Spectrum Virtualize for Public Cloud
IBM FlashSystem 9100 Multi-Cloud Solution for Business Continuity and Data Reuse	“IBM Spectrum CDM”	Packages containing the software related to IBM Spectrum Copy Data Management
	“IBM SV Cloud”	Packages containing the software related to IBM Spectrum Virtualize for Public Cloud

Executive summary

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

In a recent study by Technology Business Research (TBR), organizations ranked backup and disaster recovery as their two most-deployed hybrid-cloud data-storage infrastructure applications.[1] And in the same study, 56 percent of respondents said that they expected, within a year, to increase deployment of these cloud services.[1]

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. And it can add these capabilities to storage you already own.

Modern DR solutions may 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. Either equivalent storage is needed at each location, or significant compute resources may 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 that clients may have. Any of the data copying features of IBM Spectrum Virtualize on-premises and IBM Spectrum Virtualize for Public Cloud, including both synchronous and asynchronous replication capabilities, can be used between the two products.

IBM FlashSystem 9100 is built with IBM Spectrum Virtualize, which provides a unique capability to bridge the gap in real-time replication.

- 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, opening up new opportunities for flexibility and cost savings.

Scope

This blueprint guide provides:

- A solutions architecture and related solution configuration workflows, with the following essential software and hardware components:
  - IBM FlashSystem 9100
  - IBM Spectrum Virtualize for Public Cloud
  - IBM Spectrum Copy Data Management
  - VMware Site Recovery Manager
- Detailed technical configuration steps for building an end-to-end business continuity and data re-use solution

This technical report does not:

- Provide performance analysis from a user perspective
- Replace any official manuals and documents issued by IBM
- Explain installation and configuration of VMware vSphere

Prerequisites

This technical paper assumes the following prerequisites:

- Basic knowledge of IBM FlashSystem 9100
- Basic knowledge of IBM Spectrum Copy Data Management
- Basic knowledge of VMware vSphere version 6.0 or later
- Basic knowledge of IP networking

Getting started: IBM FlashSystem 9100 Multi-Cloud Solution for Business Continuity and Data Reuse

This section describes the essential end-to-end business continuity and data reuse solution building materials in detail.

IBM Spectrum Virtualize for Public Cloud

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

Designed for software-defined environments, IBM Spectrum Virtualize for Public Cloud represents a solution for public cloud implementations and includes technologies that both complement and enhance public cloud offering capabilities.

IBM Spectrum Virtualize for Public Cloud provides for the deployment of IBM Spectrum Virtualize software in public clouds, starting with IBM Cloud. This new offering provides a monthly license to deploy and use IBM Spectrum Virtualize for Public Cloud in IBM Cloud to enable hybrid cloud solutions, offering the ability to transfer data between on-premises data centers using any IBM Spectrum Virtualize-based appliance and IBM Cloud.

IBM Spectrum Virtualize for Public Cloud at a glance:

Storage supported	IBM Cloud Performance and Endurance block storage
Licensing approach	Simple, flat cost per managed terabyte Monthly licensing
Platform	IBM Cloud bare-metal server infrastructure

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

IBM FlashSystem 9100

IBM FlashSystem 9100 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.

The NVMe-optimized all-flash arrays are offered in two basic models—IBM FlashSystem 9110 and IBM FlashSystem 9150. The compact 2U enclosures feature dual array controllers, dual power supplies, redundant cooling and full hot-swap capabilities. Both models have two Intel Skylake CPUs per array controller with the IBM FlashSystem 9110 model offering eight cores per CPU, while the 9150 model comes with 14 cores per CPU for higher throughput and performance. Essentially, two rack units of space can provide the performance and efficiency of more than a terabyte of memory and up to two petabytes of effective storage—all moving at NVMe speeds to tackle even the most demanding real-time analytics or artificial intelligence applications and workloads.

IBM FlashSystem 9100 also provides the software-defined, modern data protection and multi-cloud capabilities of several members of the IBM Spectrum Storage™ family. Chief among these is IBM Spectrum Virtualize, the system foundation that provides a broad set of enterprise-class data services—such as dynamic tiering, replication, IBM FlashCopy® management, data mobility, transparent cloud tiering and high-performance data-at-rest encryption, among many others. The arrays also leverage innovative new data reduction pools (DRP) that incorporate deduplication and hardware-accelerated compression technology, as well as SCSI UNMAP support and all the thin-provisioning and data-efficiency features you expect from IBM Spectrum Virtualize-based storage to potentially reduce your capital and operating expenses. Additionally, IBM FlashSystem 9100 solution features virtualization capabilities, which can be used to virtualize more than 440 IBM and non-IBM heterogeneous storage systems.

IBM Spectrum Copy Data Management

IBM Spectrum Copy Data Management makes copies of data available when and where needed, without creating unnecessary copies or leaving unused copies on valuable storage. It catalogs copy data from across your local and hybrid cloud and off-site cloud infrastructure, identifies duplicates, and compares copy requests to existing copies. Data consumers can use the self-service portal to create the copies they need, enabling business agility. Copy processes and work flows are automated to ensure consistency and reduce complexity.

VMware infrastructure

VMware ESXi is a hypervisor that runs on the system hardware directly without the need for any other software. It provides the necessary hypervisor functions to host multiple guest operating systems, such as Microsoft Windows or Linux, on the physical server.

VMware vCenter

VMware vCenter is the management software suite that is used to manage the virtual machines (VMs) within an ESXi host. When you allocate resources—such as memory, storage, networking or processors—to a VM, a vCenter server manages the way these resources are allocated and maintained. vCenter can manage a single ESXi host or a cluster of hosts. vCenter includes several features enabling mobility of VMs between ESXi hosts and storage. These features can add to the availability of the VMs that are run in a cluster

VMware Site Recovery Manager

VMware Site Recovery Manager is a DR management product from VMware that provides automated failover and DR testing. Site Recovery Manager speeds DR and can prioritize the recovery processes, specifying the order in which VMs are restarted. Site Recovery Manager automates the process of synchronizing recovery data between the primary and recovery data center sites.

IBM Storwize Family Storage Replication Adapter

IBM Storwize Family Storage Replication Adapter (SRA) is software that integrates with the Site Recovery Manager solution. The Storwize Family SRA extends Site Recovery Manager capabilities to simplify the automation of storage-system failover from the protected Site Recovery Manager site to a recovery Site Recovery Manager site.

End-to-end business continuity and data re-use solution architecture

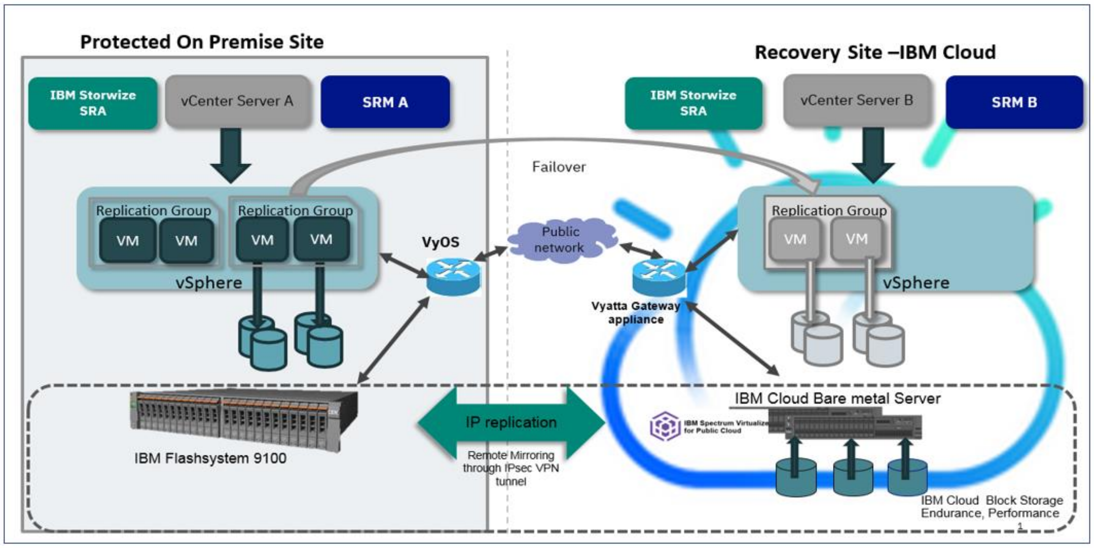

Figure 1: High-level architecture diagram

Figure 1 illustrates the end-to-end business continuity solution architecture of the IBM Spectrum Virtualize for Public Cloud and IBM FlashSystem 9100 for VMware environments using integration with Site Recovery Manager.

In this test environment, VMware vSphere 6.5 is installed at a protected on-premises site; there is also an IBM FlashSystem 9100 system connected to the VMware vSphere server host. The on-premises VMware vCenter Server 6.5 instance is equipped with the Site Recovery Manager plug-in and Storwize Family SRA.

The recovery site is in IBM Cloud and features VMware vSphere 6.5 installed on a bare-metal server purchased in IBM Cloud along with IBM Spectrum Virtualize for Public Cloud installed on bare-metal servers deployed in IBM Cloud. The IBM Cloud site also has its own instance of vCenter Server 6.5 and has Site Recovery Manager and Storwize Family SRA installed.

The protected on-premises site and IBM Cloud recovery site are connected to each other using IPsec site-to-site VPN tunneling over the public network.

Refer to Table 2 for a list of VMs configured for the solution lab test environment.

Site/Location	**Purpose of VM**	**Number of VMs**	**Operating system**
Protected on-premises site	vCenter appliance	1	vCenter appliance
Protected on-premises site	Site Recovery Manager plug-in and Storwize Family SRA	1	Microsoft Windows Server 2008
Protected on-premises site	SQL database Windows test VM for failover /failback	1	Microsoft Windows Server 2012 and Microsoft Windows SQL Server Express
Recovery IBM Cloud site	Site Recovery Manager plug-in and Storwize Family SRA	1	Microsoft Windows Server 2008
Recovery IBM Cloud site	vCenter appliance	1	vCenter appliance

Table 2: Lab solution hardware details and VM list

DR automation for a non-VMware environment in a hybrid cloud scenario can be achieved by using IBM Spectrum Copy Data Management. The high-level architecture diagram shown in Figure 2 illustrates DR automation using an IBM Spectrum Copy Data Management workflow to manage Global Mirror with change volume replication between on-premises and IBM Cloud environments.

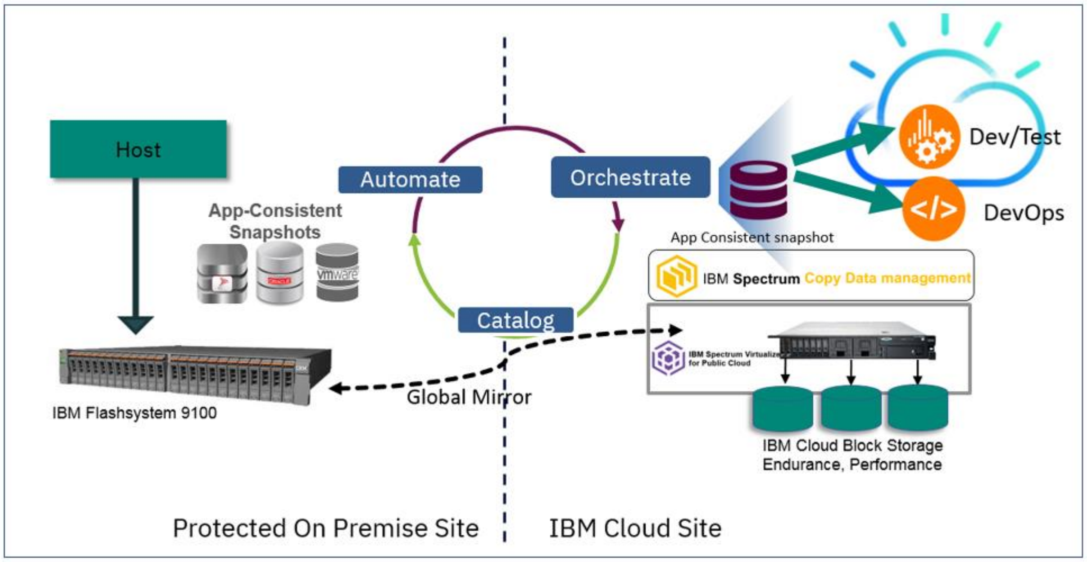
Host
App-Consistent Snapshots
Automate
Orchestrate
Catalog
Dev/Test
DevOps
App Consistent snapshot
IBM Spectrum Copy Data management
IBM Spectrum Virtualize for Public Cloud
IBM Flashsystem 9100
Global Mirror
IBM Cloud Block Storage Endurance, Performance
Protected On Premise Site
IBM Cloud Site

Figure 2: High-level architecture diagram for data re-use

The same solution is extended for data re-use in IBM Cloud. The snapshot copies created with FlashCopy technology in IBM Cloud can be leveraged for DevOps, analytics and reporting. This document describes the use of IBM Spectrum Copy Data Management to leverage use cases such as automated DR and provisioning of automated on-demand test/development environment for VMware environments. The solution presented provides full lifecycle management through automated workflows that allow you to streamline the creation, management and use of copies of data.

Configuring protected on-premises site

This section describes the benefits, features and configuration overview of IBM FlashSystem 9100.

IBM FlashSystem 9100 supports petabytes of effective data storage in a very efficient two-rack-unit chassis. IBM FlashSystem 9100 uses IBM FlashCore technology in a standard 2.5-inch SSD form factor with NVMe interfaces, so 24 IBM FlashCore modules (FCMs) can form the basis of a storage array delivering consistent microsecond latency, extreme reliability, and a wide range of operational and cost efficiencies.

The IBM FlashSystem 9100 architecture allows you to choose IBM FCMs in multiple capacities, or you can opt for industry-standard NVMe-enabled flash drives, with the capability to support both drive types simultaneously within the array.

This means that with inline high-performance data compression, effective capacities can range up to two petabytes in a single 2U enclosure, with the ability to cluster, scale out, or scale up capacity and performance to many petabytes and millions of IOPS.

IBM FlashSystem 9100 is offered in two basic models—IBM FlashSystem 9110 and IBM FlashSystem 9150. Both solutions feature dual controller canisters, dual power supplies and redundant cooling. Both models have two Intel Skylake CPUs per controller canister, with IBM FlashSystem 9110 offering eight cores per CPU while the 9150 model comes with 14 cores per CPU. Up to 768 GB of memory can be configured per controller, so, with a single 2U storage array, you can leverage the performance and efficiency of more than a terabyte of memory, and multiple petabytes of storage—all moving at NVMe speeds—to tackle even the most demanding real-time analytics or cognitive application workloads.

IBM Spectrum Virtualize software is the foundation for every IBM FlashSystem 9100 solution. With virtualization capabilities that can be extended to more than 440 IBM and non-IBM heterogeneous storage systems, it provides features such as automated data movement; synchronous and asynchronous copy services either on-premises or to the public cloud; high-availability configurations; storage tiering; and data reduction technologies.

IBM FlashSystem 9100 solutions can function as IT infrastructure modernization and transformation engines, with inbuilt IBM Spectrum Virtualize capabilities that allow you to virtualize all legacy external heterogeneous storage systems under management, reducing both capital and operational costs while helping to increase the return on your investments in legacy infrastructure.

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 using IBM FlashSystem 9100 and the public cloud. It enables real-time storage-based data replication and DR, as well as data migration between local storage and the IBM Cloud.

The IBM Spectrum Virtualize technology within IBM FlashSystem 9100 arrays offers powerful data reduction pool capabilities that include block deduplication that works to minimize the number of data copies stored and hardware-accelerated data compression technology that provides consistent, high-performance results across all application workload patterns.

The data reduction pools use a log-structured design built on top of the more efficient, distributed RAID 6 provided by IBM FlashCore technology. IBM FlashSystem 9100 data reduction pool capabilities support the SCSI UNMAP command, which allows software to tell the storage system when it’s no longer using portions of storage. This capacity is then returned to the pool to be used to satisfy other requirements. Previously, storage would stay assigned even if it was no longer being used, which wastes capacity.

IBM FlashSystem 9100 at a glance		
Models	IBM FlashSystem 9110, model AF7	IBM FlashSystem 9150, model AF8
System size	Single 2U enclosure	Clustered 4-way x 2U enclosures
Flash type	IBM-enhanced 3D TLC	
Supported drives	2.5-inch NVMe IBM FCMs 4.8 TB, 9.6 TB and 19.2 TB compressing IBM FCMs 2.5-inch NVMe flash drives 1.92 TB, 3.84 TB, 7.68 TB and 15.36 TB	
Maximum NVMe flash capacity	461 TB raw 379 TB usable, DRAID 6 758 TB effective (2:1 reduction)	1.8 PB raw 1.5 PB usable, DRAID 6 3.0 PB effective (2:1 reduction)
Maximum external storage capacity	External virtualization: Up to 32 PB usable capacity	
Drives, canisters, fans and power supplies	Fully redundant, hot-swappable	
Management software	IBM Spectrum Virtualize software	
	Deduplication and compression FlashCopy Remote mirroring External virtualization IBM Easy Tier® Data migration	

IBM FlashSystem 9100 at a glance (cont.)		
Encryption	Data-at-rest AES-XTS 256	
NVMe-oF hardware ready connectivity	Up to: 24 ports 16 GB Fibre Channel 8 ports 10 GBe iSCSI 12 ports 25 GBe iWARP or RoCE	Up to: 96 ports 16 GB Fibre Channel 16 ports 10 GBe iSCSI 48 ports 25 GBe iWARP or RoCE
SAS Expansion enclosures	Model AFF 2U 24 drive Model A9F 5U 92 drive 2.5-inch flash drives supported: 1.92 TB, 3.2 TB, 3.84 TB, 7.68 TB and 15.36 TB	
Controller CPU	Model AF7 - Four 8-core Model AF8 - Four 14-core	Sixteen 8-core Sixteen 14-core
Cache	128 GB standard; up to 1,536 GB	512 GB standard; up to 6,144 GB
Dimensions	Control enclosure: Width: 48.3 cm (19.0 in) Depth: 85.0 cm (33.5 in) Height: 8.8 cm (3.5 in)	
Weight	Control enclosure: Drive-ready (without drive modules installed): 38.5 kg (84.7 lb) Fully configured (24 IBM FCMs installed): 46.6 kg (102.5 lb)	

Table 3: IBM FlashSystem 9100 data sheet

Configuration of IBM FlashSystem 9100

The IBM FlashSystem 9100 system used in the lab setup is an IBM FlashSystem 9110 model AF7, with 24 3.2 TB flash drives. The drives are configured as Tier 0 with Easy Tier. To view the system overview page, log in to the IBM FlashSystem 9110 GUI, then log in to the cluster IP address using a supported web browser and click on **System**.

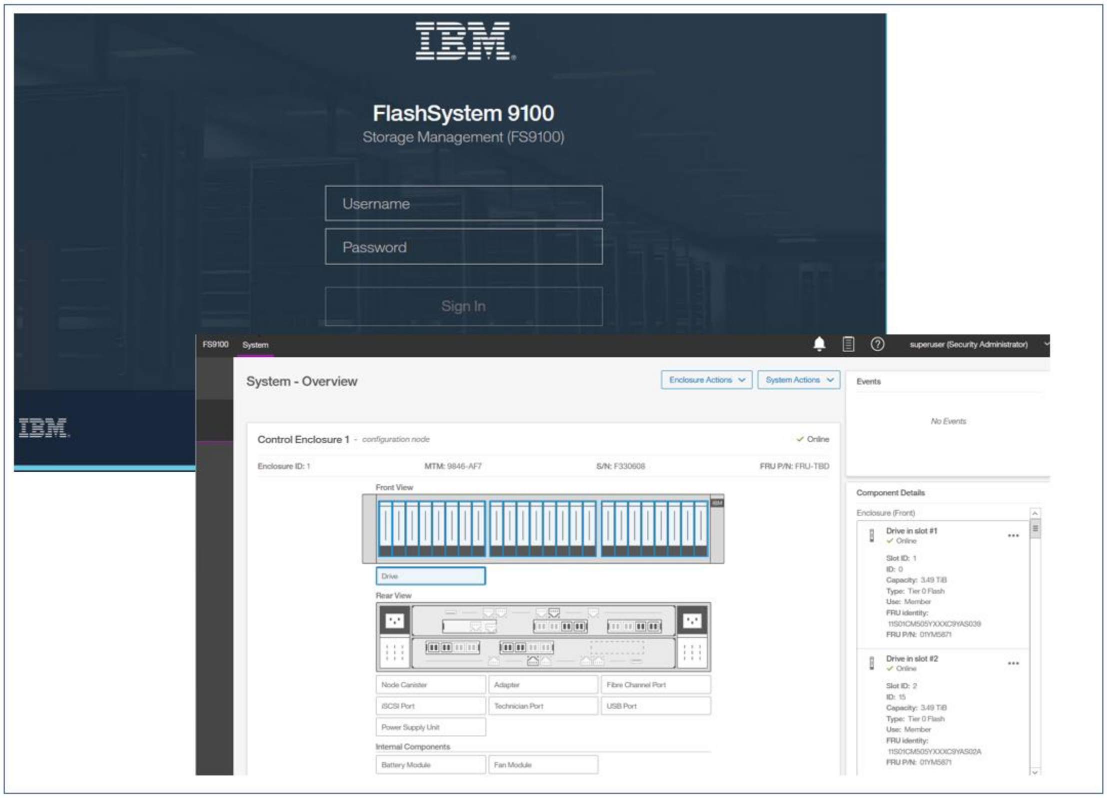

Enclosure ID: 1 MTM: 9846-AF7 S/N: F330608 FRU P/N: FRU-TBD

Drive in slot #1
✓ Online
Slot ID: 1
ID: 0
Capacity: 3.49 TiB
Type: Tier 0 Flash
Use: Member
FRU identity: 11S01CM505YXXXC9YAS039
FRU P/N: 01YM5871

Drive in slot #2
✓ Online
Slot ID: 2
ID: 15
Capacity: 3.49 TiB
Type: Tier 0 Flash
Use: Member
FRU identity: 11S01CM505YXXXC9YAS02A
FRU P/N: 01YM5871

Figure 3: IBM FlashSystem 9100 login screen

Click **Pools → Internal Storage** as shown in Figure 4.

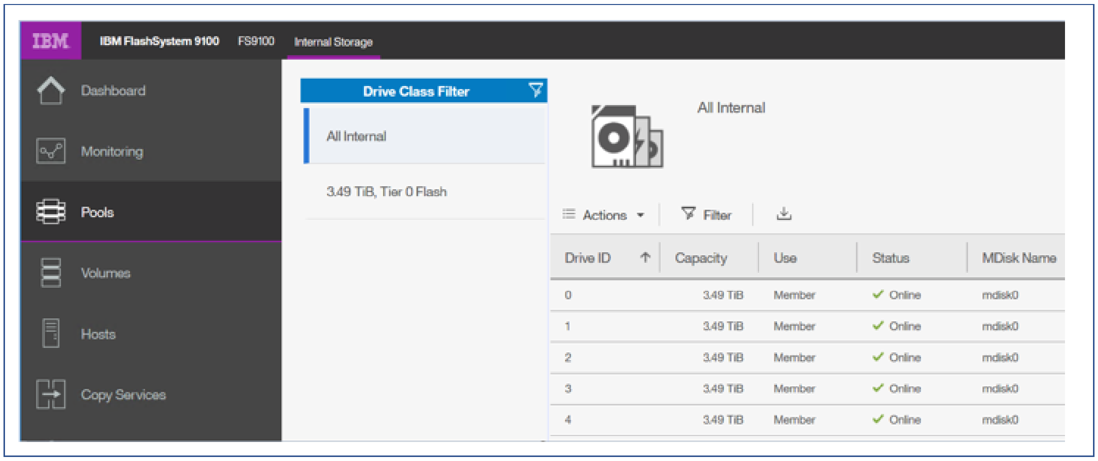

Drive ID	Capacity	Use	Status	MDisk Name
0	3.49 TiB	Member	✓ Online	mdisk0
1	3.49 TiB	Member	✓ Online	mdisk0
2	3.49 TiB	Member	✓ Online	mdisk0
3	3.49 TiB	Member	✓ Online	mdisk0
4	3.49 TiB	Member	✓ Online	mdisk0

Figure 4: IBM FlashSystem 9100 disk information

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 5.

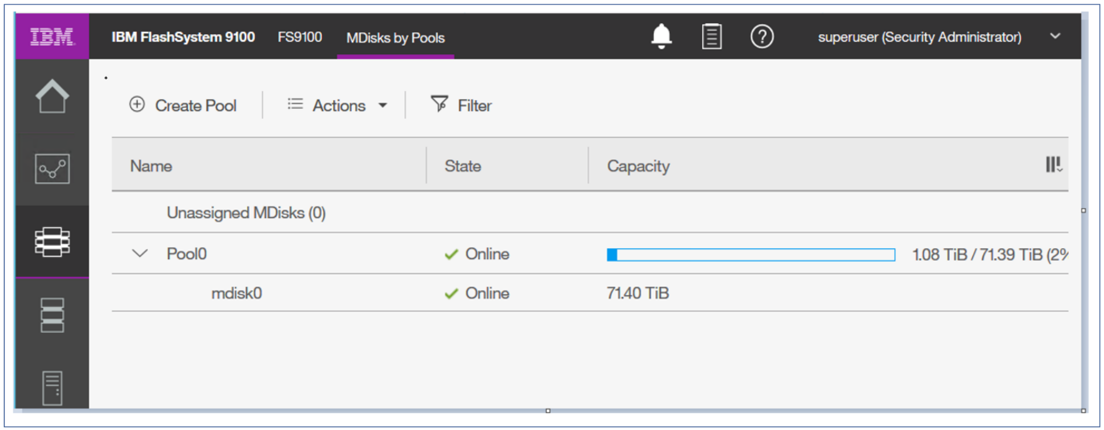

Name	State	Capacity
Unassigned MDisks (0)		
Pool0	✓ Online	1.08 TiB / 71.39 TiB (2%
mdisk0	✓ Online	71.40 TiB

Figure 5: IBM FlashSystem 9100 pool creation

Once the pool is created, create a vdisk and map the vdisk to the ESXi host.

