




DataProtect as a Service User Guide

October 23, 2024



© 2024 Cohesity, Inc. All rights reserved.

Cohesity, the Cohesity logo, SnapTree, SpanFS, DataPlatform, DataProtect, Helios, the Helios logo, DataGovern, SiteContinuity, DataHawk, and other Cohesity marks are trademarks or registered trademarks of Cohesity, Inc. in the US and/or internationally. Other company and product names may be trademarks of the respective companies with which they are associated. This material (a) is intended to provide you information about Cohesity and our business and products; (b) was believed to be true and accurate at the time it was written, but is subject to change without notice; and (c) is provided on an "AS IS" basis. Cohesity disclaims all express or implied conditions, representations, warranties of any kind.

No part of this documentation or any related software may be reproduced, stored, transmitted, or otherwise distributed in any form or by any means (electronic or otherwise) for any purpose other than the purchaser's personal use without the prior written consent of Cohesity, Inc. You may not use, modify, perform or display this documentation or any related software for any purpose except as expressly set forth in a separate written agreement executed by Cohesity, Inc., and any other use (including without limitation for the reverse engineering of such software or creating compatible software or derivative works) is prohibited, except to the extent such restrictions are prohibited by applicable law.

Published on October 23, 2024

Contents

Cohesity Data Cloud	7
Pillars	7
Protection	7
Security	8
Mobility	9
Access	10
Insights	10
Switch Between Apps	11
Set Default Landing Page	12
Breadcrumbs	15
Set User Preferences	16
Global Dashboard	18
Cohesity DataProtect Delivered-as-a-Service	23
What's New	24
October 2024	24
September 2024	24
July 2024	26
June 2024	27
May 2024	27
April 2024	28
March 2024	28
February 2024	28
January 2024	28
December 2023	32
October 2023	34
September 2023	35
August 2023	35
June 2023	36
April 2023	36
March 2023	36
January 2023	37
December 2022	37
August 2022	38
July 2022	38
March 2022	39
February 2022	39
October 2021	40

September 2021	40
August 2021	40
July 2021	40
June 2021	41
May 2021	41
April 2021	41
March 2021	42
Supported Software for Cohesity DataProtect as a Service	43
VMware	43
VMware Cloud Infrastructure	44
Microsoft Hyper-V	44
Physical Servers	45
Microsoft 365 Editions	45
Microsoft SQL Server	47
Oracle	47
SAP HANA	48
NAS	48
Cloud Services	49
Supported Workloads and Cloud Regions	50
Supported Cloud Regions	50
Supported Workloads and Cloud Providers	51
Supported Regions for AWS SaaS Connectors	53
Supported Regions for Azure SaaS Connectors	53
Get Started	54
Sign in to Cohesity DataProtect as a Service	54
Select Regions and Encryption Key Management System	57
Add Users	70
Add On-Premise SaaS Connections	72
Register a Source	73
Protect a Source	75
Recover Protected Objects & Files	76
Deploy SaaS Connector	82
User-Deployed SaaS Connectors	82
Cohesity-Deployed SaaS Connectors	110
Manage Network Bandwidth Usage	126
Configure SaaS Connector Alert Notifications	128
Enable SaaS Connector Support Channel Access	129
On-Demand Upgrade of SaaS Connectors	130
Access Management	133

Manage Users & Groups	133
Access Scope	135
Add a Single Sign-on Provider	137
Add API Keys	167
Policies	169
Create a Policy	169
Add or Remove Legal Hold to Snapshot	171
Virtual Machines	173
VMware	173
Hyper-V	188
VMware Cloud (VMC) on AWS	196
Azure VMware Solution (AVS)	206
Physical Servers	217
Physical Server Requirements	217
Register Physical Server Sources	225
Protect Physical Servers	227
Recover Physical Servers	230
Upgrade Agents	232
NAS	234
Register Generic NAS Sources	234
Verify the Kerberos Configuration in Rigel	236
Configure and Register Isilon NAS	239
Configure and Register NetApp ONTAP	243
Protect NAS Sources	250
Recover NAS Data	257
Microsoft 365	261
Microsoft 365 Requirements	261
Microsoft 365 Backup Storage	280
Register Microsoft 365 Sources	298
Explore Microsoft 365 Sources	307
Exchange Online Mailboxes	309
OneDrive for Business	329
SharePoint Online	339
Microsoft Teams	349
Microsoft Groups	362
Microsoft Azure	369
Microsoft Azure Virtual Machines	369

Microsoft Azure SQL Database	391
Amazon Web Services	407
AWS Ports and Account Requirements and Considerations	407
Register Your AWS Account	422
Amazon EC2 Instances	426
Amazon RDS Instance	450
Amazon S3 Buckets	466
Databases	475
Microsoft SQL Server	475
Oracle Database	488
SAP HANA	500
Monitoring	518
Reports	518
Detect Ransomware Attacks	538
Alerts	539
Audit Logs	557
Subscription Status	564
Banner Messages	564
Sample Banner Messages	565
How-To Videos	567
Cohesity Support	568
Reach Cohesity Support	568
Support/Service Assistance	568
Cohesity Software Running on Partner Hardware	569

Cohesity Data Cloud

Cohesity Data Cloud is a unified cloud data management platform for securing, managing, and extracting value from your data, available as self-managed software and SaaS. The following are the key features of Cohesity Data Cloud:

- **Scale and simplicity**—Manage your entire data estate easily across data centers, edge sites, and public cloud environments.
- **Zero Trust Security**—Keep your data safe with in-flight and at-rest encryption, immutability, Write Once Read Many (WORM), Role-based Access Control (RBAC), and Multi-factor Authentication (MFA).
- **AI/ML Powered**—Streamline operations and defend against ransomware with Machine Learning (ML) and Artificial Intelligence (AI)-powered recommendations.
- **3rd Party Extensibility**—Connect Cohesity Data Cloud to your other IT investments to improve visibility and streamline operations.

Pillars

Cohesity Data Cloud includes five pillars. Each pillar encompasses a set of features and functionalities tailored to a specific aspect of data management. Each pillar contains one or more specialized apps. These apps are tailored to provide you with a focused and streamlined experience for achieving your goals within that particular area. Following are the five pillars:

- [Protection](#)
- [Security](#)
- [Mobility](#)
- [Access](#)
- [Insights](#)

Protection

The **Protection** pillar offers the most comprehensive backup and recovery solution to protect cloud-native, SaaS, and on-premises data at scale. You can simplify and accelerate the backup and recovery of enterprise workloads across on-premises and cloud with a secured unified platform for data resilience.



The **Protection** pillar includes the following apps:

- **DataProtect**—Offers a unified view and global management of all your Cohesity clusters, whether on-premises, in the cloud, or as Virtual Editions, regardless of the cluster size. You can easily connect your clusters to Helios and access them from anywhere using an internet connection and your Cohesity Support Portal credentials. It simplifies cluster management and enables efficient monitoring and control across your entire infrastructure.

Important: The previous **Cluster Manager** app has been integrated into the **Protection** pillar and it is now known as **DataProtect**.

- **DataProtect as a Service**—With Cohesity DataProtect delivered as a service, you can embrace a more predictable OpEx cost model, streamline backup operations across the hybrid cloud, and harness the power of your data for greater possibilities. By signing up, connecting, and initiating data protection, you can get started within minutes, ensuring your valuable data is safe and secure. Experience the convenience and efficiency of our seamless cloud-based solution for all your backup needs.

Important: The previous **DataProtect** app has been integrated into the **Protection** pillar and it is now known as **DataProtect as a Service**.

Security

The **Security** pillar empowers you to mitigate the risks posed by ransomware and other threats through an intelligent data security and management platform, purpose-built to safeguard your data and ensure its utmost security. You can boost cyber resiliency with ransomware recovery capabilities. These solutions help enterprises identify, protect, and recover data and processes from sophisticated cybersecurity threats.



The **Security** pillar includes the following apps:

- **Security Center**—Provides a comprehensive suite of security features, including DataHawk Threat Protection, Data Classification, Cyber Vaulting, and Platform Security, all conveniently accessible from a single unified platform.
- **FortKnox**—An award-winning cyber-vaulting solution that offers a SaaS-based data isolation and recovery platform that securely stores an immutable copy of data in a Cohesity-managed cloud vault.

Mobility

The automated disaster recovery solution in the **Mobility** pillar empowers you to achieve near-zero application downtime and zero data loss through unified backup and automated disaster recovery capabilities. You can eliminate secondary data centers and reduce the complexity of your on-premises operations.

SiteContinuity simplifies business continuity and disaster recovery with automated failover and failback orchestration for your mission-critical workloads.



Access

SmartFiles, the unified file and object services solution in the **Access** pillar, enables you to manage, secure, and do more with your data with software-defined file and object services for the hybrid cloud.

SmartFiles enables seamless data access for your users and applications with simultaneous multiprotocol support for NFS, SMB, and S3. You can also manage your data efficiently from a single console, giving you global control over data on-premises, at the edge, and in the cloud.



Insights

The **Insights** pillar empowers you to engage with and uncover valuable insights from your data:

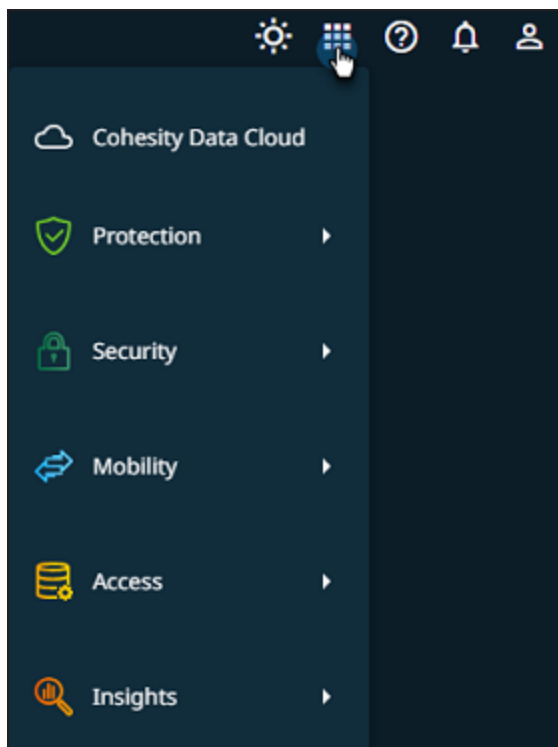


The **Insights** pillar includes the following apps:

- **Data Insights**—Harness the potential of your most important enterprise data and gain deep meaningful insights and learnings into your organization and data with Cohesity’s AI-powered conversational search solution. For more information, see [Cohesity Gaia](#).
- **Platform Insights**—Optimize, plan, and scale your Cohesity Data Cloud using AI-driven analytics. For more information, see [Platform Insights](#).

Switch Between Apps

You can use the app-selector menu to navigate between different apps:



Do one of the following:

- Click **Cohesity Data Cloud** to navigate to the Cohesity Data Cloud landing page. On this page, the easy navigation options allow you to explore the five pillars provided by Cohesity.
- Hover over **Protection** and select one of the following apps:
 - **DataProtect**—Offers you a unified view and global management of all your Cohesity clusters, whether on-premises, in the cloud, or as Virtual Editions, regardless of the cluster size.
 - **DataProtect as a Service**—Embrace a more predictable OpEx cost model,


streamline backup operations across the hybrid cloud, and harness the power of your data for greater possibilities.

- Hover over **Security** and select the following apps:
 - **Security Center**—Provides a comprehensive suite of security features, including DataHawk Threat Protection, Data Classification, Cyber Vaulting, and Platform Security, all conveniently accessible from a single unified platform.
 - **FortKnox**—A SaaS-based data isolation and recovery platform that securely stores an immutable copy of data in a Cohesity-managed cloud vault.
- Hover over **Mobility** and click **SiteContinuity**. The automated disaster recovery solution empowers you to achieve near-zero application downtime and zero data loss through unified backup and automated disaster recovery capabilities.
- Hover over **Access** and click **SmartFiles**. The unified file and object services solution enables you to manage, secure, and do more with your data with software-defined file and object services.
- Hover over **Insights** and select **Platform Insights**. Platform Insights offers a predictive and planning model that can make projections on cluster utilization and storage consumption and a set of 17 built-in reports.

Set Default Landing Page

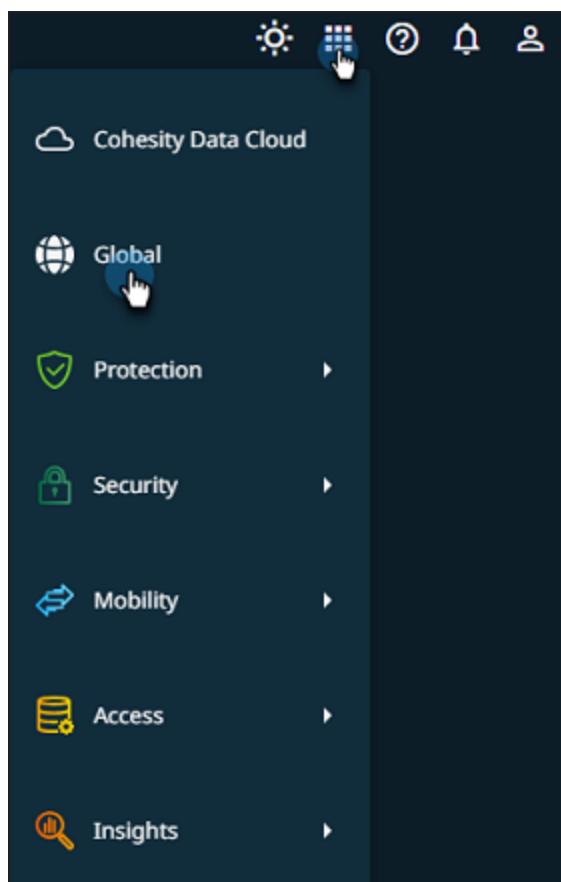
When you log in to Cohesity Data Cloud, all five pillars and apps are displayed by default. You have the ability to view all the pillars and apps, regardless of whether you have subscribed to them or not. If you have not subscribed to the app, an **Upgrade Now** option is displayed.

Click **Upgrade Now** to easily upgrade your subscription and gain access to the additional features and capabilities offered by the app:

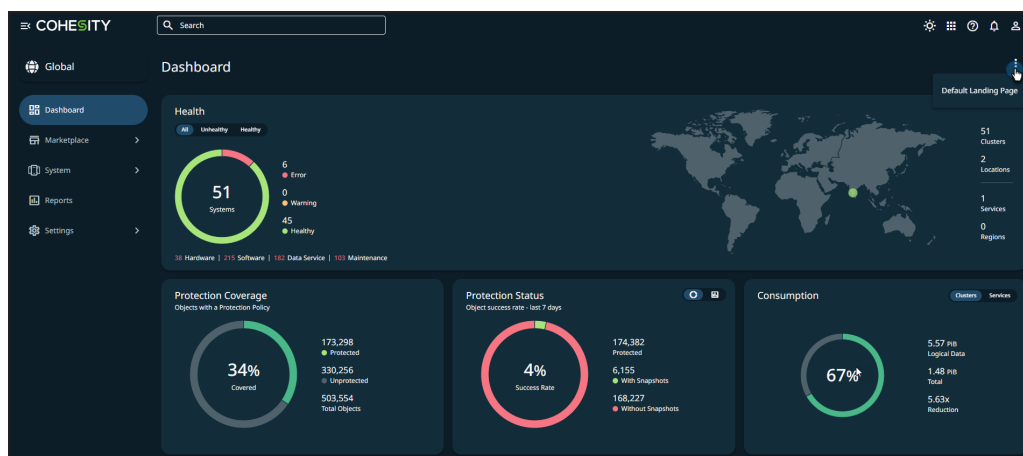


To set a specific page as the default landing page when accessing Cohesity Data Cloud, follow these steps:

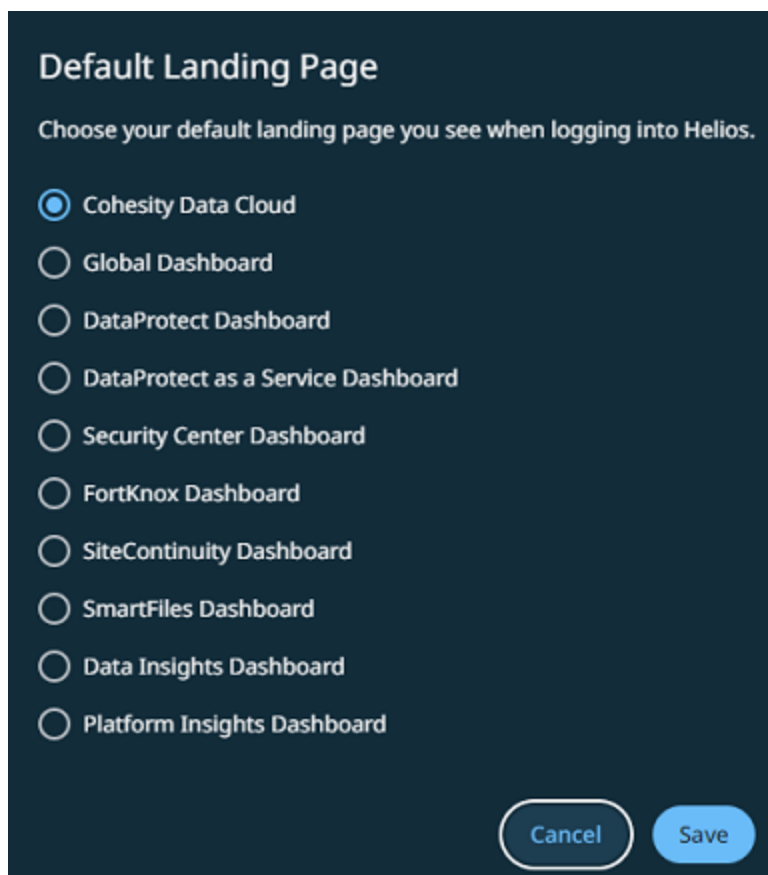
1. Log in to Cohesity Data Cloud.
2. Click any pillar and select an app.
For example, you can click the **Protection** pillar and select **DataProtect**.
3. Click the app-selector menu and select **Global**:



4. On the **Global > Dashboard** page, click the vertical ellipsis icon and click **Default Landing Page**:



5. On the **Default Landing Page** dialog, choose your default landing page and click **Save**:



Default Landing Page

Choose your default landing page you see when logging into Helios.

- ☒ Cohesity Data Cloud
- ☐ Global Dashboard
- ☐ DataProtect Dashboard
- ☐ DataProtect as a Service Dashboard
- ☐ Security Center Dashboard
- ☐ FortKnox Dashboard
- ☐ SiteContinuity Dashboard
- ☐ SmartFiles Dashboard
- ☐ Data Insights Dashboard
- ☐ Platform Insights Dashboard

[Cancel](#) [Save](#)

To view the changes to the default landing page in Cohesity Data Cloud, follow these steps:

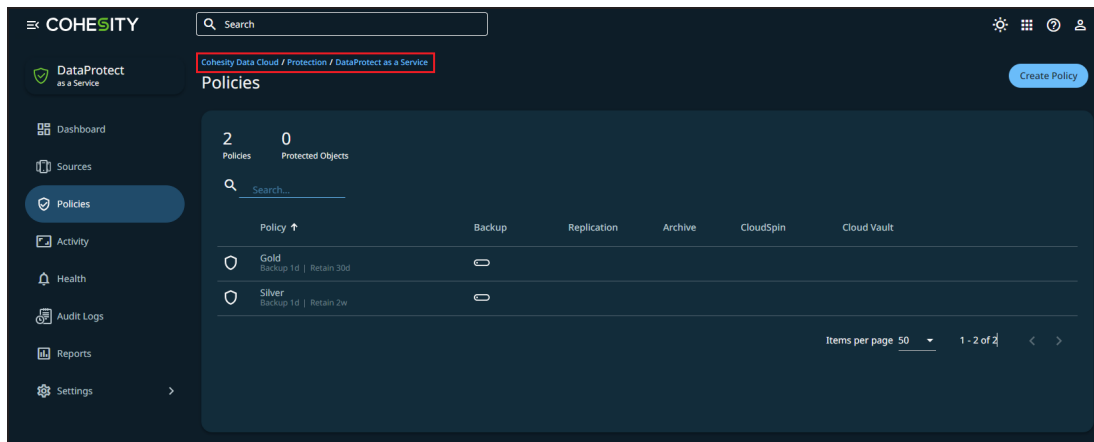
1. Log out of Cohesity Data Cloud.
2. After logging out, navigate back to the **Cohesity Data Cloud** login page.
3. Enter your credentials and log back in to Cohesity Data Cloud.

After logging back in, you can notice that the default landing page has been updated as per your preference.

Breadcrumbs

Cohesity Data Cloud introduces support for breadcrumbs, a user-friendly and efficient navigation aid. Breadcrumbs enable you to easily track your path and quickly navigate between pages within Cohesity Data Cloud. By understanding how to use breadcrumbs effectively, you can streamline your workflow and enhance your overall experience.

Breadcrumbs appear below the search bar and show the sequence of steps taken to arrive at the current location. Breadcrumbs consist of clickable links, allowing you to easily navigate back to previously visited pages:



Set User Preferences

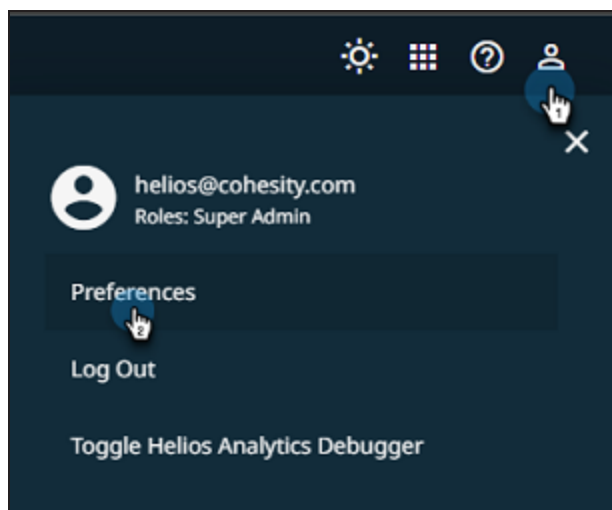
The **User Preferences** page in Cohesity Data Cloud allows you to customize various settings and options to tailor your experience according to your personal preferences. You can modify settings related to your account, user interface, and interactions with the Cohesity platform.

To set user preferences:

1. Log in to Cohesity Data Cloud.
2. Click any pillar and select an app.

For example, you can click the **Protection** pillar and select **DataProtect**.

3. Click the user icon in the upper-right corner and click **Preferences**:



The **User Preferences** dialog is displayed.

4. You can customize the following:

- **Language**—Select the language. Cohesity Data Cloud supports the following languages:
 - English
 - Japanese
- **Theme**—Select the theme. The theme you choose remains consistent across all Cohesity Data Cloud applications. Cohesity Data Cloud supports the following themes:
 - Dark
 - Light
- **Default Landing Page**—Select the default landing page that appears upon logging into Helios:
 - Cohesity Data Cloud
 - Global Dashboard
 - DataProtect Dashboard
 - DataProtect as a Service Dashboard
 - Security Center Dashboard
 - FortKnox Dashboard
 - SiteContinuity Dashboard
 - SmartFiles Dashboard
 - Data Insights Dashboard
 - Platform Insights Dashboard
- **Unsubscribed Services**—Opt to display or hide navigation items and content for any services that you have not subscribed to:
 - **Show**—Displays all the five pillars and all available services.
 - **Hide**—Displays only the pillar(s) and service(s) that you have subscribed to.
- **Byte Scaling**—Adjust the scale or size of data in terms of bytes. Cohesity Data Cloud offers the following byte scaling options:
 - Base 1024 (1 KiB = 1024 bytes)
 - Base 1000 (1 KB = 1000 bytes)
- **Time Format**—Select how time should be represented in Cohesity Data Cloud:
 - 12-hour clock
 - 24-hour clock

- Time Zone—Displays the time zone.
- **Persist Snack Bars**—Choose whether to keep messages, alerts, or notifications visible until you interact with them or let them disappear automatically.
 - **Persist Snackbar Messages**—Messages, alerts, or notifications stay visible until you acknowledge the message or dismiss it manually.
 - **Disappearing Snackbar Messages**—Messages, alerts, or notifications disappear automatically after a few seconds.

5. Click **Save**.

Global Dashboard

If you manage your Cohesity clusters through Cohesity DataProtect as a Service and if you have subscribed to any service, the **Global** dashboard provides a consolidated view of your cluster(s) and service(s).

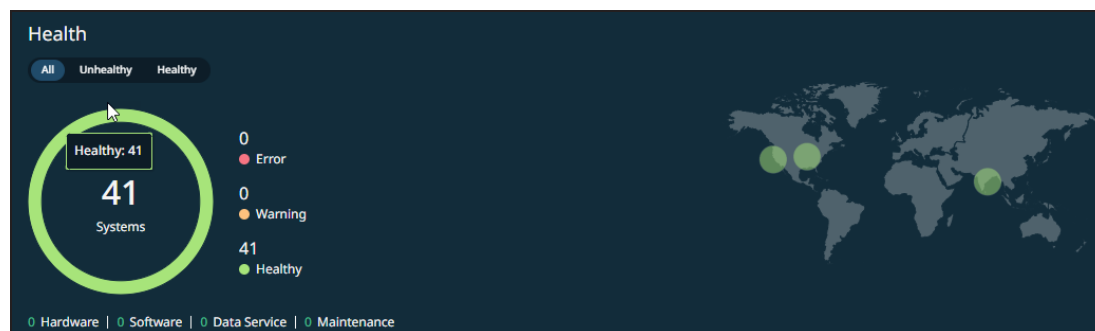
On the **Cohesity Data Cloud** landing page, click the **Cohesity Data Cloud** icon to navigate to the **Global** dashboard.

The **Global** dashboard provides a comprehensive overview of various aspects, including the health of managed clusters, protection status of objects, posture advisor score, discovered threats, and consumption metrics. The dashboard includes the following cards:

- [Health](#)
- [Protection Status](#)
- [Posture Advisor Score](#)
- [Threats Discovered](#)
- [Consumption](#)

Health

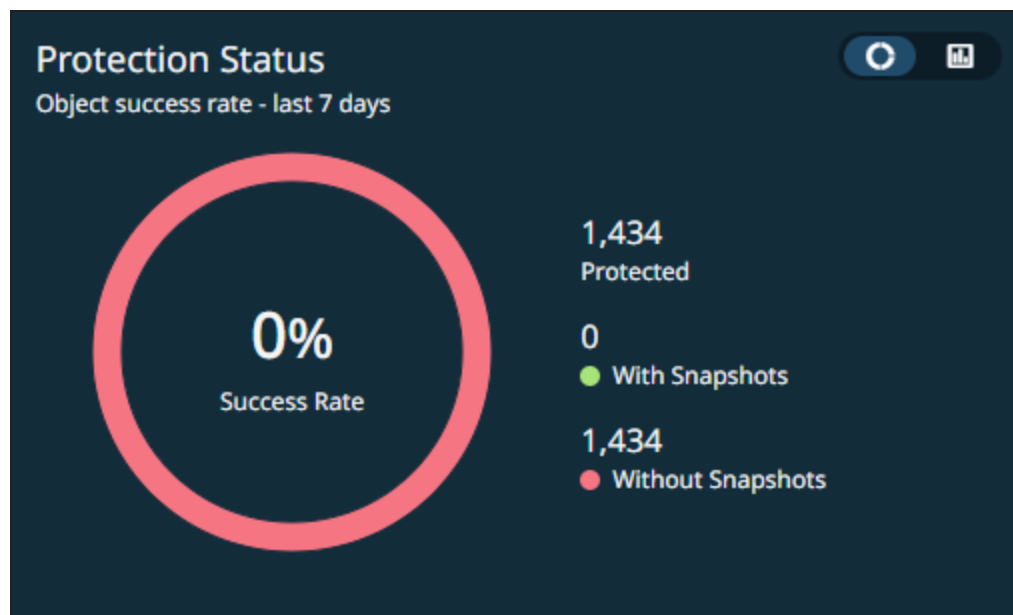
The **Health** card summarizes the health of clusters managed in Cohesity DataProtect as a Service. It displays the following details:



- The number of healthy and unhealthy clusters
- Summary of alerts generated by the Cohesity cluster(s)
- Geographical locations of the Cohesity cluster(s)

Protection Status

The **Protection Status** card provides a summary of all protected objects that had a backup run. You can view a summary of the following:




- Backup success rate
- Objects with and without snapshots
- Protected objects by type

Click on the card to navigate to the [Protected Objects](#) report page, where you can access detailed and granular information about the protected objects.

Posture Advisor Score

The **Posture Advisor Score** card allows you to get a global view of the security posture across all clusters managed in Cohesity DataProtect as a Service:



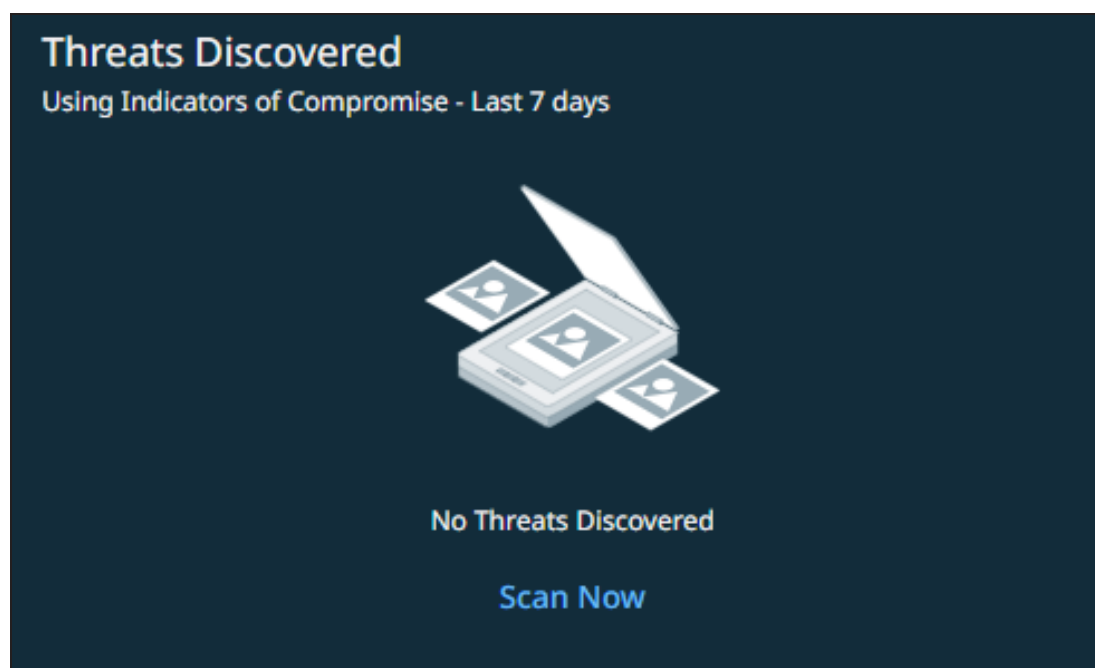
The card categorizes the score into the following categories:

- Less than 70—High risk
- 70 to 90—Medium risk
- Greater than 90—Low risk

For more information, see [Posture Advisor](#).

Threats Discovered

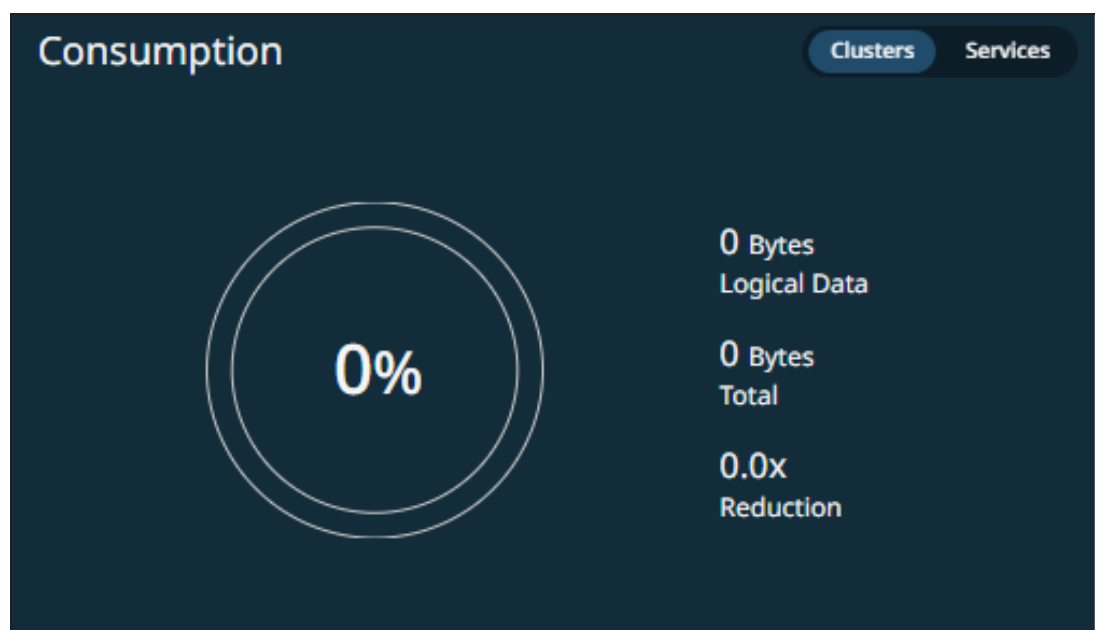
The **Threats Discovered** card summarizes the threats found during scans for malware and cyber threats using Indicators of Compromise (IOCs). You can click **Scan Now** and perform a threat scan:



For more information, see [Threat Detection](#).

Consumption


The **Consumption** card provides the storage statistics across all clusters managed in Cohesity DataProtect as a Service. You can view the following details related to clusters:



- Logical data consumed
- Total capacity available

- Storage reduction

On the **Consumption** card, click **Services** to view:



- Current plan—Free trial or a paid plan.
- Details about your subscription plan.
- Storage consumed by the protected objects in DataProtect as a Service. Click **Used** to access the Service Consumption report, which provides detailed consumption statistics.
- The amount of storage included in your subscription.
- The remaining duration of your subscription.

Cohesity DataProtect Delivered-as-a-Service

Today's companies and organizations are overwhelmed with the exponential growth in the amount of data they collect, manage, and store. You need to be able to focus on managing your data without worrying about additional hardware in your data center.

We designed Cohesity Data Cloud as a platform to host a series of Software-as-a-Service (SaaS) applications for data management. The first in the series is Cohesity DataProtect as a Service, Cohesity's SaaS offering that provides protection for your virtual and physical workloads, databases, and applications. You can sign up and start backing up your data today.

Log in to Cohesity DataProtect as a Service to protect data sources from your data center and SaaS applications in just a few steps:

1. Select a cloud region for your backups.
2. Register a source.
3. Select the objects on that source to protect.
4. Protect those objects.

Ready? > [Get started!](#)

What's New

Cohesity DataProtect as a Service keeps evolving. We're adding new features and supporting additional types of sources that you can protect in the service.

October 2024

Support for Bulk Agent Upgrade. You can now [upgrade multiple physical agents](#) at once, either immediately or at a scheduled time.

September 2024

- **Support for VMs on Azure VMware Solution (AVS).** You can now protect [VMs on Azure VMware Solution \(AVS\)](#).
- **Support for Setting Last Backup Time And Backup Status for VMware sources.** Cohesity now supports setting the [Set Last Backup detail](#) for a VM in vCenter.
- **Multi-Streamed SQL Server Log Backup.** The backup and recovery of SQL Server T-log with multi-backup files has been enhanced on DataProtect as a Service for better performance. The T-log backup settings now include options to specify the [WITH Clause and Number of Streams](#) for log backups.
- **Customize Dashboards.** Cohesity now provides [new widgets](#) to customize the DataProtect as a Service dashboards. You can add widgets that suit your needs to the dashboard.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- **Access Scope.** Cohesity DataProtect as a Service now allows users to control [access to specific sources](#). You can assign multiple access scopes to different users to manage their permissions.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- **Protection for vSphere DataSets files.** The vSphere DataSets feature offers a simple way to share small, rarely changing data between the vSphere management

layer and a guest operating system within a virtual machine that has VMware Tools installed. You can create multiple DataSets for a single VM, with each DataSet capable of storing numerous key-value entries. Multiple DataSets can be created for a single VM, and each DataSet can store multiple key-value entries. This functionality is available in vSphere 8.0 and requires the VM to run the virtual hardware version 20. Cohesity clusters now automatically back up the VM's DataSets file while [protecting a VM and also allow you to recover the DataSets file](#).

- **Auto-protect Azure VMs.** Cohesity now supports the [auto-protection](#) of Azure VMs at various levels:
 - Tenant level
 - Subscription level
 - Resource Group level

Additionally, you can also perform tag-based auto-protection of Azure VMs.

- **Tag-based exclusion of Azure VM.** You can now exclude Azure VMs from [Protection](#) based on tags.
- **Exclude disks of Azure VMs from Protection.** All the volumes of an Azure VM selected for protection are protected by default. Cohesity now supports you to exclude the disks of Azure VMs from [Protection](#).
- **File-level Recovery of Azure VM.** Cohesity now supports [file-level recovery](#) of Azure VMs.
- **Kerberos and IAM Authentication Support for RDS Ingest-Based Protection.** Cohesity DataProtect as a Service now supports Kerberos and IAM authentication methods for [protecting RDS](#) using the RDS Ingest Protection Type.
- **SAP HANA Data Protection.** Cohesity DataProtect as a Service now integrates with SAP HANA to provide [backup and recovery for SAP HANA deployments](#).
- **Microsoft 365 Backup Storage Support.** Cohesity DataProtect as a Service now supports [Microsoft 365 Backup Storage service](#) for Exchange Online, OneDrive for Business, and SharePoint Online. This feature provides an ultra-fast backup and recovery capability by creating backups within the M365 account using the Microsoft 365-protected services.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- **Backup and Recover Ownerless Teams and Groups.** Backup and recovery of Teams and Groups (Public and Private) with no owners is supported when at least one Exchange Online licensed member is available in the Teams/Groups.

If no owners/members are available in the Teams/Groups, you can contact your Cohesity account team to configure a service account with an Exchange Online license. This service account will be added as a member of the Teams/Groups before backup/recovery and removed once the process is complete.

- **Download Microsoft 365 Teams Content.** You can now download the Microsoft 365 Teams content as part of [Teams recovery](#).

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- **Nested Security Groups.** Security groups-based user import supports nested security groups for Microsoft 365 Mailboxes and OneDrives. For example, if a security group (AA) includes members (X and Y) and security groups (BB and CC), the import process includes the members (X and Y) and the members of security groups (BB and CC).

July 2024

- **Support for Customer-Managed Azure Keys for Data-at-Rest Encryption.** In Cohesity DataProtect as a Service, all data is encrypted in transit and at rest. You can now use [your own Azure Vault keys for encryption](#) when storing data in the Azure data plane(s).

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- **Backup Ownerless Teams and Groups.** Backup of Teams and Groups (Public and Private) with no owners is supported when at least one Exchange Online licensed member is available in the Teams/Groups.

If no owners or members are available in the Teams/Groups, you can contact your Cohesity account team to configure a service account (with an Exchange Online license). This service account will be added as a member of the Teams/Groups before backup or recovery and removed after the backup or recovery is completed respectively.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- **Protect SubstrateHolds Folder.** You can now protect the SubstrateHolds folder under Recoverable Items and Archive Recoverable Items root folders in the Microsoft 365 Mailboxes.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- **Protect vSphere DataSet Files.** The vSphere DataSets feature offers a simple way to share small, rarely changing data between the vSphere management layer and a guest operating system in a virtual machine with VMware tools installed. Multiple DataSets can be created for a single VM, and each DataSet can store multiple key-value entries. This functionality is available in vSphere 8.0 and requires the VM to run the virtual hardware version 20. Cohesity clusters now automatically back up the VM's DataSets file while protecting a VM and also allow you to recover the DataSets file. For more information, see [Protection for vSphere DataSets files](#).

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

June 2024

- **Legal Hold Support on Protection Runs.** Cohesity DataProtect as a Service now supports [enabling and releasing legal holds on protection runs](#). Users assigned the Data Security role can enable (and release) legal hold on existing snapshots to preserve them for legal purposes. You can also filter objects that have at least one protection run under legal hold.

May 2024

In this release, we've focused on enhancing the backend to improve scalability and system performance. Some of the changes lay the groundwork for future enhancements. While there are no new user-facing features, you can expect a smoother experience overall.

April 2024

- **On-Demand Upgrade of SaaS Connectors.** Cohesity periodically auto-upgrades SaaS connectors to ensure they run on the latest version. If the auto-upgrade fails, you can now perform an [on-demand upgrade](#) of the SaaS connectors.
- **VMware Recovery Enhancements.** The VMware recovery workflow has been enhanced to improve usability and avoid data-destructive operations.

March 2024

- **Leverage Amazon S3 Inventory Report to Protect Amazon S3 Buckets.** Cohesity now leverages the Amazon S3 inventory report to protect the Amazon S3 bucket. The inventory report contains the list of all the objects available on the Amazon S3 bucket you selected for protection. Cohesity uses this report to perform the first full backup and to periodically reconcile the list of objects. You can provide the details for creating an inventory report when [registering the AWS account](#). The subsequent incremental backups are performed by using the AWS EventBridge capability.
With this protection approach, Cohesity can back up multi-billion objects at a faster rate.
- **Enhanced X.509 Certificates Support for AWS Files and Folders Recovery Workflows.** File or folder-level recovery of AWS EC2 instances requires a Cohesity agent installed on the AWS EC2 target. For enhanced security, when installing the agent on the target EC2, Cohesity now automatically deploys an X.509 certificate.

February 2024

- **Recover Deleted Teams Private Channels.** You can now recover the deleted private channels to the original Microsoft 365 domain. If you are recovering to an alternate Microsoft 365 domain, you can create a new private channel and recover the data.
- **Enhanced Activity Page.** The Activity page in Cohesity DataProtect as a Service now provides consistent and timely information about the protected objects like protection status, timeline, and so on.
- **Deprecation of Basic Auth.** The Basic Auth authentication method is no longer supported for Microsoft 365 source registration.

January 2024

- **Subscription Banners:** Cohesity Helios now displays [banners on the UI](#), providing details on your Cohesity DataProtect delivered as a Service subscription status,

allowing you to take necessary actions.

- **Pillars:** Cohesity Data Cloud now includes five pillars. Each pillar encompasses a set of features and functionalities tailored to a specific aspect of data management. Each pillar contains one or more specialized apps. These apps are tailored to provide you with a focused and streamlined experience for achieving your goals within that particular area. Following are the five pillars:
 - [Protection](#)
 - [Security](#)
 - [Mobility](#)
 - [Access](#)
 - [Insights](#)

If you are an existing user, refer to the table below to identify the pillar to which the app belongs now and its updated name:

Note: The table excludes **Access** and **Insights** pillars since these pillars do not contain pre-existing apps.

Existing App Name	Pillar	Updated App Name	Description	Navigation
Cluster Manager	Protection	DataProtect	<p>The previous Cluster Manager app has been integrated into the Protection pillar and it is now known as DataProtect.</p> <p>DataProtect allows you to efficiently manage your Cohesity clusters.</p>	<ol style="list-style-type: none">1. Log in to Cohesity Data Cloud.2. Click the Protection pillar.3. Click DataProtect.

Existing App Name	Pillar	Updated App Name	Description	Navigation
DataProtect	Protection	DataProtect as a Service	<p>The previous DataProtect app has been integrated into the Protection pillar and it is now known as DataProtect as a Service.</p> <p>You can utilize DataProtect as a Service, an enterprise-grade Backup as a Service (BaaS) solution, to safeguard your critical SaaS, cloud-native, and on-premises data sources.</p>	<ol style="list-style-type: none"> 1. Log in to Cohesity Data Cloud. 2. Click the Protection pillar. 3. Click DataProtect as a Service.
DataHawk	Security	Security Center	<p>The previous DataHawk and Security Center apps have been unified and integrated into the Security pillar, now collectively known as Security Center.</p>	<ol style="list-style-type: none"> 1. Log in to Cohesity Data Cloud. 2. Click the Security pillar. 3. Click Security Center.
Security Center	Security	Security Center	<p>The app name remains unchanged.</p> <p>Security Center provides you with the capability to monitor the security posture of your Cohesity clusters, perform threat scans, and classify your critical data.</p>	

Existing App Name	Pillar	Updated App Name	Description	Navigation
FortKnox	Security	FortKnox The app name remains unchanged.	The FortKnox app has been integrated into the Security pillar. Enhance your cyber resiliency with FortKnox, a robust SaaS data isolation and recovery solution that ensures the safety of your data by maintaining an immutable copy in a Cohesity-managed cloud vault.	<ol style="list-style-type: none"> 1. Log in to Cohesity Data Cloud. 2. Click the Security pillar. 3. Click FortKnox.
SiteContinuity	Mobility	SiteContinuity The app name remains unchanged.	The SiteContinuity app has been integrated into the Mobility pillar. Simplify business continuity and disaster recovery with automated failover and failback orchestration for your mission-critical workloads.	<ol style="list-style-type: none"> 1. Log in to Cohesity Data Cloud. 2. Click the Mobility pillar. 3. Click SiteContinuity.

- **New Landing Page:** A newly introduced Cohesity Data Cloud landing page now presents a consolidated view of the five pillars. This user-friendly interface enables you to effortlessly navigate into the diverse pillars provided by Cohesity. For more information, see [Sign in to Cohesity DataProtect as a Service](#).
- **Application Switcher Changes:** The application switcher has undergone an update to align with the five pillars. Consequently, this modification has brought about changes to the existing navigation. For more information, see [Switch Between Apps](#).
- **Default Landing Page:** When you log in to Cohesity Data Cloud, all five pillars are displayed by default. However, you can set a specific page as the default landing page. For more information, see [Set Default Landing Page](#).

- **User Preferences:** Customize various settings and options to tailor your experience according to your personal preferences. You can modify settings related to your account, user interface, and interactions with the Cohesity platform. For more information, see [Set User Preferences](#).
- **Global Dashboard:** The Global dashboard has been revamped to provide a comprehensive overview of various aspects, including the health of managed clusters, protection status of objects, posture advisor score, discovered threats, and consumption metrics. For more information, see [Global Dashboard](#).
- **Breadcrumbs:** Cohesity Data Cloud introduces support for breadcrumbs, a user-friendly and efficient navigation aid. For more information, see [Breadcrumbs](#).

December 2023

Note: This is a phased rollout. These features may not be available instantly to all customers. For any queries, contact your Cohesity account team.

- **Protect Azure VM.** You can now protect Azure VMs in your Azure source. You can backup and recover Azure VMs using public or private endpoints. For more details, see [Microsoft Azure Virtual Machines](#).
- **Protect Azure SQL Server.** You can now protect your Azure SQL Server databases using Cohesity DataProtect as a Service. This feature can be used for long-term archival, immutable data protection, regulatory offsite requirements, and more. For more details, see [Microsoft Azure SQL Database](#).
- **Protect Microsoft SQL Server AG.** You can now protect Microsoft SQL Server Always On Availability Groups (AGs), which is a High Availability (HA) and Disaster Recovery (DR) solution. For more details, see [Microsoft SQL Server](#).
- **Protect Lists in SharePoint Online.** You can now protect Lists on Microsoft 365 SharePoint Online. Lists contain various types of data such as links, announcements, contacts, issue trackers, surveys, and more. For more details, see [SharePoint Online](#).

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- **Backup and Download Teams Posts and Private Chats.** You can now backup and download the following in Microsoft 365 Teams and Mailbox:
 - Posts from all channels in Teams
 - Posts from a single channel in Teams

- Chats from a specific user in the Mailbox

For more details, see [Exchange Online Mailboxes](#) and [Microsoft Teams](#).

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- **Microsoft 365 Recovery Self-Service.** Cohesity provides a self-service workflow to help the end users recover the Microsoft 365 Mailbox and OneDrive items.

Users can recover the following using the Microsoft 365 Self-Service Portal:

- **Microsoft 365 Mailbox** - Emails, Folders, Calendars, Contacts, Tasks, and Notes.
- **Microsoft 365 OneDrive** - Files and Folders.

For more details, see [Mailbox Items Recovery Self-Service](#) and [OneDrive Content Recovery Self-Service](#).

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- **Granular Recovery of AWS S3 Objects.** You can now perform granular recovery of AWS S3 objects by providing the prefix. For more details, see [Recover Your Amazon S3 Buckets](#).
- **Add Tags During Recovery of EC2 Instances.** You can now add custom tags to the EC2 instances during the recovery task. The recovered EC2 instances are updated with the new and existing tags. For more details, see [Recover Your Amazon EC2 Instances](#).
- **Protect PostgreSQL and Aurora (PostgreSQL Compatible).** In addition to the snapshot management support for all Amazon RDS Database engine types, you can now leverage ingest-based backup for RDS PostgreSQL and Aurora (PostgreSQL compatible). For more details, see [Protect Your Amazon RDS Databases](#).
- **Protect Mailbox Recoverable Items.** You can now protect the Recoverable Items folder in the Microsoft 365 Mailboxes. These folders preserve the items that are soft deleted or deleted from the Deleted Items folder. For more information, see [Protect Microsoft 365 Mailboxes](#) and [Recover User Mailboxes](#).

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- **Protect PHL in SharePoint Online.** Cohesity supports the protection of Preservation Hold Library (PHL) for SharePoint Online which is used to store the files needed for compliance reasons. For more information, see [Protect Microsoft 365 SharePoint Online Sites](#) and [Recover SharePoint Sites](#).

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- **Protect PHL in OneDrive.** Cohesity supports the protection of Preservation Hold Library (PHL) for OneDrive which is used to store the files needed for compliance reasons. For more information, see [Protect Microsoft 365 OneDrives](#) and [Recover User OneDrives](#).

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

October 2023

Support for VMware Cloud (VMC) on AWS

You can now protect VMware Cloud on AWS vCenter. Review the [requirements](#), [register VMC on AWS as a source](#), and [start protecting your VMC on AWS](#)!

Enhancements to vCenter and ESXi Host Configuration

When registering a vCenter or Standalone ESXi host, you can now:

- Limit the number of concurrent streams per data store
- Limit the number of concurrent backups per vCenter (not applicable for ESXi host)
- Set the minimum free space that must always be available in the datastore

By configuring the above settings that directly impact how protection runs perform, you can optimize the backup process and achieve better performance. For more information, see [Register VMware Sources](#).

SharePoint Online Sites Protection in Multi-Geo Locations

You can now discover and protect SharePoint Online sites in the satellite storage locations of the Microsoft 365 tenant along with the Central storage locations. For more information, see [SharePoint Online](#).

Note: This is an Early Access feature. Contact your Cohesity account team to enable the feature.

PST Download Support for Mailbox Items

You can now export and download the Exchange Online Mailbox items like emails, folders, calendar invites, contacts, notes, and tasks in the Portable Storage Table (PST) format. For more information, see [Recover Mailbox Items](#).

Note: This is an Early Access feature. Contact your Cohesity account team to enable the feature.

Physical Server Host Configuration Enhancements

The wizard to register a Physical Server is simplified and now allows you to specify a nickname for your source. For more information, see [Register Physical Server Sources](#).

AWS SaaS Connector Enhancements

When creating AWS SaaS Connections, you can now:

- Select multiple network security groups for associating with SaaS Connectors.
- Specify NTP servers to synchronize the time on the Cohesity DataProtect as a Service.
- Specify the IP addresses of the Domain Name System (DNS) servers that Cohesity DataProtect as a Service must use.

For more information, see [Create AWS SaaS Connection](#).

September 2023

Protection for VMware Cloud (VMC) on AWS

You can now protect VMware Cloud on AWS vCenter. Review the [requirements](#), [register VMC on AWS as a source](#), and [start protecting your VMC on AWS](#)!

Note: This is an Early Access feature. Contact your Cohesity account team to enable the feature.

August 2023

AWS SaaS Connection Enhancements.

You can now:

- Identify the reason for the failure of the AWS SaaS Connection deployment by reviewing the error messages.
- Monitor the progress of AWS SaaS Connection deployment.
- Retry the AWS SaaS Connection if the connection fails.
- Retry the AWS SaaS Connection Deletion.
- Forcefully delete the AWS SaaS Connection.
- View the number of AWS SaaS Connectors deployed.

For more information, see [Manage User-Deployed SaaS Connections](#).

June 2023

New Onboarding Wizard. Cohesity DataProtect as a Service has a new wizard to simplify the onboarding process and assist you with registering your data sources. For more information, see [Get Started](#).

April 2023

- **Alert notification by email.** You can now create [alert notification rules](#) in Cohesity DataProtect as a Service that send emails based on the alert categories, severities, and names.
- **Support for file and folder inclusion rules.** Cohesity DataProtect as a Service now supports the addition of file and folder inclusion rules while protecting physical servers.

March 2023

- **Support for Microsoft Azure Cloud Regions.** Cohesity DataProtect as a Service is now available in Microsoft Azure Cloud. You can now choose Microsoft Azure cloud to back up the [supported workloads](#).

Note: The cloud regions visible in your subscription are based on your purchased Cohesity DataProtect as a Service entitlements.

- **Recover Deleted Teams Private Channels.** You can now recover the [deleted private channels](#) to the original Microsoft 365 domain. If you are recovering to an alternate Microsoft 365 domain, you can create a new private channel and recover the data.

Note: This is an Early Access feature. Contact your Cohesity account team to enable the feature for your tenant.

- **Support Microsoft 365 GCC Government Editions.** Cohesity DataProtect as a Service now supports [Microsoft 365 Government Community Cloud \(GCC\) editions](#).

Note: Cohesity DataProtect as a Service does not support GCC High.

- **KMS Encryption Key Settings for AWS EC2 Instances Recovery.** When you recover AWS EC2 instances to a new AWS account, you can now change the KMS encryption key settings of EC2 instances. For more information, see [Recover Your Amazon EC2 Instances](#).
- **Amazon S3 Protection.** You can now protect Amazon S3 buckets in your AWS account, in addition to protecting your Amazon EC2 instances and Amazon RDS databases. For more information, see [Amazon S3 Buckets](#).

Note: This is an Early Access feature. Contact your Cohesity account team to enable the feature for your tenant.

January 2023

Dual Network Support for VMware SaaS Connector. You can now [deploy the VMware SaaS connector](#) with dual network configuration in deployments where the data sources are in a private non-routable VLAN. For example, you can configure a data (primary) network for communication with Cohesity Data Cloud (SaaS) and a different secondary network for communication with your data sources.

December 2022

Microsoft 365 Teams Alternate Recovery. You can now recover the [whole Teams](#) instance or [specific Teams content items](#) to an alternate Microsoft 365 domain.

Security Groups for Mailbox and OneDrive Protection. Cohesity now supports protecting Microsoft 365 Mailbox and OneDrive data of Security Groups. Security Groups are Microsoft native groups that are used for granting access to Microsoft 365 resources which can contain users or devices.

For more information, see [Protect Microsoft 365 Mailboxes](#) and [Protect Microsoft 365 OneDrives](#).

August 2022

Audit Logs. DataProtect now provides audit information for the events generated on the registered regions through DataProtect. For more information, see [Audit Logs](#).

July 2022

- **SaaS Connector Groups.** If you use one vCenter to manage multiple ESXi clusters in different geographic locations, you can group the local SaaS connectors at each location into [SaaS Connector Groups](#), followed by associating these Connector Groups to vCenter resources in that location. It helps you ensure efficient routing of your backup and recovery data traffic through SaaS Connectors that operate in that same location.
- **SaaS Connector Alert.** A **Critical** alert, **SaaSConnectorStatusAlert** is triggered when the SaaS connector is not reachable due to a network connection issue or is down. You can configure [alert email notifications in DataProtect](#) to receive this alert and take appropriate action.
- **Granular Recovery for Amazon EC2.** You can now perform a [granular file and folder recovery](#) for Amazon EC2. This feature is available for Cohesity snapshots and not AWS snapshots.
- **Recover Mailbox items.** In addition to recovering [individual emails](#) and [folders](#), you can now recover [calendar invites](#), [contacts](#), [notes](#), or [tasks](#).
- **Add Multiple Microsoft 365 Service Accounts.** To manage Exchange Online throttling mailbox protection on tenants where OAuth is not enabled, you can add multiple Microsoft 365 service accounts during the [source registration](#) or edit the source configuration and add multiple Microsoft 365 Service User Accounts.
- **Download multiple OneDrive or SharePoint Site files and folders.** As part of the recovery workflow, you can now download [multiple files and folders from a user's OneDrive backup](#) or [document libraries and files from the SharePoint site backup](#).
- **Microsoft 365 Protection for Groups.** In addition to protecting Microsoft 365 user Mailboxes, OneDrives, SharePoint Online Sites, and Teams, you can now protect your [Groups](#) data as well.

Note: This is an Early Access feature. Contact your Cohesity account team to enable the feature for your tenant.

March 2022

- **Microsoft 365 Mailboxes and OneDrives.** UI enhancements:
 - Global search for a Microsoft 365 User now correctly displays the matching Mailbox and OneDrive objects separately using the correct Icons. (Earlier, Microsoft 365 User icon was incorrectly displayed for both Mailbox and OneDrive objects, making it difficult to interpret search results.)
 - Protection Coverage, Status, and Last Backup widgets on the Dashboard now treat Microsoft 365 Mailboxes and OneDrives as separate objects.
 - Sources page now treats Microsoft 365 Mailboxes and OneDrives as separate objects when displaying Protected and Unprotect Object counts.
 - Protected Objects Report now lists Mailbox and OneDrive objects separately.
- **Video Examples.** We've recorded several [how-to videos](#) to help you learn some of the key DataProtect tasks in step-by-step examples.

February 2022

- **Granular Microsoft 365 Teams Recovery.** You can now recover [specific Teams content items](#), in addition to [whole Teams](#).
- **Microsoft 365 Express Registration.** You can now let [Cohesity create the Azure application](#) you need to register your Microsoft 365 sources. And if your business requires it, [you can still enter your specific Azure application details manually](#) as well.
- **Ransomware Detection for More Workloads.** During protection runs, the Cohesity DataProtect as a Service [detects anomalies](#) in your data and triggers a specific critical alert, DataIngestAnomalyAlert. You can now check for these anomalies, inspect any that occur, and when necessary, recover the object from the latest clean backup.

Note:

This feature is now available for the following data sources:

- VMware VMs
- Hyper-V VMs
- Generic NAS
- NetApp
- Physical (file-based)
- Amazon EC2 (Cohesity snapshots)
- SQL VDI

- **Streamlined SaaS Connector Firewall Port Requirements.** You no longer need to open outgoing firewall ports 11117 and 29991 for your [SaaS Connectors](#).

October 2021

- **Granular Microsoft 365 OneDrive Recovery.** You can now recover [specific contents from a user's OneDrive](#), in addition to [whole OneDrives](#).

September 2021

- **Granular Microsoft 365 SharePoint Sites Recovery.** When recovering Microsoft 365 SharePoint Online sites, you can now recover specific [document library items](#) as well as [whole sites](#).

August 2021

- **Amazon RDS Protection.** You can now [protect the Amazon RDS databases](#) in your AWS account, in addition to [protecting your Amazon EC2 instances](#).

July 2021

- Cohesity DataProtect as a Service now supports this additional [cloud region](#) to store your data:
 - **Europe (London)**
- Cohesity DataProtect delivered as a Service is now SOC 2 Type II certified.

June 2021

- **Granular Microsoft 365 Mailbox Recovery.** We've added indexing to [Microsoft 365 Mailbox protection](#), so that you can recover [individual emails](#) and [folders](#), in addition to [whole Mailboxes](#).
- **Hyper-V Protection.** You can now register your SCVMM server and Standalone Hyper-V hosts to [protect your Hyper-V VMs](#).

Note: The Hyper-V recovery workflow currently only supports granular (file- & folder-level) recovery. VM-level recovery is coming soon.

- **AWS EC2 Data Ingest.** Now you have [two options for protecting your AWS EC2 instances](#): **AWS snapshots** are saved to the same account and region as your EC2 instances, while **Cohesity snapshots** are saved to your Cohesity DataProtect as a Service [cloud region](#).
- **Oracle Database Protection.** Register your Oracle servers and hosts to [protect your Oracle Databases](#).
- **Source-Specific SaaS Connectors.** We've updated our [SaaS Connectors](#) with specific choices for your data source types: VMware, AWS, and Hyper-V.
- **Cohesity DataProtect as a Service** now supports this additional cloud region to store your data:
 - **Europe (Frankfurt)**

May 2021

- **Microsoft 365 Protection for OneDrives, SharePoint Online Sites, and Teams.** Now, in addition to protecting Microsoft 365 user Mailboxes, you can protect your [Microsoft 365 OneDrives](#), [SharePoint Online Sites](#), and [Teams](#).
- **Physical Server Protection.** Use the Cohesity DataProtect as a Service to [protect your Linux and Windows servers](#).
- **AWS EC2 Instances.** You can now [protect the EC2 instances](#) in your AWS account.

April 2021

- **Bandwidth Throttling.** If you need to manage the network bandwidth consumption of your backup and recovery tasks, you can now [schedule bandwidth usage limits](#) in your SaaS Connections.
- **Differential Restores for VMware VMs.** When time is of the essence when you're recovering VMs, you can now take advantage of [VMware differential restores](#) when

recovering VMs to their original locations.

- **Ransomware Detection.** During protection runs, the Cohesity DataProtect as a Service [detects anomalies](#) in your data and triggers a specific critical alert, **DataIngestAnomalyAlert**. You can now check for these anomalies, inspect any that occur, and when necessary, recover the object from the latest clean backup.
- Cohesity DataProtect as a Service now supports these additional [cloud regions](#) to store your data:
 - **US East (N. Virginia)**
 - **US West (N. California)**
 - **Asia Pacific (Sydney)**

March 2021

- **Microsoft 365 Mailbox Protection.** You can now protect your Microsoft 365 user Mailboxes. Check the [requirements](#), [register your Microsoft 365 sources](#), and start protecting user [Mailboxes](#)!
- **SQL Server Protection.** Now you can protect your SQL databases with Cohesity DataProtect as a Service. Make sure you meet the [SQL requirements](#), [register your SQL Server sources](#), and get started [protecting your SQL databases](#)!
- **Reporting.** Cohesity DataProtect as a Service reports give you useful insights on your data protection trends. [Inspect and share](#) your data protection and recovery results.

Supported Software for Cohesity DataProtect as a Service

VMware

vCenter, vSphere, ESXi versions	Virtual Machine Hardware	Guest OS
8.0 U3, 8.0 U2, 8.0 U1, 8.0	9, 10, 11, 13, 14, 15, 17, 18, 19, 20, 21	CentOS 5.10, 6.6+, 7.0 - 7.9, 8.0, 8.3 Debian 9.6, 10, 11.x openSUSE 15.1 Oracle Linux (OEL) 5.8 - 5.11, 6.x, 7.0 - 7.9, 8.0, 8.1, 8.2 - 8.4, 8.5 Red Hat Enterprise Linux (RHEL) 6.6+, 7.0 - 7.9, 8.0 - 8.2 Solaris 10, 11 (x86) SUSE Linux Enterprise Server 11 SP4, 12 SP4, 12.3, 15.0, 15.3 Ubuntu 14.x, 16.x, 18, 19.x, 20.x, 22.04 Windows 7, 8, 10
7.0 U3	9, 10, 11, 13, 14, 15, 17, 18, 19	Windows 2008 R2 Windows 2012, 2012 R2 Windows 2016 Windows Server 2016 Core Windows 2019 Windows Server 2019 Core Windows 2022 Windows Server 2022 Core
7.0 U2	9, 10, 11, 13, 14, 15, 17, 18, 19	
7.0 U1	9, 10, 11, 13, 14, 15, 17, 18	
7.0	9, 10, 11, 13, 14, 15, 17	
6.7 U3	9, 10, 11, 13, 14, 15	
6.7 U2	9, 10, 11, 13, 14, 15	
6.7	9, 10, 11, 13, 14	
6.5	9, 10, 11, 13	

VMware Cloud Infrastructure

VMware Cloud Infrastructure
VMware Cloud (VMC) on AWS
Azure VMware Solution (AVS)

Microsoft Hyper-V

Hyper-V	Versions	Guest OS	Notes
Hyper-V Standalone, SCVMM , Cluster	2016, 2019, 2022	<p>All the operating systems supported by Hyper-V that are compatible with Cohesity's Physical Servers category. The following guest OS are also supported:</p> <p>FreeBSD 8.4,9.1-9.3, 10.0-10.3, 11.0, 11.1-11.3, 12-12.1</p>	<ul style="list-style-type: none"> • Cohesity supports Hyper-V in the Desktop and Server Core installation modes. • Supported VM configuration versions are 8.0 or later. • The Cohesity does not support backing up of Hyper-V VMs with shared disks. • For agent-based file level restores only 64-bit guest operating systems are supported. • Cohesity supports Hyper-V in the Nano Server installation mode with Windows Server 2016. Nano Server 2019 is used as a container and you cannot install Microsoft Hyper-V inside the container. It is primarily used for non-containerized installations.

Physical Servers

Operating System	CPU Architecture
Windows Server 2016 Core, 2019 Core, 2022 Core	64-bit
Windows 10 Desktop Edition, Windows 11	64-bit
Windows 2008 R2, 2012, 2012 R2, 2016, 2019, 2022	64-bit
CentOS 6.0+, 7.0 - 7.9, 8.0, 8.3, 9.0	64-bit
Oracle Linux (OEL) 6.x, 7.0 - 7.7, 7.9, 8.0 - 8.10, 9.0 - 9.4	64-bit
Red Hat Enterprise Linux (RHEL) 6.7+, 7.0 - 7.9, 8.0 - 8.10, 9.0 - 9.2	64-bit
Rocky Linux 8.x - 8.10, 9.x - 9.4	64-bit
SUSE Linux Enterprise Server (SLES) 11 SP4, 12 SP3 - 12 SP5, 15 SP3, 15 SP4, 15 SP5, 15 SP6	64-bit
openSUSE 15.1, 15.3, 15.4	64-bit
Ubuntu 14.04, 16.04, 18.04, 19.04, 20.04, 22.04, 23.04	64-bit
Debian 9.6, 10. 11.2	64-bit
Amazon Linux 2	64-bit

Microsoft 365 Editions

Cohesity DataProtect as a Service supports the following Microsoft 365 editions:

Note: Cohesity DataProtect as a Service does not support GCC High.

Microsoft Plans	Editions
Microsoft 365 For Business	<ul style="list-style-type: none"> • Microsoft 365 Business Basic • Microsoft 365 Business Standard • Microsoft 365 Business Premium
Microsoft 365 For Enterprise	<ul style="list-style-type: none"> • Microsoft 365 E3 • Microsoft 365 E5 • Microsoft 365 F3
Office 365 For Enterprise	<ul style="list-style-type: none"> • Office 365 E1 • Office 365 E3 • Office 365 E5
Microsoft 365 GCC (Government Community Cloud)	<ul style="list-style-type: none"> • Office 365 Government G1 • Office 365 Government G3 • Office 365 Government G5 • Microsoft 365 Government G3 • Microsoft 365 Government G5 • Office 365 F3

Microsoft SQL Server

Database	Notes
Microsoft SQL Server 2022	<ul style="list-style-type: none"> MS SQL Server 32-bit deployment is not supported with the agent-based approach. It is recommended to use the Native SQL dump to Cohesity SMB for backups.
Microsoft SQL Server 2019	<ul style="list-style-type: none"> MS SQL Server 64-bit running on Windows 64-bit is supported.
Microsoft SQL Server 2017	<ul style="list-style-type: none"> SQL databases and logs residing on Clustered Shared Volumes (CSV) can be protected using Cohesity's VDI based backup method.
Microsoft SQL Server 2016	<ul style="list-style-type: none"> FILESTREAM databases can be protected using Cohesity's VDI.
Microsoft SQL Server Express 2019	<ul style="list-style-type: none"> Memory-optimized databases can be protected using Cohesity's VDI based backup method. <p>Supported on the following Operating Systems:</p> <ul style="list-style-type: none"> Windows 2022 64-bit
Microsoft SQL Server Express 2017	<ul style="list-style-type: none"> Windows 2019 64-bit
Microsoft SQL Server Express 2016	<ul style="list-style-type: none"> Windows 2016 64-bit

Oracle

Database	Oracle Release	Supported OS Versions
Oracle Database (Standalone)	11gR2	Oracle Enterprise Linux (OEL) 6.x, 7.x, 8.x
	12cR1	Red Hat Enterprise Linux (RHEL) 6.x, 7.x, 8.x
	12cR2	
	18c	
	19c	

SAP HANA

Database	Supported OS Versions
SAP HANA 2.0 SPS07 or earlier Note: <ul style="list-style-type: none"> Only SAP HANA standalone deployments are supported. SAP HANA 1.0 is not supported. 	Supported on x86-64-based systems running on RHEL and SLES.

NAS

Vendor	Version	Notes
Generic NAS	All vendors supporting NFS (v3) or SMB (v2, v3+)	SMB v1 is not supported in DataProtect delivered as a Service.
Dell EMC Isilon	8.0.x, 8.1, or 8.2.x, 9.0.x - 9.4.x, 9.5.x, 9.7.x, 9.8.x	
NetApp ONTAP Cluster-Mode	8.2, 8.3, 9.3, 9.5, 9.6, 9.7, 9.8, 9.9.1, 9.10, 9.11.x, 9.12.x, 9.13, 9.13x, 9.14x	

Cloud Services

Cloud Service Provider	Supported Services
AWS	EC2
	<p>RDS</p> <p>Supported RDS Databases Using Ingest-Based Protection :</p> <ul style="list-style-type: none">• PostgreSQL: 11.22-R2- 15.6.• Aurora (PostgreSQL Compatible): 11.9 - 15.6
	S3
Azure	Virtual Machine
	Azure SQL Database and Azure SQL Managed Instance

Supported Workloads and Cloud Regions

You can use Cohesity DataProtect as a Service to store your backups on the Cohesity-managed SaaS platform in Amazon Web Services (AWS) or Microsoft Azure.

Supported Cloud Regions

The following is the list of supported regions and countries to store your backups on the Cohesity-managed SaaS platform:

Cloud Provider	Cloud Regions
AWS	Americas: <ul style="list-style-type: none">• US-East-1 (N. Virginia)• US-East-2 (Ohio)• US-West-1 (N. California)• US-West-2 (Oregon)• Canada-Central (Quebec)
	Asia Pacific: <ul style="list-style-type: none">• Asia-Pacific (Singapore)• Asia-Pacific (Sydney)• Asia-Pacific (Mumbai)• Asia-Pacific (Tokyo)
	Europe, Middle East, Africa: <ul style="list-style-type: none">• Europe (Frankfurt)• Europe (London)• Europe (Paris)• Middle East (Bahrain)

Cloud Provider	Cloud Regions
Microsoft Azure	Americas: <ul style="list-style-type: none">• Central US (Iowa)• West US 2• East US• Canada (Toronto) Europe: <ul style="list-style-type: none">• France Central (Paris)• UK South (London) Australia: <ul style="list-style-type: none">• Australia East (New South Wales)

Supported Workloads and Cloud Providers

The following table lists the supported workloads on AWS and Azure cloud regions:

Cloud Provider	Supported Workloads
AWS	VMware Virtual Machines
	VMware Cloud (VMC)
	Hyper-V Virtual Machines
	Physical Servers (Windows and Linux)
	NAS (NetApp, Isilon, and generic files)
	<div> Note: You can protect AWS solutions such as AWS FSx by registering them as Generic NAS, provided the AWS solutions use NFS (v3) or SMB (v2, v3)+ protocol. </div>
	Microsoft 365 (Exchange, OneDrive, SharePoint, Groups, and Teams)
	Amazon EC2
	Amazon RDS
	Amazon S3
	Microsoft SQL Server
	Oracle
	SAP HANA
Microsoft Azure	Microsoft 365 (Exchange, OneDrive, SharePoint, Groups, and Teams)
	Azure Virtual Machines
	Azure SQL Database and Azure SQL Managed Instance
	NAS - Generic files <div> Note: You can protect Azure Files and Azure NetApp Files by registering them as Generic NAS, provided the files use NFS (v3) or SMB (v2, v3)+ protocol. </div>

Supported Regions for AWS SaaS Connectors

Cohesity supports the backup of AWS services in the regions where you can deploy AWS SaaS Connectors. You can deploy AWS SaaS connectors in all the regions AWS supports.

Supported Regions for Azure SaaS Connectors

Cohesity supports the backup of Azure VMs or SQL in the regions where you can deploy Azure SaaS Connectors. You can deploy Azure SaaS connectors in the following Azure regions.

Note: Contact your Cohesity Account Team to deploy SaaS Connectors in a region that is not listed below.

Americas:

- Central US (Iowa)
- South Central US (Texas)
- North Central US (Illinois)
- Canada Central (Toronto)
- East US (Virginia)
- West US (West US)
- East US 2 (Virginia)
- West US 2 (Washington)
- West US 3 (Phoenix)

Europe

- France Central (Paris)
- North Europe (Ireland)
- Germany West Central (Frankfurt)
- UK South (London)
- West Europe (Netherlands)

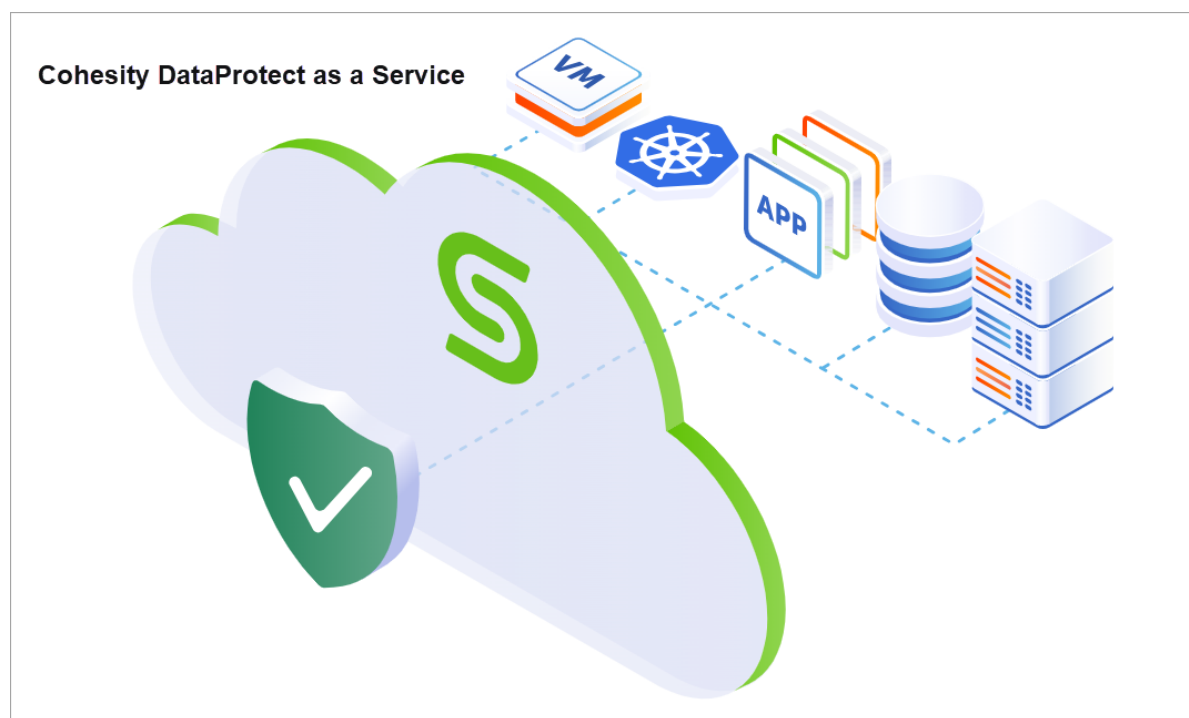
Australia

- Australia East (New South Wales)

Get Started

To get started:

1. [Sign in to the Helios account](#) that has Cohesity DataProtect as a Service enabled.
2. [Select a cloud region](#) for your backups and choose a [Key Management System](#) for your data encryption.
3. [Add users](#) to access Cohesity DataProtect as a Service.
4. [Register](#) your source.
5. Select the objects on that source to protect.
6. [Protect](#) those objects.



