COHESITY

Use Cohesity Multi-tenancy to Offer Disaster Recovery as a Service (DRaaS)

Cohesity Platform's Multi-tenancy Enables DRaaS for Service Providers

Version 1.3

December 2025

ABSTRACT

The document is focused on service providers who manage a multi-tenant IT infrastructure and want to provide DRaaS for their customers using Cohesity Platform. The primary audience for this document within service providers is individuals serving in roles such as Product Managers, Service Managers, Enterprise Architects, and Operations, who will use the guide as a reference. The secondary audience for this guide is Solution Development teams, including Solution and Portfolio Architects and Production Change Management organizations, and anyone who needs to understand the technical aspects of the solution offerings and capability delivered by Cohesity.

Table of Contents

Figures

Tables

Introduction to Disaster Recovery as a Service with Cohesity Platform

Replication is the fundamental operation that makes any Disaster Recovery as a Service (DRaaS) solution possible. Cohesity's DRaaS refers to the approach of maintaining a regularly updated replica of organizational backup data in another Cohesity cluster in the same or different location for data retention and offsite storage. Cohesity leverages service providers' multi-tenant infrastructure wherein they replicate data from multiple customers on shared but logically isolated infrastructure.

Secure multi-tenancy is an important component of planning IT infrastructure today. Multi-tenancy has many advantages, including investment efficiencies, security, and data isolation. With Cohesity, service providers can create an "organization" corresponding to each customer or tenant on a Cohesity cluster. An Organization Name and Organization ID serve as the multi-tenancy identifier for each tenant.

Figure 1: Use Cohesity Multi-tenancy to Offer DRaaS

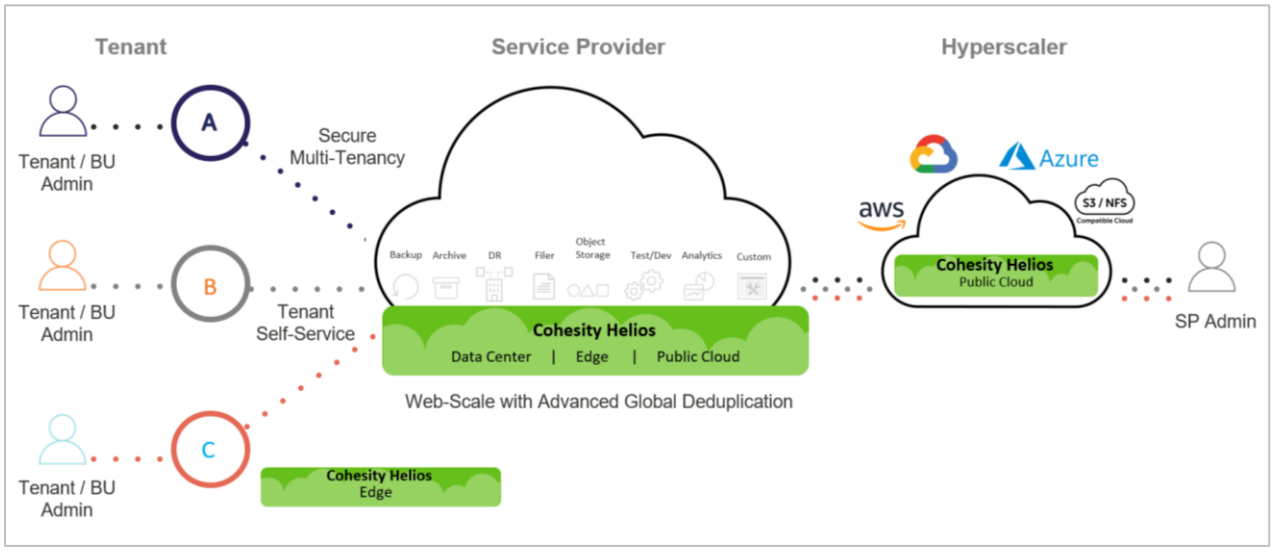

Benefits of Using Cohesity for Disaster Recovery

Cohesity Helios Platform is a purpose-built solution for service providers. Figure 2 below illustrates the features provided by Cohesity that complement a typical service provider deployment.

Figure 2: Benefits of Using Cohesity for Disaster Recovery

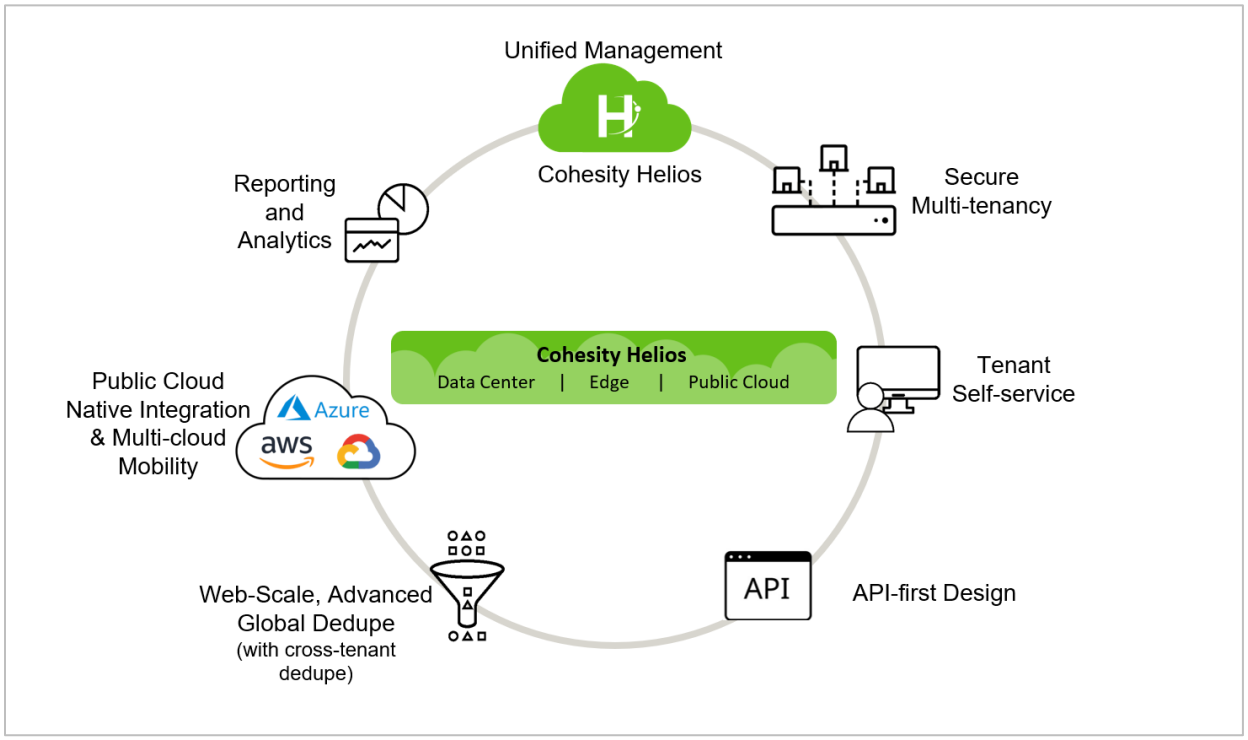

- **Unified Management**. Helios is Cohesity's SaaS-based management platform that provides a single view and global management of all your Cohesity clusters, whether on-premises, cloud, or Virtual Edition, regardless of cluster size. You can quickly connect your clusters to Helios and then access them from anywhere using an internet connection and your Cohesity Support Portal credentials.

  Helios provides:

  - **Multi-cluster management.** Actively manage all your primary and DR clusters with multi-cluster monitoring, reporting, and orchestrated upgrades, from a single dashboard.
  - **Global actionable search**. Search across clusters and take actions right from the search results page. For example, search for all unprotected VMs and create Protection Groups to protect them.
  - **SmartAssist**. Automatically schedule and orchestrate Protection Runs for tenants and workloads to help meet SLAs. Get recommendations based on capacity forecasting and disk failure prediction. View important Field Notices from Cohesity Support.
  - **Cloud Edition**: Deploy NGCE clusters using Helios.

- **Secure Multi-tenancy**. Secure multi-tenancy can be enabled on Cohesity clusters to provide a logical separation for multiple tenants hosted on the same physical cluster. Each tenant is represented as an ‘organization’ in a Cohesity cluster.

- **Tenant Self-service**. Self-service means that tenant users can independently access the DR cluster to perform a recovery when their primary site is partially or fully unavailable. With role-based access control (RBAC), tenant users can have role-based access to the organization, which facilitates self-service data restore without requiring assistance from the service provider.
- **API-first Design**. API-first design enables automation and orchestration for almost all the workflows available on the cluster. You can, for example, achieve DR orchestration and automation using the API framework.
- **Web-scale, Advanced Global Deduplication**. Cohesity Platform is built on a web-scale architecture. It provides great operational scalability and improved storage efficiency thanks to advanced global deduplication.
- **Public Cloud-native Integration and Multi-cloud Mobility**. Cohesity natively integrates with all major public cloud vendors and provides multi-cloud mobility, which, in turn, allows cloud DR facilities to offer DRaaS on their customers' preferred cloud platforms
- **Efficient Data Replication to a Remote Cluster**. Take advantage of Cohesity Platform's deduplication and compression technologies to replicate data from a source cluster to a remote DR cluster efficiently.
- **Reporting and Analytics:** Cohesity gives service providers access to usage metrics to help them implement chargeback. Service providers can retrieve and analyze these metrics via:
  - **Built-in Reports**. Accessible in the Cohesity browser UI. See Reports in the online Help.
  - **Custom Reports**. Generated using:
    - The Cohesity REST API.
    - The Custom Reporting Database. See the Cohesity Custom Reporting Solution Guide.
    - API Access from the cluster and Helios.

DRaaS Terminology

Before you proceed to set up DRaaS with Cohesity, you want to get familiar with a few common terms used in the various scenarios of configuration and deployment of DRaaS infrastructure.

Table 1: DRaaS Terminology

TERMINOLOGY	DESCRIPTION
AD	Active Directory
BaaS	Backup as a Service
Cross-connect	A cross-connect is a point-to-point cable link between the customer and service providers.
DNS	Domain Name Services
Hybrid Extender	Hybrid Extender is a proxy that is deployed on the tenant vCenter and helps to set up a TCP/IP secure channel from the tenant's local network to the Cohesity cluster in the service provider environment.

TERMINOLOGY	DESCRIPTION
Hybrid Extender V2 (HyXV2)	Hybrid Extender V2 launched in version 6.6, and it has dual home configuration.
Network Realm	Network Realm ensures source-level selection in a multi-tenant environment and allows you to create and manage multiple isolated networks within a tenant.
IaaS	Infrastructure as a Service
Multi-tenancy	Multi-tenancy is an architecture in which a single instance is shared across multiple clients that are logically isolated from each other while being physically integrated.
Replication	Organizations can achieve enterprise-level resiliency with site-to-site replication between Cohesity clusters.
NTP	Network Time Protocol
Organization	Organization is a multi-tenancy identifier for each tenant.
RBAC	Role-based access control
TCP	Transmission Control Protocol
Virtual IPs (VIPs)	Virtual IP address. * *Cohesity recommends you use one virtual IP for each node* in a Cohesity *cluster.*
VLAN	Virtual Local Area Network
VPN tunnel	VPN tunnels are used to connect two private networks across public networks.
Primary Cluster	Cohesity cluster on which the primary workload is being backed up.
DR Cluster	Cohesity cluster on the DR site to which the data is being replicated from the primary cluster.
Primary Workload Infrastructure	Primary workload infrastructure is the infrastructure from which the primary cluster is backing up the data.
DR Workload Infrastructure	DR workload infrastructure is the infrastructure where the workloads are recovered when primary infrastructure is down.

Explore the DRaaS Capabilities of Cohesity Platform

Service providers require certain crucial capabilities to be able to offer "as-a-service" solutions. Table 2 outlines these key requirements and explains how to achieve them using the native features in Cohesity.

Table 2: Native Cohesity Features for DRaaS

FEATURE	DETAIL
Per-tenant Replication	Remote replication on Cohesity clusters is secure multi-tenancy aware, that is, replication is managed securely on a per-tenant basis, a core requirement for service providers offering DRaaS.
Network Isolation	Isolate replication traffic for each tenant using dedicated VLAN/IP address ranges in Local Backup with Offsite Replication deployments.
Data Isolation	Offer isolation of replicated data among multiple tenants using tenant-specific Storage Domains hosted on remote clusters paired with their sources.
Per-tenant Reporting	• Organizations see their storage consumption. • For chargeback purposes, service providers can access consumption information including tenant organizations, Storage Domains, and data consumption per tenant with the new "Storage Consumed by Organizations" report.
Storage Efficiency	Service providers can achieve even higher storage efficiency rates by choosing the option to configure tenants to share the same Storage Domain, which acts as a single deduplication domain, across tenant data, while access is still limited to each.
Tenant Impersonation	• Impersonate organization users to preview and verify what the tenant sees for improved troubleshooting and debugging. Examples of impersonation include: ○ Service provider administrator acquires tenant administrator privileges during impersonation. ○ Service provider operator acquires tenant operator privileges during impersonation. ○ Service provider viewers have tenant *viewer* privileges during impersonation. * *To ensure maximum security, all actions taken under impersonation are audit-logged with the authentic user's identity (not the impersonated identity) and are visible to the tenants and the service provider*.
Single Sign-on (SSO)	• Cohesity uses an Identity Provider (IdP) for single sign-on (SSO) access to the Cohesity cluster.

FEATURE	DETAIL

	• For a multi-tenancy enabled Cohesity cluster, configure RBAC (role-based access control) using SSO, LDAP, or Active Directory for any organization defined in the cluster.
	• Supported SSO Identity Providers:
	  ○ Active Directory Federation Services
	  ○ Azure Active Directory
	  ○ Duo
	  ○ Ping Identity
	  ○ Okta
	  ○ Auth0

Self-service Data Recovery	Achieve self-service data recovery using a combination of SSO and RBAC.
Data encryption and key management	• Cohesity has a built-in key manager that generates keys and stores them internally on the SSDs in encrypted form. • Data on the cluster is encrypted using a Data Encryption Key (DEK) and the DEK is in turn encrypted using a Key Encryption Key (KEK). The encryption/decryption process is done within the Cohesity cluster and is transparent to all inbound/outbound protocols and applications, such as backup, archiving, and file services. Also, the keys are stored in a distributed fashion to be resilient in the rare event of a hardware failure. • The default key rotation policy for KEK is 90 days, but the administrator can configure a different value.

Third-party Integration	• Cohesity's API-first architecture offers seamless automation, consumable APIs, error-free consistency, and self-service manageability.
	• Supported automation frameworks include:
	  ○ VMware vCloud Director
	  ○ VMware vRealize Automation (vRA) and Orchestration (vRO)
	  ○ ServiceNow
	  ○ Python/Powershell SDK

Determine Your Network Topology

To be able to set up and offer DRaaS, a service provider must first set up Backup as a Service (BaaS), which can be deployed in different network topologies. Because the procedure to set up DRaaS depends on BaaS, you need to first determine which network topology is deployed in your infrastructure.

Two of the most common are:

- **Scenario 1**: Hosted Backup with Offsite Replication
- **Scenario 2**: Local Backup with Offsite Replication

Figure 3: DRaaS Deployment Topologies

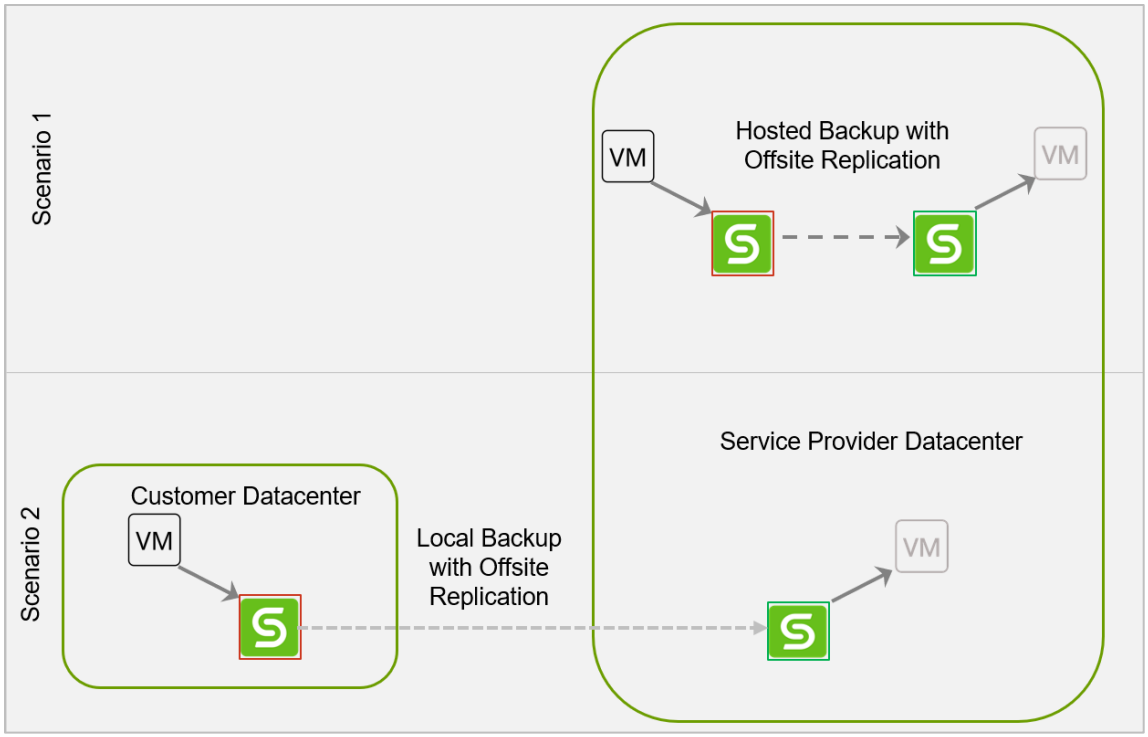

Hosted Backup with Offsite Replication

Service providers often provide shared hosted infrastructure as a service to their customers wherein they host the customer's IT infrastructure. Cohesity offers them the ability to provide their customers with both backup and disaster recovery as a service in a secure multi-tenant environment.