To create a vdisk, click **Volumes by Pool → Create Volumes** as shown in Figure 6.

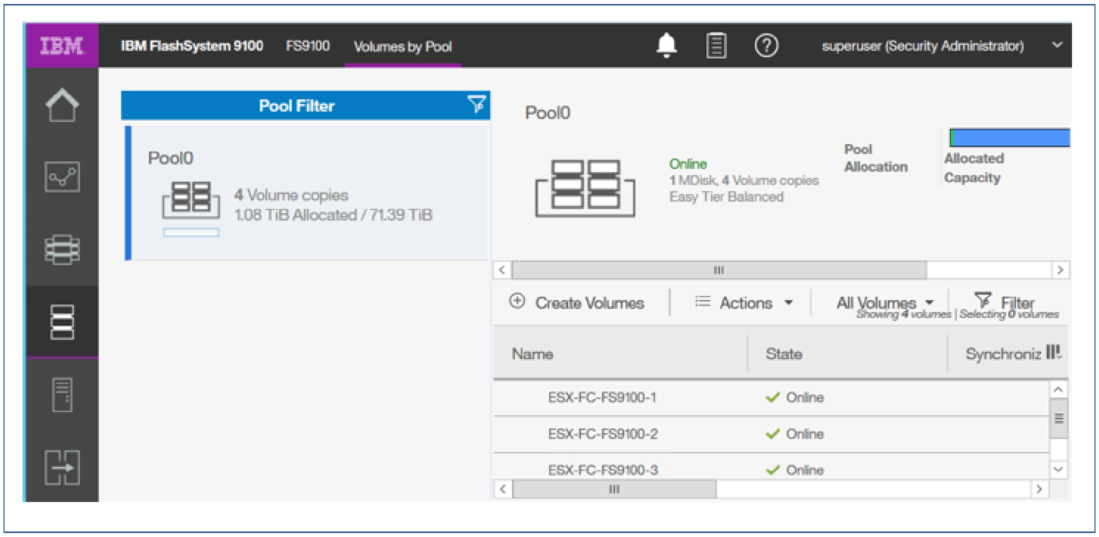

Name	State	Synchroniz
ESX-FC-FS9100-1	✓ Online	
ESX-FC-FS9100-2	✓ Online	
ESX-FC-FS9100-3	✓ Online	

Figure 6: IBM FlashSystem 9100 volume creation

On the pool-creation screen, as shown in the figure above, 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 needs to be enabled or disabled. Then click on **Create and Map** as shown in Figure 7. Follow the wizard and map the volume to the ESXi host at the on-premises site.

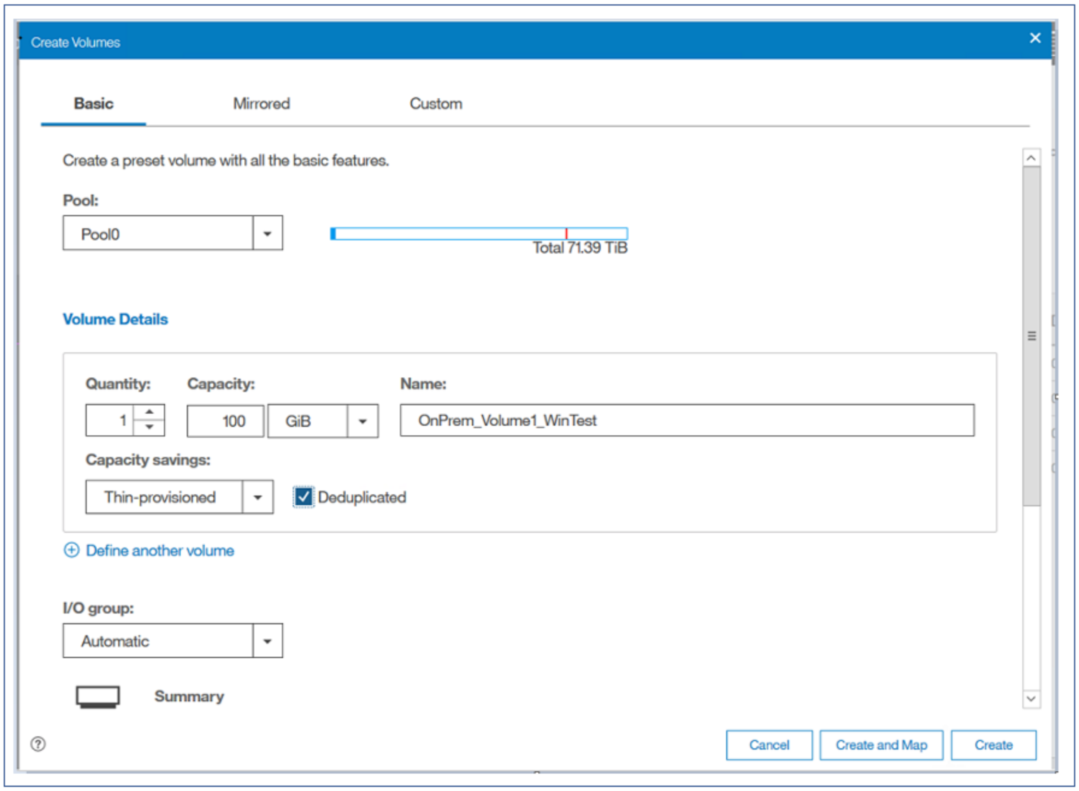

Create a preset volume with all the basic features.

Pool: Pool0 — Total 71.39 TiB

Quantity: 1 Capacity: 100 GiB Name: OnPrem_Volume1_WinTest

Capacity savings: Thin-provisioned ☑ Deduplicated

I/O group: Automatic

Figure 7: Create and map volume

Configuring recovery site at IBM 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 described, a two-node IBM Spectrum Virtualize for Public Cloud cluster is configured.

Installation of IBM Spectrum Virtualize for Public Cloud

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

- Rent bare-metal server in IBM Cloud
- Order appropriate VLAN and networking in IBM Cloud
- Order IBM Cloud Block Storage (Endurance or Performance)
- Purchase IBM Spectrum Virtualize for Public Cloud activation key

To order the infrastructure in IBM Cloud:

- Sign up for IBM Cloud
- IBM Spectrum Virtualize for Public Cloud uses the IBM Cloud computing platform that provides servers, networking and storage. To create an account for IBM Cloud, click the following link to get started: https://control.bluemix.net/

  For detailed information and steps for ordering bare-metal servers in IBM Cloud, refer to IBM Knowledge Center at: https://www.ibm.com/support/knowledgecenter/en/STHLEK_8.1.3/spectrum.virtualize.cloud.813.doc/svcl_plan_provision.html

The bare-metal server ordered will have 64 GB of RAM with Red Hat Enterprise Linux 7.x pre-installed and will come with one public IP address and one private IP address. IBM Cloud portal shows the ordered bare-metal servers as shown in Figure 8.

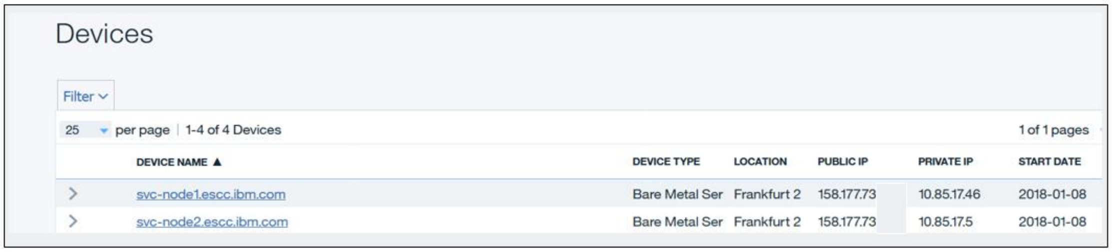

Devices

	DEVICE NAME ▲	DEVICE TYPE	LOCATION	PUBLIC IP	PRIVATE IP	START DATE
>	svc-node1.escc.ibm.com	Bare Metal Ser	Frankfurt 2	158.177.73	10.85.17.46	2018-01-08
>	svc-node2.escc.ibm.com	Bare Metal Ser	Frankfurt 2	158.177.73	10.85.17.5	2018-01-08

Figure 8: IBM Cloud bare-metal server details

The next step is to order the appropriate networking infrastructure in IBM Cloud.

- Order two VLANs—one each for public IP subnet and private IP subnet
- Order a private IP subnet block, which will be part of the private VLAN; for the lab solution in the example, that is *`VLAN ID 818`*
- Order a portable private IP subnet block, which will be part of the same private VLAN; for the lab solution in this example, the private IP block is *`10.85.17.0/26`* (64 IPs) and the portable IP block is *`10.85.33.64/26`* (64 IPs), part of same *`VLAN ID: 846`*

- Order the network gateway appliance from https://control.bluemix.net/ as shown in Figure 9.

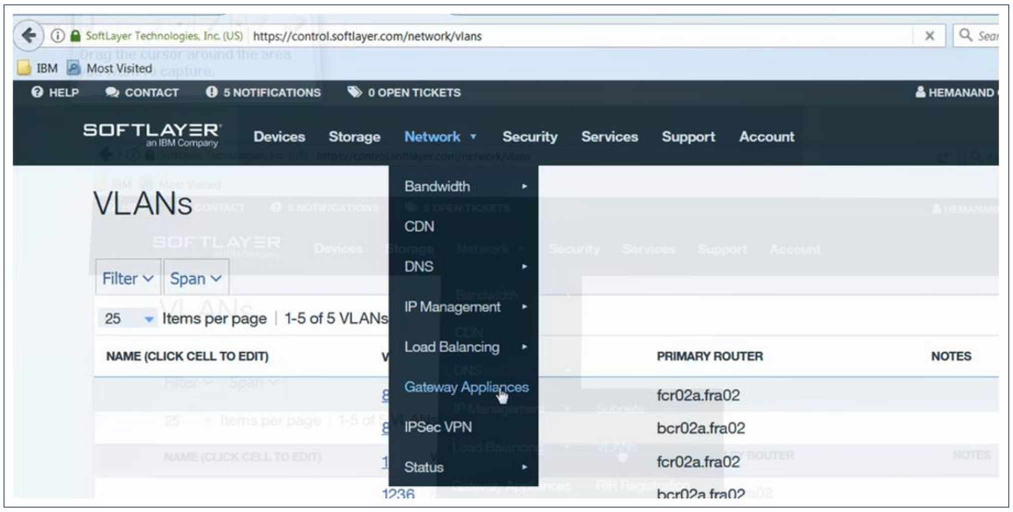
SoftLayer Technologies, Inc. (US) https://control.softlayer.com/network/vlans
IBM Most Visited
HELP CONTACT 5 NOTIFICATIONS 0 OPEN TICKETS HEMANAND
SOFTLAYER an IBM Company
Devices Storage Network Security Services Support Account
VLANs
Filter Span
25 Items per page | 1-5 of 5 VLANs
NAME (CLICK CELL TO EDIT)
PRIMARY ROUTER
NOTES
Bandwidth
CDN
DNS
IP Management
Load Balancing
Gateway Appliances
IPSec VPN
Status
1236
fcr02a.fra02
bcr02a.fra02
fcr02a.fra02
bcr02a.fra02

Figure 9: IBM Cloud network infrastructure ordering

The ordered gateway appliance comes with one public IP, one private IP and the public bandwidth for use. The gateway appliance of IBM Cloud portal is as shown in Figure 10.

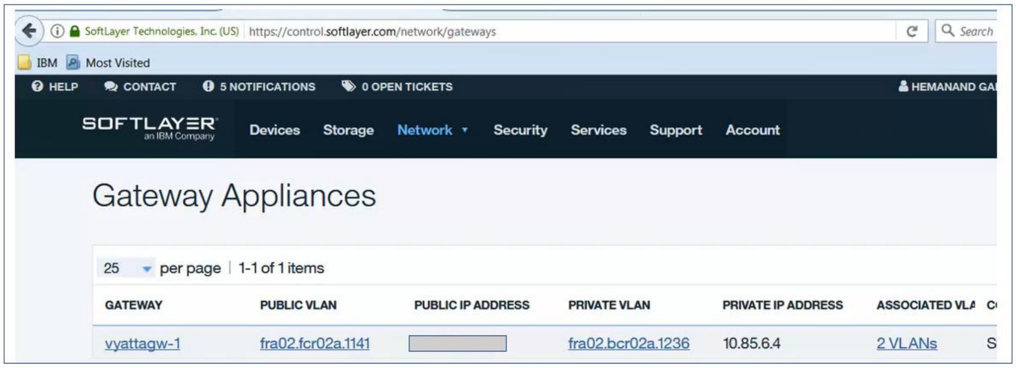
SoftLayer Technologies, Inc. (US) https://control.softlayer.com/network/gateways
Search
IBM Most Visited
HELP CONTACT 5 NOTIFICATIONS 0 OPEN TICKETS HEMANAND GA
SOFTLAYER an IBM Company
Devices Storage Network Security Services Support Account

Gateway Appliances

25 per page | 1-1 of 1 items

GATEWAY	PUBLIC VLAN	PUBLIC IP ADDRESS	PRIVATE VLAN	PRIVATE IP ADDRESS	ASSOCIATED VLA	C
vyattagw-1	fra02.fcr02a.1141	[redacted]	fra02.bcr02a.1236	10.85.6.4	2 VLANs	S

Figure 10: IBM Cloud network gateway appliance

There is one public VLAN and one private VLAN associated with the rented network gateway appliance. As shown in Figure 11, there are two VLANs associated to the gateway.

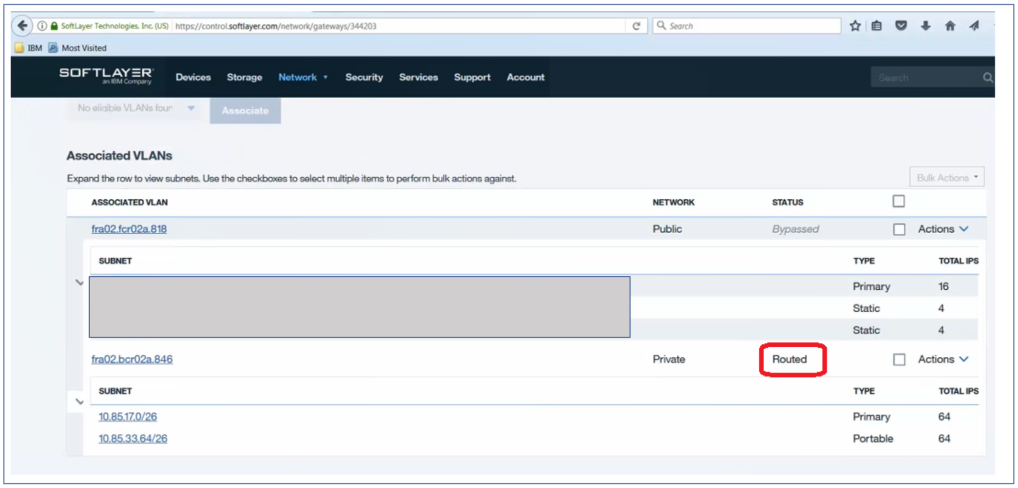

Associated VLANs

Expand the row to view subnets. Use the checkboxes to select multiple items to perform bulk actions against.

ASSOCIATED VLAN	NETWORK	STATUS	☐
fra02.fcr02a.818	Public	*Bypassed*	☐ Actions

SUBNET	TYPE	TOTAL IPS
[redacted]	Primary	16
[redacted]	Static	4
[redacted]	Static	4

ASSOCIATED VLAN	NETWORK	STATUS	☐
fra02.bcr02a.846	Private	Routed	☐ Actions

SUBNET	TYPE	TOTAL IPS
10.85.17.0/26	Primary	64
10.85.33.64/26	Portable	64

Figure 11: IBM Cloud VLAN and subnets

Set the private VLAN status to “Routed,” as shown in Figure 11.

Implementation of IBM Spectrum Virtualize for Public Cloud software

All the cloud resources required for an IBM Spectrum Virtualize for Public Cloud implementation on the IBM Cloud must be provisioned before the IBM Spectrum Virtualize for Public Cloud software can be deployed. This section contains detailed instruction for implementing IBM Spectrum Virtualize in IBM Cloud. IBM Spectrum Virtualize for Public Cloud implementation starts from the following assumptions:

- The required bare-metal server resources are in place and active
- The appropriate network infrastructure is available in IBM Cloud

Download the one-click installer from the location http://5.10.112.93/one-click-install-Win.zip. (The one-click installer is also available for Apple macOS and for Red Hat Enterprise Linux.)

IBM Spectrum Virtualize for Public Cloud can be installed in two different modes:

- Fully automated installation
- Semi-automated installation

This document explains the steps involved in installing IBM Spectrum Virtualize for Public Cloud in semi-automated mode.

For installing in semi-automated mode, it could be valuable for you to produce a spreadsheet in advance to assign a specific IP address to each specific role, as shown:

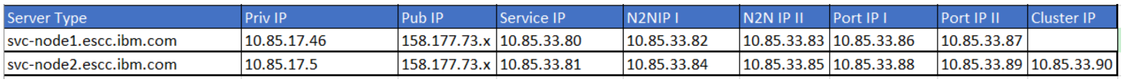

Server Type	Priv IP	Pub IP	Service IP	N2NIP I	N2N IP II	Port IP I	Port IP II	Cluster IP
svc-node1.escc.ibm.com	10.85.17.46	158.177.73.x	10.85.33.80	10.85.33.82	10.85.33.83	10.85.33.86	10.85.33.87	
svc-node2.escc.ibm.com	10.85.17.5	158.177.73.x	10.85.33.81	10.85.33.84	10.85.33.85	10.85.33.88	10.85.33.89	10.85.33.90

Table 4: IP planning for IBM Spectrum Virtualize for Public Cloud

Each bare-metal node requires five IPs to be assigned, from the portable private IP subnet block `10.85.33.64/26.` The purpose of the various IPs are as follows:

Service IP: Required for service management for the IBM Spectrum Virtualize for Public Cloud node

Port IPs: Used for back-end iSCSI Storage virtualization, IP replication for remote copy and iSCSI connection to the host

Node IPs: Used for IP clustering between two IBM Spectrum Virtualize for Public Cloud nodes

Cluster IP: For IBM Spectrum Virtualize for Public Cloud cluster management and to access IBM Spectrum Virtualize for Public Cloud

Once the IP plan is ready, create a YAML file with your IP configuration, as shown below, for both the nodes:

```
    # version=8.1.3.0-180623_1711
cluster:
  ipAddress: 10.85.33.90  # cluster ip
  gateway: 10.85.33.65    #
  netmask: 255.255.255.192
  site1:
    BareMetalServers:
    - servername: svc-node1  # the name showed in cloud web portal
      publicIpAddress: x.x.x.x
      privateIpAddress: 10.85.17.46
      user: root  # username with root privilege
      password: xxxxxxx  # login password for user
      privateKey: C:\Users\ADMIN\.ssh\bm01_private_key
      serial: SL000001 # Bare Metal server serial number
      id: 123456
      serviceIp:
        netmask: 255.255.255.192  # add in runtime
        ipAddress: 10.85.33.80
        gateway: 10.85.33.65
      portIp:
      - netmask: 255.255.255.192
        ipAddress: 10.85.33.86
        gateway: 10.85.33.65
      - netmask: 255.255.255.192
        ipAddress: 10.85.33.87
        gateway: 10.85.33.65
      nodeIp:
      - netmask: 255.255.255.192
        ipAddress: 10.85.33.82
        gateway: 10.85.33.65
      - netmask: 255.255.255.192
        ipAddress: 10.85.33.83
        gateway: 10.85.33.65

servername: svc-node2  # the name showed in cloud web portal
      publicIpAddress: x.x.x.x
      privateIpAddress: 10.85.17.46
      user: root  # username with root privilege
      password: xxxxxxx  # login password for user
      privateKey: C:\Users\ADMIN\.ssh\bm01_private_key
      serial: SL000002  # Bare Metal server serial number
      id: 123456
      serviceIp:
        netmask: 255.255.255.192  # add in runtime
        ipAddress: 10.85.33.81
        gateway: 10.85.33.65
      portIp:
      - netmask: 255.255.255.192
        ipAddress: 10.85.33.88
        gateway: 10.85.33.65
      - netmask: 255.255.255.192
        ipAddress: 10.85.33.89
        gateway: 10.85.33.65
      nodeIp:
      - netmask: 255.255.255.192
        ipAddress: 10.85.33.84
        gateway: 10.85.33.65
      - netmask: 255.255.255.192
        ipAddress: 10.85.33.85
        gateway: 10.85.33.65
```

Listing 1: YAML file for two-node cluster

Save the YAML file where the installed software is kept and run the `installer.exe` from the Windows machine.

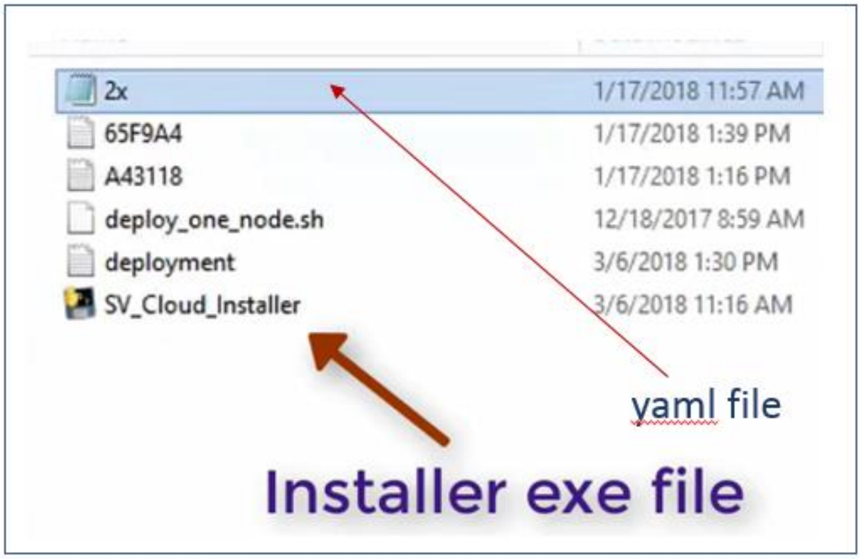

2x	1/17/2018 11:57 AM
65F9A4	1/17/2018 1:39 PM
A43118	1/17/2018 1:16 PM
deploy_one_node.sh	12/18/2017 8:59 AM
deployment	3/6/2018 1:30 PM
SV_Cloud_Installer	3/6/2018 11:16 AM

Figure 12: Installer for IBM Spectrum Virtualize for Public Cloud

Once the installer is running, it will generate a nonce—an ID unique to each node. This is a six-digit number which is used to generate the activation key for cluster activation. The unique ID generated for each node can be seen in Figure 13.

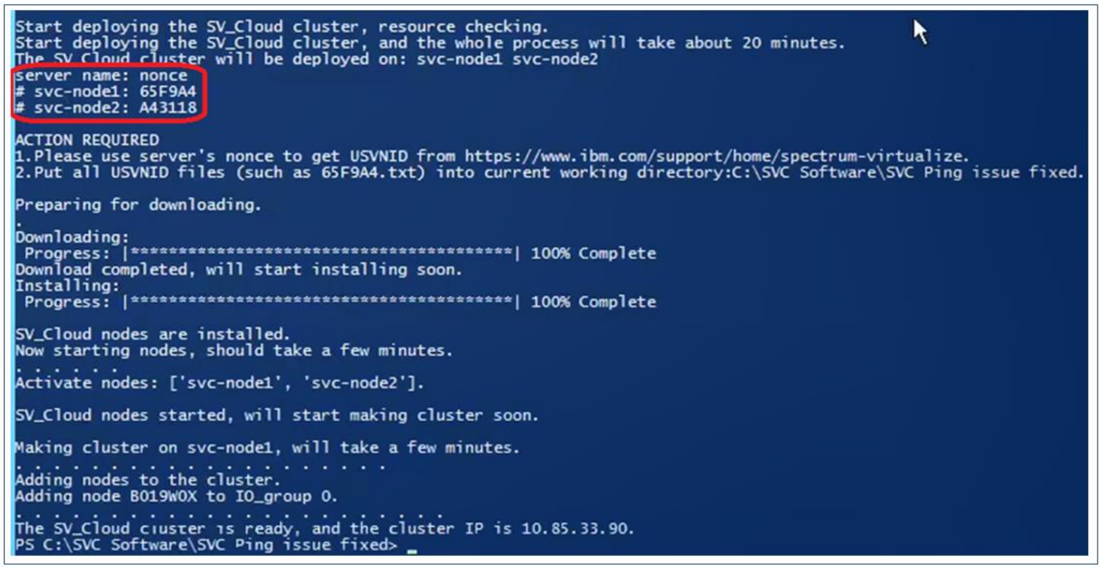

```
Start deploying the SV_Cloud cluster, resource checking.
Start deploying the SV_Cloud cluster, and the whole process will take about 20 minutes.
The SV_Cloud cluster will be deployed on: svc-node1 svc-node2
server name: nonce
# svc-node1: 65F9A4
# svc-node2: A43118

ACTION REQUIRED
1.Please use server's nonce to get USVNID from https://www.ibm.com/support/home/spectrum-virtualize.
2.Put all USVNID files (such as 65F9A4.txt) into current working directory:C:\SVC Software\SVC Ping issue fixed.

Preparing for downloading.
.
Downloading:
 Progress: |****************************************| 100% Complete
Download completed, will start installing soon.
Installing:
 Progress: |****************************************| 100% Complete

SV_Cloud nodes are installed.
Now starting nodes, should take a few minutes.
. . . . . .
Activate nodes: ['svc-node1', 'svc-node2'].

SV_Cloud nodes started, will start making cluster soon.

Making cluster on svc-node1, will take a few minutes.
. . . . . . . . . . . . . . . . . . . .
Adding nodes to the cluster.
Adding node B019W0X to IO_group 0.
. . . . . . . . . . . . . . . . . . . . . . .
The SV_Cloud cluster is ready, and the cluster IP is 10.85.33.90.
PS C:\SVC Software\SVC Ping issue fixed> _
```

Figure 13: IBM Spectrum Virtualize for Public Cloud installation

To complete the installation, the activation key is required. Go to the IBM support web site at: https://www.ibm.com/support/home/spectrum-virtualize

Download the activation key for each node in the IBM Spectrum Virtualize cluster and save the key to the same location where it installer is kept; the installer will activate and create the cluster. In the above example, the cluster IP is `10.85.33.90.`

Logging in to an IBM Spectrum Virtualize for Public Cloud cluster is almost the same as logging in to a node. Just replace the service IP with the cluster IP. Log in to the cluster with a GUI by using your browser as shown in Figure 14.