Sign in to Cohesity DataProtect as a Service

To access the Cohesity DataProtect as a Service, you'll need the Helios username from the welcome email and the password you set when you activated your Helios account.

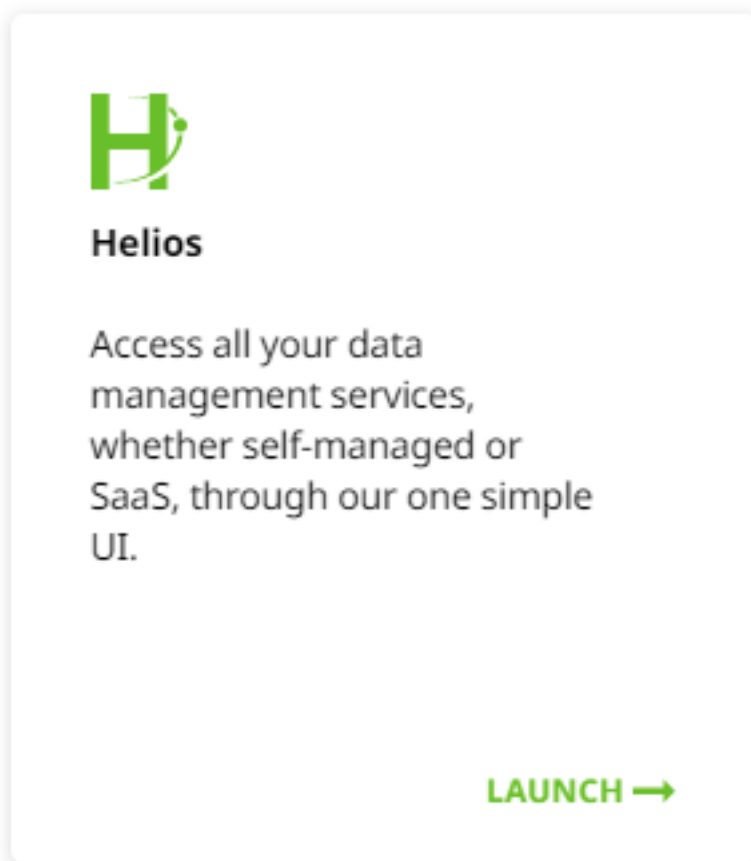
You must sign in to Helios through [MyCohesity](#). MyCohesity is a secure, single sign-on (SSO) portal that provides fast and easy access to all of your Cohesity resources. If you do not have a MyCohesity account, [sign up](#) for an account to access all your Cohesity resources from a single dashboard. For more information about MyCohesity, review [this page](#).

To sign in to Helios:

1. Go to the [MyCohesity](#) website.
2. Enter your MyCohesity username and password and click **Log in**.

Note: The MyCohesity homepage displays all tiles when you are not logged in. When you log in, you can only see the tiles you are allowed to access. If you do not see a tile, you do not have access to that resource. For more information, see this [knowledge base article](#).

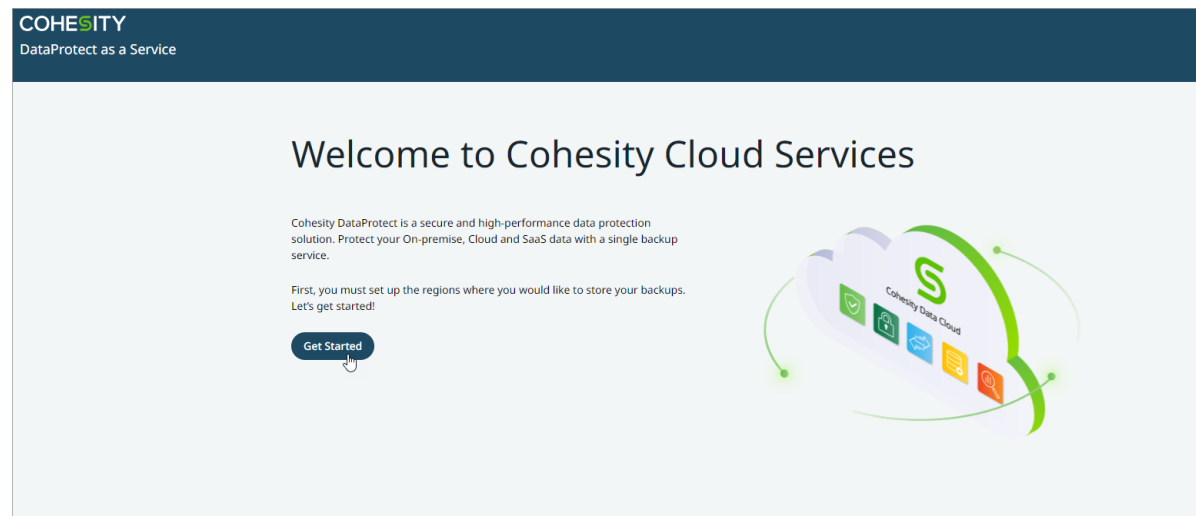
3. On the **Helios** tile, click **Launch**:



4. On the **Cohesity Data Cloud** landing page, click the **Protection** solution area and then select **DataProtect as a Service**.

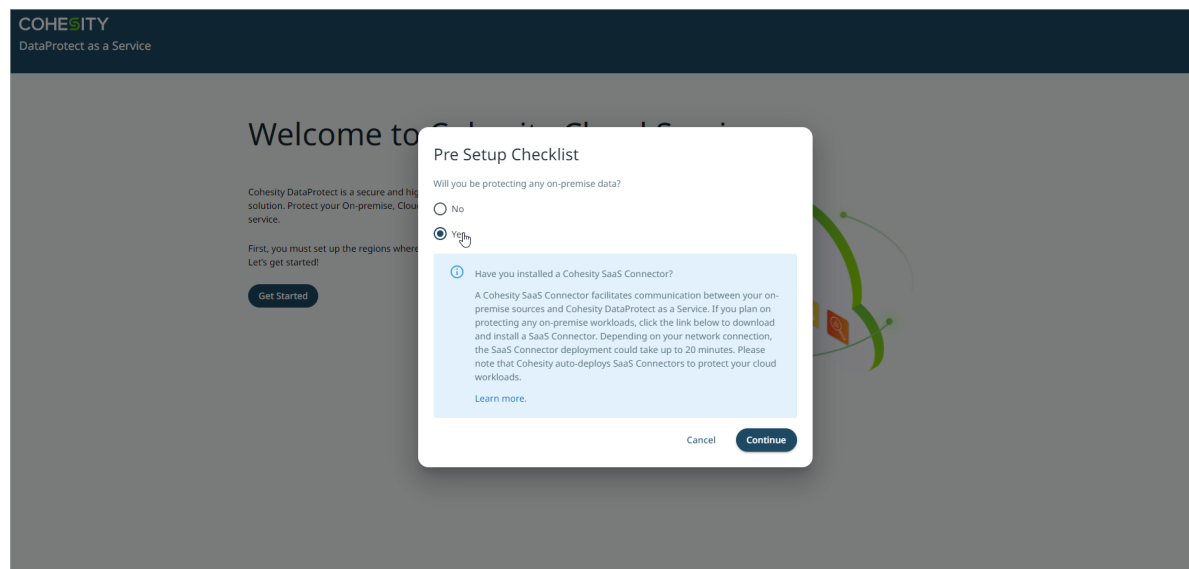


On the **Welcome to Cohesity Cloud Services** page, click **Get Started** to protect data sources from your data center and SaaS applications.



When prompted, if you are protecting on-premises data, select **Yes**. You must install a Cohesity SaaS Connector in your environment to protect on-premises data. A Cohesity SaaS Connector facilitates communication between your on-premise sources and Cohesity DataProtect as a Service. For more information, see [Deploy SaaS Connector](#).

Select **No**, if you are protecting SaaS workloads like Microsoft 365. And, click **Continue**.

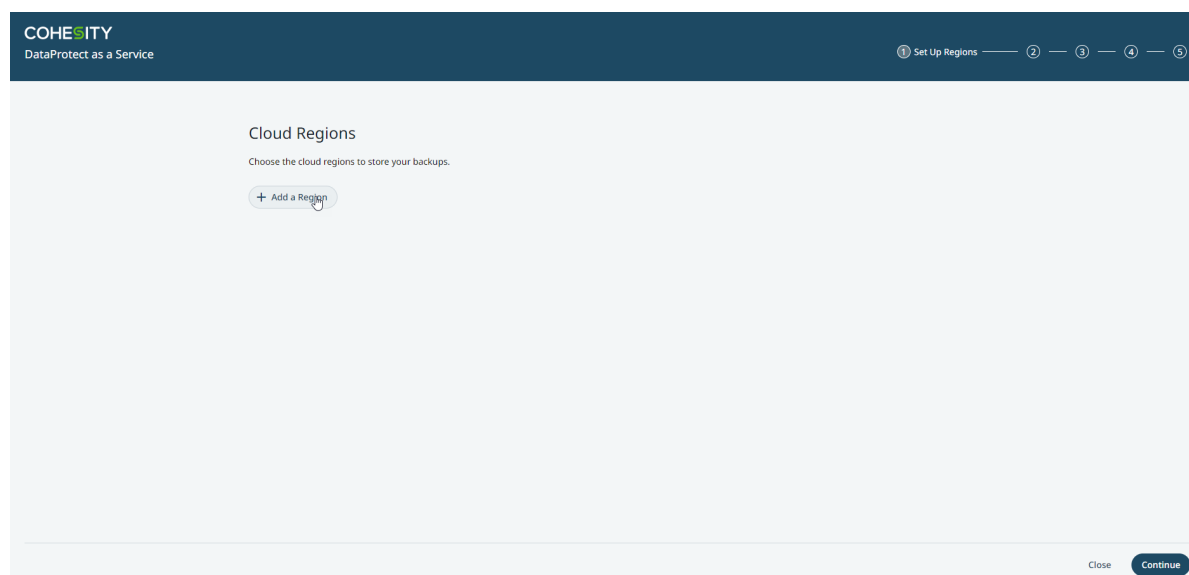


Select Regions and Encryption Key Management System

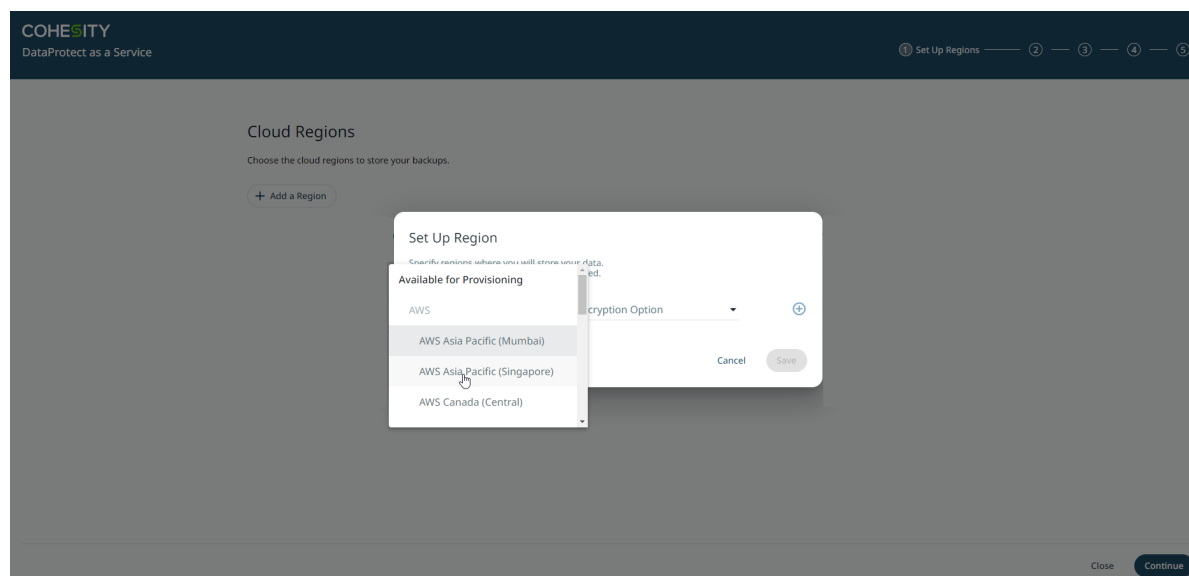
Before you can use Cohesity DataProtect as a Service, you need to select at least one cloud region for your data backups and choose your desired [encryption options](#) for securing your backups.

Important: Once data is backed up to one region, you cannot move it to another. To back your data up in another region, you can add that region and start protecting your data there.

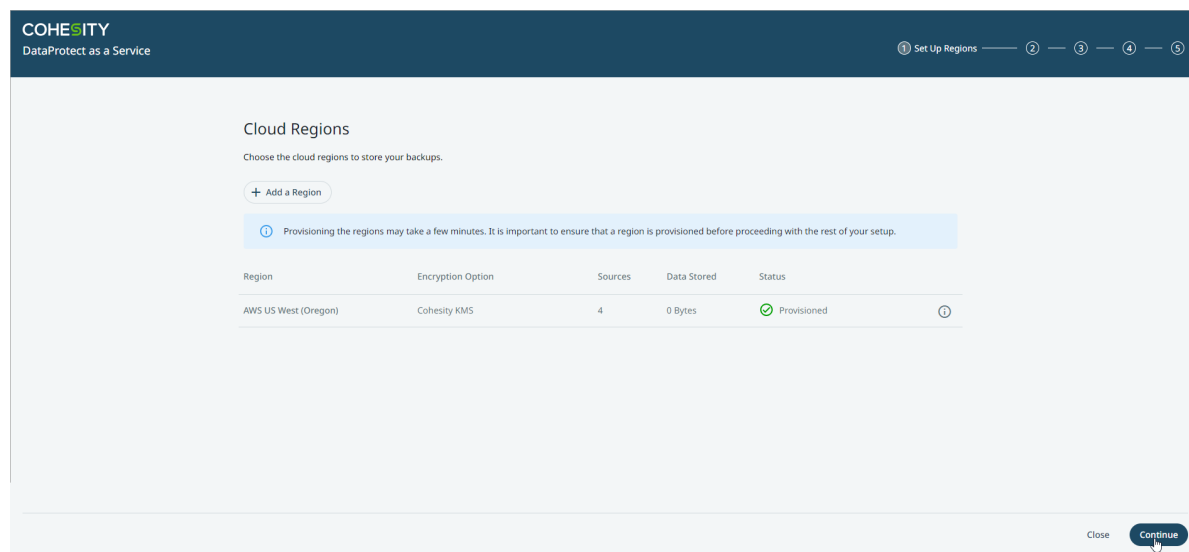
On the **Cloud Regions** page, click **Add a Region**.



From the **Set Up Region** dialog, select the **cloud region** for your data backups and choose the **encryption option**. For more information on the encryption options, see [Choose Key Management System \(KMS\)](#).



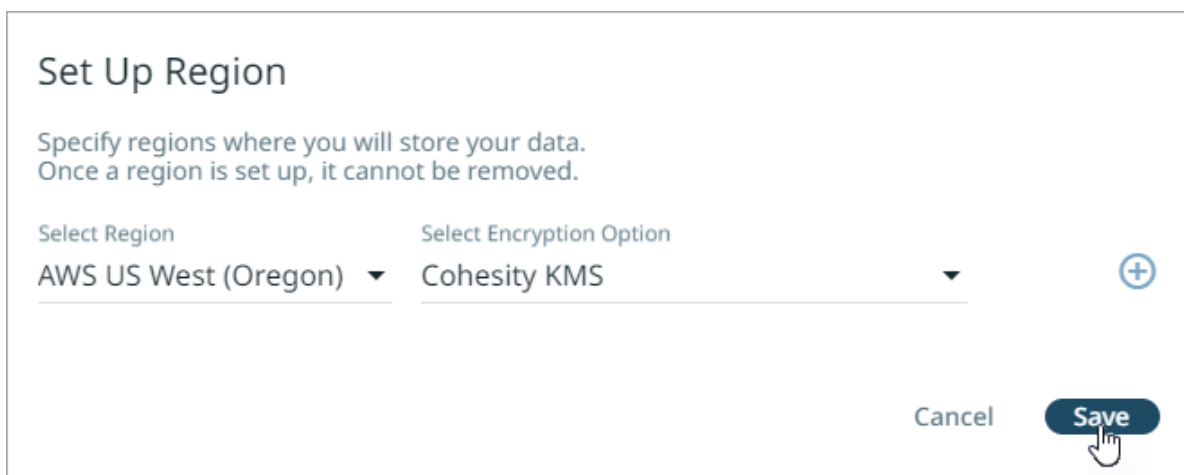
Once the cloud region is provisioned, click **Continue**.



Choose Key Management System (KMS)

In Cohesity DataProtect as a Service, all the data is encrypted both in flight and at rest. The encryption keys used for at-rest data encryption are the AWS Key Management System (KMS) or Azure Key Vault keys. Customers can choose to encrypt their data using Cohesity-generated AWS KMS or Azure Key Vault keys or bring their own AWS KMS or Azure Vault keys:

- **Cohesity KMS.** Depending on the region you select to store the data, Cohesity generates and uses unique AWS KMS keys or Azure Key Vault keys for each customer to encrypt their data.



Set Up Region

Specify regions where you will store your data.
Once a region is set up, it cannot be removed.

Select Region Select Encryption Option

AWS US West (Oregon) ▼ Cohesity KMS ▼

Cancel Save

- **Self-Managed KMS.** You can also use your own AWS encryption keys (Customer Master Keys) or Azure Key Vault keys.
 - To use your own AWS encryption keys (Customer Master Keys), see [Self-Managed KMS for AWS](#).
 - To use your own Azure Key vault keys, see [Self-Managed KMS for Azure](#).

In both encryption options, Cohesity uses AES-256 encryption keys called DEKs (Data Encryption Keys) to encrypt the data at rest. DEKs are generated using the AWS CMK and rotated every 4 hours. The Data Encryption Key is encrypted with the AWS CMK and stored along with the data—it is never stored in plain text.

Self-Managed KMS for AWS

To encrypt your data, you can use your own AWS encryption keys (Customer Managed Keys).

Prerequisites

Review and understand the following high level process of using your own AWS encryption keys (Customer Managed Keys):

1. You provide the CMK Amazon Resource Name (ARN) for the cloud region you selected.

Note: Cohesity supports both single and multi region self-managed KMS keys.

2. Cohesity generates the JSON for a key policy document that allows the DataProtect service to make API calls to your CMK.
3. You add the generated JSON contents to your AWS CMK's Policy in your AWS account.

Important: Cohesity recommends using the Cohesity-managed KMS for data encryption. If you choose the self-managed KMS, you are responsible for protecting the CMK keys used for data encryption. Note that if the CMS keys are compromised, then the data stored on Cohesity DataProtect as a Service will not be recoverable.

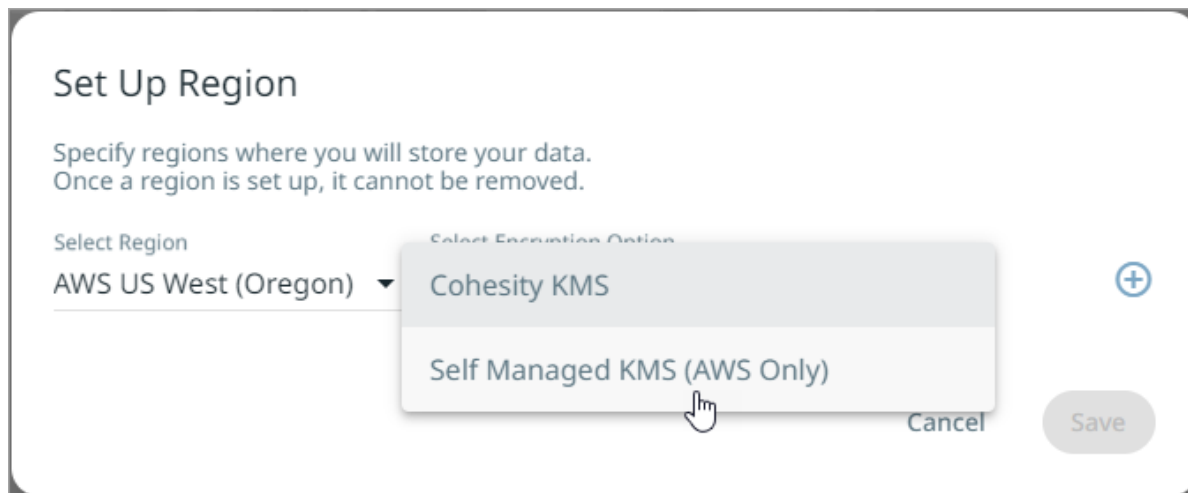
With the Self-Managed KMS option, you can audit the access calls made to your CMK to find important information, including when the CMK was used, the operation that was requested, the identity of the requester, and the source IP address. For more, see [Logging AWS KMS API calls with AWS CloudTrail](#) and [What Is AWS CloudTrail?](#) in the AWS documentation.

Note that you can also revoke CMK access to Cohesity at any time, after which Cohesity cannot decrypt the data stored in Cohesity DataProtect as a Service and all backup & recovery operations will fail.

Configure Self-Managed KMS for AWS

To use your own AWS KMS keys, on the **Set Up Region** dialog, perform the following:

1. Choose the **Region** and select the **Self Managed KMS (AWS Only)** as the **Encryption Option**.



2. Enter your **AWS Key ARN** for the selected region and click **Get JSON**.

Set Up Region

Specify regions where you will store your data.
Once a region is set up, it cannot be removed.

Select Region

AWS US West (Oregon) ▼

Select Encryption Option

Self Managed KMS (AWS Only) ▼

+

Enter your AWS Key ARN for the selected region below and copy the generated JSON script into your AWS account.

Enter Your AWS Encryption Key ARN

arn:aws:kms:us-east-2: [redacted]

Get JSON

Waiting for AWS Encryption Key ARN

Add the above JSON to the Key Policy section of your AWS KMS Key.
This step allows Cohesity to use your keys for data encryption.

Cancel

Save

3. Copy the generated JSON script.

Set Up Region

Specify regions where you will store your data.
Once a region is set up, it cannot be removed.

Select Region ▼ Select Encryption Option ▼ ⊕

AWS US West (Oregon) Self Managed KMS (AWS Only)

Enter your AWS Key ARN for the selected region below and copy the generated JSON script into your AWS account.

Enter Your AWS Encryption Key ARN
arn:aws:kms:us-east-2: Get JSON

```
{
  "Sid": "Allow use of the key",
  "Effect": "Allow",
```

Add the above JSON to the Key Policy section of your AWS KMS Key.
This step allows Cohesity to use your keys for data encryption.

Cancel Save

Go to your **AWS CMK** and add the copied JSON script under the **"Statement"** element in the **Key Policy** section as shown below:

Key Management Service (KMS)

AWS managed keys
Customer managed keys

Custom key stores
AWS CloudHSM key stores
External key stores

KMS > Customer managed keys > Key ID: 00000000-0000-0000-0000-000000000000 Key actions ▼ Edit

General configuration

Alias arn:aws:kms:us-east-2:123456789012:key/alias-123456789012	Status Enabled	Creation date 2021-10-17T10:10:10Z
ARN arn:aws:kms:us-east-2:123456789012:key/alias-123456789012	Description -	Regionality us-east-2

Key policy

Edit

```
{
  "Version": "2012-10-17",
  "Id": "key-default-1",
  "Statement": [
    {
      "Sid": "Enable IAM User Permissions",
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam:636428481998:root"
      },
      "Action": "kms:*",
      "Resource": "*"
    }
  ]
}
```

4. Click **Save**.

Set Up Region

Specify regions where you will store your data.
Once a region is set up, it cannot be removed.

Select Region

AWS US West (Oregon) ▼

Select Encryption Option

Self Managed KMS (AWS Only) ▼

+

Enter your AWS Key ARN for the selected region below and copy the generated JSON script into your AWS account.

Enter Your AWS Encryption Key ARN

arn:aws:kms:us-east-2: [redacted]

Get JSON

```
{
  "Sid": "Allow use of the key",
  "Effect": "Allow",
```

Add the above JSON to the Key Policy section of your AWS KMS Key.
This step allows Cohesity to use your keys for data encryption.

Cancel

Save

Self-Managed KMS for Azure

To encrypt your data, you can use your own Azure Key vault keys.

Prerequisites

Before you use your own Azure Key vault keys for Cohesity DataProtect as a Service to encrypt the data in flight and at rest.

Ensure the following prerequisites:

1. Copy the details of your Azure Key Vault:
 - **Directory ID**
 - **Vault URL**
 - **Key Identifier**

For detailed instructions, see [Copy the Required Details of your Azure Key Vault on the Azure Portal](#).

2. Register an app on Azure Portal for Cohesity to access your Azure Key Vault. For detailed instructions, see [Register an App on the Azure Portal](#).
3. Configure the required access policy for the app to access your key vault. For detailed instructions, see [Configure the Required Access Policy for the App to Access your Key Vault](#).

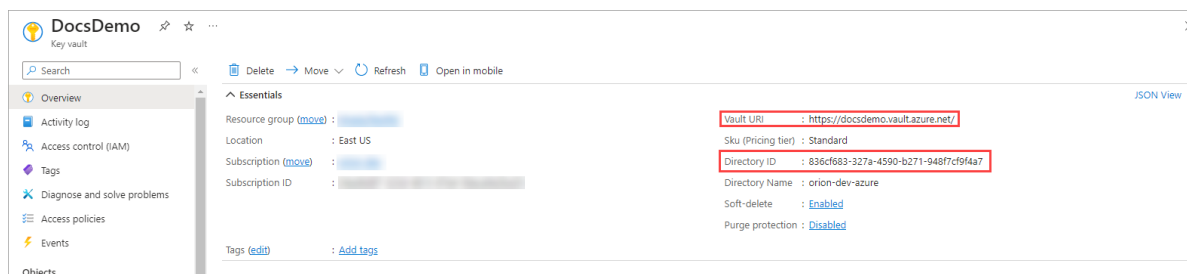
Copy the Required Details of your Azure Key Vault on the Azure Portal

From the Azure Portal, copy the following details related to your Azure Key vault:

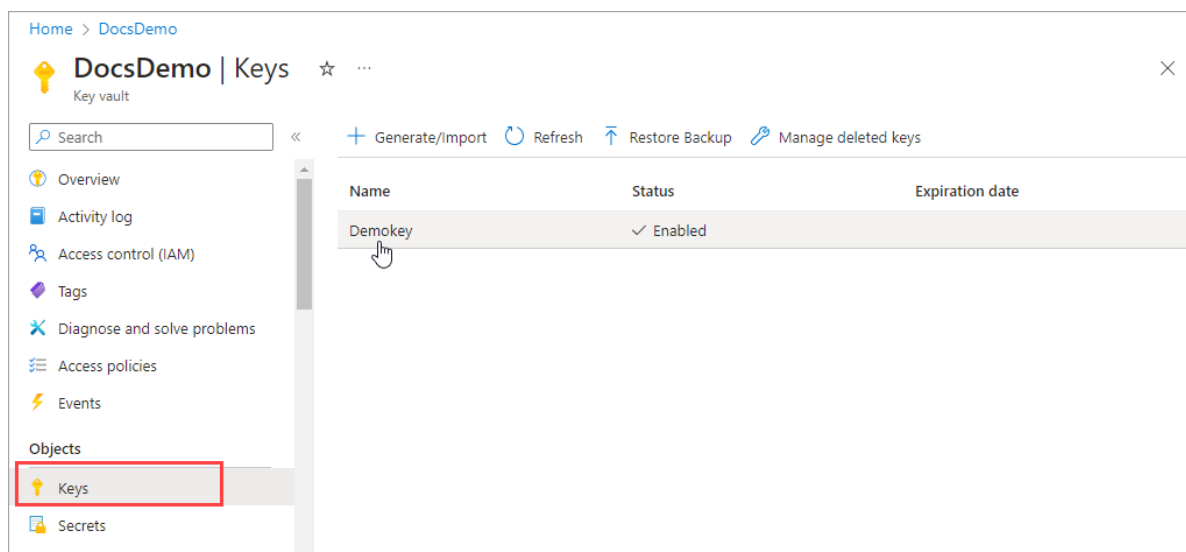
- **Directory ID**
- **Vault URL**
- **Key Identifier**

To copy the details related to your Azure Key vault:

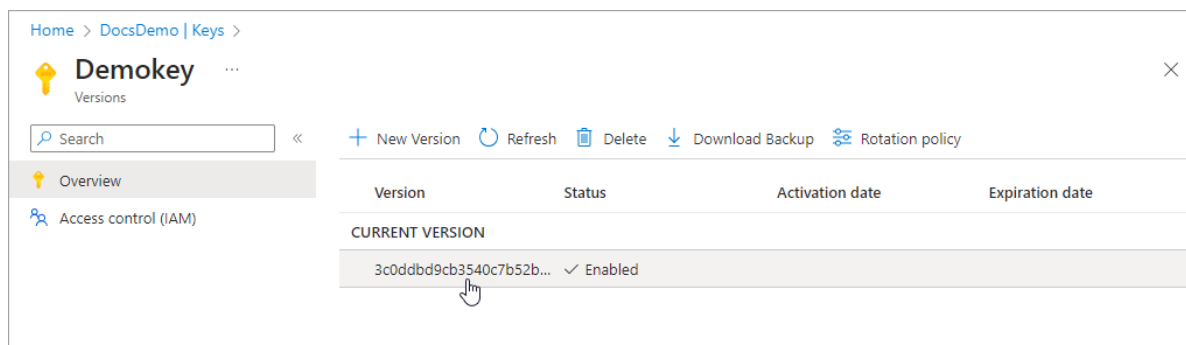
1. Log in to the Azure Portal with your Microsoft Azure administrator user credentials.
2. Under **Azure Services**, click **Key Vaults**.
3. In the **Key Vaults** page, select your Azure Key vault.
4. On the **Azure Key vault** page, perform the following:
 1. Copy the **Directory ID** and the **Vault URL**.



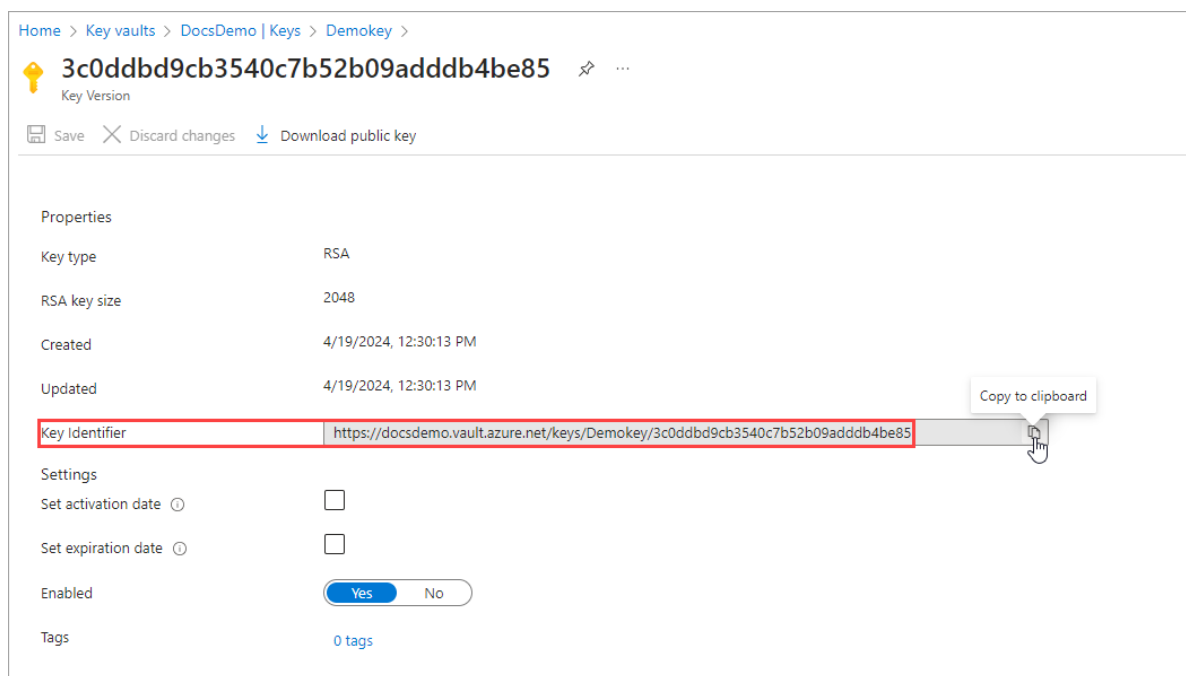
2. Under **Objects**, select **Keys** and click on your key.



3. Click on the current version of the key.



4. Copy the **Key Identifier**.



Register an App on the Azure Portal

To register an app on Azure Portal:

1. Log in to the [Azure Portal](#) with your Microsoft 365 administrator user credentials.
2. Search for and select **App registrations**.
3. Click **New registration**.
4. Enter a **Name** for your app, select the **Supported account types** that can access the app, and click **Register**.
5. On the **RegisteredApp** page, click **Overview** and copy the **Application (client) ID**. You need to use this ID when you set up the Self Managed KMS for the Azure region on Cohesity DataProtect as a Service.
6. Next, you need to add a new client secret (or "App Secret Key") that the app can use to prove its identity when requesting a token. Click **Certificates & Secret > New client secret**.
7. Enter a **Description**, select the **recommended number of days** under **Expires**, and click **Add**.
8. Under **Client secrets**, click the **Copy** button next to the string under **Value**. Also,

note down the **Expiry** date.

Home > App registrations > DocsDemoApp

DocsDemoApp | Certificates & secrets

Search Got feedback?

Overview
Quickstart
Integration assistant

Manage

Branding & properties
Authentication
Certificates & secrets
Token configuration
API permissions
Expose an API
App roles
Owners
Roles and administrators
Manifest

Support + Troubleshooting
Troubleshooting
New support request

Got a second to give us some feedback? →

Credentials enable confidential applications to identify themselves to the authentication service when receiving tokens at a web addressable location (using an HTTPS scheme). For a higher level of assurance, we recommend using a certificate (instead of a client secret) as a credential.

Application registration certificates, secrets and federated credentials can be found in the tabs below.

Certificates (0) **Client secrets (1)** Federated credentials (0)

A secret string that the application uses to prove its identity when requesting a token. Also can be referred to as application password.

+ New client secret

Description	Expires	Value	Copy to clipboard	Get ID
Demokeys	10/19/2024	pG98Q~LD7EM5G.cr0GVXN6xYACzmImr... e1fbcd7e-0303-4f6a-8281-97895914469e		

Configure the Required Access Policy for the App to Access your Key Vault

The app must have proper authentication and authorization to access a key vault. Access to your key vault can be provided using an [Azure RBAC](#) or [Key Vault access policy](#).

Note: Microsoft recommends using Azure RBAC instead of Key Vault access policy.

If you choose to use [Azure RBAC](#), ensure that you create a custom role and add the following entries under "permissions" in the JSON file of the custom role. Then, assign the custom role to the app.

```
"permissions": [
  {
    "actions": [
      "Microsoft.KeyVault/vaults/keys/read"
    ],
    "notActions": [],
    "dataActions": [
      "Microsoft.KeyVault/vaults/keys/read",
      "Microsoft.KeyVault/vaults/keys/encrypt/action",
      "Microsoft.KeyVault/vaults/keys/decrypt/action"
    ],
    "notDataActions": []
  }
]
```

If you choose to use the [Key Vault access policy](#), then ensure that you create a custom access policy with the following permissions and assign it to the app:

Type	Permissions
Key Management Operations	Get
Cryptographic Operations	<ul style="list-style-type: none"> • Decrypt • Encrypt

Sample:

[Home](#) > [DocsDemo | Access policies](#) >

Create an access policy

DocsDemo

Key Management Operations

☐ Select all

☒ Get

☐ List

☐ Update

☐ Create

☐ Import

☐ Delete

☐ Recover

☐ Backup

☐ Restore

Secret Management Operations

☐ Select all

☐ Get

☐ List

☐ Set

☐ Delete

☐ Recover

☐ Backup

☐ Restore

Privileged Secret Operations

☐ Select all

☐ Purge

Certificate Management Operations

☐ Select all

☐ Get

☐ List

☐ Update

☐ Create

☐ Import

☐ Delete

☐ Recover

☐ Backup

☐ Restore

☐ Manage Contacts

☐ Manage Certificate Authorities

☐ Get Certificate Authorities

☐ List Certificate Authorities

☐ Set Certificate Authorities

☐ Delete Certificate Authorities

Privileged Certificate Operations

☐ Select all

☐ Purge

Cryptographic Operations

☐ Select all

☒ Decrypt

☒ Encrypt

☐ Unwrap Key

☐ Wrap Key

☐ Verify

☐ Sign

Privileged Key Operations

☐ Select all

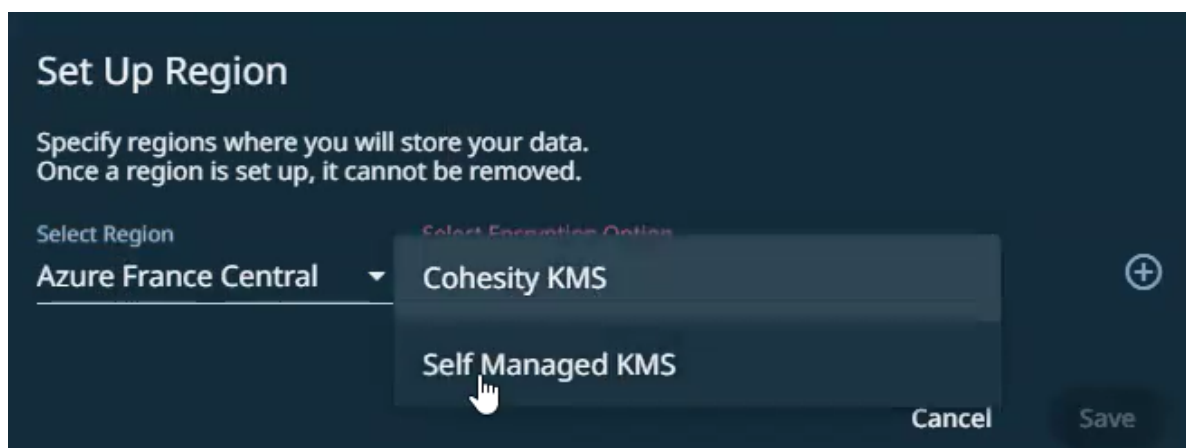
☐ Purge

☐ Release

Configure Self-Managed KMS for Azure

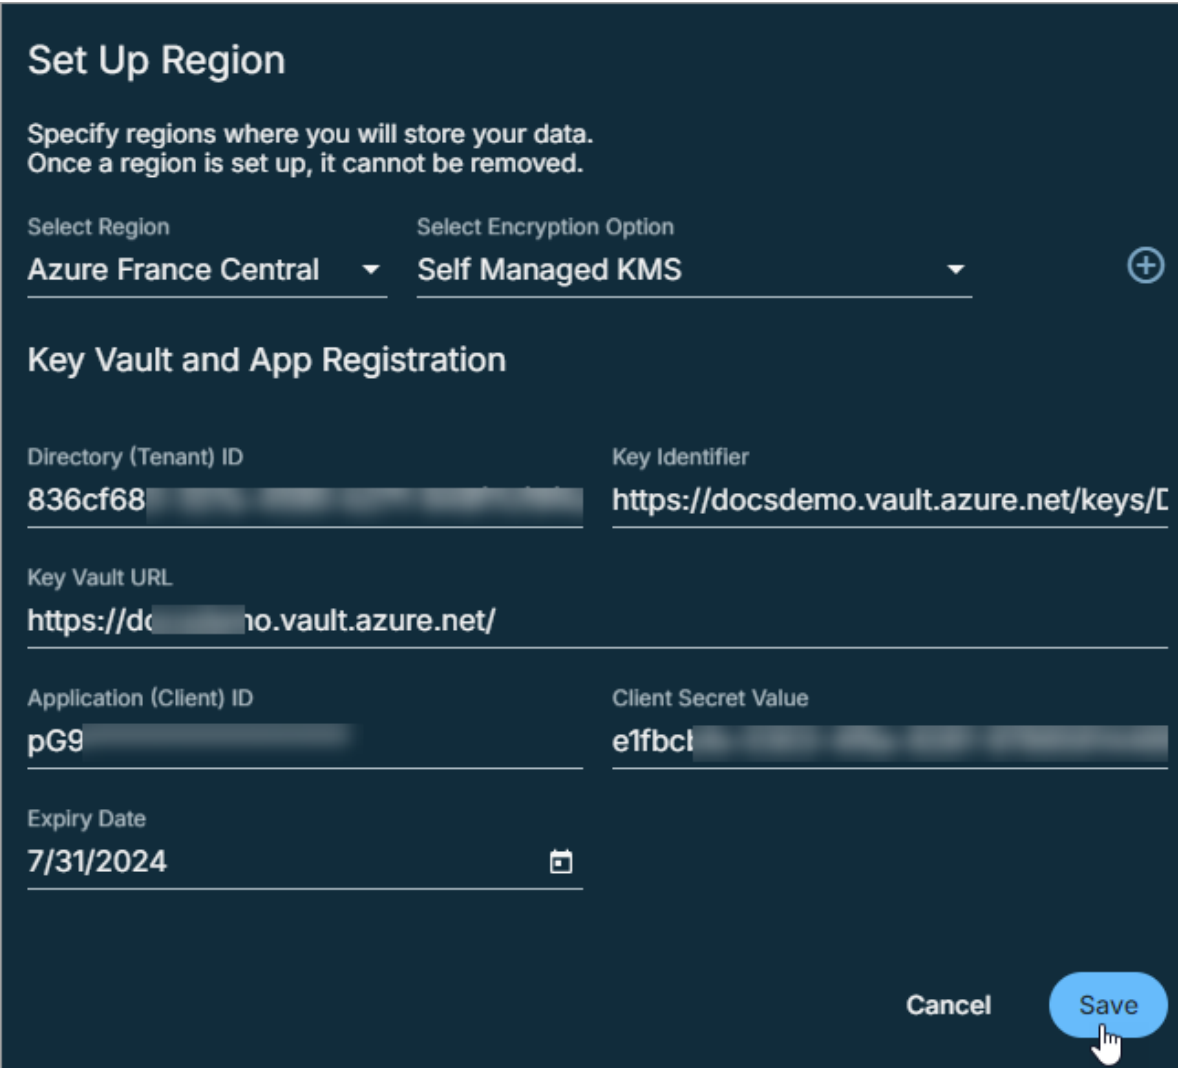
To use your own Azure Key Vault keys, on the **Set Up Region** dialog, perform the following:

1. Choose the **Region** and select the **Self Managed KMS** as the **Encryption Option**.



2. Under the **Key Vault and App Registration** sections, enter the following details:
 1. The **Directory (Tenant) ID**, **Key Identifier**, and **Key Vault URL** that you have [copied from your Azure Key Vault on Azure Portal](#).
 2. The **Application (Client) ID**, **Client Secret Value**, and the **expiry date** of the Client Secret that you have [noted down while registering the app on Azure](#)

Portal.



The screenshot shows a dark-themed configuration window. The top section, 'Set Up Region', has a title and a note: 'Specify regions where you will store your data. Once a region is set up, it cannot be removed.' Below this are two dropdown menus: 'Select Region' with 'Azure France Central' selected, and 'Select Encryption Option' with 'Self Managed KMS' selected. A plus icon is to the right. The bottom section, 'Key Vault and App Registration', contains several input fields: 'Directory (Tenant) ID' with '836cf68...', 'Key Identifier' with 'https://docsdemo.vault.azure.net/keys/E...', 'Key Vault URL' with 'https://docsdemo.vault.azure.net/', 'Application (Client) ID' with 'pG9...', 'Client Secret Value' with 'e1fbct...', and 'Expiry Date' with '7/31/2024' and a calendar icon. At the bottom right are 'Cancel' and 'Save' buttons, with a hand cursor clicking the 'Save' button.

Set Up Region

Specify regions where you will store your data.
Once a region is set up, it cannot be removed.

Select Region: Azure France Central ▼

Select Encryption Option: Self Managed KMS ▼

Key Vault and App Registration

Directory (Tenant) ID: 836cf68...

Key Identifier: https://docsdemo.vault.azure.net/keys/E...

Key Vault URL: https://docsdemo.vault.azure.net/

Application (Client) ID: pG9...

Client Secret Value: e1fbct...

Expiry Date: 7/31/2024

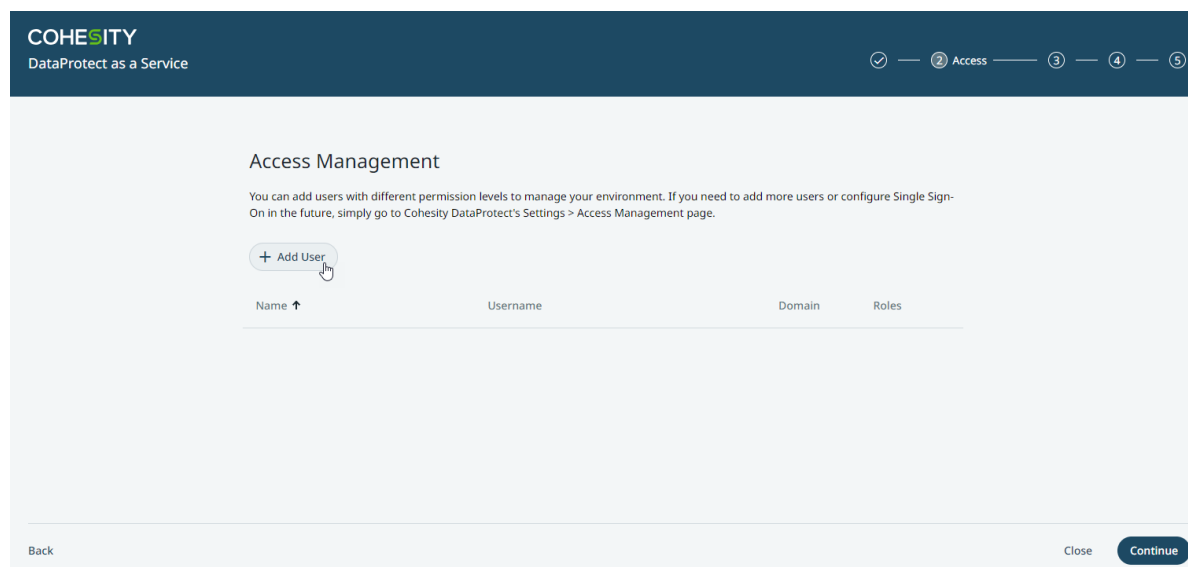
Buttons: Cancel, Save

3. Click **Save**.

Add Users

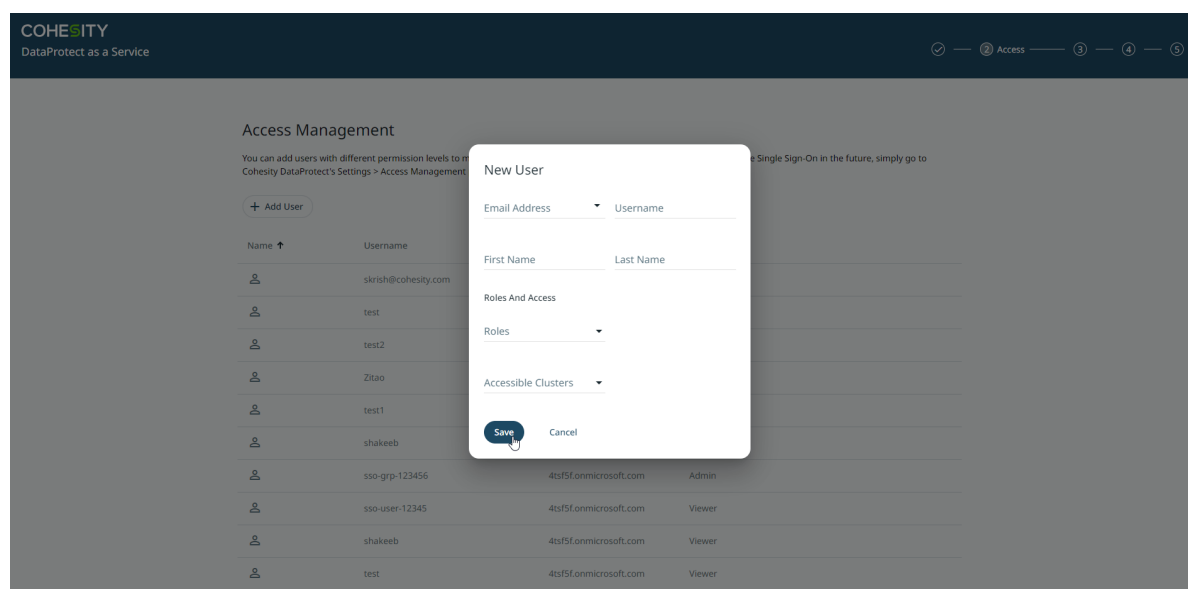
To manage user access to your Cohesity DataProtect as a Service, we recommend that you add users. Once you create them, your users can start using your Cohesity DataProtect as a Service with their own logins. You can add users with different permission levels to manage your environment. For more information, see [Access Management](#).

On the **Access Management** page, click **Add User** to add users.



On the **New User** dialog, perform the following:

1. Enter the following details:
 - **Username.** The user's email address.
 - **Email Address.** The user's email address again.
 - **First Name.** The user's first name in Cohesity DataProtect as a Service.
 - **Last Name.** Typically, the domain of your email address.
2. Under **Roles and Access**, assign an appropriate **Role** to this user and select the **Clusters** that this user can access. See [Roles](#) for more information.
3. Click **Save**.



Once you have added the users, click **Continue**.

COHESITY
DataProtect as a Service

Access Management

You can add users with different permission levels to manage your environment. If you need to add more users or configure Single Sign-On in the future, simply go to Cohesity DataProtect's Settings > Access Management page.

[+ Add User](#)

Name ↑	Username	Domain	Roles
	skrishh@cohesity.com	cohesity.com	Super Admin
	test	cohesityso.onmicrosoft.com	Super Admin
	test2	cohesityso.onmicrosoft.com	Super Admin
	Zitao	cohesity.com	Super Admin
	test1	cohesityso.onmicrosoft.com	Admin
	shakeeb	cohesityso.onmicrosoft.com	Viewer

Close [Continue](#)

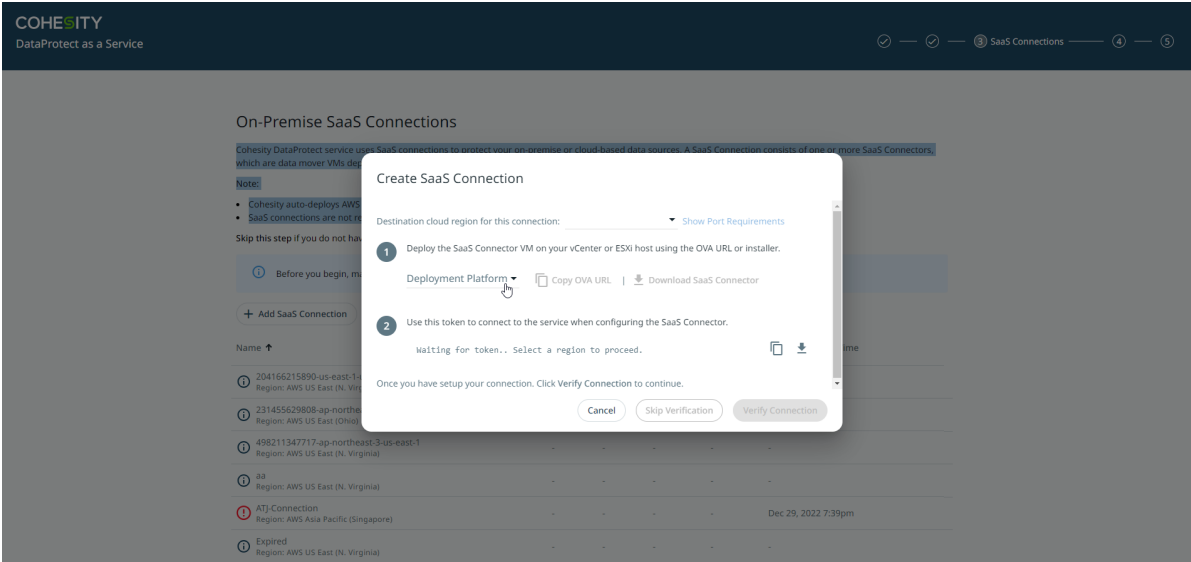
Add On-Premise SaaS Connections

Cohesity DataProtect as a Service uses SaaS connections to protect your on-premise or cloud-based data sources. A SaaS Connection consists of one or more SaaS Connectors, which are data mover VMs deployed in your VMware, Hyper-V, or AWS environment.

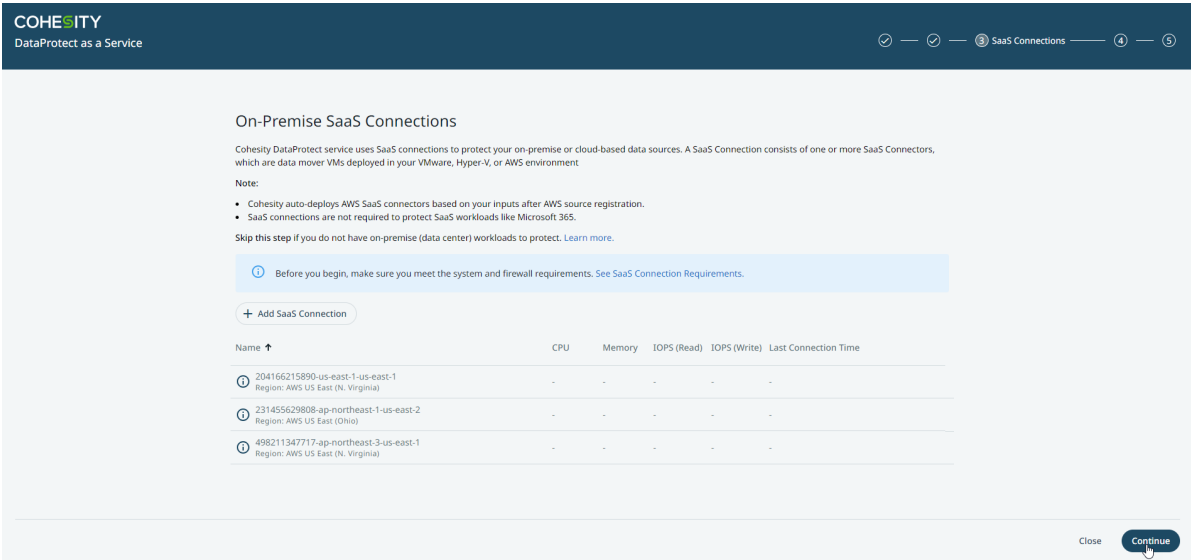
Note:

- Cohesity auto-deploys AWS SaaS connectors based on your inputs after AWS source registration.
- Cohesity auto-deploys Azure SaaS connectors based on your inputs after Azure source registration.
- SaaS connections are not required to protect SaaS workloads like Microsoft 365.

Deploy one or more SaaS Connector VMs. On the **On-Premise SaaS Connections** page, click **Add SaaS Connection** to create a SaaS Connection. Depending on the data source you want to protect, you must deploy SaaS Connectors. For more information, see [Deploy SaaS Connector](#).

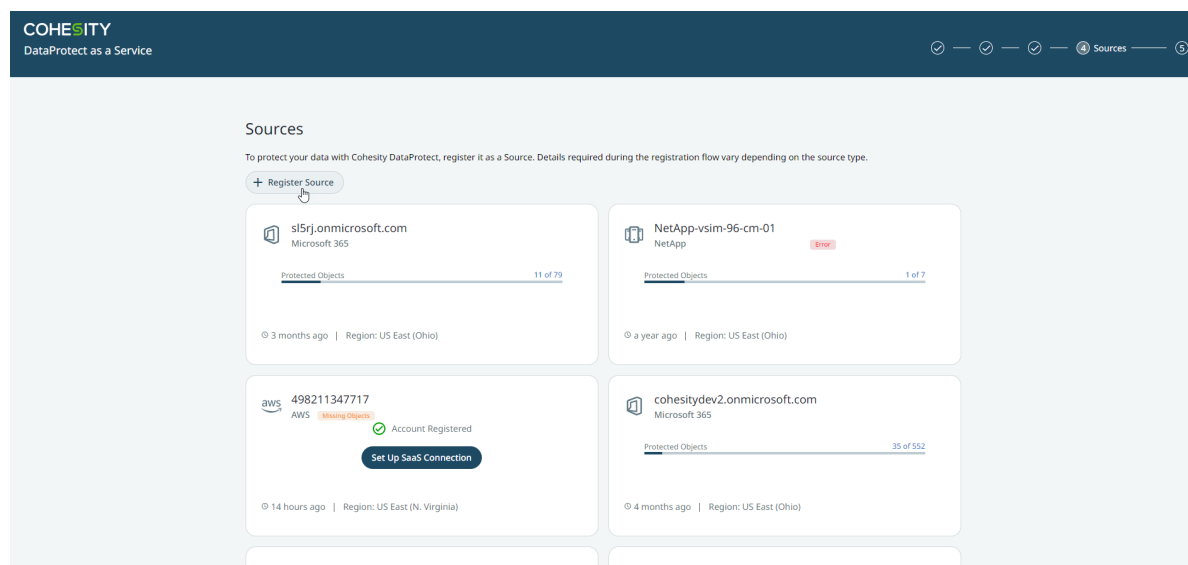


Once you have added the SaaS connections, click **Continue**.

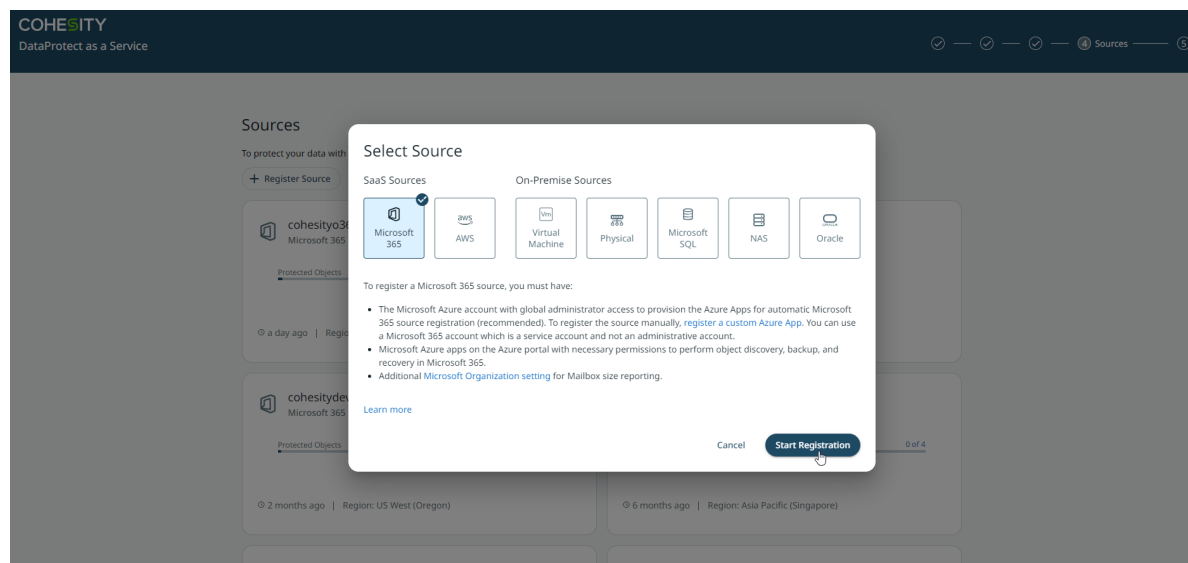


Register a Source

To start protecting your data, register your data sources. On the **Sources** page, click **Register Source** to register your data sources.



Select your data source on the **Select Source** dialog, and click **Start Registration**.

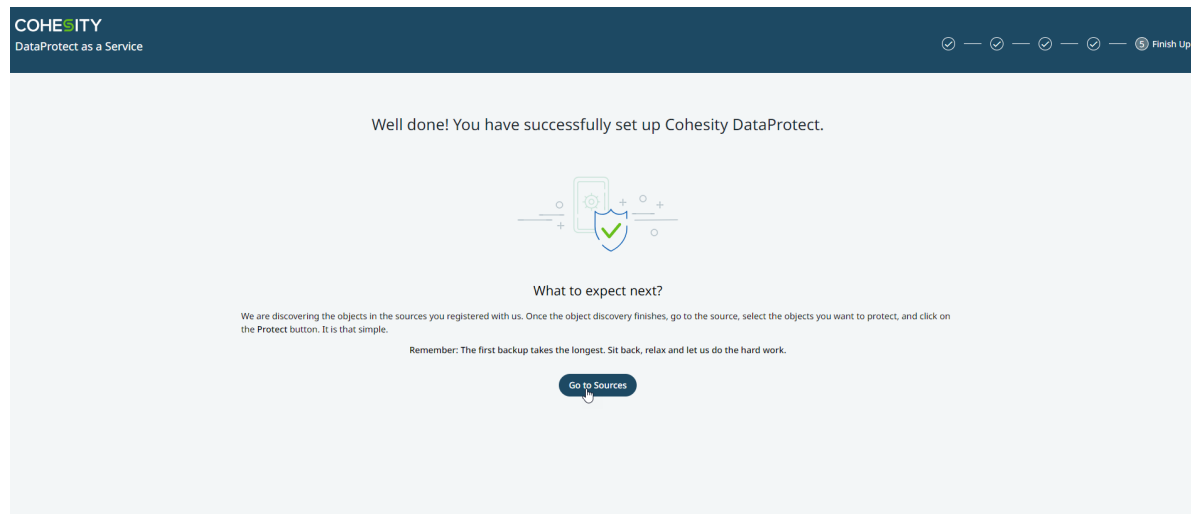


The process for registering each source is unique to the type of source. See the registration steps for:

- [VMware](#)
- [Generic NAS](#)
- [Isilon NAS](#)
- [NetApp ONTAP](#)
- [Microsoft SQL Server](#)
- [Microsoft 365](#)
- [Physical Servers](#)

- [AWS Account](#)
- [Microsoft Azure](#)
- [Hyper-V](#)
- [Oracle](#)

After you have registered your data sources, click **Continue**. Then, click **Go to Sources** to start protecting your data sources.



Protect a Source

Once you have registered a source in Cohesity DataProtect as a Service, you can start protecting the objects, volumes, and files in that source. For detailed instructions, see the respective Workload Types:

- [Protect VMware VMs](#)
- [Protect Hyper-V VMs](#)
- [Protect NAS Sources](#)
- [Protect Oracle Databases](#)
- [Protect Physical Servers](#)
- [Protect Microsoft SQL Server Databases](#)
- [Protect Your Amazon EC2 Instances](#)
- [Protect Your Amazon RDS](#)
- [Protect Your Amazon S3 Buckets](#)
- [Protect Azure Virtual Machines](#)
- [Protect your Microsoft Azure SQL Database](#)

Recover Protected Objects & Files

After you [protect a source](#), you can recover the objects and files from your backups, to their original or a new location. To get started:

- [Set Up Recovery](#)
- [Recover Objects & Volumes](#)
- [Recover Files & Folders](#)

Note: The steps in this article comprise the general recovery process. For workload-specific details, see [Supported Workload Types](#) below.

Set Up Recovery

The screenshot shows the COHESITY DataProtect interface. On the left is a sidebar with navigation links: Dashboard, Sources, Policies, Activity, Alerts, Reporting, and Settings. The main area is titled 'Sources' and features a search bar and a 'Register Source' button. Below this, there are four source categories, each with a table of sources:

- AWS (1)**: A table with one source (ID: 418011677528) in the US EAST (N Virginia) region, with 1 Protected and 12 Unprotected objects, last refreshed 8 hours ago.
- Generic NAS (1)**: A table with one source (Path: \\10.2.165.239\share) in the US East (Ohio) region, with 0 Protected and 1 Unprotected object, last refreshed 6 months ago.
- Isilon (1)**: A table with one source (Path: Ruby-Isi-Sim-New) in the US East (Ohio) region, with 3 Protected and 4 Unprotected objects, last refreshed an hour ago.
- NetApp (1)**: A table with one source (Path: Ruby-Isi-Sim-New) in the US East (Ohio) region, with 3 Protected and 4 Unprotected objects, last refreshed an hour ago.

To recover protected objects & volumes or files & folders:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click into the **Source** name.
3. Above the tree, select **Object Protection > Protected**.
4. Use the filters, search box, and views to locate the objects or files you need.
5. To recover:

- **Objects (VMs) or NAS volumes**, continue with [Recover Objects & Volumes](#) below.
- **Files and folders**, continue with [Recover Files & Folders](#) below.

Tip: You can also use Global Search to locate, filter, and select the objects you need. Click the **Global Search** box at the top or type **slash (/)** anywhere to start your search.

Recover Objects & Volumes

To recover protected objects (VMs or NAS volumes):

1. Locate and select them, and then click **Recover** at the top.
2. In the **VM Protection Group Selection** dialog box, review the list of Virtual Machines in the Protection Group and click **Confirm**. This pop-up is displayed for vCenter, Standalone ESXi Host, and vCloud Director.


The **New Recovery** form will be displayed with the **Latest** snapshot (protection run). You can click the View Selected Objects option to view the details of the Virtual Machines selected.
3. If you need to recover from an earlier snapshot, click the Edit icon to select a new recovery point.
 - For each object under **Selected**, you can click the **Edit** icon to open the **Recovery Point** calendar. Click **List** to view the available recovery points by timestamp and click one.
 - Click **Select Recovery Point**.
 - Click **Next: Recover Options** to return to the form.
4. Under **Recover To**, select **Original Location** or **New Location**.
 - **For VMs:** If you choose **New Location**, select a **Registered Source**, **Resource Pool**, **Datastores**, and the **VM Folder**.
 - **For NAS volumes:** If you choose **New Location**, select a **Registered Source** and the **Volume**.
5. Select your [Recovery Options](#).
6. Click **Recover**. Review the Recovery Summary and click **Confirm**.

Cohesity DataProtect as a Service opens the **Activity** page, showing your file recovery task as it runs, along with the recovery progress on the right.

Tip: The **Activity** page also shows the entire history of all protection runs and recovery tasks, including any that are in progress.

Recover Files & Folders

To recover a specific file or files (or the folders containing them) from a protected source:



1. Locate the source object containing the files and click **Recover Files** on the row for that object to open the **Select Files** form.
2. If you need to recover from an earlier snapshot, click the **Recovery Point** calendar drop-down to select the recovery point.
 - Click **List** to view the available recovery points by timestamp and click one.
 - Click **Apply**.
3. Click into the path to find the files and add them to the **Selected Items** list.
4. Choose how to recover your files: download locally or recover.
 - Click **Download Files** to open the **Activity** page, showing your file recovery task. Click into the recovery task and click **Download Files** a second time to save them to your local system.
 - Click **Save** to open the **New Recovery** form. Under **Recover To**, select **Original Location** or **New Location**.
 - If you choose **Original Location**, enter a **Username** and **Password** that has access to the original server. You can also enable **Recover to Alternate Path** to enter a new path on the original server.
 - If you choose **New Location**, select a registered **Source** and a **Target** (VM) or **Volume** (NAS). Enter a **Username** and **Password** that has access to that server and enter a **Recover To** path.
5. Select your [Recovery Options](#).
6. Click **Start Recovery**.

Cohesity DataProtect as a Service opens the **Activity** page, showing your file recovery task as it runs, along with the recovery progress on the right.

Tip: The **Activity** page also shows the entire history of all protection runs and recovery tasks, including any that are in progress.

Supported Workload Types

While the recovery steps outlined above are generally the same for each workload type, there are differences. For recovery details that are specific to each workload type, see:

- [Recover VMware VMs & Files](#)
- [Recover VMC \(on AWS\) VMs & Files](#)
- [Recover VMC \(on Azure\) VMs & Files](#)
- [Recover NAS Data](#)
- [Recover Microsoft SQL Server Databases](#)
- [Recover Microsoft 365:](#)
 - [Mailboxes, Emails, & Folders](#)
 - [OneDrives & OneDrive Contents](#)
 - [SharePoint Online Sites & Items](#)
 - [Teams](#)
- [Recover Physical Servers](#)
- [Recover Amazon](#)
 - [EC2 Instances](#)
 - [RDS Databases](#)
 - [S3 Buckets](#)
- [Recover Files Hyper-V VMs & Files](#)
- [Recover Oracle Databases](#)
- [Recover Microsoft Azure](#)
 - [Recover Azure Virtual Machines](#)
 - [Recover Azure SQL Databases](#)

Recovery Options

When you [recover objects & volumes](#) or [files & folders](#) in Cohesity DataProtect as a Service, you can configure many additional options. While the options differ among object types and files, they often include the options below, for:

- [Objects \(VMs & Volumes\) Options](#)
- [Files & Folders Options](#)
- [General Recovery Options](#)

Recovery Options for Objects (VMs & Volumes)

- **Overwrite Existing VM.** (*Applies when recovering to the original location*) Enable this option to recover the VM by deleting the original VM. The recovered VM will have the original VM name. Once you select this option, a pop-up dialog box will be displayed. Review the message, type **YES**, and click **Confirm**.

Important: The original VM is deleted before the recovery. Therefore a recovery failure will also lead to the loss of the original VM.

- **Attempt Differential Recovery.** By enabling this option, Cohesity attempts to recover the VM by overwriting only the difference between the original VM and the snapshot selected for recovery. Any newly added data in the original VM is deleted. This option is available only if you have selected **Overwrite Existing VM**, and you can learn more about its pros and cons in [Recover VMware VMs & Files](#).
- **Network.** By default, the VMs that are to be recovered do not have a virtual Network Interface Card (vNIC) attached. Enable the **Attach** option to attach a virtual Network Interface Card (vNIC) to each VM that is to be recovered.

With the **Attach** option enabled, the following options are displayed:

- **Network.** From the drop-down menu, select a network to attach the virtual Network Interface Card (vNIC) to a new network.
- **Start Connected.** Enable this option to connect to the new network when the VM reboots for each recovered VM. If this option is not selected, the VMs are not connected to any network on reboot.
- **Preserver MAC Address.** Enable this option to preserve the MAC address when recovering to an alternate location.
- **Rename.** Add **Prefix** and/or **Suffix** strings to the names of the new VMs created by this task.
- **Power State.** Disable **Power On** if you want the recovered VMs to remain powered off after they are created.
- **Continue on Error.** Enable **Continue recovery even if errors occur when recovering VMs** if you want the recovery task to continue even if errors occur when recovering the VMs. For example, if one of the VMs cannot be created, Cohesity will still attempt to create the other VMs.

Recovery Options for Files & Folders

- **Overwrite Existing File/Folder.** By default, this option is enabled to overwrite the existing files and folders. Disable this option to create the files and folders in the specified location. Any duplicate files are skipped.
- **Preserve File/Folder Attributes.** By default, this option is enabled and the ACLs, permissions, and timestamps are preserved for all files and folders. If you disable this option, then ACLs and permissions are not preserved. If you are recovering both folders and files, then folders will receive the new timestamps, but files retain their original timestamps. If recovering only files, then files will receive the new timestamps.

General Recovery Options

- **Continue on Error.** Enable this option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.
- **Task Name.** Change the default name of the recovery task.

Next > When you've made your choices, click **Start Recovery** to [recover the objects or files](#) to the selected location.

Deploy SaaS Connector

To register on-premises or cloud-based data sources with Cohesity DataProtect as a Service, you need to use a SaaS Connection to establish connectivity between your source and the service. A SaaS Connection consists of one or more SaaS Connectors, which are VMs that act as data movers between your data sources and the Cohesity DataProtect as a Service.

Important: SaaS Connectors are closed virtual appliances managed by Cohesity. Except as specifically instructed by Cohesity, Cohesity does not support customers installing third-party software or OS updates independent of the Cohesity release process, and any attempt to install third-party software or deploy updates sourced from sources other than Cohesity may void your existing contractual terms with Cohesity, including any applicable warranties.

To create a SaaS Connection, you must deploy one or more SaaS Connector VMs. Depending on how SaaS connectors are deployed, SaaS Connectors can be classified as:

- **User-deployed SaaS Connectors:** The user must deploy the SaaS Connectors manually on the source you want to protect.
- **Cohesity-deployed SaaS Connectors:** Cohesity will auto-deploy the SaaS Connectors on the source you want to protect.

The following table provides information about the supported SaaS Connectors.

User-Deployed SaaS Connectors	<ul style="list-style-type: none">• VMware SaaS Connectors• Hyper-V SaaS Connectors
Cohesity-Deployed SaaS Connectors	<ul style="list-style-type: none">• AWS SaaS Connectors• Azure SaaS Connectors

User-Deployed SaaS Connectors

User-deployed SaaS Connectors are the Connectors that you must deploy manually on the source you want to protect using Cohesity DataProtect as a Service. Cohesity supports the following user-deployed SaaS Connectors:

- [VMware SaaS Connector](#)
- [Hyper-V SaaS Connector](#)

User-Deployed SaaS Connector Requirements

Before deploying the SaaS Connector, review and understand the following requirements:

Supported Sources

You can deploy SaaS Connectors manually on the following sources:

- VMware VMs
- VMware Cloud on AWS
- Hyper-V VMs
- Physical Server
- NAS
- Oracle
- Microsoft SQL Server

User-Deployed SaaS Connector System Prerequisites

Ensure that the SaaS Connector VM that you deploy for your SaaS Connection meets the following system requirements:

- 4 CPUs
- 10 GB RAM
- Disk space
 - 50 GB for VMware
 - 40 GB for Hyper-V
- Outbound Internet connection

SaaS Connector Sizing Recommendations

We recommend that you have one SaaS Connector for each 160 VMs or 16 TB of source data. If you have more data, we recommend that you stagger their first full backups.

Note: These requirements are subject to change.

Check Firewall Ports

Ensure that the ports listed in the SaaS Connector Management section in the [Firewall Ports for User-Deployed SaaS Connectors](#) topic are open in your firewall to allow communication between the Cohesity SaaS Connector(s) and Cohesity Cloud Services.

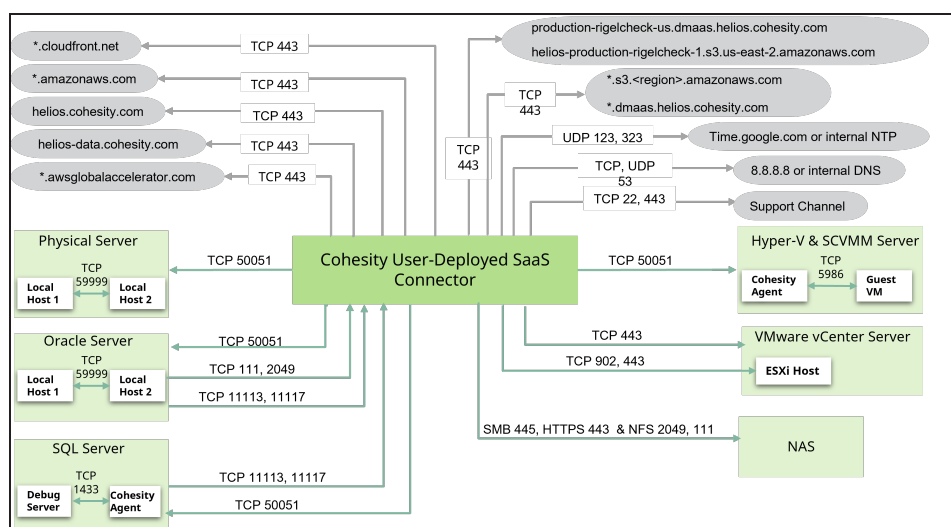
These firewall rules allow outgoing traffic from a SaaS Connector to the Cohesity DataProtect as a Service endpoint. The SaaS Connector opens a secure encrypted gRPC tunnel to the endpoint and uses it for both backup and recovery traffic.