In Figure 4 below, Tenant 1 infrastructure (blue) is co-hosted with Tenant 2 infrastructure (orange), and both are backed up to a Cohesity cluster. Cohesity's DRaaS solution replicates this backup data to the remote DR cluster via another VLAN channel.

Figure 4: Data Backup to Source Multi-tenant Cluster and Replication to DR

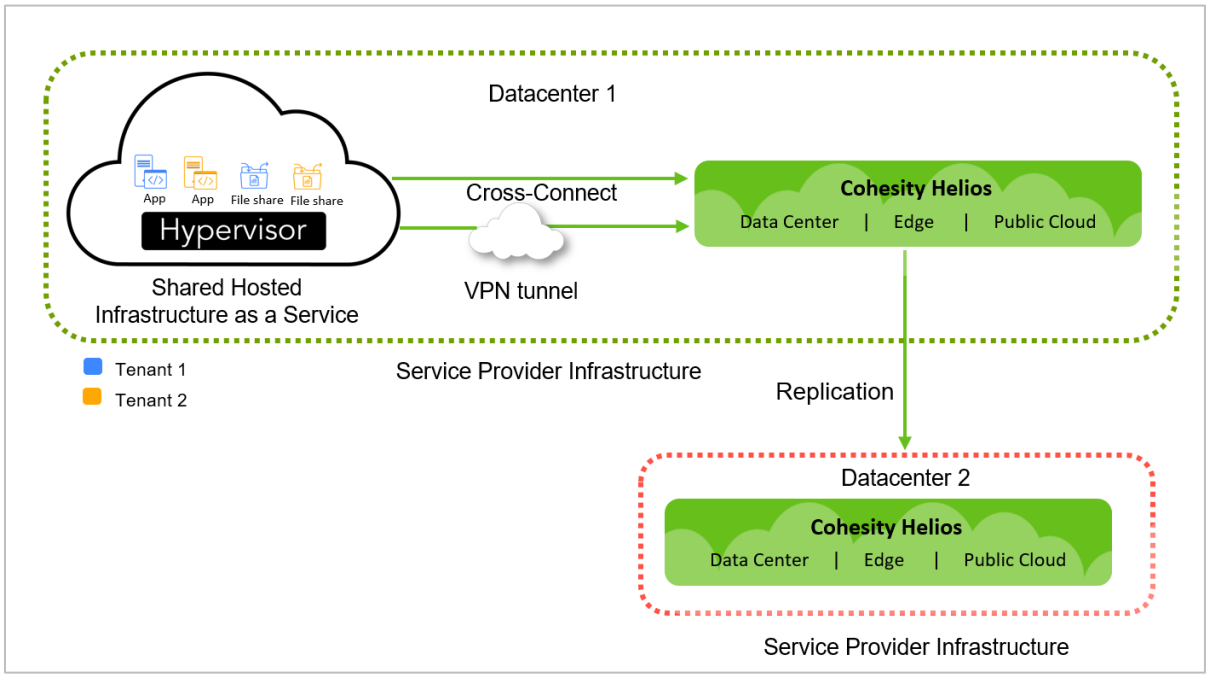

Local Backup with Offsite Replication

In this scenario, service providers provide managed backup as a service for customers who require the primary copy of backup data to remain on their premises and then offsite the backup (similar to tape offsiting) to a central service provider location for DR purposes by means of replication.

To that end, service providers can deploy Cohesity Virtual Edition (VE) or Physical Edition in the tenant environment to take the backups locally and then replicate or archive the backup snapshots to their environment.

Local Backup with *Offsite Replication* can be implemented in different ways, depending on whether the tenant and service provider data centers have a VPN connection or not.

Replicate Clusters over WAN Using VPN

Two sites connected over WAN using VPN act as a stretched network. Hence, the replication works seamlessly without special configuration.

Figure 5 below illustrates how service providers can use Cohesity to achieve local backup with offsite replication. Each tenant backs up their data to a local, on-premises physical or Virtual Edition, Cohesity cluster, and replicates the backed up data to the service provider's remote DR cluster.

Figure 5: Replicate Clusters over WAN Using VPN

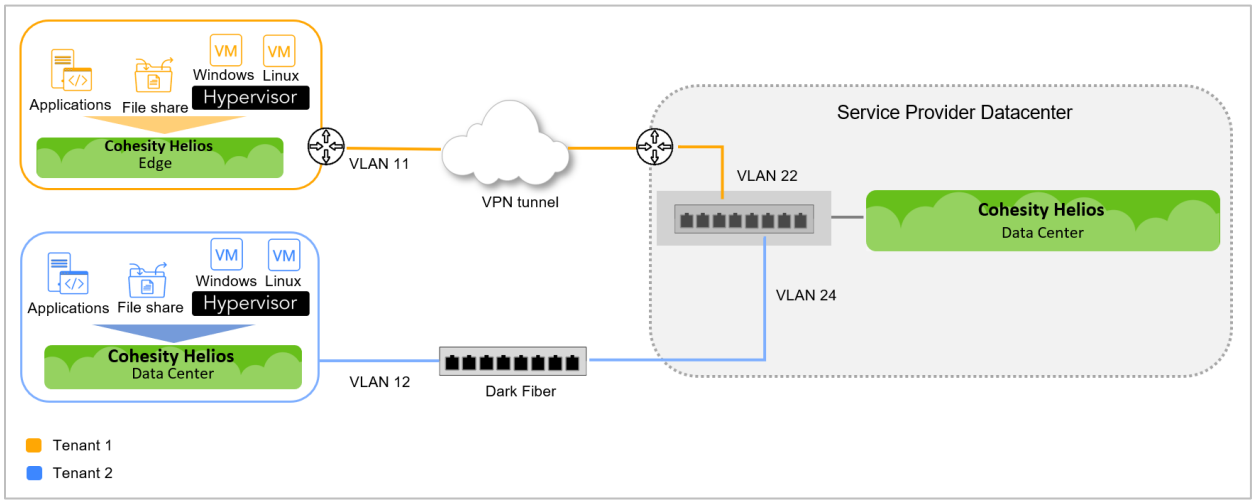

For more details on adding and assigning VLANs per organization, see Multi-tenancy Deployment Guide.

Replicate Cluster over WAN without VPN

To replicate data over WAN without using a VPN, you can use a network address translation (NAT) gateway on both the source and DR sites, to map the private IPs for Cohesity node VIPs to NAT public IPs.

POC Setup for Replication over WAN

For proof of concept (POC) setup, the NAT can be a 1:1 NAT between one of the Cohesity VIPs and a NAT public IP on both sites. Even though the connection is between one VIP on the source cluster to one VIP on the DR cluster, the replication workload is distributed across the cluster nodes internally.

Figure 6: Proof of Concept Setup for Replication over WAN

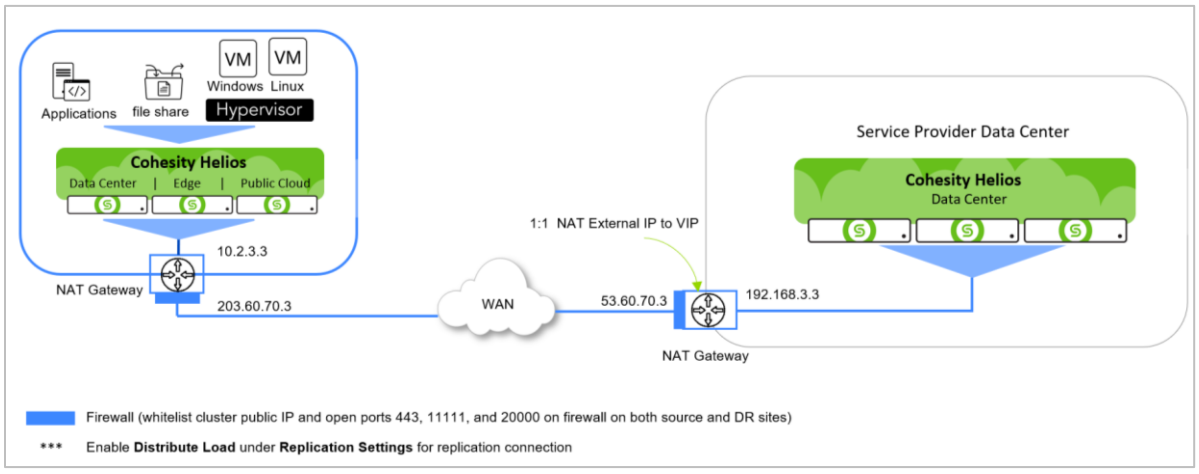

Firewall: Whitelist public IPs and open ports 443, 11111, and 20000 on the firewall for both source and DR.

* Use the public IPs to create the replication connection.

Production Setup for Replication over WAN

For a production setup, the NAT load balancer should be used in round-robin mode on the service provider side. All the Cohesity VIPs should be NATed to a single NAT public IP. The NAT gateway takes care of load-balancing the incoming request on the public IP to the multiple Cohesity VIPs.

Figure 7: Production Setup for Replication over WAN

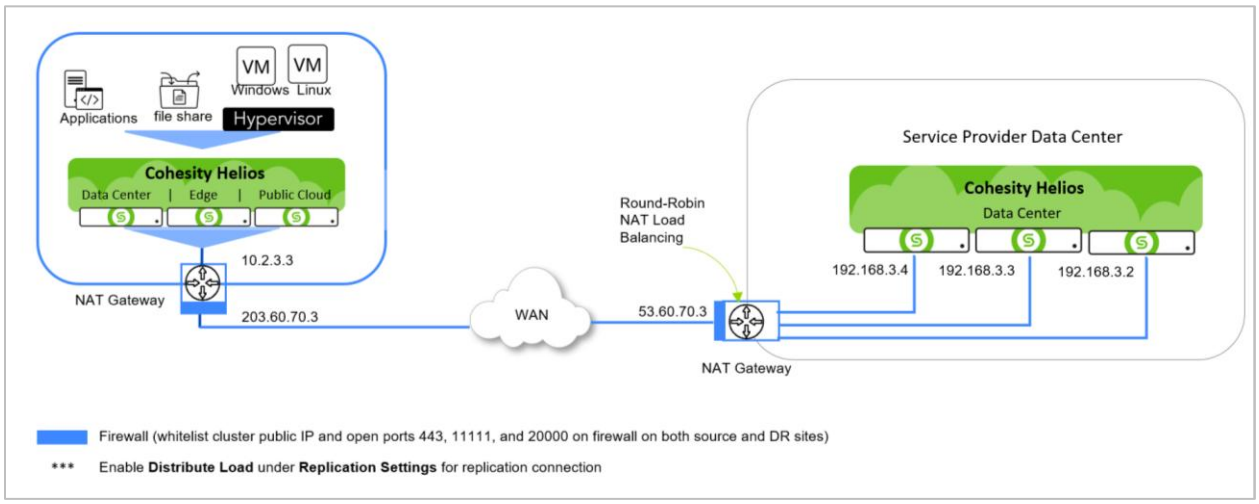

Firewall: Whitelist public IPs and open ports 443, 11111, and 20000 on the firewall for both source and DR.

* Use the public IPs to create the replication connection.

For more on Cohesity networking, see Cohesity Networking Quick Start Guide and Optimal Network Designs with Cohesity.

NOTE: You cannot use FQDN to create the remote connection between source and DR cluster, as this setup requires public IPs.

Deploy DRaaS for Your Network Topology

As mentioned earlier, the procedure to set up DRaaS for your Backup as a Service infrastructure and then recover data using a DR Cohesity cluster is dependent on your deployed network topology.

Deploy DRaaS in 'Hosted Backup with Offsite Replication' Topology

If your backup infrastructure topology is Hosted Backup with Offsite Replication, the procedure to set up the DR cluster for replication and recovery is:

1. Prepare Your Clusters — see Common Prerequisites for Replication and Prerequisites for 'Hosted Backup with Offsite Replication' Topology below.
2. Create Connection with Remote DR Cluster.
3. Add DR Cluster to Protection Policy.
4. Add Source to Remote DR Cluster.
5. Perform Data Recovery
    a) Recover Data from a DR Protection Group.
    b) Failover to a DR Protection Group and Recover Data.

Deploy DRaaS in 'Local Backup with Offsite Replication' Topology

If your backup infrastructure topology is Local Backup with Offsite Replication, the procedure to set up the DR cluster for replication and recovery is:

1. Prepare Your Clusters.
2. Create Connection with Remote DR Cluster.
3. Add DR Cluster to Protection Policy.
4. Add Source to Remote DR Cluster.
5. Perform Data Recovery
    a) Recover Data from a DR Protection Group.
    b) Failover to a DR Protection Group and Recover Data.

Prepare Clusters for Replication

Replication is the fundamental operation that facilitates a DRaaS solution. The source and remote clusters need to be prepared for replication, especially when multi-tenancy is enabled. To ensure a successful replication, you need to satisfy a few prerequisites, as listed below.

Common Prerequisites for Replication

The prerequisites that apply to both network topologies (Hosted and Local) are:

- Ensure that both the source and DR clusters are running Cohesity version 6.6 or later.
- Open/allow the firewall ports in Table 3 to establish a connection between the source and remote clusters.

Table 3: List of Firewall Ports to Open for Primary-DR Cluster Connection

PORT	SOURCE	TARGET	DIRECTION	NETWORK PROTOCOL	USAGE NOTES	TYPE OF TRAFFIC
443	Source Cohesity Cluster	Disaster Recovery Cluster	Bidirectional	TCP	Required for remote access to the cluster.	Replication
11111	Source Cohesity Cluster	Disaster Recovery Cluster	Bidirectional	TCP	Data path handling	Management
20000	Source Cohesity Cluster	Disaster Recovery Cluster	Bidirectional	TCP	Interaction handling between Cohesity and target backup systems	Management

- Add Replication to the Protection Policy assigned to the tenant on the source cluster.
- If you are replicating over WAN without VPN, whitelist the replication IPs on the firewall.

Prerequisites for ‘Hosted Backup with Offsite Replication’ Topology

For a Hosted Backup with Offsite Replication topology, there are a few prerequisites to satisfy in addition to those mentioned in Table 3 above.

- Create an organization on the DR cluster corresponding to the source. Use the same Organization ID (case-sensitive) on both the source and DR clusters.
- Create a Storage Domain Pairing between Source and DR Cluster.

Prerequisites for ‘Local Backup with Offsite Replication’ Topology

For a Local Backup with Offsite Replication topology, there are a few prerequisites to satisfy in addition to those already mentioned in Table 3 above.

- In a *Local Backup with Offsite Replication* topology, the source cluster is not enabled for multi-tenancy but the DR cluster is enabled for multi-tenancy. In this scenario, create organizations on the DR cluster and register the DR cluster on the source cluster using the organization’s username and password.
- Create a Storage Domain Pairing between Source and DR Cluster (where Organization is created)

Set Up Clusters for Replication

To implement Cohesity's DRaaS solution in a tenant's backup infrastructure, either the service provider administrator or the tenant administrator first needs to set up data replication between the source and DR Cohesity clusters. This involves two steps:

1. Create a connection with the remote cluster.
2. Add the remote cluster to the Protection Policy under replication configuration.

Create Connection with DR Cohesity Cluster

When you create a replication connection, you authenticate the source and DR cluster on each other, pair the corresponding Storage Domains, and select the IPs for replication traffic along with other replication settings.
To set up a replication run between a source and remote Cohesity cluster, both clusters must have authenticated communication between each other and for that, a remote connection must be set up. The procedure to set up a remote cluster connection depends on the tenant's DRaaS network topology.

To create a connection between:

- Hosted backup cluster and remote DR cluster, see Connect Your Hosted Backup Cluster to a DR Cluster.
- Local backup cluster and the remote DR cluster, see Connect Your Local Backup Cluster to a DR Cluster.

Connect Your Hosted Backup Cluster to a DR Cluster

In a *Hosted Backup with Offsite Replication* topology, the service provider administrator should create the remote connection. To perform that operation, the service provider administrator must have the following information:

- VIP address for the source and remote cluster.
- Organization pairs created on both source and remote cluster with the same Organization IDs.
- Administrator credentials for source and remote cluster.
- Interface group name to be used for replication traffic on source and remote cluster.
- Storage Domain mapping between Storage Domains on source and remote cluster.