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

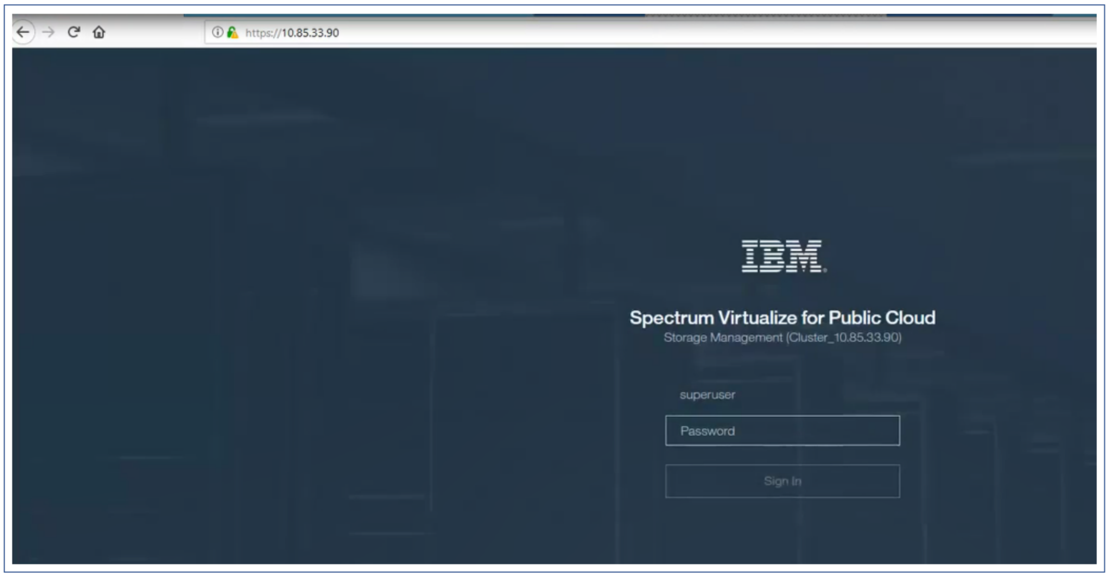
https://10.85.33.90
IBM
Spectrum Virtualize for Public Cloud
Storage Management (Cluster_10.85.33.90)
superuser
Password
Sign In

Figure 14: IBM Spectrum Virtualize for Public Cloud login screen

Configuration of the cloud quorum

IP quorum applications are used in Ethernet networks to resolve failure scenarios in which half the nodes on the system become unavailable. These applications determine which nodes can continue processing host operations to avoid a "split-brain" scenario in which both halves will attempt to independently service I/O, causing corruption. Since IBM Spectrum Virtualize for Public Cloud cluster relies on nodes identified by IP address, it requires at least one IP quorum application on either a bare-metal or virtual server in IBM Cloud.

The IP quorum application is required for both two- and four-node systems in IBM Spectrum Virtualize for Public Cloud configurations. In two-node systems, the IP quorum application maintains availability after a node failure. In systems with four nodes, an IP quorum application is necessary to handle other failure scenarios. The IP quorum application is a Java application that runs on a separate bare-metal or virtual server in IBM Cloud.

There are strict requirements on the IP network when using IP quorum applications. All IP quorum applications must be reconfigured and redeployed to hosts when certain aspects of the system configuration change. These aspects include adding or removing a node from the system, or when node service IP addresses are changed.

Please refer to the IBM Knowledge Center for details of cloud quorum configuration: https://www.ibm.com/support/knowledgecenter/STHLEK_8.1.3/spectrum.virtualize.cloud.813.doc/svc_ipquorumconfig.html

Configuration of back-end storage

IBM Spectrum Virtualize for Public Cloud uses the back-end storage provided by IBM Cloud as external MDisk.

To order back-end storage in IBM Cloud:

Log in using https://control.bluemix.net

Click on **Storage → Block Storage → Order Block Storage**.

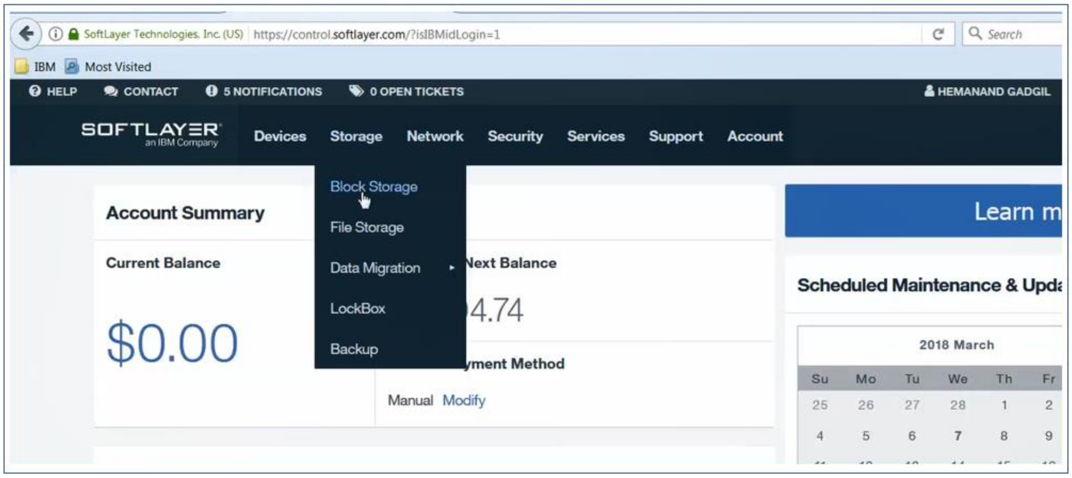

Figure 15: IBM Cloud block storage ordering

Then, select **Endurance** or **Performance** storage depending upon the requirement and the capacity.

Plan details are described in greater detail at: https://www.ibm.com/support/knowledgecenter/en/STHLEK_8.1.3/spectrum.virtualize.cloud.813.doc/svcl_plan_block_storage.html

Once the appropriate storage is ordered, the details of the IBM Cloud block storage will be visible, as shown in Figure 16.

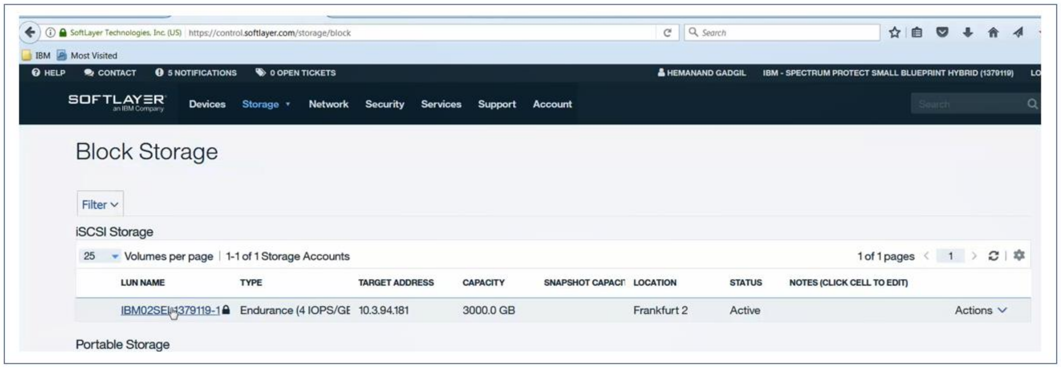

Block Storage

iSCSI Storage

25 Volumes per page | 1-1 of 1 Storage Accounts 1 of 1 pages 1

LUN NAME	TYPE	TARGET ADDRESS	CAPACITY	SNAPSHOT CAPACI	LOCATION	STATUS	NOTES (CLICK CELL TO EDIT)
IBM02SEL1379119-1	Endurance (4 IOPS/GE	10.3.94.181	3000.0 GB		Frankfurt 2	Active	

Portable Storage

Figure 16: IBM Cloud block storage

IBM Spectrum Virtualize for Public Cloud uses iSCSI virtualization for back-end storage. To configure the back-end, you must identify the port IP to be used for virtualization.

The host iSCSI Qualified Name (IQN) details need to be updated for iSCSI virtualization; to query IBM Spectrum Virtualize for Public Cloud IQN, use the `lsnode` command. The details of port IP for iSCSI virtualization and the host IQN can be seen in the GUI as shown in Figure 17.

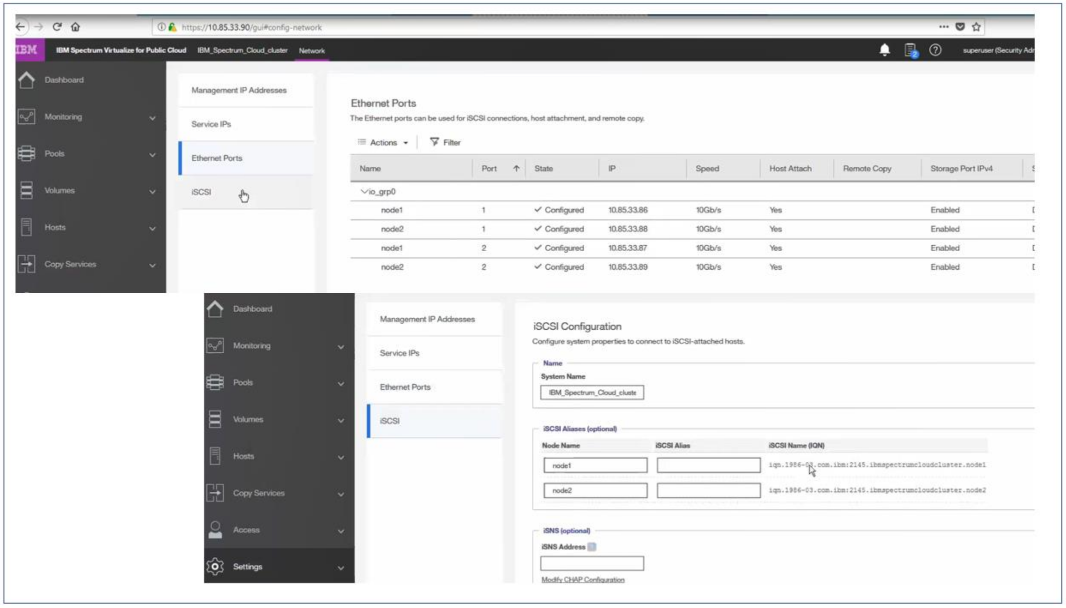

Ethernet Ports

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

Name	Port	State	IP	Speed	Host Attach	Remote Copy	Storage Port IPv4
io_grp0							
node1	1	✓ Configured	10.85.33.86	10Gb/s	Yes		Enabled
node2	1	✓ Configured	10.85.33.88	10Gb/s	Yes		Enabled
node1	2	✓ Configured	10.85.33.87	10Gb/s	Yes		Enabled
node2	2	✓ Configured	10.85.33.89	10Gb/s	Yes		Enabled

iSCSI Configuration

Configure system properties to connect to iSCSI-attached hosts.

Name

System Name: IBM_Spectrum_Cloud_cluste

iSCSI Aliases (optional)

Node Name	iSCSI Alias	iSCSI Name (IQN)
node1		iqn.1986-03.com.ibm:2145.ibmspectrumcloudcluster.node1
node2		iqn.1986-03.com.ibm:2145.ibmspectrumcloudcluster.node2

iSNS (optional)

iSNS Address

Figure 17: IQN for IBM Spectrum Virtualize for Public Cloud

Now go to the IBM Cloud IaaS portal.

1. Click **Block Storage → Storage** as shown in Figure 16 above, then scroll down the page to find and click on the appropriate node in the "Device Name" column as shown in Figure 18.

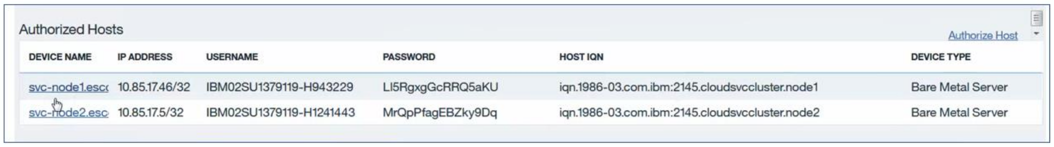

Authorized Hosts

Authorize Host

DEVICE NAME	IP ADDRESS	USERNAME	PASSWORD	HOST IQN	DEVICE TYPE
svc-node1.escc	10.85.17.46/32	IBM02SU1379119-H943229	LI5RgxgGcRRQ5aKU	iqn.1986-03.com.ibm:2145.cloudsvccluster.node1	Bare Metal Server
svc-node2.esc	10.85.17.5/32	IBM02SU1379119-H1241443	MrQpPfagEBZky9Dq	iqn.1986-03.com.ibm:2145.cloudsvccluster.node2	Bare Metal Server

Figure 18: IBM Cloud Portal host authorization

2. Click on the **Storage** tab and update the host IQN as shown in Figure 19.

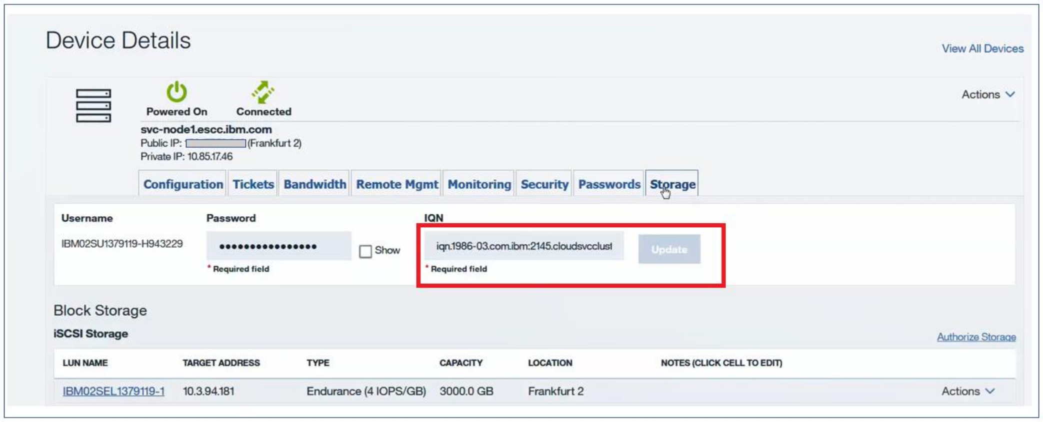

Device Details

View All Devices

svc-node1.escc.ibm.com
Public IP: [redacted] (Frankfurt 2)
Private IP: 10.85.17.46

Username: IBM02SU1379119-H943229

Password: •••••••••••••••• □ Show
* Required field

IQN: iqn.1986-03.com.ibm:2145.cloudsvcclust Update
* Required field

Block Storage

iSCSI Storage

Authorize Storage

LUN NAME	TARGET ADDRESS	TYPE	CAPACITY	LOCATION	NOTES (CLICK CELL TO EDIT)	
IBM02SEL1379119-1	10.3.94.181	Endurance (4 IOPS/GB)	3000.0 GB	Frankfurt 2		Actions

Figure 19: Host IQN update on IBM Cloud Portal

3. Go to IBM Spectrum Virtualize for Public Cloud GUI and move to **Pool** → **External Storage** menu → **Add Cloud Storage** in the GUI as shown Figure 20.

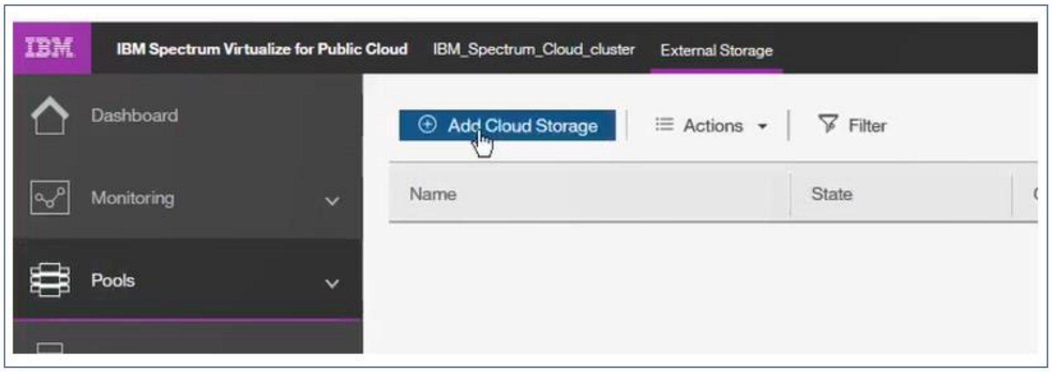

Figure 20: External storage virtualization

1. Choose either the Automatic or Manual option. In our example, we choose the Automatic option:

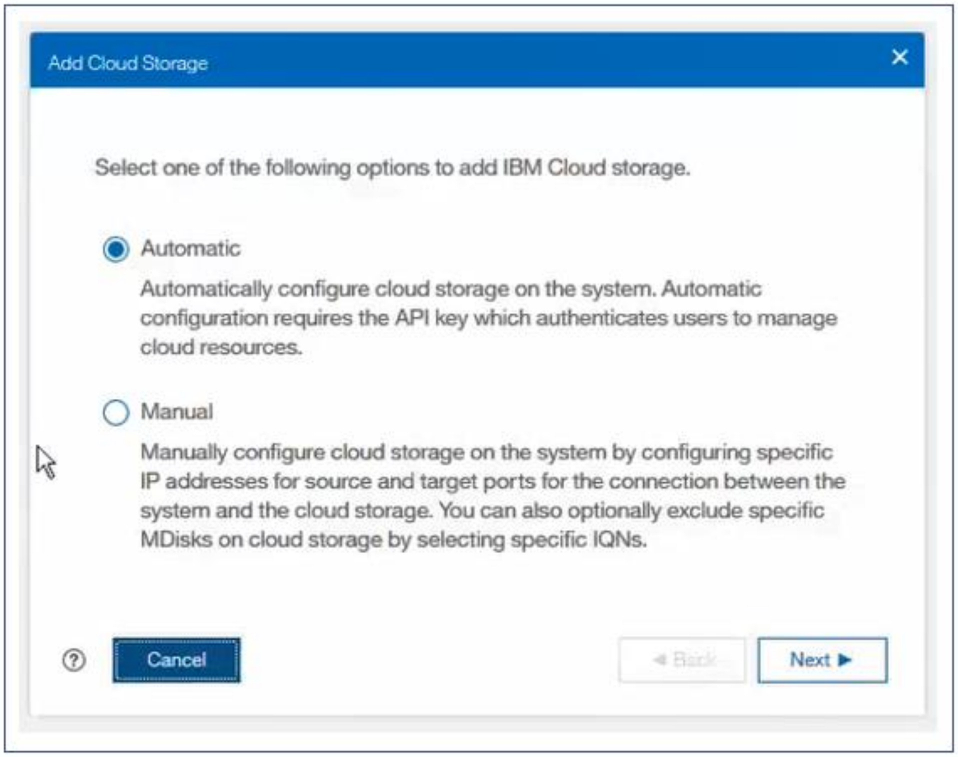
Add Cloud Storage

Select one of the following options to add IBM Cloud storage.

● Automatic

Automatically configure cloud storage on the system. Automatic configuration requires the API key which authenticates users to manage cloud resources.

○ Manual

Manually configure cloud storage on the system by configuring specific IP addresses for source and target ports for the connection between the system and the cloud storage. You can also optionally exclude specific MDisks on cloud storage by selecting specific IQNs.

Figure 21: Add cloud storage

2. Fill in the required fields shown in Figure 22.

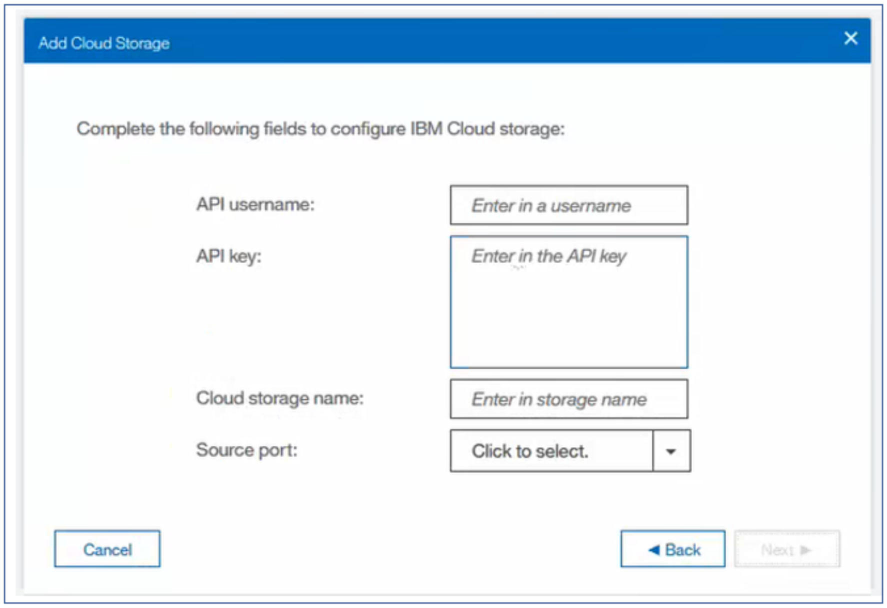

Add Cloud Storage

Complete the following fields to configure IBM Cloud storage:

API username:	Enter in a username
API key:	Enter in the API key
Cloud storage name:	Enter in storage name
Source port:	Click to select.

Figure 22: Add cloud storage configuration

Use the following information when filling out the required fields:

- API username = your IBM Cloud VPN username
- API key = your IBM Cloud key, an example of which is shown in Figure 23

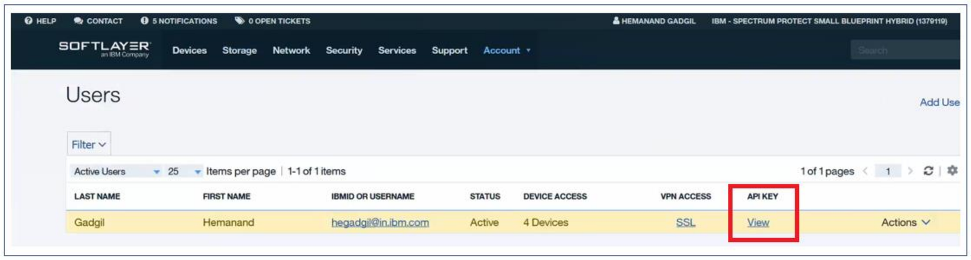

Users

LAST NAME	FIRST NAME	IBMID OR USERNAME	STATUS	DEVICE ACCESS	VPN ACCESS	API KEY	
Gadgil	Hemanand	hegadgil@in.ibm.com	Active	4 Devices	SSL	View	Actions

Figure 23: API key on IBM Cloud Portal

- Cloud storage name = your IBM Cloud Storage logical unit number (LUN), as illustrated in Figure 19 above
- Source port = your IBM Spectrum Virtualize node port ID. You can select from ports ID 1 and ID 2. It is recommended that both be used in round-robin fashion, both for better workload balance and to achieve redundancy for each LUN (MDisk)

Now your storage LUN (MDisk) has been discovered and configured as shown in Figure 24.