The connectivity status between a SaaS Connection and the Cohesity DataProtect as a Service is displayed both in the SaaS Connection and the Cohesity DataProtect as a Servicedashboard.

Firewall Ports for User-Deployed SaaS Connectors

A typical SaaS Connector connects with the Cohesity DataProtect as a Service and the Data Sources. The following diagram shows the source, destination, ports, and protocols for traffic flow between the user-deployed SaaS Connector and the Data Sources and the user-deployed SaaS Connector and Cohesity DataProtect as a Service.

More information is provided in the sections that follow the diagram.



Legend

- Traffic flow between the SaaS Connector and Data Sources
- Traffic flow between the SaaS Connector and Cohesity Data Protect Service

SaaS Connector Management

Ensure that the following ports are open to allow communication between the Cohesity SaaS Connector(s) and Cohesity Cloud Services:

Source	Destination	Port	Protocol	Purpose
SaaS Connector	helios.cohesity.com	443	TCP	Connection used for control path
SaaS Connector	*.awsglobalaccelerator.com	443	TCP	Connection used for control path
SaaS Connector	*.s3.<region>.amazonaws.com *.dmaas.helios.cohesity.com	443	TCP	Connection used for data path
SaaS Connector	helios-data.cohesity.com	443	TCP	Used to send telemetry data
SaaS Connector	*.cloudfront.net	443	TCP	To download upgrade packages
SaaS Connector	production-rigelcheck-us.dmaas.helios.cohesity.com helios-production-rigelcheck-1.s3.us-east-2.amazonaws.com	443	TCP	Required to perform connectivity checks with the Cohesity Cloud Services.
SaaS Connector	8.8.8.8 or internal DNS	53	TCP, UDP	Host resolution.

Source	Destination	Port	Protocol	Purpose
SaaS Connector	time.google.com or internal NTP	123, 323	UDP	<p>Incoming NTP requests are detected by port 123.</p> <p>Chrony is the default implementation of NTP used by recent versions of CentOS and RHEL. Open port 323 if you want to use the Chronyc tool to monitor the synchronization status of Chrony and make changes if necessary.</p>
SaaS Connector	rt.cohesity.com	22 or 443	TCP	<p>The Cohesity Support Channel uses Secure Shell (SSH) and listens through port 22 or 443. Port 22 is used by default and can be updated to 443 using the Cohesity CLI. For more information, see Manage the Support Channel.</p>

Virtual Machines

VMware

Ensure that the following ports are open to allow communication between the Cohesity SaaS Connector(s) and VMware environment:

Source	Destination	Port	Protocol	Purpose
SaaS Connector	VMware vCenter	443	TCP	Required for making VMware API calls for backup and recovery over HTTPS/HTTPS (TLS).
SaaS Connector	ESXi Host(s)	443	TCP	Required for VMware Tools-based file and folder recoveries. Allow communication to each ESXi host over port 443 for VMware tools-based file and folder recovery, irrespective of whether the vCenter or Standalone ESXi host is registered with the Cohesity DataProtect as a Service.
SaaS Connector	ESXi Host(s)	902	TCP	Needs to be open on each ESXi host for VADP (vSphere Storage APIs for Data Protection), a vSphere API, that enables backup and restore operations via port 902.

Microsoft SCVMM and Hyper-V Servers

Ensure that the following ports are open to allow communication between the Cohesity SaaS Connector(s) and Hyper-V environment:

Source	Destination	Port	Protocol	Purpose
Cohesity Agent running on Standalone Hyper-V and SCVMM server	Guest VM (local host) running on Standalone Hyper-V and SCVMM Server	5986	TCP	Required for file and folder recovery operations.
SaaS Connector	Standalone Hyper-V and SCVMM Server	50051	TCP	Required for backup and recovery operations..

VMC on AWS

Ensure that the following ports are open to allow communication between the Cohesity SaaS Connector(s) and the VMC in the AWS environment:

Source	Destination	Destination Port	Protocol	Purpose
SaaS Connector	VMware vCenter	443	TCP	Required for making VMware API calls for backup and recovery over HTTPS/HTTPS (TLS). Note: Needs to be configured as a Management Gateway firewall rule in the VMC UI.

Source	Destination	Destination Port	Protocol	Purpose
SaaS Connector	ESXi Hosts	443	TCP	<p>Required for VMware Tools-based file and folder recoveries. Allow communication to each ESXi host over port 443 for VMware tools-based file and folder recovery, irrespective of whether the vCenter or Standalone ESXi host is registered with the Cohesity cluster.</p> <div>Note: Needs to be configured as a Management Gateway firewall rule in the VMC UI.</div>

Source	Destination	Destination Port	Protocol	Purpose
SaaS Connector	Any	Any	TCP	<p>Required for backup and recovery operations.</p> <div> Note: Cohesity recommends selecting "Any" in the Service column when configuring this Compute Gateway firewall rule in the VMC UI. </div>

VMC on AVS

Ensure that the following ports are open to allow communication between the Cohesity SaaS Connector(s) and the VMC in the AVS environment:

Source	Destination	Destination Port	Protocol	Purpose
SaaS Connector	VMware vCenter	443	TCP	<div>Required for making VMware API calls for backup and recovery over HTTPS/HTTPS (TLS).</div> <div>Note: Refer to the VMware cloud provider's documentation for updating the gateway firewall rules.</div>

Source	Destination	Destination Port	Protocol	Purpose
SaaS Connector	ESXi Hosts	443	TCP	<p>Required for VMware Tools-based file and folder recoveries. Allow communication to each ESXi host over port 443 for VMware tools-based file and folder recovery.</p> <div>Note: Refer to the VMware cloud provider's documentation for updating the gateway firewall rules.</div>

Source	Destination	Destination Port	Protocol	Purpose
SaaS Connector	Any	Any	TCP	<div>Required for backup and recovery operations.</div> <div>Note: Refer to the VMware cloud provider's documentation for updating the gateway firewall rules.</div>

Source	Destination	Destination Port	Protocol	Purpose
SaaS Connector	ESXi Host(s) hosted in AVS environment	902	TCP	<p>Each ESXi host must have port 902 open for VADP (vSphere Storage APIs for Data Protection), a vSphere API, allowing backup and restoring operations through port 902.</p> <div> Note: Refer to the VMware cloud provider's documentation for updating the gateway firewall rules. </div>

Physical Servers

Ensure that the following ports are open to allow communication between the Cohesity SaaS Connector(s) and Physical Servers:

Source	Destination	Port	Protocol	Purpose
SaaS Connector	Physical Windows or Linux Server	50051	TCP	Required for Backup and Recovery operations.

Source	Destination	Port	Protocol	Purpose
Local Host (Physical Windows or Linux Server)	Local Host (Physical Windows or Linux Server)	59999	TCP	Required for local-to-local communication for self-monitoring and debugging purposes.

Agent Upgrade

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes
Client	Cohesity SaaS Connector	80,443	TCP	For Agent upgrade from UI.

Databases

Oracle Servers

Ensure that the following ports are open to allow communication between the Cohesity SaaS Connector(s) and Oracle Server:

Source	Destination	Port	Protocol	Purpose
SaaS Connector	Oracle Server	50051	TCP	Required for Backup and Recovery operations.
Oracle Server	SaaS Connector	111, 2049	TCP	Required for Backup and Recovery operations in Linux servers.

Source	Destination	Port	Protocol	Purpose
Oracle Server	SaaS Connector	11113, 11117	TCP	Required for Backup and Recovery operations in Windows servers.
Local Host (Physical Windows or Linux Server)	Local Host (Physical Windows or Linux Server)	59999	TCP	Required for local-to-local communication for self-monitoring and debugging purposes.

Microsoft SQL Servers

Ensure that the following ports are open to allow communication between the Cohesity SaaS Connector(s) and Microsoft SQL Server:

Source	Destination	Port	Protocol	Purpose
SaaS Connector	MS SQL Host	50051	TCP	Required for Backup and Recovery operations.
MS SQL Host	SaaS Connector	11113, 11117	TCP	Required for Backup and Recovery operations.
MS SQL Host	Cohesity agent running on the MS SQL Host	1433	TCP	Default TCP port for MS SQL instances. Ensure port is open to allow communication between the MS SQL instance and the Cohesity Agent.

Network Attached Storage (NAS)

Ensure that the following ports are open to allow communication between the Cohesity SaaS Connector(s) and NAS Server:

Source	Destination	Port	Protocol	Purpose
SaaS Connector	NAS Server	2049 in NFS server & 111 in portmapper	NFS	To establish connection with the NAS source and carry out the Backup and Recovery operations.
		445	SMB	To establish connection with the NAS source and carry out the Backup and Recovery operations.
		443	HTTPS	Required for snapshot-based backups of Netapp, Isilon, Pure Storage, and so on.

Deploy VMware SaaS Connectors

You can install a VMware SaaS Connector using an installer OVA in your VMware environment, on a vCenter or ESXi host in your data center that has access to your data sources and meets the [SaaS Connection system and firewall requirements](#). Once deployed, each SaaS Connector acts as a virtual machine in your data center.

Tip: For better performance and redundancy, we recommend deploying at least two SaaS Connectors for each SaaS Connection in your data center. To add (or remove) a SaaS Connector, see [Manage User-Deployed SaaS Connections](#).

Cohesity can concurrently back up individual VM disks using multiple SaaS connectors, leading to faster backups. Cohesity determine the concurrency based on various factors, such as disk size and the workload of the SaaS connectors. Cohesity cannot recover individual VM disks using multiple SaaS Connectors because VMware does not support it.

All the data that a SaaS Connection handles, from your sources to the cloud storage where your backups reside, is encrypted in flight and at rest.

Create VMware SaaS Connection

To create a VMware SaaS Connection:

1. In **DataProtect as a Service**, navigate to **Settings > SaaS Connections**.
2. Click **New Connection**.
3. In the **Create SaaS Connection** dialog, select the following:
 - a. From the **Destination Cloud Region for this Connection** drop-down, select a region for your data backups.
 - b. From the **Deployment Platform** drop-down, select **VMware**.
 - c. To deploy the SaaS Connector in your data center, do one of the following:
 - i. Copy the OVA URL.
 - ii. Download the OVA file.
4. Click **Skip Verification** or **Cancel**. After deploying the SaaS Connector OVA, you will need to verify the connection later in Step 12.

Create SaaS Connection-8

Destination cloud region for this connection: **AWS US East (Ohio)** [Show Port Requirements](#)

1 Deploy the SaaS Connector VM on your vCenter or ESXi host using the OVA URL or installer.

Deployment Platform
VMware [Copy OVA URL](#) | [Download SaaS Connector](#)

2 Use this token to connect to the service when configuring the SaaS Connector.

6fde331b-6f9c-43f4-94d7-a042cc568433-t1-e2 **020 days**

Once you have setup your connection. Click **Verify Connection** to continue.

[Cancel](#) [Skip Verification](#) [Verify Connection](#)

5. To deploy the SaaS Connector OVA in your data center:

1. Log in to your vCenter host.
2. From the **Hosts and Clusters** tab in the vSphere Web Client, right-click on any cluster that can host your VM and select **Deploy OVF Template**. The **Deploy OVF Template** wizard opens.

3. On the **Select an OVF template** page, do one of the following and click **Next**:
 - Paste the link of the OVA file you copied in Step 3 (a) in the **URL** field.
 - Select **Local file**, click **UPLOAD FILES**, and browse to the location of the OVA file you downloaded in step 3 (b).
4. On the **Select a name and folder** page, enter the following and click **Next**:
 1. In the **Virtual machine name** field, enter a unique name for your SaaS Connector.
 2. In the **Select a location for the Virtual Machine** field, select where your VM should reside from the displayed list of inventory locations.
5. On the **Compute Resources** page, select a compute resource for the SaaS Connector VM and click **Next**.
6. On the **Review details** page, verify the SaaS Connector information and click **Next**.
7. On the **Configuration** page, verify **SAAS-CONNECTOR** is selected and click **Next**.
8. On the **Select storage** page, select a datastore with at least 181 GB of free disk space and click **Next**.

9. On the **Select networks** page, select a destination network and click **Next**. You can select VLANs from both the **DataNetwork** and the **SecondaryNetwork** fields. The Data Network is used for communication with Cohesity SaaS, and the Secondary Network is used for communication with your data sources. Based on your requirements:
 - To deploy the SaaS Connector on a single network, select the same VLAN in both **DataNetwork** and **SecondaryNetwork**.
 - To deploy the SaaS Connector on a dual network, select different VLANs in **DataNetwork** and **SecondaryNetwork**, respectively.

Note:

- The SaaS Connector must have dual IP addresses if your data sources are in a private non-routable VLAN.
- Once you have deployed the SaaS Connector on a single network, you cannot modify the SaaS Connector to use a dual network or vice versa.

10. On the **Customize template** page, enter the network settings: **Network IP Address**, **Network Netmask**, and **Default Gateway**. If you have selected a different VLAN for the secondary network, enter the **Network IP Address**, **Network Netmask**, and **Default Gateway** for the secondary network, as well. Click **Next**.

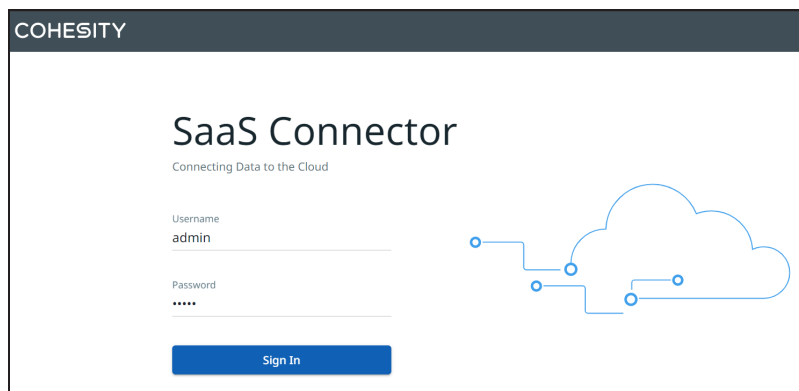
Note:

- To set the network settings using static IP addresses, manually enter the details in the respective fields for both **DataNetwork** and **SecondaryNetwork**.
- To set the network settings using DHCP, leave the fields blank in both the **DataNetwork** and **SecondaryNetwork** sections.
- **Data Network** and **Secondary Network** must be configured using the same network configuration method. That is static IP addresses or DHCP.

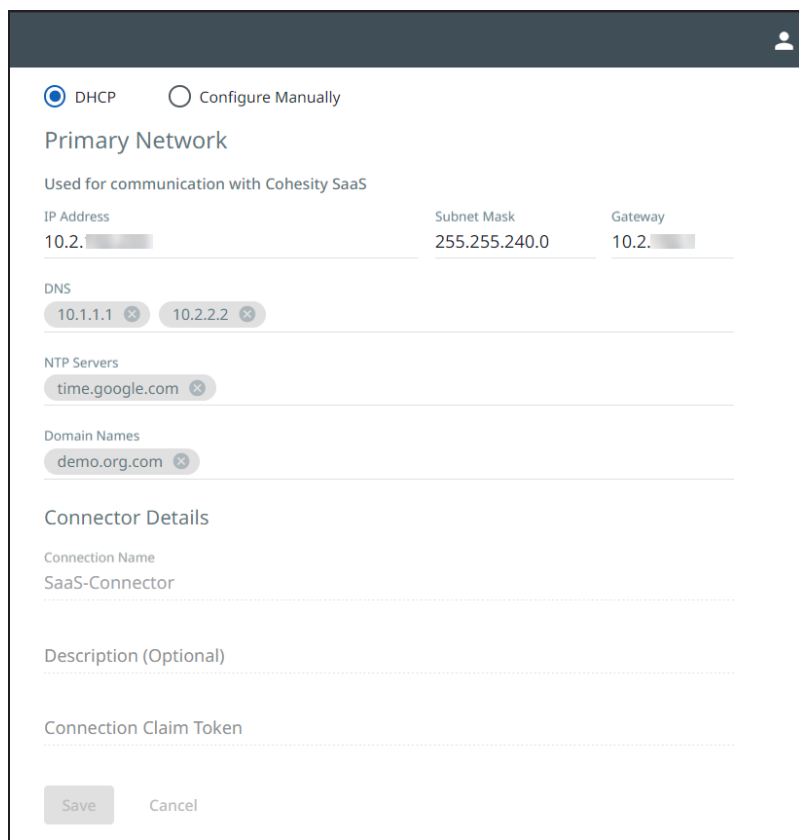
11. Review the summary on the **Ready to complete** page and click **Finish**.
12. Once the VM is created, power it on.

After it boots, the services in the SaaS Connector VM (including the UI) can take 4-5 minutes to start.

6. Enter the IP address of the SaaS Connector VM in the address bar of your browser and click **Enter**.
7. On the SaaS Connector's User Interface, enter "admin" in the **Username** and **Password** fields to log into the SaaS Connector.





8. On the next screen, you are prompted to change your password. Change your default password and log in again with your new password.
9. Verify the network configuration settings, make necessary changes, and click **Continue**.



10. In Cohesity DataProtect as a Service, copy the Connection token from the **Create New Connection** dialog.

Create Saas_Connection-99

Destination cloud region for this connection: **AWS US East (Ohio)** [Show Port Requirements](#)

- 1 Deploy the SaaS Connector VM on your vCenter or ESXi host using the OVA URL or installer.
Deployment Platform: **VMware** [Copy OVA URL](#) | [Download SaaS Connector](#)
- 2 Use this token to connect to the service when configuring the SaaS Connector.
db8a0b58-6d23-43e3-b6ef-e62825e2fe9a-13-e2 19 days  

Once you have setup your connection. Click Verify Connection to continue.

[Cancel](#) [Skip Verification](#) [Verify Connection](#)

11. On the **SaaS Connector** page, paste the Connection token in the **Connection Claim Token** field and click **Save**.

It can take another few minutes for the SaaS Connector to authenticate to the Cohesity DataProtect as a Service.

Network is configured [Edit Network Settings](#)

Connector Details

Connection Name
SaaS-Connector

Description (Optional)

Connection Claim Token
84c61...

[Save](#) [Cancel](#)

12. Once the SaaS Connector authenticates successfully, return to the **Create New Connection** dialog and click **Verify Connection**.

Connector Groups

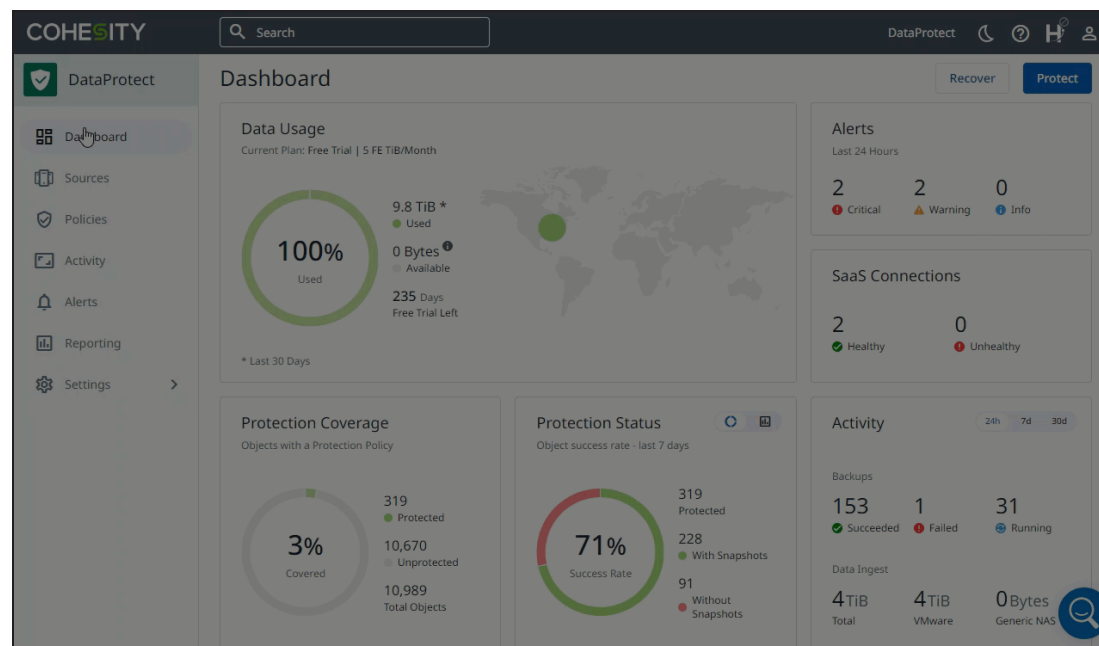
If you use one vCenter to manage multiple ESXi clusters in different geographic locations, you can group the local SaaS Connectors at each location into SaaS Connector Groups, and then associate these Connector Groups to vCenter resources in that location. SaaS

Connector Groups help you ensure efficient routing of your backup and recovery data traffic through SaaS Connectors that operate in the same location as the data sources.

For example, if your vCenter manages two data centers, one in New York and another in San Francisco, the SaaS Connectors in both locations can be grouped into two separate Connector Groups. These Connector Groups can then be associated with the respective data centers in each of those regions.

Note: You can only group SaaS Connectors in a [VMware SaaS Connection](#).

Create Connector Groups



To create and manage Connector Groups in a SaaS Connection:

1. In **DataProtect as a Service**, navigate to **Settings > SaaS Connections**.
2. From the **Actions** menu (:) of a SaaS Connection, select **Manage Connectors**.
3. Select **Group Connectors**. The **Group Connectors** page displays all SaaS Connectors under **Ungrouped**.
4. Click **New Group**.
5. In the **Edit SaaS Connector Group** dialog, enter a name for the new Connector Group and click **Update**. The new Connector Group is displayed on the **Manage Connectors** page.
6. Drag a SaaS Connector from **Ungrouped** to the new Connector Group. You can add more SaaS Connectors until you have all of the Connectors you want in that Connector Group.
7. Click **Done** to save your changes.

To create more Connector Groups, repeat steps 3 to 7.

Once you've created a Connector Group, you're ready to associate the Connector Group to the data center or ESXi host of a specific geographic region. For details, see [Manage Network Traffic](#).

Manage Connector Groups

To view the details of an existing Connector Group:

1. In **DataProtect as a Service**, navigate to **Settings > SaaS Connections**.
2. Click the name of a SaaS Connection you need to explore. All the Connector Groups in that SaaS Connection are displayed under **Connector Details**.
3. To rename a Connector Group, click the **Actions** menu (:) and select **Rename**.
4. To adjust a Connector Group's bandwidth usage, click the **Actions** menu (:) and select **Bandwidth Usage**.


Note: Ungrouped SaaS Connectors inherit the bandwidth settings of the SaaS Connection.

5. To remove a SaaS Connector from the Connector Group, click the **Actions** menu (:) and select **Ungroup**.

Edit SaaS Connectors

A SaaS Connector can belong to only one Connector Group. To move a SaaS Connector to another Connector Group:

1. In **DataProtect as a Service**, navigate to **Settings > SaaS Connections**.
2. From the **Actions** menu (:) of a SaaS Connection, select **Manage Connectors**.
3. To move a single SaaS Connector to another Connector Group, perform one of the following:

1. Drag the SaaS Connector to the other Connector Group.
2. Click the **Move** () icon and select a Connector Group from the list.

The SaaS Connector is moved to the selected Connector Group.

4. To move all the SaaS Connectors from one Connector Group to another, click the **Actions** menu (:) and select a Connector Group from the list.

The SaaS Connectors are moved to the selected Connector Group.

5. Click **Done** to save your changes.

Ungroup SaaS Connectors

To remove all the SaaS Connectors from a Connector Group:

1. In **DataProtect as a Service**, navigate to **Settings > SaaS Connections**.
2. From the **Actions** menu (:) of a SaaS Connection, select **Manage Connectors**.
3. From the **Actions** menu (:) of that Connector Group, select **Ungroup**.

The removed SaaS Connectors are displayed under **Ungroup**.

4. Click **Done** to save your changes.

Delete Connector Groups

You can only delete Connector Groups that do not have any SaaS Connectors and are not associated with a vCenter source.

To delete a Connector Group:

1. In **DataProtect as a Service**, navigate to **Settings > SaaS Connections**.
2. From the **Actions** menu (:) of a VMware SaaS Connector, select **Manage Connectors**.
3. From the **Actions** menu (:) of the Connector Group, select **Delete**.
4. Click **Done** to save your changes.

Manage Network Traffic

After you [create Connector Groups](#), you can associate them with specific data centers or ESXi clusters.

To associate a Connector Group with the desired vCenter resources:

1. In **DataProtect as a Service**, navigate to **Settings > SaaS Connections**.
2. Select a VMware SaaS Connection.
3. Click the **Linked Sources** tab.
4. Click **Get Started**.
5. Select **Manage Network Traffic** from the **Actions** menu (:) of a VMware source.

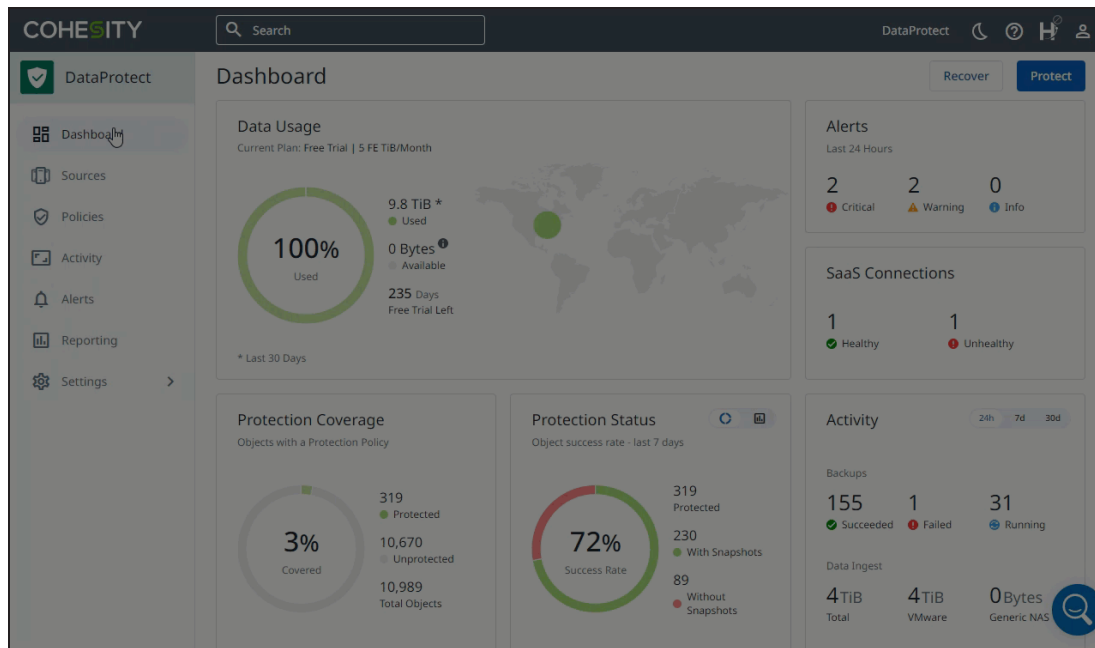
The **Manage Traffic Sources** page displays the Connector Groups you have created and the hierarchy of the data centers, clusters, ESXi hosts, and folders in that vCenter.

6. Drag the data center or ESXi host of a specific geographic region to the Connector Group. The Cohesity cluster chooses the Connector Group associated with the vCenter source closest to the VM in the vCenter hierarchy.

For example, if CG1 is connected to the ESXi host A and CG2 is connected to a folder inside ESXi host A, Cohesity will choose CG2 to protect a VM inside that folder. For this reason, Cohesity recommends that you associate Connector Groups with data centers or ESXi hosts.

Note: Protection runs fail if Connector Groups with no SaaS Connectors are added to the SaaS Connection.

- Enter your credentials in the **Username** and **Password** fields and click **Save**. The Connector Group is displayed on the **Manage Traffic Routes** page.



Once a Connector Group is successfully associated with a vCenter source, network traffic for future **VMware VM protection runs** (in-progress protection runs, if any, are not affected) is steered through the Connector Group to the SaaS Connectors, thereby containing the traffic within a data center or geographical area.

Deploy Hyper-V SaaS Connectors

You can deploy a Hyper-V SaaS Connector using a template VHD in your Hyper-V data center environment that has access to your data sources and meets the [SaaS Connection system and firewall requirements](#).

Once deployed, each SaaS Connector is a virtual machine that runs on a Hyper-V host in your data center.

Tip: For better performance and redundancy, we recommend that you deploy at least two SaaS Connectors for each SaaS Connection in your data center. To add (or remove) a SaaS Connector, see [Manage User-Deployed SaaS Connections](#).

Note: Deploying Hyper-V SaaS Connectors as Generation 2 VMs is not supported.

All the data that a SaaS Connection handles, from your sources to the cloud storage where your backups reside, is encrypted in flight and at rest.

Create Hyper-V SaaS Connection

To create a Hyper-V SaaS Connection:

1. Navigate to **Sources** and click **Register Source**.
2. Select any workload type.
3. In the form, click **Create New Connection**.
4. Under **Deployment Platform** and select **Hyper-V**, then select a **Connection Region** for your data backups.
5. Prepare to deploy the SaaS Connector in your data center:
 - **Copy** the VHD URL.
 - **Copy** or **Download** the **Connection Token**.
6. To deploy the SaaS Connector VHD in your data center:
 1. Log in to your SCVMM server or Standalone Hyper-V host.
 2. Download the VHD file to the SCVMM server or Standalone Hyper-V host using the VHD URL.
 3. From the Hyper-V Manager, open the **New (Create) Virtual Machine** wizard. (For detailed instructions, see [Create a virtual machine in Hyper-V](#) in the Microsoft documentation.)
 1. **Configure** the name, location, generation, hardware for the VM.
 2. **Select Virtual Hard Disk**. Select 'Use existing virtual hard disk' and choose the downloaded VHD file.
 3. **Configure Networking**. Select an operational Virtual Switch to connect the VM to.
 4. **Review**. Review the configuration from the Summary section and click Finish.

Note: After it boots, the services in the SaaS Connector VM (including the UI) can take 4-5 minutes to start.

7. Browse to the SaaS Connector IP address that is assigned to the SaaS Connector VM. On initial login, change the default password and log in again with your new password. Enter the **Connection Token** and common configuration settings and click

Save.

Note: It can take another few minutes for the SaaS Connector to authenticate to the Cohesity DataProtect as a Service.

8. Once the SaaS Connector authenticates successfully, return to the **Create New Connection** dialog and click **Verify Connection**.

Manage User-Deployed SaaS Connections

To optimize performance, we recommend that you use at least two SaaS Connectors in each SaaS Connection you create, and that you have one SaaS Connector for each 160 VMs or 16 TB of source data. (If you have more VMs, we recommend that you stagger their first full backups.)

You can also manage the network bandwidth consumption of your backup and recovery tasks in your SaaS Connections.

Cohesity clusters can concurrently back up individual VM disks using multiple SaaS connectors, leading to faster backups. Cohesity clusters determine the concurrency based on various factors, such as disk size and the workload of the SaaS connectors. Cohesity clusters cannot recover individual VM disks using multiple SaaS Connectors because VMware and HyperV do not support it.

Add SaaS Connector

To add a SaaS Connector to an existing user-deployed SaaS Connection:

1. In **DataProtect as a Service**, navigate to **Settings > SaaS Connections**.
2. Click the Actions menu (:) next to the SaaS Connection and select **Download Installer** to save the OVA to your data center.
3. To deploy the OVA or VHD, follow the instructions in **Step 6** in the respective [SaaS connector topics](#) for VMware and HyperV.
4. Back in the **SaaS Connections** page, click the Actions menu (:) next to the SaaS Connection again and select **Connection Token**. In the dialogue, click the **Copy to Clipboard** button.
5. Browse to the SaaS Connector IP and log in as admin/admin. On initial login, change the default password and log in again with your new password. Enter the **Connection Token** and common configuration settings and click **Save**.
6. Once the SaaS Connector authenticates successfully to the Cohesity DataProtect as a Service, click the **Expand (v)** button next to the SaaS Connection to confirm that the new SaaS Connector is listed.

To add more SaaS Connectors to the same SaaS Connection, repeat the steps above.

Remove SaaS Connector

To remove a SaaS Connector from one of your SaaS Connections:

1. In **DataProtect as a Service**, navigate to **Sources** and click into a source that uses the SaaS Connection.
2. Click the **Connection** tab.
3. Under **Connection Details**, click the Actions menu (:) next to the SaaS Connector and select **Remove from Connection**.

The SaaS Connector is removed from the SaaS Connection. If other healthy Connectors remain in the SaaS Connection, it will continue to function over those Connectors.

Modify SaaS Connector Network Settings

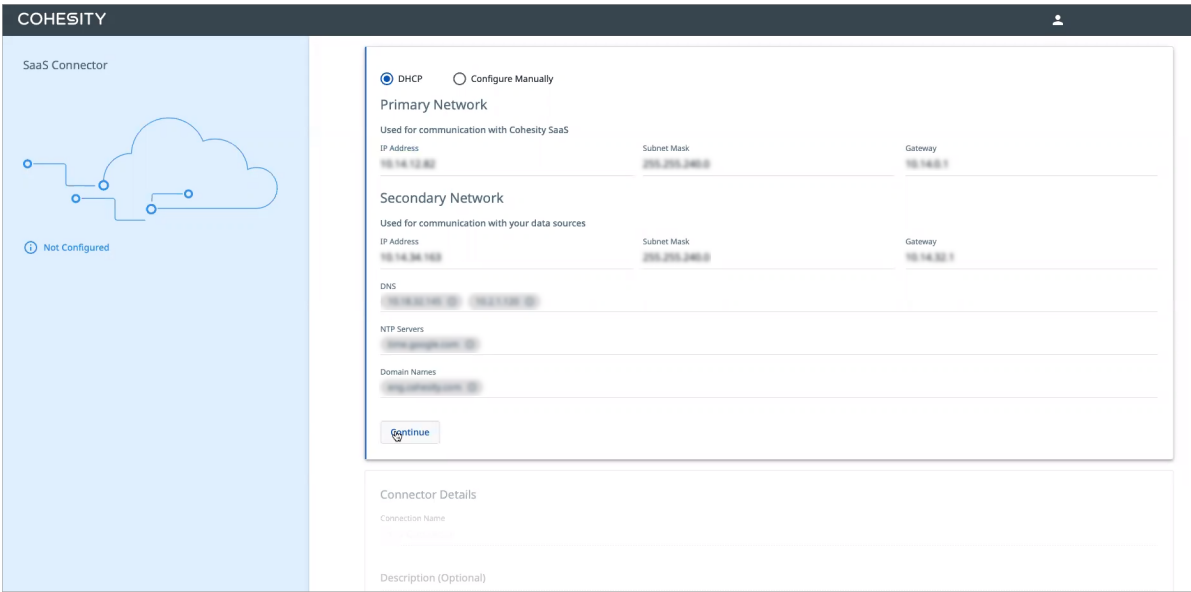
After you have deployed your SaaS Connector, you can use the SaaS Connector UI to:

- Modify the SaaS Connector's existing network configuration.
- Change the network settings of the SaaS Connector from manual (static IP addresses) to DHCP, or from DHCP to manual configuration.

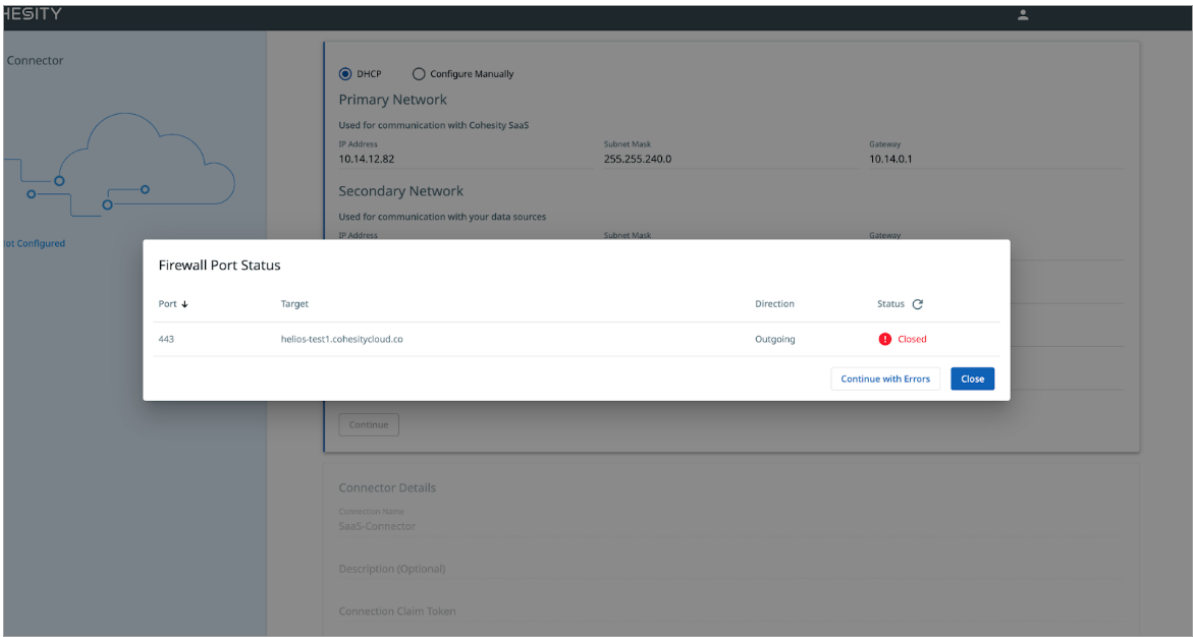
To modify the SaaS Connector network settings:

1. Using a browser, type the IP address of the SaaS Connector VM and log in to the SaaS Connector UI.
2. On the SaaS Connector UI:
 1. Select **DHCP** to change the network configuration from static IP addresses to DHCP, then click **Continue**.
 2. Select **Configure Manually** to change the network configuration from DHCP to

static IP addresses, then click **Continue**.



- 3. The SaaS Connector verifies the network connection to the Cohesity SaaS service using the primary network. Click **Close** if the verification is successful or if errors are found during the verification process. You can also click **Continue with Errors** to ignore the errors and continue using the SaaS Connector.



Cohesity-Deployed SaaS Connectors

Cohesity-deployed SaaS Connectors are the Connectors that Cohesity deploys on the source you want to protect using Cohesity DataProtect as a Service. Cohesity supports the following Cohesity-deployed SaaS Connectors:

- AWS SaaS Connector
- Azure SaaS Connector

Cohesity-Deployed SaaS Connector Requirements

Before deploying the SaaS Connector, review and understand the following requirements:

Supported Sources

Following are the sources on which Cohesity can deploy the SaaS Connectors:

Connectors	Supported Sources
AWS SaaS Connector	<ul style="list-style-type: none">• AWS<ul style="list-style-type: none">• AWS EC2• AWS RDS instance• AWS RDS database (PostgreSQL and Aurora (PostgreSQL Compatible))
Azure SaaS Connector	<ul style="list-style-type: none">• Azure<ul style="list-style-type: none">• Azure VM• Azure SQL

Check Firewall Ports

Ensure that the ports listed in the SaaS Connector Management section in the [Firewall Ports for Cohesity-Deployed SaaS Connectors](#) topic are open in your firewall to allow communication between the Cohesity SaaS Connector(s) and Cohesity Cloud Services.

These firewall rules allow outgoing traffic from a SaaS Connector to the Cohesity DataProtect as a Service endpoint. The SaaS Connector opens a secure encrypted gRPC tunnel to the endpoint and uses it for both backup and recovery traffic.

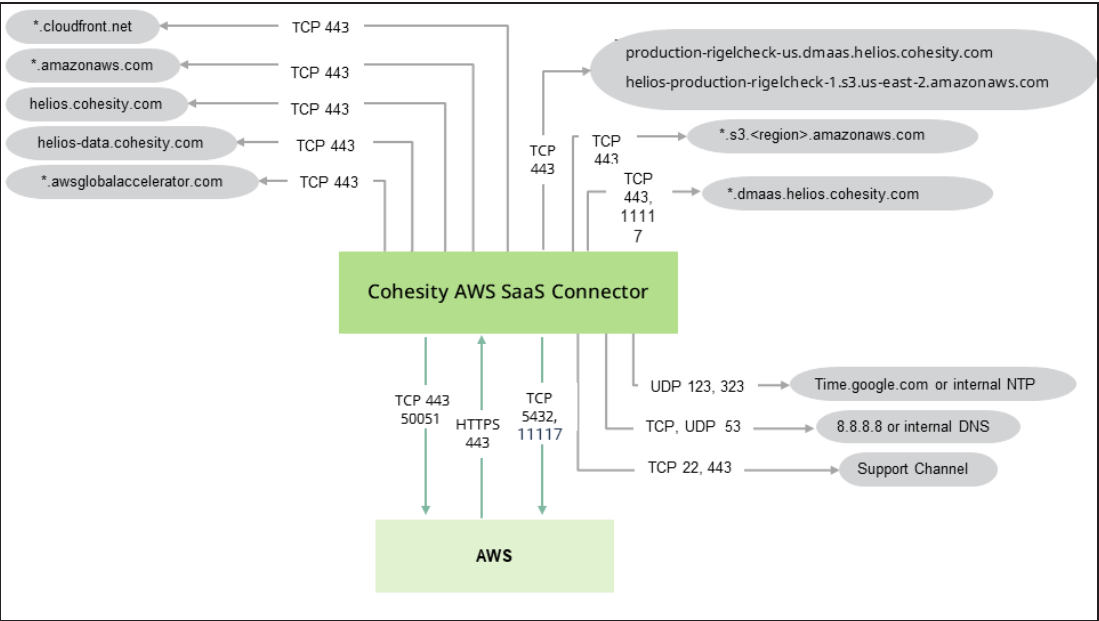
The connectivity status between a SaaS Connection and the Cohesity DataProtect as a Service is displayed both in the SaaS Connection and the Cohesity DataProtect as a Service dashboard.

Firewall Ports for Cohesity-Deployed SaaS Connectors

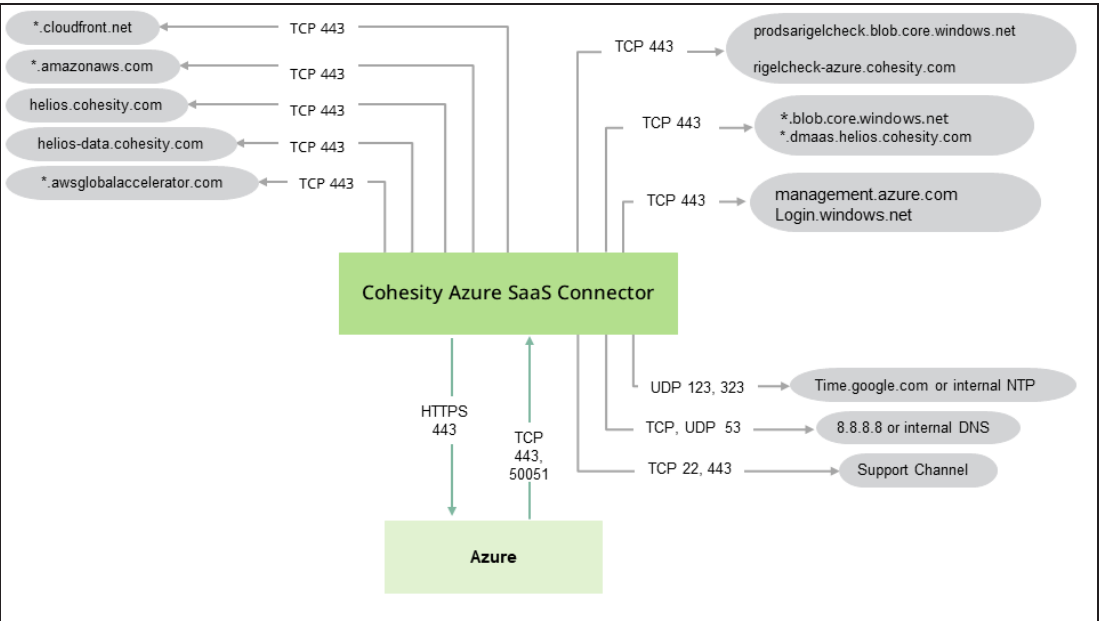
A typical SaaS Connector connects with the Cohesity DataProtect as a Service and the Data Sources. The following diagrams show the source, destination, ports, and protocols for traffic flow between the Cohesity-deployed SaaS Connector and the Data Sources, and the Cohesity-deployed SaaS Connector and Cohesity DataProtect as a Service.

More information is provided in the sections that follow the diagrams.

Firewall Port Requirements for AWS SaaS Connector



Firewall Port Requirements for Azure SaaS Connector



Legend

- Traffic flow between the SaaS Connector and Data Sources
- Traffic flow between the SaaS Connector and Cohesity Data Protect Service

SaaS Connector Management

Ensure that the following ports are open to allow communication between the Cohesity-deployed SaaS Connector(s) and Cohesity DataProtect as a Service:

Source	Destination	Port	Protocol	Purpose
SaaS Connector	helios.cohesity.com	443	TCP	Connection used for control path.
SaaS Connector	*.awsglobalaccelerator.com	443	TCP	Connection used for control path.
SaaS Connector	helios-data.cohesity.com	443	TCP	Used to send telemetry data.
SaaS Connector	*.cloudfront.net	443	TCP	To download upgrade packages.
SaaS Connector	8.8.8.8 or internal DNS	53	TCP, UDP	Host resolution.
SaaS Connector	time.google.com or internal NTP	123, 323	UDP	<p>Incoming NTP requests are detected by port 123.</p> <p>Chrony is the default implementation of NTP used by recent versions of CentOS and RHEL. Open port 323 if you want to use the Chronyc tool to monitor the synchronization status of Chrony and make changes if necessary.</p>

Source	Destination	Port	Protocol	Purpose
SaaS Connector	rt.cohesity.com	22 or 443	TCP	The Cohesity Support Channel uses Secure Shell (SSH) and listens through port 22 or 443. Port 22 is used by default and can be updated to 443 using the Cohesity CLI. For more information, see Manage the Support Channel .
AWS SaaS Connector Specific				
SaaS Connector	production-rielcheck-us.dmaas.helios.cohesity.com helios-production-rigelcheck-1.s3.us-east-2.amazonaws.com	443	TCP	Required to perform connectivity checks with the Cohesity Cloud Services.
SaaS Connector	*.s3.<region>.amazonaws.com *.dmaas.helios.cohesity.com	443	TCP	Connection used for data path.
SaaS Connector	*.dmaas.helios.cohesity.com	11117	TCP	Connection used for data path.
Azure SaaS Connector Specific				
SaaS Connector	prodsarigelcheck.blob.core.windows.net rigelcheck-azure.cohesity.com	443	TCP	Precheck endpoints for connectivity.
SaaS Connector	*.blob.core.windows.net *.dmaas.helios.cohesity.com	443	TCP	Connection used for data path.
SaaS Connector	management.azure.com Login.windows.net	443	TCP	

AWS

Ensure that the following ports are open to allow communication between the Cohesity SaaS Connector(s) and AWS account:

Source	Destination	Port	Protocol	Purpose
SaaS Connector	AWS EC2 and RDS Ingest	443	TCP	Required for Backup and Recovery operations.
SaaS Connector	AWS RDS instance	5432 or User-configured-Postgres port	TCP	Required for communication with Postgres server.
SaaS Connector	AWS RDS instance	11117	TCP	Required for Postgres backup and recovery.
SaaS Connector	AWS EC2	50051	TCP	Required for EC2 file-level recovery.
AWS EC2 and RDS Ingest	SaaS Connector	443	TCP	Required for Backup and Recovery operations.

Azure

Ensure that the following ports are open to allow communication between the Cohesity SaaS Connector(s) and Azure Source:

Source	Destination	Port	Protocol	Purpose
SaaS Connector	Azure VM and SQL	443	TCP	Required for Backup and Recovery operations.
Azure VM and SQL	SaaS Connector	443	TCP	
Azure VM and SQL	SaaS Connector	1443	TCP	

Source	Destination	Port	Protocol	Purpose
SaaS Connector	Azure VM	50051	TCP	Required for Azure VM file-level recovery.

Deploy AWS SaaS Connectors

If you want Cohesity DataProtect as a Service to protect your AWS EC2 instances using [Cohesity Snapshots](#), you need to set up a SaaS Connection for each AWS region where you have EC2 instances to protect. Each SaaS Connector is an m5.xlarge AWS EC2 instance.

Note: To prepare your AWS account for Cohesity SaaS Connector deployment in a Public or Private subnet, see [AWS SaaS Connector Deployment Guide](#).

Create AWS SaaS Connection

To create an AWS SaaS Connection:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the **Actions** menu (:) next to the AWS account and select **Setup SaaS Connection**.
3. In the **Configure SaaS Connection for AWS** dialog, provide:
 1. **Region.** Select the AWS region where you have EC2 instances to protect.
 2. **Number of Connectors.** Enter the number of SaaS Connectors you want to deploy in the region.

Tip: For better performance and redundancy, we recommend that you deploy at least two SaaS Connectors for each SaaS Connection. To add (or remove) a SaaS Connector, see [Manage User-Deployed SaaS Connections](#).

3. **Subnet.** Select the subnet where you want the SaaS Connectors to be launched. Using a secured public subnet is more cost-efficient than a private subnet.
4. **Network Security Groups.** Select the network security group to be associated with SaaS Connectors. You can provide multiple network security groups to be associated with SaaS Connectors. Make sure the network security group follows the firewall rules.

5. **DNS:** By default, 8.8.8.8 is used as the Domain Name System (DNS) server. You can optionally, remove this default value and specify the IP addresses of the DNS servers that the Cohesity DataProtect as a Service should use. Separate multiple IPs with commas. Ensure the Active Directory DNS IP address (if applicable) is listed first. Verify that the NTP servers and other entities in the system can be resolved by the specified DNS server.
6. **NTP Servers:** By default, time.google.com is used as the NTP server. You can optionally remove this default value and specify the IP address or the Fully Qualified Domain Name of the NTP server(s) that must be used to synchronize the time on the Cohesity DataProtect as a Service.
7. **Tags.** Specify the tags to be used for your SaaS Connectors. (Optional)
8. To create a SaaS connection for each region in your AWS account, click **Add another SaaS Connection** and provide the above details.

Manage SaaS Connection for AWS Account: 226926521492

Region AWS US East (Ohio)	Number of Connectors 3
Subnet subnet-0e0378137d21412c4 (DMAas)	Network Security Groups sg-0114282dc0634bf4e (cohesity_ce_cluster... sg-00c1ff3abf55e3aba (cohesity_control) sg-0a9247a68da6208fc (cohesity_fleet_sg)
DNS 8.8.8.8	NTP Servers time.google.com
Tags team=eng	

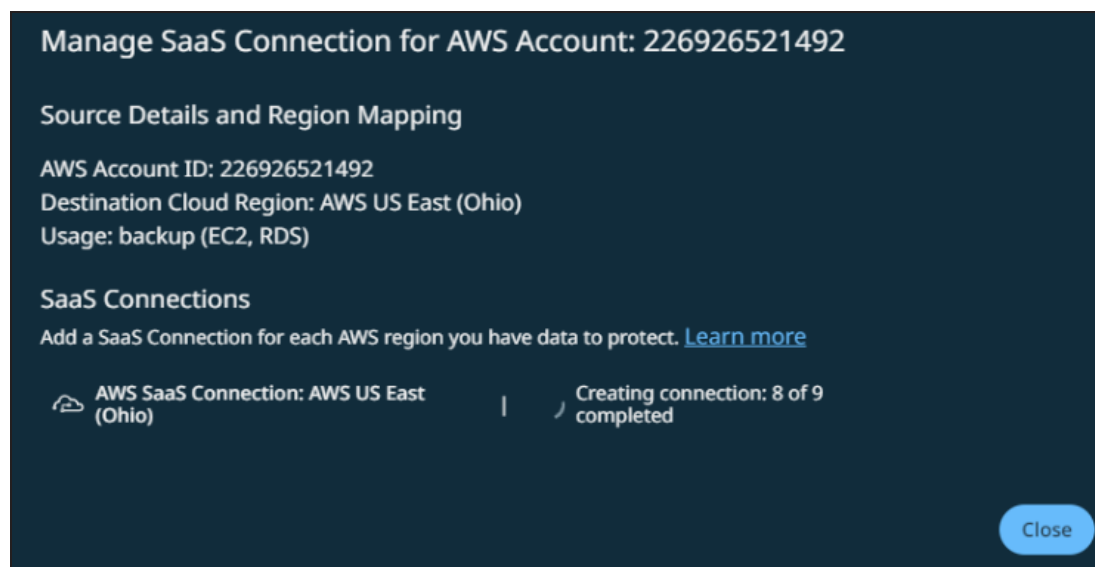
Format: label=value. Comma separated multiple tags.

+ Add another SaaS Connection

Cancel Update

4. Click **Create Connections**.

The progress status of the AWS SaaS Connection will be displayed in the UI:



If the connection fails, then the corresponding error message is displayed.

Next > Your new AWS SaaS Connection is now available to use when you [protect your AWS EC2 instances](#).

Deploy Azure SaaS Connectors

Once you register your Azure source, you must set up a SaaS Connection for each region under each Azure subscription in your Azure source. A SaaS Connection consists of one or more SaaS Connectors, which are VMs that act as data movers between your data sources and the Cohesity DataProtect as a Service. Each Azure SaaS Connector is an Standard_D8s_v3 instance.

Create Azure SaaS Connection

To create an Azure SaaS Connection:

1. In **DataProtect as a Service**, navigate **Sources**.
2. Click the Actions menu (:) next to the Azure source and select **Setup SaaS Connection**.
3. Click the > icon next to subscription for the subscription you want to set up the SaaS Connection.

Manage SaaS Connections

Source Details

Domain Name: 75818451-2edd-4f92-8f36-47882b1a59b5
 Destination Cloud Region: Azure Central US
 Azure Platforms: Azure VM, Azure SQL

Subscriptions

With SaaS Connection	Without SaaS Connection
1	0

3215d0a4-e99d-4b0c-b106-e2a7708e0024
 2 SaaS Connections

Close

The Azure Connection form appears. In the Azure Connection form, the Azure Source and the Subscription ID are selected by default.

4. Click **Add SaaS Connection**.
5. In the **Connection Details** section, provide the following details:
 - a. **SaaS Connection Name**: Provide a name for the SaaS Connection.
 - b. **Location**: From the drop-down list, select the Azure region where you have the Azure cloud services to protect.
 - c. **Resource Group**: From the drop-down list, select a resource group that will hold the resources related to the SaaS Connection.
 - d. **Number of Connectors**: Enter the number of SaaS Connectors you want to deploy in the region.
6. In the Network Settings section, provide the following details:
 - a. **Network Resource Group**: From the drop-down list, select the resource group for the virtual network.
 - b. **Virtual Network**: From the drop-down list, select a virtual network to which you want to connect the SaaS connections.
 - c. **Subnet**: Select the subnet where you want the SaaS Connectors to be launched
7. In the **Other Settings (Optional fields)** section, provide the following details:
 - a. **Network Security Group**: From the drop-down list, select a security group that will be associated with the specified subnet. You can select multiple network security groups.

b. **Application Security Group:** Select the application security groups you want to attach to the SaaS connector.

c. **Azure Managed Identity:** Enter the managed identity that must be attached to the SaaS Connectors. This setting cannot be edited later. Example:

```
/subscriptions/1234ab56-a2b2-a1b1-a12b-  
abc12345c678d/resourcegroups/example-  
rg/providers/Microsoft.ManagedIdentity/userAssignedIdentities/myM  
anagedIdentity.
```

For Azure SQL, the SaaS Connector's managed identity will be used to authenticate to the SQL server for export/import if the SQL server source's credential setting is set to "Managed Identity".

For more information on managed identity, see [Microsoft Azure documentation](#).

d. **DNS Servers:** Enter the IP addresses of the DNS servers that the SaaS Connectors should use. Separate multiple IPs with commas. Ensure the Active Directory DNS IP address (if applicable) is listed first. Verify that the NTP servers and other entities in the system can be resolved by the specified DNS server.

e. **NTP Servers:** Enter the IP addresses or the Fully Qualified Domain Name of the NTP server(s) that must be used to synchronize the time on the SaaS Connector.

f. **Tags:** Specify the tags to be used for your SaaS Connectors.

8. Click **Save**.

T

Azure Connections

X

Source and Subscription

Source

Subscription

Connection Details

SaaS Connection Name

Location

Default1211

centralus

Resource Group

Number of Connectors

azure-vm-automation

3

Network Settings

Network Resource Group

Virtual Network

Subnet

azure-vm-automation

azure-vm-automation-vnet

azure-vm-automation-subnet

Choose the resource group for the subnet

Choose a subnet to deploy the SaaS connector

Other Settings (Optional fields)

Network Security Group

Application Security Group

Azure Managed Identity

/subscriptions/3215d0a4-e99d-4b0c-b106-e2a7708e0024/resourceGroups/sql-automation-rg-2/providers/Microsoft.ManagedIdentity/userAssignedIdentities/sql-auto-managedidentity

DNS Servers

NTP Servers

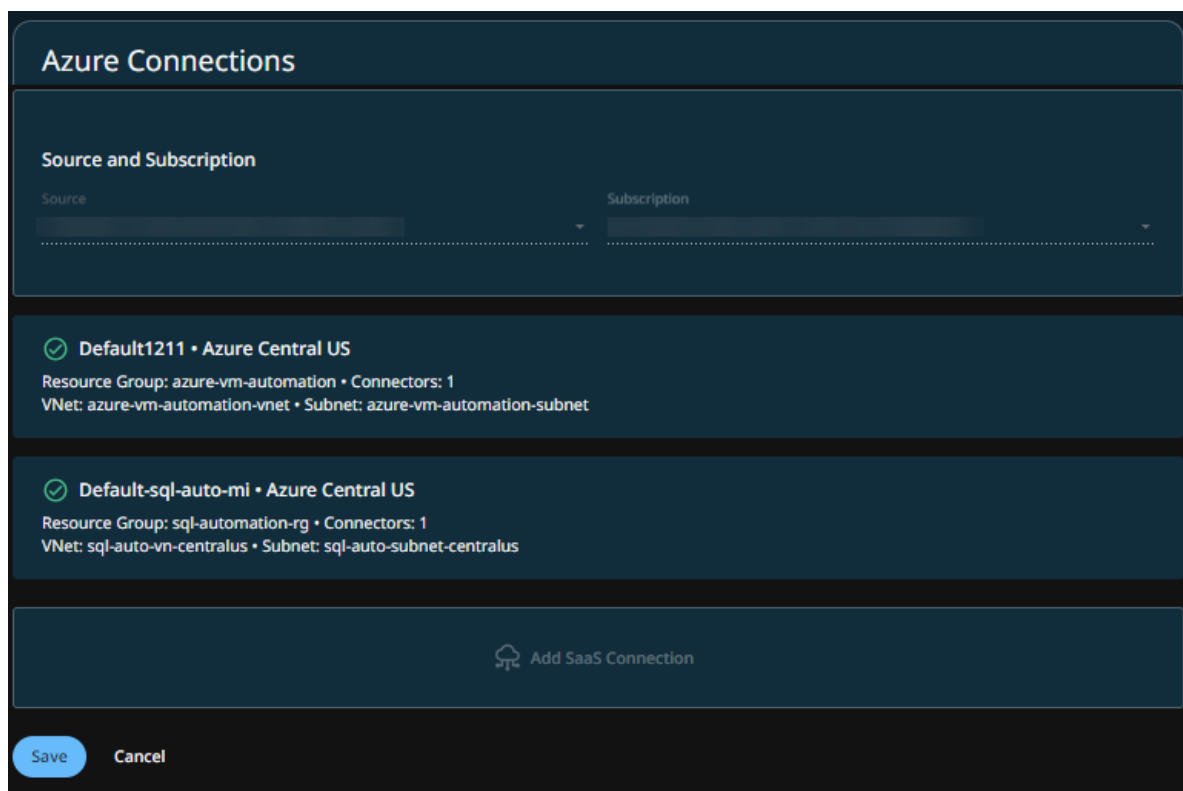
Tags

Format: label=value. Comma separated multiple tags.

Save

Cancel

- To create another SaaS connection for each region in the selected subscription, click **Add SaaS Connection** and provide the above details.



10. Click **Create Connections**.

Repeat the steps above to set up SaaS Connections for each subscription and its regions in your Azure source.

Once you set up SaaS Connection, you can protect the Azure services of your Azure source.

Manage Cohesity-Deployed SaaS Connections

Add AWS SaaS Connector

To add an AWS SaaS Connector to an existing AWS SaaS Connection:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the Actions menu (:) next to the AWS account and select **Setup SaaS Connection**.
3. Click the edit icon next to the SaaS Connection.
4. Update the **Number of Connectors**.
5. Click **Update**.

Add Azure SaaS Connector

To add an Azure SaaS Connector to an existing Azure SaaS Connection:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the Actions menu (:) next to the Azure source and select **Setup SaaS Connection**.
3. Click the edit icon next to the SaaS Connection.
4. Update the **Number of Connectors**.
5. Click **Save**.

Retry AWS SaaS Connection

Note: This workflow is applicable only for AWS SaaS Connection.

If the AWS SaaS Connection fails with an error, you can fix the error and then retry to add the SaaS Connections:

To retry the AWS SaaS Connection:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the Actions menu (:) next to the AWS account and select Setup SaaS Connection.
3. Click the **retry** icon (↺) for the AWS SaaS Connection that failed.
4. In the **Retry SaaS Connection** dialog, you can optionally update the number of SaaS Connectors you want to deploy in the region.
5. Click **Retry**.

Cohesity will retry to create the AWS SaaS Connection.

Remove SaaS Connector

To remove a SaaS Connector from one of your SaaS Connections:

1. In **DataProtect as a Service**, navigate to **Sources** and click into a source that uses the SaaS Connection.
2. Click the **Connection** tab.
3. Under **Connection Details**, click the Actions menu (:) next to the SaaS Connector and select **Remove from Connection**.

The SaaS Connector is removed from the SaaS Connection. If other healthy Connectors remain in the SaaS Connection, it will continue to function over those Connectors.

Retry AWS SaaS Connection Deletion

Note: This workflow is applicable only for AWS SaaS Connection.

If you are not able to delete or remove an AWS SaaS Connection due to an error, then you can fix the error and retry to delete the connection:

To retry the AWS SaaS Connection deletion:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the AWS Source for which you are creating or deploying the SaaS Connection.
3. Click the **Connections** tab.
4. Click the Actions menu (:) next to the AWS SaaS Connection, and then click **Delete Connection**.
5. Type **Yes** to confirm, and then click **Retry Deletion**:

Remove SaaS Connections

You can retry or force remove the connection. If you select the force delete option, the connection will be deleted from UI but the cloud resources will remain and will require manual removal from cloud consoles.

Type 'YES' to confirm

This field is required

CancelForce RemoveRetry Deletion

Cohesity retries to delete the AWS SaaS Connection.

Forcefully Delete AWS SaaS Connection

Note: This workflow is applicable only for AWS SaaS Connection.

If you are not able to delete an AWS SaaS Connection due to an unknown error or if the retry of the AWS SaaS Connection deletion fails, then you can forcefully delete the AWS SaaS Connection:

To forcefully delete the AWS SaaS Connection:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the AWS Source for which you are creating or deploying the SaaS Connection.
3. Click the **Connections** tab.
4. Click the Actions menu (:) next to the AWS SaaS Connection, and then click **Delete Connection**.

5. Type **Yes** to confirm, and then click **Force Remove**:

Remove SaaS Connections

You can retry or force remove the connection. If you select the force delete option, the connection will be deleted from UI but the cloud resources will remain and will require manual removal from cloud consoles.

Type 'YES' to confirm

This field is required

CancelForce RemoveRetry Deletion

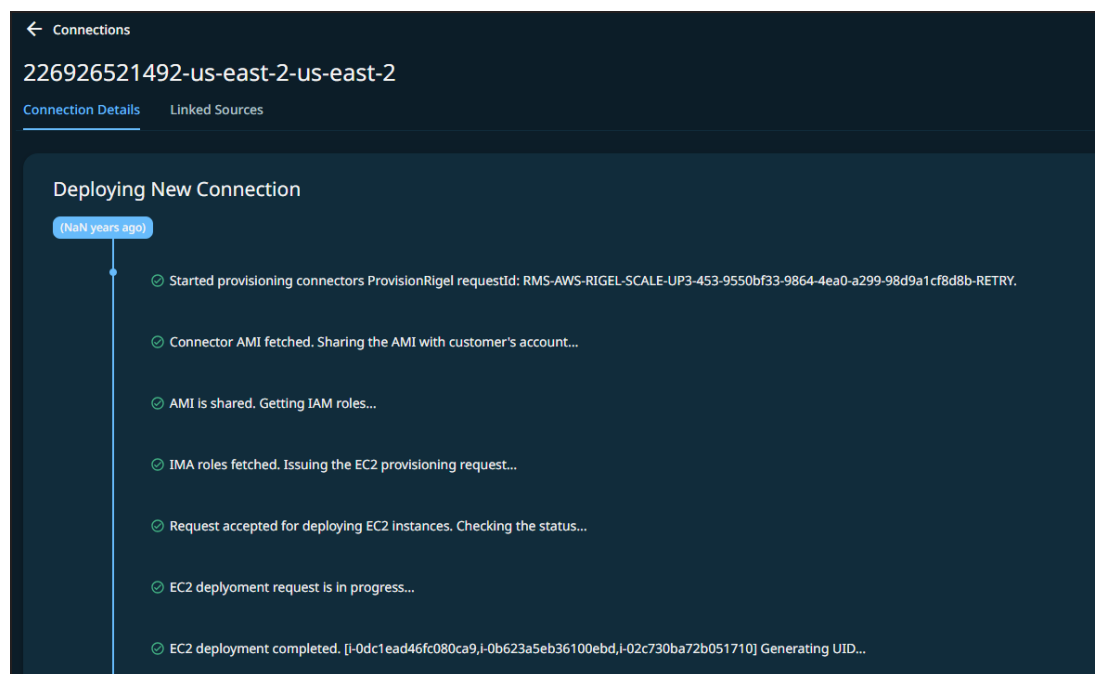
The AWS SaaS Connection will be removed with warning messages.

View AWS SaaS Connection Creation or Deletion Progress

To view the step-by-step progress of the AWS SaaS Connection creation or deletion:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the AWS Source for which you have deployed the SaaS Connection.
3. Click the **Connections** tab.
4. Click **View Details** of the SaaS Connection for which you want to view the status of deployment or deletion.

The **Connection Details** tab displays the step-by-step progress of the AWS SaaS Connection deployment or deletion:



Manage Network Bandwidth Usage

In Cohesity DataProtect as a Service, network bandwidth usage is automatically balanced among the SaaS Connectors within each SaaS Connection. However, if you need to contain the amount of network bandwidth consumed by your backup and recovery tasks at different times and days of the week, the Cohesity DataProtect as a Service allows you to throttle your bandwidth consumption in your SaaS Connections.

The bandwidth usage options in each SaaS Connection allow you to choose the days of the week and set the start and end times to limit bandwidth usage to a specific value in bytes per second.

Important:

- If the defined start time and end time are the same, then the bandwidth limit is applied for the day from 12:00 AM till 11:59 PM.
- If the defined start time is greater than the end time, then the interval is split across days. For example, if **9:00 PM** and **5:00 AM** are set as the start and end times on Monday, then two intervals are set: **9 PM-11:59 PM** on Monday and **12:00 AM-5 PMon Tuesday**.
- When time intervals overlap, each new interval overrides the one above it in the list.
- Bandwidth usage limits are only applicable for backup and volume-level recovery tasks and are not applicable for file-level recovery.

To configure a bandwidth usage limit:

1. In **DataProtect as a Service**, navigate to **Settings > SaaS Connections**.
2. Click the Actions menu (:) next to the SaaS Connection and select **Bandwidth Usage Options**.

or

Under **Sources**, click into a source. In the **Connections** tab, click the Actions menu (:) in the top-right corner of the page and select **SaaS Connection > Bandwidth Usage Options**.

3. In the **Bandwidth Usage Options** dialog:
 1. From the drop-down list, select **Upload** (for backup traffic) or **Download** (for recovery traffic)
 2. Select the **Days** of the week.
 3. Set the **Start Time** and **End Time**.
 4. Specify the bandwidth usage **limit**. Consider the following factors while setting up the bandwidth:
 - The default setting does not limit bandwidth usage.
 - Once you set a bandwidth limit, you cannot revert to the default setting.
 - Setting the parameter to a negative number throws the following error message: "Bandwidth bps cannot be less than zero".
 - Setting the parameter to zero is the opposite of unlimited, as it disables the bandwidth entirely.

- Setting the parameter to a positive number higher than int64 throws the following error message: "Unmarshal type error".

Tip: Click the plus (+) to add multiple schedules.

4. Click **Save**.

Configure SaaS Connector Alert Notifications

Cohesity DataProtect as a Service creates a **Critical** alert, **SaaSConnectorStatusAlert**, when the SaaS connector is not reachable due to a network connection issue or is down. A critical alert signifies that immediate action is required because Cohesity detected a severe problem that might be imminent or major functionality is not working.

You can configure alert email notifications in Cohesity DataProtect as a Service to receive the alerts you need.

Note: The alert, **SaaSConnectorStatusAlert**, is not displayed in the **Alerts** tab on the Alerts Dashboard.

To configure email notification for SaaS Connector alerts:

1. In **DataProtect as a Service**, navigate to **Health**, and select the **Notification** tab.
2. Select **Create > New Alert Notification Rule**.
3. In the **Create Alert Notification Rule** dialog, enter:
 1. **Notification Name**. The name for the notification, for example, SaaS Connection Failure.
 2. **Alert Source**. The source of the Alert.
 3. **Alert Severity**. Select **Critical** from the drop-down.
 4. **Alert Type**. Select **Maintenance** from the drop-down.
 5. **Alert Category**. *Optional*. Select one or more categories from the drop-down. Otherwise, all alerts in any category trigger the notification.
 6. **Alert Name**. *Optional*. Select one or more names from the drop-down. Otherwise, any Alert name will trigger the notification.
7. In the **Create Notifications via** section, select **Email**.
 1. Select **To** and type an email address or distribution list of the recipients who need to receive the email notifications.

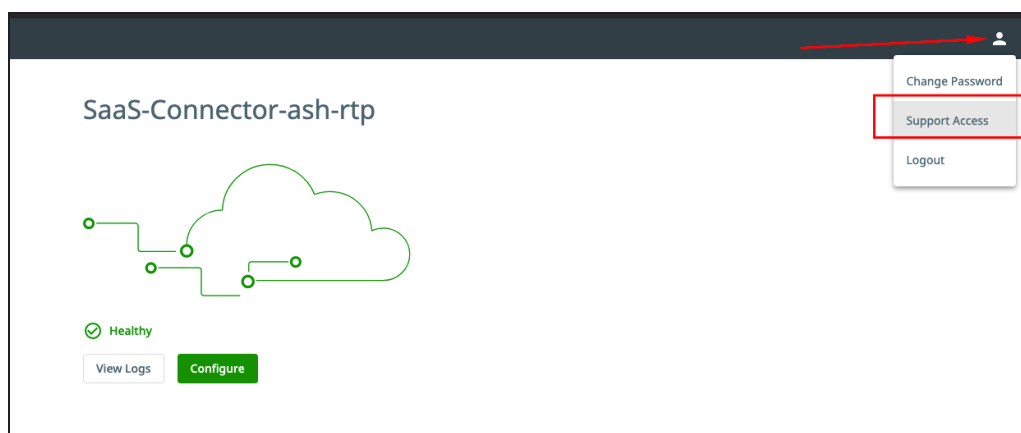
2. Select **CC** and type an email address or distribution list of the recipients who need to be copied on the email notifications.
4. Click **Create**.

Enable SaaS Connector Support Channel Access

Cohesity recommends that you enable the Support Channel for SaaS Connector when you need assistance from Cohesity employees.

To enable the SaaS Connector support channel:

1. Login to the SaaS Connector UI (<https://<saas-connector-ip>/>) using the credentials you have set while creating the SaaS connector.
2. Click the user icon on the top-right corner of the SaaS connector UI, and then click **Support Access**.



3. On the support Access page, enable the **Support Channel** option to enable the support channel access.
4. From the channel active until field, you can select the date till you want the support channel to be active.
5. Copy the **Support Access Token** and **Cluster ID** from the UI and send these details to [Cohesity Support](#).



Cohesity will enable support access for the SaaS Connector.

On-Demand Upgrade of SaaS Connectors

If there are errors during the auto-upgrade process of SaaS Connectors, the upgrade may fail. You can resolve these errors and perform an on-demand upgrade of the SaaS Connectors.

To perform on-demand upgrade of the SaaS Connectors:

1. In DataProtect as a Service, navigate to **Settings > SaaS Connections**.

Cohesity DataProtect as a Service will display a **Needs Upgrade** tag next to the SaaS Connections that require an upgrade.

Name	CPU	Memory	IOPS (Read)	IOPS (Write)	Sources	Connectors	Last Connection Time
226926521492-ap-northeast-1-us-east-2 Region: AWS US East (Ohio)	2%	26%	-	-	1	1	Mar 21, 2024 1:31pm
Geo Stretched E2E Connection Region: AWS US East (Ohio)	2%	40%	-	-	0	3	Mar 21, 2024 1:31pm
geo-stretched-auto-connection Region: AWS US East (Ohio)	8%	39%	-	-	0	2	Mar 21, 2024 1:31pm

2. Click **Needs Upgrade**.

The connectors in the SaaS Connection that require an upgrade are displayed with the **Needs Upgrade** status tag.

The SaaS connectors in this SaaS connection are running an older version. Upgrade them as soon as possible.

1 Connectors need upgrade.


Connectors

172.31.12.253 Needs Upgrade

Upgrade View Details


3. To view the details of the SaaS Connectors, click **View Details**.
4. To upgrade the SaaS Connectors, click **Upgrade**.

The **Upgrade Connectors?** dialog appears indicating that the SaaS Connectors will be rebooted during the upgrade.



5. Click **Yes** to confirm.

The upgrade process will be initiated, and the status tag will transition to **upgrading**.



Note: The upgrade process will not be initiated if you click No.

Once the upgrade is successful, there will be no status tag for the SaaS Connector.

If the upgrade fails, Cohesity DataProtect as a Service will display **Upgrade Failed** as the status tag.

Troubleshoot SaaS Connector Upgrade Failure

If the upgrade you initiated fails, then you can troubleshoot the upgrade failure by performing the following steps:

1. In DataProtect as a Service, navigate to **Settings > SaaS Connections**.
2. Click the **Upgrade Failed** tag displayed next to the SaaS Connection.
3. Click **View Details**.

The Connection Details page appears.

4. Click the **Upgrade Failed** status tag displayed for the SaaS Connectors that failed to upgrade.

The error message corresponding to the upgrade failure is displayed. You can take the corrective actions and then retry the upgrade.

Error Message	Corrective Action
Failed to query package URL.	Ensure TCP port 443 is opened between the SaaS Connectors and *cloudfront.net.
Http error querying the package URL.	
Error downloading software package.	