Once the service provider has that information, they can create the remote connection.
To register a connection to a remote Cohesity cluster for replication:

1. Log in to the Primary Cohesity cluster.
2. Go to **Infrastructure > Remote Clusters**.

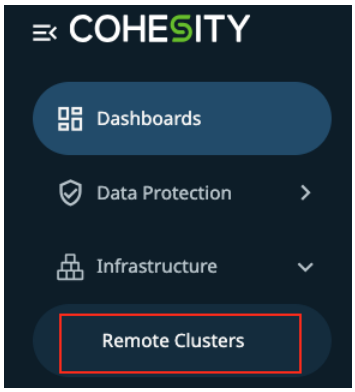

3. Click on **Register Remote Cluster**.

4. Enter **VIP or Node IP Addresses** for the remote cluster, the service provider administrator's **Username** and **Password**, and select **Interface Group** to enter the **Interface Group** name**.** Click **Continue** to validate the connection**.**

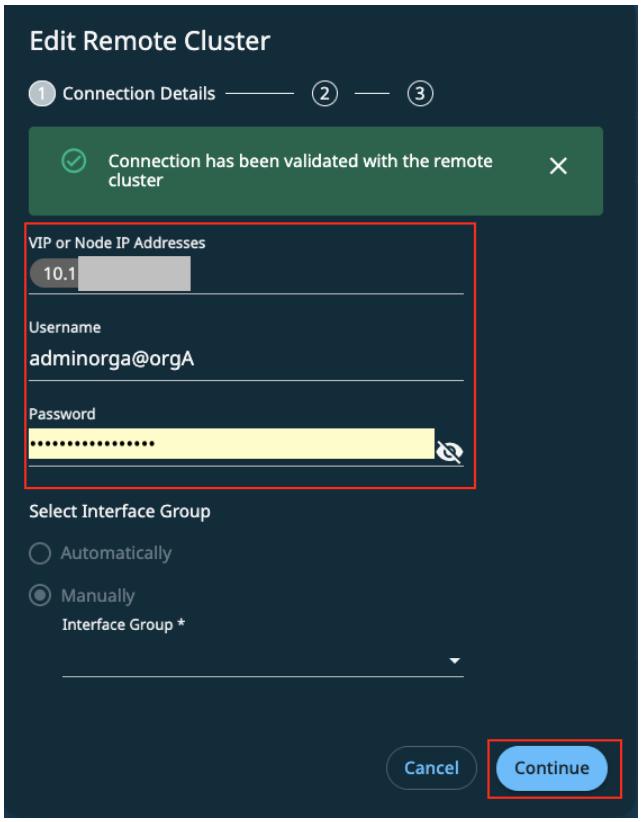

5. Under **Replication Settings**, **Pair Storage Domains** for Source and DR Cluster, enable **Compression,** if all IPs are reachable, enable this option.

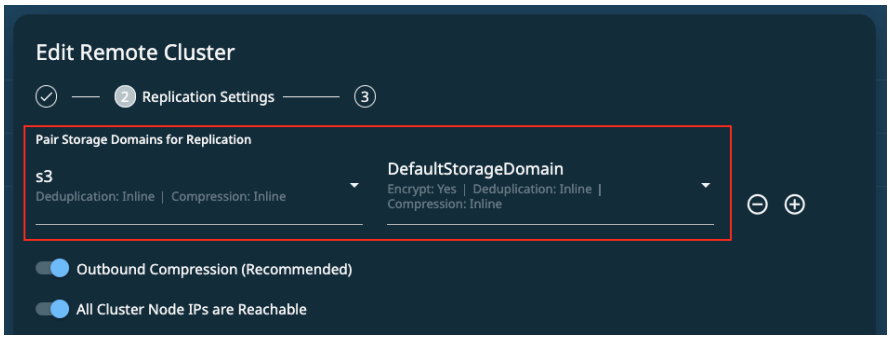

6. Under **Replication settings**, enable **Encryption In-flight**, click **Generate New Key.** The encryption key is displayed. The Cohesity cluster that is receiving encrypted data must enable encryption and specify the same encryption key as the Cohesity cluster that sent the replication data. Copy the encryption key to the clipboard, so you can paste it while creating the connection from the capturing Cohesity cluster back to the remote Cohesity cluster as described in Create Connection Back to the Capturing Cohesity cluster.

7. You can enable Throttle if you need to throttle replication traffic. To select custom times and transfer limits, enable Quiet Time and Throttle Overrides**.** Click **Continue** and **Save** to create the replication connection.

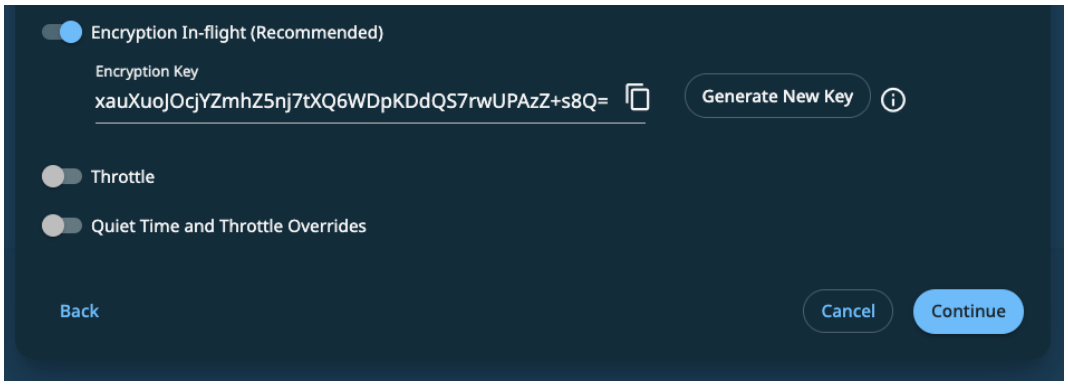

For more on Replication Settings, see Create a Connection to the Remote Cluster.

8. Be sure to set up a replication connection from the DR cluster, as well.

> **NOTE:**
>
> - If your Cohesity version is 6.5 or later, only a unidirectional replication setup is required from the source cluster to the DR cluster. Vice-versa is not required.
> - Because the source and the remote cluster have the same Organization IDs, Cohesity automatically detects the Protection Groups for an organization on the source and replicates its snapshots to the remote cluster under the same *Organization ID*.

For more, see Replication and Remote Access Setup in the online Help.

Once the connection is established between your primary cluster and DR cluster, proceed to Add Remote Cluster to Protection Policy Under Replication.

Connect Your Local Backup Cluster to a DR Cluster

In a *Local Backup with Offsite Replication* topology, both the tenant administrator as well as the service provider administrator have to perform a series of steps to create a remote cluster connection, as illustrated below.

Figure 8: Connect Local Primary Cluster to Remote Cluster

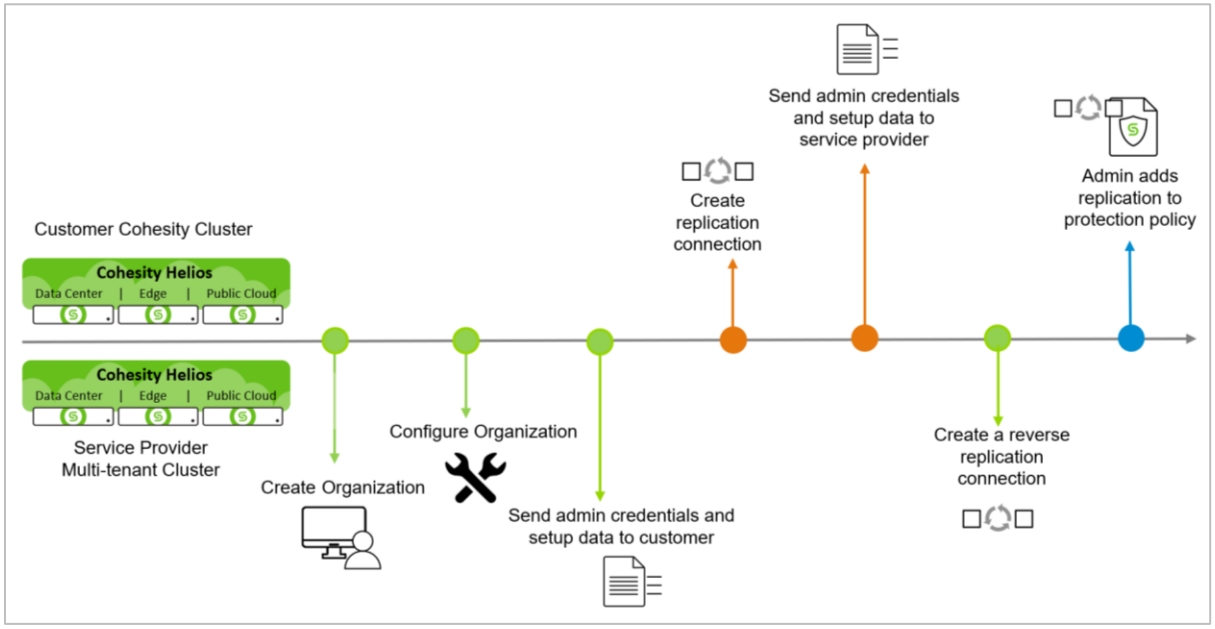

Tenant Cluster Administrator Tasks

Before proceeding to register a connection from the primary (tenant) cluster to a DR Cohesity cluster in the service provider's data center for replication, the tenant administrator must have the following details:

- VIP addresses for the DR cluster.
- Tenant administrator credentials for the organization created on the service provider cluster.
- Storage Domain name assigned to the tenant.

Once the tenant administrator has that information, they can register the remote cluster connection using the steps below:

1. Log in to the primary Cohesity cluster.
2. Go to **Settings > Targets > Remote Clusters**.
3. Click **Add Cluster** in the top right of the page.
4. Enter **VIP or Node IP Addresses** for the remote cluster**,** the tenant administrator's **Username** and **Password**, and select **Interface Group** to enter the **Interface Group** name. Click **Connect**.

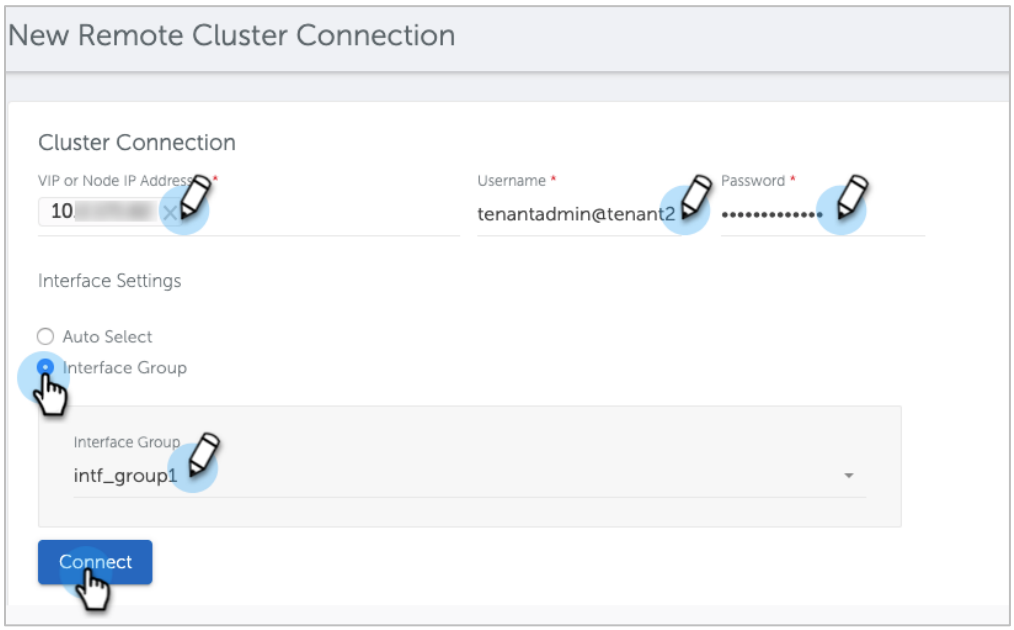

5. Under **Cluster Options**, enable **Replication** and click **Add Storage Domain Pairing**.

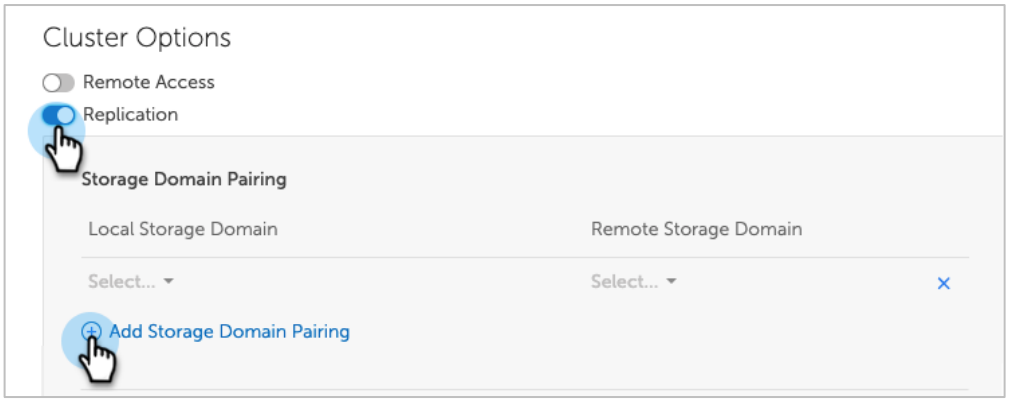

6. Select the **Local Storage Domain** and the **Remote Storage Domain** that you intend to connect for replication.

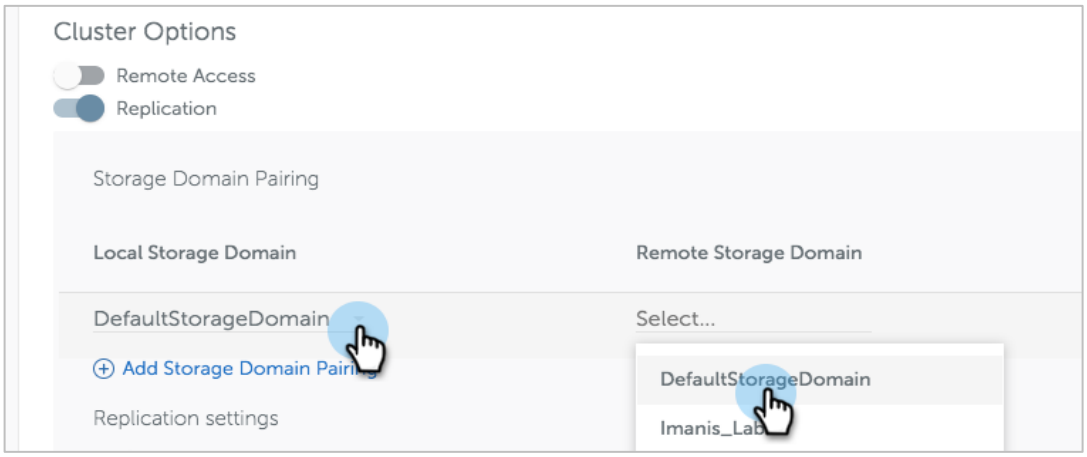

7. Under **Replication settings**, select **Distribute Load**, **Outbound Compression**, **Enabled Compression**, enter the **Encryption Key**, and enable **Data Transfer Rate Limit** if you need to throttle replication traffic. To select custom times and transfer limits, enable **Blackout and Data Transfer Rate Limit Overrides**. Click **Create** to create the replication connection.

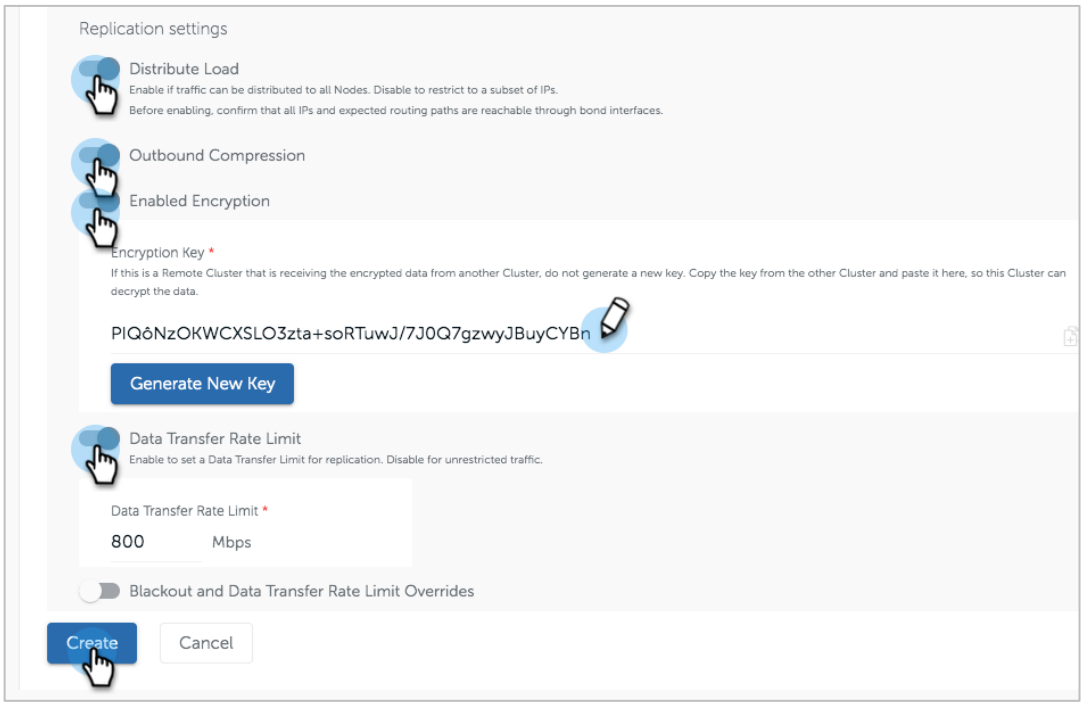

Service Provider Administrator Tasks

Once the primary (tenant) cluster administrator creates a replication connection from the primary cluster to the service provider (DR) cluster, (prior to Cohesity version 6.5) the service provider must create a reverse replication connection from the DR cluster to the primary cluster as well. For this, the service provider administrator must have:

- VIP address for the primary cluster.
- Cluster administrator credentials for the primary cluster.
- Storage Domain name that is used by the Protection Group to be replicated.

Once the service provider administrator has that information, they can establish a reverse replication connection using the steps below:

1. Log in to the DR Cohesity cluster as the tenant administrator (The service provider administrator can impersonate a tenant administrator user).