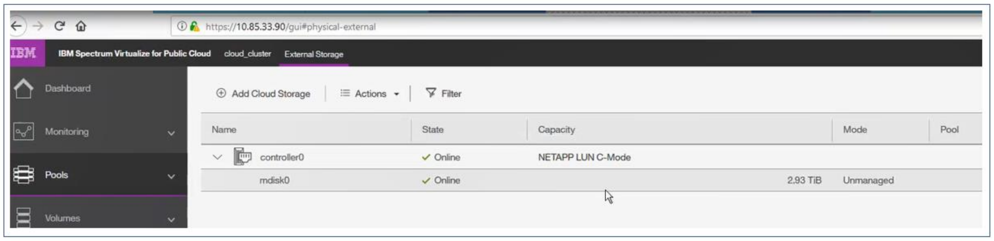

Name	State	Capacity		Mode	Pool
controller0	✓ Online	NETAPP LUN C-Mode			
mdisk0	✓ Online		2.93 TiB	Unmanaged	

Figure 24: External virtualized storage details

6. Now create the Storage Pool. Click **Pool → Create Pool** as shown in Figure 25.

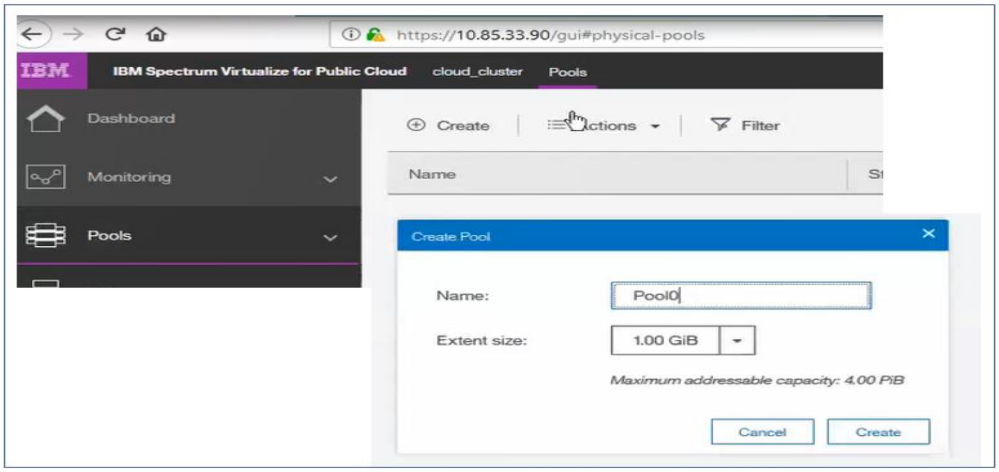

Name: Pool0

Extent size: 1.00 GiB

Maximum addressable capacity: 4.00 PiB

Figure 25: Creating a pool

7. Assign the externally virtualized MDisk to the pool. To assign the disk, right-click **Pool → Add Storage**.

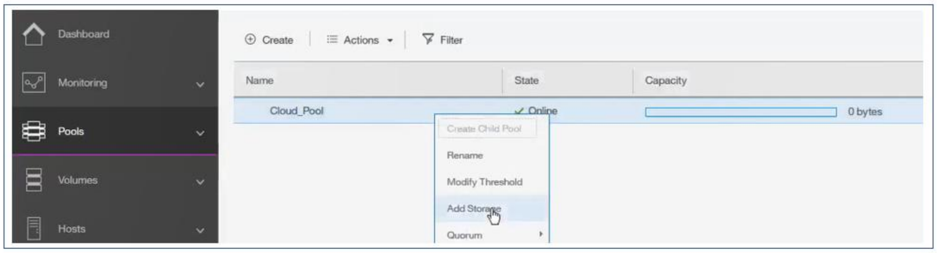

Name	State	Capacity
Cloud_Pool	✓ Online	0 bytes

Figure 26: Adding MDisk to the pool

8. Click on **External** and select the external controller and MDisk as shown in Figure 27.

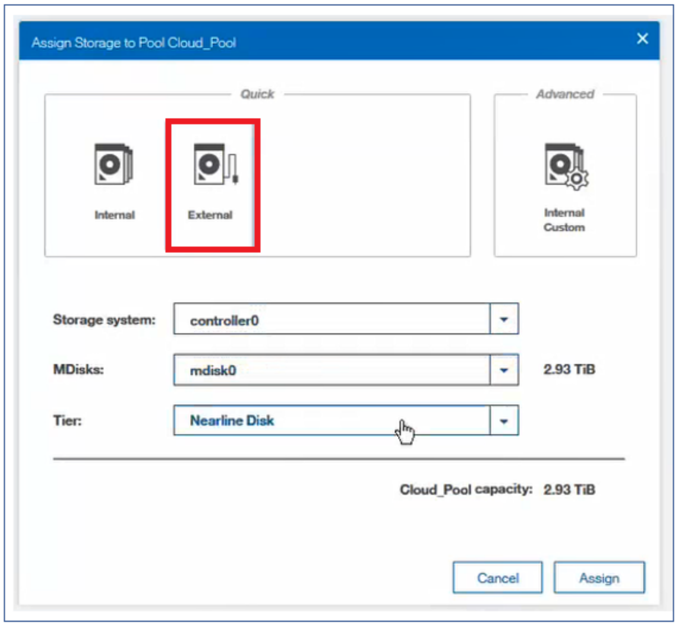

Storage system: controller0

MDisks: mdisk0 2.93 TiB

Tier: Nearline Disk

Cloud_Pool capacity: 2.93 TiB

Figure 27: Adding an external MDisk to the pool

Then you can create a vdisk and assign the volume for host access using iSCSI.

Configuring site-to-site IPsec VPN for hybrid cloud connectivity

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

NOTE: This section describes the logical steps for the use case illustrated, but the on-premises network configuration, infrastructure and security policy may vary on a case-by-case basis. This section is intended to give a high-level logical example.

The high-level logical network architecture diagram for hybrid cloud connectivity is shown in Figure 28.

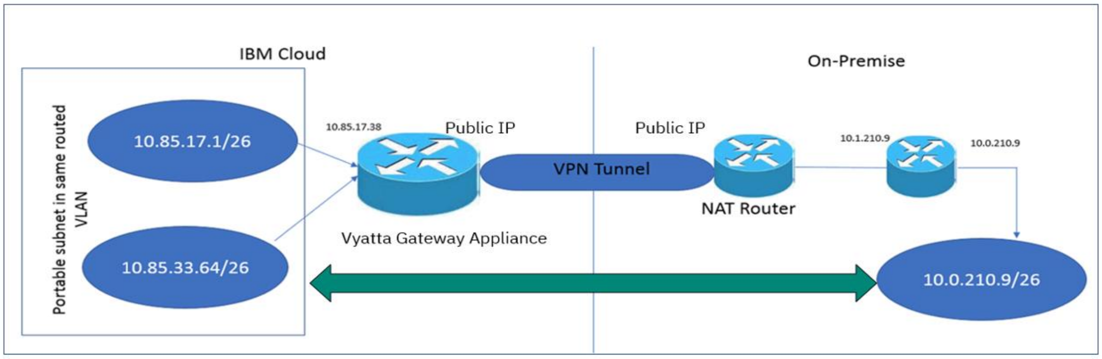

Figure 28: Hybrid cloud network connectivity topology

As shown in Figure 28, the network gateway appliance at the IBM Cloud site is a Vyatta Gateway Appliance, which acts a default router for both the private and portable private subnet IP blocks (`10.85.17.1/26` and `10.85.33.64/26`, respectively).

All of the ESXi host and IBM Spectrum Virtualize for Public Cloud service IPs, port IPs, node IPs and hosted VMs are in the portable IP subnet `10.85.33.64/26`.

At the on-premises site, there is a network address translation (NAT) router, which is the core router, with a public IP address. That public IP address is NAT'ed 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 will create a secure communication network between your IBM Cloud infrastructure and on-premises infrastructure. Network communication between the private subnets is controlled by the access control list populated at the creation of the VPN IPsec site-to-site tunnel.

Configuration steps for IPsec site-to-site VPN tunnel

On the IBM Cloud site's Vyatta Gateway Appliance:

1. Set up the virtual interface (VIF) to all subnets in the VLAN. Get the list of subnets in the VLAN from IBM Cloud: https://control.bluemix.net/network/vlans

   Otherwise, the different subnets can't access each other. In our example, we have `VLAN 846`, which has two private IP subnets, `10.85.17.0/26` and `10.85.33.64/26`.

   To set the VIF, use command:

   ```
   set interfaces bonding dp0bond0 vif 846 address 10.85.33.65/26
   set interfaces bonding dp0bond0 vif 846 address 10.85.17.1/26
   set interfaces bonding dp0bond0 vif 846 vlan 846
   ```

   Also, set the VIF for the public IP subnet on a different bond; in the lab setup, the VLAN ID for the public IP subnet block is 818.

   ```
   set interfaces bonding dp0bond1 vif 818 address dhcp
   set interfaces bonding dp0bond1 vif 818 address x.x.x.x/28
   set interfaces bonding dp0bond1 vif 818 vlan 818
   ```

2. Create the Internet Key Exchange (IKE) and Encapsulating Security Payload (ESP) groups for VPN configuration:

```
set security vpn ipsec ike-group IKE-CLOUDDC proposal 1
set security vpn ipsec ike-group IKE-CLOUDDC proposal 1 encryption
aes256
set security vpn ipsec ike-group IKE-CLOUDDC proposal 1 hash
sha2_256
set security vpn ipsec ike-group IKE-CLOUDDC proposal 1 dh-group 5
set security vpn ipsec ike-group IKE-CLOUDDC lifetime 86400
```

3. Create an ESP group with name CLOUDDC:

```
set security vpn ipsec esp-group ESP-CLOUDDC proposal 1
set security vpn ipsec esp-group ESP-CLOUDDC proposal 1 encryption
'aes256'
set security vpn ipsec esp-group ESP-CLOUDDC proposal 1 hash
'sha2_256'
set security vpn ipsec esp-group ESP-CLOUDDC lifetime '3600'
set security vpn ipsec esp-group ESP-CLOUDDC compression 'disable'
set security vpn ipsec esp-group ESP-CLOUDDC mode 'tunnel'
set security vpn ipsec esp-group ESP-CLOUDDC pfs 'dh-group5'
```

4. Create site-to-site VPN IPSec configuration and set the pre-shared secret key:

```
set security vpn ipsec site-to-site peer <remote on-premise public
IP>
set security vpn ipsec site-to-site peer <remote on-premise public
IP> authentication mode pre-shared-secret
set authentication pre-shared-secret TEST1
set default-esp-group ESP-CLOUDDC
set ike-group IKE-CLOUDDC
set local-address <Local IBM Cloud side public IP >
```

5. Create a tunnel with local private IP subnet and remote private IP subnet. Once the tunnel is created, all the machines that are part of these local and remote subnets can communicate with each other.

```
set tunnel 1 local prefix 10.85.33.64/26
set tunnel 1 remote prefix 10.0.210.0/24
```

Note: In the solution lab environment, the on-premises private subnet is behind the NAT, so NAT traversal (NAT-T) needs be additionally enabled as a part of VPN tunnel configuration.

6. Enable NAT-T and allow the required private subnet to communicate across the VPN tunnel:

```
set security vpn ipsec nat-traversal enable
set security vpn ipsec nat-networks allowed-network 10.0.210.0/24
set security vpn ipsec nat-networks allowed-network 10.1.210.0/24
```

Similarly configure the router for the VPN tunnel at the on-premises site. For this example, we have configured the IP addresses for the VyOS router as follows:

`eth0: 10.0.210.9/24`
`eth1: 10.1.210.9/24` – NAT rules are created to translate this private address to the public address.

a. Configure an IKE group, naming it in our example ONPREMDC:

```
set vpn ipsec ike-group IKE-ONPREMDC proposal 1
set vpn ipsec ike-group IKE-ONPREMDC proposal 1 encryption aes256
set vpn ipsec ike-group IKE-ONPREMDC proposal 1 hash sha256
set vpn ipsec ike-group IKE-ONPREMDC proposal 1 dh-group 5
set vpn ipsec ike-group IKE-ONPREMDC lifetime 86400
```

b. Configure an ESP group, naming it in our example ONPREMDC:

```
set vpn ipsec esp-group ESP-ONPREMDC proposal 1
set vpn ipsec esp-group ESP-ONPREMDC proposal 1 encryption
'aes256'
set vpn ipsec esp-group ESP-ONPREMDC proposal 1 hash 'sha256'
set vpn ipsec esp-group ESP-ONPREMDC lifetime '3600'
set vpn ipsec esp-group ESP-ONPREMDC compression 'disable'
set vpn ipsec esp-group ESP-ONPREMDC mode 'tunnel'
set vpn ipsec esp-group ESP-ONPREMDC pfs 'dh-group5'
```

c. Set the interface that is used for the VPN tunnel; in this example it is eth1:

```
set vpn ipsec ipsec-interfaces interface eth1
```

d. Set the authentication key and create a tunnel between the local (on-premises) private subnet and the remote subnet (cloud site private portable subnet):

```
set vpn ipsec site-to-site peer <Public IP@Cloud Router>
authentication mode pre-shared-secret
set vpn ipsec site-to-site peer <Public IP@Cloud Router>
authentication pre-shared-secret TEST1
set vpn ipsec site-to-site peer <Public IP@Cloud Router> default-
esp-group ESP-ONPREMDC
set vpn ipsec site-to-site peer <Public IP@Cloud Router> ike-
group IKE-ONPREMDC
set vpn ipsec site-to-site peer <Public IP@Cloud Router> local-
address 10.1.210.9
set vpn ipsec site-to-site peer <Public IP@Cloud Router> tunnel 1
local prefix 10.0.210.0/24
set vpn ipsec site-to-site peer <Public IP@Cloud Router> tunnel 1
remote prefix 10.85.33.64/26
```

e. Enable NAT and NAT-allowed networks for the private subnet blocks at the cloud site:

```
set vpn ipsec site-to-site peer <Public IP@Cloud Router> nat-
traversal enable
set vpn ipsec site-to-site peer <Public IP@Cloud Router> nat-
networks allowed-network 10.85.17.0/26
set vpn ipsec site-to-site peer <Public IP@Cloud Router> nat-
networks allowed-network 10.85.33.64/26
```

This step will activate the IPsec VPN tunnel.

To check whether the VPN tunnel is functioning, type the command `show vpn ipsec sa` as shown in Figure 29. In the example, the VPN tunnel is up between the on-premises site and the IBM Cloud site.

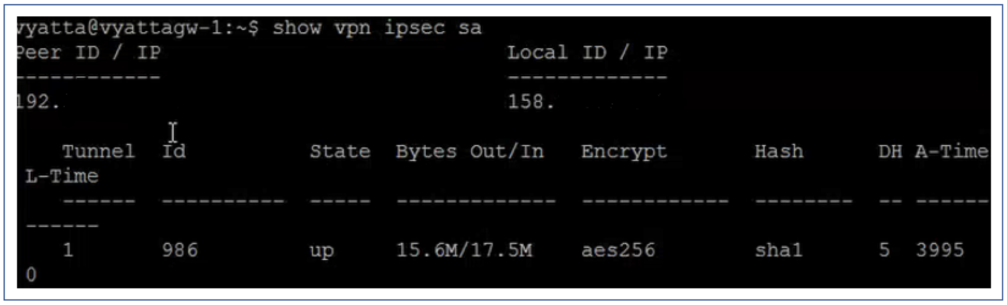

```
vyatta@vyattagw-1:~$ show vpn ipsec sa
Peer ID / IP                               Local ID / IP
------------                               -------------
192.                                       158.

   Tunnel  Id          State  Bytes Out/In   Encrypt      Hash      DH A-Time
 L-Time
   ------  ----------  -----  -------------  ------------ --------  -- ------
 ------
   1       986         up     15.6M/17.5M    aes256       sha1      5  3995
0
```

Figure 29: VPN IPsec tunnel status

Setting up Global Mirror relationship for storage

This section describes the essential details for creating the Global Mirror IP replication relationship between on-premises IBM FlashSystem 9100 and IBM Spectrum Virtualize for Public Cloud cluster running in IBM Cloud.

Once the VPN tunnel is up and running, the machines in the on-premises private subnet `10.0.210.9/26` can connect and ping any machine running in the IBM Cloud private subnet `10.85.33.64/26.`

Steps to create Global Mirror IP replication between on-premises IBM FlashSystem 9100 and IBM Spectrum Virtualize for Public Cloud cluster are as follows.

1. On the IBM Spectrum Virtualize for Public Cloud GUI, go to **Settings → Network → Ethernet** and select the IP port, then right-click to select **Modify Remote Copy** as shown in Figure 30.

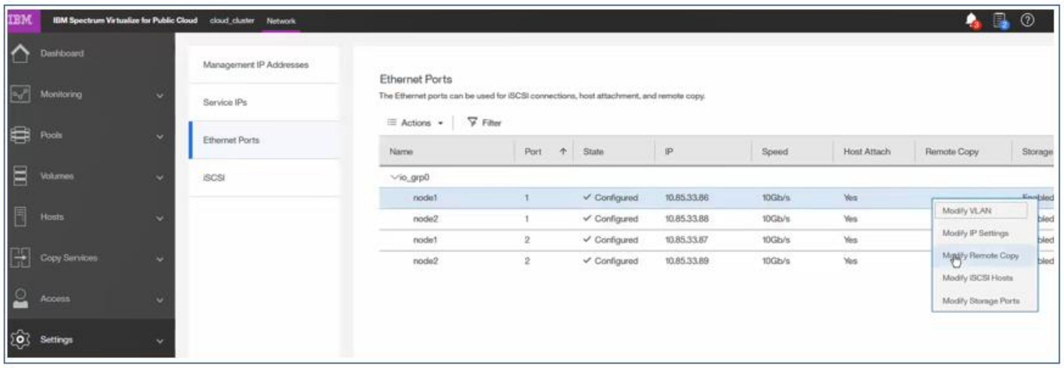

Ethernet Ports

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

Name	Port	State	IP	Speed	Host Attach	Remote Copy	Storage
io_grp0							
node1	1	✓ Configured	10.85.33.86	10Gb/s	Yes		Enabled
node2	1	✓ Configured	10.85.33.88	10Gb/s	Yes		bled
node1	2	✓ Configured	10.85.33.87	10Gb/s	Yes		bled
node2	2	✓ Configured	10.85.33.89	10Gb/s	Yes		bled

Figure 30: Setting ports for remote replication

2. Enable IPv4 for the remote group and select the group, as shown in Figure 31.

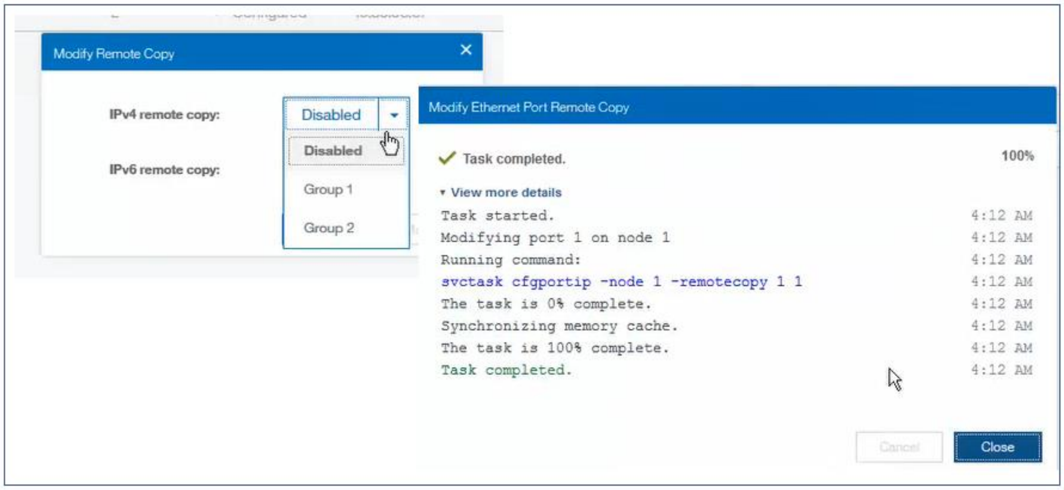

IPv4 remote copy: Disabled

IPv6 remote copy:

✓ Task completed. 100%

▾ View more details

```
Task started.                                              4:12 AM
Modifying port 1 on node 1                                 4:12 AM
Running command:                                           4:12 AM
svctask cfgportip -node 1 -remotecopy 1 1                  4:12 AM
The task is 0% complete.                                   4:12 AM
Synchronizing memory cache.                                4:12 AM
The task is 100% complete.                                 4:12 AM
Task completed.                                            4:12 AM
```

Figure 31: Setting remote copy port

Once you have enabled the port, as shown in Figure 31, you will be able to set the remote copy group, as shown in Figure 32.

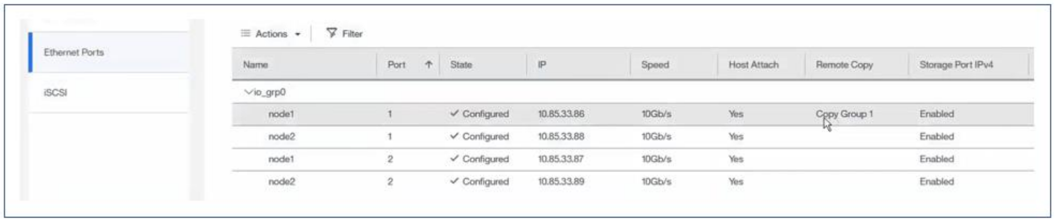

Name	Port	State	IP	Speed	Host Attach	Remote Copy	Storage Port IPv4
io_grp0							
node1	1	✓ Configured	10.85.33.86	10Gb/s	Yes	Copy Group 1	Enabled
node2	1	✓ Configured	10.85.33.88	10Gb/s	Yes		Enabled
node1	2	✓ Configured	10.85.33.87	10Gb/s	Yes		Enabled
node2	2	✓ Configured	10.85.33.89	10Gb/s	Yes		Enabled

Figure 32: Setting remote copy group

4. Repeat steps 1 and 2 for the on-premises IBM FlashSystem 9100 system.

5. Select **Copy services → Partnership → Create Partnership** as in Figure 33.

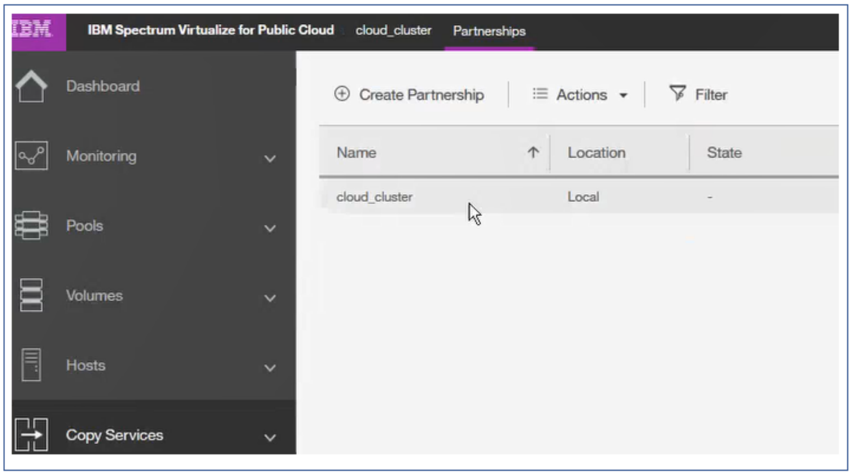

Name	Location	State
cloud_cluster	Local	-

Figure 33: Create partnership

6. Provide the cluster IP address of the remote system, as shown in Figure 34.

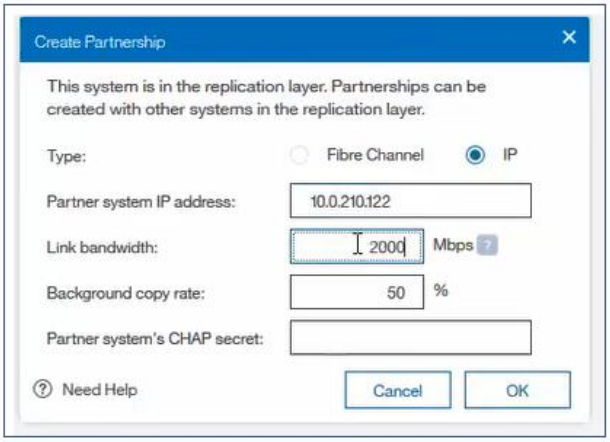

This system is in the replication layer. Partnerships can be created with other systems in the replication layer.

Type: ○ Fibre Channel ● IP

Partner system IP address: 10.0.210.122

Link bandwidth: 2000 Mbps

Background copy rate: 50 %

Partner system's CHAP secret:

Figure 34: Setting IP address for remote replication

7. This sets up the Global Mirror IP replication relationship between the on-premises IBM FlashSystem 9100 system and the IBM Spectrum Virtualize for Public Cloud cluster running in IBM Cloud.

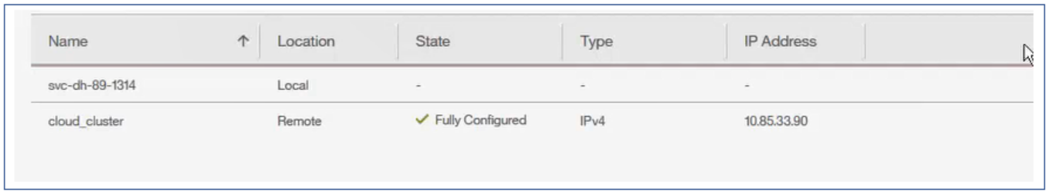

Name	Location	State	Type	IP Address
svc-dh-89-1314	Local	-	-	-
cloud_cluster	Remote	✓ Fully Configured	IPv4	10.85.33.90

Figure 35: IP replication

8. Once the Global Mirror relationship is created, set the consistency group and replicate the volumes for DR between the on-premises site and IBM Cloud as shown in Figure 36.