For all other error messages related to upgrade failures, contact your Cohesity Account Team.

Access Management

On logging into Cohesity DataProtect as a Service, the admin can add other users, define roles, specify cluster access, and generate API keys to access . To manage users, roles, and define their access, in the Helios dashboard, navigate to **Settings > Access Management**.

Manage Users & Groups

To manage user access to your Cohesity DataProtect as a Service, we recommend that you add users and groups. Once you create them, your users can start using your Cohesity DataProtect as a Service with their own logins.

Add Users

To add a user:

1. In **DataProtect as a Service**, navigate to **Settings > Access Management** and click the **Users** tab.
2. Click **Add User**.

Note: Only the user with Admin privileges will be able to add a new user.

3. In the dialog, select **Add User** and enter:
 - **Username.** The user's email address.
 - **Email Address.** The user's email address again.
 - **First Name.** The user's first name in Cohesity DataProtect as a Service.
 - **Last Name.** Typically, the domain of your email address.
4. Under **Roles and Access**, assign an appropriate **Role** to this user. See [Roles](#) for more information.
5. Click **Save**.

The new user receives a welcome email with a link to reset their password, and appears in the list on the **Users** tab. From there, you can edit or delete the user, or prompt them to reset their password.

Roles

Roles	Description
Cohesity Support Admin	This role allows Cohesity Support to create a Super Admin user for the customer. Only Cohesity Support has access to this role, and it is typically used when the customer has lost access to a Super Admin user due to turnover and other events.
Data Security	Data Security users have Self Service Data Protection role privileges and can create DataLock Views and set DataLock expiration dates.
High Classified	User who has High classified role can fetch cluster details needed for specific API calls.
Operator	Operator users have Viewer role privileges and can run existing Protection Groups and create Recover Tasks.
SMB Backup Operator	SMB Backup Operators have privilege to perform SMB backup and SMB restore.
Super Admin	Super Admin users have full access to all actions and workflows within the Cohesity Dashboard. They can manage other Super admins and admins.
Viewer	Viewer users have read-only access for all workflows within the Cohesity Dashboard.

Manage Users

To change a user's settings, click the Actions menu (⋮) next to the user and select:

- **Edit.** To update their Email Address, First Name, and/or Last Name.
- **Delete.** To delete the user from your Cohesity DataProtect as a Service.
- **Reset Password.** To send the user an email with a link to reset their password.

Change Password

To change your Cohesity DataProtect as a Service password:

1. In **DataProtect as a Service**, navigate to **Settings > Access Management** and click the user to open the User Details page.
2. Click **Reset Password** and follow the prompts.

Add SSO Users & Groups

If you have added Single Sign-on (SSO) to Cohesity DataProtect as a Service, you can add users and groups from your SSO domain for additional user management.

To add SSO users and groups:

1. In **DataProtect as a Service**, navigate to **Settings > Access Management**.
2. Click **Add User** on the **Users** tab.
3. In the dialog, select **Add SSO Users & Groups** and enter:
 - **SSO Domain**. The domain you used to add SSO.
 - **SSO Users**. The users in your SSO domain who need access to Cohesity DataProtect as a Service.
 - **SSO Groups**. The groups in your SSO domain who need access Cohesity DataProtect as a Service.
4. Click **Save**.

The new SSO users and groups you entered appear in the list on the **Users** tab. To group them, click the **Domain** column sort them by your SSO domain.

Click the **Actions** menu (:) next to the SSO user or group to **Edit** or **Delete** them.

Access Scope

Access Scope enables users to manage access to the resources in the Cohesity DataProtect as a Service, restricting their capabilities to only those resources. It helps the Managed Service Providers (MSPs) who manage multiple tenants and need to enforce strict access controls for different users.

Access Levels within the Access Scope

- **Source Level Access**—Restricts user access to specific resources within the system. It focuses on individual resources.
- **Source Type Level Access**—Grants access to all resources within a specific source type (e.g. VMware and O365 resources). It enables access to all resources of a particular source type.
- **Region Level Access**—Defines access based on geographic regions. It enables access to resources within a specific region.
- **Service Level Access**—Broadens access to all resources within a specific cloud service. It enables access across all resources of a particular cloud service.

Add Access Scope

1. In **DataProtect as a Service**, navigate to **Settings > Access Management** and click the **Access Scope** tab.
2. In the **Access Scope** tab, click **Add Access Scope**.

3. Enter a name under the **Name** Column.
4. Enter a description for the Access Scope under the **Description** column.
5. Click the **Edit** icon.
6. Select the resources that require Access Scope.
7. Enable the **Auto Assign** option if required.

Note: The Auto Assign option ensures users access newly added resources without additional configuration.

8. Click **Continue**.
9. Click **Add** to add the Access Scope.

Assigning Access Scope to Users

Important: Only users with the Helios Access Scope Management privilege can create or manage access scopes.

For New Users,

1. In **DataProtect as a Service**, navigate to **Settings > Access Management** and click the **Users** tab.
2. Click **Add User**.
3. Turn on the **Enable Access Scope** option.
4. Select an Access Scope under the **Name** drop-down and click **Save**. For more information on adding users, see [Manage Users & Groups](#).

Note: When a user has multiple access scopes, their permissions are combined, giving them access to all resources covered by those scopes.

For Existing Users,

1. In **DataProtect as a Service**, navigate to **Settings > Access Management** and click the **Users** tab.
2. Click the **Actions** menu (:) and then click **Edit**.
3. Turn on the **Enable Access Scope** option.

Note: Access Scopes cannot be deleted if there are active members (users) associated with them. This prevents accidental loss of access controls for users.

Difference Between Roles and Access Scopes

Roles and access scopes both serve to manage user access within the system, they operate differently and can be used together for more granular control.

Roles	Access Scope
Roles define what actions a user can perform across the entire system, such as viewing data, running tasks, or managing users. For example, a "Super Admin" role grants full access to all actions and workflows, while a "Viewer" role limits a user to read-only access.	Access Scopes restrict access to specific resources or objects, such as particular sources, adapters, regions, or services. This is particularly useful for Managed Service Providers (MSPs) who need to enforce strict access controls across multiple tenants.

Add a Single Sign-on Provider

You can now configure Cohesity DataProtect as a Service to use an Identity Provider (IdP), such as Okta, for single sign-on (SSO) access. Cohesity DataProtect as a Service must be added as an application to your IdP such as Okta. The SSO must then be configured along with the SSO URL and certificate file in Cohesity DataProtect as a Service. After the integration, users can sign in to Cohesity DataProtect as a Service using either the IdP sign in page or sign in with the SSO link in the Cohesity DataProtect as a Service login page.

The following identity providers are supported:

Identity Provider	Documentation Link
Active Directory Federation Services (AD FS)	Configure SSO with Active Directory Federation Services (AD FS)
Azure	Configure SSO with Azure
Duo Single Sign-on	Integration with Duo for SSO
Ping Identity	Integration with Ping Identity for SSO
Okta Single Sign-on	Configure SSO with Okta

Configure SSO

To configure SSO:

1. In **DataProtect as a Service**, navigate to **Settings > Access Management > Single Sign-On**.
2. Click **Configure SSO**.
3. Select one of the following options:
 - **SAML**: Security Assertion Markup Language (SAML) is an XML-based protocol used for SSO login.
 - **OpenID Connect**: OpenID Connect is an open authentication protocol that uses OAuth2.0 framework.
4. If you select **SAML**, then refer to the following table:

Name	Description
SSO Domain	<p>Unique domain name that will differentiate this IdP from others. As Cohesity DataProtect as a Service supports multiple IdPs, this has to be a unique string (usually company domain). For a user to be redirected to this IdP, the user will need to log in via SSO using <code>username@SSO_DOMAIN</code>.</p> <p>When a user logs in to Cohesity DataProtect as a Service using SSO and enters the email address as <code>foo@bar.com</code>, Cohesity DataProtect as a Service looks for the IdP that has the SSO Domain configured as <code>bar.com</code> and redirects this user <code>foo</code> to the matching IdP. This is how Cohesity DataProtect as a Service determines which IdP the user needs to be forwarded to.</p>
SSO Provider	From the drop-down, select the SSO provider name of your choice. Select the I have read the SSO documentation provided by <SSO provider name> check box. Cohesity recommends reading the SSO documentation before proceeding to the next step.
Assign to Organization	Optional. In a multitenant-enabled cluster, you can configure SSO for an organization that has been added to the Cohesity cluster. Select an organization from the drop-down.
Single Sign-on URL	Paste the URL that you copied from your IdP.
Provider Issuer ID	Paste the issuer ID that you copied from your IdP.
X.509 Certificate	Click Select File and browse to the location to select the file that you downloaded and renamed previously.

5. If you select **OpenID Connect** (OIDC), perform the following steps and then refer to the table:

Prerequisites:

1. Create the OIDC app within your Identity Provider (IdP). For more information, see [Create OIDC app integrations](#).

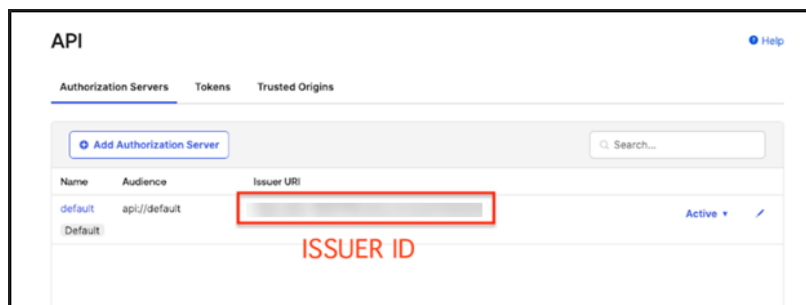
Note: OIDC is an open standard and Single-Sign On with Cohesity DataProtect as a Service is intended to work with any OpenID Connect supported Identity Provider. For setup details, refer to your Identity Provider's documentation.

2. Map the OIDC configuration details from Okta IDP to Cohesity DataProtect as a Service side configurations:

1. To get the Issuer ID:

1. Navigate to **Security > API**.
2. On the **API** page, click **Authorization Servers**.

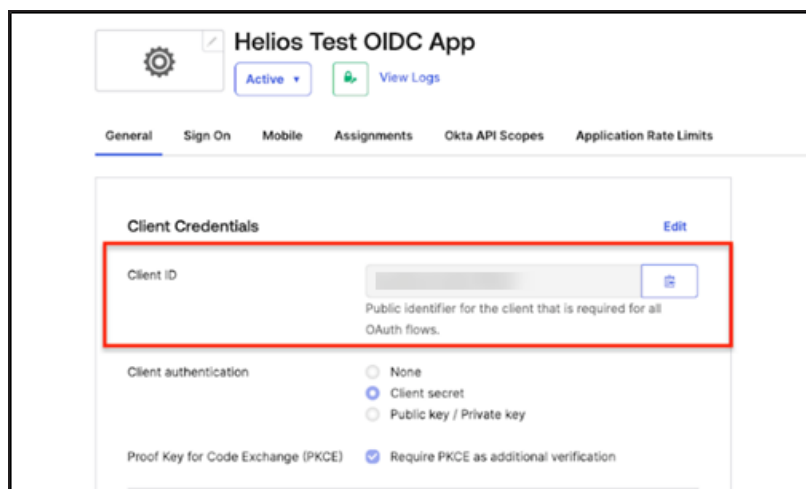
You can find the issuer ID in the Issuer URI section.



2. To get the Client ID:

1. Navigate to Applications.
2. On the Helios Test OIDC App page, click **General**.

You can find the Client ID in the Client Credentials window.



3. To generate the **JSON Web Key Set (JWKS)** URL:

1. Construct the URL as follows:

Format: <issuer ID>/well-known/openid-configuration

For example: `https://***-00000000.okta.com/oauth2/default/.well-known/openid-configuration`.

2. Enter the constructed URL in the address bar of the browser and JSON output will be displayed.

Name	Description
OpenID Server Domain	Enter a unique domain name.
OpenID Server URL for the public (JWKS)	Enter the JSON Web Key Set (JWKS) URL. You can get this URL from your identity provider.
Client ID	Enter the ID of the application created in the identity provider.
Issuer ID	Enter the Issuer ID URL. You can get the URL from your identity provider.
Public Key Expiration (Seconds)	Specifies the time in seconds before which Cohesity starts fetching for new public keys from the identity provider. The default value is 86400 seconds (24 hours).

Name	Description
Public Key Refresh Interval (Seconds)	Specifies the cache refresh interval in seconds to limit the requests to the OIDC server and also to refresh the public key, in case of token signature validation failure. The default value is 600 seconds (10 minutes).
Token Validity (Seconds)	Specifies the validity time in seconds for the token. The validity check is done only if the token is not expired. If it's expired, then the 401 unauthorized or invalid token error is displayed. The default value is 15 minutes.

6. Enter the following details:

Name	Description
Default Role for all SSO Users	Select a role to use as the default role for users signing on with SSO. Typically, you would select this option only during the initial SSO configuration. You can change this option later.
Access to All Clusters or Limited Clusters	Select if the identity provider users can have access to all clusters or limited clusters.

Name	Description
Sign Auth Request	<p>Optional. Enable this option if you want authorization requests to be signed with the Cohesity DataProtect as a Service public key. The Cohesity DataProtect as a Service public key must be uploaded to the IdP site.</p> <p>Note: This option is not available if you select the OpenID Connect protocol.</p> <p>Perform the following steps to obtain the Cohesity DataProtect as a Service public certificate:</p> <ol style="list-style-type: none"> 1. Log in to Cohesity DataProtect as a Service. 2. Start a browser and enter https://helios.cohesity.com/v2/mcm/sslCertificate in the browser address bar. 3. Copy-paste the certificate to Notepad or Word Processor. 4. In the copied certificate, replace \n with a new line. <p>Click to view a sample of the Cohesity DataProtect as a Service public certificate</p> <pre> -----BEGIN CERTIFICATE----- MIIG1zCCBb+gAwIBAgIJAIuZz4iuB+NVMA0GCSqGSIb3DQEBCwUAMIG0MQswCQYD VQQGEwJVUzEQMA4GA1UECBMHQXJpem9uYTETMBEGA1UEBxMKU2NvdHRzZGFsZTEa MBGGA1UEChMRR29EYWRkeS5jb20sIEluYy4xLTArBgNVBAsTJGh0dHA6Ly9jZXJ0 cy5nb2RhZGR5LmNvbS9yZXBvc2l0b3J5LzEzMDEGA1UEAxMQR28gRGFkZHZHkgU2Vj dXJlIENlcnRpZmljYXRlIEF1dGhvcml0eSAtIEcyMB4XDTIwMDEyOTIwMzYzNFoX DTIyMDcyOTIwMzYzNFowRjEhMB8GA1UECxMYRG9tYWluIENvbnRyb2wgVmFsaWRh dGVkMSEwHwYDVQQDExhoZWxpbnM3ZGF0YS5jb2hlc2l0eS5jb20wggeiMA0GCSqG SIb3DQEBAQUAA4IBDwAwggEKAoIBAQDSToInp3D+wBCvJuHfhQwfl8qFr2aWe5rA tu6TV5udPCq+ORqC2UZ05HtLnv9NTXLJtISpH208fJmMBIsmQL6u6LgQ3bA7B3w5 q9e+Q/nsvDUS1MI0wjJsdVb96UZJHU4hRFeFm2seMB1jhscOOaWBdcP3wEaSum8O oSqc7Gs1UGZImxJrNmC0ikCOH9kDK8qj9Bie05CQUM4nGhpzjr3zgGte1MvGBxji GOOW/dW/qB5lmScndAoXmMzwyTQVWxHasXRpYCawGEuG0+V4iGVJs14dSvKT8o4b JOHFwXHcU8mesdfPvq9YTkH6TkYdl5S4WFYygR5rltwzDCc4NmH/AgMBAAGjggNX MIIDUzAMBgNVHRMBAf8EAjAAMB0GA1UdJQQWMBQGCCsGAQUFBwMBBggrBgEFBQcD AjA0BgNVHQ8BAf8EBAMCBaAwOAYDVR0fBDEwLzAtCugKYyNaHR0cDovL2NybC5n b2RhZGR5LmNvbS9nZGlhMnMxLTlxNjcuY3J5SjMF0GA1UdIARWMFQwSAYLYIZIAYb9 </pre>

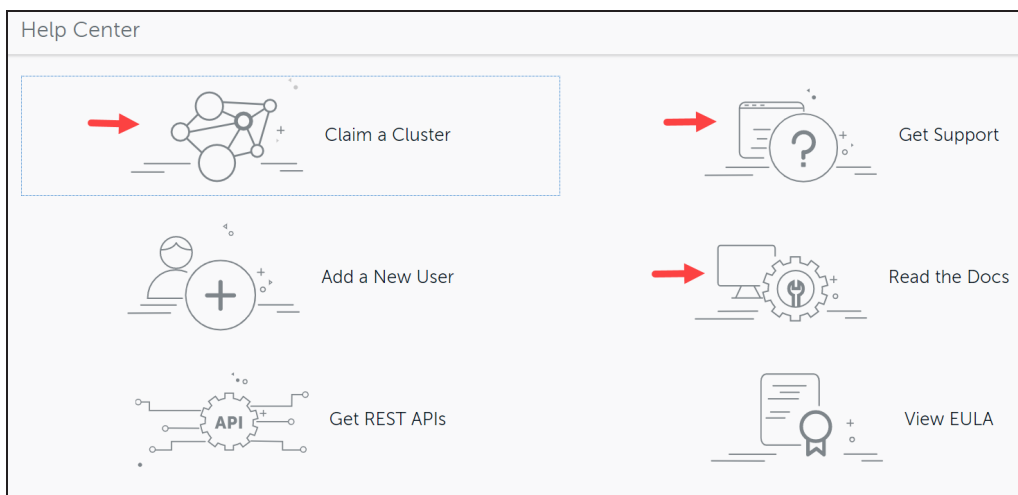
Name	Description
	<pre>bQEhFwEwOTA3BggrBgEFBQCcCARYraHR0cDovL2N1cnRpZmljYXR1cy5nb2RhZGR5 LmNvbS9yZXBvc2l0b3J5LzAIBgZngQw -----END CERTIFICATE-----</pre> <p>5. Save the Notepad or Word Processor as .pem or .crt format.</p> <p>6. The Cohesity DataProtect as a Service public key must be uploaded to the IdP site.</p>

7. Click **Save**.

Cohesity DataProtect as a Service validates the connection to the IdP. If the connection succeeds, the SSO provider is added to the provider list and you can edit, delete or deactivate the provider. Users can start accessing Cohesity DataProtect as a Service through their IdP home page or the Cohesity DataProtect as a Service sign-in page by clicking the **Sign in with SSO** link.

Considerations

- If you have logged into Cohesity DataProtect as a Service using Okta credentials (or any other IdP), you will not be able to directly access some of the portals in the Help Center such as Claim a Cluster, Get Support, and Read the Docs as these portals require Cohesity Support portal credentials to log in.



- If no default role is assigned to a user in the IdP entry, then such users will be rejected. Users will need to have an explicit entry.
- If the SAML assertions are to be signed and encrypted, then the Cohesity DataProtect as a Service certificate must be used.

Next > Add Cohesity DataProtect as a Service **users and groups** from your SSO domain.

Configure SSO with Active Directory Federation Services (AD FS)

This topic provides step-by-step instructions on configuring and using Active Directory Federation Services (AD FS) on Cohesity SSO.

Prerequisites

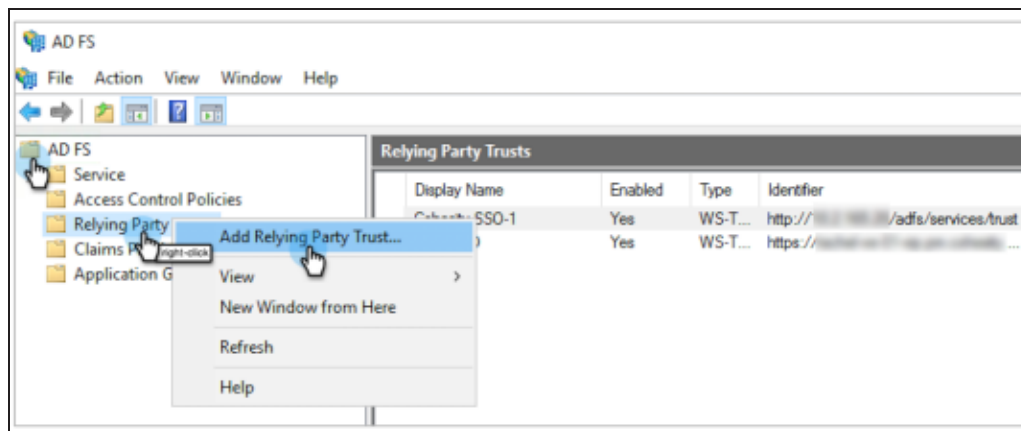
- Install AD FS on the server. For more information, see [Deploy and configure AD FS](#).
- An Active Directory instance where all users have an email address attribute.
- A server running Microsoft Server 2016, 2012, or 2008.
- An SSL certificate to sign your AD FS login page and the Signing Certificate for that certificate.
- An installed certificate for hosted SSL.

Add a Relying Party Trust (RPT)

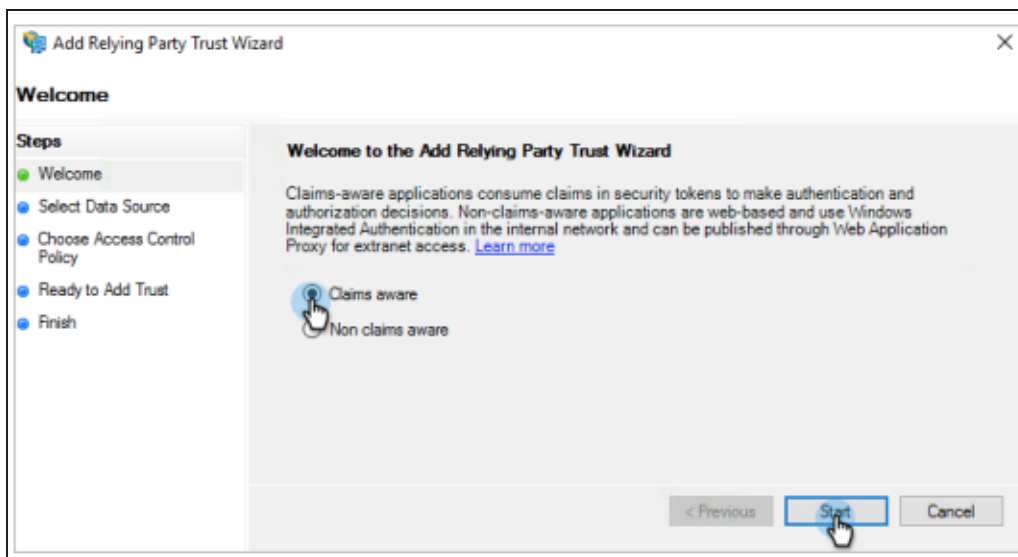
Perform the following steps to add a Relying Party Trust (RPT) to enter the Cohesity SSO authenticate URL via the SAML 2.0 WebSSO protocol.

1. Log in to the server and open **AD FS**.
2. Under **AD FS**, right-click **Relying Party Trusts** and select **Add Relying Party Trust**.

The **Add Relying Trust Party Wizard** page is displayed.



3. Select **Welcome**, select **Claims aware**, and then click **Start**.



4. Under **Select Data Source**, select **Enter data about the relying party manually** and click **Next**.



5. Under **Specify Display Name**, in the **Display name** field, enter a display name and click **Next**.

Add Relying Party Trust Wizard

Specify Display Name

Enter the display name and any optional notes for this relying party.

Display name: CohesitySSOSetup

Notes:

< Previous **Next >** Cancel

6. Under **Configure Certificate**, leave the default certificate settings and click **Next**.
7. Under **Configure URL**, do the following:
 1. Select the **Enable Support for the SAML 2.0 WebSSO protocol** check box.
 2. In the **Relying party SAML 2.0 SSO service URL** field, enter :
<https://helios.cohesity.com/v2/mcm/idp/authenticate>

Add Relying Party Trust Wizard

Configure URL

AD FS supports the WS-Trust, WS-Federation and SAML 2.0 WebSSO protocols for relying parties. If WS-Federation, SAML, or both are used by the relying party, select the check boxes for them and specify the URLs to use. Support for the WS-Trust protocol is always enabled for a relying party.

☐ Enable support for the WS-Federation Passive protocol

The WS-Federation Passive protocol URL supports Web-browser-based claims providers using the WS-Federation Passive protocol.

Relying party WS-Federation Passive protocol URL:

Example: <https://fs.contoso.com/adfs/ls/>

☒ Enable support for the SAML 2.0 WebSSO protocol

The SAML 2.0 single-sign-on (SSO) service URL supports Web-browser-based claims providers using the SAML 2.0 WebSSO protocol.

Relying party SAML 2.0 SSO service URL:

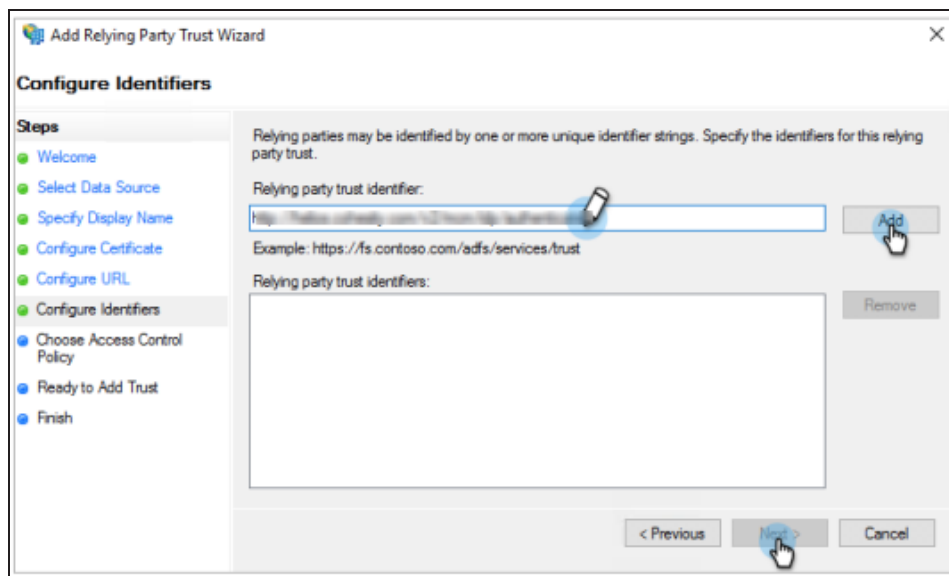
<https://helios.cohesity.com/v2/mcm/idp/authenticate>

Example: <https://www.contoso.com/adfs/ls/>

< Previous **Next >** Cancel

8. Under **Configure Identifiers**, do the following:

1. In the **Relying party trust identifier** field, enter
https://helios.cohesity.com/v2/mcm/idp/authenticate
2. Click **Add** and then click **Next**.




9. Under **Choose Access Control Policy**, you can optionally configure multi-factor authentication (MFA) and click **Next**. For more information, see [Configure Additional Authentication Methods for AD FS](#).
10. Under **Ready to Add Trust**, see an overview of the settings and click **Next**.
11. Under **Finish**, click **Close**.

Create Claim Rules

Cohesity looks for SAML attributes to identify users and assign roles.

Perform the following steps to pass SAML attributes:

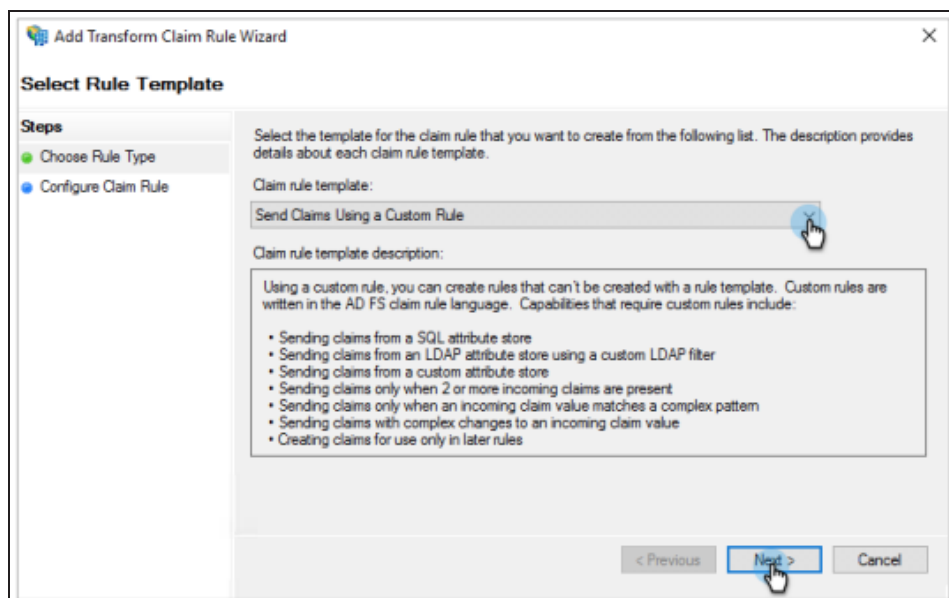
1. Log in to the server and open **AD FS**.
2. Under **AD FS**, select **Relying Party Trusts** and select the RPT that you added.
3. On the right, click **Edit Claim Issuance Policy**.



- Click **Add Rule**.

The **Add Transform Claim Rule Wizard** page is displayed.

- Under **Select Rule Template**, do the following:
 - From the **Claim rule template** drop-down, select **Send LDAP Attributes as Claims**.
 - Click **Next**.



- Under **Edit Rule**, do the following:
 - In the **Claim rule name** field, enter a name.
 - From the **Attribute store** drop-down, select **Active Directory**.
 - In the **Mapping of LDAP attributes to outgoing claim types** table:

1. Under **LDAP Attribute (Select or type to add more)**, from the drop-down, select **User-Principal-Name**.
2. Under **Outgoing Claim Type**, from the drop-down, select **E-Mail Address**.
3. Click **OK**.

Edit Rule - UPN to Email

You can configure this rule to send the values of LDAP attributes as claims. Select an attribute store from which to extract LDAP attributes. Specify how the attributes will map to the outgoing claim types that will be issued from the rule.

Claim rule name:

Rule template: Send LDAP Attributes as Claims

Attribute store:

Mapping of LDAP attributes to outgoing claim types:

	LDAP Attribute (Select or type to add more)	Outgoing Claim Type (Select or type to add more)
▶	User-Principal-Name	E-Mail Address
*		

7. Click **Add Rule** to create another rule.
8. From the **Claim rule template** drop-down, select **Transform an Incoming Claim**.
9. Click **Next**.
10. Under **Edit rule**, do the following:
 1. In the **Claim rule name** field, enter a name.
 2. From the **Incoming claim type** drop-down, select **E-Mail Address**.
 3. From the **Outgoing claim type** drop-down, select **email**.

4. Click **OK**.

Edit Rule - Email address to Cohesity email Attribute

You can configure this rule to map an incoming claim type to an outgoing claim type. As an option, you can also map an incoming claim value to an outgoing claim value. Specify the incoming claim type to map to the outgoing claim type and whether the claim value should be mapped to a new claim value.

Claim rule name:

Rule template: Transform an Incoming Claim

Incoming claim type:

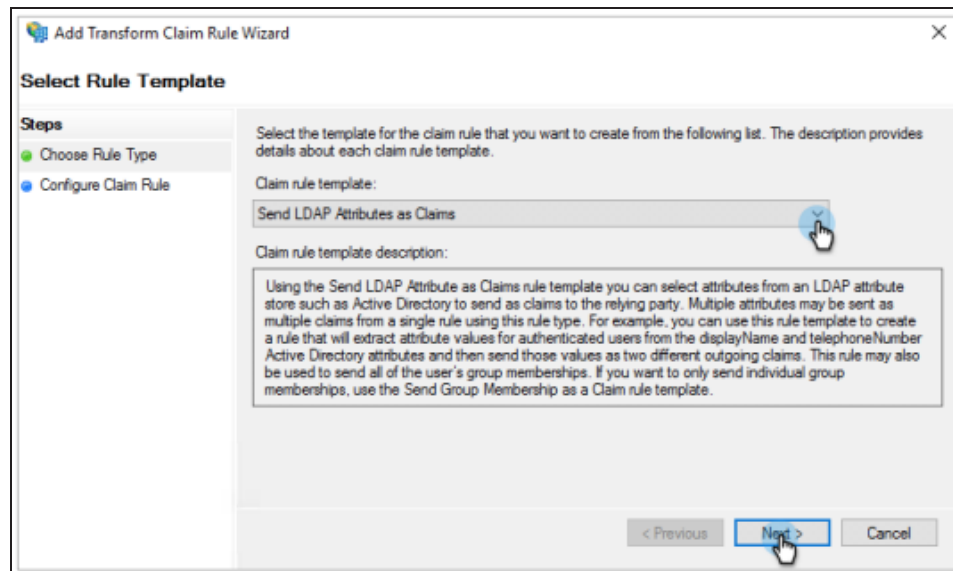
Incoming name ID format:

Outgoing claim type:

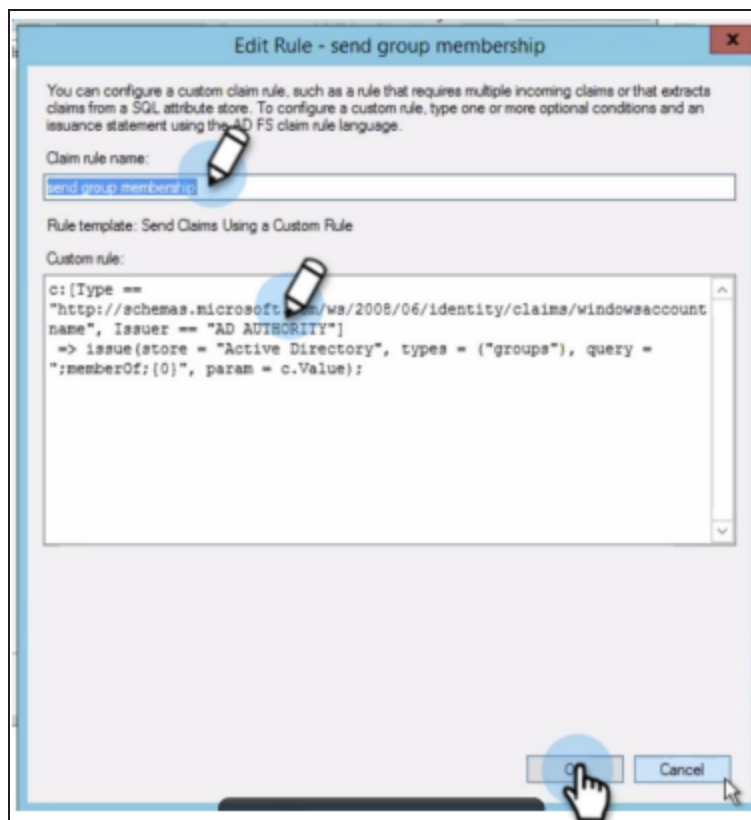
Outgoing name ID format:

☒ Pass through all claim values
☐ Replace an incoming claim value with a different outgoing claim value
 Incoming claim value:
 Outgoing claim value:
☐ Replace incoming e-mail suffix claims with a new e-mail suffix
 New e-mail suffix:
 Example: fabrikam.com

11. Follow the steps above to pass group SAML attributes.
12. To extract the user group name and send it to Cohesity, you need to create a custom rule in AD FS:
 1. Click **Add Rule** to create the custom rule.
 2. From the **Claim rule template** drop-down, select **Send Claims Using a Custom Rule**.
 3. Click **Next**.



4. Under **Edit rule**, do the following:
 1. In the **Claim rule** name field, enter a name.
 2. In the **Custom rule** field, create and enter a custom rule. For more information, see [Understanding Claim Rule Language in AD FS](#).
 3. Click **OK**.



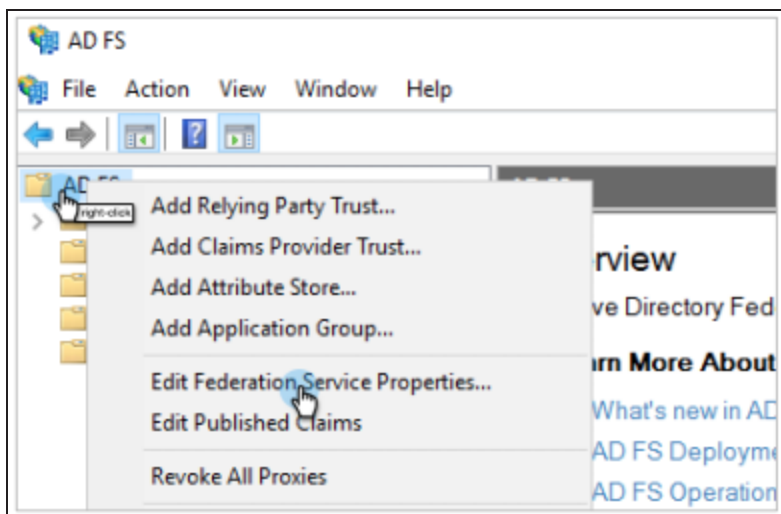
Note: This rule might be different for different AD FS configurations. Make sure to edit the custom rule accordingly. For more information, see [When to Use a Custom Claim Rule](#).

Retrieve the SSO URL, Provider Issuer ID, and Certificate

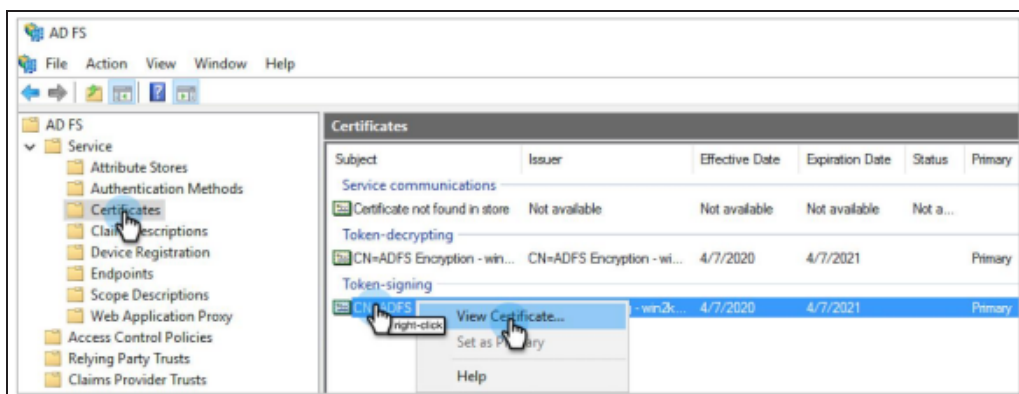
You need to retrieve the Federation Service name and Federation Service Identifier which is required when adding AD FS as an SSO provider to Cohesity.

Perform the following steps to retrieve the Federation Service name and Federation Service Identifier:

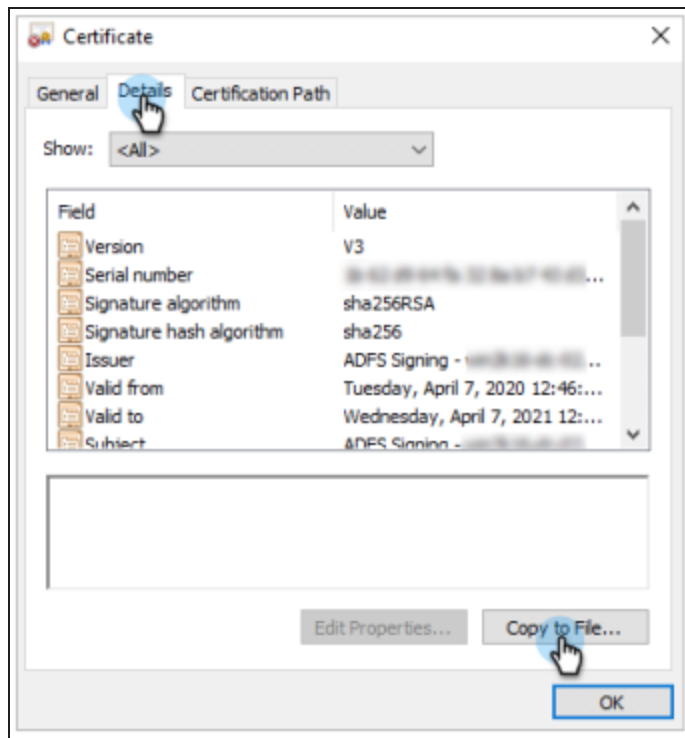
1. Log in to the server and open **AD FS**.
2. Right-click **AD FS** and select **Edit Federation Service Properties**.



3. Copy the **Federation Service name** and the **Federation Service Identifier** and save it for later use. You will need these when you [Configure SSO](#) to Cohesity.
4. To download the certificate, navigate to **AD FS > Service > Certificates**.
5. Under **Token-signing**, right-click the certificate and select **View Certificate**.



6. Click the **Details** tab and then click **Copy to File**.
The **Certificate Export Wizard** page is displayed.



7. Select **Base-64 encoded X.509 (.CER)**, click **Next**, and follow the instructions to download the certificate (.cer).
8. Convert certificate file from the .cer to the .pem format.

To convert the file:

- On Mac/Linux, rename the file with the `.pem` filename extension.
- On Windows, run the following command:

```
openssl x509 -in mycert.crt -out mycert.pem -outform PEM
```

Consideration


Cohesity DataProtect as a Service does not support **Sign Auth Requests** to sign the SAML requests to the ADFS server.

Configure SSO with Azure

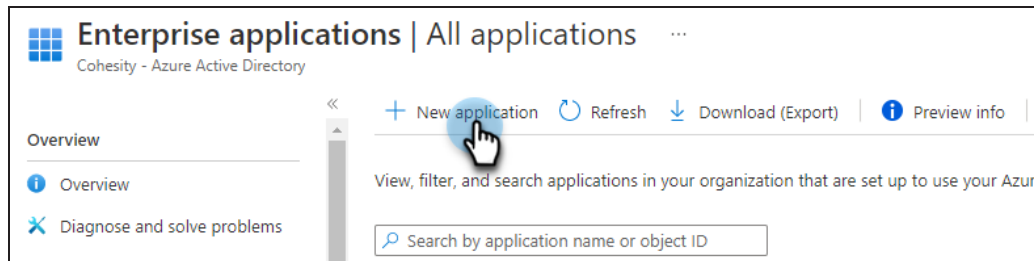
This topic provides step-by-step instructions on creating an Azure Active Directory application.

Perform the following steps to create an Azure AD SSO:

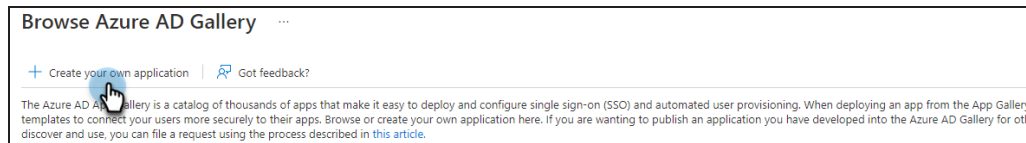
1. Log in to [Azure portal](#).
2. Under **Azure services**, click **Azure Active Directory**. If Azure Active Directory is not listed, click **More Services** and select **Azure Active Directory**.



3. On the left, click **Enterprise applications**.
4. Under **All applications**, click **New Application**.




5. On the **Browse Azure AD Gallery** page, click **Create your own application**.




6. In the **What's the name of your app**, enter a display name for your application.
7. Select **Integrate any other application you don't find in the gallery (Non-gallery)** and click **Create**.

Create your own application ✕

 Got feedback?

If you are developing your own application, using Application Proxy, or want to integrate an application that is not in the gallery, you can create your own application here.


What's the name of your app?




What are you looking to do with your application?

☐ Configure Application Proxy for secure remote access to an on-premises application

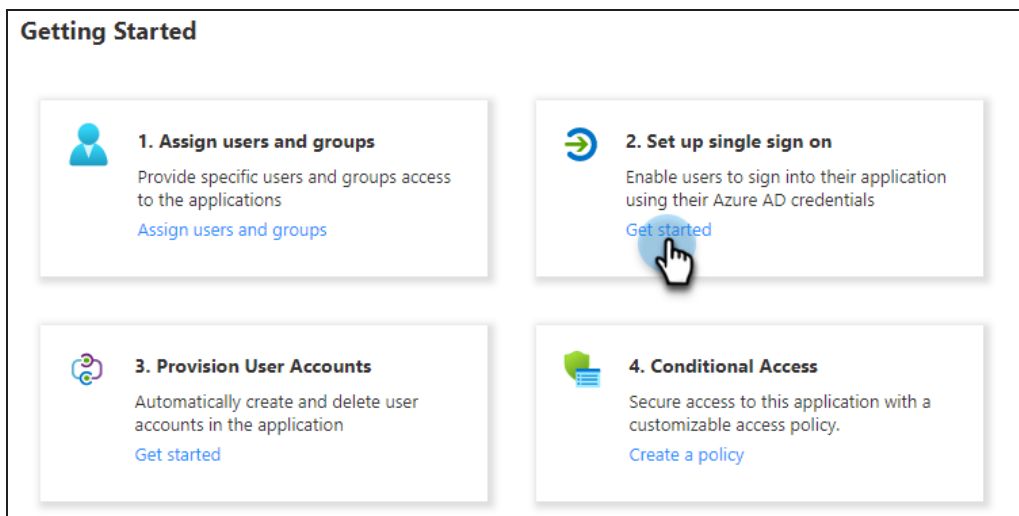
☐ Register an application to integrate with Azure AD (App you're developing)


☒ Integrate any other application you don't find in the gallery (Non-gallery) 

Create



8. On the **<app> Overview** page, under **General Settings**, on the **Set up single sign on** tile, click **Get Started**.



9. Under **Select a single sign-on method**, click the **SAML** tile.
10. Under **Set up Single Sign-On with SAML**, do the following:
 1. In the **Basic SAML Configuration** section, click the edit  icon and do the following:
 1. Under **Identifier (Entity ID)**, click **Add identifier**.
For example,
`https://helios.cohesity.com/v2/mcm/idp/authenticate`
 2. Under **Reply URL (Assertion Consumer Service URL)**, click **Add reply URL**.
For example,
`https://helios.cohesity.com/v2/mcm/idp/authenticate`
 3. Click **Save**.

Note: If you have multiple Cohesity clusters and you want to use this Azure AD application for all of them, you can use the additional cluster FQDNs to enter multiple **Identifiers** and

Reply URLs in this step.

Basic SAML Configuration

Save

Identifier (Entity ID) * ⓘ
The default identifier will be the audience of the SAML response for IDP-initiated SSO

https://helios.cohesity.com/v2/mcm/idp/authenticate ✓

Reply URL (Assertion Consumer Service URL) * ⓘ
The default reply URL will be the destination in the SAML response for IDP-initiated SSO

https://helios.cohesity.com/v2/mcm/idp/authenticate ✓

2. In the **Attributes & Claims** section, click the edit icon and do the following:
 1. Click **Add new claim**.
The **Manage claim** page is displayed.
 2. **Name**: Enter a name for the attribute.
 3. **Source**: Select Attribute.
 4. **Namespace**: Optional. Enter a namespace URI.
 5. **Source attribute**: From the drop-down, select the source attribute.
 6. Click **Save**.

Manage claim ...

Save Discard changes Got feedback?

Name * email ✓

Namespace Enter a namespace URI ✓

Source * ☒ Attribute ☐ Transformation

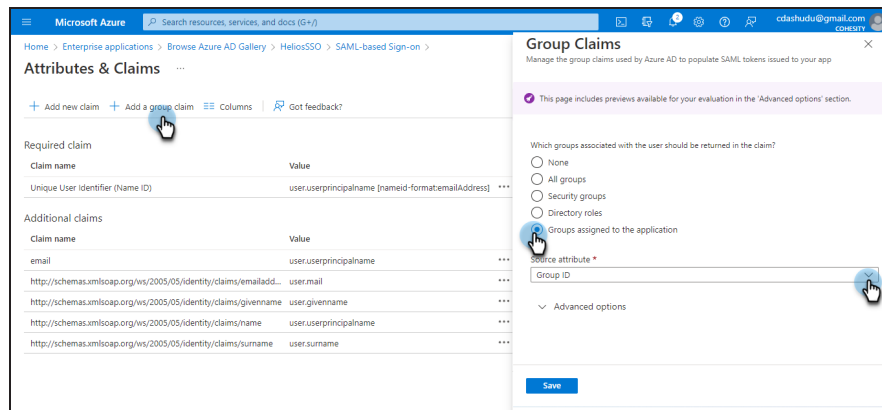
Source attribute * user:principalname

Claim conditions

3. If you plan to use user groups-based RBAC, you need to pass the "Groups" SAML attribute to Cohesity. Perform the following steps:
 1. Under **User Attributes & Claims**, click **Add a group claim**.
 2. For **Which groups associated with the user should be returned in the claim?**, select **Groups assigned to the application**.

Note: Groups must be directly assigned to the application. Azure will not send the groups attribute that are a subgroup of a group which is assigned to the application.

- From the **Source attribute** drop-down, select the source attribute.



- Under **Advanced options**:
 - Select the **Customize the name of the group claim** check box.
 - Name:** Enter a name as groups.
 - Namespace:** Enter the namespace URI. This is optional.

d. Click **Save**.

Group Claims
Manage the group claims used by Azure AD to populate SAML tokens issued to your app

This page includes previews available for your evaluation in the 'Advanced options' section.

Advanced options

☒ Customize the name of the group claim

1 Name (required)
groups

Namespace (optional)

☐ Emit groups as role claims ⓘ

☐ Apply regex replace to groups claim content (Preview)

Save

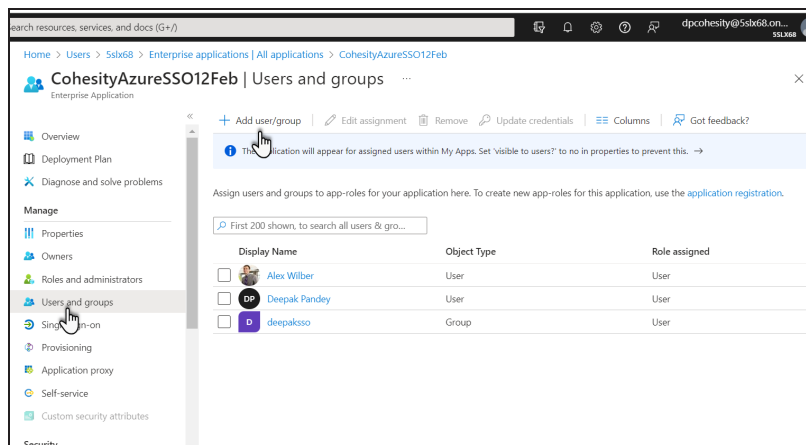
2

Note: To use source attributes like sAMAccountName to pass the user group name in the "Groups" SAML attribute make sure that Azure AD groups are synchronized from an on-premises Active Directory using Azure AD Connect Sync 1.2.70.0 or above. For more information, see [Azure AD Connect: Upgrade from a previous version to the latest](#).

If you don't have an on-prem Active Directory synced with Azure AD, in the **Source** attribute drop-down, select **Group ID**.

4. Depending on the value of the Source attribute you selected, you need to create the corresponding **users and groups**. For example, if you use:
 1. **sAMAccountName**, you need to create groups with the SSO Group value as the AD groups name.
 2. **Group ID**, you need to create SSO groups using **Azure AD's Group ID**. To obtain the Azure AD's Group ID:

- a. Click the application name
- b. Under **Manage**, click **Users and groups**.





- c. Click **Add user/group** to assign a user or a group who should be able to access Cohesity DataProtect as a Service using this Azure AD application.
- d. From the list of users, click a user.

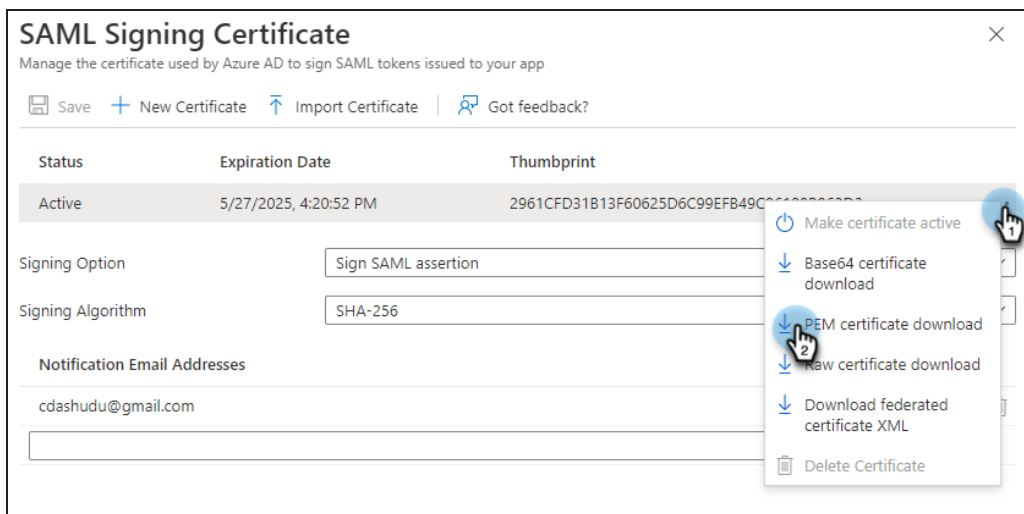
Note: Nested groups are not supported and will not be passed under the Groups SAML attributes.

Retrieve the SSO URL, Provider Issuer ID, and Certificate

You need to retrieve Azure AD information to configure SSO on Cohesity DataProtect as a Service for the IdP (Azure AD).

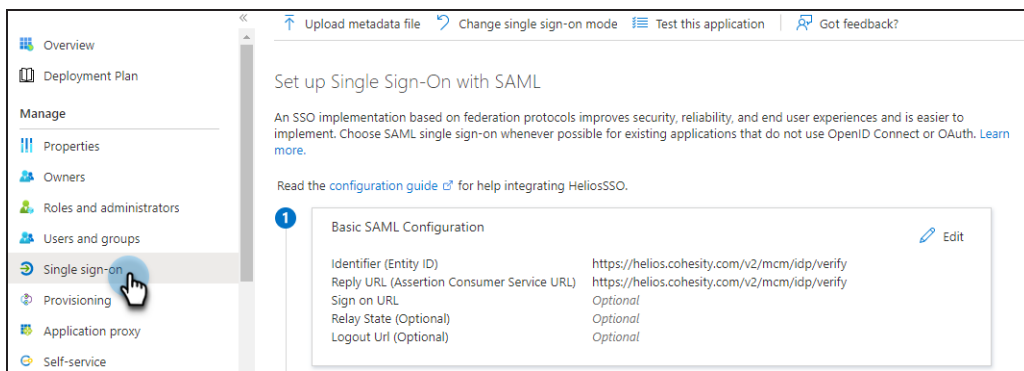
Perform the following steps to retrieve the SSO URL, Entity ID, and certificate from the Azure AD application:

1. Log in to [Azure portal](#).
2. Under **Azure services**, click **Azure Active Directory**. If Azure Active Directory is not listed, click **More Services** and select **Azure Active Directory**.
3. On the left, click **Enterprise applications**.
4. Click the application name and under **Manage**, click **Single sign-on**.
5. Under **Set up Single Sign-On with SAML**, in the **SAML Signing Certificate** section, click the edit  icon.
6. On the **SAML Signing Certificate**, click the ellipsis () icon and select **PEM certificate download**.



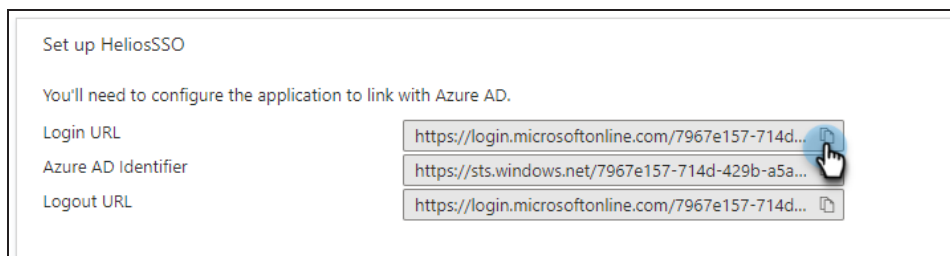
Note: Cohesity SSO only accepts *.pem format certificate.

7. Under **Manage**, click **Single sign-on**.



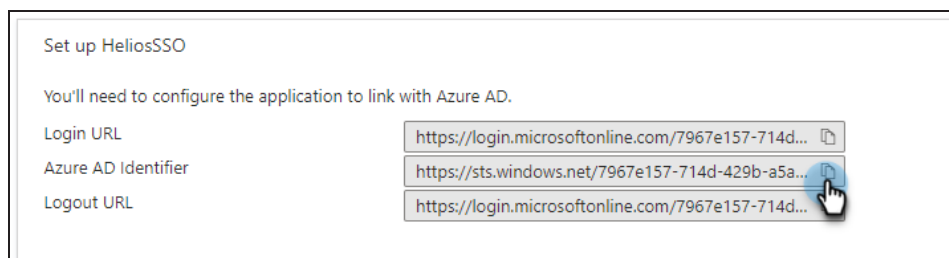
8. Under **Set up Single Sign-On with SAML**, in the **Set up <application name>** section, do the following:

1. Copy the **Login URL** and save it for later use. You will use this URL to enter the Cohesity Single-Sign-On URL when you **Configure SSO** to Cohesity.



2. Copy the **Azure AD Identifier** URL and save it for later use. You will use this URL to enter the Cohesity Provider Issuer ID when you **Configure SSO** to

Cohesity.



Set up HeliosSSO

You'll need to configure the application to link with Azure AD.

Login URL	https://login.microsoftonline.com/7967e157-714d...
Azure AD Identifier	https://sts.windows.net/7967e157-714d-429b-a5a...
Logout URL	https://login.microsoftonline.com/7967e157-714d...

You need to add the SSO provider in Cohesity DataProtect as a Service. For more information, see [Configure SSO](#).

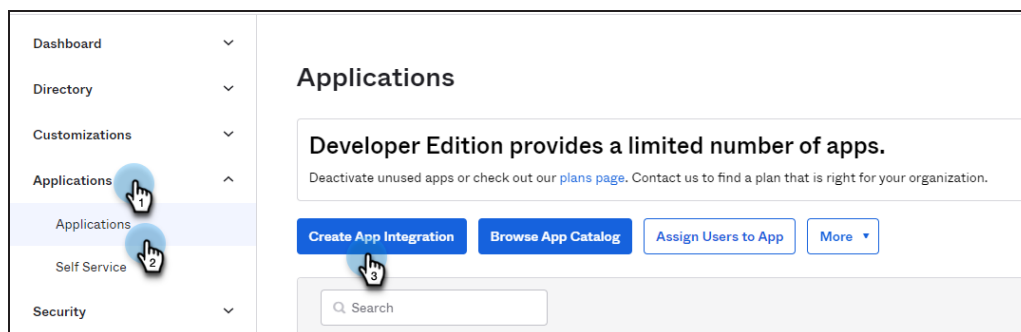
Configure SSO with Okta

This topic provides step-by-step instructions on adding Cohesity DataProtect as a Service as an application to Okta.

Perform the following steps to add Cohesity DataProtect as a Service as an application to Okta:



1. Log in to Okta as an Okta administrator.
2. Navigate to **Applications > Applications** and click **Create App Integration**.

The **Create a New Application Integration** page is displayed.



3. For the **Sign on method**, select **SAML 2.0** and click **Next**.

The **Create SAML Integration** page is displayed.

4. Click the **General** tab and for **General Settings** do the following:
 1. **App Name**: Specify an app name of your choice to display in the Cohesity DataProtect as a Service tile on the SSO page.
 2. **App logo (optional)**: Click  > **Browse files** and navigate to the location of the logo and select the logo. Click **Apply** to upload the logo. Click  to

delete the logo.

3. **App Visibility:** Leave the default settings for **Do not display application icon for users** and **Do not display application icon in the Okta Mobile app**.
4. Click **Next**.

1 General Settings

App name: CohesitySSO

App logo (optional): [Gear icon] [Delete icon]

App visibility:

- ☐ Do not display application icon to users
- ☐ Do not display application icon in the Okta Mobile app

Cancel Next

5. Click the **Configure SAML** tab and for **SAML Settings** do the following:

1. **Single sign on URL:** Specify the application URL followed by `/idps/authenticate`.

For example: `https://<cluster_fqdn>/idps/authenticate`.

For Helios use, `https://helios.cohesity.com/v2/mcm/idp/authenticate`.


Note: To find the FQDN and VIP address, log in to Cohesity Data Cloud (Self-managed) and navigate to **Settings > Cluster > Networking > VIPs**.

The **Use this for Recipient URL and Destination URL** check box is selected by default.

2. **Audience URI (SP Entity ID):** Specify the same URL as above.
3. **Application username:** Select your preference.


A SAML Settings

General

Single sign on URL ⓘ 

☒ Use this for Recipient URL and Destination URL


☐ Allow this app to request other SSO URLs

Audience URI (SP Entity ID) ⓘ 

Default RelayState ⓘ

If no value is set, a blank RelayState is sent





Name ID format ⓘ



Application username ⓘ 

4. Under **Attribute Statements**, map the Email and/or Login SAML attributes to the Okta user profile attributes. If the value is not available in the drop-down list, type it as shown in the table. You can map either or both attributes.

SAML Attribute	Okta User Profile Attribute Value
Email	user.email
Login	user.login

Attribute Statements (optional) [LEARN MORE](#)

Name	Name format (optional)	Value
email 	Unspecified ▼	user.email 
login 	Unspecified ▼	user.login 

[Add Another](#)

5. Under **Group Attribute Statements (Optional)**, map the groups attribute to the Okta Filter attribute. (For example, select **Starts with** and enter **cohesity_** to pass any group name that starts with 'cohesity_' to Cohesity.) If you want

to use an existing group, use a regex to pass all groups.

Note: You should enter "groups" in the name field to map the groups attribute to the Okta Filter attribute.

Name	Name format (optional)	Filter
groups	Unspecified	Matches regex (.*)

[Add Another](#)

6. Click **Next**.
7. Click **Finish** to add the application.
6. Click the **Sign On** tab and do the following:
 1. Under **SAML Setup**, located at the right side, click **View SAML setup instructions**.

The **How to Configure SAML 2.0 for <application name>** page is displayed.

Type	Created	Expires	Status	Actions
SHA-1	Today	Sep 2028	Inactive ⚠	Actions
SHA-2	Today	May 2032	Active	Actions

SAML Setup
Single Sign On using SAML will not work until you configure the app to trust Okta as an IdP.

[View SAML setup instructions](#)

2. Copy the **Identity Provider Single Sign On URL** and save it for later use. You will use this URL to enter the Cohesity Single Sign-On URL when you [Configure SSO](#) to Cohesity.

A sample URL is shown below.

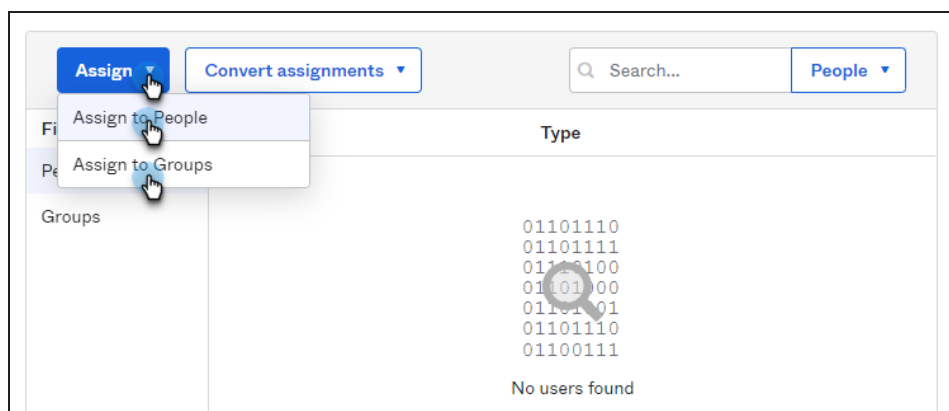
```
https://mycompany.okta.com/app/cohesitymycompany_heliosapp/exkhhbyzrgu0YvJFk0h7/sso/saml
```

3. Copy the **Identity Provider Issuer** and save it for later use. You will use this URL to enter the Cohesity Provider Issuer ID when you [Configure SSO](#) to Cohesity.

A sample URL is shown below.

`http://okta.com/exkhhbyzrgu0YvJFk0h7`

4. Click **Download certificate** to download the `okta.cert` file and note its download location.
5. Convert the downloaded `okta.cert` file to `okta.pem`. You will upload this file to Cohesity DataProtect as a Service later.
7. Click the **Assignments** tab and do the following:
 1. From the **Assign** drop-down, select **Assign to People** to assign users to your Cohesity Okta application.
 2. From the **Assign** drop-down, select **Assign to Groups** to assign groups to the app.



You have now configured the Okta application for Cohesity. You need to add the SSO provider in Cohesity DataProtect as a Service. For more information, see [Configure SSO](#).

Add API Keys

You can add your Cohesity API keys to your Cohesity DataProtect as a Service to:

- Authenticate an application or script for reporting and workflow automation via Cohesity's REST API calls for Cohesity DataProtect.

To add your API key:

1. In **DataProtect as a Service**, navigate to **Settings > Access Management**, and click the **API Keys** tab.
2. Click **Add API Key**.
3. Enter a **Name** for the API key.
4. Click **Save** to advance to the **API Key Details** page, where you can:
 - **View or Copy API Key Token**. To use with the application or script you wish to authenticate.

- **Scan QR Code.** Scan the QR code that is displayed with your Helios Mobile App to monitor your Cohesity DataProtect as a Service in the mobile app.

When you return to the **API Keys** tab, your new key appears in the list.

Note: The API keys you add are available only to you.

Click the **Actions** menu (:) next to the API key to **Delete** it.

Sample API Keys



Once you have [added an API Key](#), you can start making API calls.

Below is a sample API key,

API Key Details

The API Key Token will be available only once on creation. Please store it in a secure location.

Use API Keys to authenticate an application or script to Helios for management by APIs. Refer to Helios REST API documentation for details about using these keys.

Name	Sample
API Key Token	*****  

The API Key Token is only available once upon creation. Make sure to store it in a secure location. Use this key to authenticate an application or script to Helios for API management. For the detailed list of APIs, see <https://api.cohesity.com>.

Policies

In Cohesity DataProtect as a Service, a policy is a reusable collection of settings that define how and when the objects & files in a source are protected. You can create as many policies with specific settings for different use cases as you need.

In a policy, you set the frequency (**Backup every**) and retention period (**Keep for**) for each protection run. You can also add a Periodic Full Backup, Quiet Times, and Log Backup schedules — see [More Options](#).

Create a Policy

To create a policy:

1. In **DataProtect as a Service**, navigate to **Policies**.
2. Click **Create Policy**.
3. Enter a **Policy Name**, choose a **Backup every** interval and a **Keep for** retention period.
4. If you wish to add a DataLock, Periodic Full Backup, Quiet Times, or schedule database Log Backups, click [More Options](#).
5. Click **Create**.

More Options

Settings	Descriptions
DataLock	<p>Typically used for compliance and regulatory purposes, DataLock is a protection policy option that can only be enabled by a user with the Data Security role. Use it when you need to prevent the deletion of backup snapshots for a specified duration. You can set the DataLock duration to the same period as your backup retention, or to a shorter period.</p> <div> <p>Note: Only a user with the Data Security role can enable or disable DataLock on a policy, or delete or edit a DataLocked policy. Disabling a DataLock does not unlock any previously DataLocked snapshots.</p> </div>
Periodic Full Backup	<p>After the first Protection Run, Cohesity DataProtect as a Service backs up only the data that changed with <i>incremental</i> backups. Use this option to add a <i>full</i> backup run at regular intervals.</p>

Settings	Descriptions
Quiet Times	<p>If there are times you need to protect your network from too much traffic, add a Quiet Time period to define the times when new Protection Runs do not start. (Note that those already running at the beginning of a Quiet Time will still complete the run.) By default, a Quiet Time period is set in your browser's time zone.</p> <div>Tip: To add more Quiet Time periods, click Quiet Times again.</div>
Log Backup	<p>If you are protecting databases, you can set a separate frequency and retention period for your log backups.</p>

Add or Remove Legal Hold to Snapshot

Users who are assigned the Data Security role can add legal hold on existing snapshots to preserve them for legal purposes. Legal hold snapshots cannot be deleted until the legal hold is removed. Legal hold can be added to unlocked and DataLocked snapshots. Cohesity recommends that the data security user enables Legal Hold for business-critical workloads after a ransomware attack is confirmed.

You can add a legal hold to the snapshots of object protection runs.

Currently, Cohesity supports the following:

- Legal hold on single run of a single object (both UI and API),
- Legal hold on multiple runs of a single object (API only),
- Legal hold on multiple runs of multiple objects (API only).

Note: Legal hold prevents snapshots from being deleted until the legal hold is removed.

Using legal hold for long periods of time may result in the cluster running out of space.

To add a legal hold for the backup targets in a protection run:

1. In DataProtect as a Service, under **Sources**, find the required source and click on it.
2. Under **Activity**, click a protection run.
3. Click the **Actions** menu (⋮) and select **Add Legal Hold**. This option is enabled only if you have the Data Security role assigned. Users assigned with other roles can only view the Legal Hold status.
4. The snapshot details are displayed in the **Add Legal Hold** window. Click **Save**.

To remove legal hold from the backup targets in a protection run:

1. In DataProtect as a Service, under **Sources**, find the required source and click on it.
2. Under **Activity**, click a protection run.
3. Click the Actions menu (⋮) and select **Remove Legal Hold**.
4. The snapshot details will be displayed in the **Remove Legal Hold** window. Click **Save**.

To filter the legal hold enabled protection runs:

1. In DataProtect as a Service, under **Sources**, find the required source and click on it.
2. In the **Task Status** drop-down, select **Legal Hold**. The protection runs for which legal hold is enabled will be displayed.

Virtual Machines

Cohesity DataProtect as a Service unifies fragmented data protection solutions for virtualized environments. With Cohesity DataProtect as a Service, organizations no longer need to deal with complex and expensive protection solutions that result in multiple infrastructure silos and copies of data.

VMware

Cohesity DataProtect as a Service provides a simple, fast, cost-effective backup, recovery, and data management solution for VMware environments.

VMware Requirements and Considerations

To register VMware VMs, ensure your vCenter or standalone ESXi host meets these software versions and user privilege requirements. Check your software versions and the user role privileges you'll need on [vCenter](#) or [standalone ESXi](#) below.

Note: For information on the supported cloud regions where you can back up this source, see [Supported Workloads and Cloud Regions](#).

Firewall Ports

Ensure that the ports listed in the VMware section of the [Firewall ports](#) topic are open to allow communication between the Cohesity SaaS Connector(s) and VMware environment.

Support Matrix

Before you register your VMware sources, ensure that you have the supported VMware environments. For more information, see [Supported Software for Cohesity DataProtect as a Service](#).

Add User Privileges for vCenter Sources

If the VMware source is vCenter, ensure that the user account has the role privileges listed for each category below.

Note: Starting with VMware vSphere 8.0 version, the **Profile-driven Storage** privilege level is replaced with **VM Storage Policies**.

Category	Privileges	Notes
Cryptographic Operations*	<ul style="list-style-type: none"> Add Disk Direct Access Encrypt New 	* Required only for encrypted VMs
Datastore	<ul style="list-style-type: none"> Allocate space Browse datastore Configure datastore* Low level file operations Move datastore Remove file 	* Required only if Source Datastore throttling is enabled.
Folder	<ul style="list-style-type: none"> Create folder Delete folder 	
Global	<ul style="list-style-type: none"> Disable Methods Enable Methods Licenses Log event Manage custom attributes Set custom attribute 	
Host > Configuration	<ul style="list-style-type: none"> Maintenance Query patch Storage partition configuration 	
Host > Local operations	Reconfigure virtual machine	
Network	<ul style="list-style-type: none"> Assign network 	
Profile-driven Storage	<ul style="list-style-type: none"> Profile-driven storage update Profile-driven storage view 	

Category	Privileges	Notes
Resource	<ul style="list-style-type: none"> Assign virtual machine to resource pool Migrate powered off virtual machine Migrate powered on virtual machine 	
Session	<ul style="list-style-type: none"> View and stop sessions 	
Virtual Machine > Change Configuration	<ul style="list-style-type: none"> Acquire disk lease Add existing disk Add new disk Add or remove device Advanced configuration Change Settings Change Swapfile placement Configure Raw device Display connection settings Remove disk Rename Toggle disk change tracking UpgradeVirtualHardware 	
Virtual Machine > Change Operations (For Runbook)	<ul style="list-style-type: none"> Change CPU count Change Memory Change Settings Change resource Modify device settings Rename* 	* Rename permission is required for a copy recovery.

Category	Privileges	Notes
Virtual Machine > Edit Inventory	<ul style="list-style-type: none"> • Create new • Register • Remove • Unregister 	
Virtual Machine > Guest Operations	<ul style="list-style-type: none"> • Guest operation modifications • Guest operation program execution • Guest operation queries 	
Virtual Machine > Interaction	<ul style="list-style-type: none"> • Connect devices • Guest operating system management by VIX API • Power off • Power on 	
Virtual Machine > Provisioning	<ul style="list-style-type: none"> • Allow disk access • Allow read-only disk access • Allow virtual machine download • Customize guest* • Mark as template 	*Required for Runbook
Virtual Machine > Snapshot Management	<ul style="list-style-type: none"> • Create snapshot • Remove snapshot • Revert snapshot 	
vApp	<ul style="list-style-type: none"> • Add virtual machine • Assign resource pool • Unregister 	

Category	Privileges	Notes
vSphere Tagging	<ul style="list-style-type: none"> Assign or Unassign vSphere Tag Assign or Unassign vSphere Tag on Object 	

Add User Privileges for Standalone ESXi Sources

If the VMware source is standalone ESXi, ensure that the user account has the role privileges listed for each category below.

Category	Privileges	Notes
dvPort Group	<ul style="list-style-type: none"> Create Modify 	
dvSwitch	<ul style="list-style-type: none"> Create Delete 	
Datastore	<ul style="list-style-type: none"> AllocateSpace Browse Config* Delete* DeleteFile FileManagement Move* Rename* UpdateVirtualMachineFiles* UpdateVirtualMachineMetadata* 	* Required only if Source Datastore throttling is enabled
Folder	<ul style="list-style-type: none"> Create Delete 	

Category	Privileges	Notes
Global	<ul style="list-style-type: none"> • DisableMethods • EnableMethods • Licenses • LogEvent • Manage custom attributes • Set custom attribute 	
Host > Configuration	Storage	
Host > Local operations	Delete virtual machine	
Network	Assign	
Resource	<ul style="list-style-type: none"> • AssignVMToPool • ColdMigrate • HotMigrate 	
System	<ul style="list-style-type: none"> • Anonymous • Read • View 	
vApp	<ul style="list-style-type: none"> • AssignResourcePool • AssignVM • Unregister 	
Session	View and stop sessions	

Category	Privileges	Notes
Virtual machine > Configuration	<ul style="list-style-type: none"> AddExistingDisk AddNewDisk AddRemoveDevice AdvancedConfig CPUCount ChangeTracking DiskLease EditDevice HostUSBDevice Memory RawDevice ReloadFromPath RemoveDisk Rename ResetGuestInfo Resource Settings SwapPlacement UpgradeVirtualHardware 	
Virtual machine > Guest Operations	<ul style="list-style-type: none"> Execute Modify Query 	
Virtual machine > Interact	<ul style="list-style-type: none"> GuestControl PowerOff PowerOn 	

Category	Privileges	Notes
Virtual machine > Inventory	<ul style="list-style-type: none"> • Create • Delete • Register • Unregister 	
Virtual machine > Provisioning	<ul style="list-style-type: none"> • DiskRandomRead • GetVmFiles 	
Virtual Machine > State	<ul style="list-style-type: none"> • Create snapshot • Remove snapshot • Revert to snapshot 	
Cryptographic Operations	<ul style="list-style-type: none"> • Add Disk • Direct Access • Encrypt • Migrate 	

Considerations

- If you are protecting a VM equipped with vTPM, enable the *Cryptographer.ManageKeyServers* privilege in the vCenter for the user account registered with Cohesity.
- If you are recovering a VM equipped with vTPM using a different KMS or replacing the current KMS with a new one, the vTPM VMs may not start after the recovery process, even if the new KMS has the same name.
- The supported maximum size of the NVRAM file is 10 MB. Contact your Cohesity account team if your NVRAM files are larger than 10 MB.
- Restoring the NVRAM file from the encrypted VMs is not supported, except for vTPM encryption.
- The maximum size of a DataSets file can be up to 300 MB. In most cases, the DataSets file size will be a few MBs. Contact your Cohesity account team if the files exceed 300 MB.
- Restoring the DataSets file from the encrypted VMs is not supported, except for vTPM encryption.