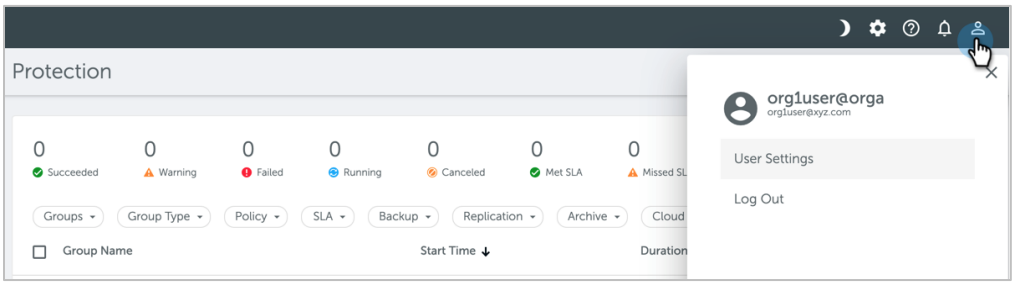

2. Go to **Settings > Targets > Remote Clusters**.

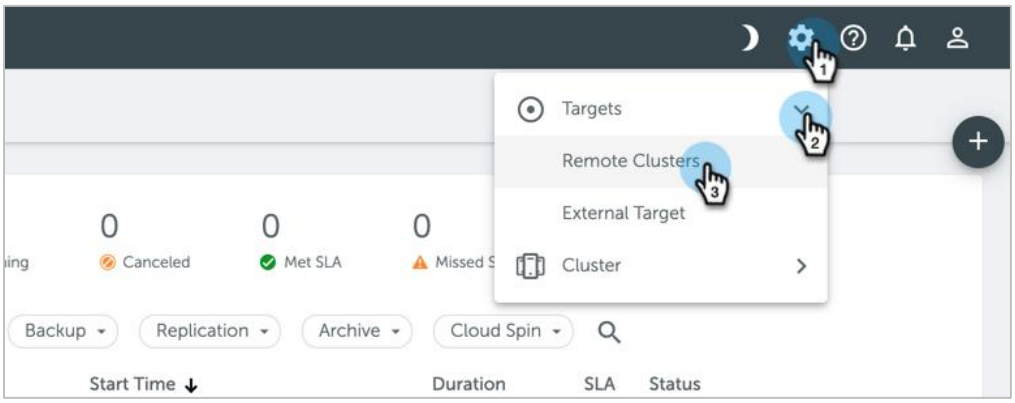

3. Click **Add Cluster**.

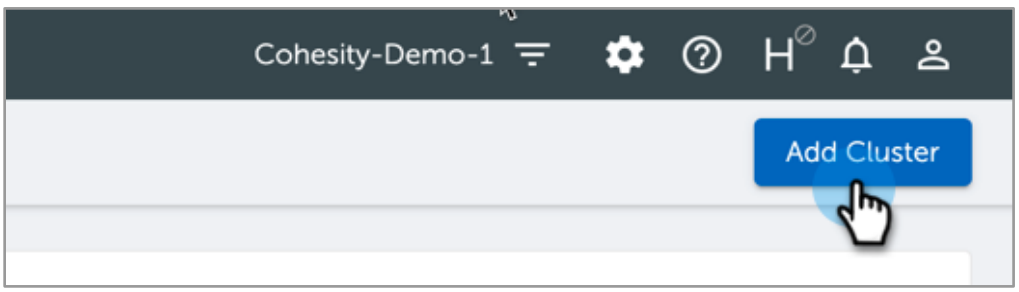

4. Enter **VIP or Node IP Address** for the tenant's primary cluster, the tenant administrator's **Username** and **Password,** and select **Interface Group** to enter the **Interface Group** name. Click **Connect** to validate the connection.

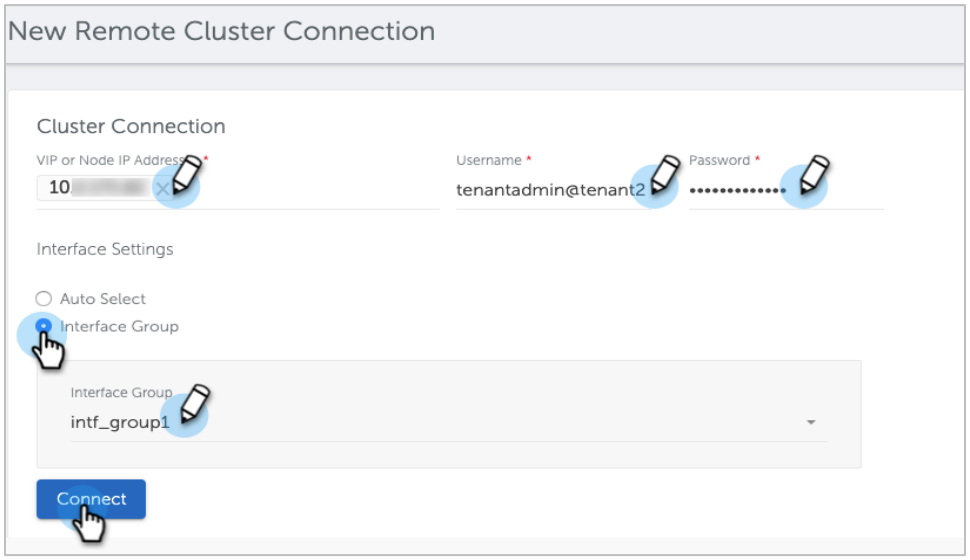

5. Under **Cluster Options**, enable **Replication** on and click **Add Storage Domain Pairing**.

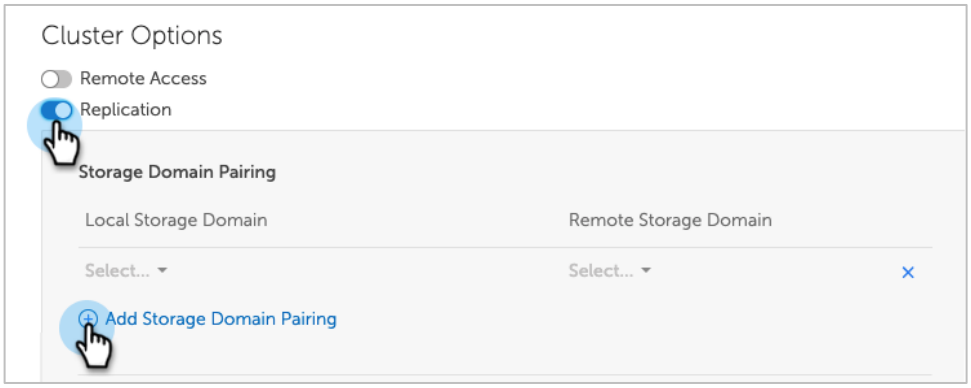

6. Select the **Local Storage Domain** and **Remote Storage Domain** that you intend to connect for replication.

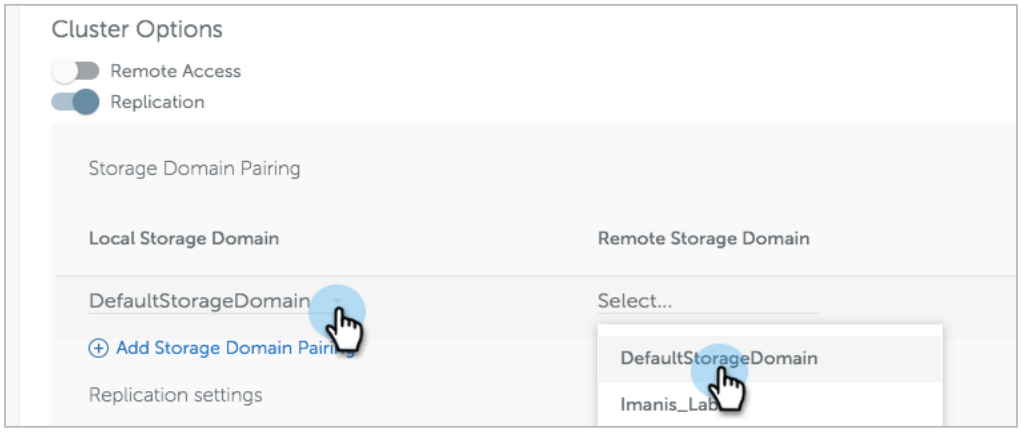

7. Under **Replication settings,** select **Distribute Load** and **Outbound Compression**. Click **Create** to create the replication connection.

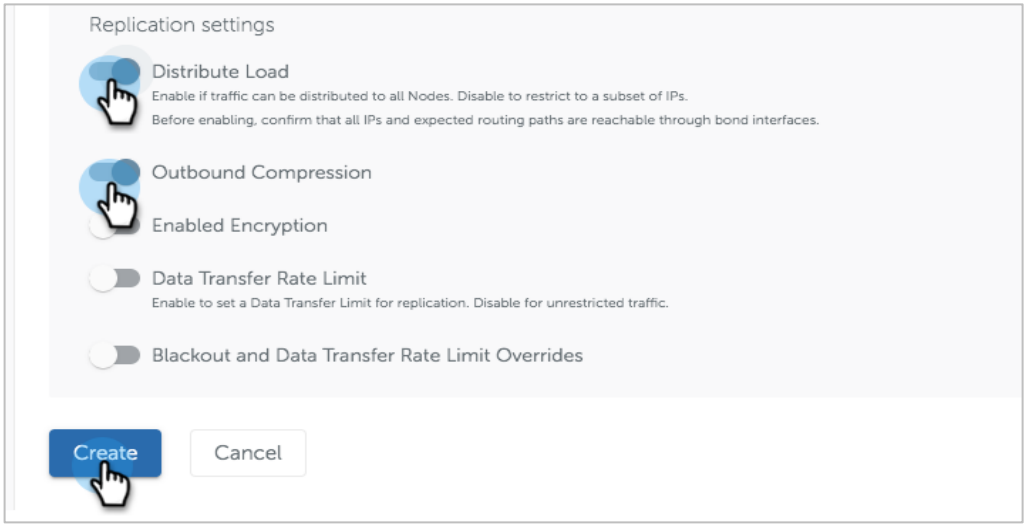

Once the connection is established between your backup cluster and DR cluster, proceed to Add Remote Cluster to Protection Policy Under Replication.

Add Remote Cluster to Protection Policy Under Replication

To replicate the source Cohesity cluster to the remote DR cluster, add the required replication settings to the Protection Policy on the source cluster.

1. Log in to the source Cohesity cluster and go to **Data Protection > Policies**.

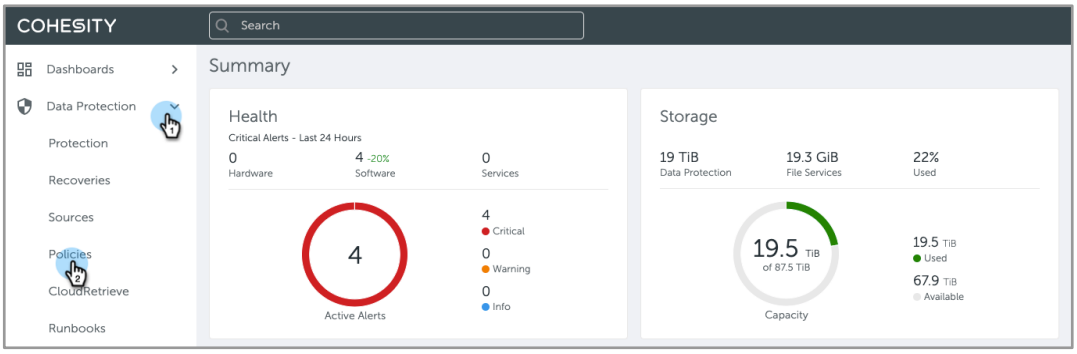

2. Under **Policies**, select the Policy to which you need to add replication and click the **Actions** menu to select **Edit Policy**.

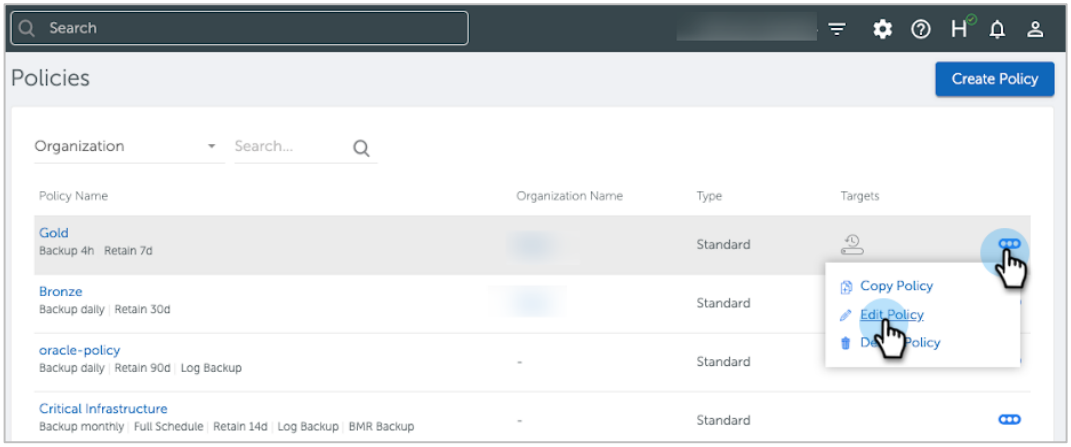

3. Select **Remote Cluster** and select the remote cluster from the drop-down.

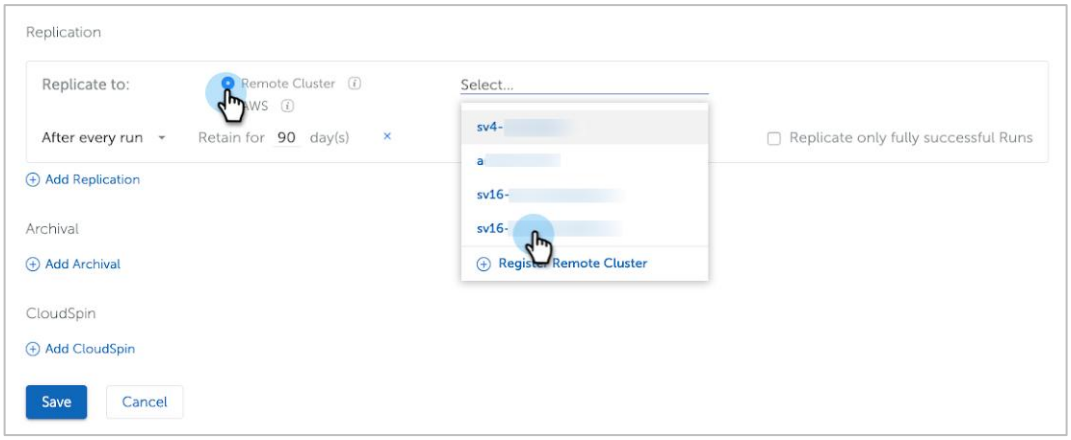

4. Select the schedule for replication and click **Save**.

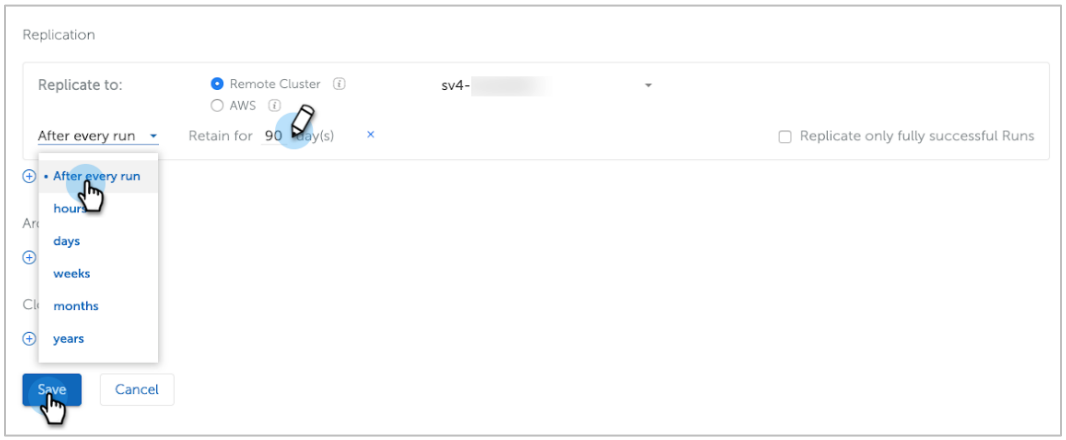

Assess Common Failure Scenarios and Disaster Recovery Workflows

You can recover data through Cohesity's DRaaS solution whenever necessary when the source cluster is unavailable. There are two workflows by which you can use the DR Cohesity cluster to recover data, depending upon the disaster scenario and where you want to recover the data.

1. **Recover Data**. To recover data from the DR cluster to either the primary workload infrastructure or DR workload infrastructure.
2. **Failover Protection Group and Recover Data**. Failover the Protection Group and recover the data, to provide ongoing protection. The failed-over Protection Group continues to back up the recovered data as per the schedule defined under the attached Protection Policy.

Table 3 below describes the DR scenarios that you might encounter, along with the DR options provided by Cohesity.

Table 4: Disaster Scenarios and Possible Outcomes with DR Procedure

SC.	SOURCE WORKLOAD INFRASTRUCTURE	SOURCE COHESITY CLUSTER	OUTCOME	DR PROCEDURE
1	No Failure	Failure	Failover Job (Skip Recovery): Back up source workload data when source Cohesity cluster is unavailable.	1. Add the source workload as source on DR cluster. 2. Failover the Protection Group 3. Add Policy to the DR Protection Group. 4. Skip the recovery flow.
2	No Failure but partial data loss	Failure	Recover to Source	Recover data to source workload when source Cohesity cluster is unavailable. 1. Add the source workload as source on DR cluster. 2. Use recovery workflow to recover the data.
			Failover the Protection Group and Recover to Source	Recover data to source workload and backup recovered data. 1. Add the source workload as source on DR cluster. 2. Failover the Protection Group 3. Add Policy to the DR Protection Group. 4. Continue to recover

SC.	SOURCE WORKLOAD INFRASTRUCTURE	SOURCE COHESITY CLUSTER	OUTCOME	DR PROCEDURE
3	Failure	Failure	Recover to DR	Recover data to DR workload Infrastructure. 1. Add the source workload as source on DR cluster. 2. Use recovery workflow to recover the data.
			Failover the Protection Group and Recover to DR workload Infrastructure	Recover data to DR Infrastructure and backup the recovered data. 1. Add the DR workload as a source on the DR cluster. 2. Failover the Protection Group 3. Add Policy to the DR Protection Group. 4. Continue to recover

Scenario 1: Active Primary Workload Infrastructure and Inactive Primary Cohesity Cluster

In a scenario when the primary workload infrastructure is active but the primary Cohesity cluster to which data is backed up is inactive, the DRaaS response is to failover the Protection Group and back up data from the primary workload to the remote DR cluster.