9. Go to **Copy Services → Remote Copy → Create Consistency Group**.

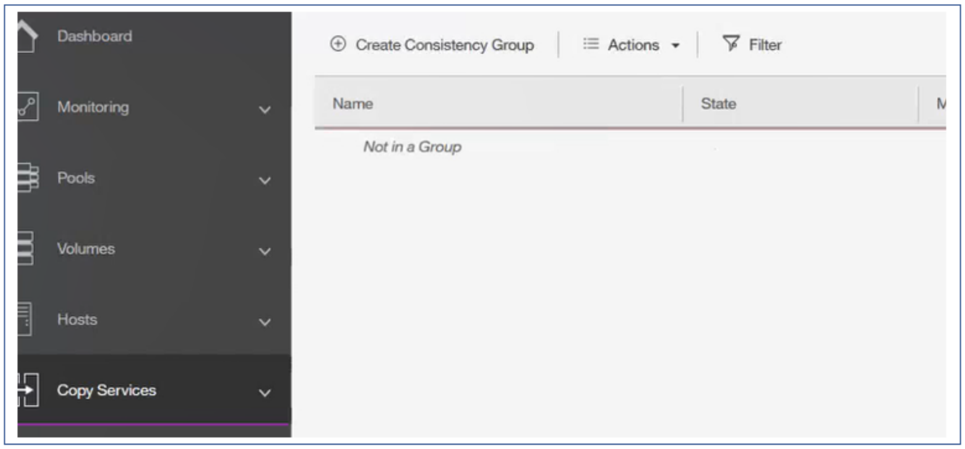

Figure 36: Create consistency group

10. Select the type of copy, as shown in Figure 37.

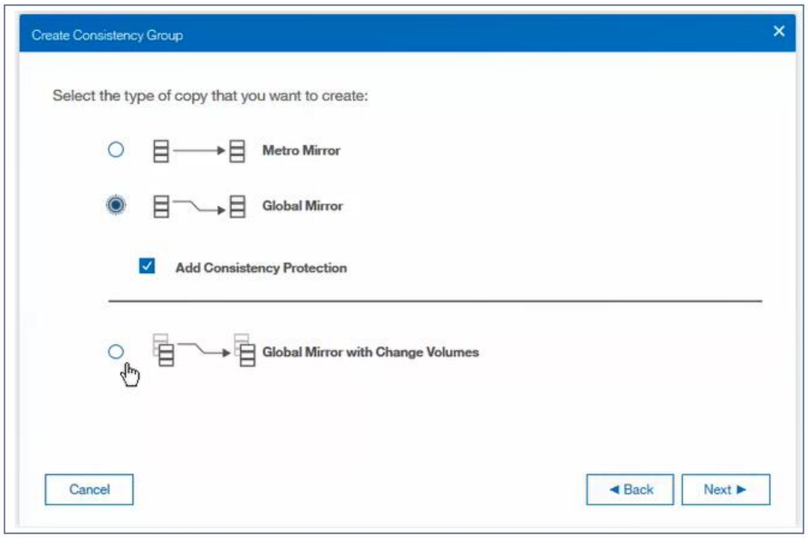

Create Consistency Group

Select the type of copy that you want to create:

○ Metro Mirror

● Global Mirror

☑ Add Consistency Protection

○ Global Mirror with Change Volumes

Figure 37: Select replication type

11. Add local and remote volume for the copy and begin copying.

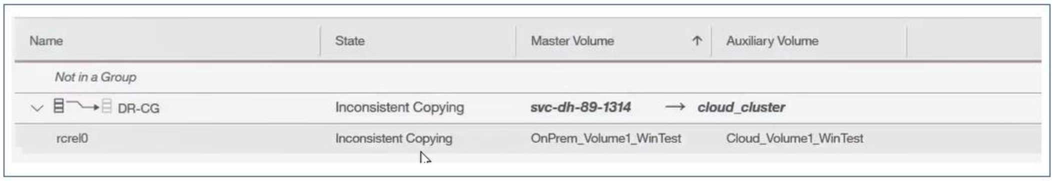

Name	State	Master Volume	Auxiliary Volume
Not in a Group			
DR-CG	Inconsistent Copying	***svc-dh-89-1314*** →	***cloud_cluster***
rcrel0	Inconsistent Copying	OnPrem_Volume1_WinTest	Cloud_Volume1_WinTest

Figure 38: Local and remote copy volume

VMware configuration

This section describes the essential details for setting up the Site Recovery Manager configuration.

For setting up Site Recovery Manager, the vCenter Server instances must initially be set up at the protected and recovery sites. Site Recovery Manager takes advantage of the vSphere ability to perform a vMotion migration across the sites and across the vCenter Server instances. Therefore, the two vCenter Server instances need to be configured.

For detailed installation of vCenter, refer to:
https://docs.vmware.com/en/VMware-vSphere/6.5/vsphere-esxi-vcenter-server-65-installation-setup-guide.pdf

The primary storage is configured on IBM FlashSystem 9100 at the on-premises site and volumes are configured and presented to the on-premises ESXi host.

Similarly, the replicated DR volume in IBM Cloud is presented to the ESXi host in IBM Cloud from IBM Spectrum Virtualize running in IBM Cloud.

Installation of Site Recovery Manager server at protected site

This section provides the high-level steps for installation of the Site Recovery Manager at the primary protected site.

Installation of Site Recovery Manager at the protected site requires the entry of the IP address or fully qualified domain name of the protected site’s VMware vSphere Platform Service controller during the registration of Site Recovery Manager.

Similarly, install and register Site Recovery Manager for the recovery site, which is in IBM Cloud. This requires the entry of the IP address or fully qualified domain name of the recovery site’s vSphere Platform Service controller.

For detailed installation steps, refer to the VMware Site Recovery Manager installation guide.

Installation of Storwize Family SRA

Storwize Family SRA is software that integrates with the Site Recovery Manager solution. Storwize Family SRA extends Site Recovery Manager capabilities to benefit the automation of the failover of storage systems at the protected Site Recovery Manager site to a recovery Site Recovery Manager site.

You can download a supported version of Storwize Family SRA from the VMware website at:
https://my.vmware.com/web/vmware/info/slug/infrastructure_operations_management/vmware_site_recovery_manager/8_1#drivers_tools

Storwize Family SRA works with a Site Recovery Manager instance at protected and recovery sites. Hence, it needs to be installed on the same server where the Site Recovery Manager instance has been installed at the protected and recovery sites.

For more information regarding Storwize Family SRA, refer to the Storwize Family SRA supported version release notes at:
https://www.ibm.com/support/knowledgecenter/en/SSEQ4E_3.3.0/UG/SVC_SRA_SRM5_1_install.html

Configuration of Site Recovery Manager

Site Recovery Manager can be configured using VMware vCenter Web Client. After Site Recovery Manager and Storwize Family SRA are installed successfully on both the sites, the Site Recovery plug-in will be visible on vCenter Web Client, as shown Figure 39.

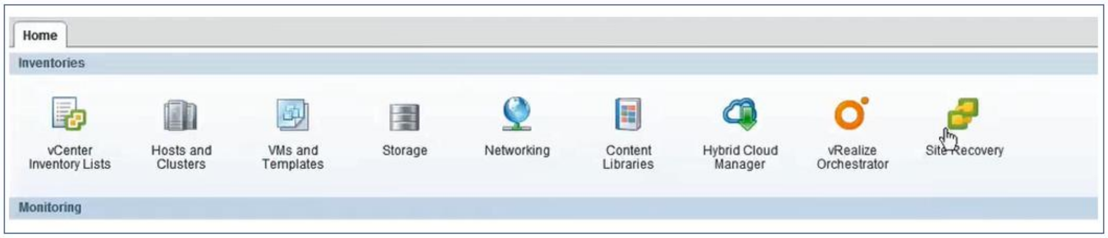

Figure 39: Site Recovery Manager plug-in

Site pairing

Before using Site Recovery Manager, the Site Recovery Manager instances at both sites should be connected to each other. To do site pairing, perform the following steps:

1. Log in to vCenter Web Client and click the **Site Recovery** icon, as shown in Figure 40.

2. Click **Sites**. Both sites—on-premises and cloud—will be visible.

3. Right-click on on-premises and click **Pair Site**. The pair site wizard is displayed.

4. Provide the Platform Services Controller address at the secondary site and select the vCenter Server with the registered Site Recovery Manager server instance at the cloud site.

5. When these steps are completed, the Platform Services Controller and vCenter Server are paired. Notice that the protected site now relates to the recovery site.

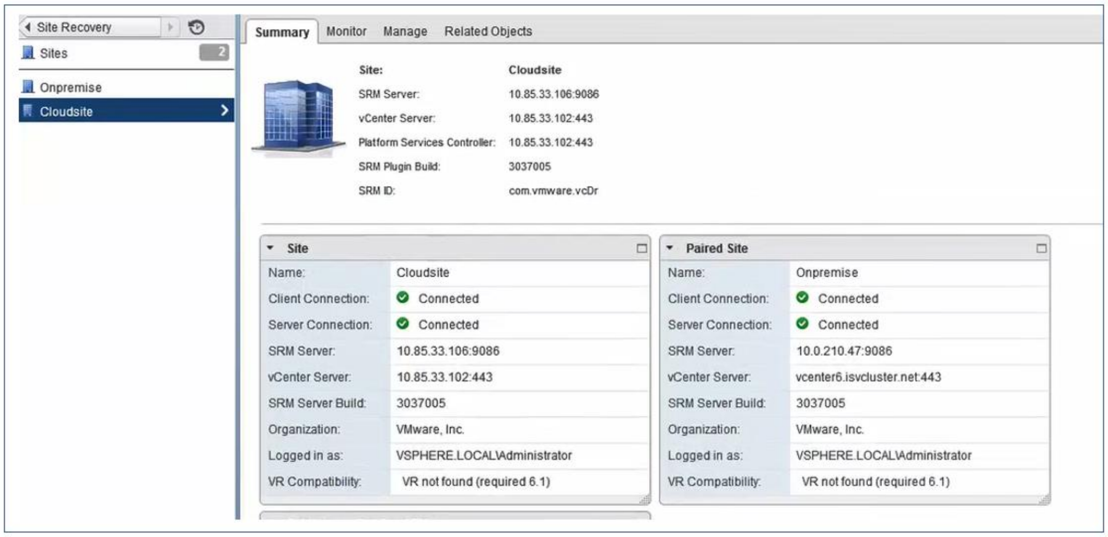

Site:	Cloudsite
SRM Server:	10.85.33.106:9086
vCenter Server:	10.85.33.102:443
Platform Services Controller:	10.85.33.102:443
SRM Plugin Build:	3037005
SRM ID:	com.vmware.vcDr

Site

Name:	Cloudsite
Client Connection:	✓ Connected
Server Connection:	✓ Connected
SRM Server:	10.85.33.106:9086
vCenter Server:	10.85.33.102:443
SRM Server Build:	3037005
Organization:	VMware, Inc.
Logged in as:	VSPHERE.LOCAL\Administrator
VR Compatibility:	VR not found (required 6.1)

Paired Site

Name:	Onpremise
Client Connection:	✓ Connected
Server Connection:	✓ Connected
SRM Server:	10.0.210.47:9086
vCenter Server:	vcenter6.isvcluster.net:443
SRM Server Build:	3037005
Organization:	VMware, Inc.
Logged in as:	VSPHERE.LOCAL\Administrator
VR Compatibility:	VR not found (required 6.1)

Figure 40: Site pairing

Configuration of array manager in Site Recovery Manager

The next step after site pairing is to configure the array manager. This requires adding one array manager for each site. It will enable the discovered array pairs for use with Site Recovery Manager. Perform the following steps to configure an array manager:

1. Click **Site Recovery → Array Based Replication**.

2. Click **Objects → Add Array Manager** and follow the wizard.

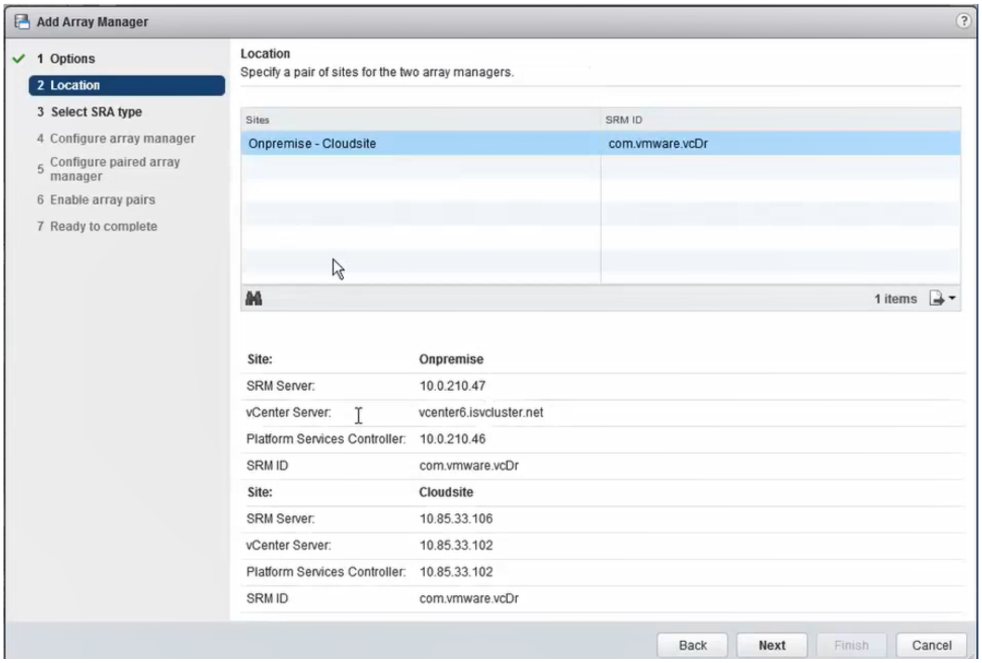

Location
Specify a pair of sites for the two array managers.

Sites	SRM ID
Onpremise - Cloudsite	com.vmware.vcDr

Site:	**Onpremise**
SRM Server:	10.0.210.47
vCenter Server:	vcenter6.isvcluster.net
Platform Services Controller:	10.0.210.46
SRM ID	com.vmware.vcDr
Site:	**Cloudsite**
SRM Server:	10.85.33.106
vCenter Server:	10.85.33.102
Platform Services Controller:	10.85.33.102
SRM ID	com.vmware.vcDr

Figure 41: Add array manager

After the configuration of the array manager, notice that the secondary site is visible, as shown in Figure 42.

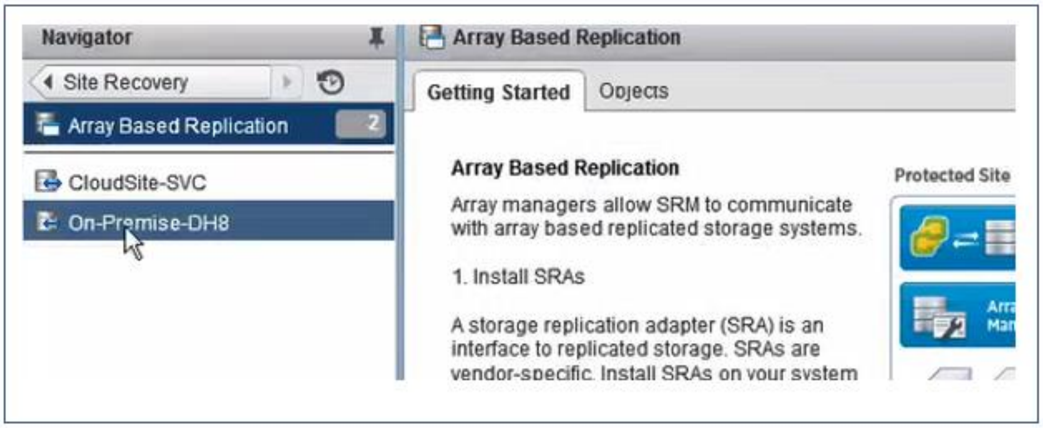

Array Based Replication

Array managers allow SRM to communicate with array based replicated storage systems.

1. Install SRAs

A storage replication adapter (SRA) is an interface to replicated storage. SRAs are vendor-specific. Install SRAs on your system

Figure 42: Array-based replication

Configuration of Site Recovery Manager resource mappings

Site Recovery Manager mapping configuration allows mapping of resources at the protected site to resources at the recovery site. During a recovery, when a VM starts at the recovery site, the VM uses the resources on the recovery site specified in the mappings. To enable bidirectional protection and to protect VMs and their resource information, reverse mappings can be configured to map the objects on the recovery site back to their corresponding objects on the protected site.

To complete the resource mapping, click **Site Recovery Sites**. On the Summary tab, go to “Guide to Configuring SRM,” as shown in Figure 43.

Configure bidirectional resource mappings, network mappings, folder mappings and placeholder datastore mappings between protected and recovery sites.

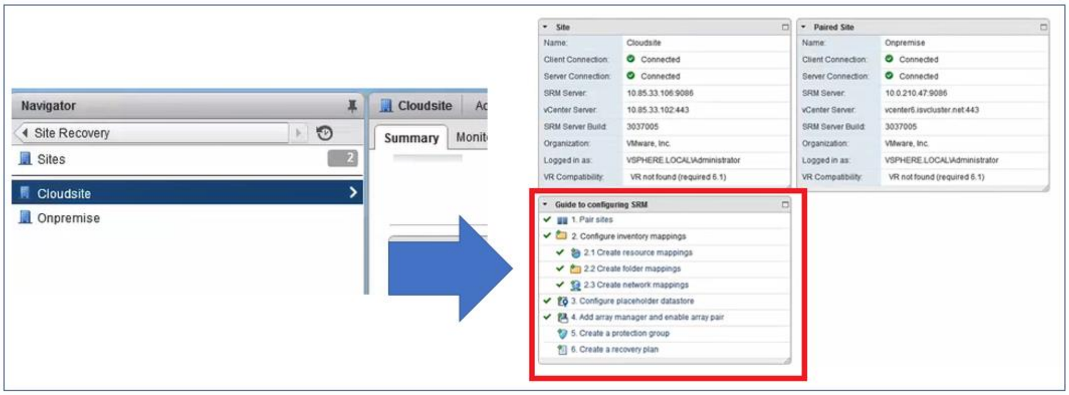

Site	
Name:	Cloudsite
Client Connection:	✓ Connected
Server Connection:	✓ Connected
SRM Server:	10.85.33.106:9086
vCenter Server:	10.85.33.102:443
SRM Server Build:	3037005
Organization:	VMware, Inc.
Logged in as:	VSPHERE.LOCAL\Administrator
VR Compatibility:	VR not found (required 6.1)

Paired Site	
Name:	Onpremise
Client Connection:	✓ Connected
Server Connection:	✓ Connected
SRM Server:	10.0.210.47:9086
vCenter Server:	vcenter6.isvcluster.net:443
SRM Server Build:	3037005
Organization:	VMware, Inc.
Logged in as:	VSPHERE.LOCAL\Administrator
VR Compatibility:	VR not found (required 6.1)

Guide to configuring SRM

- ✓ 1. Pair sites
- ✓ 2. Configure inventory mappings
  - ✓ 2.1 Create resource mappings
  - ✓ 2.2 Create folder mappings
  - ✓ 2.3 Create network mappings
- ✓ 3. Configure placeholder datastore
- ✓ 4. Add array manager and enable array pair
- 5. Create a protection group
- 6. Create a recovery plan

Figure 43: Site Recovery Manager configuration

Protection group configuration using vCenter Server

Using Site Recovery Manager, create the protection group with the protection type as datastore groups (array-based replication). The array-based replication protection groups automate the process of protecting and unprotecting the VMs depending on the data stores in the protection groups. As shown in Figure 44, follow the wizard to complete the creation of protection groups.

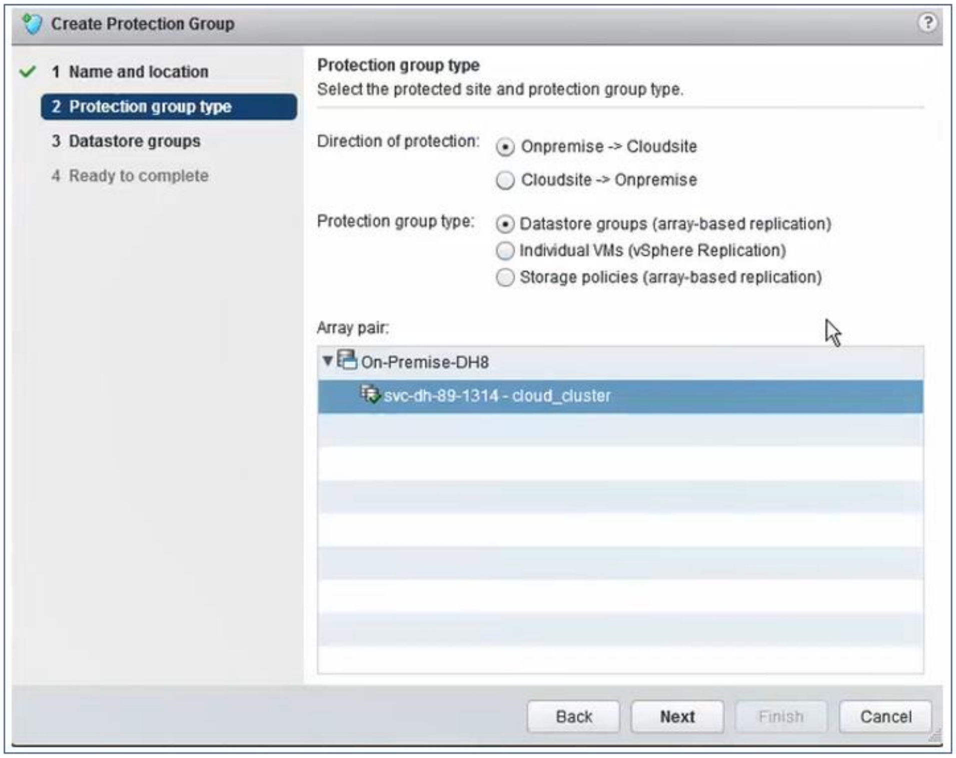

Create Protection Group

1 Name and location
2 Protection group type
3 Datastore groups
4 Ready to complete

Protection group type
Select the protected site and protection group type.

Direction of protection: ● Onpremise -> Cloudsite ○ Cloudsite -> Onpremise

Protection group type: ● Datastore groups (array-based replication) ○ Individual VMs (vSphere Replication) ○ Storage policies (array-based replication)

Array pair:
On-Premise-DH8
svc-dh-89-1314 - cloud_cluster

Figure 44: Create protection group

The following steps are required in the creation of protection group:

- Specify which sites will be protected and recovery sites
- Select the datastores that need to be protected in the datastore groups
- Review the protection group settings and click **Finish** to create a storage policy-based protection group

In the lab environment, the protection group is created to protect a Windows VM which has Microsoft SQL Server Express database running.

The datastore group has three volumes, as shown in Figure 45:

- The C drive, which is volume one, has the VM’s Virtual Machine Disk (VMDK) files.
- The E drive, which is volume two, has an SQL database running.
- The F drive has log files.

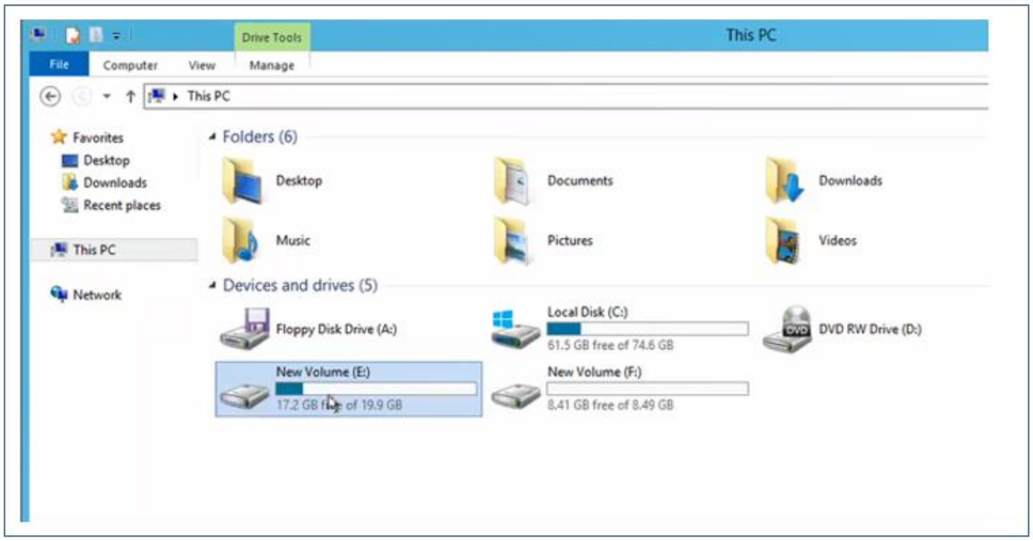

Figure 45: Test VM disk layout

A test database (testDB) is created at the on-premises site, as shown in Figure 46.

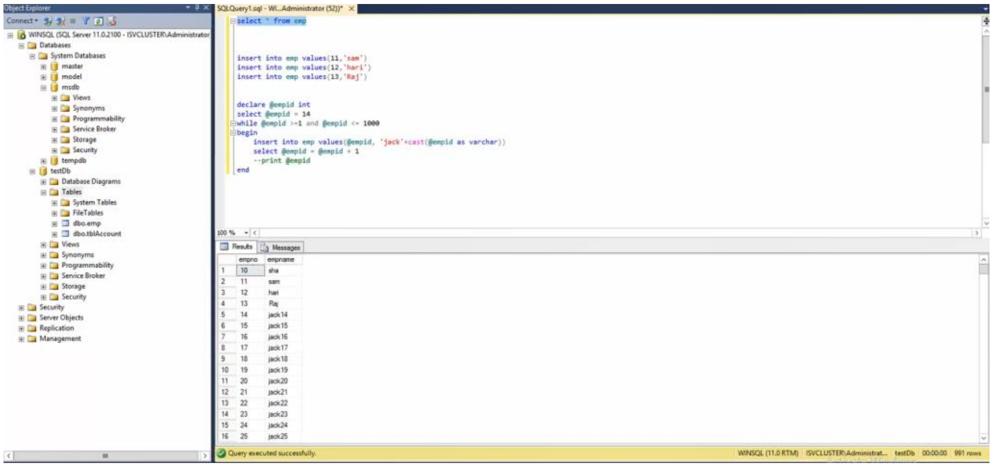

```
select * from emp



insert into emp values(11,'sam')
insert into emp values(12,'hari')
insert into emp values(13,'Raj')


declare @empid int
select @empid = 14
while @empid >=1 and @empid <= 1000
begin
    insert into emp values(@empid, 'jack'+cast(@empid as varchar))
    select @empid = @empid + 1
    --print @empid
end
```

	empno	empname
1	10	sha
2	11	sam
3	12	hari
4	13	Raj
5	14	jack14
6	15	jack15
7	16	jack16
8	17	jack17
9	18	jack18
10	19	jack19
11	20	jack20
12	21	jack21
13	22	jack22
14	23	jack23
15	24	jack24
16	25	jack25

Query executed successfully. WINSQL (11.0 RTM) ISVCLUSTER\Administrat... testDb 00:00:00 991 rows

Figure 46: testDB row insertion

The next step is to create a recovery plan for failover and failback of the VM running an SQL database.

Creating a recovery plan

In Site Recovery Manager, a recovery plan is a set of automated rules. It controls the steps of the recovery process, including storage replication and data synchronization, network configuration at the secondary site, and spin-up of the VM at the secondary site. A recovery plan can address a planned recovery or an unplanned recovery.

To create a recovery plan in vCenter Web Client, click **Site Recovery**. Then click **Recovery Plans → Create a Recovery Plan**.

Follow the wizard to create a recovery plan by specifying the recovery plan name, recovery site and protection group that were created earlier.

As shown in Figure 47, the wizard displays the protected site and recovery site vCenter details and the protection group associated with the plan.