Next > [Register your VMware source](#) to protect it!

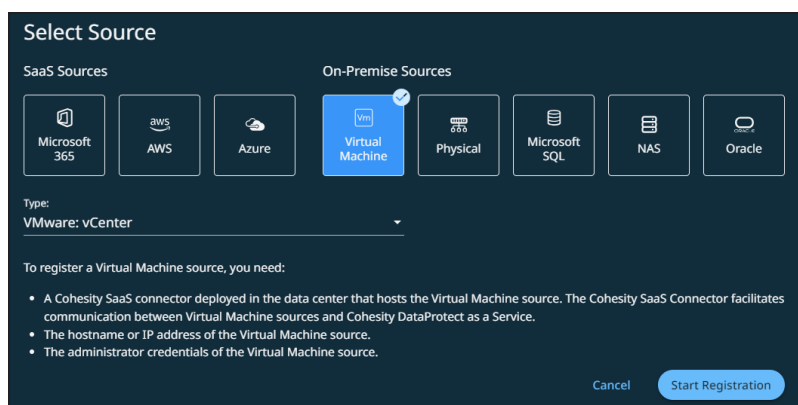
Register VMware Sources

To start protecting your VMware VMs, you need to register your data sources.

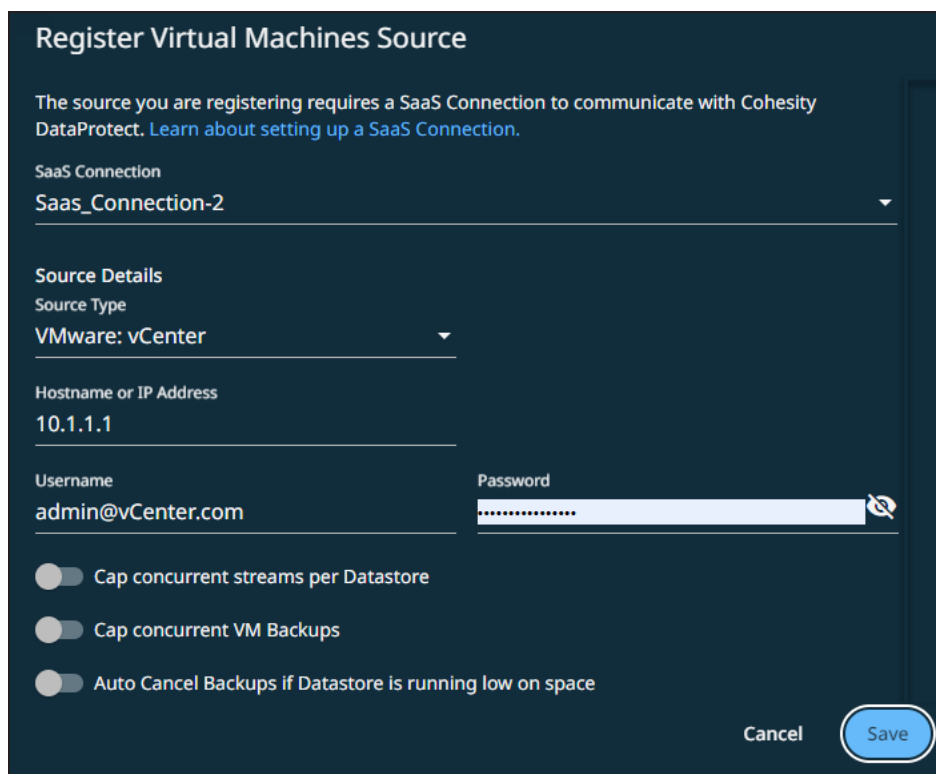
Note: To connect with sources in your data center, you'll need to use a SaaS Connection (or [create one](#)) to establish connectivity between the sources and the Cohesity DataProtect as a Service.

To register your VMware sources:

1. Ensure the [VMware requirements](#), such as Software versions, Firewall Ports, and the vCenter user account role privileges, are met.
2. Navigate to the **Sources** page and click **+ Register Source** in the upper-right corner of the page.
3. In the **Select Source** dialog box, select **Virtual Machine**.
4. In the **Type:** drop-down, select **VMware: vCenter** or **VMware: Standalone ESXi Host** and click **Start Registration**.



5. Select an existing SaaS connection marked as **Unused**, or click **Create SaaS Connection** and follow the instructions in [Create a SaaS Connection](#).



The dialog box is titled "Register Virtual Machines Source". It contains a message: "The source you are registering requires a SaaS Connection to communicate with Cohesity DataProtect. [Learn about setting up a SaaS Connection.](#)". Below this, there is a "SaaS Connection" dropdown menu with "Saas_Connection-2" selected. Under "Source Details", the "Source Type" is "VMware: vCenter". The "Hostname or IP Address" is "10.1.1.1". The "Username" is "admin@vCenter.com" and the "Password" is masked with dots. There are three toggle switches: "Cap concurrent streams per Datastore" (off), "Cap concurrent VM Backups" (off), and "Auto Cancel Backups if Datastore is running low on space" (off). At the bottom right are "Cancel" and "Save" buttons.

6. Click **Continue**.
7. Enter the **Hostname or IP Address**.
8. Enter the **Username** and **Password**.
9. Enter the following (optional):
 - a. **Cap concurrent streams per Datastore**. Turn this toggle ON and set the number of concurrent backup streams for each datastore **Number of Streams** field.

When registering a VMware vCenter, the value specified in the **Number of Streams** field will be applicable to all datastores within that vCenter. You can override this value for specific datastores after completing the VMware vCenter registration. When the registration is complete, Cohesity discovers all the datastores within the vCenter and displays them in the **Add Datastore Override** field. To override the global value, select a datastore, enable **Cap**, and enter the number of streams. If you disable **Cap** for a data store, the limit you set in the **Number of Streams** field will not be applied. The default limit set by Cohesity is applied instead.
 - b. **Cap concurrent VM Backups**. Turn this toggle ON, and the number of VMs concurrently backed up per vCenter in the **Number of VM Backups** field.

Note: This field is not applicable if you register an ESXi host.

- c. **Auto Cancel Backups if Datastore is running low on space.** Turn this toggle ON if you want to set the minimum free space that should always be available in the datastore. Once set, Cohesity automatically prevents backups from starting if the datastore does not meet this requirement.

Enter the minimum amount of free space required in either the **GiB** or **%** fields. For example, if you enter 10 in the **%** field, 10% of the datastore space is available at all times.

The percentage option is more flexible as it auto-adjusts when the datastore shrinks or expands, even if a VM has multiple data stores. For example, if two virtual VM disks reside in separate datastores of size 500 GB and 1 TB each, 10% of free space will apply to both the datastores.

- d. **Set Last Backup detail.** For VMware vCenter only: Toggle this setting ON to set the last backup time and backup status for a VM in vCenter. These are published as custom attributes. By default, the option is turned OFF. When a backup is performed on a vCenter with the feature enabled, the custom attributes "Last Cohesity Backup Attempt Time" and "Last Cohesity Backup Status" will be set or updated to reflect the backup time and status.

- **Last Cohesity Backup Attempt Time** - The exact timestamp when the last backup attempt succeeded.
- **Last Cohesity Backup Status** - Status of the last backup for each VM.

Note: When performing VM-based backups, if the backup is successful but has warnings caused by implicit fallbacks such as failure to utilize storage snapshots or issues with deleting the VM snapshot (which is a time-consuming operation), we will set the last backup status as 'Success' since the VM was successfully backed up. However, if there are errors or warnings due to explicit fallbacks (for example, transitioning from SAN to NBDSSL, or from app consistent to crash consistent), we will update the backup status to accurately reflect the error or warning.

10. Click **Complete**.

Next > You are now ready to [protect your VMs](#).

Protect VMware VMs

Once you have [registered vCenter Server or ESXi host](#) as sources, you're ready to protect the VMs on those ESXi hosts.

Cohesity supports the backup and recovery of NVRAM files when backing up VMware VMs.

The DataSets files are automatically protected. You can backup and recover the DataSets files when backing up VMware VMs.

Protection for NVRAM files

The NVRAM (non-volatile random-access memory) file of a VMware VM stores the boot drive setting and, in the case of multiple disks, the boot order. As VMware recommends, Cohesity automatically backs up the NVRAM file when backing up its VM. NVRAM backup and recovery is available for VMC.

Note: It is important to back up the NVRAM file if you created the VM with the Virtual Trusted Platform Module (vTPM).

Protection for vSphere DataSets files

The VM configuration files and vSphere DataSets files are stored together, and a reference to the DataSet (.dsd) is written into the virtual machine configuration file (.vmx). The datasets of the VM are stored within the DataSets file. The DataSets files are denoted with a .dsd extension and are stored alongside the VM configuration files. The maximum size of the DataSets file is 300 MB.

You can automatically back up the VM's DataSets files while protecting a VM using Cohesity clusters. You can also recover the DataSets files.

DataSets file functionality is available in vSphere 8.0 and requires the VM to be running virtual hardware version 20.

To protect a VMware source:

1. In **DataProtect as a Service**, under **Sources**, find the VMware source name and click into it.
2. Use the filters and search box at the top to narrow your search.
3. Use the checkboxes to select the objects for protection. To protect the whole source, click the checkbox above the column.
4. Click the **Protect** icon above the checkboxes.
5. In the **New Protection** dialog, select a **Policy** that matches the schedule and retention period you need. If the existing policies do not meet your needs, you can [create a new policy](#) with the settings you need.
6. To change or configure any of the additional settings, select **More Options** and perform the below steps or else, click **Protect**.

7. In the **Start Time** field, enter the time the protection run should start. The default time zone is the browser's time zone. You can change the time zone of the job by selecting a different time zone.
8. In the **SLA** field, define how long the administrator expects a protection run to take. Enter:
 - **Full.** The number of minutes you expect a full protection run, which captures all the blocks in an object, to take.
 - **Incremental.** The number of minutes you expect an incremental protection run, which captures only the changed blocks in an object, to take.
9. If you need to change any of the additional settings, click the down arrow icon next to **Additional Settings** and click **Edit**.
10. Click **Protect**.

Cohesity DataProtect as a Service starts backing up the VMs you selected. You can monitor the status of the backup on the **Activity** page.

Note: The backups start immediately after you protect the objects, regardless of the time you set for the protection run.

Also, the **Activity** tab of a specific VM instance shows the history of all protection runs, including the one in progress.

Additional Settings

Advance Settings	Description
End Date	If you need to end protection on a specific date, enable this to select the date.
Exclusions	Enable Exclude Disks to select the disks to exclude for all VMs in this object's protection. Enter the Controller Type , Controller Bus Number , and Unit Number for each disk to exclude. Excluded disks are not backed up and are not recovered during VM recovery.

Advance Settings	Description
App Consistent Backups	<p>Enable App Consistent backups if you want the guest operating systems of all the protected VMs to be quiesced before snapshots of these VMs are created. Quiescing of VMs prior to capturing snapshots ensures the integrity of the data saved in the snapshots.</p> <p>With the App Consistent backups enabled, the following options are available:</p> <ul style="list-style-type: none"> • Take a Crash Consistent backup if unable to perform an App Consistent backup. Enable this option if you want Cohesity to capture a crash-consistent snapshot if Cohesity fails to capture an app-consistent snapshot. If this option is disabled and Cohesity is unable to perform an app-consistent backup of a VM, a snapshot is not captured. <p>Note: This option is applicable only for VSS copy backup.</p>
Cancel Runs at Quiet Time Start	<p><i>(Available only if the selected policy has at least one Quiet Time.)</i></p> <p>When enabled, all the protection runs that are currently executing will cancel when the Quiet Time period starts. By default, this setting is disabled, meaning that after a protection run starts, it continues to execute even when a Quiet Time period starts. However, new protection runs will not start during a Quiet Time.</p>

Next > When the first protection run completes, you will be ready to [recover the protected VMs and files](#) when and if you need to.

Recover VMware VMs & Files

After you [protect your VMware sources](#), you can recover VMs and files from your backups, to their original or a new location.

Recover VMware VMs or Files

To recover VMware VMs or files:

1. In **DataProtect as a Service**, go to **Sources** to set up your recovery task.
2. Click on the **Source** name.
3. Above the tree, select **Protection Status > Protected**.
4. Use the filters, search box, and views to locate the objects or files you need.
5. To recover:
 - **VMs**, continue with the [Recover Objects & Volumes](#) procedure in [Recover Protected Objects & Files](#).

- **Files and folders**, continue with the [Recover Files & Folders](#) procedure in [Recover Protected Objects & Files](#).
6. Select your [Recovery Options](#) and click **Start Recovery**.

Note: If you are recovering a VM to the original location and enable **Overwrite Existing VM**, you can choose to take advantage of **Attempt Differential Recovery** to shorten your recovery time, after considering the [implications](#) below.

Cohesity DataProtect as a Service begins to restore the selected VMs or files to the selected location.

Accelerate VM Recoveries with Differential Restore

In Cohesity DataProtect as a Service, you can recover the VM by overwriting only the difference between the original VM and the snapshot selected for recovery. This option is available only if you have selected to recover to the original location and enabled **Overwrite Existing VM** in the [VM recovery options](#) in your recovery task.

Differential recovery substantially reduces the amount of data transfer in a recovery process. In the task activity log (under **Activity**), you can view the amount of data transfer saved by selecting differential recovery.

However, there are several important implications to consider before choosing to **Attempt Differential Recovery**:

- Any newly added data in the original VM is deleted.
- The recovered VM will have the existing VM name.
- You can choose this option if there are no hardware configuration changes involved in the original VM.
- If you want to reclaim free space for thin-provisioned disks, then Cohesity recommends not to attempt differential recovery and only perform **Overwrite Existing VM** recovery.
- If the original VM is not present or if the attempt at differential recovery fails, then Cohesity will perform an **Overwrite Existing VM** recovery.
- In the original VM, if there are any newly added disks or any disks that were excluded during backup, then the recovered VM will not have these newly added disks, nor any disks excluded during backup.
- All the snapshots present on the original VM are consolidated and removed as part of differential recovery.

Hyper-V

Cohesity DataProtect as a Service provides a simple, fast, cost-effective backup, recovery, and data management solution for Hyper-V environments.

Hyper-V Requirements

To register your Hyper-V sources, ensure you meet the requirements and install the Cohesity Agents on your SCVMM server and Standalone Hyper-V hosts.

Before you register your Hyper-V sources, confirm that you meet the software version, [firewall](#), and [permissions](#) requirements below, [install the Cohesity Agent on your SCVMM server](#), and then [install it on your Standalone Hyper-V hosts](#).

Also, be sure to review the [best practice](#) recommended below.

Note: For information on the supported cloud regions where you can back up this source, see [Supported Workloads and Cloud Regions](#).

Support Matrix

Before you register your Hypervisors, ensure that you have the supported Hypervisor versions. For more information, see [Supported Software for Cohesity DataProtect as a Service](#).

Firewall Ports

Ensure that the ports listed in the Microsoft SCVMM and Hyper-V Servers section in the [Firewall Ports for User-Deployed SaaS Connectors](#) topic are open to allow communication between the Cohesity SaaS Connector(s) and Hyper-V environment.

Minimum Permissions

To be able to register your Hyper-V SCVMM (System Center Virtual Machine Manager) server and Standalone Hyper-V hosts as sources, you need to first install the Cohesity Agent on that source. To install the Cohesity Agent, you can use either the LOCAL SYSTEM account or a domain user with administrative privileges on the SCVMM application.

For Hyper-V standalone clusters, add:

1. All hosts' machine accounts:
 - Start the **Failover Cluster Manager**.
 - From the clusters list, right-click the standalone cluster and select **Properties > Cluster Permissions > Add > Object Types > Check Computers > OK**.
 - Type `<hostname>$` in **"Enter the object names to select"** and select **Check Names > OK > Allow "Read" > OK**.

2. Add all machine accounts to the Administrators group of each host in the standalone cluster.

Download and Install the Cohesity Agent on Your SCVMM Server

Before you can register your SCVMM, you need to install the Cohesity Agent on the SCVMM server, or on an existing proxy endpoint that is connected to the SCVMM server.

To install the Cohesity Agent on your SCVMM server:

1. In **DataProtect as a Service**, navigate to **Sources**, and select **+Register Source> Virtual Machines**.
2. Select **HyperV: SCVMM Server** as the **Source Type**.
3. Click **Download Cohesity Agent**. Ensure the Agent has been downloaded to the appropriate SCVMM server.
4. As an administrator with local system privileges, run the executable and complete the installation wizard. Install the Agent without additional components.

The Agent starts automatically. Next, you'll need to install the Agent on the Standalone Hyper-V hosts that you plan to protect.

Download and Install the Cohesity Agent on Your Standalone Hyper-V Hosts

Now install the Cohesity Agent on the Standalone Hyper-V hosts that you want to protect.

To install the Cohesity Agent on your Standalone Hyper-V hosts:

1. In **DataProtect as a Service**, navigate to **Sources**, and select **+Register Source> Virtual Machines**.
2. Select **HyperV: Standalone Host** as the **Source Type**.
3. Click **Download Cohesity Agent**. Ensure the Agent has been downloaded to the appropriate Standalone Hyper-V hosts.
4. As an administrator with local system privileges, run the executable and complete the installation wizard on each host. Install the Agent without additional components.

The Agent starts automatically.

Note: The minimum recommended specification for Guest Windows VMs is: 2 GB RAM and the equivalent of a 1 GHz processor.

Best Practice

For Hyper-V 2016 and 2019, configure all VMs' **Automatic Stop Action** to shut down or turn off, instead of save. This results in all powered-on VMs having minimal size `.vmrs` files. VMs in the saved state, by contrast, generally have `.vmrs` files greater than 10 MB. Though Cohesity supports the backup of `.vmrs` files greater than 10 MB, we recommend that you back up `.vmrs` files with minimal size.

Next > Register your SCVMM server and Standalone Hyper-V hosts!

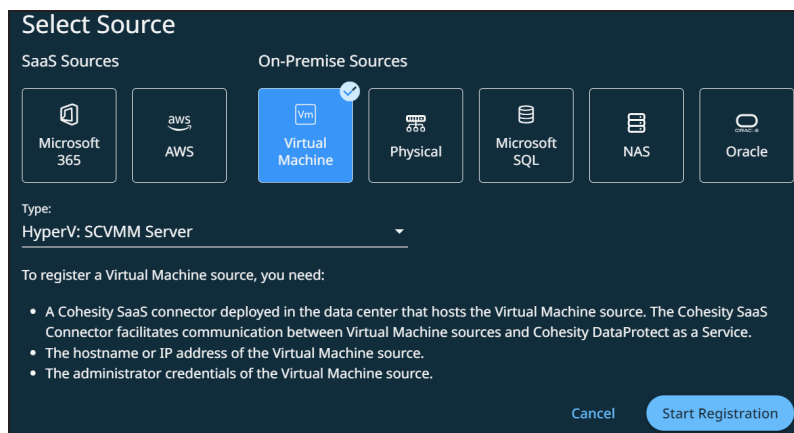
Register Hyper-V Sources

To start protecting your Hyper-V VMs, you need to register your SCVMM server and Standalone Hyper-V hosts as Cohesity DataProtect as a Service sources.

Note: To connect with Hyper-V sources in your data center, you'll need to use a SaaS Connection (or [create one](#)) to establish connectivity between the sources and the Cohesity DataProtect as a Service.

To register your Hyper-V sources:

1. Ensure the [Hyper-V requirements](#), such as Software versions, Firewall Ports, and the vCenter user account role privileges, are met.
2. In **DataProtect as a Service**, navigate to the **Sources** page and click **+ Register Source** in the upper-right corner of the page.
3. In the **Select Source** dialog box, select **Virtual Machine**.
4. From the **Type** drop-down, select one of following and click **Start Registration**:
 - **HyperV: SCVMM Server**
 - **HyperV: Standalone Host**
 - **HyperV: Failover Cluster**



5. Select an existing SaaS connection marked **Unused**, or click **Create SaaS Connection** and follow the instructions in [Create a SaaS Connection](#).
6. Enter the **Hostname or IP Address**.
7. Enter the **Username** and **Password**.
8. Click **Complete**.

Note:

If you are planning to edit the registered source (**Actions** menu (:) > **Edit**) for moving the SCVMM proxy agent endpoint to a different proxy or to the SCVMM cluster, or for moving from the SCVMM cluster to a proxy, then you must also:

1. Copy the old agent registry values from the Cohesity folder and subfolders.
2. When moving to:
 - A proxy, paste the registry values to the new endpoint.
 - SCVMM, paste the registry values to the active SCVMM node. (RDP'ing into the SCVMM cluster redirects to the active master node.)

Best Practices

- Allocate at least 4 CPUs and 10 GB RAM is recommended for your Hyper-V SaaS Connector.
- Deploy your Hyper-V SaaS Connector VMs onto failover clusters in a highly available manner.
- Ensure your Hyper-V SaaS Connector VMs do not contain stateful data that would require backup and restore upon recovery. In case of disaster, simply [deploying a new Hyper-V SaaS Connector VM](#) is enough. It is unnecessary to back up SaaS Connectors, and doing so can degrade performance.
- Create copies of the golden VHD for multiple SaaS Connectors. Do not create differencing disks on top of a SaaS Connector.
- You need only a single networking switch; additional networking switches will not be consumed.
- Cohesity DataProtect as a Service supports both Hyper-V Generation 1 and 2. Select the one that best meets your internal best practices.
- You can convert from VHD to VHDX if you prefer. This, again, depends on your organization's internal best practices.

Next > You're ready to [protect your Hyper-V VMs](#).

Protect Hyper-V VMs

Once you have [registered your SCVMM server and Standalone Hyper-V hosts as sources](#), you're ready to protect the VMs on those Hyper-V hosts.

To protect your Hyper-V VMs:

1. In **DataProtect as a Service**, under **Sources**, find the Hyper-V source name and click into it.
2. Use the filters and search box at the top to narrow your search.
3. Use the checkboxes to select the objects for protection. To protect the whole source, click the checkbox above the column.
4. Click the **Protect** icon above the checkboxes.
5. In the **New Protection** dialog, select a **Policy** that matches the schedule and retention period you need. If the existing policies do not meet your needs, you can [create a new policy](#) with the settings you need.
6. To change or configure any of the additional settings , select **More Options** and perform the below steps or else, click **Protect**.
7. In the **Start Time** field, enter the time the protection run should start. The default time zone is the browser's time zone. You can change the time zone of the job by selecting a different time zone.
8. In the **SLA** field, define how long the administrator expects a protection run to take. Enter:
 - **Full**. The number of minutes you expect a full protection run, which captures all the blocks in an object, to take.
 - **Incremental**. The number of minutes you expect an incremental protection run, which captures only the changed blocks in an object, to take.
9. If you need to change any of the additional settings, click the down arrow icon next to [Additional Settings](#) and click **Edit**.
10. Click **Protect**.

Cohesity DataProtect as a Service starts backing up the Hyper-V VMs you selected. You can monitor the status of the backup on the **Activity** page.

Note: The backups start immediately after you protect the objects, regardless of the time you set for the protection run.

Also, the **Activity** tab of a specific Hyper-V VM instance shows the history of all protection runs, including the one in progress.

Additional Settings

Advance Settings	Description
End Date	If you need to end protection on a specific date, enable this to select the date.

Advance Settings	Description
App Consistent Backups	<p>Enable App Consistent backups if you want the guest operating systems of all the protected VMs to be quiesced before snapshots of these VMs are created. Quiescing of VMs prior to capturing snapshots ensures the integrity of the data saved in the snapshots.</p> <p>With the App Consistent backups enabled, the following options are available:</p> <ul style="list-style-type: none"> • Take a Crash Consistent backup if unable to perform an App Consistent backup. Enable this option if you want Cohesity to capture a crash-consistent snapshot if Cohesity fails to capture an app-consistent snapshot. If this option is disabled and Cohesity is unable to perform an app-consistent backup of a VM, a snapshot is not captured. • Backup application data and truncate their log files. Enable this option if you want to back up applications (Microsoft SQL Server, Exchange Server) that are running on the Hyper-V server and truncate the logs of applications. <p>Note: This option is applicable only for VSS copy backup.</p>
Cancel Runs at Quiet Time Start	<p><i>(Available only if the selected policy has at least one Quiet Time.)</i></p> <p>When enabled, all the protection runs that are currently executing will cancel when the Quiet Time period starts. By default, this setting is disabled, meaning that after a protection run starts, it continues to execute even when a Quiet Time period starts. However, new protection runs will not start during a Quiet Time.</p>

Next > When the first protection run completes, you will be ready to [recover files from your protected Hyper-V VMs](#) if and when you need to.

Recover Hyper-V VMs & Files

Once you have [protected your Hyper-V VMs](#), you can [recover Hyper-V VMs or files](#), to their original or a new location.

Prerequisite

Before recovering files to a target VM, depending on the guest OS system on the target VM, be sure to [install the Cohesity Windows or Linux Agent](#) on the target VM.

Recover Hyper-V VMs

To recover Hyper-V VMs from your protected Hyper-V VMs:

1. In **DataProtect as a Service**, go to **Sources** to set up your recovery task.
2. Click on the **Source** name.

3. Above the tree, select **Protection Status > Protected**.
4. Use the filters, search box, and views to locate and select the Hyper-V source you want to recover from.

Tip: You can also use Global Search to locate, filter, and select the objects you need. Click the **Global Search** box at the top or type **slash (/)** anywhere to start your search.

5. Locate and select the Hyper-V VMs you need, and then click **Recover** at the top to open the **New Recovery** form with the **Latest** snapshot (protection run).

Note: If you do not see the option to recover VMs from your Hyper-V backups, please contact [Cohesity Support](#) to request it.

6. If you need to recover from an earlier snapshot, click the **Edit** icon to select a new recovery point.
 - For each VM under **Selected**, you can click the **Edit** icon to open the **Recovery Point** calendar. Click **List** to view the available recovery points by timestamp and click one.
 - Click **Select Recovery Point**.
 - Click **Next: Recover Options** to return to the form.
7. Under **Recover To**, select **Original Location** or **New Location**.
 - If you choose **New Location**, select a **Registered Source, Resource Pool, Datastores**, and the **VM Folder**.
8. Select your [Recovery Options](#) (for object recovery).
9. Click **Start Recovery**.

Cohesity DataProtect as a Service opens the **Activity** page, showing your file recovery task as it runs, along with the recovery progress on the right.

Tip: The **Activity** page also shows the entire history of all protection runs and recovery tasks, including any that are in progress.

Recover Hyper-V Files

To recover Hyper-V files from your protected Hyper-V VMs:

1. In **DataProtect as a Service**, go to **Sources** to set up your recovery task.
2. Click into the **Source** name.

3. Above the tree, select **Protection Status > Protected**.
4. Use the filters, search box, and views to locate and select the Hyper-V source you want to recover from.

Tip: You can also use Global Search to locate, filter, and select the objects you need. Click the **Global Search** box at the top or type **slash (/)** anywhere to start your search.

5. Locate the source object containing the files you want to recover and click the **Recover Files** icon on that row.

By default, the latest snapshot is selected for recovery. To recover from a different snapshot, click the snapshots drop-down in the top-right corner and select the snapshot you need.

Note: Changing the snapshot after selecting the items (files or folder) removes the selected items from the cart.

6. Browse to the file or folder that you want to recover by clicking folders and their subfolders.
7. Select the files to recover and choose one of the following options:
 - **Next.** If you select this option, then continue to the next step to configure the file recovery options.
 - **Download Files.** If you are recovering a single file, this option downloads the file to your browser's download folder. For all other selections, this creates a recovery task. When the task completes, from the **Activity** page, click the task name and then click **Download Files** to download the generated zip file.
8. Under **Recover To**, select **Original Server** or **New Server**. For:
 - **Original Server**, by default, the files are received to the original path. If you want to recover to an alternate path, then toggle off **Recover to Original Path** and enter the path. The default alternate path is **/tmp/Recover-*<date_time>***.
 - **New Server**, select a **Registered Source**. You also have the option to register a new source. Select the target VM, username, and password. By default, the files are recovered to the **/tmp** directory, but you can provide a different directory if needed.
9. Select your file [Recovery Options](#) and click **Recover**.

Cohesity DataProtect as a Service opens the **Activity** page, showing your file recovery task as it runs, along with the recovery progress on the right.

Tip: The **Activity** page also shows the entire history of all protection runs and recovery tasks, including any that are in progress.

VMware Cloud (VMC) on AWS

Cohesity DataProtect as a Service provides a simple, fast, cost-effective backup, recovery, and data management solution for VMware Cloud (VMC) on AWS environments.

To protect VMC on AWS using Cohesity DataProtect as a Service:

1. Ensure the [requirements](#) for VMC on AWS are met.
2. [Register your VMC on AWS SDDC \(software-defined data center\)](#) as a source in Cohesity DataProtect as a Service.
3. [VMC on AWS source](#).
4. [Recover VMs or files & folders](#).

VMC on AWS Requirements

Before [registering your VMC on AWS vCenter as a data source](#), ensure that it meets the software versions, user privilege, and firewall port requirements.

Note: For information on the supported cloud regions where you can back up this source, see [Supported Workloads and Cloud Regions](#).

Supported Software

SDDC Version 1.22 is supported.

User Privileges for VMware Cloud Sources

To protect your VMC on AWS SDDC, you need a user account with the CloudAdmin role. For information on the privileges of a CloudAdmin, see [CloudAdmin Privileges](#).

Firewall Ports

Ensure the [firewall ports](#) are open to allow communication between the Cohesity SaaS Connector(s) and the VMC in the AWS environment.

Considerations

- Protection of the following is not supported:
 - VMs on NFS datastores.
 - VMs with disks on the IDE controller.
- Recovery from on-premises vCenter to VMC on AWS SDDC or VMC on AWS SDDC to on-premises vCenter is not supported.
- We recommend that you have one SaaS Connector for each ESXi host.

Note: These requirements are subject to change.

Next> [Register VMware Cloud on AWS Source.](#)

Register VMware Cloud on AWS Source

To start protecting your VMware Cloud on AWS SDDC (software-defined data center), you need to register it as a source on Cohesity DataProtect as a Service.

Before you Begin

- Create a VMware SaaS Connection in your VMC environment (using your VMware Cloud Services console) to establish connectivity between your sources and the Cohesity DataProtect as a Service. For details, see [Deploy VMware SaaS Connectors](#).
- Ensure that the [requirements](#) for the VMware Cloud on AWS source, such as software version, user account role privileges, and firewall ports, are met.

Register

To register your VMware Cloud on AWS source:

1. In **DataProtect as a Service**, navigate to **Sources** and click **+ Register Source > Virtual Machines**.
2. In the **Register Virtual Machines Source** form, select an existing SaaS connection from the **SaaS Connection** drop-down, or click **Create SaaS Connection** and follow the instructions in the [Deploy VMware SaaS Connectors](#) section.

3. Select the Source Type as **VMware: vCenter**.
4. Enter the **Hostname** or **IP Address** of the VMware Cloud Services console.
5. Enter the **Username** and **Password** you use to log into the VMware Cloud Services console.
6. Click **Save**.

The VMware Cloud Services source is registered and displayed under the **Sources** page. For information on managing your VMC on AWS source, see [Manage the VMC on AWS Source](#).

Create SaaS Connector Groups

You must create at least one SaaS Connector group for each ESXi cluster that needs to be protected if there is more than one ESXi cluster in the SDDC. When creating the SaaS Connector Groups:

1. [Group the SaaS Connectors at each ESXi cluster](#) into two separate SaaS Connector Groups.
2. [Associate the Connector Groups with the respective ESXi clusters](#).

For example, to protect two ESXi clusters (EC-1 and EC-2) in your SDDC, add the SaaS Connectors of EC-1 to SaaS Connector Group CG-1 and EC-2 to SaaS Connector Group CG-2. Then associate CG1 to EC-1 and CG2 to EC-2.

Note: If there is only one ESXi cluster in the SDDC and no plans to add more ESXi clusters, SaaS Connector Groups are not required.

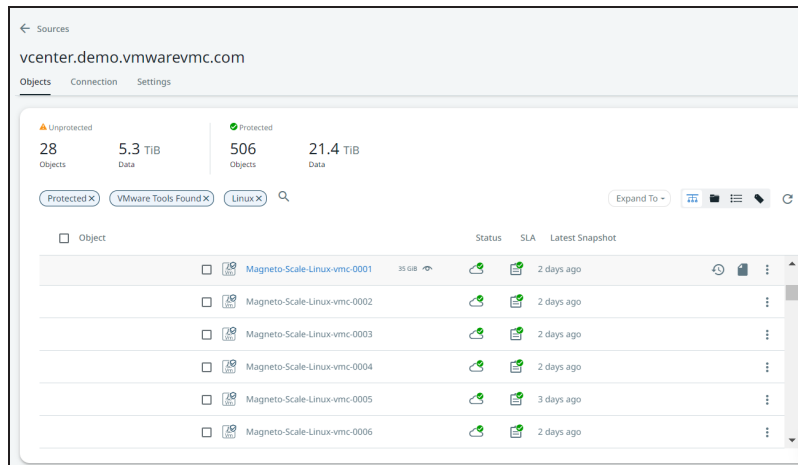
Next > You are now ready to [protect your VMC on AWS source](#).

Manage the VMC on AWS Source






View the Registered VMC on AWS Source

To view a registered VMC on AWS Source:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. The **Sources** page displays all the sources registered with Cohesity DataProtect as a Service. Click the Actions menu (⋮) of the VMware Cloud source to [Edit](#), [Unregister](#), or [Manage Network Traffic](#).



3. Click on a source. The details of the source are displayed in the **Objects**, **Connection**, and **Settings** tab. The **Objects** tab displays the hierarchy of the data centers, ESXi hosts, clusters, folders, and VMs in that VMC on AWS source.

- A selected check box (☒) indicates the object is selected to be backed up.
-  - Indicates a Windows VM.
-  - Indicates a Linux VM.
-  - A blue shield indicates a protected VM.
-  - Indicates an excluded object that is not protected. An ancestor of the object is auto-protected, but this object is excluded through inheritance (an ancestor is explicitly excluded).
-  - Click to view the general information—type, server size, and version of the selected object.

Filters on the Objects tab help you display only the objects you want. The filtering options are:

- **Protection Status.** Protected or Unprotected.
- **VM Attributes.** VMware Tools Found and Cohesity SaaS Connector.
- **Host Type.** Linux or Windows.

Protection Status ▾

VM Attributes ▾


Host Type ▾

Enter the name of an object in the **Search** box. You can even enter a partial name or an asterisk * wildcard. As you type, the objects that contain your search term appear.

Search

🔍

Display modes can show the objects in a hierarchy or a flat list. The object hierarchy is automatically refreshed every four hours. To manually refresh the object hierarchy, select **Data Protection > Sources**, find the source in the list, and click the icon.



View the Connection Status

To view the connection status of a VMC on AWS Source:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. On the **Sources** page, select the **Connection** tab.









The **Connection** tab displays the status of the connection between the VMC on AWS source and SaaS Connection and between the SaaS Connection and the Cohesity Cloud Service in the following ways:

- The **Glance bar** displays the performance of the SaaS connection with respect to the attributes such as CPU, memory, and IOPS (Read and Write).

Average (last 24 hours)			
3%	29%	15.81	0.31
CPU	Memory	IOPS (Read)	IOPS (Write)

- The **Connection Activity** graph displays the read/write activity on the DataSource in the past 24 hours.
- The **Connector Details** section displays the connectivity status and other details of

the individual SaaS connectors.

Connector Details							
Connector	IP	Status	Release	CPU	Memory	IOPS (Read)	IOPS (Write)
 SaaS-Connector	192.1.1.0	 Healthy	hedp2023may	12%	34%	-	4.22
 SaaS-Connector	192.1.1.1	 Healthy	hedp2023may	18%	26%	-	11.14
 SaaS-Connector	192.2.2.2	 Healthy	hedp2023may	44%	27%	-	12.36
 SaaS-Connector	192.3.3.3	 Healthy	hedp2023may	20%	33%	-	11.82

View the Details of the Source

To view the connection status of a VMC on AWS Source:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. On the **Sources** page, select the **Settings** tab.

The **Settings** tab displays the details of the VMC on AWS source, such as the Username, Registered Date, Refreshed Date, and so on.

vcenter.sddc-demo.vmwarevmc.com	
Objects	Connection
Settings	
Type	VMware Cloud vCenter
IP or FQDN	10.1.111.1
Username	cloudadmin@vmc.local
CA Certificate	No
Registered	Jun 7, 2023 6:18pm
Refreshed	Jun 8, 2023 7:34am
Use VM Bios UUID	No
Cap concurrent streams per Datastore	Uncapped
Auto Cancel Backups if Datastore is running low on space	Off
Cap concurrent VM Backups	Uncapped
Detect VM migration across vCenter to preserve backup chain	No

Update the VMC on AWS Source Configuration

You can update the VMC on AWS configuration you provided during the registration process with the latest configuration. To edit the VMC on AWS source configuration:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the Actions menu (:) next to the VMC on AWS source and select **Edit**.
3. In the **Edit Source** page, update the respective configurations.
4. Click **Save**.

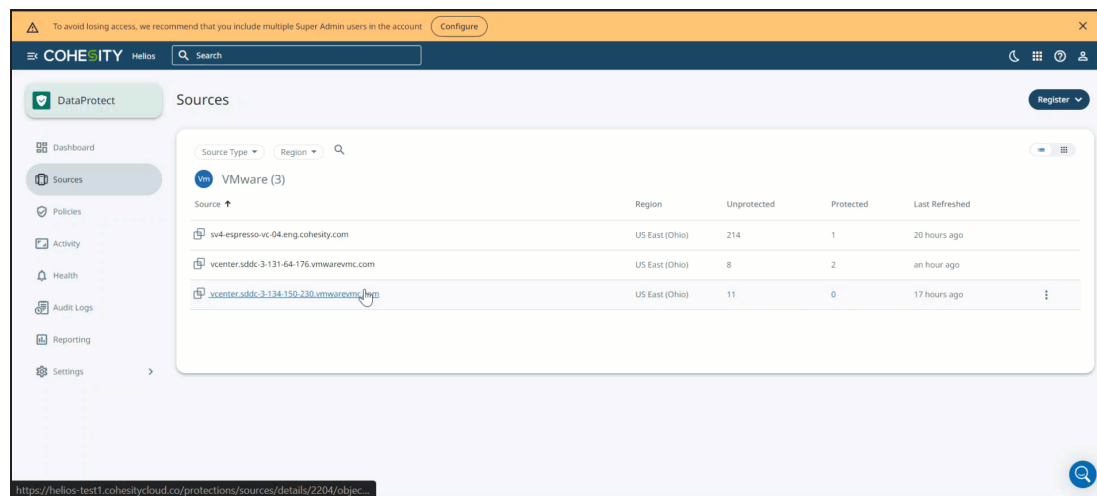
Unregister VMC on AWS Source

If you plan to stop protecting of your VMC on AWS source, you can unregister the VMC on AWS source from the Cohesity cluster. To unregister the VMC on AWS source:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the Actions menu (⋮) next to the VMC on AWS source and select **Unregister**.


Protect VMC on AWS

Once you have [registered your VMC on AWS source](#), you're ready to protect the VMs in your VMC on AWS SDDC (software-defined data center).



To protect a VMware source:

1. In **DataProtect as a Service**, under **Sources**, find the VMC on AWS source name and click on it.
2. Use the filters and search box at the top to narrow your search.
3. Use the checkboxes to select the objects for protection. To protect the whole source, click the checkbox above the column.
4. Click the **Protect** button above the checkboxes.
5. In the **New Protection** dialog, select a **Policy** that matches the schedule and retention period you need. If the existing policies do not meet your needs, you can [create a new policy](#) with the settings you need.
6. To change or configure any of the additional settings, select **More Options** and perform the below steps, or else, click **Protect**.
7. In the **Start Time** field, enter the time the protection run should start. The default time zone is the browser's time zone. You can change the time zone of the job by selecting a different time zone.

8. In the **SLA** field, enter the expected completion time for a protection run using the following options:
 - **Full.** Enter the duration (in minutes) within which you expect a complete protection run to finish running. A full protection run captures all the blocks in an object.
 - **Incremental.** Enter the duration (in minutes) within which you expect an incremental protection run to finish running. An incremental protection run captures only the changed blocks in an object.
9. If you need to change any additional settings, click the down arrow icon next to **Additional Settings** and click **Edit** icon () next to the required field.
10. Click **Protect**.

Cohesity DataProtect as a Service immediately starts backing up the VMs you selected. You can monitor the status of the backup on the **Activity** page. Also, the **Activity** tab of a specific VM instance shows the history of all protection runs, including the one in progress.

Additional Settings

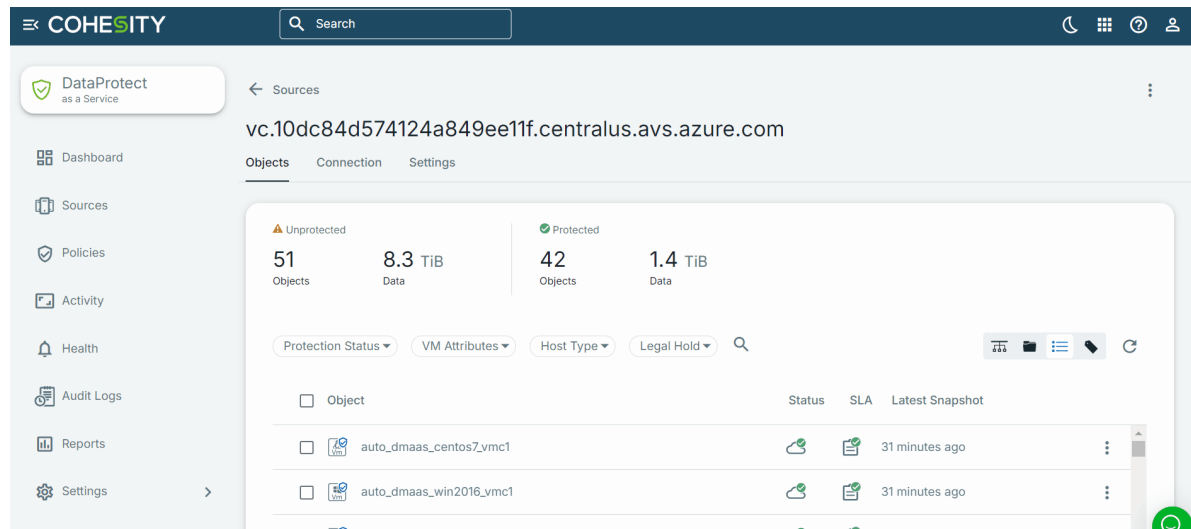
Advance Settings	Description
End Date	If you need to end protection on a specific date, enable this option and select the date.
Exclusions	Enable Exclude Disks to select the disks that are to be excluded for all VMs in the object's protection. Enter the Controller Type , Controller Bus Number , and Unit Number for each disk to exclude. Excluded disks are not backed up and are not recovered during VM recovery.

Advance Settings	Description
App Consistent Backups	<p>Note: This option is applicable only if VMware Tools are running in that VM. If VMware Tools are not found in the VM, Cohesity displays a warning, and selecting this option will be ineffective.</p> <p>Enable App Consistent backups if you want the guest operating systems of all the protected VMs to be quiesced before snapshots of these VMs are created. Quiescing of VMs prior to capturing snapshots ensures the integrity of the data saved in the snapshots.</p> <p>With the App Consistent backups enabled, the following option is available:</p> <p>Take a Crash Consistent backup if unable to perform an App Consistent backup. Enable this option if you want Cohesity to capture a crash-consistent snapshot if Cohesity fails to capture an app-consistent snapshot. If this option is disabled and Cohesity cannot perform an app-consistent VM backup, a snapshot is not captured.</p>
Cancel Runs at Quiet Time Start	<p>Note: Field is displayed only if the selected policy has at least one Quiet Time defined. By default, this option is disabled.</p> <p>When disabled, all in-progress protection runs continue to execute even when the Quiet Time starts. However, a new protection run will not start during the Quiet Time. All in-progress protection runs will abort (or pause based on your selection) once the Quiet Time starts when this option is enabled.</p>

Next > When the first protection run completes, you will be ready to [recover the protected VMs and files](#) if needed.

Recover VMC on AWS Sources

After you protect your VMs in the AVS environment, you can recover VMs and files from your backups to their original or new location.



To recover VMware VMs or files:

1. In **DataProtect as a Service**, go to **Sources** to set up your recovery task.
2. Click on the **Source** name.
3. Above the tree, select **Protection Status > Protected**.
4. Use the filters, search box, and views to locate the objects or files you need.
5. To recover:
 - **VMs**, continue with the [Recover Objects & Volumes](#) procedure in [Recover Protected Objects & Files](#).
 - **Files and folders**, continue with the [Recover Files & Folders](#) procedure in [Recover Protected Objects & Files](#).
6. Select your [Recovery Options](#) and click **Recover**.

Cohesity DataProtect as a Service restores the selected VMs or files to the selected location.

Troubleshooting

The following sections discuss possible scenarios for recovering protection runs from a failed state.

Disk Consolidation Failure

Cause

If you are protecting VMs in the VMC environment, Cohesity will use HotAdd transport mode for backup and recovery. During the backup or recovery operations, virtual disks might be left behind on the SaaS Connector/Hybrid Extender due to environmental issues. As a result, the following errors might occur in the VMware vCenter:

- VMware might not process the command to release the virtual machines' disk from the SaaS Connector or Hybrid Extender VM.
- The VM requires disk consolidation, but consolidation fails.

Solution

For the steps to resolve the error, see the [KB article](#).

Note: You can manually test if disks can be HotAdded to the SaaS connector or Hybrid Extender VM. For more details, see [KB article](#).

No Valid Bifrost is Available Error

Backup or recovery fails with the “no valid Bifrost is available” error.

Cause

All the SaaS Connectors in the Connector Group are inoperable due to misconfiguration, manual deletion, or other potential reasons.

Solution

Follow these steps to resolve the error:

1. Remove the inoperable SaaS Connector from the Connector Group. For details, see [Create Connector Groups](#).
2. Add the healthy SaaS Connectors to the Connector Group. For details, see [Manage Connector Groups](#).

Unable to Pick a SaaS Connector Error

Backup or recovery fails with the “Unable to pick a SaaS Connector for recovering VM to VMC source”.

Cause

There is more than one ESXi cluster in the VMC on AWS SDDC (software-defined data center), but you have not created SaaS Connector groups for the ESXi cluster.

Solution

Create SaaS Connector Groups for each ESXi cluster that needs to be protected. For details, see [Create SaaS Connector Groups](#).

Azure VMware Solution (AVS)

Cohesity DataProtect as a Service provides a simple, fast, cost-effective backup, recovery, and data management solution for VMs on AVS environments.

To protect VMs on AVS using Cohesity DataProtect as a Service:

1. Ensure the [requirements](#) for VMs on AVS are met.
2. [Register your AVS source](#) in Cohesity DataProtect as a Service.
3. [Protect VMs on AVS source](#).
4. [Recover VMs or files and folders](#).

VMs on AVS Requirements

Before [registering your VMs on AVS source](#), ensure that they meet the user privileges and firewall port requirements.

User Privileges for VMware Cloud Sources

To protect your VMs on the AVS solution software-defined data center (SDDC), you need a user account with the [Add User Privileges for vCenter Sources](#).

Firewall Ports

Ensure the [firewall ports](#) are open to allow communication between the Cohesity SaaS Connector(s) and the VMs in the AVS environment.

Next> [Register VMs on AVS Source](#).

Register VMs on AVS Source

To start protecting your VMs on Azure VMware solution software-defined data center (SDDC), register it as a source on Cohesity DataProtect as a Service.

Before you Begin

- Create a VMware SaaS Connection in your AVS environment (using your AVS console) to establish connectivity between your AVS source and the Cohesity DataProtect as a Service. For more information, see [Deploy VMware SaaS Connectors](#).
- Ensure that the [requirements](#) for the AVS source, such as user account role privileges and firewall ports are met.

Register

To register your AVS source:

1. In **DataProtect as a Service**, navigate to **Sources** and click **+ Register Source > Virtual Machines**.
2. Select the Source Type as **VMware: vCenter** and click the **Start Registration** button.

3. Select an existing SaaS connection from the **SaaS Connection** drop-down, or click **Create SaaS Connection** and follow the instructions in the [Deploy VMware SaaS Connectors](#) section.

Register VMware vCenter ✓ — 2 Source Details — 3

Hostname or IP Address
84d574124a849ee11f.centralus.avs.azure.com

Username
cloudadmin@vsphere.local

Password
.....

Back Cancel Continue

4. Enter the **Hostname** or **IP Address** of the AVS Cloud Services console.
5. Enter the **Username** and **Password** you use to log into the AVS console.
6. Click **Continue** and complete the source registration.

The AVS source is registered and displayed under the **Sources** page. For information on managing your VMs on AVS source, see [Manage VMs on AVS Source](#).

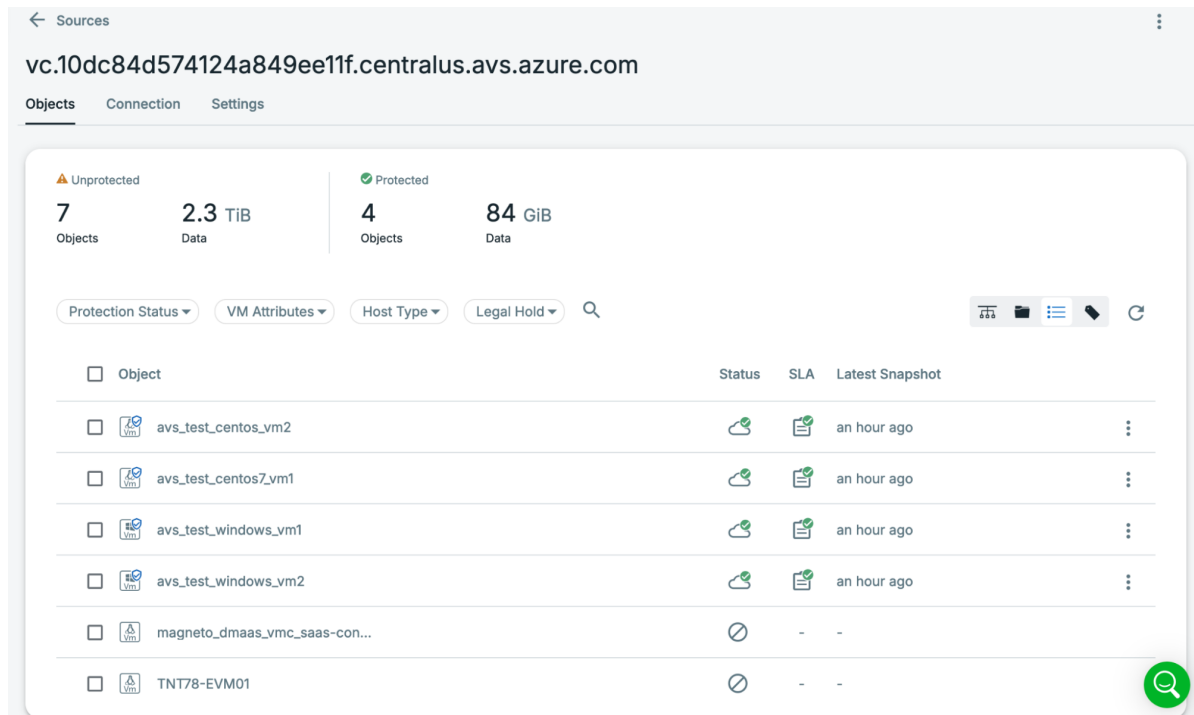
Next > You are now ready to [protect your VMs on AVS source](#).

Manage VMs on AVS Source







[View the Registered VMs on AVS Source](#)

To view the registered VMs on AVS Source:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. The Sources page displays all the sources registered with Cohesity DataProtect as a Service. Click the Actions menu (:) of the VMware Cloud source to [Edit](#), [Unregister](#), or [Manage Network Traffic](#).



- Click on a source. The details of the source are displayed in the **Objects**, **Connection**, and **Settings** tab. The **Objects** tab displays the hierarchy of the data centers, ESXi hosts, clusters, folders, and VMs in the AVS source.

- A selected check box () indicates the object is selected to be backed up.
-  - Indicates a Windows VM.
-  - Indicates a Linux VM.
-  - A blue shield indicates a protected VM.
-  - Indicates an excluded object that is not protected. An ancestor of the object is auto-protected, but this object is excluded through inheritance (an ancestor is explicitly excluded).
-  - Click to view the general information—type, server size, and version of the selected object.

Filters on the Objects tab help you display only the objects you want. The filtering options are:

- **Protection Status.** Protected or Unprotected.
- **VM Attributes.** VMware Tools Found and Cohesity SaaS Connector.
- **Host Type.** Linux or Windows.

Protection Status ▾

VM Attributes ▾


Host Type ▾

Enter the name of an object in the **Search** box. You can even enter a partial name or an asterisk * wildcard. As you type, the objects that contain your search term appear.

Search

🔍

Display modes can show the objects in a hierarchy or a flat list. The object hierarchy is automatically refreshed every four hours. To manually refresh the object hierarchy, select **Data Protection > Sources**, find the source in the list, and click the icon.



View the Connection Status

To view the connection status of an AVS Source:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. On the **Sources** page, select the **Connection** tab.









The **Connection** tab displays the status of the connection between the AVS source and SaaS Connection and between the SaaS Connection and the Cohesity Cloud Service in the following ways:

- The **Glance bar** displays the performance of the SaaS connection with respect to the attributes such as CPU, memory, and IOPS (Read and Write).

Average (last 24 hours)			
3%	29%	15.81	0.31
CPU	Memory	IOPS (Read)	IOPS (Write)

- The **Connection Activity** graph displays the read/write activity on the DataSource in the past 24 hours.
- The **Connector Details** section displays the connectivity status and other details of

the individual SaaS connectors.

Connector Details							
Connector	IP	Status	Release	CPU	Memory	IOPS (Read)	IOPS (Write)
 SaaS-Connector	192.1.1.0	 Healthy	hedp2023may	12%	34%	-	4.22
 SaaS-Connector	192.1.1.1	 Healthy	hedp2023may	18%	26%	-	11.14
 SaaS-Connector	192.2.2.2	 Healthy	hedp2023may	44%	27%	-	12.36
 SaaS-Connector	192.3.3.3	 Healthy	hedp2023may	20%	33%	-	11.82

View the Details of the Source

To view the details of an AVS Source:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. On the **Sources** page, select the **Settings** tab.

The **Settings** tab displays the details of the AVS source, such as the Username, Registered Date, Refreshed Date, and so on.

← Sources

vc.10dc84d574124a849ee11f.centralus.avs.azure.com

Objects
Connection
Settings

Type	VMware Cloud vCenter
VMware Cloud Provider	Azure VMware Solution (AVS)
IP or FQDN	10.0.0.2
Username	cloudadmin@vsphere.local
Registered	Jul 31, 2024 3:10pm
Refreshed	Jul 31, 2024 3:33pm
Cap concurrent streams per Datastore	Uncapped
Auto Cancel Backups if Datastore is running low on space	Off
Cap concurrent VM Backups	Uncapped
Maintenance Mode	Off

Update the AVS Source Configuration

You can update the AVS source configuration you provided during the registration process with the latest configuration. To edit the AVS source configuration:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the Actions menu (:) next to the AVS source and select **Edit**.
3. In the **Edit Source** page, update the respective configurations.
4. Click **Save**.

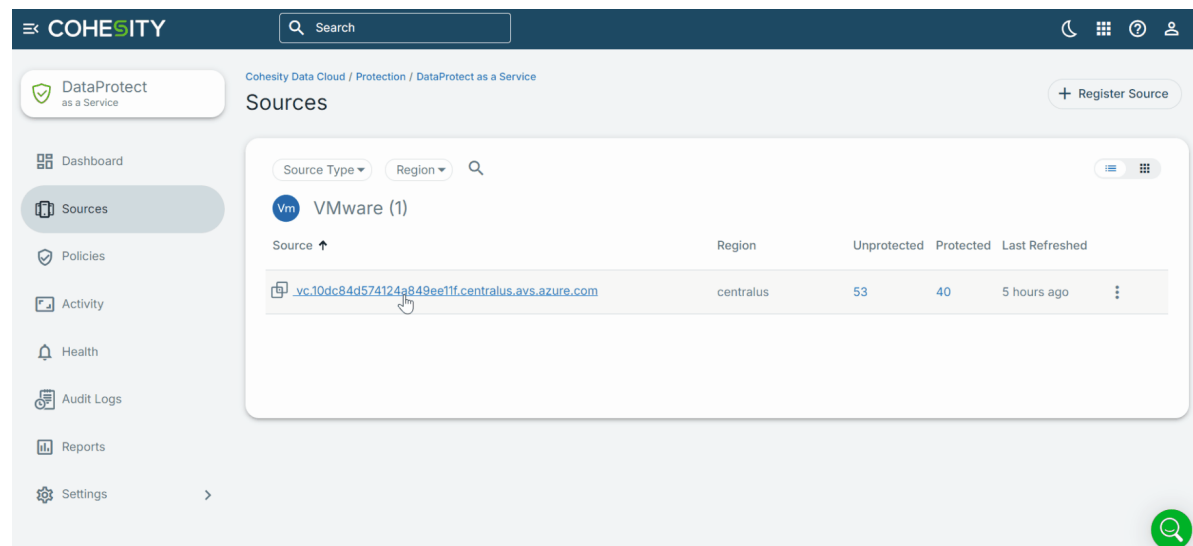
Unregister AVS Source

If you plan to stop protecting your AVS source, you can unregister the AVS source from the Cohesity cluster. To unregister the AVS source:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the Actions menu (:) next to the AVS source and select **Unregister**.


Protect VMs on AVS Source

Once you have [registered your AVS source](#), you're ready to protect the VMs in your Azure VMware solution software-defined data center (SDDC).



To protect a VMware source:

1. In **DataProtect as a Service**, under **Sources**, find the AVS source name and click on it.
2. Use the filters and search box at the top to narrow your search.
3. Use the checkboxes to select the objects for protection. To protect the whole source, click the checkbox above the column.
4. Click the **Protect** button above the checkboxes.
5. In the **New Protection** dialog, select a **Policy** that matches the schedule and retention period you need. If the existing policies do not meet your needs, you can [create a new policy](#) with the settings you need.
6. To change or configure any of the additional settings, select **More Options** and perform the below steps, or else, click **Protect**.
7. In the **Start Time** field, enter the time the protection run should start. The default time zone is the browser's time zone. You can change the time zone of the job by selecting a different time zone.
8. In the **SLA** field, enter the expected completion time for a protection run using the following options:

- **Full.** Enter the duration (in minutes) within which you expect a complete protection run to finish running. A full protection run captures all the blocks in an object.
 - **Incremental.** Enter the duration (in minutes) within which you expect an incremental protection run to finish running. An incremental protection run captures only the changed blocks in an object.
9. If you need to change any additional settings, click the down arrow icon next to [Additional Settings](#) and click **Edit** icon () next to the required field.
 10. Click **Protect**.

Cohesity DataProtect as a Service immediately starts backing up the VMs you selected. You can monitor the status of the backup on the **Activity** page. Also, the **Activity** tab of a specific VM instance shows the history of all protection runs, including the one in progress.

Additional Settings

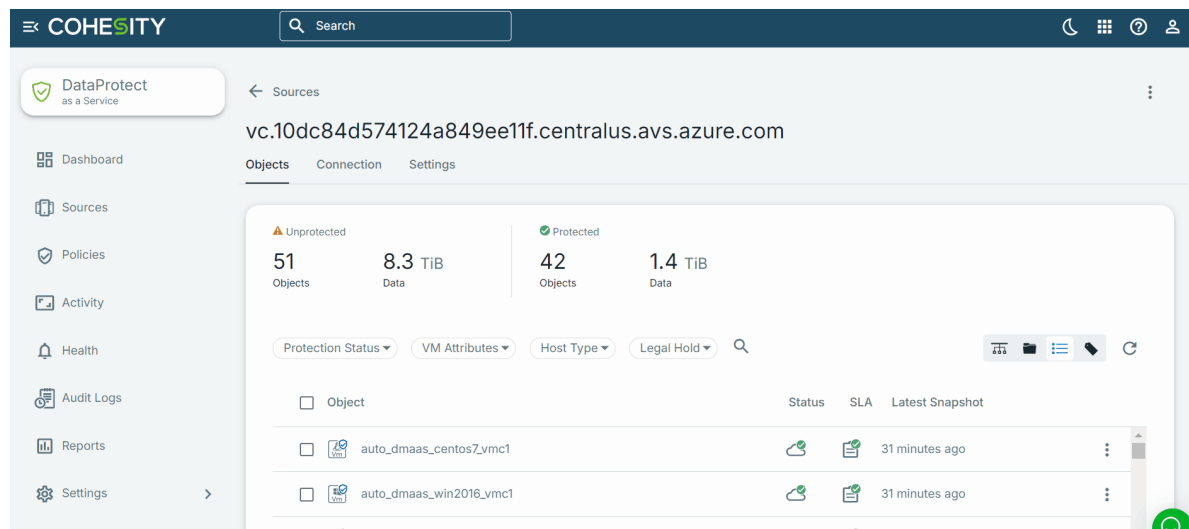
Advance Settings	Description
End Date	If you need to end protection on a specific date, enable this option and select the date.
Exclusions	Enable Exclude Disks to select the disks that are to be excluded for all VMs in the object's protection. Enter the Controller Type , Controller Bus Number , and Unit Number for each disk to exclude. Excluded disks are not backed up and are not recovered during VM recovery.

Advance Settings	Description
App Consistent Backups	<div data-bbox="618 338 1377 510"> <p>Note: This option is applicable only if VMware Tools are running in that VM. If VMware Tools are not found in the VM, Cohesity displays a warning, and selecting this option will be ineffective.</p> </div> <p>Enable App Consistent backups if you want the guest operating systems of all the protected VMs to be quiesced before snapshots of these VMs are created. Quiescing of VMs prior to capturing snapshots ensures the integrity of the data saved in the snapshots.</p> <p>With the App Consistent backups enabled, the following option is available:</p> <p>Take a Crash Consistent backup if unable to perform an App Consistent backup. Enable this option if you want Cohesity to capture a crash-consistent snapshot if Cohesity fails to capture an app-consistent snapshot. If this option is disabled and Cohesity cannot perform an app-consistent VM backup, a snapshot is not captured.</p>
Cancel Runs at Quiet Time Start	<div data-bbox="618 1024 1377 1161"> <p>Note: Field is displayed only if the selected policy has at least one Quiet Time defined. By default, this option is disabled.</p> </div> <p>When disabled, all in-progress protection runs continue to execute even when the Quiet Time starts. However, a new protection run will not start during the Quiet Time. All in-progress protection runs will abort (or pause based on your selection) once the Quiet Time starts when this option is enabled.</p>

Next > When the first protection run completes, you will be ready to [recover the protected VMs and files](#) if needed.

Recover VMs on AVS Sources

After you [protect your VMs in the AVS environment](#), you can recover VMs and files from your backups to their original or new location.



To recover VMware VMs or files:

1. In **DataProtect as a Service**, go to **Sources** to set up your recovery task.
2. Click on the **Source** name.
3. Above the tree, select **Protection Status > Protected**.
4. Use the filters, search box, and views to locate the objects or files you need.
5. To recover:
 - **VMs**, continue with the [Recover Objects](#) procedure in [Recover Protected Objects & Files](#).
 - **Files and folders**, continue with the [Recover Files & Folders](#) procedure in [Recover Protected Objects & Files](#).
6. Select your [Recovery Options](#) and click **Recover**.

Cohesity DataProtect as a Service restores the selected VMs or files to the selected location.

Physical Servers

Cohesity DataProtect as a Service provides a simple, fast, and cost-effective backup, recovery, and data management solution for Physical Servers.

Physical Server Requirements

To register your physical servers, ensure your servers meet the OS version & other requirements, then download & install the Cohesity Agent.

Before you [register your physical server sources](#), confirm that the server is on a supported OS version and meets the [disk](#) and [ports](#) requirements below, then [download & install the Cohesity Agent](#) on each server you want to protect.

Note: For information on the supported cloud regions where you can back up this source, see [Supported Workloads and Cloud Regions](#).

Support Matrix

Before you register your physical servers, ensure that you have the supported physical server versions. For more information, see [Supported Software for Cohesity DataProtect as a Service](#).

Disk Requirements

To install the Cohesity Agent, you'll need at least 56 MB of disk space on Windows systems and 360 MB on Linux systems.

Ports Requirements

Ensure that the ports listed in the Physical Servers section of the [Firewall Ports for User-Deployed SaaS Connectors](#) topic are open to allow communication between the Cohesity SaaS Connector(s) and Physical Servers.

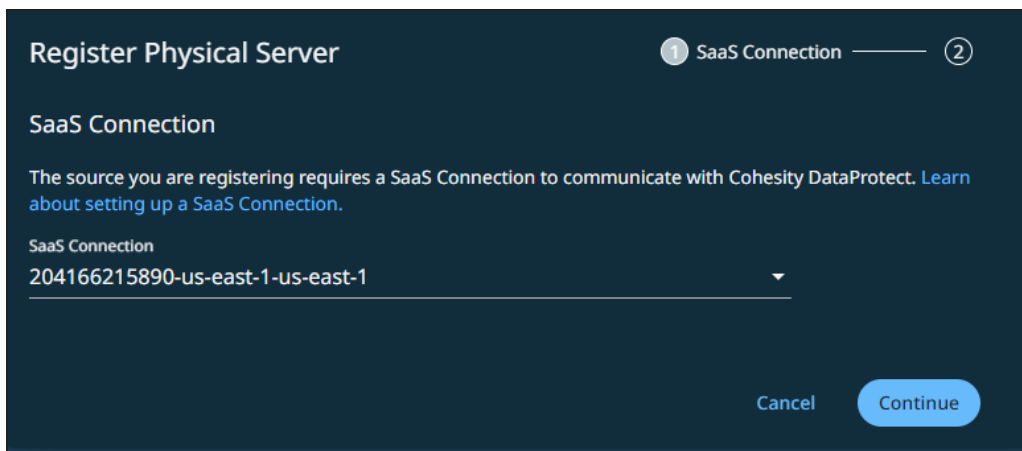
Download and Install the Cohesity Agent

Install the Cohesity Agent on each [Windows](#) and [Linux](#) physical server that you want to protect.

Install the Cohesity Windows Agent

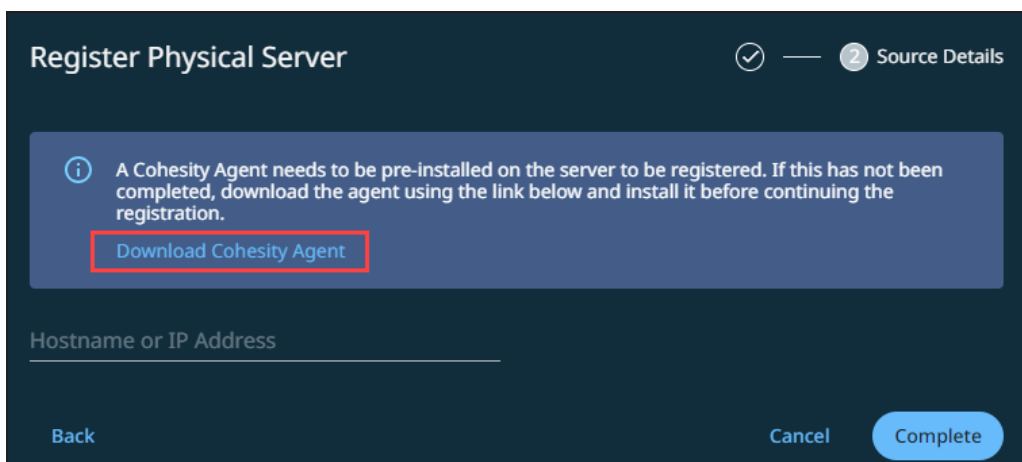
To download and install the Cohesity Windows Agent:

1. In **DataProtect as a Service**, navigate to the **Sources** page and click **+ Register Source** in the upper-right corner of the page.
2. Navigate to **Sources** and select **Register Source > Physical > Start Registration**.
3. In the Register Physical dialog box, select an existing SaaS connection marked Unused or click Create SaaS Connection and follow the instructions in [Create a SaaS Connection](#), and then click **Continue**.

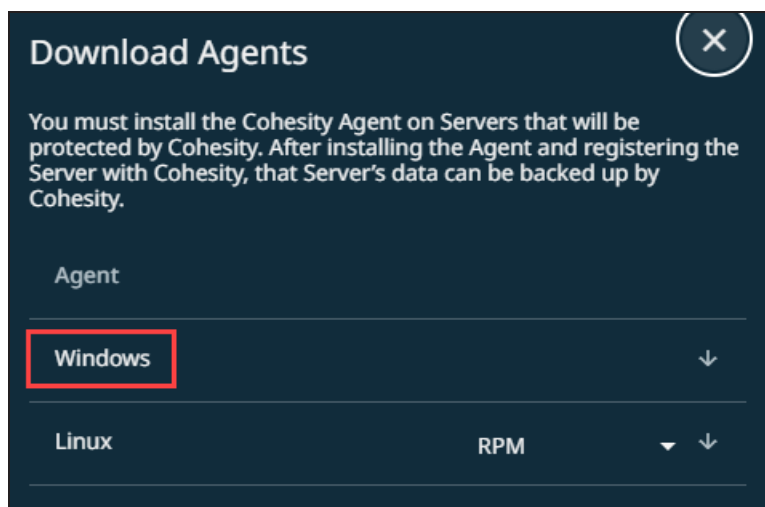


The dialog box is titled "Register Physical Server" and shows a progress bar with two steps: "1 SaaS Connection" (active) and "2". Below the title, the section "SaaS Connection" contains a message: "The source you are registering requires a SaaS Connection to communicate with Cohesity DataProtect. [Learn about setting up a SaaS Connection.](#)" Below this is a dropdown menu labeled "SaaS Connection" with the value "204166215890-us-east-1-us-east-1". At the bottom right are "Cancel" and "Continue" buttons.

4. Click **Download Cohesity Agent** and download it to the appropriate server.



The dialog box is titled "Register Physical Server" and shows a progress bar with two steps: "1" (completed) and "2 Source Details" (active). Below the title, a blue information box contains a message: "A Cohesity Agent needs to be pre-installed on the server to be registered. If this has not been completed, download the agent using the link below and install it before continuing the registration." Below this message is a button labeled "Download Cohesity Agent" which is highlighted with a red rectangle. Below the information box is a text input field labeled "Hostname or IP Address". At the bottom left is a "Back" button, and at the bottom right are "Cancel" and "Complete" buttons.



5. As an administrator with local system privileges on that server, run the executable and complete the installation wizard.

If you have only Windows servers, you're ready to [register them](#). If you have Linux servers to protect, see [Install the Cohesity Linux Agent](#).

Note: Cohesity Windows Agent does not support file path names that are not compliant with UNICODE or UTF-8 encoding.

The table below lists the encodings that are not supported.

Language	Locale	Windows Code Page
Simplified Chinese	zh_CN	936
Traditional Chinese.	zh_TW	950
Japanese-SJIS	ja_JP	932
Japanese-EUC	ja_JP	20932
Korean	ko_KR	949

Install the Cohesity Linux Agent

The Cohesity Linux Agent is available with different installer packages, providing support on multiple Linux distributions. You'll need to [install the appropriate package](#) (RPM, Debian, or SUSE RPM) for your Linux distribution or [install the script installer package](#).

The installer packages and Linux distributions on which the installer package is supported are:

Installer Package	Linux Distribution
(Default) RPM	RHEL and its derivatives
Suse RPM	SUSE
Debian	Debian and Ubuntu
Script Installer	All supported Linux operating systems

The Cohesity Linux Agent has dependencies on the following packages, which must be installed on the Linux server:

Command/Package	RHEL and its derivatives	SUSE	Debian/Ubuntu
rsync	rsync	rsync	rsync
mount	util-linux	util-linux	mount
lvm2	lvm2	lvm2	lvm2
sudo	sudo	sudo	sudo
coreutils	coreutils	coreutils	coreutils
util-linux	util-linux	util-linux	util-linux
nfs client	nfs-utils	nfs client	nfs-common
lsof	lsof	lsof	lsof
wget	wget	wget	wget

Install RPM, Debian, or SUSE RPM Installer Package

To install the RPM, Debian, or SUSE installer package:

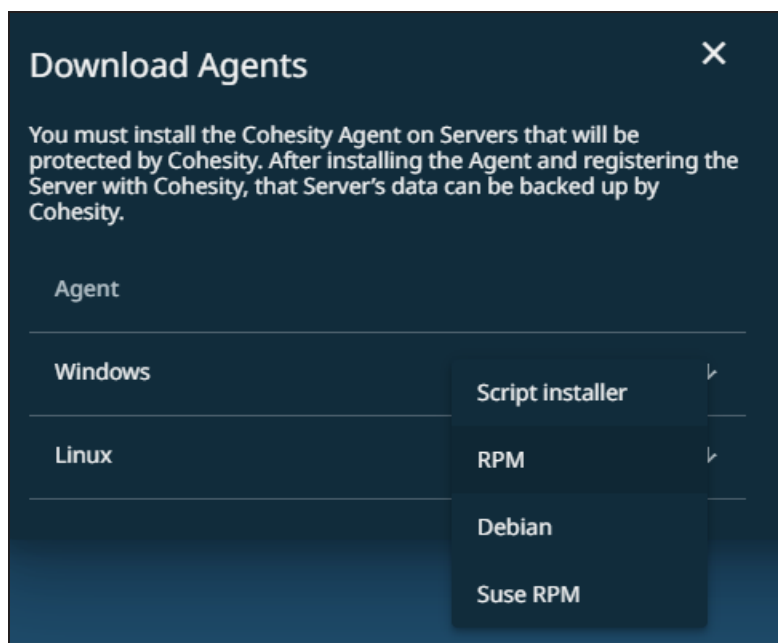
1. In **DataProtect as a Service**, navigate to the **Sources** page and click **+ Register Source** in the upper-right corner of the page.
2. Navigate to **Sources** and select **Register Source > Physical > Start Registration**.
3. In the Register Physical dialog box, select an existing SaaS connection marked Unused or click Create SaaS Connection and follow the instructions in [Create a SaaS](#)

[Connection](#), and then click **Continue**.