Failover Protection Group to DR Cluster and Back up the Source Workload on DR

When the primary Cohesity cluster is unavailable, failover the Protection Group to the remote DR cluster and resume the data backup. You can skip the recovery process and choose the objects to back up.
To address this failure scenario:

1. Add the primary workload as a source on the DR cluster.
2. Add Protection Policy and failover the Protection Group.
3. Add objects to the Protection Group.

Figure 9: Hosted Backup with Offsite Replication when Primary Cluster is Inactive

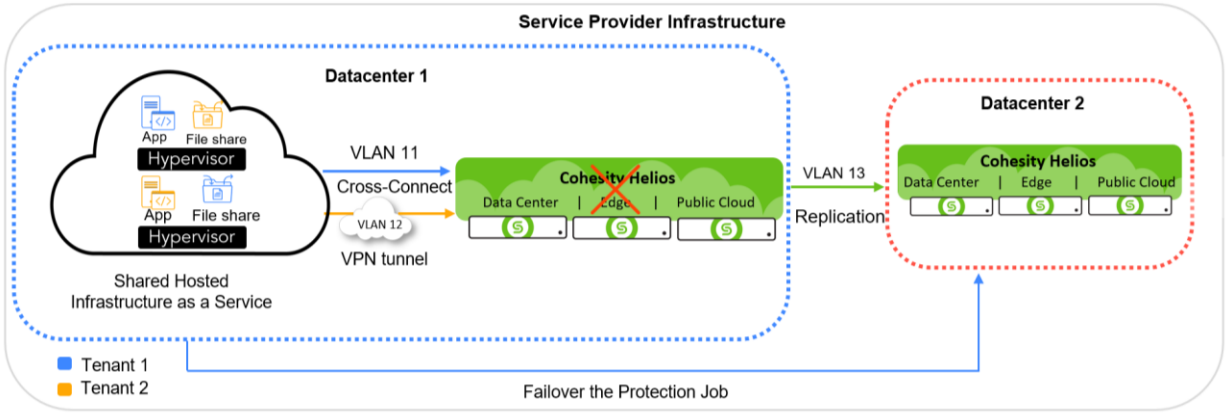

Figure 10: Local Backup with Offsite Replication when Primary Cluster is Inactive

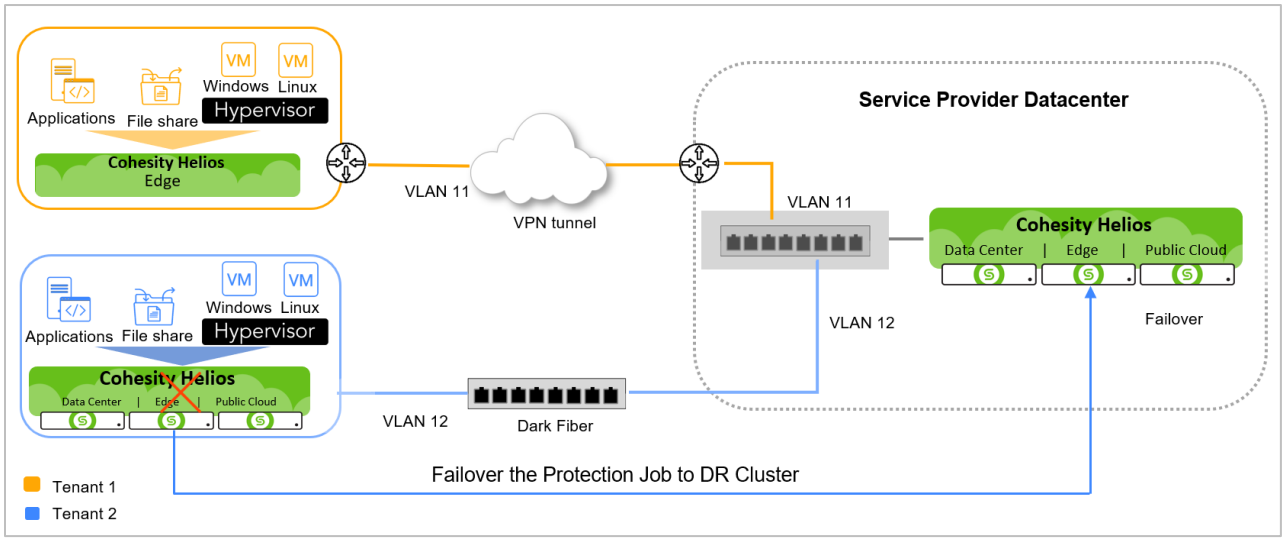

Scenario 2: Active Primary Workload Infrastructure with Partial Data Loss and Inactive Primary Cohesity Cluster

In a scenario when the primary workload infrastructure is active but the primary Cohesity cluster to which data is backed up is inactive, the DRaaS response is to failover the Protection Group and back up data from the primary workload to the remote DR cluster.

In this scenario, the two possible DRaaS responses are:

- Recover data to source without failover Protection Group.
- Failover Protection Group and recover data to source.

Recover Data to Primary Workload Infrastructure without Failover

To recover the data on the primary workload infrastructure using DR cluster:

1. Add the primary workload as a source on the DR cluster.
2. Use the recovery workflow to recover the data.

Figure 11: Hosted Backup with Offsite Replication to Recover Data to Primary Workload Infrastructure

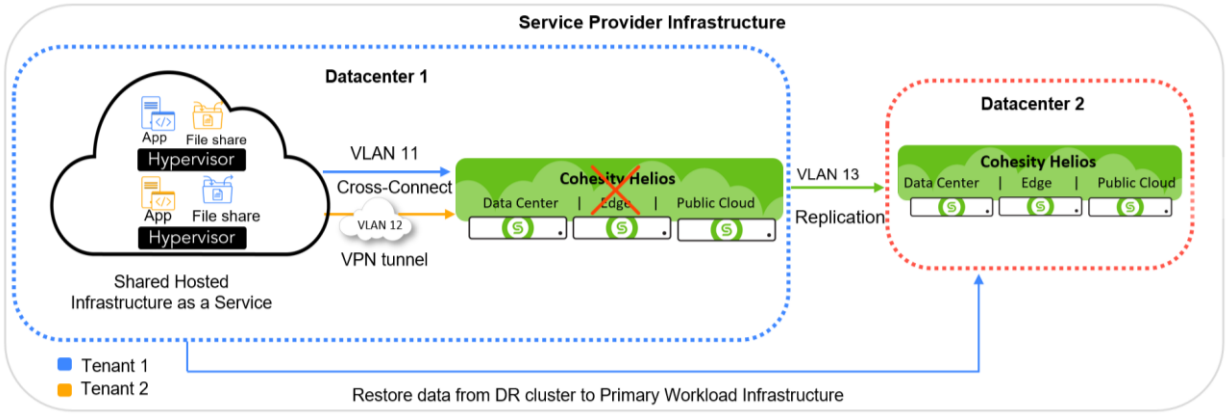

Figure 12: Local Backup with Offsite Replication to Recover Data to Primary Workload Infrastructure

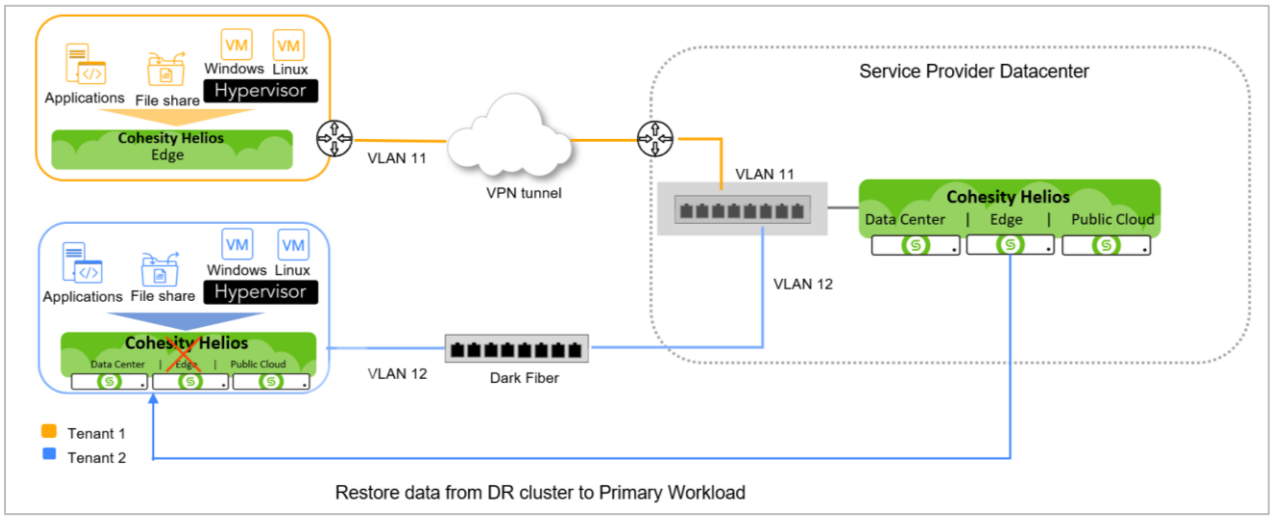

Failover to DR Cluster, Recover Data to Primary Workload Infrastructure and Continue to Back up

In this DRaaS outcome, failover the Protection Group on the remote DR cluster and recover data to the source. The failed-over Protection Group continues to back up the recovered data.

To achieve this outcome, the administrator needs to:

1. Add the primary workload as a source on the DR cluster.
2. Add a Protection Policy and failover the Protection Group.
3. Continue to recover data.

Figure 13: Hosted Backup with Offsite Replication to Recover Data to Primary Workload Infrastructure and Continue to Back up

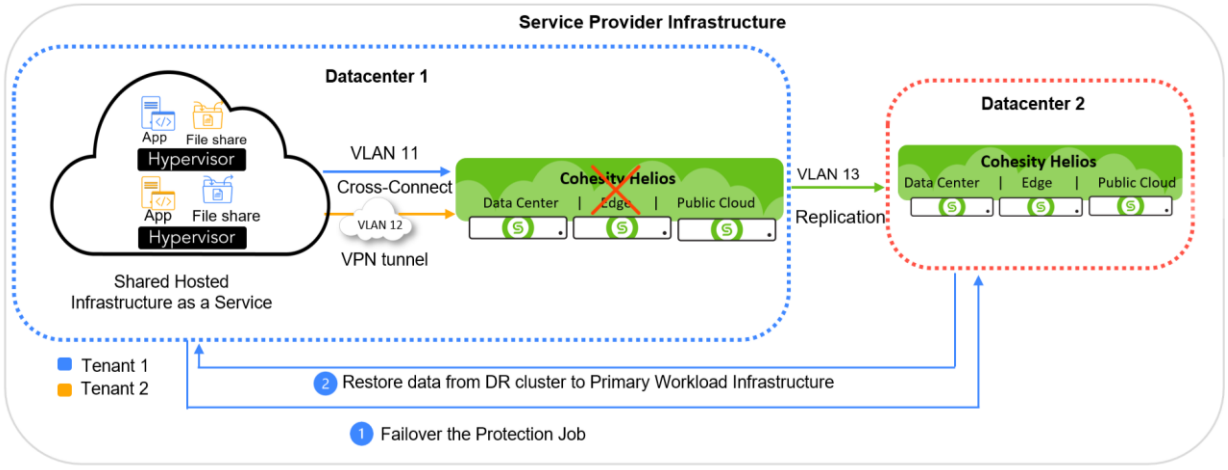

Figure 14: Local Backup with Offsite Replication to Recover Data to Primary Workload Infrastructure and Continue to Back up

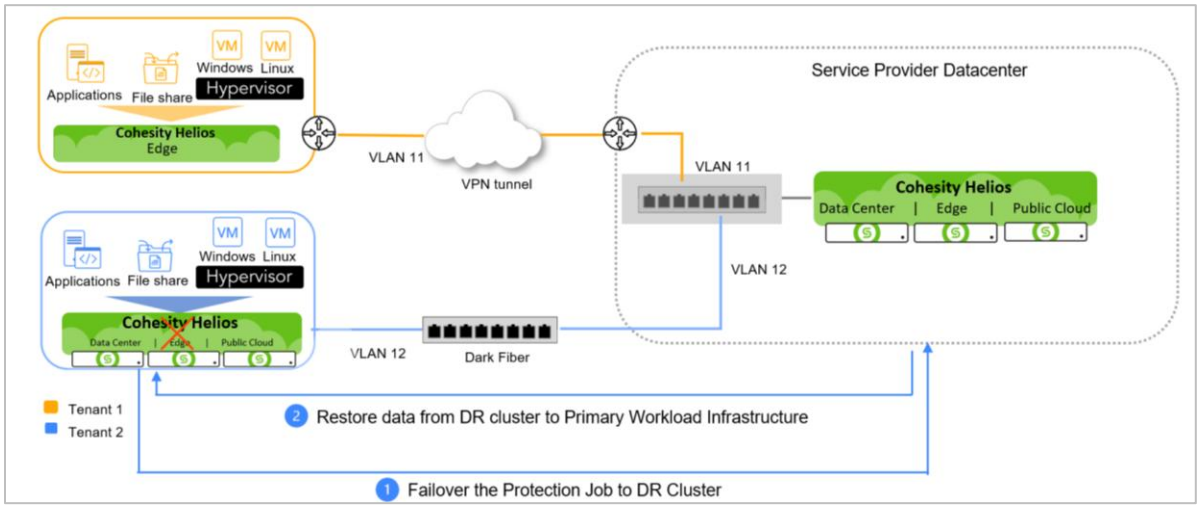

Scenario 3: Inactive Primary Workload Infrastructure and Primary Cohesity Cluster

In this scenario, the two possible DRaaS responses are:

- Recover data to a DR workload infrastructure.
- Recover data to DR workload infrastructure and continue backup.

Recover Data to DR Workload Infrastructure

To recover the data to the DR workload infrastructure using the DR cluster.

1. Add the DR workload as a source on the DR cluster.
2. Use the recovery workflow to recover the data.

Figure 15: Recover Data to DR Workload Infrastructure — Hosted Backup with Offsite Replication

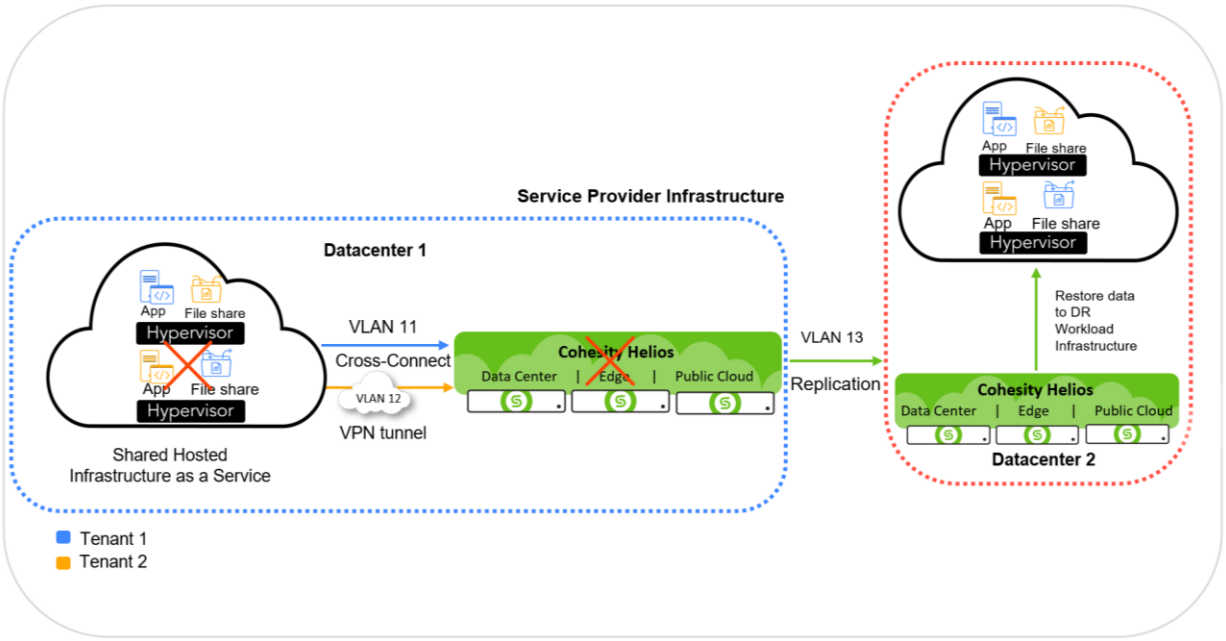

Figure 16: Recover Data to DR Workload Infrastructure — Local Backup with Offsite Replication

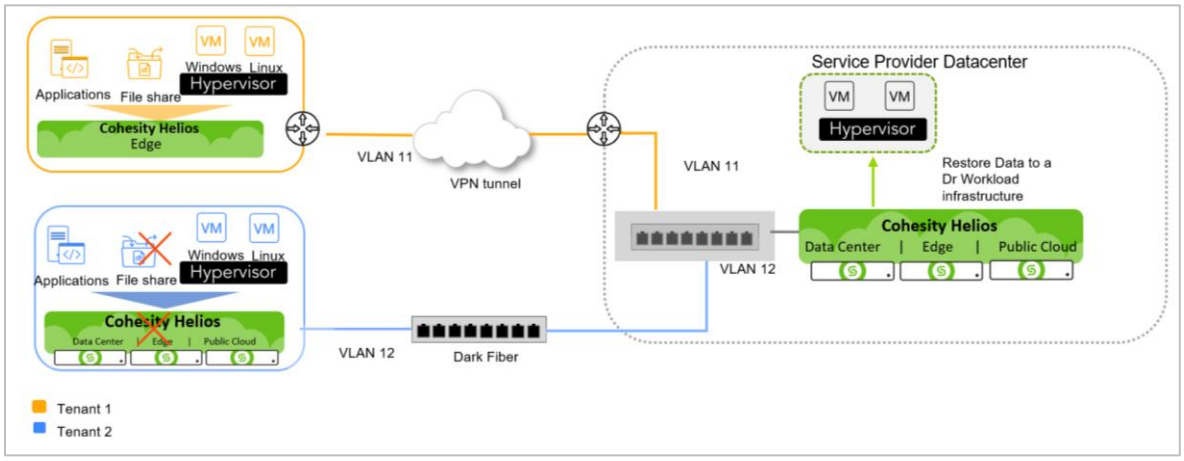

Recover Data to DR Workload Infrastructure and Continue to Backup

In this DRaaS outcome, failover the Protection Group to the DR Cohesity cluster and recover data on the DR workload infrastructure. Add the objects to the failed-over Protection Group and it continues to back up the recovered data.