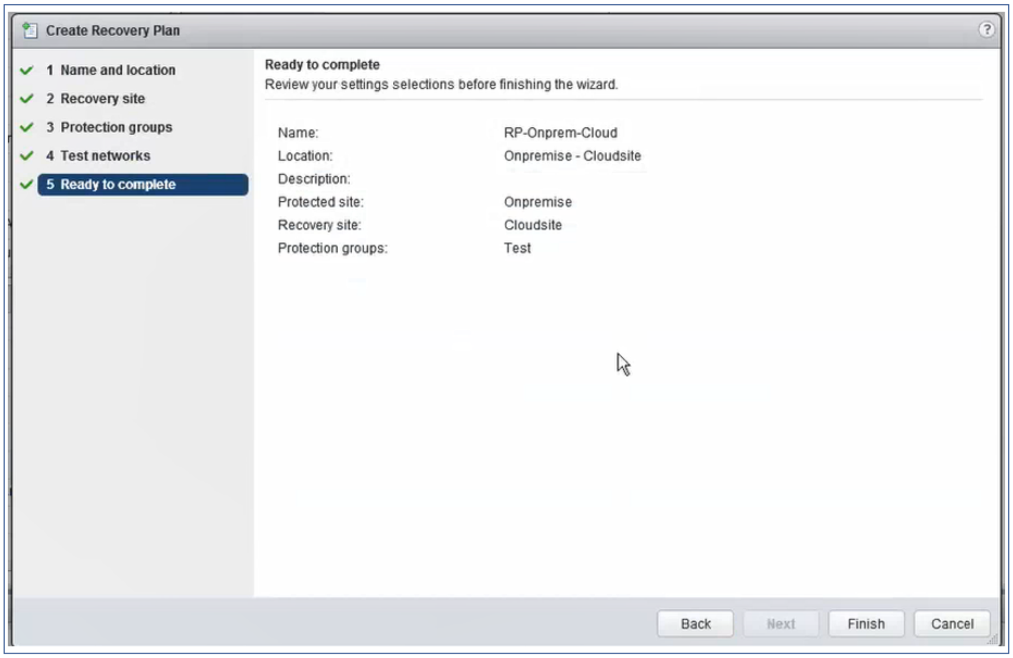

Create Recovery Plan

- ✓ 1 Name and location
- ✓ 2 Recovery site
- ✓ 3 Protection groups
- ✓ 4 Test networks
- ✓ 5 Ready to complete

Ready to complete
Review your settings selections before finishing the wizard.

Name:	RP-Onprem-Cloud
Location:	Onpremise - Cloudsite
Description:	
Protected site:	Onpremise
Recovery site:	Cloudsite
Protection groups:	Test

Figure 47: The Create Recovery Plan screen

The recovery plan is now ready, and you can view that the status is Ready (as shown in Figure 48).

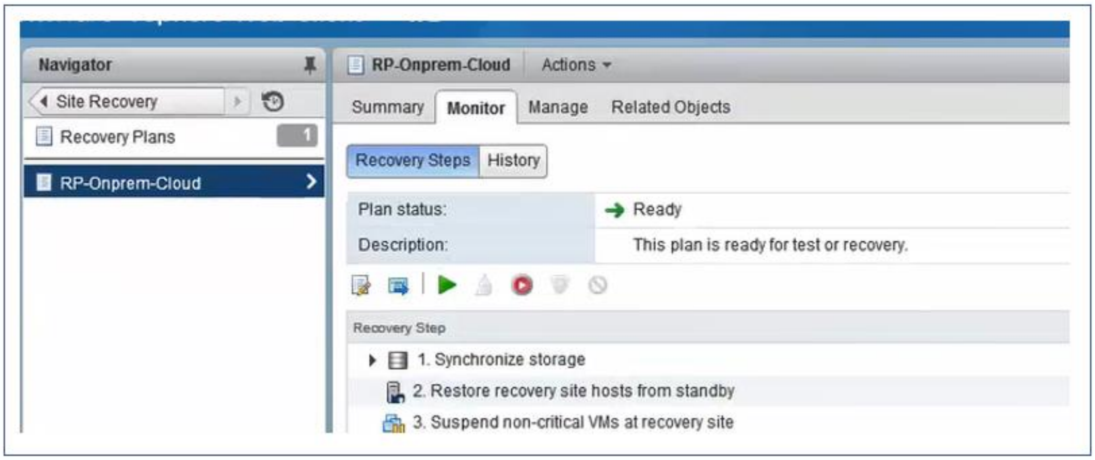

Plan status:	Ready
Description:	This plan is ready for test or recovery.

Recovery Step

1. Synchronize storage
2. Restore recovery site hosts from standby
3. Suspend non-critical VMs at recovery site

Figure 48: Configured recovery plan

The next step is to set up the IP address properties for the VM once failed over to the cloud site.

To set up the IP address, click on **Recovery Plan → Related Objects → Virtual Machines**, as shown in Figure 49. The VM that is protected under this particular recovery plan is visible. Right-click the name of the VM → **Edit properties**, then select **Manual IP Customization** from the drop-down menu, as shown in Figure 49.

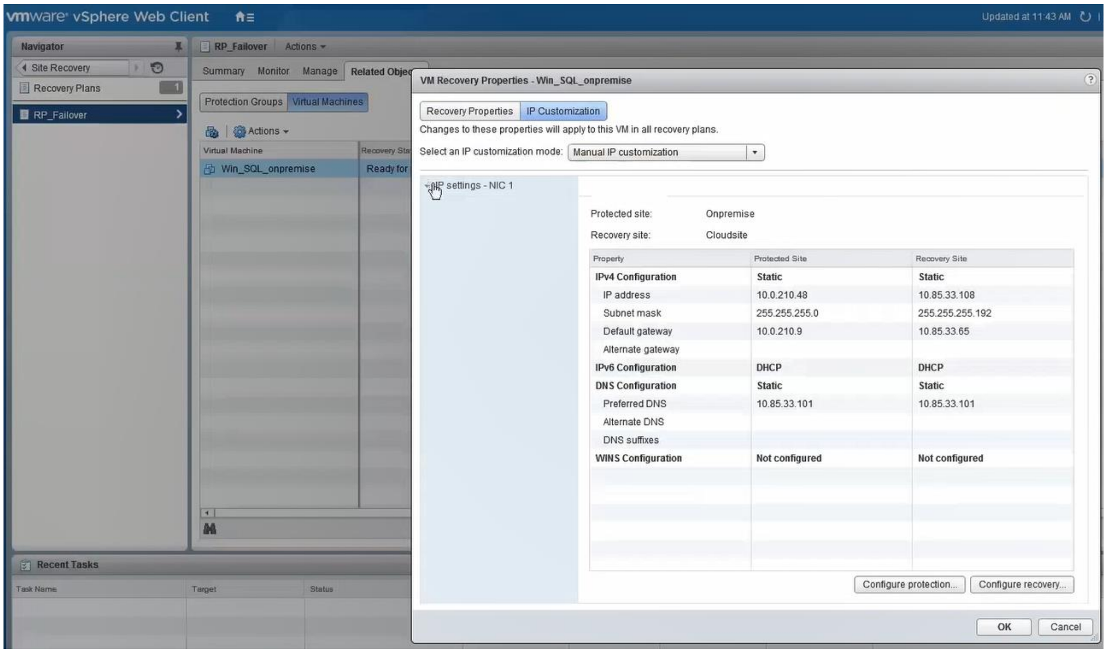
vmware vSphere Web Client
Updated at 11:43 AM
Navigator
Site Recovery
Recovery Plans 1
RP_Failover
RP_Failover Actions
Summary Monitor Manage Related Objec
Protection Groups Virtual Machines
Actions
Recent Tasks
Task Name Target Status
Recovery Properties IP Customization
IP settings - NIC 1
Configure protection... Configure recovery...
OK Cancel

VM Recovery Properties - Win_SQL_onpremise

Changes to these properties will apply to this VM in all recovery plans.

Select an IP customization mode: Manual IP customization

Virtual Machine	Recovery Sta
Win_SQL_onpremise	Ready for

Protected site: Onpremise

Recovery site: Cloudsite

Property	Protected Site	Recovery Site
IPv4 Configuration	**Static**	**Static**
IP address	10.0.210.48	10.85.33.108
Subnet mask	255.255.255.0	255.255.255.192
Default gateway	10.0.210.9	10.85.33.65
Alternate gateway		
IPv6 Configuration	**DHCP**	**DHCP**
DNS Configuration	**Static**	**Static**
Preferred DNS	10.85.33.101	10.85.33.101
Alternate DNS		
DNS suffixes		
WINS Configuration	**Not configured**	**Not configured**

Figure 49: Setting IP customization

Update the static IP address that needs to be assigned to the VM after failover and failback between on-premises and IBM Cloud sites.

Testing a recovery plan

After a recovery plan is created, you can test the recovery of VMs in a protection group to ensure that the VMs are correctly recovered at the recovery sites. This solution uses the array-based replication of Site Recovery Manager. A snapshot of the volume hosting the VMs' disk files will be created while testing the recovery plan. Snapshots will be removed after running a cleanup operation.

For testing the recovery plan, click **Test Recovery Plan**.

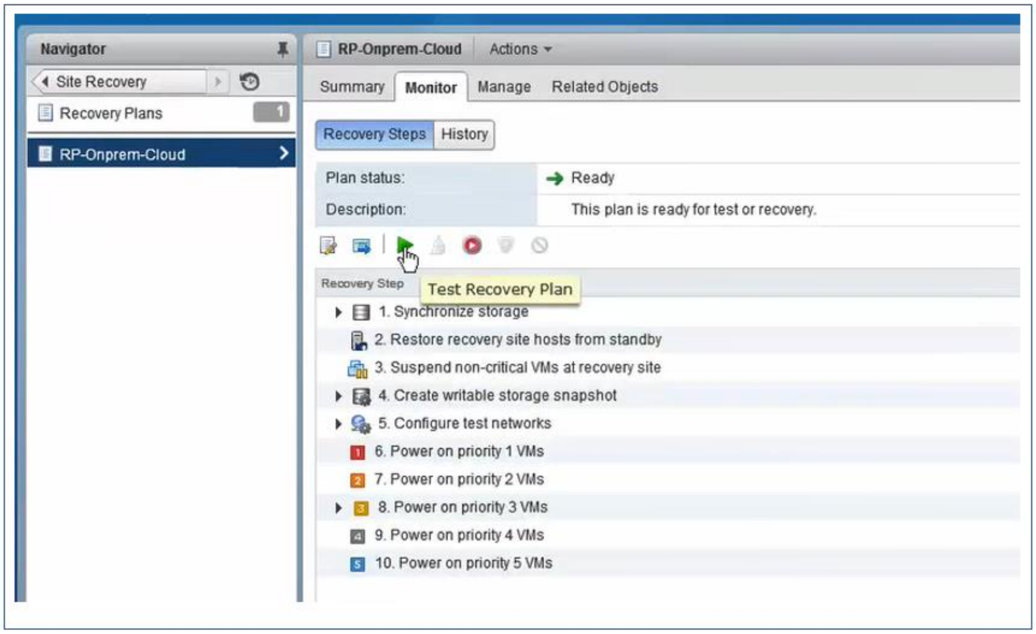

Plan status: Ready

Description: This plan is ready for test or recovery.

Recovery Step

1. Synchronize storage
2. Restore recovery site hosts from standby
3. Suspend non-critical VMs at recovery site
4. Create writable storage snapshot
5. Configure test networks
6. Power on priority 1 VMs
7. Power on priority 2 VMs
8. Power on priority 3 VMs
9. Power on priority 4 VMs
10. Power on priority 5 VMs

Figure 50: Test recovery plan

As a part of test recovery, the recovery plan runs the following steps:

- Synchronize the storage volumes that are part of the recovery plan
- Create a writeable snapshot at the secondary site
- Configure the test network
- Fail over the VM at the primary site and power on the VM at the secondary site

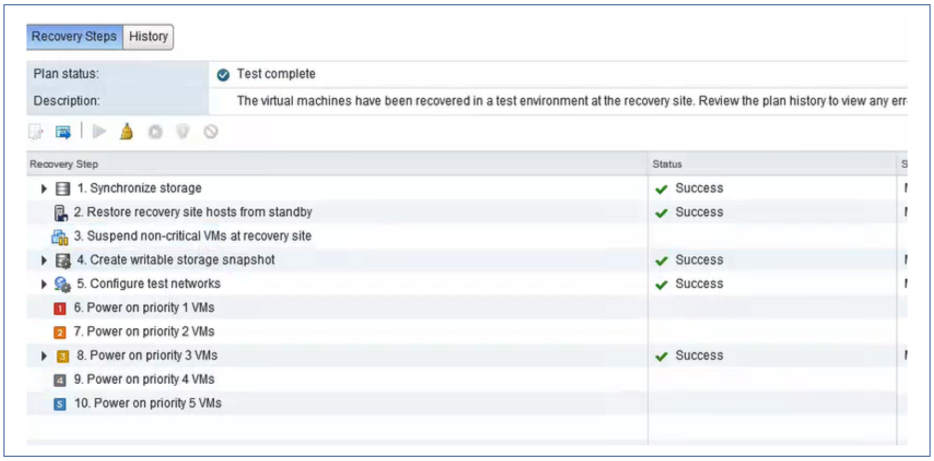

Plan status: Test complete

Description: The virtual machines have been recovered in a test environment at the recovery site. Review the plan history to view any err

Recovery Step	Status
1. Synchronize storage	✓ Success
2. Restore recovery site hosts from standby	✓ Success
3. Suspend non-critical VMs at recovery site	
4. Create writable storage snapshot	✓ Success
5. Configure test networks	✓ Success
6. Power on priority 1 VMs	
7. Power on priority 2 VMs	
8. Power on priority 3 VMs	✓ Success
9. Power on priority 4 VMs	
10. Power on priority 5 VMs	

Figure 51: Steps performed during test of recovery plan

After successfully running a test recovery plan, run the cleanup recovery plan to return the recovery plan to the Ready state. To clean up after testing the recovery plan, click **Cleanup Recovery Plan**.

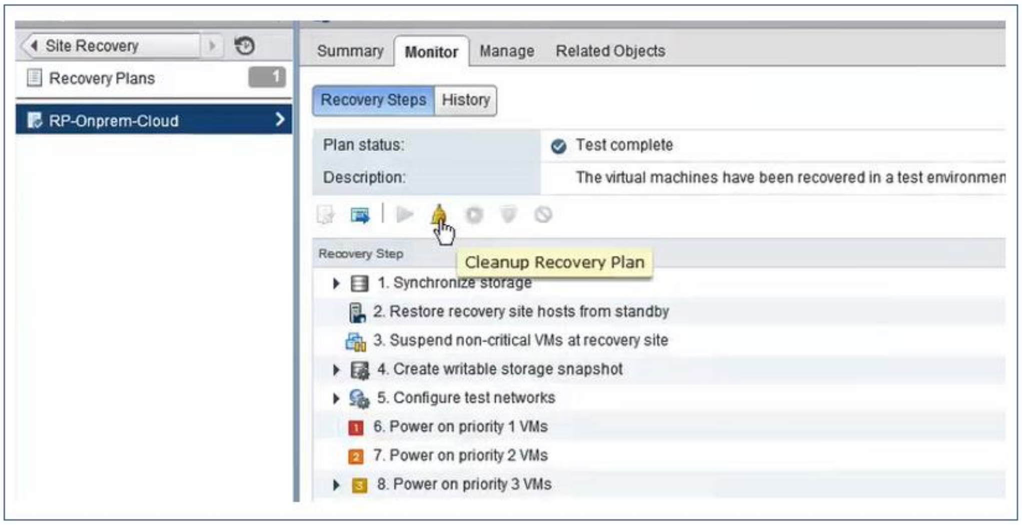

Plan status: Test complete

Description: The virtual machines have been recovered in a test environmen

Recovery Step

1. Synchronize storage
2. Restore recovery site hosts from standby
3. Suspend non-critical VMs at recovery site
4. Create writable storage snapshot
5. Configure test networks
6. Power on priority 1 VMs
7. Power on priority 2 VMs
8. Power on priority 3 VMs

Figure 52: Clean up the recovery plan

Failover using a recovery plan

To do the actual failover of the VM using the recovery plan, click **Run Recovery Plan**.

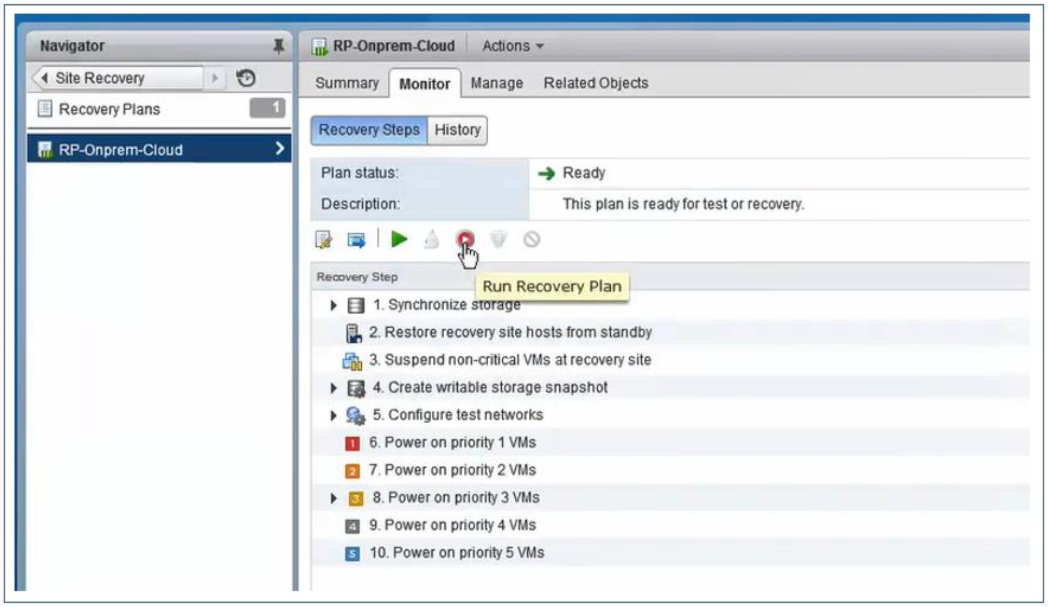

Plan status: Ready

Description: This plan is ready for test or recovery.

1. Synchronize storage
2. Restore recovery site hosts from standby
3. Suspend non-critical VMs at recovery site
4. Create writable storage snapshot
5. Configure test networks
6. Power on priority 1 VMs
7. Power on priority 2 VMs
8. Power on priority 3 VMs
9. Power on priority 4 VMs
10. Power on priority 5 VMs

Figure 53: Run actual recovery plan

The recovery plan has two options: Planned migration and DR.

In planned migration, the recovery of VMs happens when both the sites are running, and—as the name suggests—it is a planned activity.

In the case of DR, the recovery of VMs at the recovery site is done only when the protected site experiences a disaster.

Select the recovery type and click **Finish** to run the recovery plan.

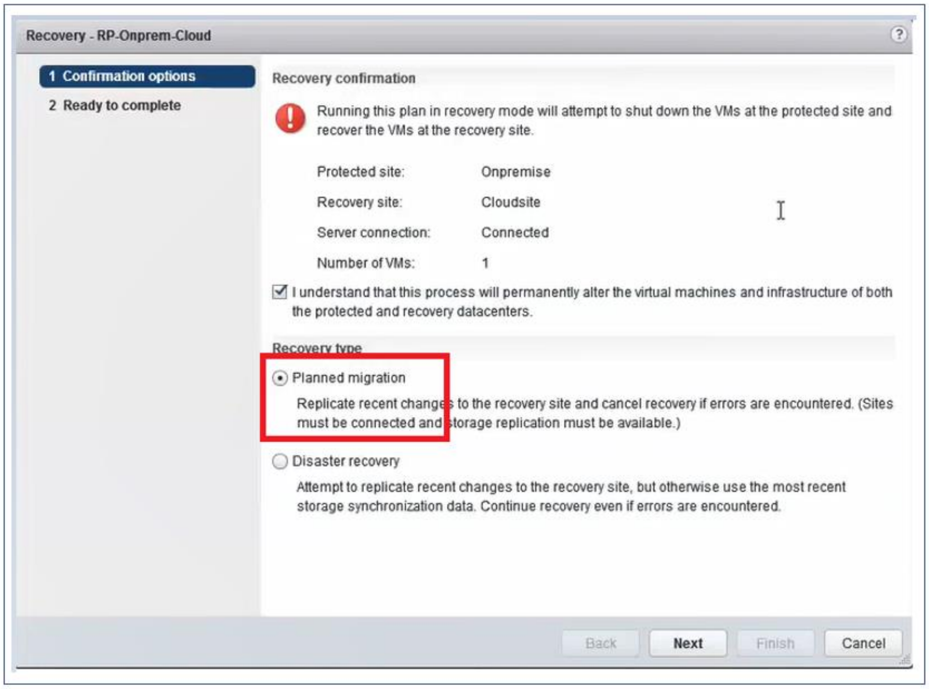

Recovery - RP-Onprem-Cloud

1 Confirmation options

2 Ready to complete

Recovery confirmation

Running this plan in recovery mode will attempt to shut down the VMs at the protected site and recover the VMs at the recovery site.

Protected site:	Onpremise
Recovery site:	Cloudsite
Server connection:	Connected
Number of VMs:	1

☑ I understand that this process will permanently alter the virtual machines and infrastructure of both the protected and recovery datacenters.

Recovery type

● Planned migration

Replicate recent changes to the recovery site and cancel recovery if errors are encountered. (Sites must be connected and storage replication must be available.)

○ Disaster recovery

Attempt to replicate recent changes to the recovery site, but otherwise use the most recent storage synchronization data. Continue recovery even if errors are encountered.

Figure 54: Recovery plan wizard

Once recovery is completed, the VM will be powered on automatically with the IP address `10.85.33.108` as defined in the recovery plan at the cloud site.

Log in to the VM at the cloud site, and then connect to the VM and check the consistency of the testDB, as shown in Figure 55.

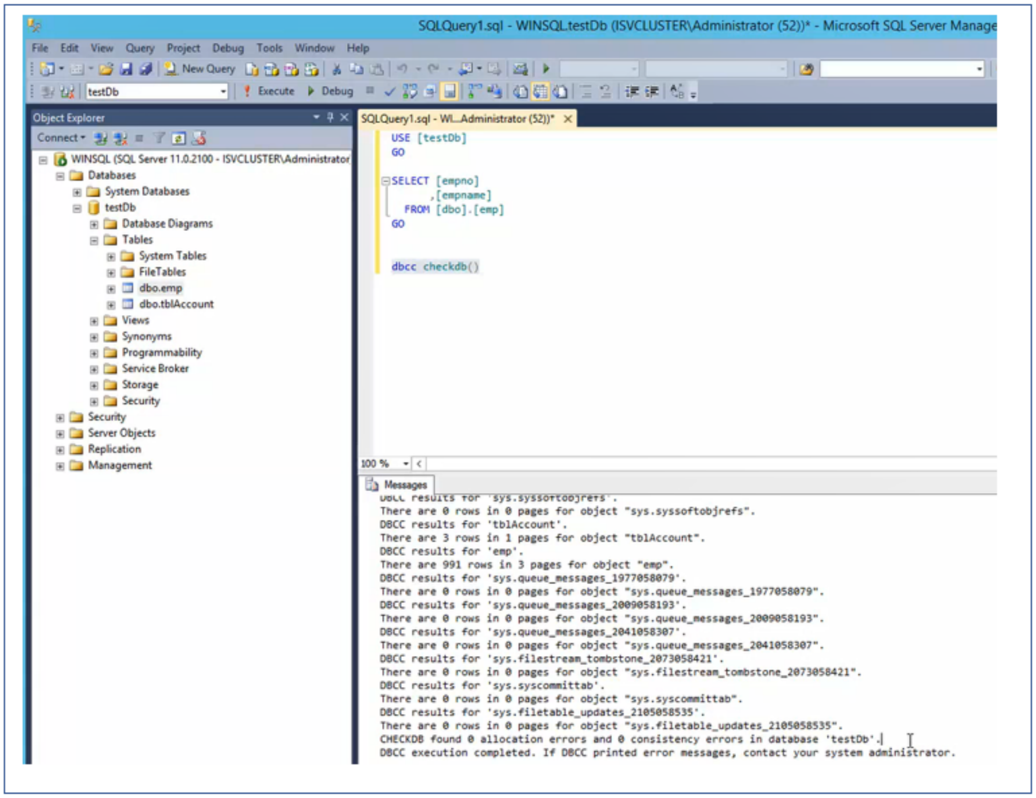

```
USE [testDb]
GO

SELECT [empno]
      ,[empname]
  FROM [dbo].[emp]
GO


dbcc checkdb()
```

```
DBCC results for 'sys.syssoftobjrefs'.
There are 0 rows in 0 pages for object "sys.syssoftobjrefs".
DBCC results for 'tblAccount'.
There are 3 rows in 1 pages for object "tblAccount".
DBCC results for 'emp'.
There are 991 rows in 3 pages for object "emp".
DBCC results for 'sys.queue_messages_1977058079'.
There are 0 rows in 0 pages for object "sys.queue_messages_1977058079".
DBCC results for 'sys.queue_messages_2009058193'.
There are 0 rows in 0 pages for object "sys.queue_messages_2009058193".
DBCC results for 'sys.queue_messages_2041058307'.
There are 0 rows in 0 pages for object "sys.queue_messages_2041058307".
DBCC results for 'sys.filestream_tombstone_2073058421'.
There are 0 rows in 0 pages for object "sys.filestream_tombstone_2073058421".
DBCC results for 'sys.syscommittab'.
There are 0 rows in 0 pages for object "sys.syscommittab".
DBCC results for 'sys.filetable_updates_2105058535'.
There are 0 rows in 0 pages for object "sys.filetable_updates_2105058535".
CHECKDB found 0 allocation errors and 0 consistency errors in database 'testDb'.
DBCC execution completed. If DBCC printed error messages, contact your system administrator.
```

Figure 55: testDB consistency check

After the recovery is completed, the recovery site becomes the primary site. The VM that is protected under the recovery plan is failed over to, and turned on at, the recovery IBM Cloud site. As shown in the above figure, now the database is on and running at the IBM Cloud recovery site.