The screenshot shows the 'Register Physical Server' window with a progress bar at the top indicating '1 SaaS Connection' and '2' for the next step. The main heading is 'SaaS Connection'. Below it, a message states: 'The source you are registering requires a SaaS Connection to communicate with Cohesity DataProtect. [Learn about setting up a SaaS Connection.](#)' A dropdown menu labeled 'SaaS Connection' shows the selected value '204166215890-us-east-1-us-east-1'. At the bottom right, there are 'Cancel' and 'Continue' buttons.

4. Click **Download Cohesity Agent**. Based on your Linux distribution, from the **Download Agents** window, select RPM, Debian, or SUSE RPM and download it to the server you want to protect.

The screenshot shows the 'Register Physical Server' window with a progress bar at the top indicating a completed step '1' and '2 Source Details'. A blue information box contains the text: 'A Cohesity Agent needs to be pre-installed on the server to be registered. If this has not been completed, download the agent using the link below and install it before continuing the registration.' Below this text is a button labeled 'Download Cohesity Agent' which is highlighted with a red rectangle. Below the information box is a text input field labeled 'Hostname or IP Address'. At the bottom left is a 'Back' button, and at the bottom right are 'Cancel' and 'Complete' buttons.



5. As the root user with local system privileges on that server, change the directory to the location of the installer package.
6. Run the following command depending on the installer package:

Installer Package	Command
RPM	<code>rpm -i el-cohesity-agent-6.8.2-1.x86_64.rpm</code> or <code>yum localinstall ./el-cohesityagent-6.8.2-1.x86_64.rpm</code>
Debian	<code>dpkg -i cohesity-agent_6.8.2-1_amd64.deb</code>
Suse RPM	<code>rpm -i cohesity-agent-6.8.2-1.x86_64.rpm</code>

Note:

By default, the installation uses the root user permission for all the files, and the service is started as root. Therefore, it is necessary to add non-root users to the sudoers list by making the following changes in the `/etc/sudoers` file:

```
<username> ALL=(ALL) NOPASSWD:ALL
Defaults:<username> !requiretty
```

- To start the service as a non-root user, create a new user or use an existing user with sudo permission and run the following command:

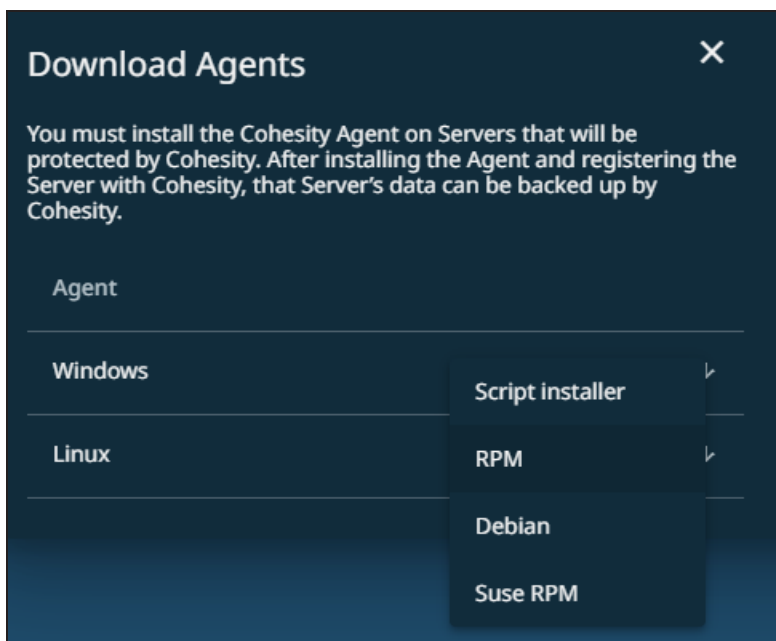
Installer Package	Command
RPM	<code>export COHESITYUSER=<username> ; rpm -i el-cohesity-agent-6.8.2-1.x86_64</code>
Debian	<code>COHESITYUSER=<username> dpkg -i cohesity-agent_6.8.2-1_amd64</code>
Suse RPM	<code>export COHESITYUSER=<username> rpm -i cohesity-agent-6.8.2-1.x86_64</code>

- Provide the location details for:
 - **Installation directory:** /opt/cohesity
 - **Log file:** /var/log/cohesity

Install Script Installer Package

To install the script installer package:

- In **DataProtect as a Service**, navigate to the **Sources** page and click **+ Register Source** in the upper-right corner of the page.
- Navigate to **Sources** and select **Register Source > Physical > Start Registration**.
- Click **Download Cohesity Agent**. Based on your Linux distribution, select **Script Installer** and download it to the server you want to protect from the **Download Agents** window.



4. As the root user with local system privileges on that server, change the directory to the location of the installer package.

Note: For SLES 11 SP4, you are required to install the Agent as the root user.

5. Make the installer executable. For example:

```
chmod +x cohesity_agent_6.8.2_linux_x64_installer
```

6. Run the executable:

```
sudo cohesity_agent_6.8.2_linux_x64_installer -- --install
```

7. Provide the location details for:

- **Installation directory:** /home/<username>/cohesityagent or /root/cohesityagent
- **Log file:** /home/cohesityagent/cohesityagent/logs

The Agent starts after installation completes, as follows:

- **CentOS and RedHat** (distributions with the "systemd" init system): The Agent starts automatically.
- **Ubuntu** (distributions with the "upstart" init system): The Agent starts automatically.

If a Linux server's /etc/sudoers file is managed by a deployment engine such as Chef, Puppet, or others, this might affect Cohesity's interaction with servers that have the Linux Agent installed. Take the corresponding actions depending on user type:

Agent Installation by User Type	Action Required
As the default cohesityagent user	<p>The Cohesity Linux Agent is installed using the cohesityagent user by default.</p> <p>For default installations, the cohesityagent user is created by the installer. During installation, the installer updates the <code>/etc/sudoers</code> file to allow cohesityagent sudo and no-tty sudo access.</p> <p>Ensure the following settings in the <code>/etc/sudoers</code> file for the cohesityagent user are preserved:</p> <pre>cohesityagent ALL=(ALL) NOPASSWD:ALL Defaults:cohesityagent !requiretty</pre> <p>For example:</p> <pre>#includedir /etc/sudoers.d dgoble ALL=(ALL) NOPASSWD:ALL cohbackup ALL=(ALL) NOPASSWD:ALL Defaults:cohbackup !requiretty</pre>
As a non-default user, for example, foo	Ensure the above settings in the <code>/etc/sudoers</code> file for the foo user are preserved by replacing the occurrences of 'cohesityagent' with 'foo'.
As root user	No changes required.

Considerations

- Currently, a source can either be protected as either a physical server or as a SQL database, but not both.
- Volume-based physical backups are not supported.

Next > Now you can [register your physical server sources](#) to protect them!

Register Physical Server Sources

Before you can [protect a physical server](#), you need to register it as a Cohesity DataProtect as a Service source.

Note: To connect with sources in your data center, you'll need to use a SaaS Connection or ([create one](#)) to establish connectivity between the sources and the Cohesity DataProtect as a Service.

To register a physical server, check that it meets the [requirements for physical servers](#) and then add it as a source in Cohesity DataProtect as a Service.

To add a physical server as a source:

1. In **DataProtect as a Service**, navigate to the **Sources** page and click **+ Register Source** in the upper-right corner of the page and then select **Physical Server**.
2. In the Register Physical dialog box, select an existing SaaS connection marked Unused or click Create SaaS Connection and follow the instructions in [Create a SaaS Connection](#), and then click **Continue**.

The screenshot shows the 'Register Physical Server' dialog box with the progress indicator at step 1, 'SaaS Connection'. The title is 'SaaS Connection'. Below the title, it says: 'The source you are registering requires a SaaS Connection to communicate with Cohesity DataProtect. [Learn about setting up a SaaS Connection.](#)' Below this is a dropdown menu labeled 'SaaS Connection' with the selected value '204166215890-us-east-1-us-east-1'. At the bottom right are 'Cancel' and 'Continue' buttons.

3. Download the Cohesity Agent by clicking the **Download Cohesity Agent** link.

The screenshot shows the 'Register Physical Server' dialog box with the progress indicator at step 2, 'Source Details'. The title is 'Source Details'. Below the title, there is an information icon and a message: 'A Cohesity Agent needs to be pre-installed on the server to be registered. If this has not been completed, download the agent using the link below and install it before continuing the registration.' Below the message is a button labeled 'Download Cohesity Agent' which is highlighted with a red rectangle. Below this is a text input field labeled 'Hostname or IP Address'. At the bottom left is a 'Back' button, and at the bottom right are 'Cancel' and 'Complete' buttons.

A Cohesity Agent needs to be pre-installed on the server to be registered. If the agent is not installed, download the agent before continuing the registration.

4. Enter the **Hostname (FQDN)** or **IP address** of the physical server you're registering. We recommend that you use the FQDN.
5. Click **Complete**.

Note: Don't run any other actions in your Cohesity DataProtect as a Service until source registration completes.

Next > You're ready to [protect your physical servers](#)!

Protect Physical Servers

Once you have [registered your physical server](#) as a source, you're ready to protect it.

To protect your physical server:

1. In **DataProtect as a Service**, navigate to **Sources**, find the Physical source name and click into it.
2. Use the filters and search box at the top to narrow your search.
3. Use the checkboxes to select the objects for protection. To protect the whole source, click the checkbox above the column. The source is automatically added as a protection object.
4. Optionally, to configure symlink, mount point, inclusion and exclusion options, click the **Edit** (pencil) icon on the right:
 - **Follow symlink NAS target** (Windows file-based backup only): Enable this option if you want to back up the symbolic link pointing to a NAS target.
 - **Protect Nested Mount Points**: Enable this option to back up the volumes that are mounted to a sub-folder within the selected directory structure.
 - **Inclusions**: This option allows you to include individual files and folders. Click to include a particular path or a particular file within the specified host.
 - **Exclusions**: This option defines how you can add exclusion entries for individual files and folders. Click to exclude a particular path or a particular file within the specified host. Such paths are children of the parent inclusion path.
5. Choose a policy to specify backup frequency and retention. If you don't have a policy, you can easily [create one](#).
6. If you wish to configure a specific **Start Time**, **End Date**, **Alerts**, and other additional settings, click **More Options**.
 - **Quiet Time**. Select this option to cancel in-progress protection runs at the start of a quiet time, as defined in the associated protection policy.
 - **Pre & Post Scripts**. Edit this option to run scripts on the protected server before and/or after a protection run. If the protection run is protecting physical servers from different hosts, then the pre and post scripts are executed for each physical server.

- **Source- Side Deduplication.** Use this option to enable source-side deduplication for all the servers that are part of the protection run.

Note: Source-side deduplication is not supported on Windows 2008 R2 servers.

- **SLA.** A service-level agreement (SLA) defines how long you expect a protection run to take. Enter:
 - **Full.** The number of minutes you expect a full protection run, which captures all the blocks in an object, to take.
 - **Incremental.** The number of minutes you expect an incremental protection run, which captures only the changed blocks in an object, to take.
- **Cache Optimization:** Toggle this option to enable cache optimization.
- **Crash Consistent Backups.** (Windows only) Enable this option to read files from the snapshots of volumes on which the files (that need backup) are residing before the protection run is executed.

7. Click **Protect**.

Cohesity DataProtect as a Service starts backing up the physical servers you selected.

Note: The backups start immediately after you protect the objects, regardless of the time you set for the protection run.

Next > When the first protection run completes, you will be ready to [recover your protected servers](#) when and if you need to.

Manage Existing Protection


Edit protection settings, change the policy, and start, stop, & pause protection.

Once you have [applied protection](#) to the objects in your sources, Cohesity DataProtect as a Service makes it easy to make changes to that protection quickly. You can:

- Edit additional settings like **End Date, Exclusions, Alerts**, and more.
- Apply a different policy.
- Start an on-demand protection run, pause and resume it, or even remove protection.

Edit Protection Settings

To edit protection settings:


1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click into the **Source** name.
3. Select **Show All > Protected** and use the other filters, search box, and views at the top to narrow your search.
4. Click the **Actions** menu () next to the object and select **Edit Protection** to open the protection settings for that object.


Apply a New Protection Policy

To change the **Policy**, click the drop-down and select a different policy. To help you choose, each policy in the list shows the **Backup** frequency and the **Retain** period for each backup.

If you don't have a policy that meets your needs, scroll to the bottom of the list and click **Create Policy** to [create your own policy](#).

Start, Stop, or Remove Protection

Click the **Actions** menu () next to the object, Cohesity DataProtect as a Service presents buttons for the actions that are possible for those objects.



With the protected objects selected, you can click:

- **Recover** to recover the object or file.
- **Unprotect** to remove protection from the object.

Tip: If a protected object is deleted from the source, you can search the object using Global Search and unprotect it.

- **Run Now** to start an on-demand protection run immediately.

Edit Additional Protection Settings

Under **Settings**, you can change the protection **Start Time** (and select the **Time Zone**).

Click the drop-down next to **Additional Settings** to change more options. See [Additional Protection Settings](#) for details.

Additional Settings

Advance Settings	Description
Start Time	Available only if the selected policy is set to Backup Daily . Indicates what time the protection run should start. Enter the Start Time and select AM or PM . The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone here.
End Date	If you need to end protection on a specific date, enable this to select the date.
Quiet Time	<p><i>(Available only if the selected policy has at least one Quiet Time)</i></p> <p>When enabled, all the protection runs that are currently executing will cancel when the Quiet Time period starts. By default, this setting is disabled, meaning that after a protection run starts, it continues to execute even when a Quiet Time period starts. However, new protection runs will not start during a Quiet Time.</p>
Pre & Post Scripts	Edit this option to run scripts on the protected server before and/or after a protection run. If the protection run is protecting physical servers from different hosts, then the pre and post scripts are executed for each physical server.
Crash Consistent Backups	<p>Crash Consistent Backups: Enable Crash Consistent backups if you want the guest operating systems of all the protected VMs to be quiesced before snapshots of these VMs are created. Quiescing of VMs prior to capturing snapshots ensures the integrity of the data saved in the snapshots.</p> <p>Continue with non consistent backup in case of failure: Enable this options to allow the backup to continue, if a crash consistent backup with VSS cannot be initiated.</p> <div> <p>Note: This option is applicable only for VSS copy backup.</p> </div>
SLA	<p>The service-level agreement (SLA) defines how long the administrator expects a protection run to take. Enter:</p> <ul style="list-style-type: none"> • Full. The number of minutes you expect a full protection run, which captures all the blocks in an object, to take. • Incremental. The number of minutes you expect an incremental protection run, which captures only the changed blocks in an object, to take.

Recover Physical Servers

After you [protect your physical servers](#), you can recover them from Cohesity DataProtect as a Service to their original or a new location.

To recover a protected physical server:


1. In **DataProtect as a Service**, navigate to **Sources**, find the Physical source name and click into it.
 2. Use the filters and search box at the top to narrow your search.
 3. Click the **Recover** icon.
 4. Select the snapshot to recover and click **Apply**.
 5. Browse the content of the backup and select files, folders, or volumes to be recovered, then click **Next**.
 6. Optionally, click **Download Files** to download the contents of the recovery task after the task is completed.
 7. Under **Recover to**, select **Original Server** or **New Server**. With:
 - **Original Server**, Cohesity will overwrite the original physical server instance. You can restore the data in the original path or provide an alternate path for the restore.

To recover to a different location in the original server, disable the **Recover to Original Path** option, and then provide the location to which the files or folders are to be recovered in the **Recover To** field. By default, the files and folders will be recovered to the original location.
- Note:** This is a destructive action that cannot be undone.
- **New Server**, select a registered **Source**, the **Target** instance, and provide the recovery **Path**.
 8. Under **Recovery Options**, you can set:
 - **Overwrite Existing File/Folder**. By default, this option is enabled to overwrite the existing files and folders. Disable this option to create the files and folders in the specified location. Any duplicate files are skipped.
 - **Preserve File/Folder Attributes**. By default, this option is enabled and the ACLs, permissions, and timestamps are preserved for all files and folders. If you disable this option, then ACLs and permissions are not preserved. If you recover both folders and files, the folders will receive the new timestamps, but the files retain their original timestamps. If you recover only files, then the files will receive the new timestamps.
 - **Continue on Error**. Enable this option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.
 - **Task Name**. Change the default name of the recovery task.
 9. Click **Recover**.

Upgrade Agents

When we release a new version of the Cohesity Agent, you will see an option to upgrade the agent on the source details page. You can either perform a bulk agent upgrade or upgrade an agent individually.

To upgrade an individual agent

1. In **DataProtect as a Service**, navigate to **Sources** and click on the physical server source name.
2. In the **Source Details** page, click the **More Options** menu () and then select **Upgrade Agent**.

Note: The Upgrade Agent option is enabled only when a new version of the Agent is available.

3. Click one of the following options:
 - **Upgrade Now** to upgrade the Agent immediately, then click **Confirm**.
 - **Schedule for Later** In the **Schedule Agent Upgrade** dialog, set the **Date & Time** for the upgrade and click **Schedule**.

To perform a bulk agent upgrade


1. In **DataProtect as a Service**, navigate to **Sources**.
2. If agent upgrades are available, click **Upgrade Agents** listed adjacent to physical sources.
3. Select the agents you want to upgrade and click **Continue**.
4. Click **Add Schedule** to schedule the agent upgrades. Options include:
 - **Now** to upgrade the Agent immediately.
 - **Later** In the **Schedule** dialog, set the **Date & Time** for the upgrade.
5. Click **Upgrade**.

To validate the current agent version

1. In **DataProtect as a Service**, navigate to **Sources** and click on the physical server source name.
2. In the **Source Details** page, click the **Settings** tab and check the agent version in the **Agent Version** field.

Agent Upgrade Status

Upgraded agents may not reflect the status immediately. To update the agent upgrade status, follow the steps below:

1. In **DataProtect as a Service**, navigate to **Sources** and click on the physical server source name.
2. In the **Source Details** page, click the **More Options** menu () and then select **Refresh**.

NAS

Cohesity DataProtect as a Service provides a simple, fast, and cost-effective backup, recovery, and data management solution for NAS environments.

Register Generic NAS Sources

Before you Begin

Ensure that the ports listed in the Network Attached Storage (NAS) section in the [Firewall Ports for User-Deployed SaaS Connectors](#) topic are open to allow communication between the Cohesity SaaS Connector(s) and NAS Server.

Note: For information on the supported cloud regions where you can back up this source, see [Supported Workloads and Cloud Regions](#).

Register

Before you can [protect a NAS device](#) you need to register it as a source in Cohesity DataProtect as a Service. You can register any generic NAS, Dell EMC Isilon NAS, or NetApp ONTAP. For:

- **Generic NAS**, see the [steps below](#).
- **Isilon NAS**, see [Configure and Register Isilon NAS](#).
- **NetApp ONTAP**, see [Configure and Register NetApp ONTAP](#).

Important: Ensure that the TCP/UDP ports 445, 8080, 111, and 2049 are open in the firewall between your SaaS Connector and data source.

For more information, see [Supported Software for Cohesity DataProtect as a Service](#).

Register Generic NAS

You can connect a generic NAS source to Cohesity DataProtect as a Service as a mount point via the NFS (v3) or SMB (v2, v3) protocol.

Note: SMB v1 is not supported in Cohesity DataProtect as a Service.

To register your generic NAS source via NFS or SMB:

1. In **DataProtect as a Service**, navigate to the **Sources** page and click **+ Register Source** in the upper-right corner of the page.
2. In the **Select Source** dialog box, select **NAS**.
3. In the **Type:** drop-down, select **Generic NAS** and click **Start Registration**.
4. In the Register Generic Nas dialog box, select an existing SaaS connection marked Unused or click Create SaaS Connection and follow the instructions in [Create a SaaS Connection](#), and then click **Continue**.

Register Generic Nas

1 SaaS Connection — 2

SaaS Connection

The source you are registering requires a SaaS Connection to communicate with Cohesity DataProtect. [Learn about setting up a SaaS Connection.](#)

SaaS Connection

Saas_Connection-21

Cancel Continue

5. Select the **Mode** as **NFS** or **SMB**.
6. Enter the **Mount Path**.
 - For **NFS**, enter the `hostname or IP:/Volume`.
 - For **SMB**, enter the `\\hostname or IP\Share Path`.
7. If you chose **SMB** above, enter the **Username** and **Password** required to access the SMB share.
8. If you are confident the mount point is correct, you can enable **Skip Mount Point validation during registration**. (Optional.)
9. Add a **Description** to make it easier to recognize this source. (Optional.)
10. Click **Complete**.

Your NAS device is now a registered source in your Cohesity DataProtect as a Service and ready to be [protected](#).

Note: If you plan to stop protecting a NAS source, you can remove it from your Cohesity DataProtect as a Service. Navigate to **Sources**, click the **Actions** menu () next to the NAS source and select **Unregister**. In the **Unregister Source** dialog, click **Unregister**.

Next > You're ready to [protect your NAS sources](#).

Verify the Kerberos Configuration in Rigel

When using Kerberos authentication for backups, it is essential to verify the Kerberos configuration. To prevent authentication issues during backup, ensure that the DNS and Key Distribution Centers (KDCs) settings are correct.

To verify the Kerberos authentication for backups:

1. Configure the DNS for Kerberos SRV records.
 - Access the DNS Management Console:
 - Depending on your DNS server (BIND, Windows DNS, etc), access the DNS management tool.
 - Create SRV records for Kerberos:

- Add SRV records for the `_kerberos._tcp` service in the relevant zone (for example, `EXAMPLE.COM`).

Example SRV Records:

Service	Protocol	Name	Priority	Weight	Port	Target
_kerberos	_tcp	_kerberos._tcp	0	100	88	kdc1.example.com
_kerberos	_tcp	_kerberos._tcp	0	100	88	kdc2.example.com
_kerberos	_tcp	_kerberos._tcp	1	50	88	kdc3.example.com

- Ensure that the DNS server is properly configured to handle requests for the queried domain.
- Ensure that the DNS server can resolve the domain names associated with the Kerberos service.

2. Ensure proper propagation.

Allow time for DNS changes to propagate, and ensure that the DNS server is functioning properly.

- Run the `nslookup` command to retrieve Kerberos SRV records{ }.
- Use the following command to retrieve the Kerberos SRV records and validate the KDCs:

```
nslookup -type=srv _kerberos._tcp.<domain>
```

Replace `<domain>` with the appropriate domain for your environment (e.g. `EXAMPLE.COM`).

Example:

```
nslookup -type=srv _kerberos._tcp.EXAMPLE.COM
```

- Check SRV records for multiple KDCs.
- The command should return a list of SRV records with details of the KDC servers, including priority, weight, port, and hostname.

Example Output:

```
_kerberos._tcp.EXAMPLE.COM    SRV service location:
```

```
priority = 0
```

```
weight = 100
```

```
port = 88
```

```
svr hostname = kdc1.example.com
```

```
_kerberos._tcp.EXAMPLE.COM    SRV service location:
```

```
priority = 0
```

```
weight = 100
```

```
port = 88
```

```
svr hostname = kdc2.example.com
```

```
_kerberos._tcp.EXAMPLE.COM    SRV service location:
```

```
priority = 1
```

```
weight = 50
```

```
port = 88
```

```
svr hostname = kdc3.example.com
```

- Verify KDC reachability.
 - Test if each KDC hostname is reachable via the network by using the ping command:

```
ping <kdc-hostname>
```

Example:

```
ping kdc1.example.com
```

```
ping kdc2.example.com
```

```
ping kdc3.example.com
```

- If all KDCs are reachable, proceed with the backup. If any KDC is unreachable, resolve network or DNS issues before starting the backup. Even if one hostname fails, the backups will fail.
- Check DNS resolution without SRV lookup.
 - As a sanity check, confirm that the DNS resolution for the domain is working by running a simple nslookup on the domain itself:

```
nslookup <domain>
```

Example:

```
nslookup QA01.ENG.COHE.COM
```

- Firewall and Network configuration.
 - Ensure that the necessary DNS and Kerberos service ports are allowed through the firewall and accessible on the network:
 - DNS (UDP/TCP port 53)
 - Kerberos (UDP/TCP port 88)

Note:

- Ensure firewalls allow traffic on UDP/TCP port 88 for Kerberos and port 53 for DNS.
- For environments using dynamic DNS, set the appropriate permissions for updating DNS records.

Configure and Register Isilon NAS

Check your Isilon requirements and minimum permissions, then register your Isilon NAS sources with Cohesity DataProtect as a Service.

Note: For information on the supported cloud regions where you can back up this source, see [Supported Workloads and Cloud Regions](#).

To add an Isilon cluster as a Cohesity DataProtect as a Service source:

1. Confirm that you have met the [Isilon requirements](#).
2. Check [the minimum Isilon user permissions](#).
3. [Register](#) your Isilon NAS source.

Note: To register other NAS types, see [Register Generic NAS Sources](#) or [Configure and Register NetApp ONTAP](#).

Isilon Requirements

- Ensure that the TCP/UDP ports 445, 8080, 111, and 2049 are open in the firewall between your SaaS Connector and data source.
- Isilon OneFS version 8.0.x, 8.1, or 8.2.x.

- NFS v3 for NFS export backups.

Note: Cohesity DataProtect as a Service uses NFS v3 and SMB v2 or v3 for data protection; SMB v1 is not supported in Cohesity DataProtect as a Service.

- On Isilon NFS shares, enable the "Mount access to subdirectories" flag. Cohesity DataProtect as a Service requires this setting to mount the **.snapshot** directory of the shared path.
- SnapshotIQ license enabled on Isilon, with these settings:

The screenshot displays the 'SnapshotIQ' configuration page. At the top, there is a navigation bar with 'Dashboard', 'Cluster Management', 'File System', and 'Data Protection'. Below this, the 'SnapshotIQ' section has three tabs: 'Snapshots', 'Snapshot Schedules', and 'Settings'. The 'Settings' tab is selected, leading to the 'Edit File System Snapshot Settings' page. This page is divided into two main sections: 'Service' and 'Visibility and Access Settings'. In the 'Service' section, there are three checked checkboxes: 'Enable snapshot service', 'Auto-create snapshots', and 'Auto-delete snapshots'. The 'Visibility and Access Settings' section contains a checked checkbox for 'Enable global visibility and access', followed by expandable sections for 'NFS Settings', 'SMB Settings', and 'Local Settings'. Each of these sections has three checked checkboxes: 'root directory accessible', 'root directory visible', and 'sub-directories accessible'. At the bottom left of the settings area is a 'Revert Changes' button.

Minimum Isilon User Permissions

Cohesity DataProtect as a Service accesses your Isilon cluster using an Isilon user account. The user account must have the following permissions to back up and restore your Isilon data via SMB or NFS.

Access-level	Command	Description
ReadOnly	Platform API	For access to Isilon's APIs.
	Auth	To verify users and passwords.
	Cluster	To obtain cluster identity and settings.
	Network	To obtain the network interfaces.
	SMB	To read the settings in the SMB server.
Read/Write	Job Engine	To read and write Changelist jobs.
	Snapshot	To fetch, create, and delete snapshots for shares and exports.
	NFS	To read and write settings to and from the NFS server. <div> Note: This setting modifies the NFS export used to mount, such as <code>/ifs</code>. </div>

Register Isilon Cluster

To register your Isilon cluster:

1. In **DataProtect as a Service**, navigate to the Sources page and click **+ Register Source** in the upper-right corner of the page.
2. In the **Select Source** dialog box, select **NAS**.
3. In the **Type:** drop-down, select **Isilon** and click **Start Registration**.
4. In the **Register Isilon** dialog box, select an existing SaaS connection marked **Unused** or click **Create SaaS Connection** and follow the instructions in [Create a SaaS Connection](#), and then click **Continue**.

The screenshot shows the 'Register Isilon' wizard with three steps: 1. SaaS Connection, 2, and 3. Step 1 is active. The title is 'SaaS Connection'. Below the title, a message states: 'The source you are registering requires a SaaS Connection to communicate with Cohesity DataProtect. [Learn about setting up a SaaS Connection.](#)' Below this, there is a dropdown menu labeled 'SaaS Connection' with the selected value 'Saas_Connection-21'. At the bottom right, there are two buttons: 'Cancel' and 'Continue'.

5. Enter the **Username** and **Password** that you configured earlier, under *Minimum Isilon User Permissions* above.
6. Enter the Isilon cluster's **Hostname or IP Address** and then click **Continue**.

The screenshot shows the 'Register Isilon' wizard with three steps: 1, 2, and 3. Step 3 is active. The title is 'Source Settings'. Below the title, there is a toggle switch for 'SMB Volumes Access' which is turned on. Below the toggle, a message states: 'Enable this if you are using SMB volumes. Note that SMB also requires a qtree with NTFS permissions or volumes with mixed-mode permissions, as well as credentials.' Below this, there are two input fields: 'Username' with the value 'smbuser' and 'Password' with a masked value '.....'. Below the 'Username' field, there is a hint: 'Username or Domain\Username'. Below the 'Password' field, there is a hint: 'Deny IP Addresses' and a value '10.101.10.1'. Below the 'Deny IP Addresses' field, there is a hint: 'Comma separate multiple IP addresses. CIDR suffix's can be used to specify IP ranges.' At the bottom left, there is a 'Back' button. At the bottom right, there are two buttons: 'Cancel' and 'Complete'.

7. If you are backing up SMB volumes or mixed-mode volumes, enable **SMB Volumes Access** and enter the local or Active Directory (AD) **Username** and **Password** required for at least **read** access to the Isilon SMB share.


Note:

- You can assign the local or AD user to the built-in "BackupAdmin" role to permit that user to read the SMB data for backup without modifying the access control lists (ACLs).
- To provide access at the share level, grant the "Run as root" and "Full Control" permissions at the share level.
- The user must have full control on the restore target during recovery.

8. To exclude IP addresses or subnets from the communications between Cohesity DataProtect as a Service and the Isilon cluster, enable **Exclude IPs** and enter those IPs.

9. Click **Complete**.

Your Isilon cluster is now a registered source in your Cohesity DataProtect as a Service and ready to be [protected](#).

Note: If you plan to stop protecting a NAS source, you can remove it from your Cohesity DataProtect as a Service. Navigate to **Sources**, click the **Actions** menu () next to the NAS source and select **Unregister**. In the **Unregister Source** dialog, click **Unregister**.

Next > You're ready to [protect your Isilon NAS sources](#).

Configure and Register NetApp ONTAP

Check your NetApp ONTAP requirements and minimum permissions, then register your ONTAP sources with Cohesity DataProtect as a Service.

Note: For information on the supported cloud regions where you can back up this source, see [Supported Workloads and Cloud Regions](#).

To add NetApp ONTAP as a Cohesity DataProtect as a Service source:

1. Confirm that you have met the [NetApp ONTAP requirements](#) below.
2. Check the supported [NetApp ONTAP versions and volumes](#).
3. Check the [minimum permissions](#).
4. [Register](#) your NetApp ONTAP source.

Note: To register other NAS types, see [Register Generic NAS Sources](#) or [Configure and Register Isilon NAS](#).

NetApp ONTAP Requirements

To register your NetApp ONTAP with Cohesity DataProtect as a Service, confirm you meet the following prerequisites:

- [SaaS Connection Requirements](#).
- Bidirectional TCP ports 111, 443, 445, 635, and 2049 are open in the firewall between your SaaS Connector and NetApp ONTAP. For details, see [Ports Used for Communication](#) below.
- The NetApp ONTAP SVM that you plan to protect have:
 - An active logical interface attached to the SVM.
 - The NFS and CIFS services configured on the SVM.
- The **Make snapshot directory (.snapshot) visible** option is enabled for all NetApp ONTAP volumes that you plan to protect.

Support Matrix

Before you register your NetApp ONTAP with Cohesity DataProtect as a Service, ensure that the Cohesity supports the NetApp ONTAP versions and volumes you want to protect.

Supported NetApp ONTAP Versions

Cohesity DataProtect as a Service supports data protection of NetApp ONTAP versions 8.2, 8.3, 9.1, 9.2, 9.3, 9.5, 9.6, 9.7, 9.8, 9.9.1, 9.10.x, 9.11.x, 9.12.x.

Supported NetApp ONTAP Volumes

The supported NetApp ONTAP versions and volume types for backup are:

Volume Type	Volume Subtype
Flex Volume	Normal Flex Volume
Data Protection Volume	SnapMirror Destination Volume
	SnapVault Destination Volume

Supported NFS and SMB versions

The supported NFS and SMB versions for backup are:

Protocol	Version	Notes
NFS	NFSv3	If NFSv4 volume backup is triggered, Cohesity DataProtect as a Service will take the backup in NFSv3 mode.
SMB	SMB v2.x and v3	SMB v1 is not supported in Cohesity DataProtect as a Service.

Minimum Permissions

Ensure the user account you use to register your NetApp ONTAP SVM or NetApp ONTAP cluster has the required permissions to communicate with the Cohesity DataProtect as a Service.

Minimum Permissions for NetApp ONTAP Cluster

Before registering a NetApp ONTAP cluster as the source type, ensure the user account has the following command permissions:

Access Level	Command	Description	Protocol
All	vserver export policy	Adds the Cohesity SaaS Connector IP to the export policy so that Cohesity DataProtect as a Service can mount volumes.	NFS
	volume snapshot	Allows fetching, creating, and deleting snapshots for volumes.	SMB / NFS
ReadOnly	vserver cifs	Fetches information about CIFS/SMB shares for volumes.	SMB / NFS
	cluster identity	Fetches information about the cluster.	SMB / NFS
	network interface	Fetches information about network interfaces that the Cohesity DataProtect as a Service connects to for mounting volumes.	SMB / NFS
	volume	Fetches information about volumes.	SMB / NFS
	vserver	Fetches information about SVM	SMB / NFS

Minimum Permissions for NetApp ONTAP SVM

When registering a NetApp ONTAP SVM as the source type, ensure the user account has the following command permissions:

Access Level	Command	Description	Protocol
All	vserver export policy	Adds the Cohesity SaaS Connector IP to the export policy so that Cohesity DataProtect as a Service can mount volumes.	SMB / NFS
	volume snapshot	Allows fetching, creating, and deleting snapshots for volumes.	SMB / NFS
ReadOnly	vserver cifs	Fetches information about CIFS/SMB shares for volumes.	SMB
	network interface	Fetches information about network interfaces to which the Cohesity DataProtect as a Service connects for mounting volumes.	SMB / NFS
	volume	Fetches information about volumes.	SMB / NFS
	vserver	Fetches information about SVM.	SMB / NFS

Minimum Permissions for SMB/CIFS Shares Backup and Recovery

To back up NetApp ONTAP SMB/CIFS shares, the user must have local or domain user credentials that allow at least read access to the SMB share.

To recover the SMB/CIFS shares, the local or domain user must have full access control on the target where the data is being restored.

Minimum Permissions for NFS Export Backup and Recovery

For Backup. To back up an NFS export, the user must have read and superuser access permissions on the NFS volume to be backed up and on the parent root volume. Before starting the backup, Cohesity DataProtect as a Service verifies that the user has these permissions and if not, Cohesity adds a new export rule for the Cohesity SaaS Connector IP with the required permissions in the export policy attached to the backup volume.

If there is already an existing export rule for the Cohesity SaaS Connector IP with a lower rule index value, then this existing export rule will override the export rule added by Cohesity for the Cohesity SaaS Connector IP. In such scenarios, you must manually update the existing export rule with the required permissions for the Cohesity SaaS Connector IP.

For the parent root volume, you must manually add the permissions for the Cohesity SaaS Connector IP.

For Recovery. To recover an NFS export, the user needs read/write and superuser access permissions on the NFS volume to be restored. Before starting the restore, you must add a new rule index for the Cohesity node subnet in the export policy attached to the source volume and parent root volume to give the necessary permissions to the Cohesity SaaS Connector IP.

Credentials for NetApp ONTAP Backup with Multiple SVMs

To register NetApp ONTAP with multiple SVMs, create a custom role with the required permissions and a local user at the SVM level. Assign the custom role to the local user. Use the respective local user account to register multiple SVMs.

Ports Used for Communication

Ensure the following ports are open in the firewall (for your backup and recovery traffic) between your SaaS Connector and NetApp ONTAP:

Port	Source	Target	Direction	Network Protocol	Usage
111	NetApp	SaaS Connector	Bidirectional	TCP/UDP	Required for RPC connection
443	NetApp	SaaS Connector	Bidirectional	TCP/UDP	Required for HTTPS connection with NetApp
445	NetApp	SaaS Connector	Bidirectional	TCP	Required for SMB
635	NetApp	SaaS Connector	Bidirectional	TCP/UDP	Required for NFS
2049	NetApp	SaaS Connector	Bidirectional	TCP/UDP	Required for NFS

Considerations

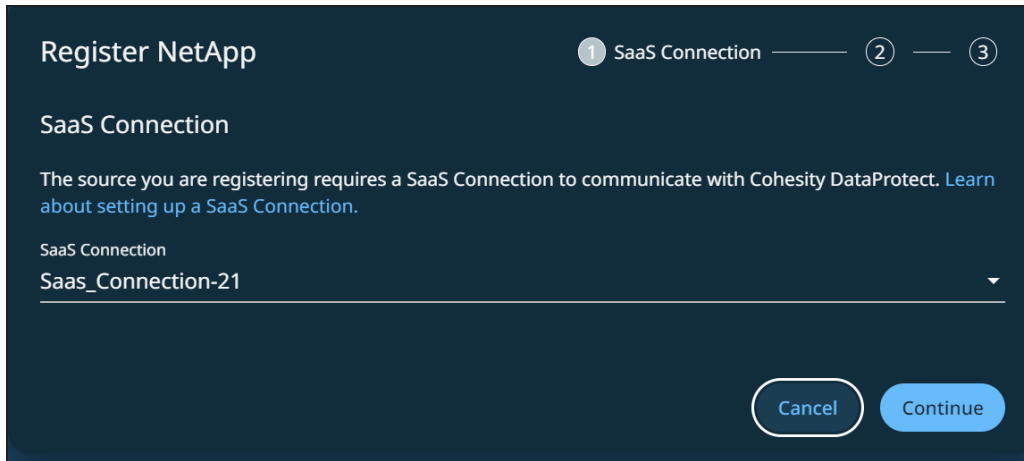
Review and understand the following limitations before you protect your NetApp ONTAP data with Cohesity DataProtect as a Service:

- Instant Volume Mount for NetApp ONTAP stub file is not supported.
- You cannot restore the NetApp Data-Protect volume to the original location or to an alternate Data-Protect volume because the Data-Protect volume is a read-only volume.
- Cohesity does not support the backup of the following NetApp ONTAP volumes:
 - FlexGroup Volume.
 - Flex Volume subtypes SnapLock Enterprise Volume, SnapLock Compliance Volume, and Encrypted Volume Storage.

Register NetApp ONTAP

To register NetApp:

1. In **DataProtect as a Service**, navigate to the Sources page and click **+ Register Source** in the upper-right corner of the page.
2. In the **Select Source** dialog box, select **NAS**.
3. In the **Type:** drop-down, select **NetApp** and click **Start Registration**.
4. In the **Register NetApp** dialog box, select an existing SaaS connection marked **Unused** or click **Create SaaS Connection** and follow the instructions in [Create a SaaS Connection](#), and then click **Continue**.



5. Choose between NetApp ONTAP cluster and SVM. Select:
 - **Cluster** to register a NetApp ONTAP cluster.
 - **VServer/SVM** to register a NetApp ONTAP SVM.
6. In the **Username** field, enter the username used to access the NetApp ONTAP cluster or SVM. Specify a user that has adequate privileges to perform actions on the source. See [Minimum Permissions](#) above for details.

Note: The username used to register the NetApp ONTAP cluster and SVM is case sensitive.

7. In the **Password** field, enter the password for the specified user.
8. In the **Hostname or IP Address** field, enter the hostname or IP address of the NetApp ONTAP cluster or SVM to register.
9. Click **Continue**.

Register NetApp

✓ — 2 Source Details — 3

Type
☒ Cluster ☐ VServer/SVM

Account Credentials

Username
Netapp

Password
.....

Hostname or IP Address
10.1.1.1

Back Cancel Continue

10. If you are backing up SMB volumes or mixed-mode volumes, enable **SMB Volumes Access** and provide the local or Active Directory user credentials that allow at least read access on the NetApp ONTAP cluster or SVM.
11. Enable **Exclude IPs** and specify the IP addresses of the NetApp ONTAP source through which the communication to the Cohesity DataProtect as a Service must *not* happen. You can enter the IP addresses in a comma-separated list or in a CIDR format.
12. Click **Complete**.

Register NetApp

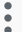
✓ — ✓ — 3 Source Settings

☐ SMB Volumes Access
 Enable this if you are using SMB volumes. Note that SMB also requires a qtree with NTFS permissions or volumes with mixed-mode permissions, as well as credentials.

☐ Exclude IPs

Back Cancel Complete

Your NetApp ONTAP is now a registered source in your Cohesity DataProtect as a Service. and ready to be **protected**.

Note: If you plan to stop protecting a NAS source, you can remove it from Cohesity DataProtect as a Service. Navigate to **Sources**, click the **Actions** menu () next to the NAS source and select **Unregister**. In the **Unregister Source** dialog, click **Unregister**.

Next > You're ready to [protect your NetApp ONTAP NAS volumes and data](#).

Protect NAS Sources

Use Cohesity DataProtect as a Service to protect the NAS volumes, files, and folders in your data center. You can protect any generic NAS source, a Dell EMC Isilon NAS cluster, or a NetApp ONTAP cluster or SVM with Cohesity DataProtect as a Service — just note that the registration process is different for each:

- [Register Generic NAS](#)
- [Configure and Register Isilon NAS](#)
- [Configure and Register NetApp ONTAP](#)

Once registered, your NAS source is ready for [protection](#)!

Important: Ensure that the TCP/UDP ports 445, 8080, 111, and 2049 are open in the firewall between Cohesity DataProtect as a Service and your NAS device.

Protect NAS

1. In **DataProtect as a Service**, navigate to **Sources**, find the NAS source name and click into it.
2. Use the filters and search box at the top to narrow your search.
3. Use the checkboxes to select the objects for protection. To protect the whole source, click the checkbox above the column.

Note:

When you check a parent object, you can choose:

- **Select All Child Objects.** To capture the tree as it currently exists, or
- **Auto Protect This <NAS source type>.** To capture the tree and any future additions.

4. Click the **Protect** icon above the checkboxes.
5. In the **New Protection** dialog, select a **Policy** that matches the schedule and retention period you need. If the existing policies do not meet your needs, you can [create a new policy](#) with the settings you need.

Note: Cohesity recommends a first full and incremental forever backup approach to protect your NAS sources.

6. To change or configure any of the additional settings , select **More Options** and perform the below steps or else, click **Protect**.
7. In the **Start Time** field, enter the time the protection run should start. The default time zone is the browser's time zone. You can change the time zone of the job by selecting a different time zone.
8. If you need to change any of the additional settings, click the down arrow icon next to [Additional Settings](#) and click **Edit**.
9. Click **Protect**.

Your selected NAS objects are backed up with the frequency and retention as defined in the policy you have selected.

Note: The backups start immediately after you protect the objects, regardless of the time you set for the protection run.

10. **Exclusions and Inclusions:** Everything is included by default. Toggle on **Exclusions and Inclusions** if you want to exclude or include locations. By creating exclusion and inclusion rules, you can limit the Protection Group to a specific set of files and directories and therefore minimize the disk space used to store the data.

Note:

To add an exclusion or inclusion, you must prefix a forward slash ('/') or suffix an asterisk ('*') to the path or to a particular file within the protected object. For example, `'/test'` or `'*.txt'`.

Add Inclusion: Click to include a particular path or a particular file within the protected object.

For example, consider four directories - **test**, **test1**, **test2**, and **test3** under the protected object with path `/ifs/TestShare1/Folder1`. The table below lists the input types for an inclusion list of folders and their respective outcome:

Input	Outcome
/	Everything inside the protection path is backed up, if exclusions are not defined.

Input	Outcome
/test	<p>When inclusion path is not suffixed by '/', Protection Group will include all folders starting from protection path.</p> <p>In this example, /test will behave as /test* and Protection Group will include 'test', 'test1', 'test2', and 'test3' from protection path for backup.</p>
/test/	<p>When inclusion path is suffixed by '/', it will only include specific folder.</p> <p>In this example, /test/ behaves as /test/* and Protection Group will only include 'test' from protection path for backup.</p>

Note:

Do not provide the folder name along with the backup path while including files and folders for a Protection Group. The backup path is considered by default. For example, /ifs/TestShare1/Folder1/test does not include any folder.

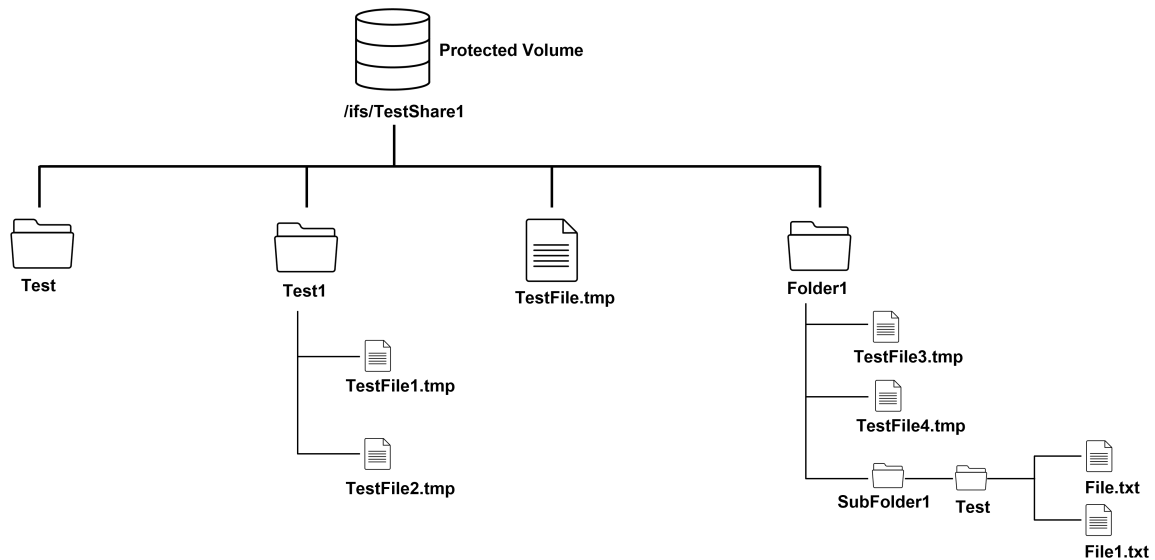
If exclusions are defined, they will take precedence over inclusions.

Add Exclusion: Click to exclude a particular path or a particular file within the protected object. You can also specify regular expressions for excluding files in the format `regex:<regex pattern>`

For example, consider there are three directories-**Test**, **Test1**, and **Folder1** and a file, **TestFile.tmp** under the protected object with path **/ifs/TestShare1**.

- **Test1** folder has **TestFile1.tmp**, **TestFile2.tmp** files.
- **Folder1** folder has **TestFile3.tmp**, **TestFile4.tmp** files, and **SubFolder1** which further has a **Test** folder with **File.txt**, **File1.txt** files in it.

The following figure illustrates the directory structure under the protected object with path **/ifs/TestShare1**.



The table below lists the input types for an exclusion list of folders and their respective outcome:

Input	Description	Outcome
/Test	<p>When the exclusion path is not suffixed by a '/', the Protection Group will exclude all files and folders with Test as the prefix. For exclusions, /Test behaves as /Test*.</p> <p>Note: Exclusion will not be effective on the sub-directory level.</p>	<p>Excludes Test and Test1 folders, and TestFile1.tmp file from the protection path.</p> <p>TestFile3.tmp and TestFile4.tmp files in the Folder1 folder will not be excluded.</p>
/Test/	<p>When the exclusion path is suffixed by '/', the Protection Group excludes only the files and folders with the name Test. For exclusions, /Test/ behaves as /Test/*.</p> <p>Note: Exclusion will not be effective on the sub-directory level.</p>	<p>Excludes the Test folder of the parent directory.</p> <p>The Test sub-folder in /Folder1/SubFolder1 will not be excluded.</p>

Input	Description	Outcome
<code>*/Test</code>	Excludes all the files and directories with the name ' Test '. Note: Exclusion will be effective on all levels of sub-directories.	Excludes the Test folder of the parent directory and the Test sub-folder in /Folder1/SubFolder1 .
<code>*.tmp</code>	Excludes all the files that with ' .tmp '. Note: Exclusion will be effective on all levels of sub-directories.	Excludes TestFile.tmp , TestFile1.tmp , TestFile2.tmp , TestFile3.tmp , TestFile4.tmp files.
<code>regex:.*\tmp</code>	Excludes the files and folders with ' .tmp '. Note: Exclusion will be effective on all levels of sub-directories.	Excludes TestFile.tmp , TestFile1.tmp , TestFile2.tmp , TestFile3.tmp , TestFile4.tmp files.
<code>regex:/Folder1/.*Test/.*\txt\$</code>	Excludes all files with .txt extension from all the Test folders directory in /Folder1/ .	Excludes File.txt and File1.txt files in the /Folder1/SubFolder/Test folder.

Note: Do not provide the folder name along with the backup path while excluding files and folders for a Protection Group. The backup path is considered by default. For example, `/ifs/TestShare1/Folder1/test` does not exclude any folder.

Cohesity automatically excludes the following NetApp system files:

- `.vtoc_internal` and `.bplustvoc_internal` files
- `.copy-offload` directory and `.tokens` file

Additional Settings

Additional Settings	Description
End Date	If you need to end protection on a specific date, enable this to select the date.
Skip Files on Errors	<i>(On by default)</i> A protection run continues even if it encounters errors on files, such as permissions errors. If files are skipped, the protection run details page indicates a Warning status and provides additional information. If toggled off, the protection run stops when it encounters an error.
Exclusions and Inclusions	By default, all files and folders are included for protection. Use this option if you want to exclude or include specific locations. By creating exclusion and inclusion rules, you can limit the protection to a specific set of files and directories and therefore minimize the disk space used to store the data.
Cancel Runs at Quiet Time Start	<i>(Available only if the selected policy has at least one Quiet Time.)</i> When enabled, all the protection runs that are currently executing will cancel when the Quiet Time period starts. By default, this setting is disabled, meaning that after a protection run starts, it continues to execute even when a Quiet Time period starts. However, new protection runs will not start during a Quiet Time.
SLA	The service-level agreement (SLA) defines how long the administrator expects a protection run to take. Enter: <ul style="list-style-type: none"> • Full. The number of minutes you expect a full protection run, which captures all the blocks in an object, to take. • Incremental. The number of minutes you expect an incremental protection run, which captures only the changed blocks in an object, to take.

Next > When the first protection run completes, you will be ready to [recover NAS volumes, files, and folders](#) when and if you need to.

Manage Existing Protection


Edit protection settings, change the policy, and start, stop, & pause protection.

Once you have [applied protection](#) to the objects in your sources, Cohesity DataProtect as a Service makes it easy to make changes to that protection quickly. You can:

- Edit additional settings like **End Date**, **Exclusions**, **Alerts**, and more.
- Apply a different policy.
- Start an on-demand protection run, pause and resume it, or even remove protection.

Edit Protection Settings

To edit protection settings:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click into the **Source** name.
3. Select **Show All > Protected** and use the other filters, search box, and views at the top to narrow your search.
4. Click the **Actions** menu () next to the object and select **Edit Protection** to open the protection settings for that object.

Apply a New Protection Policy

To change the **Policy**, click the drop-down and select a different policy. To help you choose, each policy in the list shows the **Backup** frequency and the **Retain** period for each backup.

If you don't have a policy that meets your needs, scroll to the bottom of the list and click **Create Policy** to [create your own policy](#).

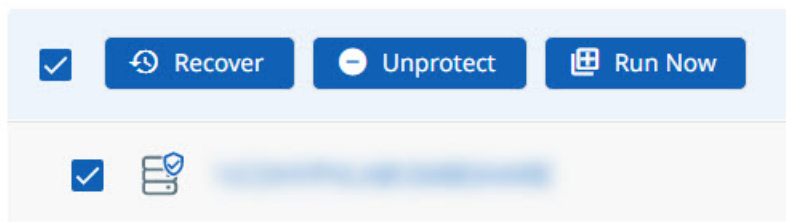
Edit Additional Protection Settings

Under **Settings**, you can change the protection **Start Time** (and select the **Time Zone**).

Click the drop-down next to **Additional Settings** to change more options. See [Additional Protection Settings](#) for details.

Start, Stop, or Remove Protection

When you select protected objects in one of your sources, Cohesity DataProtect as a Service presents buttons for the actions that are possible for those objects.



With the protected objects selected, you can click:

- **Recover** to recover the object or file.
- **Unprotect** to remove protection from the object.

Tip: If a protected object is deleted from the source, you can search the object using Global Search and unprotect it.

- **Run Now** to start an on-demand protection run immediately.

Additional Settings

Advance Settings	Description
End Date	If you need to end protection on a specific date, enable this to select the date.
Skip Files on Errors	<i>(On by default)</i> A protection run continues even if it encounters errors on files, such as permissions errors. If files are skipped, the protection run details page indicates a Warning status and provides additional information. If toggled off, the protection run stops when it encounters an error.
Exclusions and Inclusions	By default, all files and folders are included for protection. Use this option if you want to exclude or include specific locations. By creating exclusion and inclusion rules, you can limit the protection to a specific set of files and directories and therefore minimize the disk space used to store the data.
Cancel Runs at Quiet Time Start	<i>(Available only if the selected policy has at least one Quiet Time)</i> SLA When enabled, all the protection runs that are currently executing will cancel when the Quiet Time period starts. By default, this setting is disabled, meaning that after a protection run starts, it continues to execute even when a Quiet Time period starts. However, new protection runs will not start during a Quiet Time.
	The service-level agreement (SLA) defines how long the administrator expects a protection run to take. Enter: <ul style="list-style-type: none"> • Full. The number of minutes you expect a full protection run, which captures all the blocks in an object, to take. • Incremental. The number of minutes you expect an incremental protection run, which captures only the changed blocks in an object, to take.

Recover NAS Data

After you protect your NAS sources, you can recover the NAS volumes, files, and folders from your backups, to their original or a new location.

To recover protected NAS data:

1. In **DataProtect as a Service**, navigate to **Sources** to set up your NAS recovery task.
2. Follow the steps below for [Recover NAS Volumes](#) or [Recover NAS Files & Folders](#).

Set Up NAS Recovery

To recover protected NAS data:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click into the **Source** name.
3. Above the tree, select **Object Protection > Protected**.
4. Use the filters, search box, and views to locate the volumes or files you need.
5. To recover:
 - **NAS volumes**, continue with [Recover NAS Volumes](#) below.
 - **Files and folders**, continue with [Recover NAS Files & Folders](#) below.

Tip: You can also use Global Search to locate, filter, and select the objects you need. Click the **Global Search** box at the top or type **slash (/)** anywhere to start your search.

Recover NAS Volumes

To recover NAS volumes, follow these steps (from [Recover Objects & Volumes](#)):

1. Locate and select the NAS volumes you need, and then click **Recover** at the top to open the **New Recovery** form with the **Latest** snapshot (protection run).
2. If you need to recover from an earlier snapshot, click the **Edit** icon to select a new recovery point.
 - For each object under **Selected**, you can click the **Edit** icon to open the **Recovery Point** calendar. Click **List** to view the available recovery points by timestamp and click one.
 - Click **Select Recovery Point**.
 - Click **Next: Recover Options** to return to the form.
3. Under **Recover To**, select **Original Location** or **New Location**.
 - **For VMs:** If you choose **New Location**, select a **Registered Source**, **Resource Pool**, **Datastores**, and the **VM Folder**.

- For **NAS volumes**: If you choose **New Location**, select a **Registered Source** and the **Volume**.
4. Select your [Recovery Options](#).
 5. Click **Start Recovery**.

Cohesity DataProtect as a Service opens the **Activity** page, showing your file recovery task as it runs, along with the recovery progress on the right.

Tip: The **Activity** page also shows the entire history of all protection runs and recovery tasks, including any that are in progress.

Recover NAS Files & Folders

To recover files and folders from protected NAS volumes, follow these steps (from [Recover Files & Folders](#)):

1. Locate the source volume containing the files and click **Recover Files** on the row for that object to open the **Select Files** form.
2. If you need to recover from an earlier snapshot, click the **Recovery Point** calendar drop-down to select the recovery point.
 - Click **List** to view the available recovery points by timestamp and click one.
 - Click **Apply**.
3. Click into the path to find the files and add them to the **Selected Items** list.
4. Choose how to recover your files: download locally or recover.
 - Click **Download Files** to open the **Activity** page, showing your file recovery task. Click into the recovery task and click **Download Files** a second time to save them to your local system.
 - Click **Save** to open the **New Recovery** form. Under **Recover To**, select **Original Location** or **New Location**.
 - If you choose **Original Location**, enter a **Username** and **Password** that has access to the original server. You can also enable **Recover to Alternate Path** to enter a new path on the original server.
 - If you choose **New Location**, select a registered Source and a Target (VM) or Volume (NAS). Enter a **Username** and **Password** that has access to that server and enter a **Recover To** path.
5. Select your [Recovery Options](#).
6. Click **Start Recovery**.

Cohesity DataProtect as a Service opens the **Activity** page, showing your file recovery task as it runs, along with the recovery progress on the right.

Tip: The **Activity** page also shows the entire history of all protection runs and recovery tasks, including any that are in progress.

Microsoft 365

Microsoft 365 is a subscription service that bundles the traditional office productivity applications and delivers them as SaaS applications. Microsoft 365 includes Exchange Online, OneDrive for Business, SharePoint Online, Teams, and other applications. Cohesity DataProtect as a Service provides simple, fast, and cost-effective data protection solution for the following Microsoft 365 applications:

- [Exchange Online Mailboxes](#)
- [OneDrive for Business](#)
- [SharePoint Online](#)
- [Microsoft Teams](#)
- [Microsoft Groups](#)

Microsoft 365 Requirements

Before you register your Microsoft 365 sources with Cohesity DataProtect as a Service to protect your Microsoft 365 data, ensure you have met the following prerequisites:

1. In the Exchange admin center, [add these roles to the Microsoft 365 user account](#) you will use to register your Microsoft 365 sources with Cohesity DataProtect as a Service:
 - ApplicationImpersonation
 - View-Only Configuration
 - View-Only Recipients
 - MailboxSearch
 - MailRecipients
2. Update [Microsoft Organization setting](#) for Mailbox size reporting.
3. [Register a custom Azure app](#) (for manual Microsoft 365 source registration).
4. [Set additional permissions for SharePoint Online](#).

Finally, review the considerations for each supported Microsoft 365 application:

- [Exchange Online Mailboxes](#)
- [OneDrive](#)
- [SharePoint Online](#)
- [Teams](#)
- [Groups](#)

Note: For information on the supported cloud regions where you can back up this source, see [Supported Workloads and Cloud Regions](#).

Add Roles to Microsoft 365 User Account

Cohesity DataProtect as a Service accesses your Microsoft 365 domain with a user account to back up your Microsoft Exchange Online data. You can either add these roles to an existing user account or create a new user account with these roles.

Important: Ensure that multi-factor authentication is not enabled for the user account.

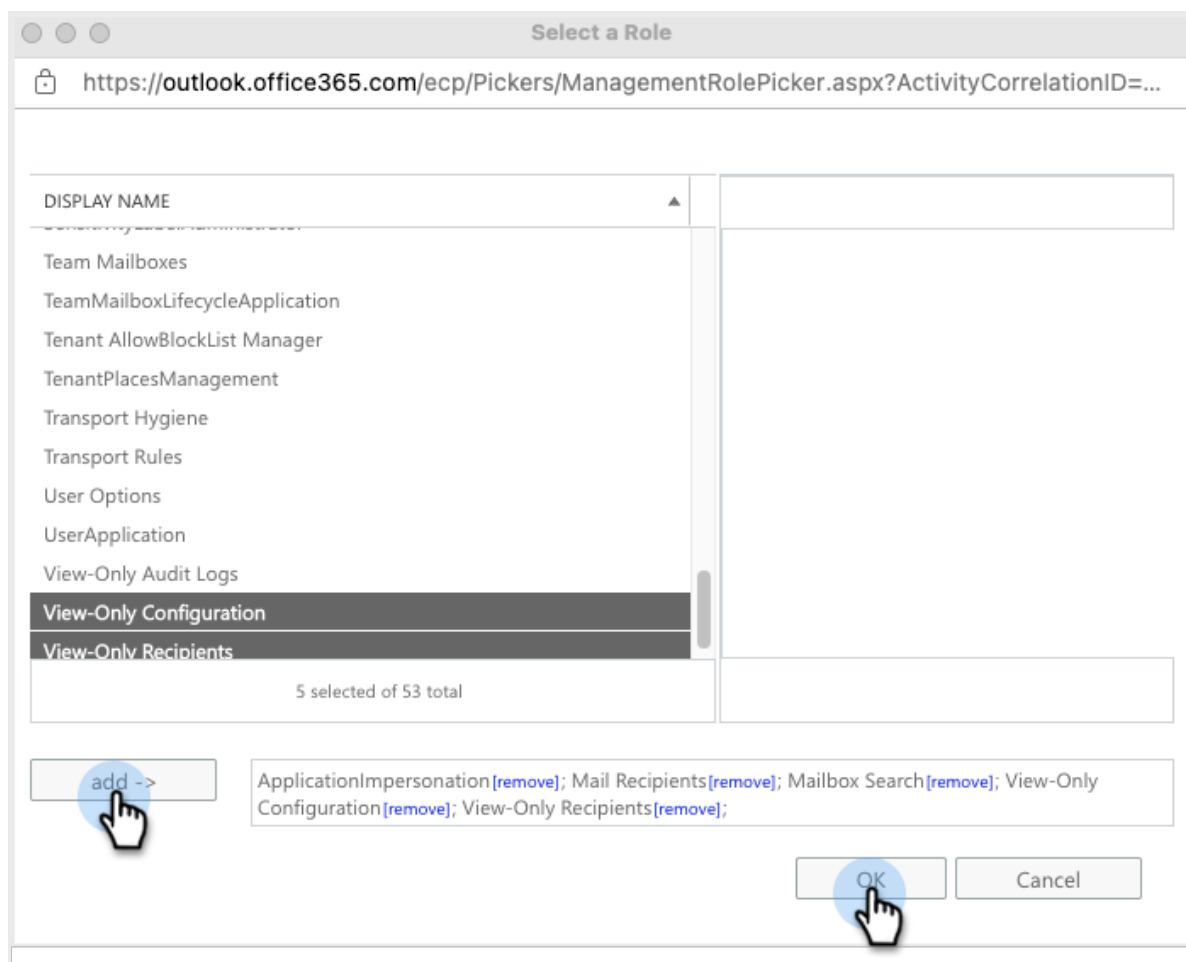
To add roles to the Microsoft 365 user account:

1. Log in to [Microsoft 365](#).
2. In the **Office 365** page, click **Admin**.
3. In the **Microsoft 365 admin center** page, select **Admin centers**, and then click **Exchange**.

Follow the steps for Classic Exchange admin center in [Step 4](#) next, or skip to Step 5 if you are in the new Exchange admin center page.

Tip: TIP: If you see a message prompting you to switch to New Exchange, you are still in classic Exchange.

4. To add roles from the Classic Exchange admin center page:
 1. Click **Permissions** and then select the **Admin roles** tab.
 2. In the **Admin roles** tab, click **+** to create a new role group.
 3. In the **new role group** page, enter a **Name** and **Description**, and under **Roles**, click **+**.
 4. In the **Write scope** drop-down, select **Default** and click **Next**.
 5. In the **Select a Role** page, select the following roles, click **Add**, and then **OK**:
 - Mail Recipients
 - Mailbox Search
 - View-Only Configuration
 - View-Only Recipients



6. Under **Members**, click **+** to add the user account you plan to use to register the Microsoft 365 domain with Cohesity DataProtect as a Service, then click **OK**.
7. Click **Save** to create the Role Group.

Role Group

https://outlook.office365.com/ecp/UsersGroups/NewAdminRoleGroup.a...

new role group

*Name:

Description:

Write scope:

Roles:
+ -

NAME
ApplicationImpersonation
Mail Recipients
Mailbox Search
View-Only Configuration
View-Only Recipients

Members:
+ -

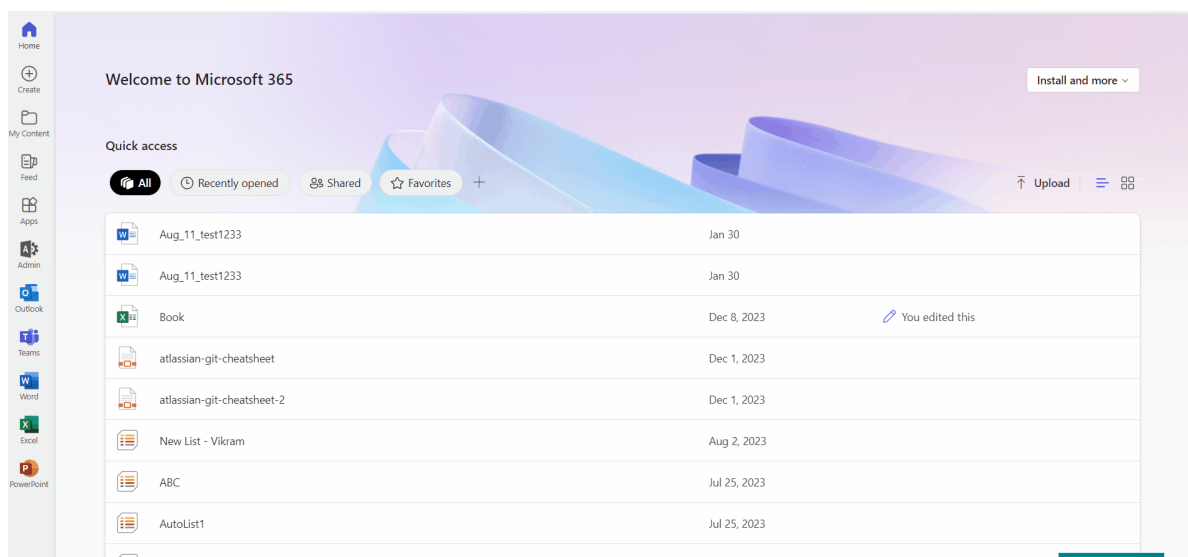
NAME	DISPLAY NAME
backupadmin	Backup Admin

Save

Cancel

You are ready to [update your Microsoft 365 Org setting](#) for Mailbox size reporting.

5. To add roles from the new Exchange admin center page:
 1. Select **Roles > Admin roles**.
 2. In the **Admin roles** page, click **Add role group**.
 3. Under **Basics**, enter a **Name** and **Description** for the admin role.
 4. In the **Write scope** drop-down, select **Default** and click **Next**.
 5. Under **Permissions**, select the following and click **Next**:
 - Mail Recipients
 - Mailbox Search
 - View-Only Configuration
 - View-Only Recipients
 6. Under **Admins**, search and select the user account you plan to use to register the Microsoft 365 domain with CCohesity DataProtect as a Service, then click **Next**.
 7. Under **Review and finish**, review the configuration and click **Add role group**.
6. After the role group is added, click **Done**.



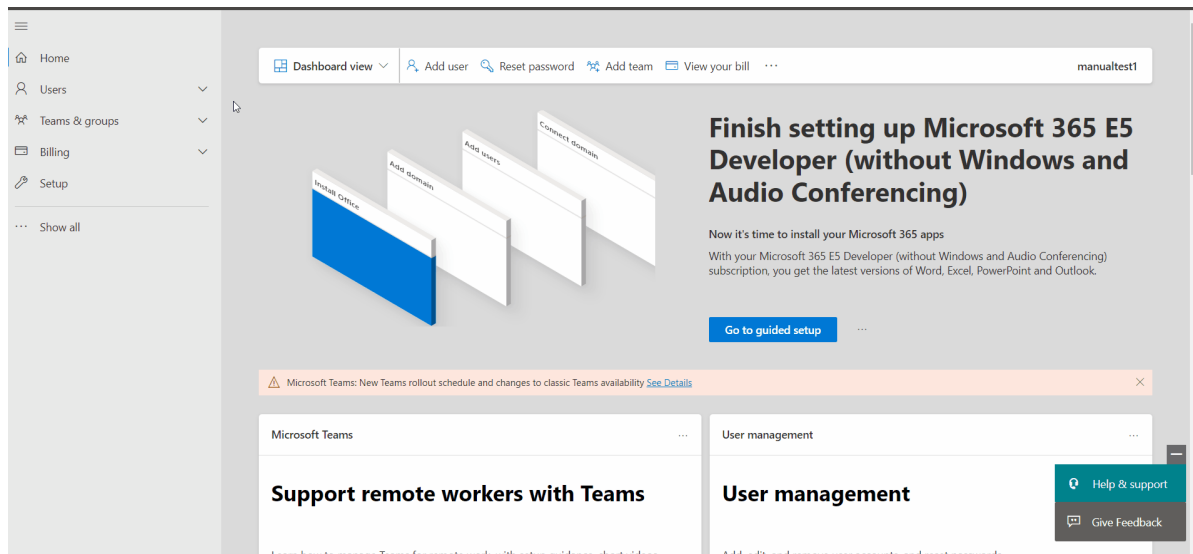
You are ready to [update your Microsoft 365 Org setting](#) for Mailbox size reporting.

Update Microsoft Organization Setting for Mailbox Size Reporting

By default, Microsoft reports, using Graph API, display information as de-identified names for users, groups, and sites. However, for Mailbox size reporting to work in Cohesity, you need to have identifiable information in the Email activity reports. To do that, you need to disable de-identified names for users, groups, and sites in Microsoft 365 reports.

Update the following organization setting in your Microsoft 365 admin center:

1. Log in to your [Microsoft 365 admin center](#) as a Microsoft 365 tenant administrator.
2. Go to **Settings > Org settings > Services > Reports**.
3. In **Reports**, ensure the information is identifiable by deselecting **Display concealed user, group, and site names in all reports**.
4. Click **Save**.



To continue, if you are using:

- Cohesity's [express registration](#) for Microsoft 365 sources, you are ready to add those sources to Cohesity DataProtect as a Service.
- The [manual registration](#) for Microsoft 365 sources, you must first [register your custom Azure app](#).

Note: For SharePoint Online data protection, ensure that you also set the required [add-in permissions](#) and [tenant permissions](#) on the Azure application.

Register Custom Azure App

To get started, you'll register a custom Azure app below to add the necessary permissions.

Go to the Azure portal, register a new app, add the permissions, and capture the App ID and Access Key. For more on registering and configuring Azure apps, see [Register an application with the Microsoft identity platform](#) and [Configure a client application to access a web API](#) in the Microsoft documentation.

Note: Make sure that you make note of the App ID and Access Key while registering the app. You'll need them to [register your Microsoft 365 domain as a source](#) in Cohesity DataProtect as a Service.

To register your custom app for Cohesity DataProtect as a Service:

1. Open Microsoft Entra ID
 1. To manage Microsoft Entra ID using the Azure Portal:
 1. Log in to the [Azure portal](#) with your Microsoft 365 administrator user credentials.
 2. Click the main menu (≡) in the top left corner and select **Microsoft Entra ID**.
 2. To manage Microsoft Entra ID using Microsoft 365:
 1. Log in to [Microsoft 365](#).
 2. On the **Microsoft 365** page, click **Admin**.
 3. On the **Microsoft 365 admin center** page, select **Admin centers** and then click **Microsoft Entra**.
2. Create a new custom app.
 1. Under the **Manage** section, select **App Registrations**, then click **New Registration**. In the **Register an application** page:
 1. Enter a **Name** for your app.
 2. Select the **Supported account types** that can access the app,
 3. In the **Redirect URI** drop-down, select **Web** and enter `https://localhost`.

4. Click **Register**.

Register an application ...

*** Name**
The user-facing display name for this application (this can be changed later).

Cohesity DataProtect App ✓

Supported account types
Who can use this application or access this API?

☒ Accounts in this organizational directory only (cohesitydmas only - Single tenant)
☒ Accounts in any organizational directory (Any Azure AD directory - Multitenant)
☐ Accounts in any organizational directory (Any Azure AD directory - Multitenant) and personal Microsoft accounts (e.g. Skype, Xbox)
☐ Personal Microsoft accounts only

[Help me choose...](#)

Redirect URI (optional)
We'll return the authentication response to this URI after successfully authenticating the user. Providing this now is optional and it can be changed later, but a value is required for most authentication scenarios.

Web ✓

Register an app you're working on here. Integrate gallery apps and other apps from outside your organization by adding from [Enterprise applications](#).

By proceeding, you agree to the [Microsoft Platform Policies](#) ↗

Register

3. After the custom app has been created, click **Overview** and copy the **Application (client) ID**. You need to use **Application (client) ID** to register Microsoft 365 as a source in Cohesity DataProtect as a Service.

Cohesity DataProtect App ...

Search (Cmd+/) << Delete Endpoints Preview features

Overview

Quickstart
Integration assistant
Manage
Branding
Authentication

Essentials

Display name	: Cohesity DataProtect App	Client credentials	: Add a certificate or secret
Application (client) ID	: 36744479-4114-4c50-8799-f588472f8e4f	Redirect URIs	: 1 web, 0 spa, 0 public client
Object ID	: e25393d3-0cbc-4f6a-b7a0-9b754afd1566	Application ID URI	: Add an Application ID URI
Directory (tenant) ID	: 942464e4-30bd-40a1-b631-cac735352ef6	Managed application in l...	: Cohesity DataProtect App
Supported account types	: My organization only		

4. Add API permissions to the custom app:

1. **Add OAuth API permission** if the Microsoft 365 source tenant has OAuth enabled for secure communication:
 1. Under the **Manage** section, select **App Registrations** and click **Add a permission**.
 2. In the **Request API permissions** page, click the **APIs my organization uses** tab.
 - a. In the search bar, enter **Office 365 Exchange Online** then **click the API**. (Use the complete app name.)

Request API permissions

Select an API

[Microsoft APIs](#)
[APIs my organization uses](#)
[My APIs](#)

Apps in your directory that expose APIs are shown below

Name	Application (client) ID
Office 365 Exchange Online	00000002-0000-0ff1-ce00-000000000000

- b. In the Office 365 Exchange Online API, click **Application Permissions**.

Request API permissions

[← All APIs](#)

03

Office 365 Exchange Online
https://ps.outlook.com

What type of permissions does your application require?

Delegated permissions
 Your application needs to access the API as the signed-in user.


Application permissions
 Your application runs as a background service or daemon without a signed-in user.

- c. Under **Other Permissions**, select `full_access_as_app` to enable OAuth and click **Add Permissions**.

App Permissions	Permission Type	Mailboxes
full_access_as_app	Application	Y

Request API permissions

[← All APIs](#)

 Office 365 Exchange Online
https://ps.outlook.com

What type of permissions does your application require?

Delegated permissions

Your application needs to access the API as the signed-in user.

Application permissions

Your application runs as a background service or daemon without a signed-in user.

Select permissions

Permission	Admin consent required
<p>Other permissions (1)</p> <p><input checked="" type="checkbox"/> full_access_as_app ⓘ Use Exchange Web Services with full access to all mailboxes</p>	Yes
<p>Calendar</p>	


2. Add Graph API permissions:
 1. Under the **Manage** section, select **App Registrations**, and then click **Add a permission**.
 2. In the **Request API permissions** page, select **Microsoft Graph API**.

Request API permissions

Select an API

[Microsoft APIs](#)
[APIs my organization uses](#)
[My APIs](#)

Commonly used Microsoft APIs



Microsoft Graph

Take advantage of the tremendous amount of data in Office 365, Enterprise Mobility + Security, and Windows 10. Access Azure AD, Excel, Intune, Outlook/Exchange, OneDrive, OneNote, SharePoint, Planner, and more through a single endpoint.