To achieve this outcome, the administrator needs to:

1. Add the DR workload as a source on the DR cluster.
2. Add a Protection Policy and failover the Protection Group.
3. Continue to recover data.

Figure 17: Recover Data to DR Workload Infrastructure and Continue to Back up — Hosted Backup with Offsite Replication

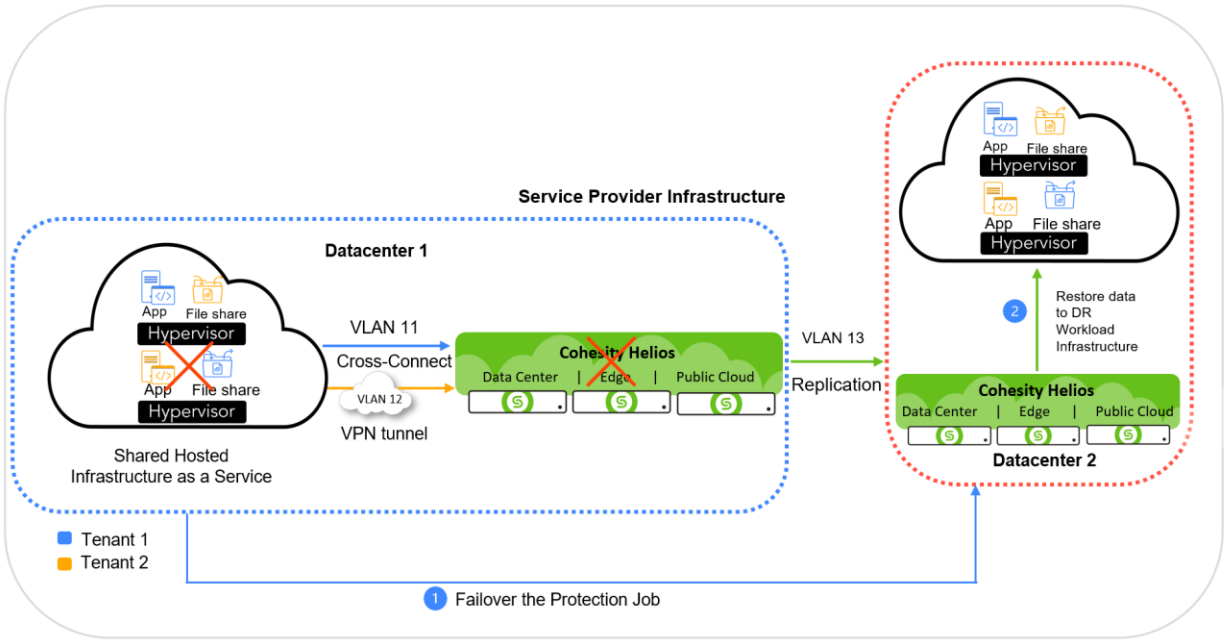

Figure 18: Recover Data to DR Workload Infrastructure and Continue to Back up — Hosted Backup with Offsite Replication

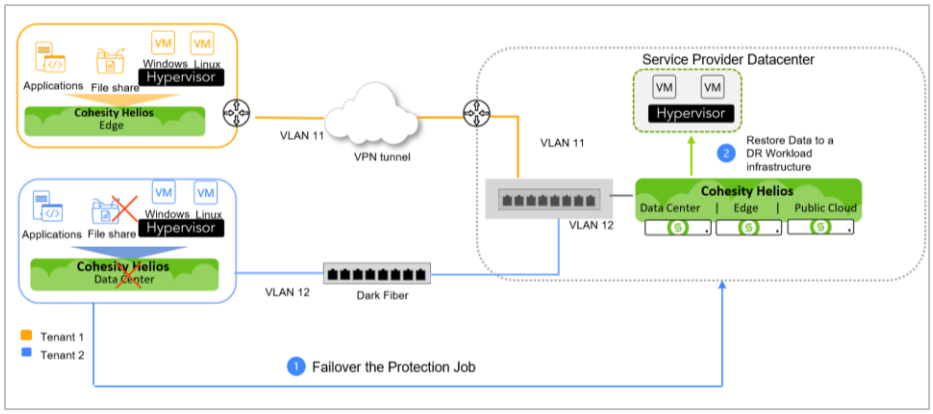

Add a Recovery Source

Once the data is secured in the DR cluster, Cohesity allows seamless and simple recovery to a specified destination cluster whenever necessary. For this, add a source cluster to the DR cluster to prepare it prior to recovery.

To prepare to recover to a primary source workload infrastructure, add your source workload to the DR Cohesity cluster. As illustrated in Figure 19, from there, you can recover to the primary workload infrastructure or failover and recover.

Figure 19: Recover Data to Primary Workload Infrastructure

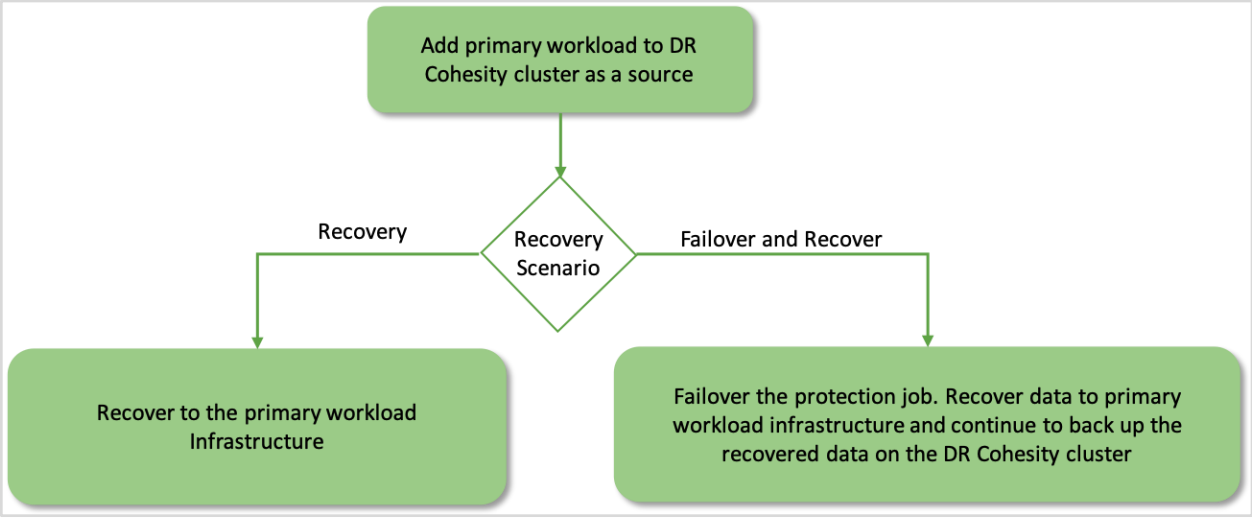

To prepare to recover to a DR workload infrastructure, add your DR workload to the DR Cohesity cluster. As illustrated in Figure 20, from there, you can recover to the DR workload infrastructure or failover and recover.

Figure 20: Recover Data to a DR Workload Infrastructure

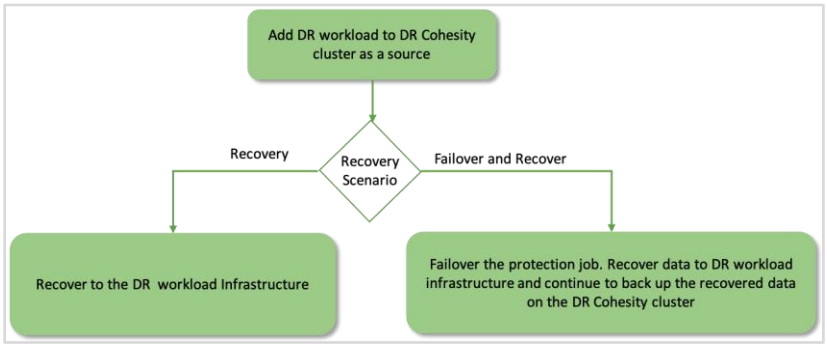

Add a Source to the DR Cohesity Cluster

Depending upon the DRaaS deployment, either the service provider administrator or tenant administrator should add the recovery source to the DR cluster.
To add a source to a DR cluster in:

- Hosted Backup with Offsite Replication topology, see Add a Source in 'Hosted Backup with Offsite Replication' Topology.
- Local Backup with Offsite Replication topology, see Add a Source in 'Local Backup with Offsite Replication' Topology.

Add a Source in 'Hosted Backup with Offsite Replication' Topology

If you have a *Hosted Backup with Offsite Replication* topology, the service provider administrator needs to assign a recovery source to the DR cluster.

To assign a recovery source:

1. Log in to the DR Cohesity cluster with a service provider administrator user.
2. Go to **Data Protection > Sources**.

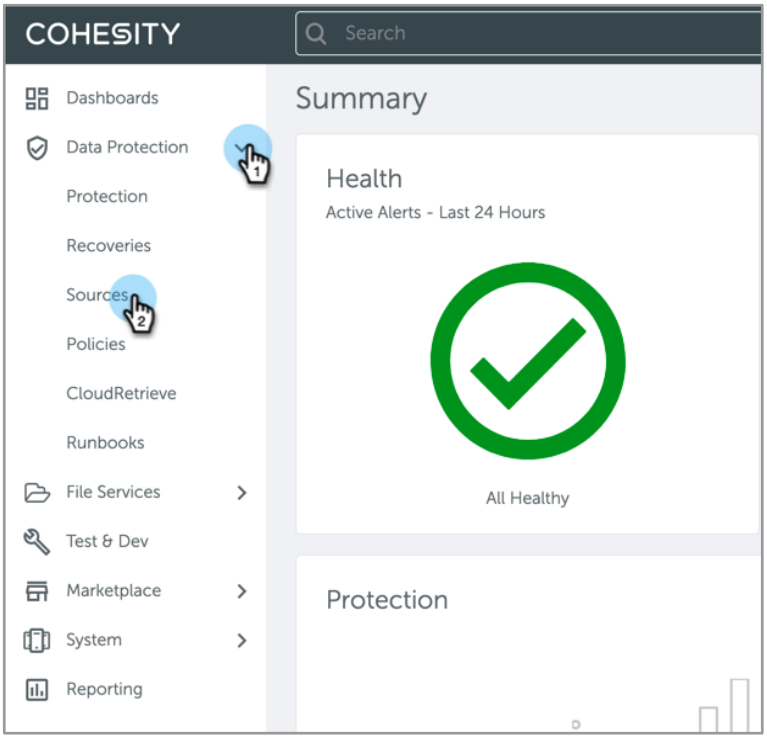

3. Click **+** and select a source application. For example, to add a vCenter, select **Virtual Server** and follow the wizard to add the source.

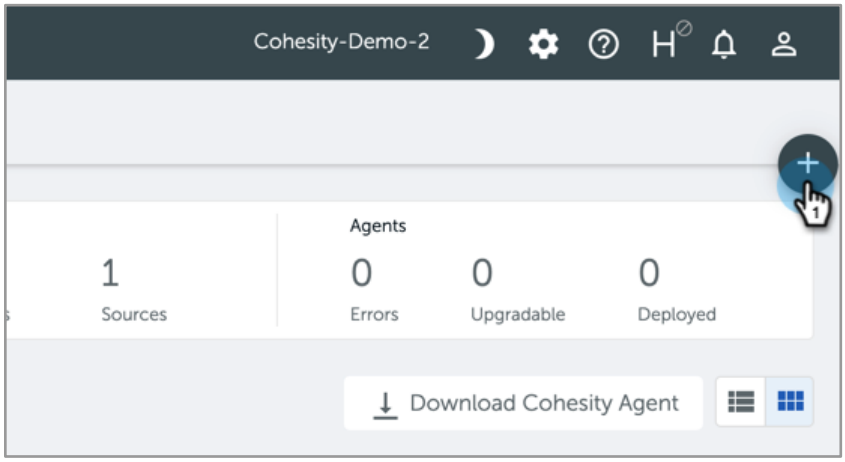

Once you add a source, add it to an organization that requires the data restore.