However, the VM running at the recovery site is not protected, so any changes made to the database or application will not be protected. To protect the VM and any changes made at the recovery IBM Cloud site, click **Reprotect Recovery Plan**.

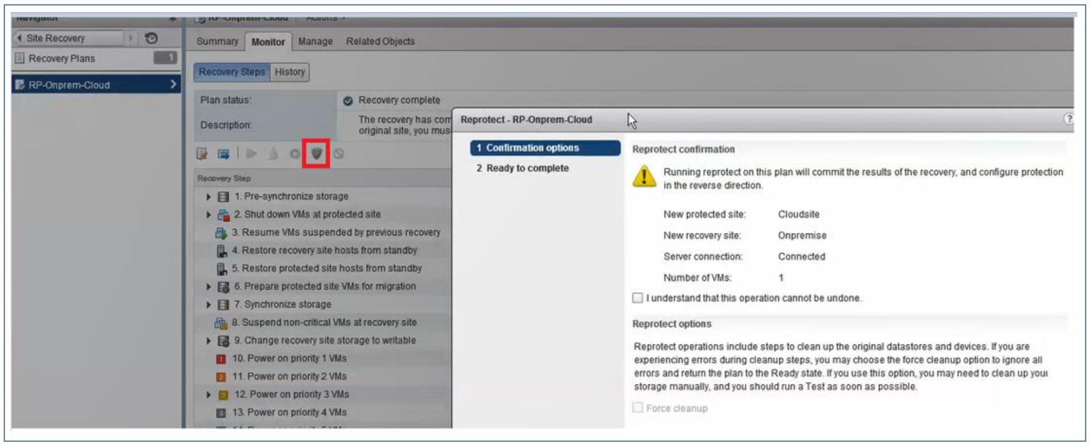

Plan status: Recovery complete

Description: The recovery has com original site, you mus

1. Pre-synchronize storage
2. Shut down VMs at protected site
3. Resume VMs suspended by previous recovery
4. Restore recovery site hosts from standby
5. Restore protected site hosts from standby
6. Prepare protected site VMs for migration
7. Synchronize storage
8. Suspend non-critical VMs at recovery site
9. Change recovery site storage to writable
10. Power on priority 1 VMs
11. Power on priority 2 VMs
12. Power on priority 3 VMs
13. Power on priority 4 VMs

Reprotect - RP-Onprem-Cloud

Reprotect confirmation

Running reprotect on this plan will commit the results of the recovery, and configure protection in the reverse direction.

New protected site: Cloudsite
New recovery site: Onpremise
Server connection: Connected
Number of VMs: 1

□ I understand that this operation cannot be undone.

Reprotect options

Reprotect operations include steps to clean up the original datastores and devices. If you are experiencing errors during cleanup steps, you may choose the force cleanup option to ignore all errors and return the plan to the Ready state. If you use this option, you may need to clean up youi storage manually, and you should run a Test as soon as possible.

□ Force cleanup

Figure 56: Reprotect Recovery Plan

After successfully running the reprotect recovery plan, the VM running at the secondary site is protected using the recovery plan. It is also possible to failback the VM back to the original on-premises site using the same recovery plan.

IBM Spectrum Copy Data Management for data re-use

Organizations are also looking for cost-efficient ways to manage their infrastructure or operations, which includes DR, managing copies of data, provisioning of development and testing infrastructure, and deploying processes and technology for DevOps.

The IBM Spectrum Copy Data Management platform, in conjunction with IBM Spectrum Storage systems, enables critical use cases by providing in-place copy data management to modernize processes within existing infrastructure.

IBM Spectrum Copy Data Management allows businesses to manage, orchestrate and analyze copy data. The solution provides full lifecycle management through automated workflows that allow you to streamline the creation, management and use of copies of data.

IBM Spectrum Copy Data Management simplifies management of copies by enabling administrators to orchestrate application-consistent copy creation, and to recover or clone data in minutes instead of hours or days.

The focus of this section is on IBM Spectrum Storage with IBM Spectrum Copy Data Management in existing storage environments to leverage use cases such as automated DR, provisioning of automated on-demand test/development environment for VMware environments by orchestrating and modernizing data copy use. The high-level architecture diagram is as shown in Figure 57.

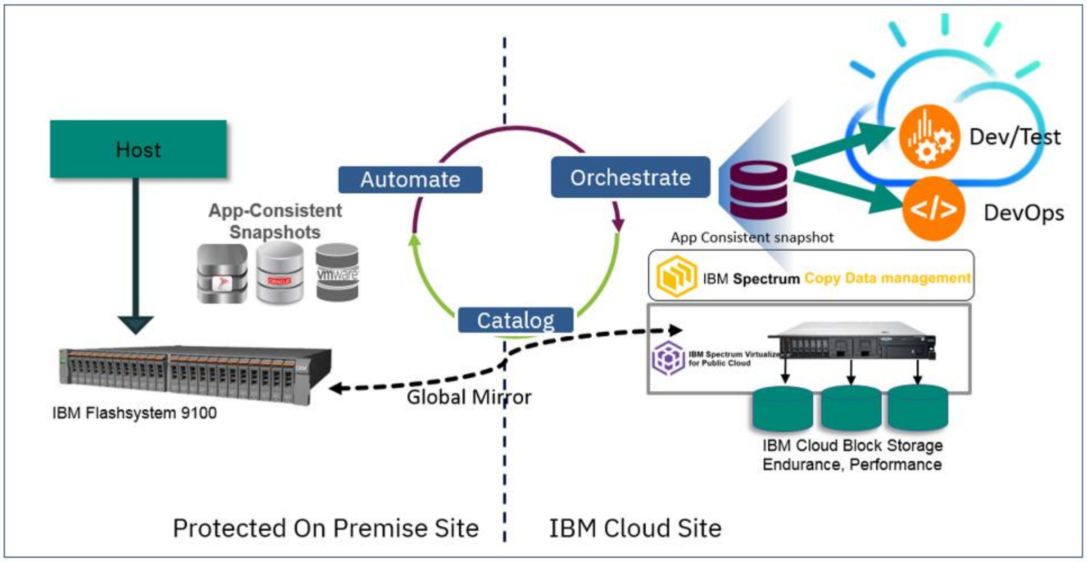
Host
App-Consistent Snapshots
Automate
Orchestrate
Catalog
Dev/Test
DevOps
App Consistent snapshot
IBM Spectrum Copy Data management
IBM Spectrum Virtualize for Public Cloud
IBM Flashsystem 9100
Global Mirror
IBM Cloud Block Storage Endurance, Performance
Protected On Premise Site
IBM Cloud Site

Figure 57: High-level architecture diagram for data re-use

The lab setup is deployed to leverage the use case of automated DR using IBM Spectrum Copy Data Management with the help of the underlying Global Mirror with change replication feature of IBM Spectrum Virtualize storage.

The same environment can be used to deploy the on-demand test/development VM environment with the refreshed copy of the production volume at the DR site in IBM Cloud using IBM Spectrum Copy Data Management and the FlashCopy feature in IBM Spectrum Virtualize.

There are multiple use cases for data re-use that can be addressed using IBM Spectrum Copy Data Management using workflows. For example:

- Implementing Global Mirror with change volume replication between an on-premises site and an IBM Cloud site
- Creating snapshots at storage level and re-using the volumes for test/dev purposes
- Creating application-consistent snapshots for supported platforms and applications such as VMware, Microsoft SQL Server and Oracle

For more supported platform for IBM Spectrum Copy Data Management please refer to: https://www.ibm.com/support/knowledgecenter/SS57AN_2.2.7/com.ibm.spectrum.cdm.doc/prd_whats_new_history.html

To demonstrate data re-use, the lab setup is deployed as explained below.

The primary and DR sites are deployed with IBM Spectrum Virtualize storage, where Global Mirror with Change Volumes is used for DR. Snapshots are created at regular intervals with the refreshed data from the production site, and the test VM is powered on using the snapshot volumes.

IBM Spectrum Copy Data Management is used to automate copy management and scheduling. It also enables on-demand access to data copies, and allows users to map and mount snapshots in minutes using management workflows.

All the infrastructure provisioning tasks as shown in Figure 58 that are performed manually are automated, providing significant reduction in infrastructure deployment time.

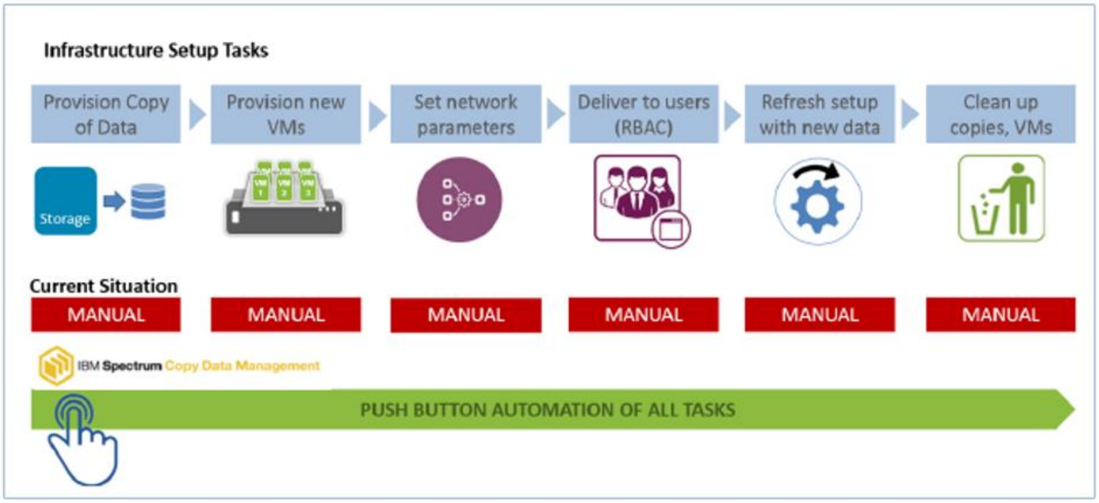

Figure 58: Infrastructure set-up tasks

Installation and configuration of IBM Spectrum Copy Data Management

IBM Spectrum Copy Data Management can be deployed within a VMware infrastructure. Use the vCenter Web Client to deploy IBM Spectrum Copy Data Management. From the **File** menu, choose **Deploy OVF Template**. If using the vCenter Web Client, click **Create/Register VM**, then select **Deploy a virtual machine from an OVF or OVA file**.

IBM Spectrum Copy Data Management comes pre-packaged with all the required software, and once powered on, the console screen points to the portal link.

Log in to the portal using a web browser: https://<hostname>:8443/portal/

The login screen appears:

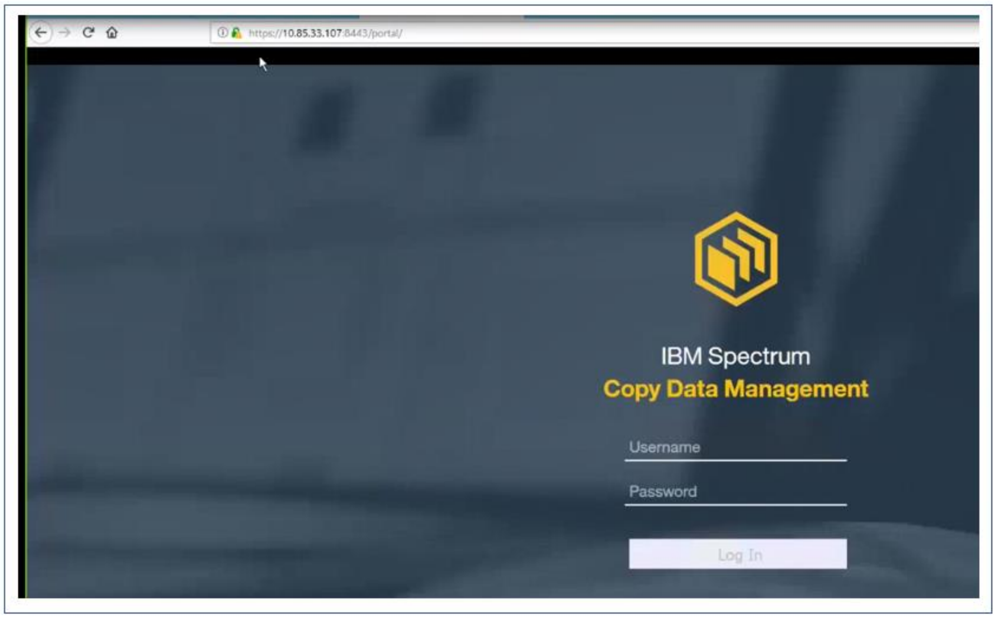

Figure 59: IBM Spectrum Copy Data Management login screen

The first step is to register the storage objects and vCenter Server. This is a one-time, agentless registration process. To register a new object, right-click the object and click **Register**, as shown in Figure 60, and then fill in the details.

IBM Spectrum Copy Data Management uses the concept of sites to identify resource locations. The registration dialog box accepts the credentials and site selection.

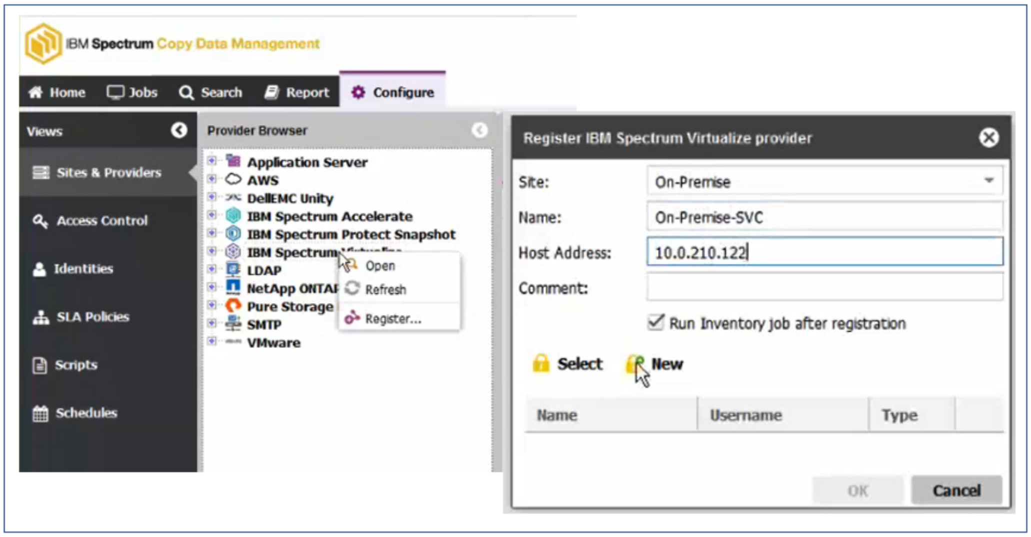

Register IBM Spectrum Virtualize provider

Site: On-Premise

Name: On-Premise-SVC

Host Address: 10.0.210.122

Comment:

☑ Run Inventory job after registration

Name	Username	Type

Figure 60: Register the objects

Similarly, register the vCenter Server as shown in Figure 61.

Provide the vCenter user name and password.

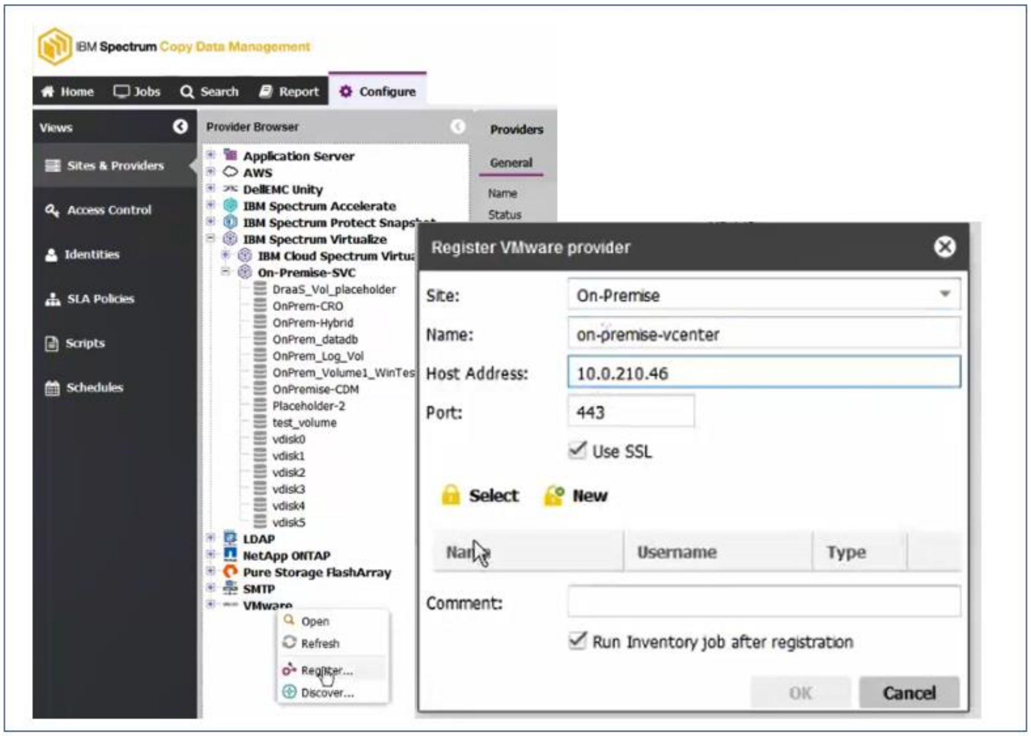

Register VMware provider

Site: On-Premise

Name: on-premise-vcenter

Host Address: 10.0.210.46

Port: 443

☑ Use SSL

Name	Username	Type

Comment:

☑ Run Inventory job after registration

Figure 61: Register vCenter

After the resources are registered, IBM Spectrum Copy Data Management automatically creates a default catalog policy. This catalog policy can discover high-level objects, such as storage volumes and MDisk information in storage arrays, and can discover high-level VMware information such as ESXi hosts managed by vCenter Server, data stores, VM details and more.

IBM Spectrum Copy Data Management policy creation

After registering the resources, the next step is to create a Service Level Agreement (SLA) policy, which will establish a storage workflow for Global Mirror and FlashCopy. To create a policy, click on **Configure → SLA Policies → New SLA Policy → IBM Spectrum Virtualize** as shown in Figure 62.

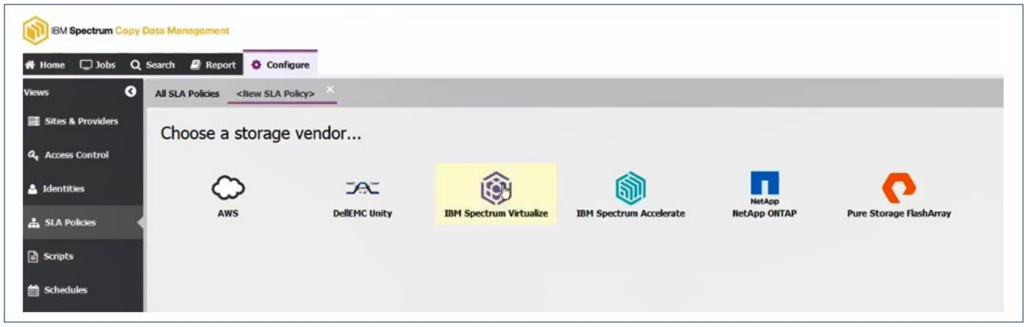

Choose a storage vendor...

Figure 62: Create SLA policies

Click **Add Global Mirror with Change Volumes** and follow the storage workflow wizard as shown in Figure 63.

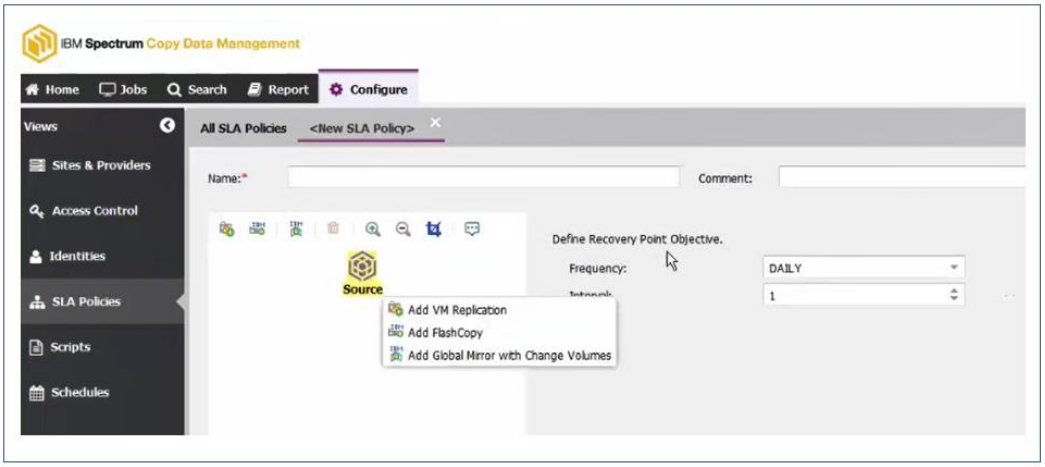

Name:*

Comment:

Define Recovery Point Objective.

Frequency: DAILY

Interval: 1

Figure 63: Create Global Mirror with Change Volumes policy

Similarly, create a FlashCopy storage workflow.

Backup job

To create the backup policy for Global Mirror with Change Volumes with IBM Spectrum Copy Data Management, click → **Jobs → Storage Controller → IBM Spectrum Virtualize → Backup**.

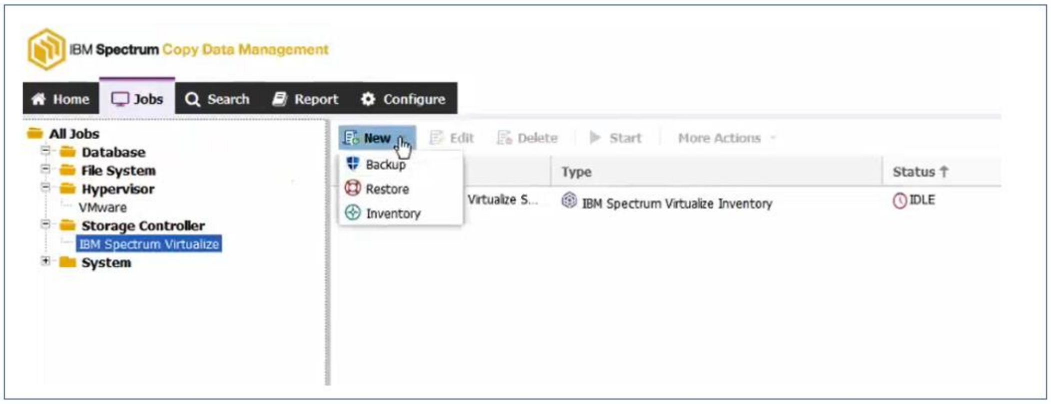

	Type	Status ↑
Virtualize S...	IBM Spectrum Virtualize Inventory	IDLE

Figure 64: Create backup job

Follow the wizard, select the source volume (in this lab environment, the source volume is an on-premises site and the target is in IBM Cloud, as shown in Figure 65), and associate the SLA policy created for Global Mirror with Change Volumes with the backup job.

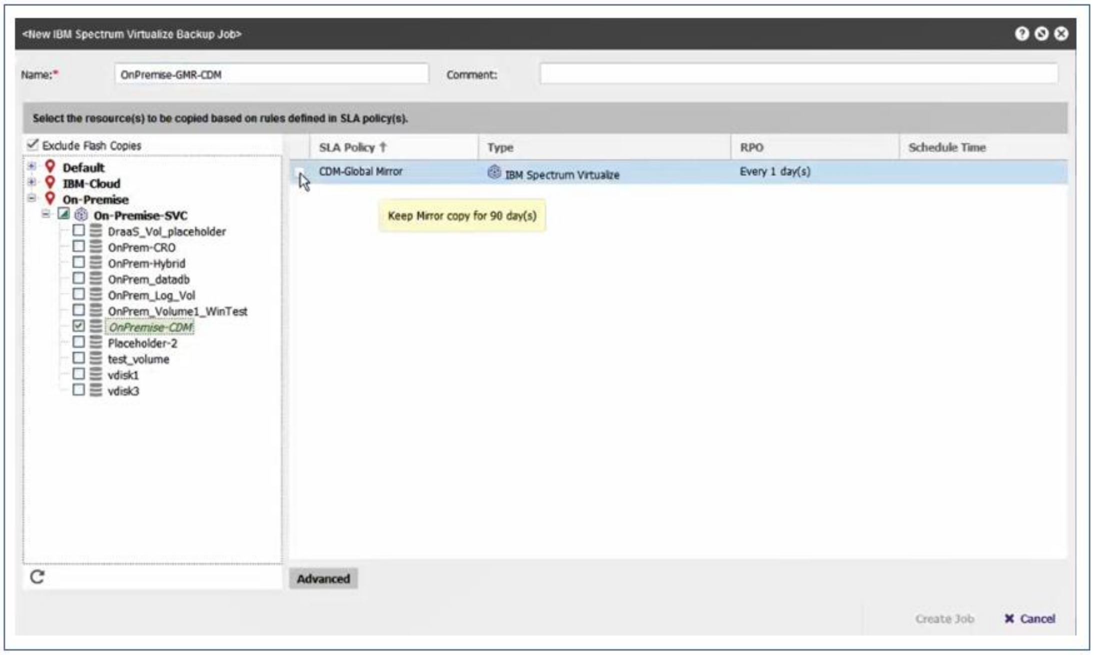
<New IBM Spectrum Virtualize Backup Job>
Select the resource(s) to be copied based on rules defined in SLA policy(s).
Default
IBM-Cloud
On-Premise
On-Premise-SVC
Advanced
Create Job
Cancel

Name:* OnPremise-GMR-CDM

Comment:

☑ Exclude Flash Copies

- ☐ DraaS_Vol_placeholder
- ☐ OnPrem-CRO
- ☐ OnPrem-Hybrid
- ☐ OnPrem_datadb
- ☐ OnPrem_Log_Vol
- ☐ OnPrem_Volume1_WinTest
- ☑ OnPremise-CDM
- ☐ Placeholder-2
- ☐ test_volume
- ☐ vdisk1
- ☐ vdisk3

SLA Policy ↑	Type	RPO	Schedule Time
CDM-Global Mirror	IBM Spectrum Virtualize	Every 1 day(s)	

Keep Mirror copy for 90 day(s)

Figure 65: Associate Global Mirror with Change Volumes SLA policy with the backup job

Once the job is successfully initiated, IBM Spectrum Copy Data Management will create the target volume and start the replication from the on-premises source site to the target IBM Cloud site, as shown in Figure 66.