3. Click **Application Permissions** and add the permissions listed below for your Microsoft 365 application.

App Permissions	Permission Type	Mailboxes	OneDrive	SharePoint Online Sites	MS Teams
Channel.Create	Application	N/A	N/A	N/A	✓
Channel.ReadBasic.All	Application	N/A	N/A	N/A	✓
ChannelMember.ReadWrite.All	Application	N/A	N/A	N/A	✓
Directory.ReadWrite.All	Application	✓	✓	✓	✓
Files.ReadWrite.All	Application	N/A	✓	✓	✓
Group.Create	Application	N/A	N/A	N/A	✓
Group.ReadWrite.All	Application	N/A	✓	✓	✓
Reports.Read.All	Application	✓	✓	✓	✓
Sites.ReadWrite.All	Application	✓	✓	✓	✓
Sites.FullControl.All	Application	N/A	N/A	✓	✓
User.Read.All	Application	✓	✓	✓	✓
User.ReadWrite.All	Application	N/A	✓	✓	✓

App Permissions	Permission Type	Mailboxes	OneDrive	SharePoint Online Sites	MS Teams
ChannelMessage.Read.All	Application	N/A	N/A	N/A	✓
Chat.Read.All	Application	✓	N/A	N/A	N/A
Mail.ReadWrite	Application	✓	N/A	N/A	N/A
MailboxSettings.Read	Application	✓	N/A	N/A	N/A

4. Click **Add permissions**.
3. Add SharePoint permissions to the custom app:
 1. Under the **Manage** section, select **App Registrations** and click **Add a permission**.
 2. In the **Request API permissions** page, select **SharePoint**. (If you don't see it, scroll further down.)
 - a. Click **Delegated Permissions** and add the permissions listed below, then click **Add permissions**.
 - b. Click **Application Permissions** and add the permissions listed below, then click **Add permissions**.

Permission Type	Permissions Name
Delegated	AllSites.FullControl
	AllSites.Manage
	AllSites.Read
	MyFiles.Read
	MyFiles.Write
	Sites.Search.All
	TermStore.ReadWrite.All
	User.ReadWrite.All
Application	Sites.FullControl.All
	Sites.Manage.All
	Sites.ReadWrite.All
	TermStore.ReadWrite.All
	User.ReadWrite.All

5. Grant admin consent for the API permissions.

1. Under **Configured permissions**, click **Grant admin consent**.
2. On the Grant admin consent confirmation, click **Yes**.

Refresh Got feedback?

Grant admin consent confirmation.

Do you **Yes** to grant consent for the requested permissions for all accounts in cohesitydmas? This will update any existing admin consent records this application already has to match what is listed below.

Yes **No**

Confirmed permissions

Applications are authorized to call APIs when they are granted permissions by users/admins as part of the consent process. The list of configured permissions should include all the permissions the application needs. [Learn more about permissions and consent](#)

+ Add a permission ✓ Grant admin consent for cohesitydmas

API / Permissions name	Type	Description	Admin consent requ...	Status
Microsoft Graph (13)				...
Channel.Create	Application	Create channels	Yes	Not granted for cohesit... ***
Channel.ReadBasic.All	Application	Read the names and descriptions of all channels	Yes	Not granted for cohesit... ***
ChannelMember.ReadWrite.All	Application	Add and remove members from all channels	Yes	Not granted for cohesit... ***
Directory.ReadWrite.All	Application	Read and write directory data	Yes	Not granted for cohesit... ***
Files.ReadWrite.All	Application	Read and write files in all site collections	Yes	Not granted for cohesit... ***
Group.Create	Application	Create groups	Yes	Not granted for cohesit... ***
Group.ReadWrite.All	Application	Read and write all groups	Yes	Not granted for cohesit... ***
Reports.Read.All	Application	Read all usage reports	Yes	Not granted for cohesit... ***
Site.FullControl.All	Application	Have full control of all site collections	Yes	Not granted for cohesit... ***

6. Create a new client secret that will be used to register Microsoft 365 as a source in Cohesity DataProtect as a Service.
 1. Under the **Manage** section, select **Certificates & secrets**.
 1. In the Client secrets section, click **New client secret**. Enter a **Description**.
 2. In the **Expires** drop-down, select how long the secret key will be valid.

3. Click **Add**.

Add a client secret

Description: Cohesity DataProtect Secret Key

Expires: Recommended: 6 months (selected)

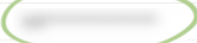


Other options: Recommended: 6 months, 3 months, 12 months, 18 months, 24 months, Custom

2. Under **Client secrets**, click the **Copy** button next to the string under **VALUE**. You need the Value key of the client secret to register Microsoft 365 as a source in Cohesity DataProtect as a Service.
3. Store the Value key in a secure location. After you exit this page, you will not be able to see the Value key again. If you lose your value key, you will need to create a new client secret.

Client secrets

A secret string that the application uses to prove its identity when requesting a token. Also can be referred to as application password.

[+ New client secret](#)

Description	Expires	Value	Secret ID
Cohesity DataProtect Secret Key	10/14/2023		aa0d893d-08d7-44df-a42b-2090bde65c43  

When you finish, your custom Azure app should include the permissions as shown below.

API / Permissions name	Type	Description	Admin consent requ...	Status
▼ Microsoft Graph (15)				
Channel.Create	Application	Create channels	Yes	✔ Granted for Cohesity ***
Channel.ReadBasic.All	Application	Read the names and descriptions of all channels	Yes	✔ Granted for Cohesity ***
ChannelMember.ReadWrite.All	Application	Add and remove members from all channels	Yes	✔ Granted for Cohesity ***
ChannelMessage.Read.All	Application	Read all channel messages	Yes	✔ Granted for Cohesity ***
Chat.Read.All	Application	Read all chat messages	Yes	✔ Granted for Cohesity ***
Directory.ReadWrite.All	Application	Read and write directory data	Yes	✔ Granted for Cohesity ***
Files.ReadWrite.All	Application	Read and write files in all site collections	Yes	✔ Granted for Cohesity ***
Group.Create	Application	Create groups	Yes	✔ Granted for Cohesity ***
Group.ReadWrite.All	Application	Read and write all groups	Yes	✔ Granted for Cohesity ***
MailboxSettings.Read	Application	Read all user mailbox settings	Yes	✔ Granted for Cohesity ***
Reports.Read.All	Application	Read all usage reports	Yes	✔ Granted for Cohesity ***
Sites.FullControl.All	Application	Have full control of all site collections	Yes	✔ Granted for Cohesity ***
Sites.ReadWrite.All	Application	Read and write items in all site collections	Yes	✔ Granted for Cohesity ***
User.Read.All	Application	Read all users' full profiles	Yes	✔ Granted for Cohesity ***
User.ReadWrite.All	Application	Read and write all users' full profiles	Yes	✔ Granted for Cohesity ***
▼ Office 365 Exchange Online (1)				
full_access_as_app	Application	Use Exchange Web Services with full access to all mailboxes	Yes	✔ Granted for Cohesity ***
▼ SharePoint (12)				
AllSites.FullControl	Delegated	Have full control of all site collections	Yes	✔ Granted for Cohesity ***
AllSites.Manage	Delegated	Read and write items and lists in all site collections	No	✔ Granted for Cohesity ***
MyFiles.Read	Delegated	Read user files	No	✔ Granted for Cohesity ***
MyFiles.Write	Delegated	Read and write user files	No	✔ Granted for Cohesity ***
Sites.FullControl.All	Application	Have full control of all site collections	Yes	✔ Granted for Cohesity ***
Sites.Manage.All	Application	Read and write items and lists in all site collections	Yes	✔ Granted for Cohesity ***
Sites.ReadWrite.All	Application	Read and write items in all site collections	Yes	✔ Granted for Cohesity ***
Sites.Search.All	Delegated	Run search queries as a user	Yes	✔ Granted for Cohesity ***
TermStore.ReadWrite.All	Application	Read and write managed metadata	Yes	✔ Granted for Cohesity ***
TermStore.ReadWrite.All	Delegated	Read and write managed metadata	Yes	✔ Granted for Cohesity ***
User.ReadWrite.All	Application	Read and write user profiles	Yes	✔ Granted for Cohesity ***
User.ReadWrite.All	Delegated	Read and write user profiles	Yes	✔ Granted for Cohesity ***

Set Additional Permissions for SharePoint Online

For SharePoint Online data protection, ensure that you set the required add-in permissions and tenant permissions below.

When you finish, your custom Azure app should include the permissions as shown below.

Add-In Permissions in SharePoint Online

Make sure that you assign the following add-in permissions to the custom app. For more information, see [Add-in permissions in SharePoint](#) in the Microsoft documentation.

Scope URI	Required Rights
http://sharepoint/content/tenant	FullControl
http://sharepoint/content/sitecollection	FullControl
http://sharepoint/content/sitecollection/web	FullControl
http://sharepoint/content/sitecollection/web/list	FullControl
http://sharepoint/taxonomy	Read,Write

Tenant Permissions

After you have [registered the custom app](#), configure the tenant permissions on the custom app.

To configure the tenant permissions:

1. Launch the **SharePoint Admin Center** using the URL: https://<your-tenant>-admin.sharepoint.com/_layouts/15/AppInv.aspx
2. In the **SharePoint Admin Center**, log in as the tenant admin.
3. In the **App ID and Title** section, perform the following:
 1. In the **App Id** field, enter the **AppID** of the custom app you have created and click **Lookup** to search for the custom app.
 2. In the **App Domain** field, enter www.localhost.com as the app domain.

Important: Do not enter any other string other than www.localhost.com in the **App Domain** field.

3. In the **Redirect URL** field, enter <https://localhost.com/> as the redirect URL.

Important: Do not enter any other URL other than <https://localhost.com/> in the **Redirect URL** field.

4. In the **Permission Request XML** field, enter the following values:

```
<AppPermissionRequests AllowAppOnlyPolicy="true">
  <AppPermissionRequest Scope="http://sharepoint/content/tenant"
    Right="FullControl" />
</AppPermissionRequest>
```

```

Scope="http://sharepoint/content/sitecollection"
Right="FullControl" />
<AppPermissionRequest
Scope="http://sharepoint/content/sitecollection/web"
Right="FullControl" />
<AppPermissionRequest
Scope="http://sharepoint/content/sitecollection/web/list"
Right="FullControl" />
<AppPermissionRequest Scope="http://sharepoint/taxonomy"
Right="Read,Write" />
</AppPermissionRequests>

```

App Configuration for SharePoint Online

App Id and Title
The app's identity and its title.

App Id:

Title:

App Domain:
Example: "www.contoso.com"

Redirect URL:
Example: "https://www.contoso.com/default.aspx"

App's Permission Request XML
The permission required by the app.

Permission Request XML:

```

<AppPermissionRequests AllowAppOnlyPolicy="true">
  <AppPermissionRequest Scope="http://sharepoint/content/tenant" Right="FullControl" />
  <AppPermissionRequest Scope="http://sharepoint/content/sitecollection" Right="FullControl" />
  <AppPermissionRequest Scope="http://sharepoint/content/sitecollection/web" Right="FullControl" />
  <AppPermissionRequest Scope="http://sharepoint/content/sitecollection/web/list" Right="FullControl" />
  <AppPermissionRequest Scope="http://sharepoint/taxonomy" Right="Read,Write" />
</AppPermissionRequests>

```

4. Click **Create**.
5. In the **Do you trust <app_title>?** page, perform the following:
 - a. From the drop-down, select **DO_NOT_DELETE_SPLIST_TENANTADMIN_AGGREGATED_SITECOLLECTIONS**.
 - b. click **Trust It**.

Important: If you have created your Microsoft 365 tenant on or after Sep 20, 2020, you must install SharePoint Online PowerShell. Using the global administrator account, run the following commands in an administrator PowerShell session:

```
Get-Module -Name Microsoft.Online.SharePoint.PowerShell -
ListAvailable | Select Name,Version
[Net.ServicePointManager]::SecurityProtocol =
[Net.SecurityProtocolType]::Tls12
Install-Module -Name Microsoft.Online.SharePoint.PowerShell -Scope
AllUsers
Connect-SPOService -Url 'https://<tenant>-admin.sharepoint.com'
Set-SPOTenant -DisableCustomAppAuthentication $False
```

Note: Custom scripts setting is not supported in SharePoint.

Enhance Backup of Large Microsoft 365 Data

When it comes to ingesting Microsoft 365 data, the size can vary from hundreds of terabytes (TBs) to a few petabytes (PBs) in some cases. One of the issues that may arise while performing the initial full backup of such large data is throttling from the Microsoft 365 APIs.

Cohesity recommends using only one application ID for Microsoft 365 backups. However, when you back up a large amount of data, using only one application ID may result in a prolonged ETA for the first full backup. Contact your Cohesity account team to determine the appropriate number of application IDs to optimize the performance while ensuring that the Microsoft Tenant Level Throttling limits are not exceeded, which may cause service failures for Microsoft 365.

After completing the first full backup, using multiple application IDs is not usually required. However, in certain situations, a single application ID may be inadequate. For guidance on the appropriate number of application IDs for incremental backups following the initial full backup, Contact your Cohesity account team.

Cohesity also supports Microsoft paid APIs to perform the first full backups faster. These APIs come with higher limits, ensuring the backups are performed faster at a higher API rate. The paid APIs are configured to be used in addition to the existing backup capacity allocated by Microsoft, rather than as a substitute, to help minimize additional costs.

Microsoft 365 Backup Storage

Note:

- This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.
- Microsoft 365 Backup Storage preview/free trial is available to protect only 75 objects (25 each for the Exchange Online Mailboxes, OneDrive, and SharePoint Online respectively) for 30 days without any cost.

Cohesity DataProtect as a Service supports Microsoft 365 Backup Storage service on Exchange Online Mailboxes, OneDrive, and SharePoint Online. This feature provides an ultra-fast backup and recovery capability by creating backups within the Microsoft 365 protected services data boundary.

For more information, see [Overview for third-party developers of Microsoft 365 Backup Storage](#).

Prerequisites

- Ensure that at least one SharePoint license is available to protect your Microsoft 365 data using the Microsoft 365 Backup Storage service.
- Ensure that the [onboarding](#) procedure is performed before using the Microsoft 365 Backup Storage service.
- Ensure that consent is provided for the [additional permissions](#) during source registration.

Considerations

- The Microsoft API limitations also apply to Cohesity DataProtect as a Service.
- Microsoft 365 Backup Storage service is only supported for Exchange Online Mailboxes, OneDrive, and SharePoint Online.
- Backed up Microsoft 365 data is stored in Microsoft.
- The following are not supported for Microsoft 365 backup storage:
 - Teams and Groups.
 - Shared Mailboxes.
 - Granular search and recovery for Exchange Online Mailboxes, SharePoint, and OneDrive.
 - Variable Retention Periods.

- Browse functionality on snapshots.
- Insights into backup compliance.
- Deletion of backed up data.
- Automatic protection for newly created Mailboxes, OneDrives, or SharePoint Sites.
- Recovery of site content from recycle bin.
- Global search.
- Latest snapshot recovery.

Additional Permissions

The following are the additional permissions required to use Microsoft 365 Backup Storage:

- [BackupRestore-Configuration.Read.All](#)
- [BackupRestore-Configuration.ReadWrite.All](#)
- [BackupRestore-Restore.Read.All](#)
- [BackupRestore-Restore.ReadWrite.All](#)
- [BackupRestore-Search.Read.All](#)
- [BackupRestore-Control.ReadWrite.All](#)

Ensure that consent is provided for these additional permissions during source registration.

Onboarding Microsoft 365 Backup Storage APIs

The onboarding process includes the following steps:

- Register Cohesity App as the M365 Backup Controller.
- Activate the M365 Backup Controller.
- Enable the billing.

The following is the procedure to onboard and offboard the Microsoft 365 Backup Storage APIs using the tool:

1. Download the **m365_backup_storage_setup.exe** file. For more information, contact your Cohesity account team.
2. Execute the **m365_backup_storage_setup.exe** file.
3. Provide the **Microsoft Tenant ID** and press **Enter**.
4. Select a workflow type to proceed: Enter A to proceed with Microsoft365 Backup Storage Onboarding.

The device code authorization flow will be used and the device code will be displayed.

5. Copy the **User Code**, click the <https://microsoft.com/devicelogin> link, and paste the code in the **Enter Code** window.
6. Click an **account** to sign in to the Cohesity M365 Backup Storage Controller, and click **Continue**.
The access token will be acquired by the tool and displayed.
7. Select a workflow to proceed:
 - Enter **1** to register a new Backup Controller for a new user
 - Enter **2** to view the registered service apps
 - Enter **3** to activate an existing Backup Controller
 - Enter **4** to deactivate an existing Backup Controller
 - Enter **5** to unregister an existing Backup Controller
8. Enter **Y** or **N** when prompted, to proceed according to the required flow.

Protect Microsoft 365 Data Using Microsoft 365 Backup Storage

The following is the procedure to protect Microsoft 365 objects using the Microsoft 365 Backup Storage APIs:

Express Registration Using Microsoft 365 Backup Storage

Before you register your Microsoft 365 domain, ensure that you have:

- [Added roles to the Microsoft 365 user account.](#)
- [Updated your Microsoft Organization setting for Mailbox size reporting.](#)

Note: Basic Auth is not supported for Microsoft 365 source registration.


To register your Microsoft 365 domain:

1. In **DataProtect as a Service**, navigate to the **Sources** page and click **+ Register Source** in the upper-right corner of the page.
2. In the **Select Source** dialog box, select **Microsoft 365** and click **Start Registration**.


Select Source

SaaS Sources

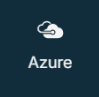
On-Premise Sources




Microsoft 365




AWS




Azure




Virtual Machine



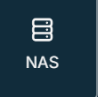
SAP HANA x86-x64



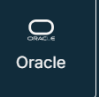
Physical



Microsoft SQL



NAS



Oracle

To register a Microsoft 365 source, you must have:

- The Microsoft Azure account with global administrator access to provision the Azure Apps for automatic Microsoft 365 source registration (recommended). To register the source manually, [register a custom Azure App](#). You can use a Microsoft 365 account which is a service account and not an administrative account.
- Microsoft Azure apps on the Azure portal with necessary permissions to perform object discovery, backup, and recovery in Microsoft 365.
- Additional [Microsoft Organization setting](#) for Mailbox size reporting.

[Learn more](#)

Cancel
Start Registration

3. In the **Source Details** section, select a cloud region for your data backups.
4. Choose the **Microsoft 365 Applications** to discover.

Note: Discovery selection change is not allowed for applications with protected objects.

Note: If the **Private Chats** and **Teams Posts** option is enabled under the **Mailbox** and **Teams** apps respectively, the Private Chats and Teams Posts will be backed up along with the corresponding Users and Teams respectively.

Private Chats and Teams Posts backup APIs are charged separately by Microsoft. All Azure apps configured by Cohesity must be linked to an [Azure subscription for billing](#). Backups of Mailboxes and Teams may fail if the Azure subscription configuration is not set appropriately.

This is an Early Access feature. Contact your Cohesity account team to enable the feature.

5. *[Optional]* You can enable the below options based on your requirement:

Note: Retaining the default values will speed up object discovery in the environment. The metadata required (which is fetched through these options) will be fetched during the backup of the objects and updated.

1. **Fetch Mailbox Info** to fetch and process the Mailbox information including the provisioning status, mailbox type, and in-place archival usage.

Note: You can enable this option to discover the Mailboxes of the users that were converted into Shared Mailboxes by revoking the user Exchange Online licenses.

2. **Fetch OneDrive Info** to fetch and process the OneDrive information including the provisioning status and storage quota.
3. **Include Users without MySite** to include users who have unprovisioned OneDrive or do not have MySite.
4. **Enable Site Tagging** to tag SharePoint Sites whether they are a Group Site or a Teams Site.

Note: Any Site that is tagged as a Group or Team will not be visible in the Sites section and these sites will be protected through the corresponding Group or Team protection.

6. In the **Account Credentials** section, enter the Microsoft 365 **Username**.
7. In the **Azure Applications** section, enter the number of Azure applications that you want to create based on your requirements and click **Create**.

Note: By default, an Azure application will be created. To better manage Microsoft 365 throttling, Cohesity recommends at least one Azure app.

8. In the **Add Azure Application** form, copy the device code and click the **Microsoft Azure App** link to open the Microsoft Azure App authorization service in a new tab.

Note: If you prefer to create your Azure apps manually, see [Manual Registration for Microsoft 365 Sources](#).

9. In the **Microsoft Azure App authorization** service, paste the copied code and click **Next**.

10. Log in to Microsoft Azure, enter the **Username** and **Password** of your Microsoft 365 account and click **Sign in**.

Note: Ensure that your Microsoft 365 account has global administrator access.

11. Follow the instructions to complete the authorization on the Microsoft Azure portal.
12. Wait for Microsoft Azure Authorization to complete.
13. To use the **Microsoft 365 Backup Storage** services, toggle *ON* the **Enable Microsoft 365 Backup Storage** option and click **Provide Consent**.

Note:

- Consent is required to connect verified Cohesity App to orchestrate backup via Microsoft 365 Backup Storage APIs and register Microsoft 365 Backup Controller.
- This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

Register Microsoft 365 Source

Username ✓
Enter full UPN (Eg: username@company.com)

Azure Applications

Number of Azure Apps to be created [Create](#)

Microsoft 365 Backup Storage ⓘ

☒ Enable Microsoft 365 Backup Storage

Consent is required to connect verified Cohesity App to orchestrate backup via Microsoft 365 Backup Storage APIs and register Microsoft 365 Backup Controller. [Learn more](#)

[Provide Consent](#)

[Cancel](#) [Register](#)

14. Sign in to Microsoft Azure, enable the **Consent on behalf of your organization** option, and click **Accept**.
15. Wait for Microsoft 365 Backup Storage service to be enabled.
16. Click **Register**.

Note: Once consent is provided and the source is registered/updated, ensure that the [onboarding procedure](#) is performed before using Microsoft 365 Backup Storage.

For SharePoint Online data protection, ensure that you set the required add-in permissions and tenant permissions on the Azure application. For more information, see [Set additional permissions for SharePoint Online](#).

You can follow the Microsoft 365 source discovery and registration progress on the **Sources** page.

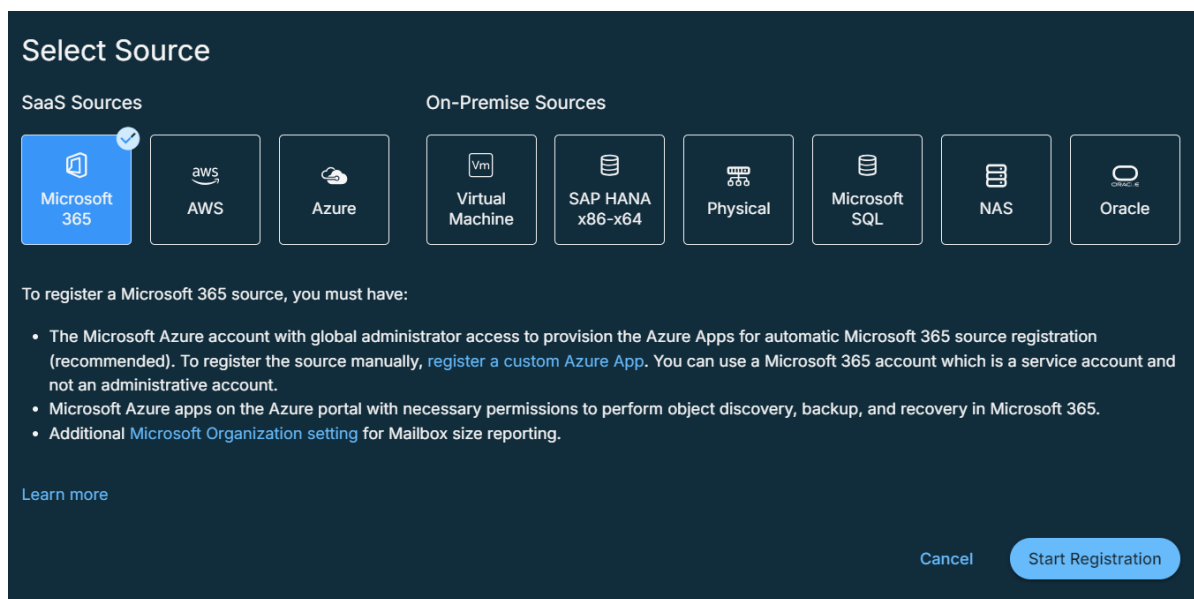
Next > You are now ready to protect your Microsoft 365 [Mailboxes](#), [OneDrives](#), [SharePoint Online Sites](#), and [Teams](#)!

Manual Registration Using Microsoft 365 Backup Storage

Note: Basic Auth is not supported for Microsoft 365 source registration.

To register your Microsoft 365 domain manually, make sure you've met all the [Microsoft 365 Requirements](#) and then:

1. In **DataProtect as a Service**, navigate to the **Sources** page and click **+ Register Source** in the upper-right corner of the page.
2. In the **Select Source** dialog box, select **Microsoft 365** and click **Start Registration**.



3. In the **Source Details** section, select a cloud region for your data backups.
4. Choose the **Microsoft 365 Applications** to discover.

Note: Discovery selection change is not allowed for applications with protected objects.

Note: If the **Private Chats** and **Teams Posts** option is enabled under the **Mailbox** and **Teams** apps respectively, the Private Chats and Teams Posts will be backed up along with the corresponding Users and Teams respectively.

Private Chats and Teams Posts backup APIs are charged separately by Microsoft. All Azure apps configured by Cohesity must be linked to an [Azure subscription for billing](#). Backups of Mailboxes and Teams may fail if the Azure subscription configuration is not set appropriately.

This is an Early Access feature. Contact your Cohesity account team to enable the feature.

5. *[Optional]* You can enable the below options based on your requirement:

Note: Retaining the default values will speed up object discovery in the environment. The metadata required (which is fetched through these options) will be fetched during the backup of the objects and updated.

1. **Fetch Mailbox Info** to fetch and process the Mailbox information including the provisioning status, mailbox type, and in-place archival usage.

Note: You can enable this option to discover the Mailboxes of the users that were converted into Shared Mailboxes by revoking the user Exchange Online licenses.

2. **Fetch OneDrive Info** to fetch and process the OneDrive information including the provisioning status and storage quota.
3. **Include Users without MySite** to include users who have unprovisioned OneDrive or do not have MySite.
4. **Enable Site Tagging** to tag SharePoint Sites whether they are a Group Site or a Teams Site.

Note: Any Site that is tagged as a Group or Team will not be visible in the Sites section and these sites will be protected through the corresponding Group or Team protection.

6. In the **Account Credentials** section, enter the Microsoft 365 **Username**.

7. In the **Azure Applications** section, enter the number of Azure applications that you want to create based on your requirements and click **Create**.

Note: By default, an Azure application will be created. To better manage Microsoft 365 throttling, Cohesity recommends at least one Azure app.

8. In the **Add Azure Application** form, click the **You can also add Azure App manually** link and then enter the **App ID** and **App Secret Key** that you noted down while registering your custom Azure app.

Tip: You can add multiple Azure apps for a Microsoft 365 source to load balance your backup and restore operations. Click **+** to add multiple Azure apps. When you do, ensure that you provide the valid **App ID** and **App Secret Key**.

9. To use the **Microsoft 365 Backup Storage** services, toggle *ON* the **Enable Microsoft 365 Backup Storage** option and click **Provide Consent**.

Note:

- Consent is required to connect verified Cohesity App to orchestrate backup via Microsoft 365 Backup Storage APIs and register Microsoft 365 Backup Controller.
- This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

Register Microsoft 365 Source

Username ✓

Enter full UPN (Eg: username@company.com)

Azure Applications

Number of Azure Apps to be created [Create](#)

Microsoft 365 Backup Storage ⓘ

☒ Enable Microsoft 365 Backup Storage

Consent is required to connect verified Cohesity App to orchestrate backup via Microsoft 365 Backup Storage APIs and register Microsoft 365 Backup Controller. [Learn more](#)

[Provide Consent](#)

[Cancel](#) [Register](#)

10. Sign in to Microsoft Azure, enable the **Consent on behalf of your organization** option, and click **Accept**.
11. Wait for Microsoft 365 Backup Storage service to be enabled.
12. Click **Register**.

Note: Once consent is provided and the source is registered/updated, ensure that the [onboarding procedure](#) is performed before using Microsoft 365 Backup Storage.

For SharePoint Online data protection, ensure that you set the required add-in permissions and tenant permissions on the Azure application. For more information, see [Set additional permissions for SharePoint Online](#).

You can follow the Microsoft 365 source discovery and registration progress on the **Sources** page.

Next > You are now ready to protect your Microsoft 365 [Mailboxes](#), [OneDrives](#), [SharePoint Online Sites](#), and [Teams](#)!

Protect Mailboxes Using Microsoft 365 Backup Storage

Once you have registered your Microsoft 365 domain as a source, you're ready to protect the user Mailboxes in your domain.

To protect your Microsoft 365 Mailboxes:

Note: If the **Private Chats** option is enabled under Mailbox during app registration, the Private Chats will be backed up along with the corresponding Users.

Mailboxes backup may fail if the Azure subscription configuration is not set appropriately.

This is an Early Access feature. Contact your Cohesity account team to enable the feature.

1. In **DataProtect as a Service**, under **Sources**, find the Microsoft 365 source and click on it.
2. Click either of the following tabs:
 - **Cohesity Backup Storage** to use the Cohesity Backup Storage services for protection.
 - **Microsoft 365 Backup Storage** to use the Microsoft 365 Backup Storage services for protection.

Note:

- Free trial subscription only supports protecting a maximum of 25 objects.
- Ensure that the [onboarding procedure](#) is performed before using Microsoft 365 Backup Storage.
- This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

3. You can click the [Settings](#) tab to update the source details.
4. Click the **Mailbox** tab.
5. Select the individual Mailboxes you wish to protect or:
 - Click **Users > Select All Child Objects** to protect all the Mailboxes in this source.

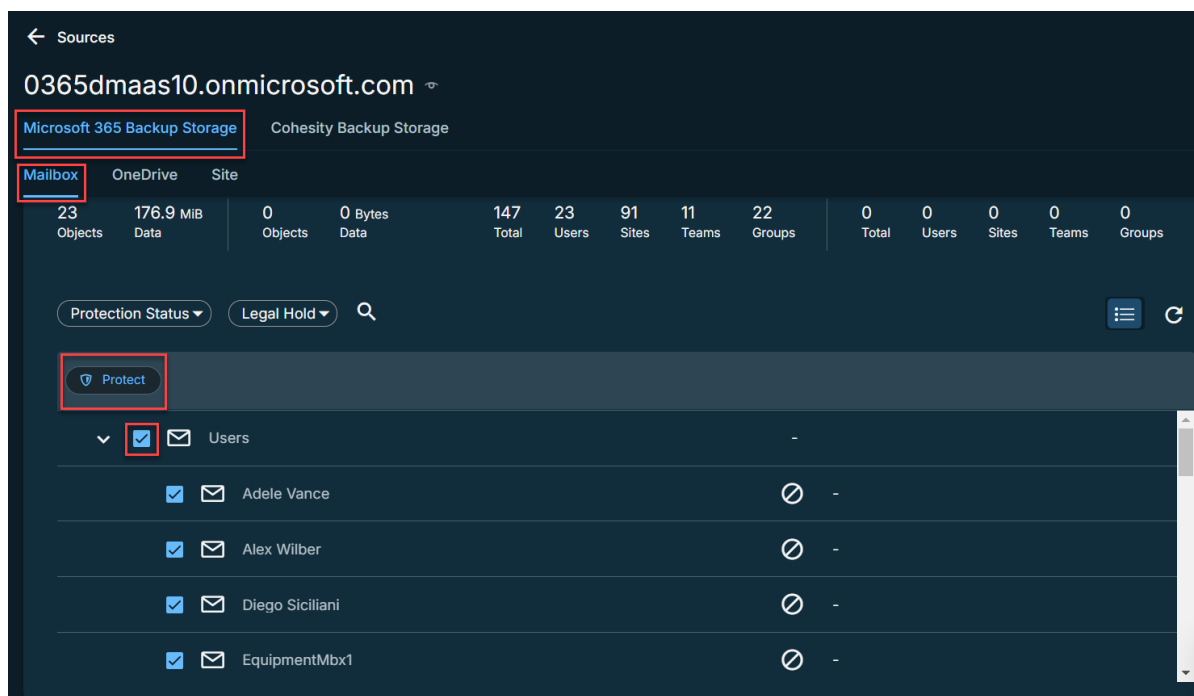
- Click **Users > Auto Protect This** to protect all the Mailboxes *plus any future additional Mailboxes* on that source.
- Click the **Security Groups** icon and select the security group to protect the Mailboxes of the users in the security group. For more information, see [Security Groups](#).

Cohesity supports only users and not devices in Security Groups.

This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

6. Click the **Protect** icon above the list.

Note: The option to select the policy is not available for protection using the Microsoft 365 Backup Storage service. This service uses the [Microsoft policies](#) by default.



7. Click **Protect**.

Note: The backups start immediately after you protect the objects, regardless of the time you set for the protection run.

Settings

Settings	Description
Type	Type of the domain.
IP or FQDN	IP address of the server or FQDN of the server.
Username	Microsoft 365 Username.
Registered	Timestamp when the source was registered.
Refreshed	Timestamp when the source was last refreshed.
Maintenance Mode	Indicates if maintenance mode is turned ON or OFF.

Next > When the first protection run completes, you will be ready to recover your protected Mailboxes when and if you need to.

Protect OneDrives Using Microsoft 365 Backup Storage

Once you have registered your Microsoft 365 domain as a source, you're ready to use Cohesity DataProtect as a Service to protect the user OneDrives in your domain.

To protect your Microsoft 365 OneDrives:

1. In **DataProtect as a Service**, under **Sources**, find the Microsoft 365 source and click on it.
2. Click either of the following tabs:
 - **Cohesity Backup Storage** to use the Cohesity Backup Storage services for protection.
 - **Microsoft 365 Backup Storage** to use the Microsoft 365 Backup Storage services for protection.

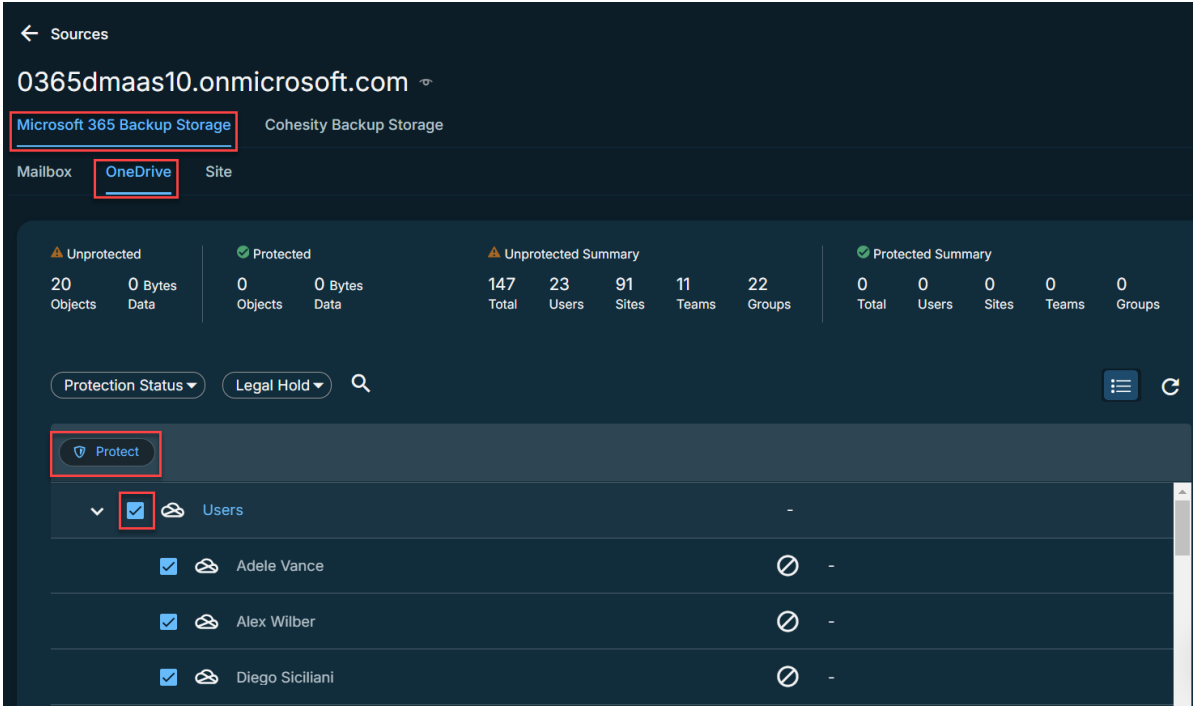
Note:

- Free trial subscription only supports protecting a maximum of 25 objects.
- Ensure that the [onboarding procedure](#) is performed before using Microsoft 365 Backup Storage.
- This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

3. You can click the [Settings](#) tab to update the source details.
4. Click the **OneDrive** tab.
5. Select the individual OneDrives you wish to protect or:
 - Click **Users > Select All Child Objects** to protect all the OneDrives in this source.
 - Click **Users > Auto Protect This** to protect all the OneDrives *plus any future additional OneDrives* on that source.
 - Click the **Security Groups** icon and select the security group to protect the OneDrives of the users in the security group. For more information, see [Security Groups](#).

Cohesity supports only users and not devices in Security Groups.
This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.
6. Click the **Protect** icon above the list.

Note: The option to select the policy is not available for protection using the Microsoft 365 Backup Storage service. This service uses the [Microsoft policies](#) by default.



7. Click **Protect**.

Note: The backups start immediately after you protect the objects, regardless of the time you set for the protection run.

Settings

Settings	Description
Type	Type of the domain.
IP or FQDN	IP address of the server or FQDN of the server.
Username	Microsoft 365 Username.
Registered	Timestamp when the source was registered.
Refreshed	Timestamp when the source was last refreshed.
Maintenance Mode	Indicates if maintenance mode is turned ON or OFF.

Next > When the first protection run completes, you will be ready to recover your protected OneDrives when and if you need to.

Protect SharePoint Sites Using Microsoft 365 Backup Storage

Once you have registered your Microsoft 365 domain as a source, you're ready to protect the SharePoint Online sites in your domain.

To protect your Microsoft 365 SharePoint Online sites:

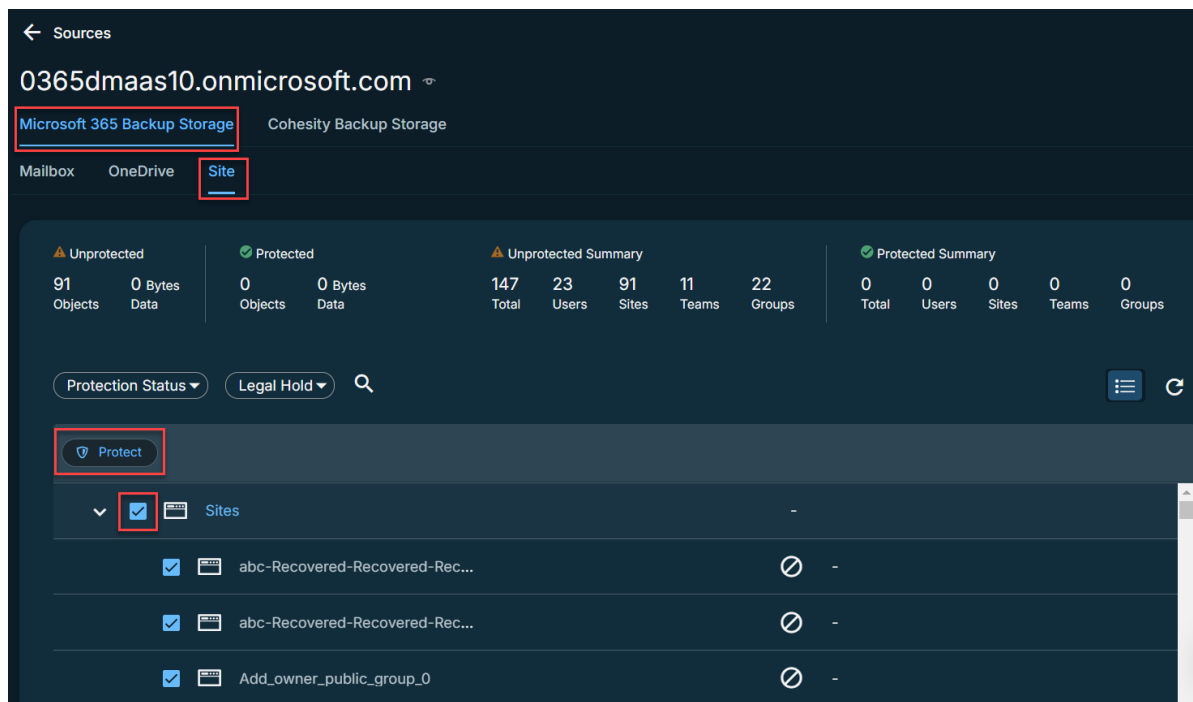
1. In **DataProtect as a Service**, under **Sources**, find the Microsoft 365 source and click into it.
2. Click either of the following tabs:
 - **Cohesity Backup Storage** to use the Cohesity Backup Storage services for protection.
 - **Microsoft 365 Backup Storage** to use the Microsoft 365 Backup Storage services for protection.

Note:

- Free trial subscription only supports protecting a maximum of 25 objects.
- Ensure that the [onboarding procedure](#) is performed before using Microsoft 365 Backup Storage.
- This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

3. You can click the [Settings](#) tab to update the source details.
4. Click the **Site** tab.
5. Select the individual SharePoint Online site you wish to protect or:
 - Click **Users > Select All Child Objects** to protect all the SharePoint Online sites in this source.
 - Click **Users > Auto Protect This** to protect all the SharePoint Online sites in this source.
6. Click the **Protect** icon above the list.

Note: The option to select the policy is not available for protection using the Microsoft 365 Backup Storage service. This service uses the [Microsoft policies](#) by default.



7. Click **Protect**.

Note: The backups start immediately after you protect the objects, regardless of the time you set for the protection run.

Settings

Settings	Description
Type	Type of the domain.
IP or FQDN	IP address of the server or FQDN of the server.
Username	Microsoft 365 Username.
Registered	Timestamp when the source was registered.
Refreshed	Timestamp when the source was last refreshed.
Maintenance Mode	Indicates if maintenance mode is turned ON or OFF.

Next > When the first protection run completes, you will be ready to recover your protected SharePoint Online sites when and if you need to.

Recover Microsoft 365 Applications Using Microsoft 365 Backup Storage

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

Recover Microsoft 365 User Mailboxes

To recover protected Microsoft 365 user Mailboxes, see [Recover User Mailboxes](#).

Note: If the mailbox is backed up using the Microsoft 365 Backup Storage service, when you choose New Location, the data will be recovered to a new folder created by Microsoft in the same mailbox. This folder name is created as per the Microsoft policy and cannot be edited.

Recover OneDrives

To recover protected Microsoft 365 user OneDrives, see [Recover User OneDrives](#).

Note: If the OneDrive is backed up using the Microsoft 365 Backup Storage service, when you choose New Location, the data will be recovered to a new OneDrive created by Microsoft. This OneDrive name is created as per the Microsoft policy and cannot be edited.

Recover SharePoint Sites

To recover protected Microsoft 365 SharePoint Online sites, see [Recover SharePoint Sites](#).

Note: If the site is backed up using the Microsoft 365 Backup Storage service, when you choose New Location, the data will be recovered to a new site created by Microsoft. This site name is created as per the Microsoft policy and cannot be edited.

Register Microsoft 365 Sources

To start protecting Microsoft 365 applications, you need to register the Microsoft 365 domain as a source in Cohesity DataProtect as a Service.

Cohesity DataProtect as a Service uses the [Microsoft Graph API](#) for object discovery, backup, and recovery in Microsoft 365. To use the Graph API, Cohesity DataProtect as a

Service uses an Azure application created and registered on the Azure portal with necessary permissions. You can either let Cohesity [create the Azure application](#) or [manually enter Azure application](#) details while registering your Microsoft 365 domain as a source in Cohesity DataProtect as a Service.

Express Registration for Microsoft 365 Sources

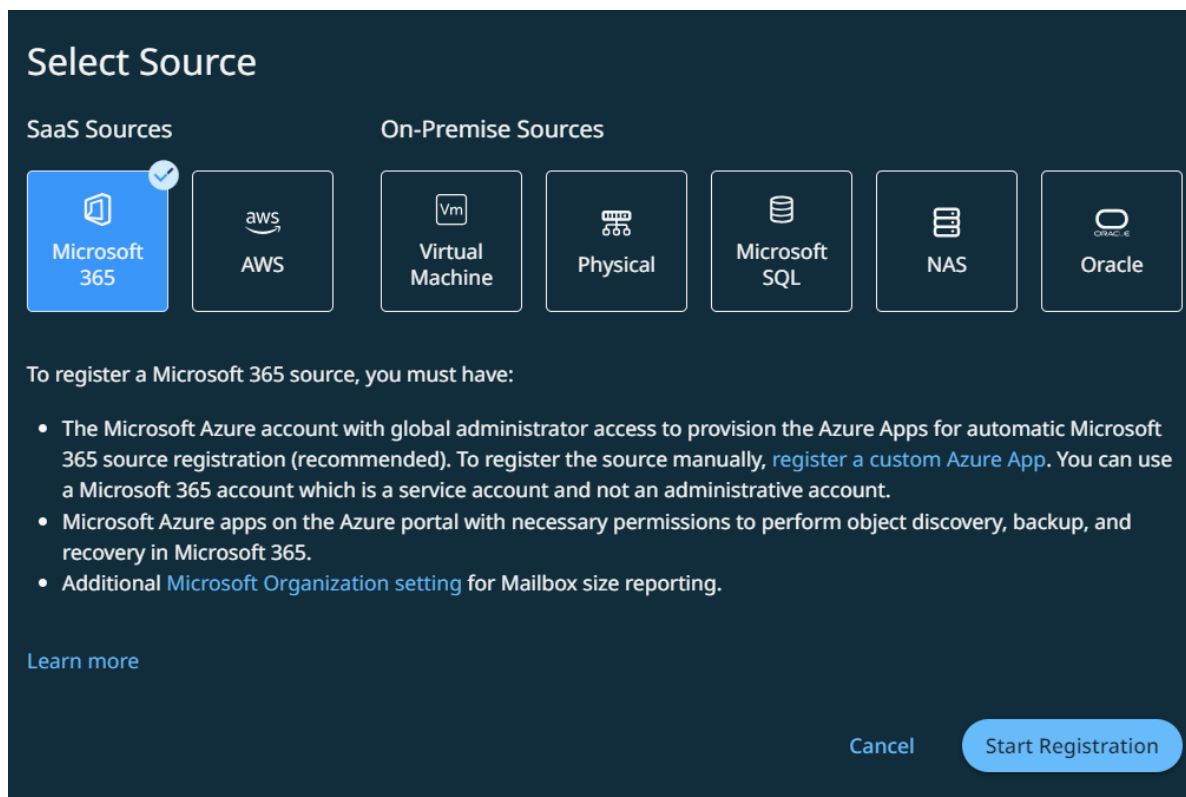
Before you register your Microsoft 365 domain, ensure that you have:

- [Added roles to the Microsoft 365 user account.](#)
- [Updated your Microsoft Organization setting for Mailbox size reporting.](#)

Note: Basic Auth is not supported for Microsoft 365 source registration.

To register your Microsoft 365 domain:

1. In **DataProtect as a Service**, navigate to the **Sources** page and click **+ Register Source** in the upper-right corner of the page.
2. In the **Select Source** dialog box, select **Microsoft 365** and click **Start Registration**.



3. In the **Source Details** section, select a cloud region for your data backups.

4. Choose the **Microsoft 365 Applications** to discover.

Note: Discovery selection change is not allowed for applications with protected objects.

Note: If the **Private Chats** and **Teams Posts** option is enabled under the **Mailbox** and **Teams** apps respectively, the Private Chats and Teams Posts will be backed up along with the corresponding Users and Teams respectively.

Private Chats and Teams Posts backup APIs are charged separately by Microsoft. All Azure apps configured by Cohesity must be linked to an [Azure subscription for billing](#). Backups of Mailboxes and Teams may fail if the Azure subscription configuration is not set appropriately.

This is an Early Access feature. Contact your Cohesity account team to enable the feature.

5. *[Optional]* You can enable the below options based on your requirement:

Note: Retaining the default values will speed up object discovery in the environment. The metadata required (which is fetched through these options) will be fetched during the backup of the objects and updated.

1. **Fetch Mailbox Info** to fetch and process the Mailbox information including the provisioning status, mailbox type, and in-place archival usage.

Note: You can enable this option to discover the Mailboxes of the users that were converted into Shared Mailboxes by revoking the user Exchange Online licenses.

2. **Fetch OneDrive Info** to fetch and process the OneDrive information including the provisioning status and storage quota.
3. **Include Users without MySite** to include users who have unprovisioned OneDrive or do not have MySite.
4. **Enable Site Tagging** to tag SharePoint Sites whether they are a Group Site or a Teams Site.

Note: Any Site that is tagged as a Group or Team will not be visible in the Sites section and these sites will be protected through the corresponding Group or Team protection.

6. In the **Account Credentials** section, enter the full **Username** of the Microsoft 365 user account with a valid SharePoint and Mailbox license.
7. In the **Azure Applications** section, enter the number of Azure applications that you want to create based on your requirements and click **Create**.

Register Microsoft 365 Source

☒ Fetch Mailbox Info ⓘ

☒ Fetch OneDrive Info ⓘ

☒ Include Users without MySite ⓘ

☒ Enable Site Tagging ⓘ

Account Credentials

Username

dmas-u-automation4-v2@cohesity.com ✓

Enter full UPN (Eg: username@company.com)

Azure Applications

Number of Azure Apps to be created 1 **Create**

Cancel Register

Note: By default, an Azure application will be created. To better manage Microsoft 365 throttling, Cohesity recommends at least one Azure app.

8. In the **Add Azure Application** form, copy the device code and click the **Microsoft Azure App** link to open the Microsoft Azure App authorization service in a new tab.

Note: If you prefer to create your Azure apps manually, see [Manual Registration for Microsoft 365 Sources](#).

9. In the **Microsoft Azure App authorization** service, paste the copied code and click **Next**.
10. Log in to Microsoft Azure, enter the **Username** and **Password** of your Microsoft 365 account and click **Sign in**.

Note: Ensure that your Microsoft 365 account has global administrator access.

11. Follow the instructions to complete the authorization on the Microsoft Azure portal.
12. Wait for Microsoft Azure Authorization to complete and then click **Register**.

For SharePoint Online data protection, ensure that you set the required add-in permissions and tenant permissions on the Azure application. For more information, see [Set additional permissions for SharePoint Online](#).

You can follow the Microsoft 365 source discovery and registration progress on the **Sources** page.

Next > You are now ready to protect your Microsoft 365 [Mailboxes](#), [OneDrives](#), [SharePoint Online Sites](#), and [Teams](#)!

Manual Registration for Microsoft 365 Sources

Note: Basic Auth is not supported for Microsoft 365 source registration.


To register your Microsoft 365 domain manually, make sure you've met all the [Microsoft 365 Requirements](#) and then:

1. In **DataProtect as a Service**, navigate to the **Sources** page and click **+ Register Source** in the upper-right corner of the page.
2. In the **Select Source** dialog box, select **Microsoft 365** and click **Start Registration**.


Select Source

SaaS Sources


On-Premise Sources




Microsoft 365




AWS




Virtual Machine




Physical



Microsoft SQL



NAS



Oracle

To register a Microsoft 365 source, you must have:

- The Microsoft Azure account with global administrator access to provision the Azure Apps for automatic Microsoft 365 source registration (recommended). To register the source manually, [register a custom Azure App](#). You can use a Microsoft 365 account which is a service account and not an administrative account.
- Microsoft Azure apps on the Azure portal with necessary permissions to perform object discovery, backup, and recovery in Microsoft 365.
- Additional [Microsoft Organization setting](#) for Mailbox size reporting.

[Learn more](#)

Cancel

Start Registration

3. In the **Source Details** section, select a cloud region for your data backups.
4. Choose the **Microsoft 365 Applications** to discover.

Note: Discovery selection change is not allowed for applications with protected objects.

Note: If the **Private Chats** and **Teams Posts** option is enabled under the **Mailbox** and **Teams** apps respectively, the Private Chats and Teams Posts will be backed up along with the corresponding Users and Teams respectively.

Private Chats and Teams Posts backup APIs are charged separately by Microsoft. All Azure apps configured by Cohesity must be linked to an [Azure subscription for billing](#). Backups of Mailboxes and Teams may fail if the Azure subscription configuration is not set appropriately.

This is an Early Access feature. Contact your Cohesity account team to enable the feature.

5. *[Optional]* You can enable the below options based on your requirement:

Note: Retaining the default values will speed up object discovery in the environment. The metadata required (which is fetched through these options) will be fetched during the backup of the objects and updated.

1. **Fetch Mailbox Info** to fetch and process the Mailbox information including the provisioning status, mailbox type, and in-place archival usage.

Note: You can enable this option to discover the Mailboxes of the users that were converted into Shared Mailboxes by revoking the user Exchange Online licenses.

2. **Fetch OneDrive Info** to fetch and process the OneDrive information including the provisioning status and storage quota.
3. **Include Users without MySite** to include users who have unprovisioned OneDrive or do not have MySite.
4. **Enable Site Tagging** to tag SharePoint Sites whether they are a Group Site or a Teams Site.

Note: Any Site that is tagged as a Group or Team will not be visible in the Sites section and these sites will be protected through the corresponding Group or Team protection.

6. In the **Account Credentials** section, enter the full **Username** of the Microsoft 365 user account with a valid SharePoint and Mailbox license.
7. In the **Azure Applications** section, enter the number of Azure applications that you want to create based on your requirements and click **Create**.

Register Microsoft 365 Source

☒ Fetch Mailbox Info ⓘ

☒ Fetch OneDrive Info ⓘ

☒ Include Users without MySite ⓘ

☒ Enable Site Tagging ⓘ

Account Credentials

Username
dmas-ul-automation4-v2@cohesity.com ✓
Enter full UPN (Eg: username@company.com)

Azure Applications

Number of Azure Apps to be created 1 **Create**

Cancel Register

Note: By default, an Azure application will be created. To better manage Microsoft 365 throttling, Cohesity recommends at least one Azure app.

- In the **Add Azure Application** form, click the **You can also add Azure App manually** link and then enter the **App ID** and **App Secret Key** that you noted down while registering your custom Azure app.

Tip: You can add multiple Azure apps for a Microsoft 365 source to load balance your backup and restore operations. Click **+** to add multiple Azure apps. When you do, ensure that you provide the valid **App ID** and **App Secret Key**.

- Click **Register**.

For SharePoint Online data protection, ensure that you set the required add-in permissions and tenant permissions on the Azure application. For more information, see [Set additional permissions for SharePoint Online](#).

You can follow the Microsoft 365 source discovery and registration progress on the **Sources** page.

Next > You are now ready to protect your Microsoft 365 [Mailboxes](#), [OneDrives](#), [SharePoint Online Sites](#), and [Teams](#)!

Manage Microsoft 365 Source Registration

After registering your Microsoft 365 domain as a source, you can:

- Update the Microsoft 365 source configuration.
- Refresh the source details.
- Unregister the Microsoft 365 domain from Cohesity.


Update the Microsoft 365 Source Configuration

After registering your Microsoft 365 source on Cohesity, you might have changed the Microsoft 365 domain configuration by:

- Changing the credentials of Microsoft 365 user account credentials.
- Updating the app secret by adding more permissions to the custom app.
- Creating a new app ID.

You can update the Microsoft 365 details provided during the registration process with the latest Microsoft 365 configuration details.

To edit the Microsoft 365 source configuration:.

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the actions menu () next to the Microsoft 365 source and select **Edit**.
3. In the **Register Microsoft 365 Source** page, update the required configurations.
4. Click **Register**.

To edit Azure App ID:

Note: Azure App ID permissions must be provided for successful Private Chats and Teams Posts backup.

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the actions menu next to the Microsoft 365 source and select **Edit**.
3. Click **Update App Permission**. This button is displayed if the **Private Chats** or **Teams Posts** discovery is enabled under the **Mailbox** or **Teams** apps respectively.
4. Copy the displayed device code to link with Microsoft Azure automatically. The code is valid for 15 minutes.
5. Open the Microsoft Azure App authorization service in a new tab and paste the copied code to complete authorization. When prompted to log in to Microsoft Azure, ensure

to use an account with global administrator access.

6. Complete the authorization and click **Update**.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

Refresh the Microsoft 365 Source

You can refresh the Microsoft 365 domain configuration and fetch the latest changes on the Microsoft 365 domain.

To refresh the Microsoft 365 Source Configuration:

1. Navigate to **Sources**.
2. Click the actions menu (⋮) next to the Microsoft 365 source and select **Refresh**.

Unregister the Microsoft 365 Domain

If you plan to stop taking the backup of your Microsoft 365 domain, you can unregister the Microsoft 365 source from the Cohesity DataProtect as a Service.

To unregister the Microsoft 365 domain:

1. Navigate to **Sources**.
2. Click the actions menu (⋮) next to the Microsoft 365 source and select **Unregister**.

Explore Microsoft 365 Sources

After you have registered your Microsoft 365 domain as a source, you can review the Users, Mailboxes, OneDrives, Sites, and Teams that Cohesity DataProtect as a Service discovered for the source.

Overview

To explore your Microsoft 365 source details, under Sources, find the Microsoft 365 source and click it.

The discovered Mailboxes, OneDrives, Sites, and Teams are listed in their respective tabs on the Microsoft 365 source details page. In addition, the source details page also displays a glance bar that communicates:

- **Object Counts.** The number of Users, Mailboxes, OneDrives, Sites and Teams discovered from the source.

- **Protected/Unprotected Objects.** The protected and unprotected count of Microsoft 365 objects in the source. For example, the number of protected and unprotected Mailboxes in the source.
- **Size.** The size (FETB) of protected and unprotected Microsoft 365 application data. For example, the amount of protected and unprotected Mailboxes data in the source.
- **Cross-App Counts.** Summary of protected and unprotected objects across all the Microsoft 365 applications in the source.

Interpret the Numbers

Every Microsoft 365 licensed user is counted as a User in Cohesity. A User might have both a Mailbox and a OneDrive. Or a User can have either a Mailbox or a OneDrive. In addition, Shared Mailboxes and Resource Mailboxes are not counted as Users. That means that the count of Mailboxes, OneDrives, and Users is not expected to be the same.

For example, in the source details page below, the right side of the glance bar lists **17** Users but the number of Mailboxes listed on the left side of the glance bar is **22**.

The screenshot displays the Cohesity DataProtect interface for a source named 'cohesitydemo3@microsoft.com'. The 'Mailbox' tab is active, showing a summary of protected and unprotected objects. The 'Unprotected' section indicates 22 Mailboxes and 0 Bytes of data. The 'Protected' section shows 0 Mailboxes and 0 Bytes of data. The 'Unprotected Summary' provides a breakdown: 58 Total, 17 Users, 22 Sites, and 19 Teams. The 'Protected Summary' shows 8 Total, 0 Users, 4 Sites, and 4 Teams. Below the summary, a table lists objects with columns for Object, Status, SLA, and Latest Snapshot. The first object is 'Users' and the second is 'Abhijeet'.

Object	Status	SLA	Latest Snapshot
> <input type="checkbox"/> <input type="checkbox"/> Users			
<input type="checkbox"/> <input type="checkbox"/> Abhijeet	⊘	-	-

Similarly, in the following details page, the right side of the glance bar lists **17** users but the number of OneDrives listed on the left side of the glance bar is **21**.

Search

DataProtect

← Sources

cohesitydataprotect.com

Mailbox **OneDrive** Site Teams

Unprotected		Protected		Unprotected Summary				Protected Summary			
21	12.9 GiB	0	0 Bytes	58	17	22	19	8	0	4	4
OneDrives	Data	OneDrives	Data	Total	Users	Sites	Teams	Total	Users	Sites	Teams

Protection Status

Object	Status	SLA	Latest Snapshot
Users			
Abhijeet		-	-

Next > You are now ready to protect your Microsoft 365 [Mailboxes](#), [OneDrives](#), [SharePoint Online Sites](#), and [Teams](#)!

Exchange Online Mailboxes

Microsoft Exchange Online is a SaaS application that is bundled in your Microsoft 365 subscription service. It is a hosted messaging solution that delivers the capabilities of Microsoft Exchange Server as a cloud-based service. It gives users access to email, calendar, contacts, and tasks from PCs, the web, and mobile devices. Using the policy-based data protection solution from Cohesity DataProtect as a Service, you can protect Exchange Online data on Microsoft 365.

Considerations

- PST download is supported only for mailbox items and not entire mailboxes.
- The exported PST of the mailbox items is valid for 72 hours. Ensure that you download the PST file within 72 hours of the recovery task completion.
- PST recovery of emails with more than 2000 recipients is not supported.
- Backup and download of the following is not supported:
 - Self-message (messages sent to self)
 - Saved or pinned messages property in the conversation
 - Meeting recordings metadata from private chats
- In Recoverable Items,

- only deletions, Purges, Discovery Holds, and SubstrateHolds folders are currently supported.

Note: SubstrateHolds folders support is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- Audits, Calendar Logging, and Versions folders are not supported.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- Granular recovery and search is not supported for Recoverable Items. The admin must recover the complete mailbox to recover the Recoverable Items.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- Converted-shared mailboxes are not backed up by default. To enable backup of converted-shared mailboxes, contact [Cohesity Support](#).
- The [Retirement of RBAC Application Impersonation in Exchange Online](#) does not impact the Microsoft 365 Exchange Online Mailboxes, Teams, and Groups protection workflow on the Cohesity cluster.
- Exchange Online Mailboxes protection is supported only for users present on Azure AD and not for users on the on-premises Active Directory.
- Exchange Online Mailboxes protection is supported for:
 - users present in Azure AD
 - hybrid setup with users present on Azure AD

It is not supported for:

- users on the on-premises Active Directory
- users natively on the on-premises Active Directory that are synchronized to Azure AD.
- By default, the Conversation History/Team Chat folder is excluded during Exchange Online backup.

- Following the Microsoft changes for EWS APIs, Cohesity no longer backs up mailbox items of the type **IPM.Teams.SkypeMessage**. For more details, see the [KB Article](#).
- [Microsoft 365 Backup Storage](#) service is supported for Mailboxes.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

Protect Microsoft 365 Mailboxes

Once you have [registered your Microsoft 365 domain](#) as a source, you're ready to protect the user Mailboxes in your domain.

To protect your Microsoft 365 Mailboxes:

Note: If the **Private Chats** option is enabled under Mailbox during app registration, the Private Chats will be backed up along with the corresponding Users.

Mailboxes backup may fail if the Azure subscription configuration is not set appropriately.

This is an Early Access feature. Contact your Cohesity account team to enable the feature.

1. In **DataProtect as a Service**, under **Sources**, find the Microsoft 365 source and click on it.
2. Click the **Mailbox** tab.
3. Select the individual Mailboxes you wish to protect or:
 - Click **Users > Select All Child Objects** to protect all the Mailboxes in this source.
 - Click **Users > Auto Protect This** to protect all the Mailboxes *plus any future additional Mailboxes* on that source.
 - Click the **Security Groups** icon and select the security group to protect the Mailboxes of the users in the security group. For more information, see [Security Groups](#).

Security groups based user import also supports nested security groups.

For example, if a security group (AA) includes members (X and Y) and security groups (BB and CC), the import process includes the members (X and Y) and the members of security groups (BB and CC).

Cohesity supports only users and not devices in Security Groups. This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

1. Click the **Protect** icon above the list.
2. Choose a policy to specify backup frequency and retention. If you don't have a policy, you can easily [create one](#).
3. Under **Settings**, edit the **Start Time** if necessary.
4. Under **Additional Settings**, you can enable **Indexing** and other [additional settings](#).

Note: Indexing is enabled by default. If you plan to [recover individual emails or folders](#), in addition to whole Mailboxes, you need to enable **Indexing** in this step. When you do, you can include or exclude specific Mailboxes from indexing.

5. Click **Protect**.

Note: The backups start immediately after you protect the objects, regardless of the time you set for the protection run.

Next > When the first protection run completes, you will be ready to [recover your protected Mailboxes](#) when and if you need to.

Additional Settings

Advance Settings	Description
Start Time	Available only if the selected policy is set to Backup Daily . Indicates what time the protection run should start. Enter the Start Time and select AM or PM . The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone here .
SLA	<p>The service-level agreement (SLA) defines how long the administrator expects a protection run to take. Enter:</p> <ul style="list-style-type: none"> • Full. The number of minutes you expect a full protection run, which captures all the blocks in an object, to take. • Incremental. The number of minutes you expect an incremental protection run, which captures only the changed blocks in an object, to take.

Advance Settings	Description
Cancel Runs at Quiet Time Start	<p><i>(Available only if the selected policy has at least one Quiet Time)</i></p> <p>When enabled, all the protection runs that are currently executing will cancel when the Quiet Time period starts. By default, this setting is disabled, meaning that after a protection run starts, it continues to execute even when a Quiet Time period starts. However, new protection runs will not start during a Quiet Time.</p>
Indexing	<p>By default, indexing is enabled.</p> <div> <p>Note: Indexing is mandatory for granular restore of an Exchange Online mailbox , such as restoring a folder or restoring an email.</p> </div>
Exclusions	<p>Select the folders that you plan to exclude from the backup or click Add to add custom folders that you want to exclude from the backup.</p> <p>To protect the Recoverable Items, ensure to deselect Archive Recoverable Items and Recoverable Items.</p> <div> <p>Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.</p> </div>

Manage Existing Protection

Edit protection settings, change the policy, and start, stop, & pause protection.


Once you have [applied protection](#) to the objects in your sources, Cohesity DataProtect as a Service makes it easy to make changes to that protection quickly. You can:

- Edit additional settings like **End Date**, **Exclusions**, **Alerts**, and more.
- Apply a different policy.
- Start an on-demand protection run, pause and resume it, or even remove protection.

Edit Protection Settings

To edit protection settings:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click into the **Source** name.
3. Select **Show All > Protected** and use the other filters, search box, and views at the top to narrow your search.

4. Click the **Actions** menu () next to the object and select **Edit Protection** to open the protection settings for that object.

Apply a New Protection Policy

To change the **Policy**, click the drop-down and select a different policy. To help you choose, each policy in the list shows the **Backup** frequency and the **Retain** period for each backup.

If you don't have a policy that meets your needs, scroll to the bottom of the list and click **Create Policy** to [create your own policy](#).

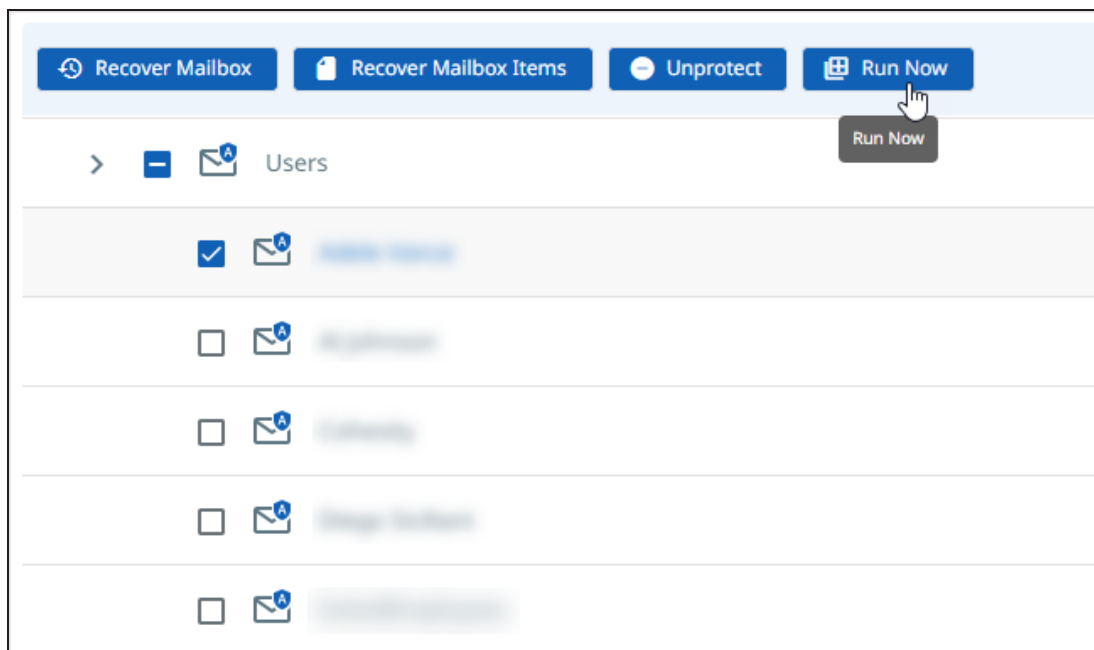
Edit Additional Protection Settings

In **DataProtect as a Service**, under **Settings**, you can change the protection **Start Time** (and select the **Time Zone**).

Click the drop-down next to **Additional Settings** to change more options. See [Additional Protection Settings](#) for details.

Start, Stop, or Remove Protection

When you select protected objects in one of your sources, Cohesity DataProtect as a Service presents buttons for the actions that are possible for those objects.



With the protected objects selected, you can click:

- **Recover Mailbox** to recover the mailbox.
- **Recover Mailbox Items** to recover the mailbox items.
- **Run Now** to start an on-demand protection run immediately.
- **Unprotect** to remove protection from the object.

Tip: If a protected object is deleted from the source, you can search the object using Global Search and unprotect it.

Additional Settings

Advance Settings	Description
Start Time	Available only if the selected policy is set to Backup Daily . Indicates what time the protection run should start. Enter the Start Time and select AM or PM . The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone here.
SLA	<p>The service-level agreement (SLA) defines how long the administrator expects a protection run to take. Enter:</p> <ul style="list-style-type: none"> • Full. The number of minutes you expect a full protection run, which captures all the blocks in an object, to take. • Incremental. The number of minutes you expect an incremental protection run, which captures only the changed blocks in an object, to take.
Cancel Runs at Quiet Time Start	<p><i>(Available only if the selected policy has at least one Quiet Time)</i></p> <p>When enabled, all the protection runs that are currently executing will cancel when the Quiet Time period starts. By default, this setting is disabled, meaning that after a protection run starts, it continues to execute even when a Quiet Time period starts. However, new protection runs will not start during a Quiet Time.</p>
Indexing	<p>By default, indexing is enabled.</p> <p>Note: Indexing is mandatory for granular restore of an Exchange Online mailbox , such as restoring a folder or restoring an email.</p>
Exclusions	<p>Select the folders that you plan to exclude from the backup or click Add to add custom folders that you want to exclude from the backup.</p> <p>To protect the Recoverable Items, ensure to deselect Archive Recoverable Items and Recoverable Items.</p> <p>Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.</p>

Recover Microsoft 365 Mailboxes and Mailbox Items

After you [protect your users' Microsoft 365 Mailboxes](#), you can recover them — as [whole Mailboxes](#) or [individual mailbox items](#) — from Cohesity DataProtect as a Service.

Note: You can recover Mailboxes to a target Mailbox as long as the Microsoft 365 domain for the target Mailbox is registered within the same [cloud region](#) as the Microsoft 365 domain of the Mailbox being recovered.

You can recover:

- [Whole Mailboxes](#)
- [Individual Mailbox Items](#)

Recover User Mailboxes

To recover protected Microsoft 365 user Mailboxes:

1. In **DataProtect as a Service**, go to **Sources** to set up your recovery task.
2. Click on the **Source** name and select the **Mailbox** tab.
3. Above the tree, select **Show All > Protected**.
4. Find the Mailbox you need and click the **Recover** icon on that row to open the **New Recovery** form with the Latest snapshot (protection run).
5. In the **New Recovery** form, if you need to add more Mailboxes and/or recover from an earlier backup, click the **Edit** icon in the top right of the form.
 - To add **Mailboxes**, enter a **Search** term on the left, locate the other Mailboxes, and select them.
 - To use a different **Recovery Point** for a Mailbox, click the **Edit** icon on the tile for that Mailbox. Find the recovery point you need and click **Select Recovery Point**.

Click **Next: Recover Options** to return to the form.

6. Under **Recover To**, select **Original Location** or **New Location**.
If you choose **New Location**, select a **Registered Source** and the **Target Mailbox**.
7. Select your **Recovery Options**:
 - **Continue on Error**. Enable to recover even if errors occur when recovering Mailboxes. For example, if one of the Mailboxes cannot be recovered, Cohesity will still attempt to recover the other selected Mailboxes.
 - **Task Name**. Change the default name of the recovery task.
 - **Include Recoverable Items**. Toggle ON to **Recover Recoverable Items**.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- **Include Archive Recoverable Items.** Toggle ON to **Recover Archive Recoverable Items**.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

8. Click **Start Recovery**.

Next > Protect your Microsoft 365 [OneDrives](#), [SharePoint Online Sites](#), and [Teams](#) so you can recover them easily when you need to, as well!

Recover Mailbox Items

After you [protect your users' Microsoft 365 Mailboxes](#), you can recover the Mailbox items such as [individual emails](#), [folders](#), [calendar invites](#), [contacts](#), [notes](#), or [tasks](#) — from Cohesity DataProtect as a Service.

Recover Emails

To recover specific emails from a protected Microsoft 365 user Mailbox:

1. In **DataProtect as a Service**, go to **Sources** to set up your recovery task.
2. Click into the **Source** name and select the **Mailbox** tab.
3. Use the search or filter options, find and select the user you need, and click the **Recover Mailbox Items** icon above the list to open the **New Recovery** form.

Tip: You can also use Global Search to locate, filter, and select the Mailbox you need. Click the Global Search box at the top or type slash (/) anywhere to start your search.

4. Select **Emails & Folders** from the **Item Type** drop-down.
5. Use the '*' wildcard character or enter the text to search for emails with a matching subject in the **Search** bar. Select the emails to recover from the search results.

Or

Click **Advanced Search** and select **Emails** to search based on these filters:

Filters	Description
Subject	Subject line in the email.
From	Mail sender email address.
To	Mail recipient email address. Use a comma or space separator to enter multiple addresses.
Date Range	Using the calendar, select a specific date range to search the emails.
Email Type	Select one of the email types: <ul style="list-style-type: none"> • All Emails • Only emails with attachments • Only emails without attachments
cc	The email address in the Cc: line of the email. Use a comma or space separator to enter multiple addresses.
bcc	The email address in the Bcc: line of the email. Use a comma or space separator to enter multiple addresses.
Search in Folder	Search for the email within the specified folder. For example, Inbox, Drafts, and so on. Use a comma or space separator to enter multiple folder names.

To use a different **Recovery Point** for a Mailbox, click the **Edit** icon on the tile for that Mailbox. Find the recovery point you need and click **Select Recovery Point**.

- Click **Next: Recover Options** to return to the form.
- Under **Recover To**, select **Original Location** or **New Location**.

If you choose **New Location**, select a **Registered Source** and the **Target Mailbox**, and specify the **Folder** name to which you plan to recover.

Note: If a folder with the specified name does not exist, Cohesity creates the folder and recovers the emails to that folder.

Select **Export as PST** to export the backed up emails as a PST file, and provide a password for the exported PST file.

Note: Once the recovery task is completed, the exported PST file is valid for 72 hours. Ensure that you download the PST within 72 hours. For more information, see [Download Exported PST File](#).

This is an Early Access feature. Contact your Cohesity account team to enable the feature.

Note: PST recovery of emails with more than 2000 recipients is not supported.

8. Select your **Recovery Options**:

- **Continue on Error.** Enable to recover even if errors occur when recovering Mailboxes. For example, if one of the emails cannot be recovered, Cohesity will still attempt to recover the other selected emails.
- **Task Name.** Change the default name of the recovery task.