4. Go to **Settings > Cluster > Organizations**.

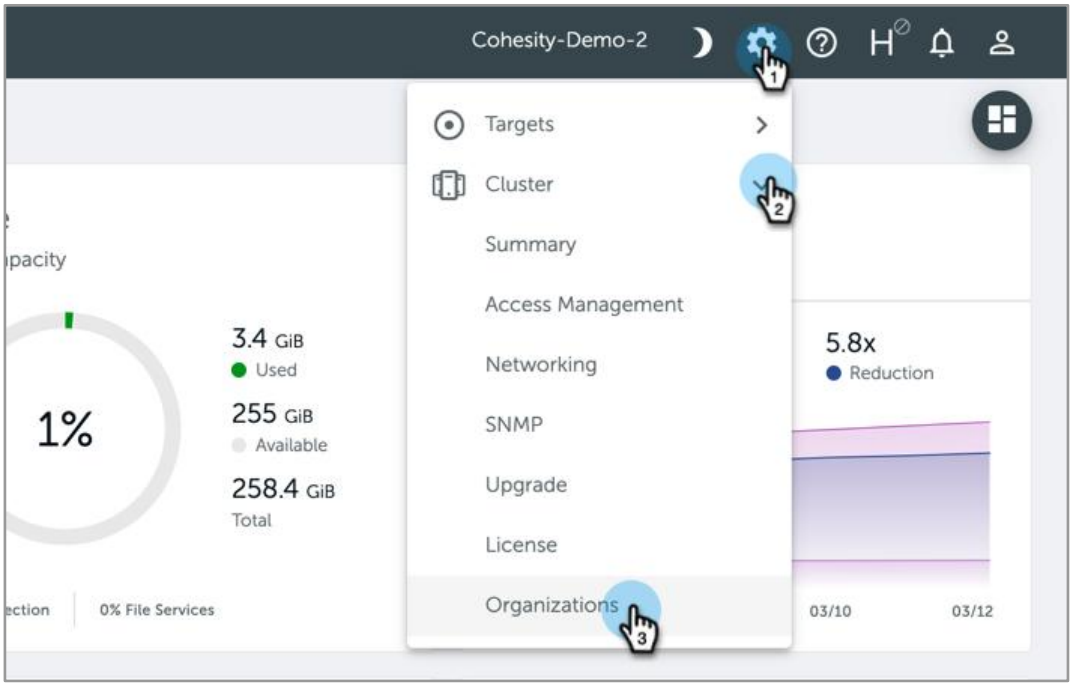

5. Select **Organizations**. Under **Organization Name,** select the organization from the list.

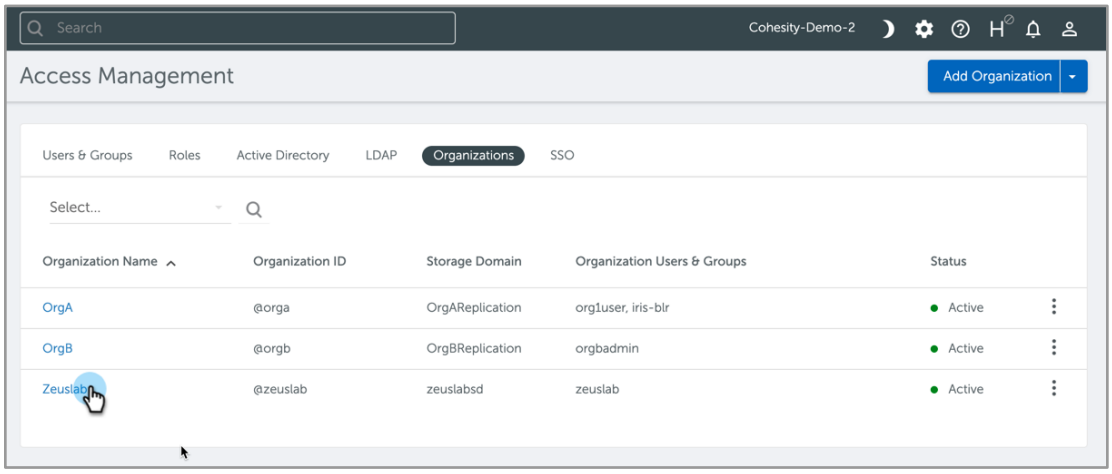

6. On the **Organization Details** page, click **Edit** on top right.

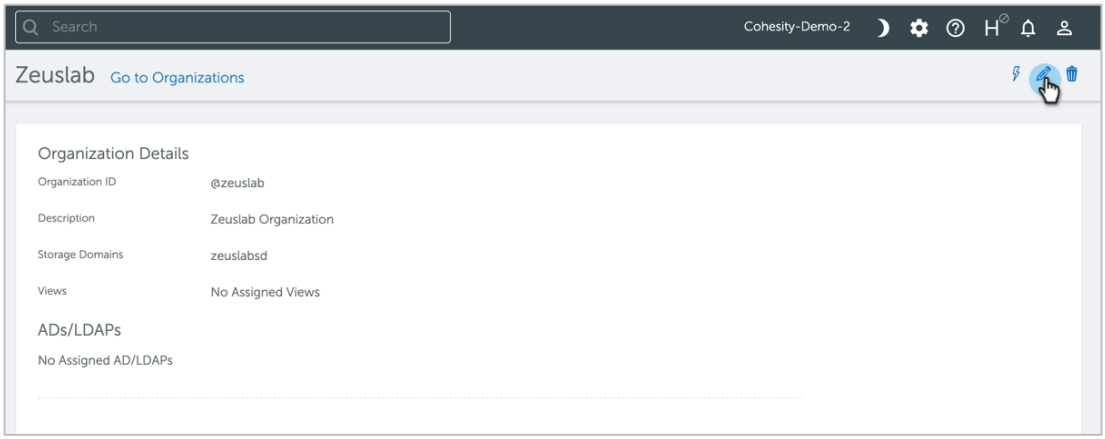

7. Click **Assign Objects**.

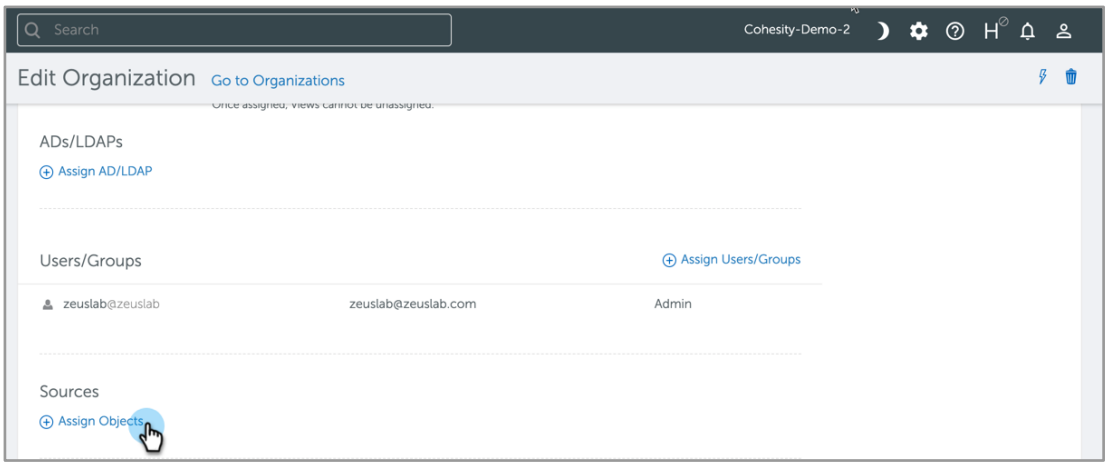

8. Select the source you added and click **Assign** to assign it to the tenant.

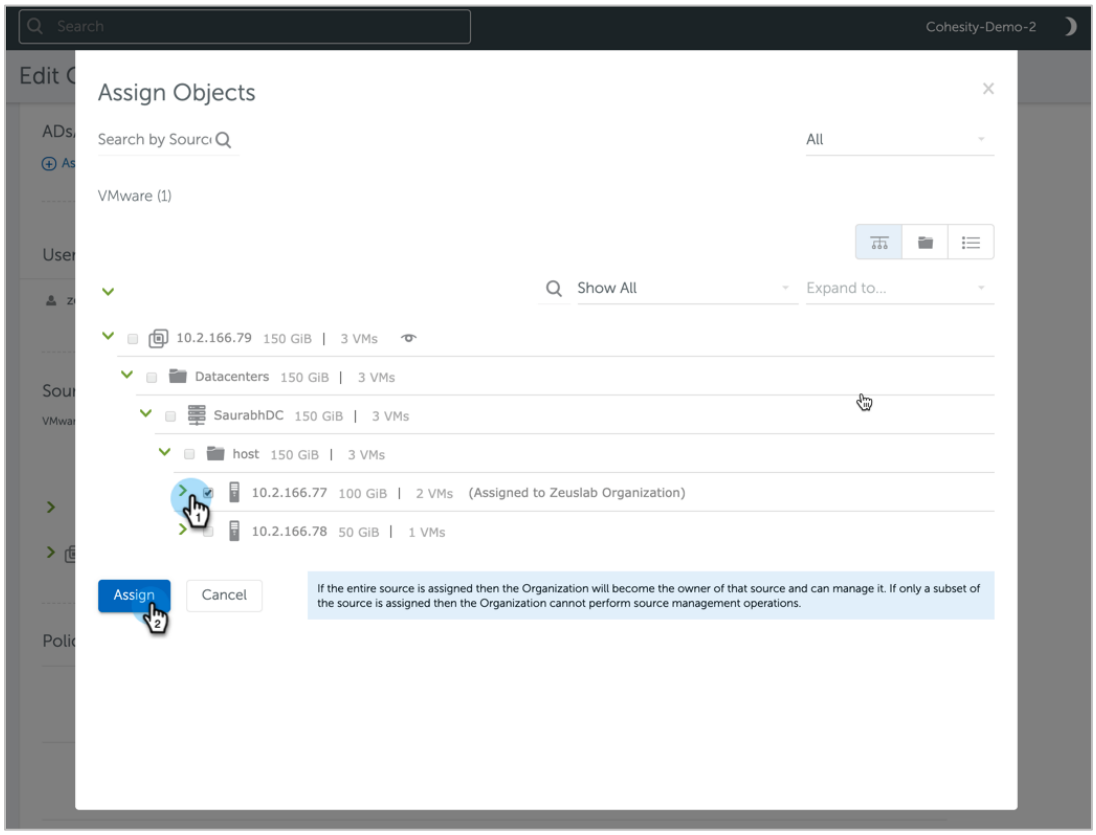

9. Once the source is registered, you can either recover the data or failover and recover data from the Protection Group.

Add a Source in 'Local Backup with Offsite Replication' Topology

For the *Local Backup with Offsite Replication* scenario, the tenant or the service provider administrator can assign a recovery source to the DR cluster.

To assign a recovery source:

1. Log in to the DR Cohesity cluster as tenant administrator or as a service provider impersonating the tenant administrator. Navigate to **Data Protection > Protection**.
2. Click **+** and select the source you want to register to recover data. For example, to add a vCenter, select **Virtual Server**.

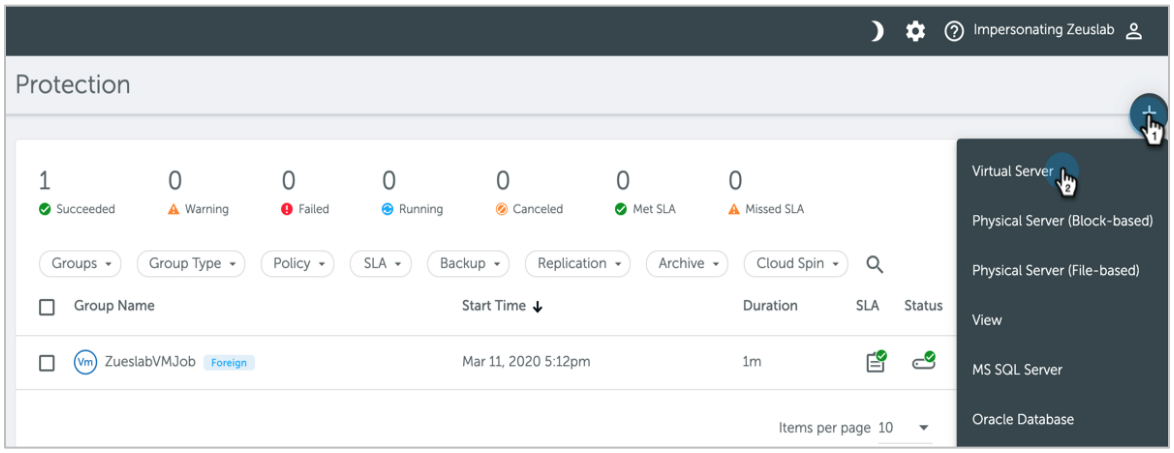

   Follow the wizard to add the source.

3. Once the source is added, you can either recover the data or failover and recover data from the Protection Group.

Recover Your Data

Depending upon your choice, the replicated data from the DR cluster can be recovered in two ways:

1. Recover the Data from a DR Protection Group.
2. Failover the DR Protection Group and recover data to resume backup of the recovered data.

Recover Data from a DR Protection Group

You can recover data from the DR cluster either to source workload or the DR workload infrastructure.

1. Log in to the DR Cohesity cluster as tenant administrator or as a service provider impersonating the tenant administrator. Navigate to **Data Protection > Protection**.
2. Click the Protection Group with your protected data.

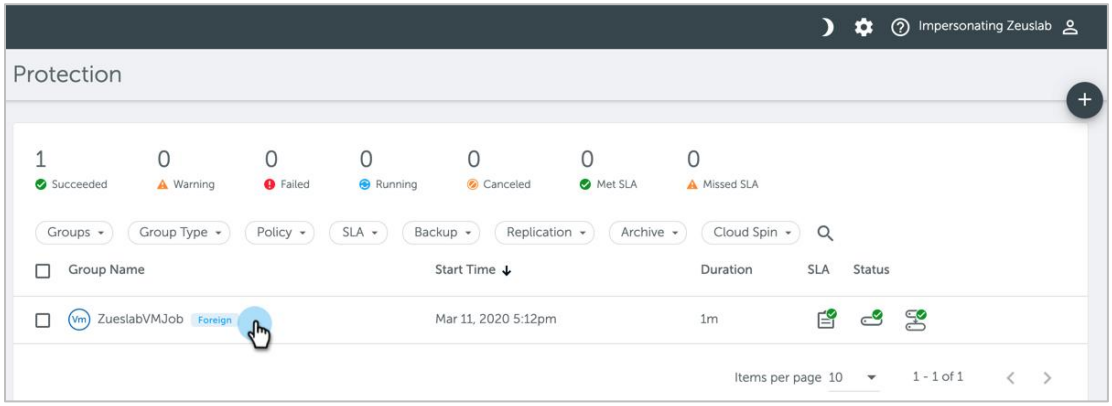

3. Select the date to which you want to restore the objects.

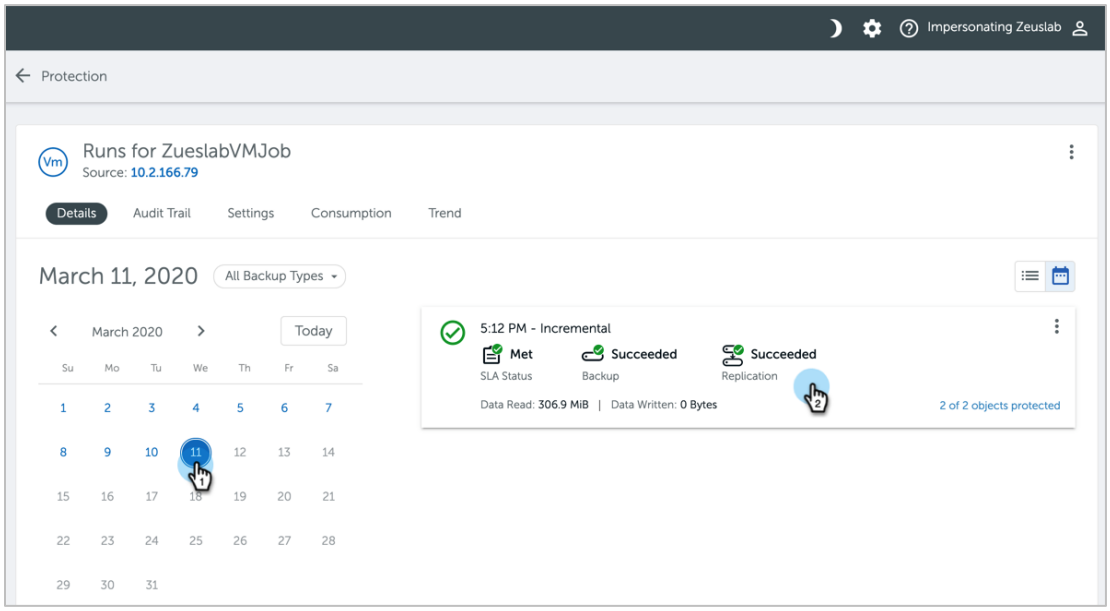

4. Select the VM to be restored and click **Recover VM**.

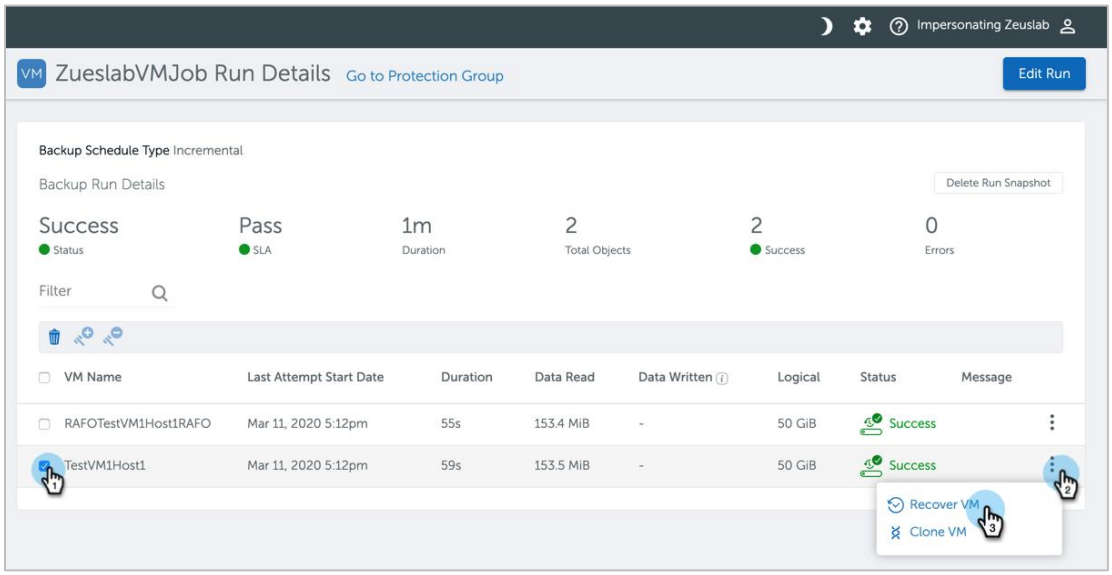

5. Select **Recovery Location,** choose whether you want to **Detach network** or **Attach to a new network**, and click **Finish**.

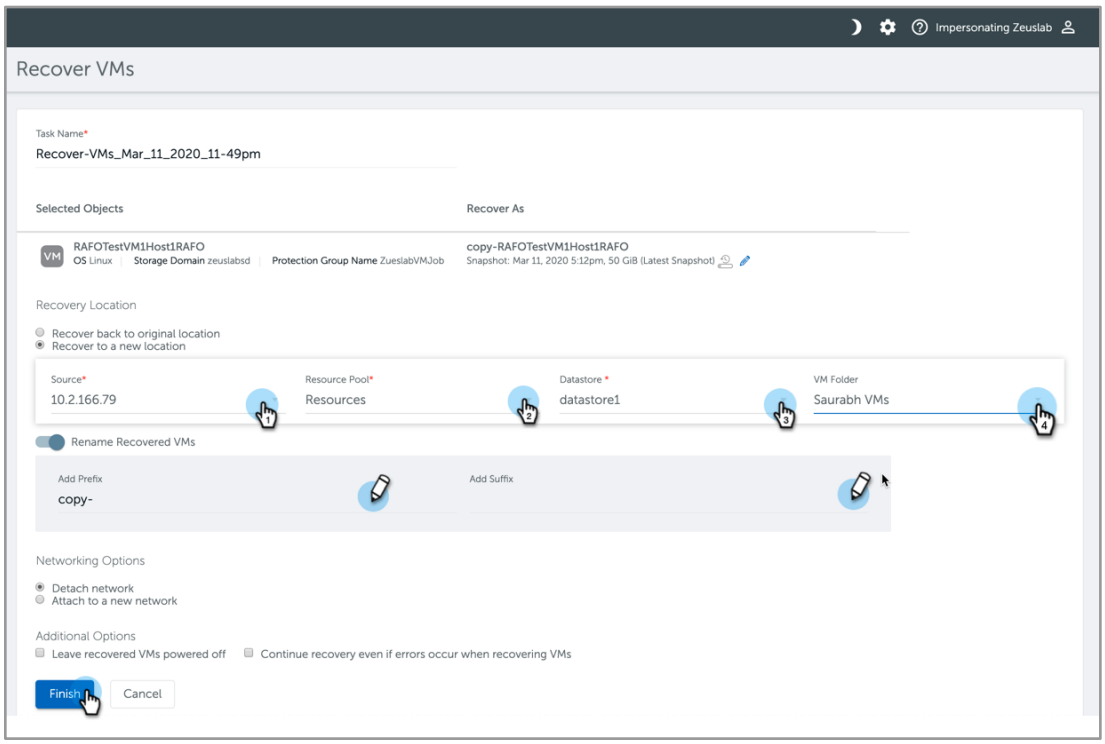

Failover to a DR Protection Group and Recover Data

To recover the data and continue to back up the recovered data, use the failover and recover data workflow.