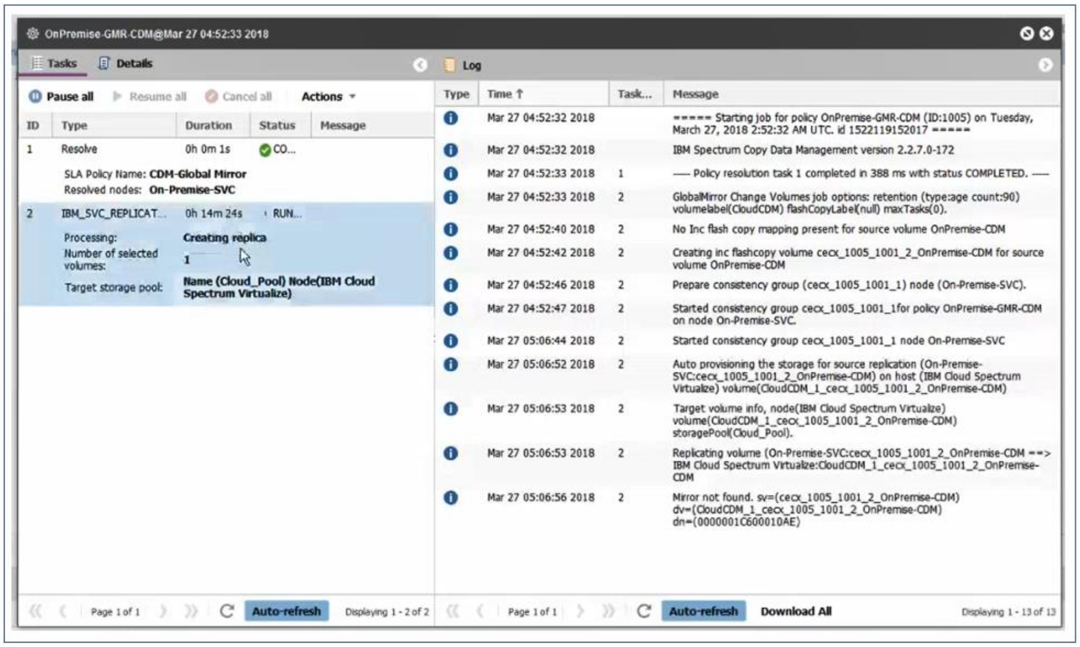

ID	Type	Duration	Status	Message
1	Resolve	0h 0m 1s	✓ CO...	
	SLA Policy Name: **CDM-Global Mirror** Resolved nodes: **On-Premise-SVC**			
2	IBM_SVC_REPLICAT...	0h 14m 24s	RUN...	
	Processing:	**Creating replica**		
	Number of selected volumes:	**1**		
	Target storage pool:	**Name (Cloud_Pool) Node(IBM Cloud Spectrum Virtualize)**		

Type	Time ↑	Task...	Message
ⓘ	Mar 27 04:52:32 2018		===== Starting job for policy OnPremise-GMR-CDM (ID:1005) on Tuesday, March 27, 2018 2:52:32 AM UTC. id 1522119152017 =====
ⓘ	Mar 27 04:52:32 2018		IBM Spectrum Copy Data Management version 2.2.7.0-172
ⓘ	Mar 27 04:52:33 2018	1	—— Policy resolution task 1 completed in 388 ms with status COMPLETED. ——
ⓘ	Mar 27 04:52:33 2018	2	GlobalMirror Change Volumes job options: retention (type:age count:90) volumelabel(CloudCDM) flashCopyLabel(null) maxTasks(0).
ⓘ	Mar 27 04:52:40 2018	2	No Inc flash copy mapping present for source volume OnPremise-CDM
ⓘ	Mar 27 04:52:42 2018	2	Creating inc flashcopy volume cecx_1005_1001_2_OnPremise-CDM for source volume OnPremise-CDM
ⓘ	Mar 27 04:52:46 2018	2	Prepare consistency group (cecx_1005_1001_1) node (On-Premise-SVC).
ⓘ	Mar 27 04:52:47 2018	2	Started consistency group cecx_1005_1001_1for policy OnPremise-GMR-CDM on node On-Premise-SVC.
ⓘ	Mar 27 05:06:44 2018	2	Started consistency group cecx_1005_1001_1 node On-Premise-SVC
ⓘ	Mar 27 05:06:52 2018	2	Auto provisioning the storage for source replication (On-Premise-SVC:cecx_1005_1001_2_OnPremise-CDM) on host (IBM Cloud Spectrum Virtualize) volume(CloudCDM_1_cecx_1005_1001_2_OnPremise-CDM)
ⓘ	Mar 27 05:06:53 2018	2	Target volume info, node(IBM Cloud Spectrum Virtualize) volume(CloudCDM_1_cecx_1005_1001_2_OnPremise-CDM) storagePool(Cloud_Pool).
ⓘ	Mar 27 05:06:53 2018	2	Replicating volume (On-Premise-SVC:cecx_1005_1001_2_OnPremise-CDM ==> IBM Cloud Spectrum Virtualize:CloudCDM_1_cecx_1005_1001_2_OnPremise-CDM
ⓘ	Mar 27 05:06:56 2018	2	Mirror not found. sv=(cecx_1005_1001_2_OnPremise-CDM) dv=(CloudCDM_1_cecx_1005_1001_2_OnPremise-CDM) dn=(0000001C600010AE)

Figure 66: Execution of backup job

With IBM Spectrum Copy Data Management, you can create application-consistent copy data policies that are aligned with the existing best practices for different applications. In this example, a VM is being protected. So, create the backup job for VMware by clicking **Jobs → VMware → Backup**.

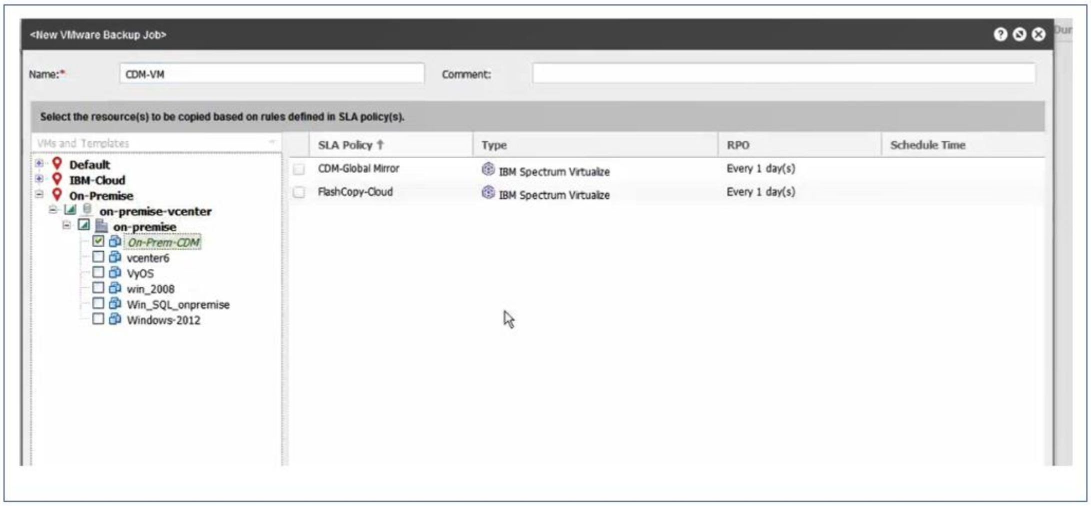

Name:* CDM-VM

Comment:

Select the resource(s) to be copied based on rules defined in SLA policy(s).

- Default
- IBM-Cloud
- On-Premise
  - on-premise-vcenter
    - on-premise
      - ☑ On-Prem-CDM
      - ☐ vcenter6
      - ☐ VyOS
      - ☐ win_2008
      - ☐ Win_SQL_onpremise
      - ☐ Windows-2012

	SLA Policy ↑	Type	RPO	Schedule Time
☐	CDM-Global Mirror	IBM Spectrum Virtualize	Every 1 day(s)	
☐	FlashCopy-Cloud	IBM Spectrum Virtualize	Every 1 day(s)	

Figure 67: Backup job for VM protection

As shown in Figure 67, select the VM that needs to be protected and associate the Global Mirror with Change Volumes SLA policy. Run the VM backup job. Once the backup job is completed, a restore job is required to use the data copies created.

Restore job

The next step in the process is to create a restore job.

The SLA policies allow you to access the data copies/snapshots for multiple use cases such as automated DR, or, as in this use case, the protected VM will be powered on with refreshed data from the production site in a snapshot created and mapped to the VMware infrastructure at the cloud site.

You must perform the following steps to create a new restore job.

On the Jobs tab, right-click **VMware → Restore**.

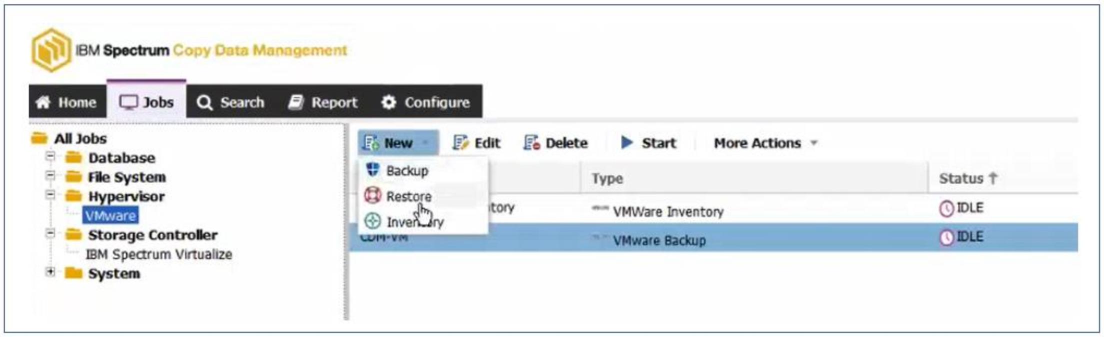

Figure 68: Creation of restore job

Click **Instant VMware restore**.

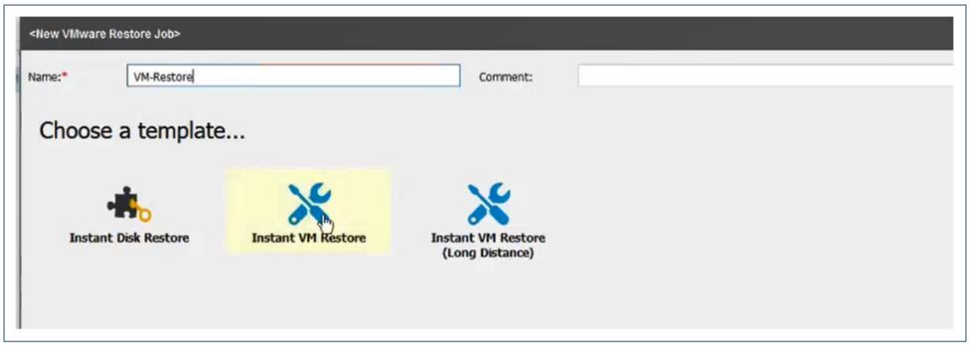

Name:* VM-Restore

Comment:

Figure 69: Restore template for VMs

Click **Source** and notice that the VM is protected in the copy policy. Click **Copy** (as shown in Figure 70) where it provides an option to use the latest snapshot version or use an older version. Click the **Advanced** button if you need a space-efficient snapshot or a complete clone of the volume.

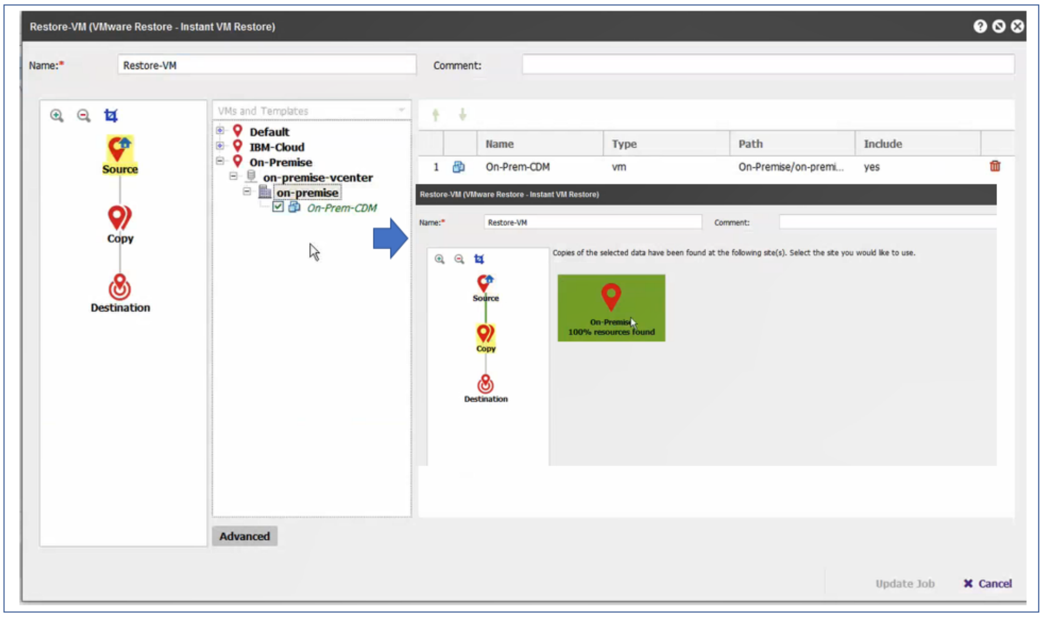

Restore-VM (VMware Restore - Instant VM Restore)

Name:* Restore-VM

Comment:

		Name	Type	Path	Include
1		On-Prem-CDM	vm	On-Premise/on-premi...	yes

Restore-VM (VMware Restore - Instant VM Restore)

Name:* Restore-VM

Comment:

Copies of the selected data have been found at the following site(s). Select the site you would like to use.

Figure 70: Configuration of restore job for VM

Select the destination ESXi host at the DR site where the test VM needs to be hosted. Provide the IP address details for the test VM, as shown in Figure 71.

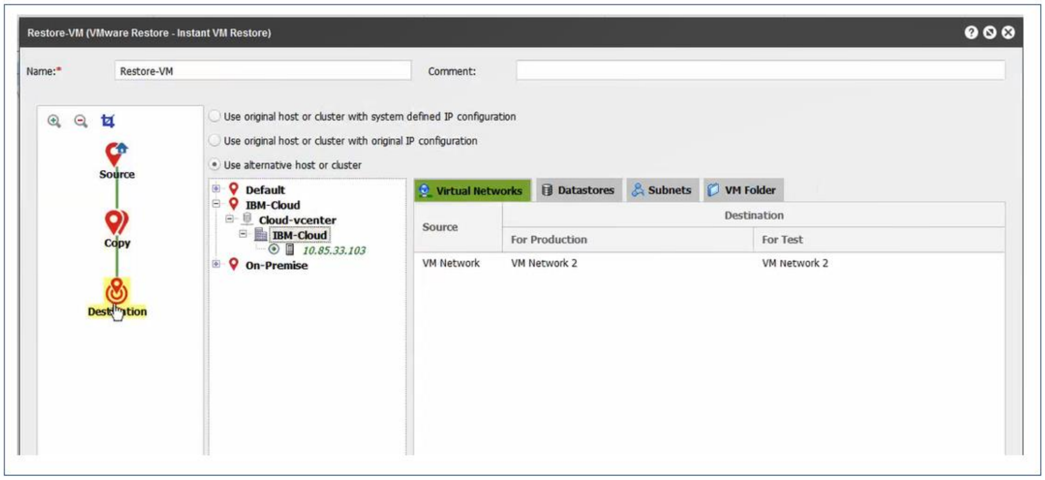

Restore-VM (VMware Restore - Instant VM Restore)

Name:* Restore-VM

Comment:

○ Use original host or cluster with system defined IP configuration

○ Use original host or cluster with original IP configuration

● Use alternative host or cluster

Source	Destination	
	For Production	For Test
VM Network	VM Network 2	VM Network 2

Figure 71: Setting parameters for target VM

The Use Data policy can be run manually, or it can be scheduled. On successful execution of the policy, the following steps are completed as part of a workflow:

- Create a snapshot
- Mount snapshot copies
- Map volumes to the destination ESXi host
- Mount the volume as a VM File System (VMFS) data store at the destination ESXi host
- Create a test VM with the network parameters provided in the policy
- Power on the VM at the destination

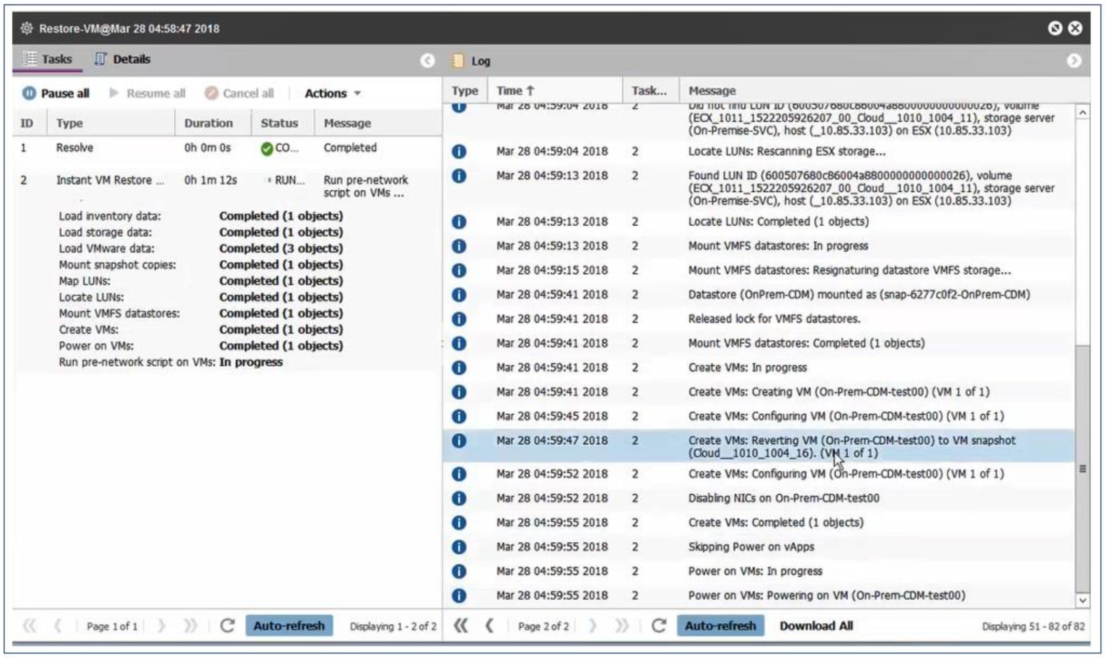

ID	Type	Duration	Status	Message
1	Resolve	0h 0m 0s	✓ CO...	Completed
2	Instant VM Restore ...	0h 1m 12s	RUN...	Run pre-network script on VMs ...

Load inventory data:	**Completed (1 objects)**
Load storage data:	**Completed (1 objects)**
Load VMware data:	**Completed (3 objects)**
Mount snapshot copies:	**Completed (1 objects)**
Map LUNs:	**Completed (1 objects)**
Locate LUNs:	**Completed (1 objects)**
Mount VMFS datastores:	**Completed (1 objects)**
Create VMs:	**Completed (1 objects)**
Power on VMs:	**Completed (1 objects)**
Run pre-network script on VMs:	**In progress**

Type	Time ↑	Task...	Message
ℹ	Mar 28 04:59:04 2018	2	Did not find LUN ID (600507680c86004a8800000000000026), volume (ECX_1011_1522205926207_00_Cloud__1010_1004_11), storage server (On-Premise-SVC), host (_10.85.33.103) on ESX (10.85.33.103)
ℹ	Mar 28 04:59:04 2018	2	Locate LUNs: Rescanning ESX storage...
ℹ	Mar 28 04:59:13 2018	2	Found LUN ID (600507680c86004a8800000000000026), volume (ECX_1011_1522205926207_00_Cloud__1010_1004_11), storage server (On-Premise-SVC), host (_10.85.33.103) on ESX (10.85.33.103)
ℹ	Mar 28 04:59:13 2018	2	Locate LUNs: Completed (1 objects)
ℹ	Mar 28 04:59:13 2018	2	Mount VMFS datastores: In progress
ℹ	Mar 28 04:59:15 2018	2	Mount VMFS datastores: Resignaturing datastore VMFS storage...
ℹ	Mar 28 04:59:41 2018	2	Datastore (OnPrem-CDM) mounted as (snap-6277c0f2-OnPrem-CDM)
ℹ	Mar 28 04:59:41 2018	2	Released lock for VMFS datastores.
ℹ	Mar 28 04:59:41 2018	2	Mount VMFS datastores: Completed (1 objects)
ℹ	Mar 28 04:59:41 2018	2	Create VMs: In progress
ℹ	Mar 28 04:59:41 2018	2	Create VMs: Creating VM (On-Prem-CDM-test00) (VM 1 of 1)
ℹ	Mar 28 04:59:45 2018	2	Create VMs: Configuring VM (On-Prem-CDM-test00) (VM 1 of 1)
ℹ	Mar 28 04:59:47 2018	2	Create VMs: Reverting VM (On-Prem-CDM-test00) to VM snapshot (Cloud__1010_1004_16). (VM 1 of 1)
ℹ	Mar 28 04:59:52 2018	2	Create VMs: Configuring VM (On-Prem-CDM-test00) (VM 1 of 1)
ℹ	Mar 28 04:59:52 2018	2	Disabling NICs on On-Prem-CDM-test00
ℹ	Mar 28 04:59:55 2018	2	Create VMs: Completed (1 objects)
ℹ	Mar 28 04:59:55 2018	2	Skipping Power on vApps
ℹ	Mar 28 04:59:55 2018	2	Power on VMs: In progress
ℹ	Mar 28 04:59:55 2018	2	Power on VMs: Powering on VM (On-Prem-CDM-test00)

Figure 72: Execution of restore job for VM for DevOps

After the test is completed, the same policy can be used and run again to clean up the test/development environment to remove the test VM and unmount the FlashCopy snapshot from the destination ESXi host.

Summary

With IBM FlashSystem 9100 and IBM Spectrum Virtualize for Public Cloud, customers are provided with the ability to optimize their existing heterogenous storage infrastructure and to plan for hybrid-cloud DR between on-premises and IBM Cloud storage. IBM Spectrum Virtualize for Public Cloud with VMware Site Recovery Manager integration offers customers planned migration of applications running on VMs across heterogenous underlying storage.

IBM Spectrum Copy Data Management offers the benefits of IBM Spectrum Storage for:

- Automatic creation and use of snapshots and replicas in existing installations based on IBM Spectrum Virtualize and IBM FlashSystem 9100 systems, to ensure application consistency
- Simplifying management of data copies by efficiently maintaining data copies residing on IBM Spectrum Virtualize storage systems on-premises or on the cloud
- Leveraging high-value use cases, such as automated DR in hybrid cloud environments
- Catering to the modern use case of using data copies in a DevOps environment with the capability of integrating application-centric VMware and IBM Spectrum Virtualize systems

Get more information

- IBM Redbooks: IBM Spectrum Virtualize for Public Cloud:
  http://www.redbooks.ibm.com/redpapers/pdfs/redp5466.pdf

- IBM Redbooks: IBM Spectrum Copy Data Management user guide:
  https://www.ibm.com/support/knowledgecenter/en/SS57AN_2.2.7/com.ibm.spectrum.cdm.doc/ss57an_pdf_files.html

- IBM Redbooks: IBM FlashSystem 9100

- IBM Spectrum Virtualize Family Storage Replication Adapter version 3.3.0:
  https://www.ibm.com/support/knowledgecenter/en/SSEQ4E_3.3.0/UG/SVC_SRA_SRM5_1_install.html

- VMware vCenter version 6.5 configuration:
  https://docs.vmware.com/en/VMware-vSphere/6.5/vsphere-esxi-vcenter-server-65-installation-setup-guide.pdf

- VMware Site Recovery Manager configuration:
  https://docs.vmware.com/en/Site-Recovery-Manager/6.5/com.vmware.srm.install_config.doc/GUID-B3A49FFF-E3B9-45E3-AD35-093D896596A0.html

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

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law IBM Japan Ltd.
1 9-2 1 , Nihonbashi-Hakozakicho, Chuo-ku Tokyo103-8510, Japan

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, MERCHANT ABILITY 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 supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact :

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119 Armonk, NY 10504-1785
US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

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.

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 the names and addresses used by an actual business enterprise 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.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_.

Trademarks

IBM, the IBM logo, ibm.com, IBM Cloud, IBM FlashSystem, IBM Spectrum, IBM Spectrum Storage, IBM Spectrum Virtualize, Easy Tier, FlashCopy, Passport Advantage and Storwize are trademarks or registered trademarks of International Business Machines Corp., 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 www.ibm.com/legal/copytrade.shtml.

Intel is a trademark of Intel Corporation in the U.S. and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

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

VMware, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions

References in the publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates.

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.

Personal use
You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

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.

August 2018.

[1] “TBR Hybrid Cloud Customer Research 1H17 | Cloud Business Quarterly,” *Technology Business Research*, 2017.