To failover a protection job, log in to the DR Cohesity cluster as tenant administrator or as a service provider impersonating the tenant administrator.

1. Go to **Data Protection** > **Protection**.
2. On the **Protection** page, select **Failover** against the protection group.

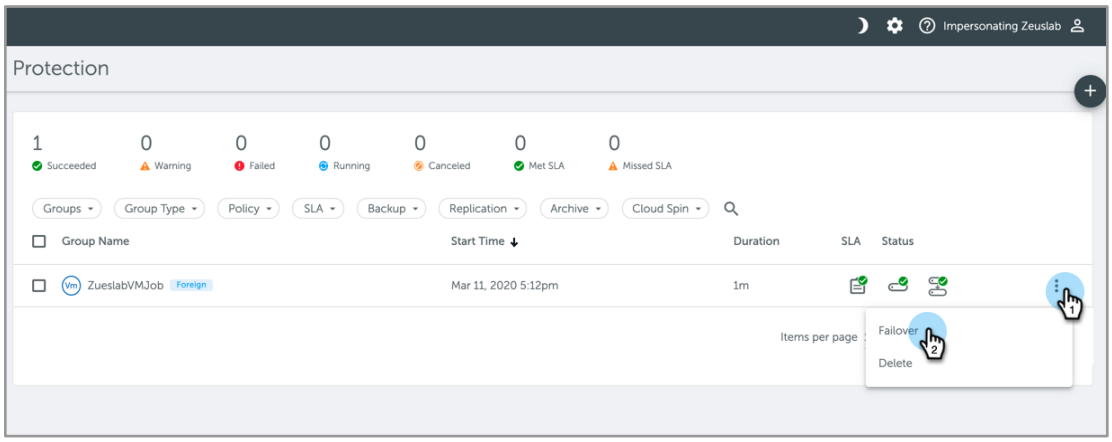

3. Select the source for **Failover to Source** and **Policy.** Click **Failover Protection Group and Continue to Recovery**.

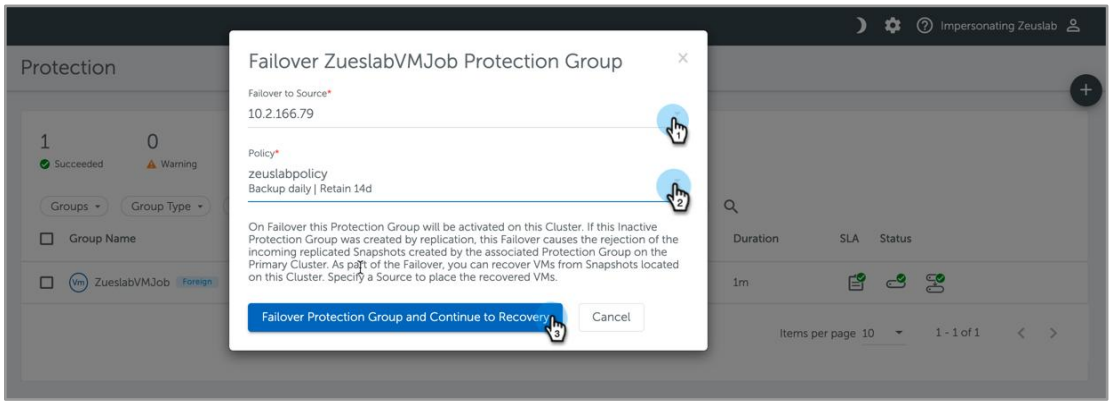

Once failover is successful, you can use the Protection Group to recover data to the added source.

4. To start recovering data, select the **Recovery Location,** and choose whether you want to **Detach network** or **Attach to a new network**, **Leave recovered VMs powered off,** and **Continue recovery even if error occurs when recovering VMs**. Once you're done, click **Finish**.

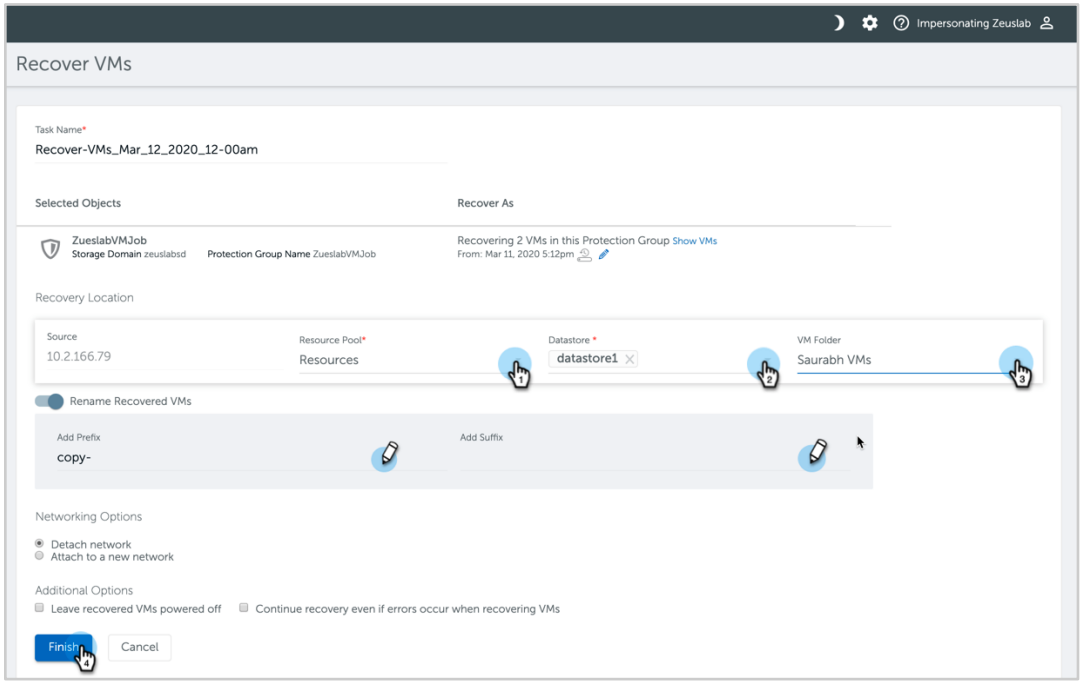

Reporting and Chargeback

Reports in Cohesity offer another good tool for monitoring, auditing, troubleshooting, and exception monitoring of the Cohesity clusters. This helps you comply with your business SLAs, manage day-to-day operations, and plan for upcoming challenges.

Cohesity provides many built-in reports that help you with planning, charge-back, compliance, and more. Some of the other benefits of Cohesity Reports include:

- The ability to schedule reports to be sent on a regular basis.
- Filter reports based on type, operation, time frame, or tenant organization.
- Integrate with third-party business intelligence tools.

Report Tenant Storage Consumed (< v6.4.1)

Prior to 6.4.1, configure a Storage Domain for each tenant specifically for replication on the DR cluster/tenant. Use **Storage Consumed by Storage Domains** report to account for the storage used for replication by each tenant.

Figure 21: Report Tenant Storage Consumed (< v6.4.1)

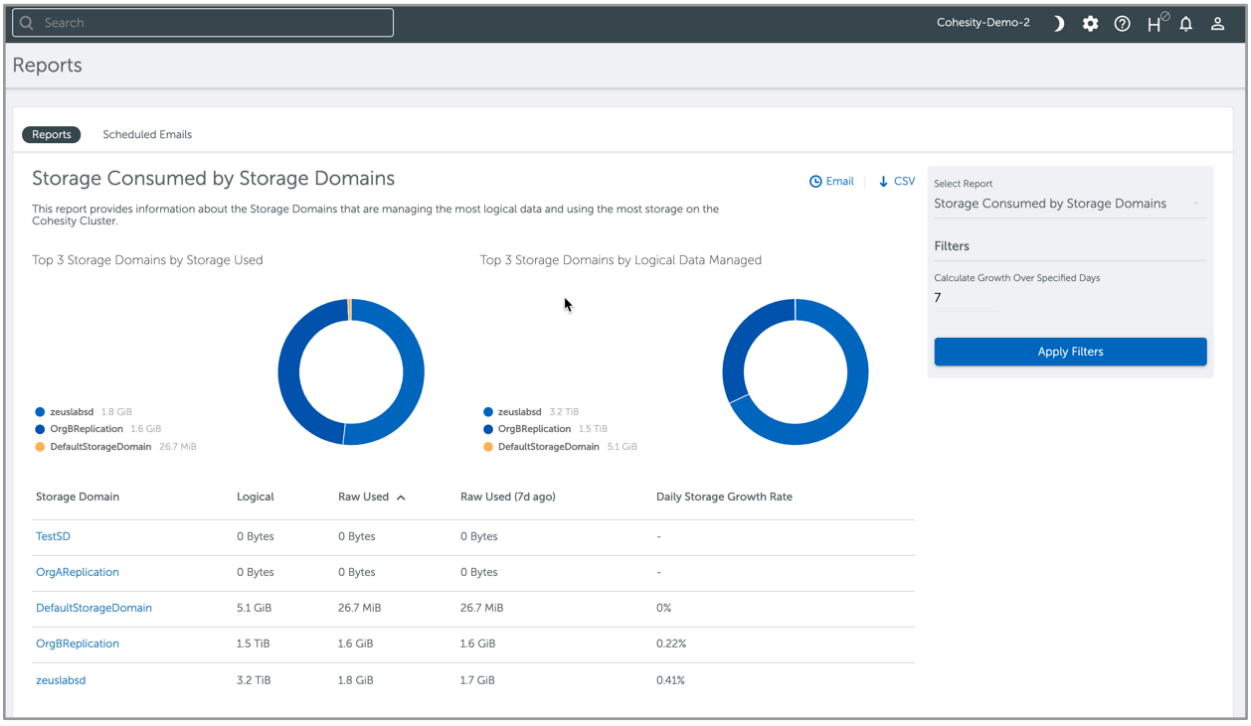

Report Tenant Storage Consumed (v6.4.1 and higher)

From 6.4.1, configure a shared Storage Domain for replication to be used across tenants or use dedicated storage domains per tenant. Service providers can then use the Storage Consumed per Organization report on the shared Storage Domain to get the information on storage consumed by each tenant for replication. If the service provider is using a dedicated Storage Domain per tenant, they can use Storage Consumed by Storage Domains.

Figure 22: Report Tenant Storage Consumed (v6.4.1 and higher)

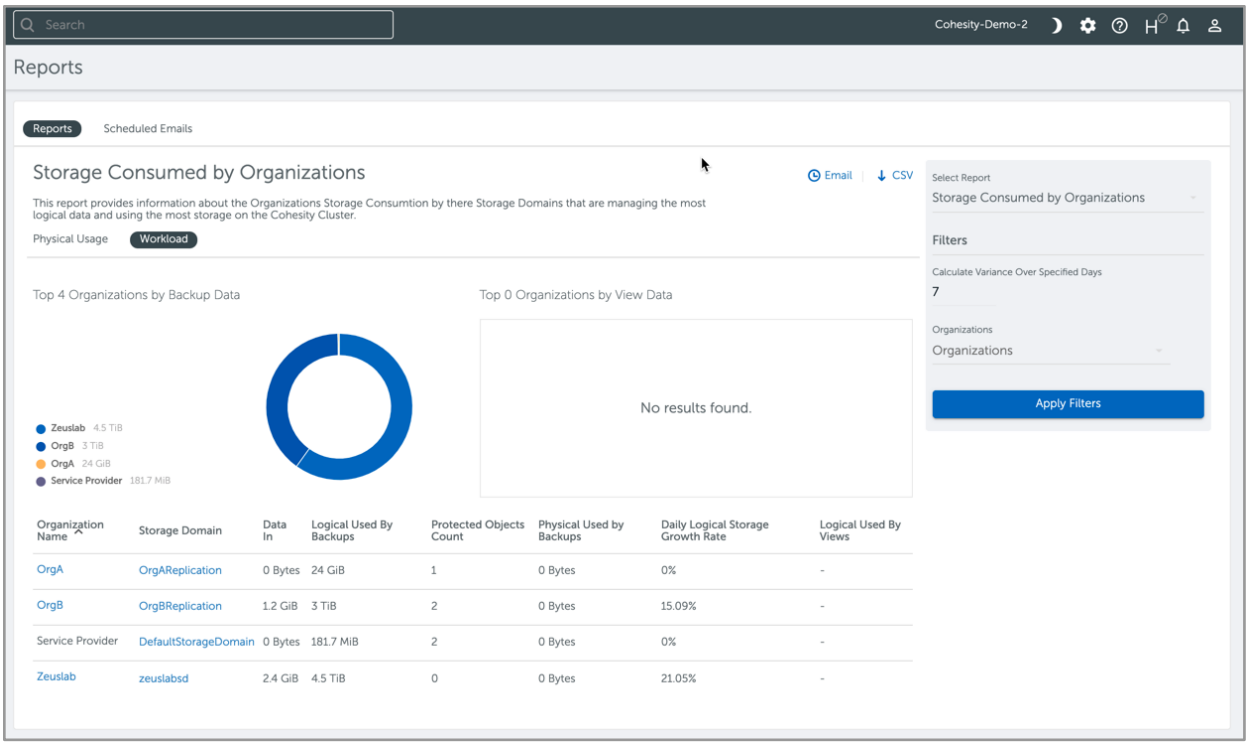

Best Practice Considerations

For service providers who are implementing a BaaS solution using Cohesity, Cohesity recommends you segregate replication and backup traffic on the source cluster to avoid bandwidth issues.

Cohesity also recommends having a VPN connection between the source and DR clusters to set up replication across the clusters if the clusters are not using public IPs for replication communication.

Alternatively, you can replicate over WAN without the need for a VPN connection.

Appendix A: Privileges for DRaaS Management

The service provider admin, as well as any user who is registered as an administrator from the tenant side, receives a certain set of default privileges, listed in the tables below.

Table 5: Administrator Privileges for Hosted Backup with Offsite Replication

RESOURCE	SP ADMINISTRATOR		TENANT ADMINISTRATOR	
	ADD	DELETE	ADD/CREATE	DELETE
Remote Cluster	Yes	Yes	No	No
Policy	Yes	Yes	No	No
Protection Group	Yes	Yes	Yes	Yes

Table 6: Administrator Privileges for Local Backup with Offsite Replication

RESOURCE	SP ADMINISTRATOR		TENANT ADMINISTRATOR	
	ADD	DELETE	ADD/CREATE	DELETE
Remote Cluster	Yes	Yes	Yes	No
Policy	Yes	Yes	Yes	No
Protection Group	Yes	Yes	Yes	Yes

Appendix B: Data Isolation

Service providers can assign either shared or dedicated Storage Domains per organization, thereby isolating the data for each organization. Service providers can also provide separate encryption keys by assigning each organization a separate Storage Domain.

> **NOTE**: If you use a dedicated Storage Domain for each tenant organization, you will not have global deduplication across tenants, as deduplication occurs at the Storage Domain level.

Appendix C: Connection Creation Checklist for ‘Hosted Backup and Offsite Replication’

Before creating a remote connection between the source and DR clusters, run this checklist through to make sure the administrator has the necessary information required.

Table 7: Connection Creation Checklist for Hosted Backup and Offsite Replication

CONNECTION CREATION CHECKLIST	
VIP address for both source and remote cluster	☑
Organization pairs created on both source and remote cluster with the same Organization IDs. For example, the source cluster has an organization (name: Firsttenant, Organization ID: tenant1) that needs to be replicated. The remote cluster should also have an organization with name: ***<any name>***, Organization ID: tenant1.	☑
Administrator credentials for both source and remote cluster	☑
Interface group name to be used for replication traffic on source and remote cluster	☑
Storage Domain mapping between Storage Domains on source and remote clusters	☑

Appendix D: Related Topics

Please visit Cohesity Partner Portal and review recommended documentation materials in the Service Provider section.

- Service Provider Solutions
- Backup as a Service Solution Guide
- Optimal Network Designs with Cohesity
- Cohesity Multi-tenancy Guide
- Multi-tenancy Best Practices
- Cohesity Encryption Best Practices

Your Feedback

Was this document helpful? Send us your feedback!

About the Authors

Jedidiah Sonavane is a Solutions Architect, STAT, and a part of Data Protect COE at Cohesity. He focuses on Service Providers/Organizations, Cloud Archive On-Prem, Gaia. His work proofs of concepts, enterprise data protection, solution validation, solution design, testing, qualification, and ensuring software usability. He collaborates closely with teams to tailor solutions that meet customer needs while adhering to industry standards and best practices.

Other essential contributors include:

- Yu-Shen Ng, Product Manager
- Palanivel Rajan, Product Manager
- Navaid Khan, Service Provider CTO
- Sourish Mazumdar, Cohesity Engineering
- Mayank Narula, Cohesity Engineering

Document Version History

VERSION	DATE	DOCUMENT HISTORY
1.3	Dec 2025	Republishing with latest template
1.2	July 2024	Republishing
1.1	Jan 2022	Rebranding updates
1.0	Mar 2020	First release

ABOUT COHESITY

Cohesity is a leader in AI-powered data security and management. Aided by an extensive ecosystem of partners, Cohesity makes it easier to protect, manage, and get value from data – across the data center, edge, and cloud. Cohesity helps organizations defend against cybersecurity threats with comprehensive data security and management capabilities, including immutable backup snapshots, AI-based threat detection, monitoring for malicious behavior, and rapid recovery at scale. Cohesity solutions are delivered as a service, self-managed, or provided by a Cohesity-powered partner. Cohesity is headquartered in San Jose, CA, and is trusted by the world's largest enterprises, including six of the Fortune 10 and 44 of the Fortune 100.

Visit our website and blog, follow us on Twitter and LinkedIn and like us on Facebook.