9. Click **Start Recovery**.

Recover Mailbox Folders

To recover specific folders from a protected Microsoft 365 user Mailbox:

1. Go to **Sources** to set up your recovery task.
2. Click into the **Source** name and select the **Mailbox** tab.
3. Use the search or filter options, find and select the user you need, and click the **Recover Mailbox Items** icon above the list to open the **New Recovery** form.

Tip: You can also use Global Search to locate, filter, and select the Mailbox you need. Click the Global Search box at the top or type slash (/) anywhere to start your search.

4. On the **New Recovery** page, select **Emails & Folders** from the **Item Type** drop-down
5. Click **Advanced Search** and select **Folders**.
6. Enter the **Folder Name** and click **Apply**. Select the folders to recover from the search results.

To use a different **Recovery Point** for a Mailbox, click the **Edit** icon on the tile for that Mailbox. Find the recovery point you need and click **Select Recovery Point**.

7. Click **Next: Recover Options** to return to the form.

- Under **Recover To**, select **Original Location** or **New Location**.

If you choose **New Location**, select a **Registered Source** and the **Target Mailbox**, and specify the **Folder** name to which you plan to recover.

Note: If a folder with the specified name does not exist, Cohesity creates the folder and recovers the data to it.

Select **Export as PST** to export the backed up mailbox folders as a PST file, and provide a password for the exported PST file.

Note: Once the recovery task is completed, the exported PST file is valid for 72 hours. Ensure that you download the PST within 72 hours. For more information, see [Download Exported PST File](#).

This is an Early Access feature. Contact your Cohesity account team to enable the feature.

- Select your **Recovery Options**:

- **Continue on Error.** Enable to recover even if errors occur when recovering Mailboxes. For example, if one of the emails cannot be recovered, Cohesity will still attempt to recover the other selected emails.
- **Task Name.** Change the default name of the recovery task.

- Click **Start Recovery**.

Recover Calendar Invites

You can recover specific calendar invites from a protected Microsoft 365 user Mailbox. However, if you plan to recover the entire calendar, then [recover the mailbox folder](#) called **Calendar**.

To recover calendar Invites:

- Go to **Sources** to set up your recovery task.
- Click into the **Source** name and select the **Mailbox** tab.
- Use the search or filter options, find and select the user you need, and click the **Recover Mailbox Items** icon above the list to open the **New Recovery** form.

Tip: You can also use Global Search to locate, filter, and select the Mailbox you need. Click the Global Search box at the top or type slash (/) anywhere to start your search.

- On the **New Recovery** page, select **Calendars** from the **Item Type** drop-down.

5. Use the '*' wildcard character or enter the text to search for calendar items with a matching subject of the calendar invite in the **Search** bar. Select the calendar invite to recover from the search results.

Or

Click **Advanced Search** and search calendar invite based on these filters and click **Apply**:

Filters	Description
Subject of Event	Subject line in the calendar invite.
Organizer	The email address of the event organizer.
Invitee	Event recipients' email addresses. Use a comma or space separator to enter multiple addresses.
Invitation Date	Using the calendar, select a specific date range to search the calendar invites.

To use a different **Recovery Point** for a Mailbox, click the **Edit** icon on the tile for that Mailbox. Find the recovery point you need and click **Select Recovery Point**.

6. Click **Next: Recover Options** to return to the form.
7. Under **Recover To**, select **Original Location** or **New Location**.

If you choose **New Location**, select a **Registered Source** and the **Target Mailbox**, and specify the **Folder** name to which you plan to recover.

Note: If a folder with the specified name does not exist, Cohesity creates the folder and recovers the calendar invite(s) to that folder.

Select **Export as PST** to export the backed up calendar invites as a PST file, and provide a password for the exported PST file.

Note: Once the recovery task is completed, the exported PST file is valid for 72 hours. Ensure that you download the PST within 72 hours. For more information, see [Download Exported PST File](#).

This is an Early Access feature. Contact your Cohesity account team to enable the feature.

8. Select your **Recovery Options**:

- **Continue on Error.** Enable to recover even if errors occur when recovering calendar invites. For example, if one of the calendar invites cannot be recovered, Cohesity will still attempt to recover the other selected calendar invite.
- **Task Name.** Change the default name of the recovery task.

9. Click **Start Recovery**.

Recover Contacts

You can recover specific contacts from a protected Microsoft 365 user Mailbox. However, if you plan to recover the complete contacts, then [recover the mailbox folder](#) called **Contacts**.

To recover specific contacts:

1. Go to **Sources** to set up your recovery task.
2. Click into the **Source** name and select the **Mailbox** tab.
3. Use the search or filter options, find and select the user you need, and click the **Recover Mailbox Items** icon above the list to open the **New Recovery** form.

Tip: You can also use Global Search to locate, filter, and select the Mailbox you need. Click the Global Search box at the top or type slash (/) anywhere to start your search.

4. On the **New Recovery** page, select **Contacts** from the **Item Type** drop-down.
5. Use the '*' wildcard character or enter the text to search for contacts with a matching contact name in the **Search** bar. Select the contact to recover from the search results.

Or

Click **Advanced Search** and search the contact based on these filters and click **Apply**:

Filters	Description
First Name	The first name of the contact.
Last Name	The last name of the contact.
Email Address	The email address of the contact.
Invitation Date	Using the calendar, select a specific date range to search the calendar invites.

To use a different **Recovery Point** for a Mailbox, click the **Edit** icon on the tile for that Mailbox. Find the recovery point you need and click **Select Recovery Point**.

6. Click **Next: Recover Options** to return to the form.
7. Under **Recover To**, select **Original Location** or **New Location**.

If you choose **New Location**, select a **Registered Source** and the **Target Mailbox**, and specify the **Folder** name to which you plan to recover.

Note: If a folder with the specified name does not exist, Cohesity creates the folder and recovers the contact(s) to that folder.

Select **Export as PST** to export the backed up contacts as a PST file, and provide a password for the exported PST file.

Note: Once the recovery task is completed, the exported PST file is valid for 72 hours. Ensure that you download the PST within 72 hours. For more information, see [Download Exported PST File](#).

This is an Early Access feature. Contact your Cohesity account team to enable the feature.

8. Select your **Recovery Options**:
 - **Continue on Error**. Enable to recover even if errors occur when recovering the contacts. For example, if one of the contacts cannot be recovered, Cohesity will still attempt to recover the other selected contacts.
 - **Task Name**. Change the default name of the recovery task.
9. Click **Start Recovery**.

Recover Notes

You can recover specific notes from a protected Microsoft 365 user Mailbox. However, if you plan to recover the complete set of notes in the user Mailbox, then [recover the mailbox folder](#) called **Notes**.

To recover specific notes:

1. Go to **Sources** to set up your recovery task.
2. Click into the **Source** name and select the **Mailbox** tab.
3. Use the search or filter options, find and select the user you need, and click the **Recover Mailbox Items** icon above the list to open the **New Recovery** form.

Tip: You can also use Global Search to locate, filter, and select the Mailbox you need. Click the Global Search box at the top or type slash (/) anywhere to start your search.

4. On the **New Recovery** page, select **Notes** from the **Item Type** drop-down.
5. Use the '*' wildcard character or enter the text to search for notes with a matching subject of the note in the **Search** bar. Select the note(s) to recover from the search results.

Or

Click **Advanced Search** and search the note based on these filters and click **Apply**:

Filters	Description
Subject	The subject of the note.
Creation Date	Using the calendar, select a specific date range to search the notes based on the creation date.
Modification Date	Using the calendar, select a specific date range to search the notes based on the modification date.

To use a different **Recovery Point** for a Mailbox, click the **Edit** icon on the tile for that Mailbox. Find the recovery point you need and click **Select Recovery Point**.

6. Click **Next: Recover Options** to return to the form.
7. Under **Recover To**, select **Original Location** or **New Location**.

If you choose **New Location**, select a **Registered Source** and the **Target Mailbox**, and specify the **Folder** name to which you plan to recover.

Note: If a folder with the specified name does not exist, Cohesity creates the folder and recovers the note(s) to that folder.

Select **Export as PST** to export the backed up notes as a PST file, and provide a password for the exported PST file.

Note: Once the recovery task is completed, the exported PST file is valid for 72 hours. Ensure that you download the PST within 72 hours. For more information, see [Download Exported PST File](#).

This is an Early Access feature. Contact your Cohesity account team to enable the feature.

8. Select your **Recovery Options**:

- **Continue on Error.** Enable to recover even if errors occur when recovering the notes. For example, if one of the notes cannot be recovered, Cohesity will still attempt to recover the other selected note.
- **Task Name.** Change the default name of the recovery task.

9. Click **Start Recovery**.

Recover Tasks

You can recover specific tasks from a protected Microsoft 365 user Mailbox. However, if you plan to recover the complete set of tasks in the user Mailbox, then [recover the mailbox folder](#) called **Tasks**.

To recover specific notes:

1. Go to **Sources** to set up your recovery task.
2. Click into the **Source** name and select the **Mailbox** tab.
3. Use the search or filter options, find and select the user you need, and click the **Recover Mailbox Items** icon above the list to open the **New Recovery** form.

Tip: You can also use Global Search to locate, filter, and select the Mailbox you need. Click the Global Search box at the top or type slash (/) anywhere to start your search.

4. On the **New Recovery** page, select **Tasks** from the **Item Type** drop-down.
5. Use the '*' wildcard character or enter the text to search for notes with a matching subject of the task in the Search bar. Select the task(s) to recover from the search results.

Or

Click **Advanced Search** and search the tasks based on these filters and click **Apply**:

Filters	Description
Subject	The subject of the task.
Creation Date	Using the calendar, select a specific date range to search the tasks based on their creation date.
Due Date	Using the calendar, select a specific date range to search the tasks based on their due date.
Status	The status of the task.

To use a different **Recovery Point** for a Mailbox, click the **Edit** icon on the tile for that Mailbox. Find the recovery point you need and click **Select Recovery Point**.

6. Click **Next: Recover Options** to return to the form.
7. Under **Recover To**, select **Original Location** or **New Location**.

If you choose **New Location**, select a **Registered Source** and the **Target Mailbox**, and specify the **Folder** name to which you plan to recover.

Note: If a folder with the specified name does not exist, Cohesity creates the folder and recovers the task(s) to that folder.

Select **Export as PST** to export the backed up tasks as a PST file, and provide a password for the exported PST file.

Note: Once the recovery task is completed, the exported PST file is valid for 72 hours. Ensure that you download the PST within 72 hours. For more information, see [Download Exported PST File](#).

This is an Early Access feature. Contact your Cohesity account team to enable the feature.

8. Select your **Recovery Options**:
 - **Continue on Error**. Enable to recover even if errors occur when recovering the tasks. For example, if one of the tasks cannot be recovered, Cohesity will still attempt to recover the other selected tasks.
 - **Task Name**. Change the default name of the recovery task.
9. Click **Start Recovery**.

Download Exported PST File

After the recovery task is completed, within 72 hours you can download the exported PST file of the mailbox items that you choose to recover.

To download the PST file:

1. Navigate to **Activity**.
2. Locate and click on the recovery task from which you want to download the exported PST file.
3. Click **Download Files**.

The PST file is downloaded to your local system.

Note: The PST file is protected with a password; you must contact the admin user and obtain the password to open the downloaded PST file.
This is an Early Access feature. Contact your Cohesity account team to enable the feature.

Download Private Chats

To download the Private Chats of a user:

1. Navigate to **Sources**.
2. Click on the required source and click the **Mailbox** tab.
3. Select **Show All > Protected** and use the other filters, search box, and views at the top to narrow your search.
4. Click the **Actions** menu next to the object and click **Download Private Chats**.
5. In the **Download Private Chats** page, select the required snapshot, provide the task name, and click **Recover**.

Note: Attachments in the Private Chats will not be downloaded.

6. Click the **View Progress** button in the pop-up message or click the **Activity** menu.
7. Once the recovery is successful, click **Download**. The Private Chats will be downloaded.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

Next > Protect your Microsoft 365 [OneDrives](#), [SharePoint Online Sites](#), and [Teams](#) so you can recover them easily when you need to, as well!

Mailbox Items Recovery Self-Service

Cohesity provides a self-service workflow to help the end users recover the Microsoft 365 Mailbox items by leveraging the Microsoft Entra ID (Azure Active Directory) login for user authentication.

Administrators can authorize the self-service workflow for users through the Security Groups.

Note: You can access the self-service portal through <https://helios.cohesity.com/#/self-service-portal/auth>.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

Mailbox Items Recovery

To recover Microsoft 365 Mailbox items from the Cohesity Self-Service Portal:

1. In the Cohesity Data Cloud login page, log in through Microsoft using the **Cohesity Self-Service Portal** link.
2. Choose **Microsoft 365 Mailbox** to recover your Emails and folders, Calendars, Contacts, Tasks, and Notes.
3. In the **Recover** page, select the required **Item Type** in the drop-down. The options include:
 - Emails & Folders
 - Calendars
 - Contacts
 - Tasks
 - Notes
4. Use the '*' wildcard character or enter the text to search for Mailbox items with a matching subject in the **Search** bar. Select the items to recover from the search results.

Or

Click **Advanced Search** and select the items to search based on the filters.

5. To use a different **Recovery Point** for a Mailbox item, click the **Edit** icon on the tile for that item. Find the recovery point you need and click **Select Recovery Point**.
6. Click **Recover**.
7. Under **Recover Type**, select the following:
 - **Recover to Original Location** to recover all the items directly to your Mailbox.
 - **Export as PST** to export all the items in the PST format. Provide a password for the exported PST file.

Note: You can download the PST file and use an agent to migrate the PST content.

Once the recovery task is completed, the exported PST file is valid for 72 hours. Ensure that you download the PST within 72 hours. For more information, see [Download Exported PST File](#).

8. Click **Finish**. You can view the recovery progress from the **Welcome** page under the **Recoveries** section or on the **Activity** page.
9. Click the action icon on the required task and click **Show Recovered Items** to view the name and size of the recovered items.

Download Exported PST File

After the recovery task is completed, within 72 hours you can download the exported PST file of the mailbox items that you choose to recover.

To download the PST file:

1. Navigate to **Activity**.
2. Locate and click on the recovery task from which you want to download the exported PST file.
3. Click **Download Files**. The PST file is downloaded to your local system.

Note: The PST file is protected with a password; you must contact the admin user and obtain the password to open the downloaded PST file.

OneDrive for Business

OneDrive for Business is a SaaS application that is bundled in your Microsoft 365 subscription service. It is an intelligent files app for Microsoft 365 connecting you to all your files so you can share and work together from anywhere while protecting your work. It enables you to easily store, access, and discover your individual and shared work files in

Microsoft 365. Using the policy-based data protection solution from Cohesity DataProtect as a Service, you can protect OneDrive for Business data on Microsoft 365.

Considerations

Review and understand the following considerations before you protect your Microsoft 365 OneDrive data:

- From the recovery workflow, you cannot download an empty folder.
- Backup and restore of OneNote files in OneDrive are not supported.
- Restoring shared permissions for files in the Preservation Hold Library (PHL) drive is not supported.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- PHL data can only be recovered using full OneDrive recovery. Granular level recovery is not supported for the PHL data.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- Data from PHL is not searchable.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- OneDrive protection is supported only for users present on Azure AD or Hybrid setup with users present on Azure AD and not for users on the on-premises Active Directory.
- File/folder permissions with no role will not be recovered since Microsoft does not allow adding such permissions via Graph APIs. For example, no role will be assigned when a Restricted View permission is added to a user from Advanced Setting.
- [Microsoft 365 Backup Storage](#) service is supported for OneDrive.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- OneDrive does not support item permissions from applications and devices.

Protect Microsoft 365 OneDrives

Once you have [registered your Microsoft 365 domain](#) as a source, you're ready to use Cohesity DataProtect as a Service to protect the user OneDrives in your domain.

To protect your Microsoft 365 OneDrives:

1. In **DataProtect as a Service**, under **Sources**, find the Microsoft 365 source and click on it.
2. Click the **OneDrive** tab.
3. Select the individual OneDrives you wish to protect or:
 - Click **Users > Select All Child Objects** to protect all the OneDrives in this source.
 - Click **Users > Auto Protect This** to protect all the OneDrives *plus any future additional OneDrives* on that source.
 - Click the **Security Groups** icon and select the security group to protect the OneDrives of the users in the security group. For more information, see [Security Groups](#).

Security groups based user import also supports nested security groups.

For example, if a security group (AA) includes members (X and Y) and security groups (BB and CC), the import process includes the members (X and Y) and the members of security groups (BB and CC).

Cohesity supports only users and not devices in Security Groups.

This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

4. Click the **Protect** icon above the list.
5. Choose a policy to specify backup frequency and retention. If you don't have a policy, you can easily [create one](#).
6. Under **Settings**, edit the **Start Time** if necessary.
7. Under **Additional Settings**, you can enable **Indexing** and other [additional settings](#).

Note: Indexing is enabled by default.

8. Click **Protect**.

Note: The backups start immediately after you protect the objects, regardless of the time you set for the protection run.

Next > When the first protection run completes, you will be ready to [recover your protected OneDrives](#) when and if you need to.

Additional Settings

Advance Settings	Description
Start Time	Available only if the selected policy is set to Backup Daily . Indicates what time the protection run should start. Enter the Start Time and select AM or PM . The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone here.
SLA	The service-level agreement (SLA) defines how long the administrator expects a protection run to take. Enter: <ul style="list-style-type: none"> • Full. The number of minutes you expect a full protection run, which captures all the blocks in an object, to take. • Incremental. The number of minutes you expect an incremental protection run, which captures only the changed blocks in an object, to take.
Cancel Runs at Quiet Time Start	<i>(Available only if the selected policy has at least one Quiet Time)</i> When enabled, all the protection runs that are currently executing will cancel when the Quiet Time period starts. By default, this setting is disabled, meaning that after a protection run starts, it continues to execute even when a Quiet Time period starts. However, new protection runs will not start during a Quiet Time.
Indexing	By default, indexing is enabled. Note: Indexing is mandatory to search for files or folders in a OneDrive.
Exclusions	Click Add to add custom folders that you want to exclude from the backup.

Manage Existing Protection


Edit protection settings, change the policy, and start, stop, & pause protection.

Once you have [applied protection](#) to the objects in your sources, Cohesity DataProtect as a Service makes it easy to make changes to that protection quickly. You can:

- Edit additional settings like **End Date**, **Exclusions**, **Alerts**, and more.
- Apply a different policy.
- Start an on-demand protection run, pause and resume it, or even remove protection.

Edit Protection Settings

To edit protection settings:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click into the **Source** name.
3. Select **Show All > Protected** and use the other filters, search box, and views at the top to narrow your search.
4. Click the **Actions** menu () next to the object and select **Edit Protection** to open the protection settings for that object.

Apply a New Protection Policy

To change the **Policy**, click the drop-down and select a different policy. To help you choose, each policy in the list shows the **Backup** frequency and the **Retain** period for each backup.

If you don't have a policy that meets your needs, scroll to the bottom of the list and click **Create Policy** to [create your own policy](#).

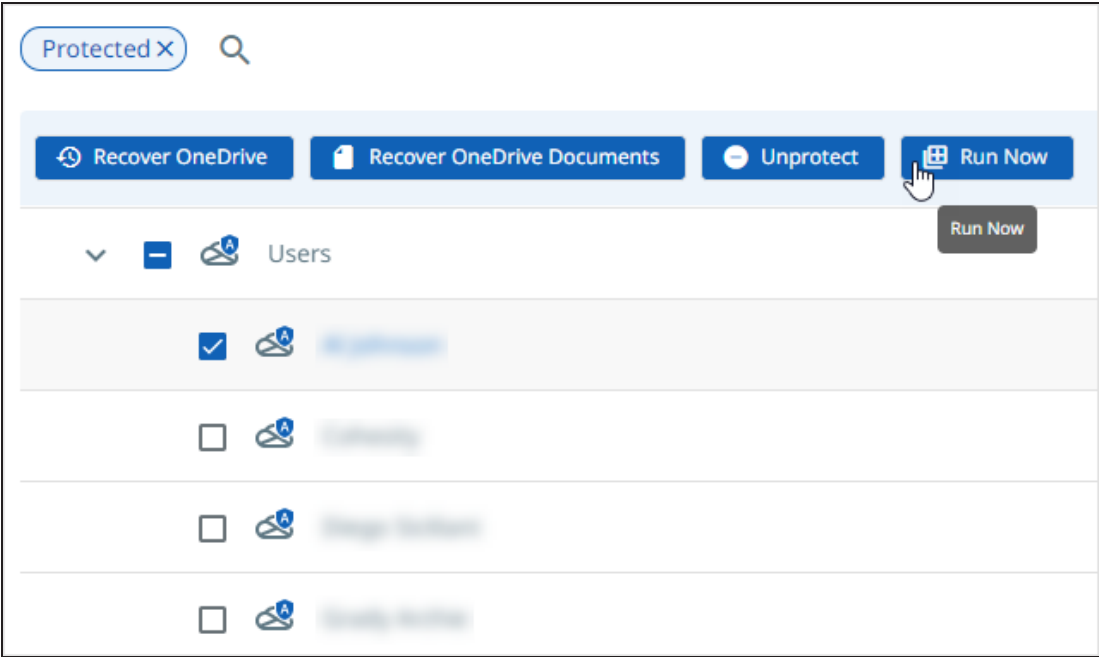
Edit Additional Protection Settings

In **DataProtect as a Service**, under **Settings**, you can change the protection **Start Time** (and select the **Time Zone**).

Click the drop-down next to **Additional Settings** to change more options. See [Additional Protection Settings](#) for details.

Start, Stop, or Remove Protection

When you select protected objects in one of your sources, Cohesity DataProtect as a Service presents buttons for the actions that are possible for those objects.



With the protected objects selected, you can click:

- **Recover OneDrive** to recover the OneDrive.
- **Recover OneDrive Documents** to recover the OneDrive documents.
- **Run Now** to start an on-demand protection run immediately.
- **Unprotect** to remove protection from the object.

Tip: If a protected object is deleted from the source, you can search the object using Global Search and unprotect it.

Additional Settings

Advance Settings	Description
Start Time	Available only if the selected policy is set to Backup Daily . Indicates what time the protection run should start. Enter the Start Time and select AM or PM . The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone here.

Advance Settings	Description
SLA	<p>The service-level agreement (SLA) defines how long the administrator expects a protection run to take. Enter:</p> <ul style="list-style-type: none"> • Full. The number of minutes you expect a full protection run, which captures all the blocks in an object, to take. • Incremental. The number of minutes you expect an incremental protection run, which captures only the changed blocks in an object, to take.
Cancel Runs at Quiet Time Start	<p><i>(Available only if the selected policy has at least one Quiet Time)</i></p> <p>When enabled, all the protection runs that are currently executing will cancel when the Quiet Time period starts. By default, this setting is disabled, meaning that after a protection run starts, it continues to execute even when a Quiet Time period starts. However, new protection runs will not start during a Quiet Time.</p>
Indexing	<p>By default, indexing is enabled.</p> <p>Note: Indexing is mandatory to search for files or folders in a OneDrive.</p>
Exclusions	<p>Click Add to add custom folders that you want to exclude from the backup.</p>

Recover OneDrives

After you [protect your users' Microsoft 365 OneDrives](#), you can recover them — as [whole OneDrives](#) or just [specific contents in a user's Microsoft 365 OneDrive](#) — from Cohesity DataProtect as a Service, to the same location, alternate location, or across Microsoft 365 domains.

Note: You can recover a OneDrive to a target OneDrive as long as the Microsoft 365 domain for the target OneDrive is registered within the same [cloud region](#) as the Microsoft 365 domain of the OneDrive being recovered. the same cloud region.

You can recover:

- [User OneDrives](#)
- [User OneDrive Contents](#)

Recover User OneDrives

To recover protected Microsoft 365 user OneDrives:

1. In **DataProtect as a Service**, go to **Sources** to set up your recovery task.
2. Click into the **Source** name and select the **OneDrive** tab.
3. Above the tree, select **Protection Status > Protected**.
4. Use the search and filter options to find and select the OneDrive you need, click the **Actions (:)** menu on that row, and select **Recover OneDrive** to open the **New Recovery** form with the **Latest** snapshot (protection run).
5. In the **New Recovery** form, if you need to add more OneDrives and/or recover from an earlier backup, click the **Edit** icon in the top right of the form.
 - To add OneDrives, enter a **Search** term on the left, locate the other OneDrives, and select them.
 - To use a different **Recovery Point** for a OneDrive, click the **Edit** icon on the tile for that OneDrive. Find the recovery point you need and click **Select Recovery Point**.

Click **Next: Recover Options** to return to the form.

6. Under **Recover To**, select **Original Location** or **New Location**.

If you choose **New Location**, select a **Registered Source** and the **Target OneDrive**.

7. Select your **Recovery Options**:
 - **Continue on Error**. Enable to recover even if errors occur when recovering OneDrives. For example, if one of the OneDrives cannot be recovered, Cohesity will still attempt to recover the other selected OneDrives.
 - **Include Preservation Hold Library**. Enable to recover the Preservation Hold Library that is part of the Cohesity snapshot. Recovering the Preservation Hold Library data may increase the recovery time substantially as it can include a large amount of data. The recovery will fail if the Target Location does not have sufficient space.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- **Task Name**. Change the default name of the recovery task.

8. Click **Start Recovery**.

Next > Protect your Microsoft 365 [Mailboxes](#), [SharePoint Online Sites](#), and [Teams](#) so you can recover them easily when you need to, as well!

Recover OneDrive Contents

Important: Before you can recover a user's OneDrive contents, you need to set up [Microsoft 365 OneDrive protection](#) with **Indexing** enabled.

To recover specific OneDrive contents from a protected Microsoft 365 OneDrive:

1. In **DataProtect as a Service**, go to **Sources** to set up your recovery task.
2. Click into the **Source** name and select the **OneDrive** tab.
3. Above the tree, select **Protection Status > Protected**.
4. Use the search and filter options to find and select the OneDrive you need, click the **Actions** (:) menu on that row, and select **Recover OneDrive Documents** to open the **New Recovery** form.
5. On the **New Recovery Microsoft 365 - OneDrive** page, in the **Recovery Type**, select any one of the following tabs to search for the file or folder:
 - **Browse OneDrive and Recover.** You can browse the individual user OneDrive to navigate and select the files or folders to be restored.
 - **Search Files and Recover.** You can use the global search to find the files and folders that need to be restored.
6. To *browse* and recover:
 1. In the **Recovery Type** section, select **Browse OneDrive and Recover**.
 2. Select the file or folder you plan to restore. Do any one of the following based on your requirements:
 - To recover the file(s) or folder(s), click **Next**.
 - To download the file(s) or folder(s), click **Download Files**.A new recovery task is created to download the file(s) or folder(s). When the task completes, from the **Activity** page, click the task name and then click **Download Files** to download the generated zip file.
 3. Click **Next: Recover Options** to return to the form and skip to **step 8**.
7. To *search* and recover:
 - a. In the **Recovery Type** section, select **Search Files and Recover**.
 - b. Use the '*' wildcard character and/or enter text to search for the folders or files with a matching folder name or file name in the **Search** bar. Select the folders or files to recover from the search results.

Or

Click **Advanced Search** and select **Both**, **Files**, or **Folder** and search based on the available filters and click **Apply**.

- c. To use a different **Recovery Point** for a folder or file, click the **Edit** icon on the tile for that folder or file. Find the recovery point you need and click **Select Recovery Point**.
 - d. Click **Next: Recover Options** to return to the form.
8. Under **Recover To**, select **Original Location** or **New Location**.
 - If you choose **Original Location**, the existing document library is overwritten.
 - If you choose **New Location**, select a **Registered Source** and the **Target Site**, and specify the **Document Library** name to which you plan to recover the document library items. Optionally, you can also enter a **new prefix for the Document Library**.

Note: If a folder with the specified name does not exist in the OneDrive, Cohesity creates the folder and recovers the OneDrive contents to that folder.

9. Select your **Recovery Options**:
 - **Continue on Error**. Enable to recover even if errors occur when recovering the document library items. For example, if a document cannot be recovered, will still attempt to recover the other selected documents from that document library.
 - **Task Name**. Change the default name of the recovery task.
10. Click **Start Recovery**.

Next > Protect your Microsoft 365 [Mailboxes](#), [SharePoint Online Sites](#), and [Teams](#) so you can recover them easily when you need to, as well!

OneDrive Content Recovery Self-Service

Cohesity provides a self-service workflow to help the end users recover the Microsoft 365 OneDrive content by leveraging the Microsoft Entra ID (Azure Active Directory) login for user authentication.

Administrators can authorize the self-service workflow for users through the Security Groups.

Note: You can access the self-service portal through <https://helios.cohesity.com/#/self-service-portal/auth>.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

OneDrive Content Recovery

To recover OneDrive content from the Cohesity Self-Service Portal:

1. In the Cohesity Data Cloud login page, log in through Microsoft using the **Cohesity Self-Service Portal** link.
2. Choose **Microsoft 365 OneDrive** to recover your OneDrive Files and Folders.
3. In the **Recover** page, browse the OneDrive or search for the OneDrive content.
4. Use the '*' wildcard character or enter the text to search for the folders or files with a matching subject in the **Search** bar. Select the folders or files to recover from the search results.

Or

Click **Advanced Search**, select **Both**, **Files**, or **Folder**, search based on the available filters, and click **Apply**.

5. To use a different **Recovery Point** for a folder or file, click the **Edit** icon on the tile for that folder or file. Find the recovery point you need and click **Select Recovery Point**.
6. Under **Recover Type**, select **Recover to Original Location** to recover all the items directly to your Mailbox.
7. Click **Finish**. You can view the recovery progress from the **Welcome** page under the **Recoveries** section or on the **Activity** page.
8. Click the action icon on the required task and click **Show Recovered Items** to view the name and size of the recovered items.

SharePoint Online

SharePoint Online is a SaaS application bundled with the Microsoft 365 service. It provides an extensive range of collaborative and creative capabilities enabling organizations to share, manage, and access information from almost any device.

Using the policy-based data protection solution from Cohesity DataProtect as a Service, you can backup and recover the SharePoint Online site templates. Thus enabling you to backup and recover the SharePoint Online sites or subsites and its contents such as document libraries and so on.

Considerations

Review and understand the following considerations before you protect your Microsoft 365 SharePoint Online data:

- Document libraries enabled with the ForceCheckout option are not recovered.
- Recovery of sites with the out-of-the-box (OOTB) modern theme or composed look is not supported.
- Backup and recovery of the site or subsite URLs with non-ANSI characters are not supported.
- Recovery of a site collection is not supported if the site URL has changed after the backup.
- From the recovery workflow, you cannot download an empty folder.
- Cohesity DataProtect as a Service discovers and protects the SharePoint Online sites created in the central storage location of your Microsoft 365 tenant.
- SharePoint Online sites created in satellite storage locations of your Microsoft 365 tenant can be discovered and protected. You can recover to the same tenant and same location or an alternate tenant and default location. For more details, see [Multi-Geo Capabilities in SharePoint Online](#).

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- Cohesity currently does not support Geo-Stretched Microsoft 365 tenants.
- Recovering shared permissions for files in the PHL drive is not supported.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- Data from PHL is not searchable.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- Backup of checked-out files in SharePoint is not supported.
- Custom scripts setting is not supported in SharePoint.

- Granular recovery of Team and Group sites is not supported when Site Tagging is enabled.
- During SharePoint granular recovery, you cannot browse or download files from the document libraries that contain a slash (/) in its name. As a workaround, you can perform a full recovery.
- File/folder permissions with no role will not be recovered since Microsoft does not allow adding such permissions via Graph APIs. For example, no role will be assigned when a Restricted View permission is added to a user from Advanced Setting.
- [Microsoft 365 Backup Storage](#) service is supported for SharePoint Online.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- SharePoint sites that are archived or locked will not be backed up and the source discovery may complete early due to these locked sites. Contact your Cohesity account team to configure the cluster to ignore the locked sites during discovery.

SharePoint Lists

Note: This is an Early Access feature. Contact your Cohesity account team to enable the feature.

- Recovery of comments does not include the commenter name and the actual time the comment was added.
- Item-level granular recovery is not supported.
- [External lists](#) are not supported.
- Hidden lists backup is supported, while system hidden lists (catalogs) recovery is not supported.
- Hidden lists (apart from catalogs) are recovered as not hidden.
- Recovery of sites (to the original or alternate location) creates new lists with the current DateTime suffix.
- Attachments with more than 4 MB size are not backed up.
- Columns of the type *User* in the recovered lists may display incorrect user for alternate restore.
- Recovered lists for the template type *Playlist* do not display the embedded video.
- For the comments that include mentions, if any of the mentioned users are deleted, the names of the users are displayed instead of the mentions.

- For embedded images, the linkage and thumbnail may be broken in the restored list.
- Lists items backup is not supported for Team Sites under Teams and Group sites under Groups.
- List views are not supported.
- Recovery of lists fails for the following:
 - if the list includes a lookup column.
 - if the list includes a managed metadata column.

Protect Microsoft 365 SharePoint Online Sites

Once you have [registered your Microsoft 365 domain](#) as a source, you're ready to protect the SharePoint Online sites in your domain.

To protect your Microsoft 365 SharePoint Online sites:

1. In **DataProtect as a Service**, under **Sources**, find the Microsoft 365 source and click into it.
2. Click the **Site** tab.
3. Select the individual SharePoint Online site you wish to protect or:
 - Click **Users > Select All Child Objects** to protect all the SharePoint Online sites in this source.
 - Click **Users > Auto Protect This** to protect all the SharePoint Online sites in this source.
4. Click the **Protect** icon above the list.
5. Choose a policy to specify backup frequency and retention. If you don't have a policy, you can easily [create one](#).
6. Under **Settings**, edit the **Start Time** if necessary.
7. Under **Additional Settings**, you can enable **Indexing**, configure a specific **End Date, Alerts**, and other [additional settings](#).

Note: If you plan to recover individual [document library items](#), in addition to [whole sites](#), you need to enable **Indexing** in this step. When you do, you can include or exclude specific sites from indexing.

8. Click **Protect**.

Note: The backups start immediately after you protect the objects, regardless of the time you set for the protection run.

Next > When the first protection run completes, you will be ready to [recover your protected SharePoint Online sites](#) when and if you need to.

Additional Settings

Advance Settings	Description
Start Time	Available only if the selected policy is set to Backup Daily . Indicates what time the protection run should start. Enter the Start Time and select AM or PM . The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone here.
SLA	The service-level agreement (SLA) defines how long the administrator expects a protection run to take. Enter: <ul style="list-style-type: none"> • Full. The number of minutes you expect a full protection run, which captures all the blocks in an object, to take. • Incremental. The number of minutes you expect an incremental protection run, which captures only the changed blocks in an object, to take.
Cancel Runs at Quiet Time Start	<i>(Available only if the selected policy has at least one Quiet Time)</i> When enabled, all the protection runs that are currently executing will cancel when the Quiet Time period starts. By default, this setting is disabled, meaning that after a protection run starts, it continues to execute even when a Quiet Time period starts. However, new protection runs will not start during a Quiet Time.
Indexing	By default, indexing is enabled. <div> Note: Indexing is mandatory for granular restore of SharePoint sites. </div>

Manage Existing Protection


Edit protection settings, change the policy, and start, stop, & pause protection.

Once you have [applied protection](#) to the objects in your sources, Cohesity DataProtect as a Service makes it easy to make changes to that protection quickly. You can:

- Edit additional settings like **End Date**, **Exclusions**, **Alerts**, and more.
- Apply a different policy.
- Start an on-demand protection run, pause and resume it, or even remove protection.

Edit Protection Settings

To edit protection settings:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click into the **Source** name.
3. Select **Show All > Protected** and use the other filters, search box, and views at the top to narrow your search.
4. Click the **Actions** menu () next to the object and select **Edit Protection** to open the protection settings for that object.

Apply a New Protection Policy

To change the **Policy**, click the drop-down and select a different policy. To help you choose, each policy in the list shows the **Backup** frequency and the **Retain** period for each backup.

If you don't have a policy that meets your needs, scroll to the bottom of the list and click **Create Policy** to [create your own policy](#).

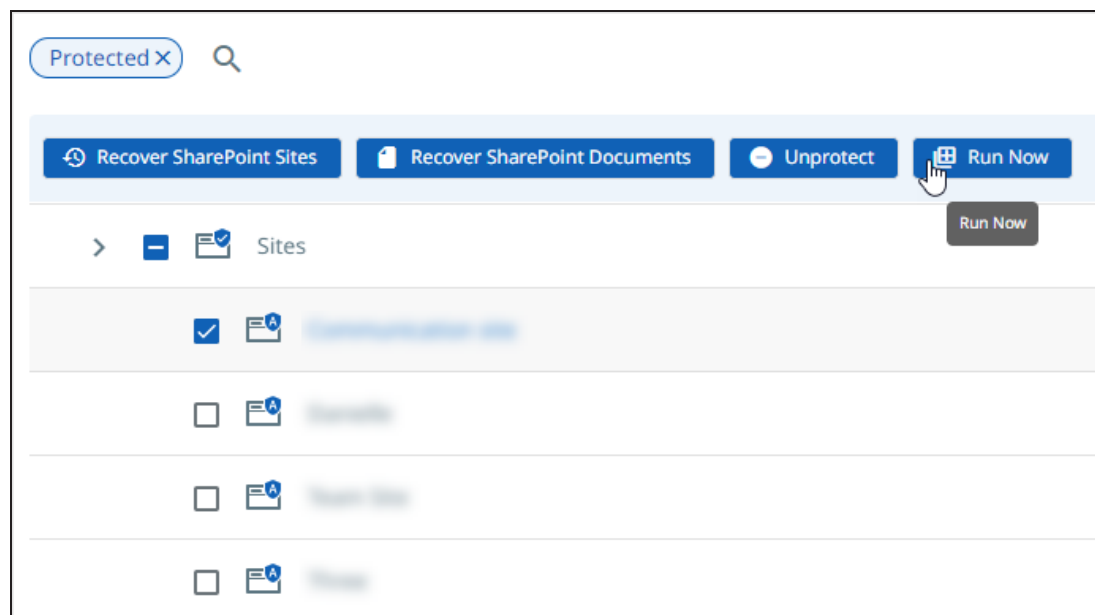
Edit Additional Protection Settings

Under **Settings**, you can change the protection **Start Time** (and select the **Time Zone**).

Click the drop-down next to **Additional Settings** to change more options. See [Additional Protection Settings](#) for details.

Start, Stop, or Remove Protection

When you select protected objects in one of your sources, Cohesity DataProtect as a Service presents buttons for the actions that are possible for those objects.



With the protected objects selected, you can click:

- **Recover SharePoint Sites** to recover the SharePoint sites.
- **Recover SharePoint Documents** to recover the SharePoint documents.

- **Run Now** to start an on-demand protection run immediately.
- **Unprotect** to remove protection from the object.

Tip: If a protected object is deleted from the source, you can search the object using Global Search and unprotect it.

Additional Settings

Advance Settings	Description
Start Time	Available only if the selected policy is set to Backup Daily . Indicates what time the protection run should start. Enter the Start Time and select AM or PM . The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone here.
SLA	The service-level agreement (SLA) defines how long the administrator expects a protection run to take. Enter: <ul style="list-style-type: none"> • Full. The number of minutes you expect a full protection run, which captures all the blocks in an object, to take. • Incremental. The number of minutes you expect an incremental protection run, which captures only the changed blocks in an object, to take.
Cancel Runs at Quiet Time Start	<i>(Available only if the selected policy has at least one Quiet Time)</i> When enabled, all the protection runs that are currently executing will cancel when the Quiet Time period starts. By default, this setting is disabled, meaning that after a protection run starts, it continues to execute even when a Quiet Time period starts. However, new protection runs will not start during a Quiet Time.
Indexing	By default, indexing is enabled. Note: Indexing is mandatory for granular restore of SharePoint sites.

Recover Microsoft 365 SharePoint Online Sites & Items

After you [protect your users' Microsoft 365 SharePoint Online sites](#), you can recover them — as [whole sites](#) or just [specific document library items](#) — from Cohesity DataProtect as a Service, to the same location, alternate location, or across Microsoft 365 domains.

Note: To recover site system files such as HTML, Javascript, and so on, ensure that you enable Custom Scripts permissions on the tenant. For more information, see [Tenant Permissions](#) in Microsoft 365 Requirements.

You can recover:

- [SharePoint Sites](#)
- [SharePoint Document Library Items](#)

Recover SharePoint Sites

To recover protected Microsoft 365 SharePoint Online sites:

1. In **DataProtect as a Service**, go to **Sources** to set up your recovery task.
2. Click on the **Source** name.
3. Above the tree, select **Show All > Protected**.
4. Find the sites you need and click the **Recover** button on that row to open the **New Recovery** form with the **Latest** snapshot (protection run).
5. In the **New Recovery** form, if you need to add more SharePoint Online sites and/or recover from an earlier backup, click the **Edit** icon in the top right of the form.
 - To add SharePoint Online sites, enter a **Search** term on the left, locate the other SharePoint Online sites, and select them.
 - To use a different Recovery Point for a site, click the **Edit** icon on the tile for that site. Find the recovery point you need and click **Select Recovery Point**.

Note: To recover a site collection and its sub-sites, search using the site collection relative URL such as `"/sites/myrootsite"` and add them to the recovery task.

Click **Next: Recover Options** to return to the form.

6. Under **Recover To**, select **Original Location** or **New Location**.

If you choose **New Location**, select a **Registered Source** and the **Target**.

Note: Sites created in satellite storage locations are recovered to the same location in the target tenant as the source tenant. If the same location is not available, the sites are recovered to the central location.
This is an Early Access feature. Contact your Cohesity account team to enable the feature.

7. Select your **Recovery Options**:

- **Continue on Error.** Enable to recover even if errors occur when recovering SharePoint Online sites. For example, if one of the sites cannot be recovered, Cohesity will still attempt to recover the other selected sites.
- **Include Preservation Hold Library.** Enable to recover the Preservation Hold Library that is part of the Cohesity snapshot. Recovering the Preservation Hold Library data may increase the recovery time substantially as it can include a large amount of data. The recovery will fail if the Target Location does not have sufficient space.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- **Task Name.** Change the default name of the recovery task.

8. Click **Start Recovery**.

Next > Protect your Microsoft 365 [Mailboxes](#), [OneDrives](#), and [Teams](#) so you can recover them easily when you need to, as well!

Recover SharePoint Document Library Items

Important: Before you can recover SharePoint document library items, you need to set up [SharePoint protection](#) with **Indexing** enabled.

To recover specific document library items from a protected Microsoft 365 SharePoint Online Site:

1. In **DataProtect as a Service**, go to **Sources** to set up your recovery task.
2. Click on the **Source** name and select the **Site** tab.
3. Use the search or filter options, find and select the site you need, click the **Actions** menu (:) on that row, and select **Recover SharePoint Documents** to open the **New Recovery** form.
4. In the **New Recovery Microsoft 365 - SharePoint Online** page, under the **Recovery Type** section, select any one of the following to search for the file or folder:
 - **Browse Site and Recover.** You can browse the individual site to navigate and select the files/document library to be restored.
 - **Search Files and Recover.** You can use the global search to find the files and document libraries that need to be restored.
5. To *browse* and recover:

- a. In the **Recovery Type** section, select **Browse Site and Recover**.
 - b. Search for the site name and click the site name to browse the site.
 - c. Select the file or document library you plan to restore. Do any one of the following based on your requirements:
 - i. To recover the file(s) or document library(s), click **Next**.
 - ii. To download the file(s) or document library(s), click **Download Files**.

A new recovery task is created to download the file(s) or document library (s). When the task completes, from the **Activity** page, click the task name and then click **Download Files** to download the generated zip file.
 - d. Click **Next: Recover Options** to return to the form and skip to **step 8**.
6. To *search* and recover:
1. In the **Recovery Type** section, select **Search Files and Recover**.
 2. Use the '*' wildcard character and/or enter text (such as '*.xlsx' or '*.pdf') to search for the folders or files with a matching folder name or file name in the **Search** bar. Select the folders or files to recover from the search results.
- Or**
- Click **Advanced Search** and select **Both, Files**, or **Folder** and search based on the available filters and click **Apply**.
3. To use a different **Recovery Point** for a folder or file, click the **Edit** icon on the tile for that folder or file. Find the recovery point you need and click **Select Recovery Point**.
 4. Click **Next: Recover Options** to return to the form.
7. Under **Recover To**, select **Original Location** or **New Location**.
- If you choose **Original Location**, the existing document library is overwritten.
 - If you choose **New Location**, select a **Registered Source** and the **Target Site**, and specify the **Document Library** name to which you plan to recover the document library items. Optionally, you can also enter a new **prefix for the Document Library**.
- Note:** If a document library with the specified name does not exist on the site, Cohesity creates the document library and recovers the folders or files to that document library.
8. Select your **Recovery Options**:

- **Continue on Error.** Enable to recover even if errors occur when recovering the document library items. For example, if a document cannot be recovered, Cohesity will still attempt to recover the other selected documents from that document library.
- **Task Name.** Change the default name of the recovery task.

9. Click **Start Recovery**.

Recover SharePoint Lists

Cohesity now supports the recovery of the Lists in Microsoft 365 SharePoint Online. Lists are a collection of data like links, announcements, contacts, issue trackers, surveys, and so on.

For more details, see [SharePoint Online](#).

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

Next > Protect your Microsoft 365 [Mailboxes](#), [OneDrives](#), and [Teams](#) so you can recover them easily when you need to, as well!

Microsoft Teams

Microsoft Teams is a collaboration solution provided by Microsoft that is bundled with the Microsoft 365 service. For more information, see [Microsoft documentation](#). Using the policy-based data protection solution from Cohesity DataProtect as a Service, you can backup and recover Teams data in Microsoft 365.

Considerations

Review and understand the following considerations before you protect your Microsoft 365 Teams data:

- Granular recovery of files and folders is supported.
- Backup and recovery of channel tabs are not supported.
- Backup and recovery of subsites of Teams site is not supported.
- If folders such as Feeds, Sync Issues, Legacy Archive Journals, Outbound, Managed Folders, Files, Yammer Root, Clutter, MeContact, and Archive, are not already present, the folders are skipped during recovery.
- Recovering the following Teams data from the Teams backup is not supported:

- Channel names and descriptions
- System Document Libraries
- Backup and download of the following is not supported:
 - Self-message (messages sent to self)
 - Saved or pinned Posts property in the conversation
 - Meeting recordings metadata
 - Shared Channels
- Granular recovery of Team sites is not supported when Site Tagging is enabled.
- Backup of Teams with no owners is supported when at least one Exchange Online licensed member is available in the Teams.

If no owners/members are available in the Teams, you can contact your Cohesity account team to configure a service account (with an Exchange Online license). This service account will be added as a member of the Teams before backup/recovery and removed after the backup/recovery is completed.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- File/folder permissions with no role will not be recovered since Microsoft does not allow adding such permissions via Graph APIs. For example, no role will be assigned when a Restricted View permission is added to a user from Advanced Setting.

Protect Microsoft 365 Teams

Once you have [registered your Microsoft 365 domain](#) as a source, you're ready to use Cohesity DataProtect as a Service to protect the Teams data in your domain.

To protect your Microsoft 365 Teams:

Note: If the Teams Posts option is enabled under the Teams during app registration, the Teams Posts will be backed up along with the corresponding Teams.

Teams backup may fail if the Azure subscription configuration is not set appropriately.

This is an Early Access feature. Contact your Cohesity account team to enable the feature.

1. In **DataProtect as a Service**, under **Sources**, find the Microsoft 365 source and click into it.
2. Click the **Teams** tab.
3. Select the individual Team you wish to protect or:
 - Click **Users > Select All Child Objects** to protect all the Teams in this source.
 - Click **Users > Auto Protect This** to protect all the Teams in this source.
4. Click the **Protect** icon above the list.
5. Choose a policy to specify backup frequency and retention. If you don't have a policy, you can easily [create one](#).
6. Under **Settings**, edit the **Start Time** if necessary.
7. Under **Additional Settings**, you can enable **Indexing**, configure a specific **End Date, Alerts**, and other [additional settings](#).

Note: If you plan to recover individual document library items (coming soon!), in addition to whole sites, you need to enable **Indexing** in this step. When you do, you can include or exclude specific sites from indexing.

8. Click **Protect**.

Note: The backups start immediately after you protect the objects, regardless of the time you set for the protection run.

Next > When the first protection run completes, you will be ready to [recover your protected Teams](#) when and if you need to.

Additional Settings

Advance Settings	Description
Start Time	Available only if the selected policy is set to Backup Daily . Indicates what time the protection run should start. Enter the Start Time and select AM or PM . The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone here .

Advance Settings	Description
SLA	<p>The service-level agreement (SLA) defines how long the administrator expects a protection run to take. Enter:</p> <ul style="list-style-type: none"> • Full. The number of minutes you expect a full protection run, which captures all the blocks in an object, to take. • Incremental. The number of minutes you expect an incremental protection run, which captures only the changed blocks in an object, to take.
Cancel Runs at Quiet Time Start	<p><i>(Available only if the selected policy has at least one Quiet Time)</i></p> <p>When enabled, all the protection runs that are currently executing will cancel when the Quiet Time period starts. By default, this setting is disabled, meaning that after a protection run starts, it continues to execute even when a Quiet Time period starts. However, new protection runs will not start during a Quiet Time.</p>
Indexing	<p>By default, indexing is enabled.</p> <div> <p>Note: Indexing is mandatory for granular restore of Teams contents.</p> </div>

Manage Existing Protection


Edit protection settings, change the policy, and start, stop, & pause protection.

Once you have [applied protection](#) to the objects in your sources, Cohesity DataProtect as a Service makes it easy to make changes to that protection quickly. You can:

- Edit additional settings like **End Date**, **Exclusions**, **Alerts**, and more.
- Apply a different policy.
- Start an on-demand protection run, pause and resume it, or even remove protection.

Edit Protection Settings

To edit protection settings:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click into the **Source** name.
3. Select **Show All > Protected** and use the other filters, search box, and views at the top to narrow your search.
4. Click the **Actions** menu () next to the object and select **Edit Protection** to open the protection settings for that object.

Apply a New Protection Policy

To change the **Policy**, click the drop-down and select a different policy. To help you choose, each policy in the list shows the **Backup** frequency and the **Retain** period for each backup.

If you don't have a policy that meets your needs, scroll to the bottom of the list and click **Create Policy** to [create your own policy](#).

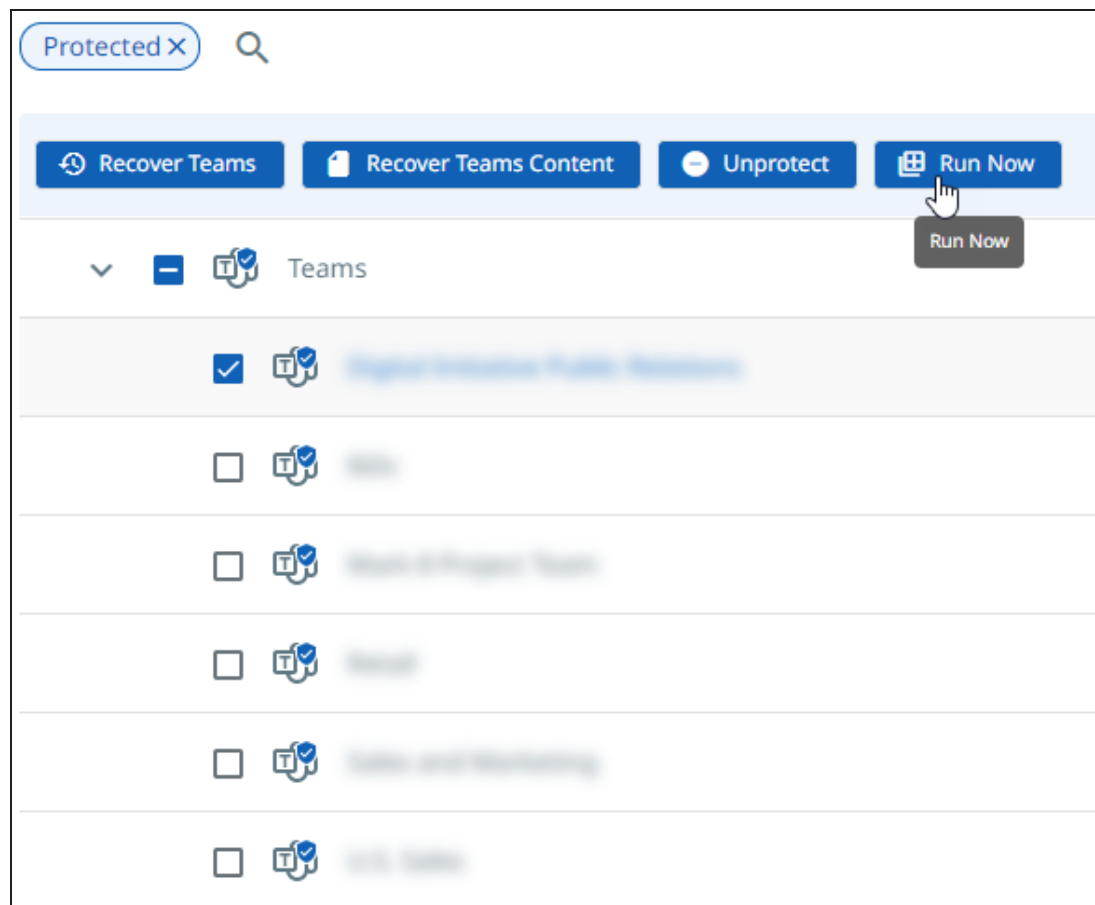
Edit Additional Protection Settings

Under **Settings**, you can change the protection **Start Time** (and select the **Time Zone**).

Click the drop-down next to **Additional Settings** to change more options. See [Additional Protection Settings](#) for details.

Start, Stop, or Remove Protection

When you select protected objects in one of your sources, Cohesity DataProtect as a Service presents buttons for the actions that are possible for those objects.



With the protected objects selected, you can click:

- **Recover Teams** to recover the Teams.
- **Recover Teams Content** to recover the Teams content.

- **Run Now** to start an on-demand protection run immediately.
- **Unprotect** to remove protection from the object.

Tip: If a protected object is deleted from the source, you can search the object using Global Search and unprotect it.

Additional Settings

Advance Settings	Description
Start Time	Available only if the selected policy is set to Backup Daily . Indicates what time the protection run should start. Enter the Start Time and select AM or PM . The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone here.
SLA	The service-level agreement (SLA) defines how long the administrator expects a protection run to take. Enter: <ul style="list-style-type: none"> • Full. The number of minutes you expect a full protection run, which captures all the blocks in an object, to take. • Incremental. The number of minutes you expect an incremental protection run, which captures only the changed blocks in an object, to take.
Cancel Runs at Quiet Time Start	<i>(Available only if the selected policy has at least one Quiet Time)</i> When enabled, all the protection runs that are currently executing will cancel when the Quiet Time period starts. By default, this setting is disabled, meaning that after a protection run starts, it continues to execute even when a Quiet Time period starts. However, new protection runs will not start during a Quiet Time.
Indexing	By default, indexing is enabled. Note: Indexing is mandatory for granular restore of Teams contents.

Recover Microsoft 365 Teams and Teams Content

After you protect your users' Teams, you can recover them — as [whole Teams](#) or just [specific Teams content](#) — from Cohesity DataProtect as a Service to the original Team in the same Microsoft 365 domain.

You can recover:

- [Microsoft 365 Teams](#)
- [Microsoft 365 Teams Content](#)

Recover Microsoft 365 Teams

To recover protected Microsoft 365 Teams:

1. In **DataProtect as a Service**, go to **Sources** to set up your recovery task.
2. Click on the **Source** name and select the **Teams** tab.
3. Above the tree, select **Show All > Protected**.
4. Find the Team you need and click the **Recover** button on that row to open the **New Recovery** form with the **Latest** snapshot (protection run).
5. In the **New Recovery** form, if you need to add more Teams and/or recover from an earlier backup, click the **Edit** icon in the top right of the form.
 - To add Teams, enter a **Search** term on the left, locate the other Teams, and select them.
 - To use a different **Recovery Point** for a Team, click the **Edit** icon on the tile for that Team. Find the recovery point you need and click **Select Recovery Point**.

Click **Next: Recover Options** to return to the form.

6. Under **Recover To**, select **Original Location** or **New Location**.
 - If you choose **Original Location**, the existing Teams content is overwritten.
 - If you choose **New Location**, select a **Registered Source** and the **Target Team** or click **Create New Team** to **create a new Team**.

For sample recovery use cases, see [Sample Teams Recovery Use Cases](#).

7. Select your **Recovery Options**:
 - **Restore Original Owner Members**. Disable the option if you do not want to restore the original owners and channel members to the newly created Team.
 - **Team Owner**. Select the target team owner that needs to be added to the original team owners.
 - **Continue on Error**. Enable to recover even if errors occur when recovering Teams. For example, if one of the Teams cannot be recovered, Cohesity will still attempt to recover the other selected Teams.
 - **Task Name**. Change the default name of the recovery task.
8. Click **Start Recovery**.

Sample Teams Recovery Use Cases

Use Case	Instructions
Restore the Team's data to the original Microsoft 365 domain.	Under Recover To , select Original Location .
Restore a Team's data to the original Microsoft 365 domain and add a new owner.	<ol style="list-style-type: none"> Under Recover To, select Original Location. In Recovery Options, disable Restore Original Owner Members and select an Additional Team Owner from the Team Owner drop-down list.
Restore the Team's data to a different Team that exists on the original Microsoft 365 domain.	<ol style="list-style-type: none"> Under Recover To, select New Location. Select the original Microsoft 365 domain as the Registered Source and select the Target Team to which you plan to restore the data.
Restore the Team's data to a different Team that exists on the original Microsoft 365 domain and add a new owner.	<ol style="list-style-type: none"> Under Recover To, select New Location. Select the original Microsoft 365 domain as the Registered Source and select the Target Team to which you plan to restore the data. Under Recovery Options, select an Additional Team Owner from the Team Owner drop-down list.
Restore the Team's data to a new Team on the original Microsoft 365 domain.	<ol style="list-style-type: none"> Under Recover To, select New Location. Select the original Microsoft 365 domain as the Registered Source, click Create New Team and specify the Team Name.
Restore the Team's data to a new Team on the original Microsoft 365 domain and add a new owner.	<ol style="list-style-type: none"> Under Recover To, select New Location. Select the original Microsoft 365 domain as the Registered Source, click Create New Team and specify the Team Name. Under Recovery Options, select an Additional Team Owner from the Team Owner drop-down list.

Use Case	Instructions
Restore the Team's data to a different Team that exists on a different Microsoft 365 domain.	<ol style="list-style-type: none"> 1. Under Recover To, select New Location. 2. Select the target Microsoft 365 domain from the Registered Source drop-down. 3. Select the Target Team to which you plan to restore the data.
Restore the Team's data to a new Team on a different Microsoft 365 domain. Also, add a new owner.	<ol style="list-style-type: none"> 1. Under Recover To, select New Location. 2. Select the target Microsoft 365 domain from the Registered Source drop-down. 3. Click Create New Team and specify the Team Name. 4. Under Recovery Options, select an Additional Team Owner from the Team Owner drop-down list.

Next > Protect your Microsoft 365 [Mailboxes](#), [OneDrives](#), and [SharePoint Online Sites](#) so you can recover them easily when you need to, as well!

Recover Microsoft 365 Teams Content

To recover specific content from a protected Microsoft 365 Team:

1. In **DataProtect as a Service**, go to **Sources** to set up your recovery task.
2. Click into the **Source** name and select the **Teams** tab.
3. Use the search or filter options, find and select the Team you need, and click **Recover Teams Content** on that row to open the **New Recovery** form.
4. Use the '*' wildcard character and/or enter the text (such as '*.xlsx' or '*.jpg') to search for the folders or files with a matching folder name or file name in the **Search** bar. Select the folders or files to recover from the search results.

Or

Click **Advanced Search** and select **Both**, **Files**, or **Folder** and search based on the available filters and click **Apply**.

5. To use a different **Recovery Point** for a folder or file, click the **Edit** icon on the tile for that folder or file. Find the recovery point you need and click **Select Recovery Point**.
6. Click **Next: Recover Options** to return to the form.
7. Under **Recover To**, select **Original Location** or **New Location**.

- If you choose **Original Location**, the existing Teams content is overwritten.
- If you choose **New Location**, select a **Registered Source** and the **Target Team** or click **Create New Team** to create a new Team. Then, select the **Target Channel** or click **Create New Channel** and select **Public** or **Private** to create a new channel in the selected Team. If you choose to create a **Private** channel then select the channel owner from the drop-down.
- Choose **Download** to download the Teams content.
A new recovery task is created to download the content. When the task is completed, from the **Activity** page, click the task name and then click **Download Files** to download the generated zip file.

Note: When you download multiple files from different snapshots, multiple recovery jobs are run and the files are downloaded separately.

For sample recovery use cases, see [Sample Teams Content Recovery Use Cases](#).

8. Select your **Recovery Options**:

- **Restore Original Owner Members.** Disable the option if you do not want to restore the original owners and channel members to the newly created Team.
- **Team Owner.** Select the target team owner that needs to be added to the original team owners.
- **Continue on Error.** Enable to recover even if errors occur when recovering Teams content. For example, if one of the Teams cannot be recovered, Cohesity will still attempt to recover the other selected Teams.
- **Task Name.** Change the default name of the recovery task.

9. Click **Start Recovery**.

[Sample Teams Content Recovery Use Cases](#)

Use Case	Instructions
Restore the Team's data to the original Microsoft 365 domain.	Under Recover To, select Original Location.
Restore a Team's data to the original Microsoft 365 domain and add a new owner.	<ol style="list-style-type: none"> 1. Under Recover To, select Original Location. 2. In Recovery Options, disable Restore Original Owner Members and select an Additional Team Owner from the Team Owner drop-down list.

Use Case	Instructions
<p>Restore the Team's data to a different Team that exists on the original Microsoft 365 domain.</p>	<ol style="list-style-type: none"> 1. Under Recover To, select New Location. 2. Select the original Microsoft 365 domain as the Registered Source and select the Target Team to which you plan to restore the data. 3. Select the Target Channel or click Create New Channel to create a new target channel to which you plan to restore the channel data.
<p>Restore the Team's data to a different Team that exists on the original Microsoft 365 domain. Also, restore the channel data to an existing channel.</p>	<ol style="list-style-type: none"> 1. Under Recover To, select New Location. 2. Select the original Microsoft 365 domain as the Registered Source and select the Target Team to which you plan to restore the data. 3. Select the Target Channel to which you plan to restore the channel data.
<p>Restore the Team's data to a different Team that exists on the original Microsoft 365 domain. Also, restore the channel data to a new channel.</p>	<ol style="list-style-type: none"> 1. Under Recover To, select New Location. 2. Select the original Microsoft 365 domain as the Registered Source and select the Target Team to which you plan to restore the data. 3. From the Target Channel drop-down, click Create New Channel to create a new channel in the selected Team.
<p>Restore the Team's data to a different Team that exists on the original Microsoft 365 domain and add a new owner.</p>	<ol style="list-style-type: none"> 1. Under Recover To, select New Location. 2. Select the original Microsoft 365 domain as the Registered Source and select the Target Team to which you plan to restore the data. 3. Select the Target Channel or click Create New Channel to create a new target channel to which you plan to restore the channel data. 4. Under Recovery Options, select an Additional Team Owner from the Team Owner drop-down list.

Use Case	Instructions
<p>Restore the Team's data to a new Team on the original Microsoft 365 domain.</p>	<ol style="list-style-type: none"> 1. Under Recover To, select New Location. 2. Select the original Microsoft 365 domain as the Registered Source, click Create New Team and specify the Team Name. 3. Select the Target Channel or click Create New Channel to create a new target channel to which you plan to restore the channel data.
<p>Restore the Team's data to a new Team on the original Microsoft 365 domain and add a new owner.</p>	<ol style="list-style-type: none"> 1. Under Recover To, select New Location. 2. Select the original Microsoft 365 domain as the Registered Source, click Create New Team and specify the Team Name. 3. Select the Target Channel or click Create New Channel to create a new target channel to which you plan to restore the channel data. 4. Under Recovery Options, select an Additional Team Owner from the Team Owner drop-down list.
<p>Restore the Team's data to a different Team that exists on a different Microsoft 365 domain.</p>	<ol style="list-style-type: none"> 1. Under Recover To, select New Location. 2. Select the target Microsoft 365 domain from the Registered Source drop-down. 3. Select the Target Team to which you plan to restore the data. 4. Select the Target Channel or click Create New Channel to create a new target channel to which you plan to restore the channel data.

Use Case	Instructions
Restore the Team's data to a new Team on a different Microsoft 365 domain. Also, add a new owner.	<ol style="list-style-type: none"> 1. Under Recover To, select New Location. 2. Select the target Microsoft 365 domain from the Registered Source drop-down. 3. Click Create New Team and specify the Team Name. 4. From the Target Channel drop-down, click Create New Channel to create a new channel in the selected Team. 5. Under Recovery Options, select an Additional Team Owner from the Team Owner drop-down list.

Download Teams Posts


To download Teams Posts from all Channels:

1. Navigate to **Sources**.
2. Click on the required source and click the **Teams** tab.
3. Select **Show All > Protected** and use the other filters, search box, and views at the top to narrow your search.
4. Click the **Actions** menu next to the object and click **Download Teams Posts**.
5. In the **Download Teams Posts** page, select the required snapshot, provide the task name, and click **Recover**.
6. Click the **View Progress** button in the pop-up message or click the **Activity** menu.
7. Once the recovery is successful, click **Download**. The Teams Posts will be downloaded by default in the **.htm** format.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

To download Teams Posts from a single Channel:

1. Navigate to **Sources**.
2. Click on the required source and click the **Teams** tab.
3. Select **Show All > Protected** and use the other filters, search box, and views at the top to narrow your search.

4. Click the **Recover Teams Content** icon () next to the object.
5. In the **Recover Teams Content** page, select the **Item Type** as **Channels**.
6. Hover over the required Channel and click the **Download Posts** button.
7. In the **Download Teams Items** page, select the required snapshot, provide the task name, and click **Download**.
8. Click the **View Progress** button in the pop-up message or click the **Activity** menu.
9. Once the recovery is successful, click **Download**. The Teams Posts will be downloaded.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

Next > Protect your Microsoft 365 [Mailboxes](#), [OneDrives](#), and [SharePoint Online Sites](#) so you can recover them easily when you need to, as well!

Microsoft Groups

Microsoft 365 groups are used for collaboration between users, both inside and outside your company. With each Microsoft 365 group, members get a group email and shared workspace for conversations, files, calendar events, and a planner.

Note: This is an Early Access feature. Contact your Cohesity account team to enable the feature for your tenant.

Using the policy-based data protection solution from Cohesity DataProtect as a Service, you can back up and recover Unified Groups data in Microsoft 365.

Considerations

Review and understand the following considerations before you protect your Microsoft 365 Groups data:

- Granular recovery of Group messages and other contents is not supported.
- Restoring system document libraries is not supported. You can restore only the non-system document libraries on a Group site.
- The entities protected for Groups include the SharePoint sites associated with the Group.

- Mail-enabled security groups, security groups, and distribution lists are not supported.
- Granular recovery of Group sites is not supported when Site Tagging is enabled.
- Backup of Groups (Public and Private) with no owners is supported when at least one Exchange Online licensed member is available in the Group.

If no owners/members are available in the Group, you can contact your Cohesity account team to configure a service account (with an Exchange Online license). This service account will be added as a member of the Group before backup/recovery and removed after the backup/recovery is completed.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

- File/folder permissions with no role will not be recovered since Microsoft does not allow adding such permissions via Graph APIs. For example, no role will be assigned when a Restricted View permission is added to a user from Advanced Setting.

Protect Microsoft 365 Groups

Once you have [registered your Microsoft 365 domain](#) as a source, you're ready to protect the Groups in your domain.

To protect your Microsoft 365 Groups:

1. In **DataProtect as a Service**, under **Sources**, find the Microsoft 365 source and click into it.
2. Click the **Group** tab.
3. Select the individual Team you wish to protect or:
 - Click **Groups** > **Select All Child Objects** to protect all the Teams in this source.
 - Click **Groups** > **Auto Protect This** to protect all the Teams in this source.
4. Click the **Protect** icon above the list.
5. Choose a policy to specify backup frequency and retention. If you don't have a policy, you can easily [create one](#).
6. Under **Settings**, edit the **Start Time** if necessary.
7. Under **Additional Settings**, you can enable **Indexing**, configure a specific **End Date**, **Alerts**, and other [additional settings](#).

Note: If you plan to recover individual document library items (coming soon!), in addition to whole sites, you need to enable **Indexing** in this step. When you do, you can include or exclude specific sites from indexing.

8. Click **Protect**.

Note: The backups start immediately after you protect the objects, regardless of the time you set for the protection run.

Additional Settings

Advance Settings	Description
Start Time	Available only if the selected policy is set to Backup Daily . Indicates what time the protection run should start. Enter the Start Time and select AM or PM . The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone here.
SLA	The service-level agreement (SLA) defines how long the administrator expects a protection run to take. Enter: <ul style="list-style-type: none"> • Full. The number of minutes you expect a full protection run, which captures all the blocks in an object, to take. • Incremental. The number of minutes you expect an incremental protection run, which captures only the changed blocks in an object, to take.
Cancel Runs at Quiet Time Start	<i>(Available only if the selected policy has at least one Quiet Time)</i> When enabled, all the protection runs that are currently executing will cancel when the Quiet Time period starts. By default, this setting is disabled, meaning that after a protection run starts, it continues to execute even when a Quiet Time period starts. However, new protection runs will not start during a Quiet Time.

Manage Existing Protection


Edit protection settings, change the policy, and start, stop, & pause protection.

Once you have [applied protection](#) to the objects in your sources, Cohesity DataProtect as a Service makes it easy to make changes to that protection quickly. You can:

- Edit additional settings like **End Date**, **Exclusions**, **Alerts**, and more.
- Apply a different policy.
- Start an on-demand protection run, pause and resume it, or even remove protection.

Edit Protection Settings

To edit protection settings:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click on the **Source** name.
3. Select **Show All > Protected** and use the other filters, search box, and views at the top to narrow your search.
4. Click the **Actions** menu () next to the object and select **Edit Protection** to open the protection settings for that object.

Apply a New Protection Policy

To change the **Policy**, click the drop-down and select a different policy. To help you choose, each policy in the list shows the **Backup** frequency and the **Retain** period for each backup.

If you don't have a policy that meets your needs, scroll to the bottom of the list and click **Create Policy** to [create your own policy](#).

Edit Additional Protection Settings

Under **Settings**, you can change the protection **Start Time** (and select the **Time Zone**).

Click the drop-down next to **Additional Settings** to change more options. See [Additional Protection Settings](#) for details.

Start, Stop, or Remove Protection

When you select protected objects in one of your sources, Cohesity DataProtect as a Service presents buttons for the actions that are possible for those objects.

With the protected objects selected, you can click:

- **Recover** to recover the Groups.
- **Run Now** to start an on-demand protection run immediately.
- **Unprotect** to remove protection from the object.

Tip: If a protected object is deleted from the source, you can search the object using Global Search and unprotect it.

Additional Settings

Advance Settings	Description
Start Time	Available only if the selected policy is set to Backup Daily . Indicates what time the protection run should start. Enter the Start Time and select AM or PM . The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone here.
SLA	The service-level agreement (SLA) defines how long the administrator expects a protection run to take. Enter: <ul style="list-style-type: none"> • Full. The number of minutes you expect a full protection run, which captures all the blocks in an object, to take. • Incremental. The number of minutes you expect an incremental protection run, which captures only the changed blocks in an object, to take.
Cancel Runs at Quiet Time Start	<i>(Available only if the selected policy has at least one Quiet Time)</i> When enabled, all the protection runs that are currently executing will cancel when the Quiet Time period starts. By default, this setting is disabled, meaning that after a protection run starts, it continues to execute even when a Quiet Time period starts. However, new protection runs will not start during a Quiet Time.

Recover Groups

After you protect the Groups in your domain, you can recover them as whole Groups from Cohesity DataProtect as a Service, to the same Microsoft 365 Group, to an alternate Microsoft 365 Group, or to a new Microsoft 365 Group in the same Microsoft 365 domain.

Points to note:

- Granular recovery of Group contents is not supported.
- If you're restoring a Group that does not exist in the Microsoft 365 domain, Cohesity creates a new Group with the metadata and data from the backup snapshot.
- If you restore to an existing Group, the group resources in the existing Microsoft 365 Group are overwritten or appended with the restored data. The following table details the group resources that are overwritten or appended:

Restore Behavior	Group Resource Type
Appended	members
	owners
	mails (data)
Overwritten	hideFromAddressLists
	hideFromOutlookClients
	displayName visibility
	securityEnabled
	description
	theme

You can restore the Microsoft 365 Group data to:

- The same Microsoft 365 Group.
- A different Microsoft 365 Group in the same Microsoft 365 domain.
- A new Microsoft 365 Group in the same Microsoft 365 domain.

To recover protected Microsoft 365 Group:

1. In **DataProtect as a Service**, go to **Sources** to set up your recovery task.
2. Click on the **Source** name and select the **Group** tab.
3. Above the tree, select **Show All > Protected**.
4. Find the Group you need and click the **Recover** button on that row to open the **New Recovery** form with the **Latest** snapshot (protection run).
5. In the **New Recovery** form, if you need to add more Groups and/or recover from an earlier backup, click the **Edit** icon in the top right of the form.
 1. To add Teams, enter a **Search** term on the left, locate the other Teams, and select them.
 2. To use a different **Recovery Point** for a Team, click the **Edit** icon on the tile for that Team. Find the recovery point you need and click **Select Recovery Point**.
6. Click **Next: Recover Options** to return to the form.
7. Under **Recover To**, select **Original Location** or **New Location**.

If you choose **New Location**, specify the **Group Name** and the **Group SMTP**.

8. Select your **Recovery Options**:

1. **Continue on Error**. Enable to recover even if errors occur when recovering Groups. For example, if one of the Groups cannot be recovered, Cohesity will still attempt to recover the other selected Groups.
2. **Task Name**. Change the default name of the recovery task.

9. Click **Start Recovery**.

Next > Protect your Microsoft 365 [Mailboxes](#), [OneDrives](#), [SharePoint Online Sites](#), and [Teams](#) so you can recover them easily when you need to, as well!

Microsoft Azure

Cohesity DataProtect as a Service provides a simple, fast, and cost-effective backup, recovery, and data management solution for Microsoft Azure Services:

To protect the Microsoft SQL Servers running in Azure VMs, see [Protect Microsoft SQL Server Running in Azure VM](#).

Microsoft Azure Virtual Machines

Cohesity DataProtect as a Service provides a simple, fast, and cost-effective backup, recovery, and data management solution for Microsoft Azure Virtual Machines in your Azure source.

Azure Requirements and Considerations

Before you register your Azure sources with Cohesity DataProtect as a Service, ensure the Azure VMs you want to backup are on the regions Cohesity supports, you've met the prerequisites and understood the considerations.

Note: For information on the supported cloud regions where you can back up this source, see [Supported Workloads and Cloud Regions](#).

Requirements

Before you register Azure with Cohesity DataProtect as a Service, ensure:

- To perform the following steps:
 - a. Register an application with Azure AD and create a service principal. For information, see the [Azure documentation](#).
 - b. Create an application secret key for setting up authentication for the service principal. For information, see the [Azure documentation](#).
 - c. Create a custom role at the subscription level with the required [permissions](#) for backup and recovery.
For information about creating a custom role, see the [Azure documentation](#).
 - d. Assign the custom role to the Azure AD application created in step a. For more information, see the [Azure documentation](#).

The application ID and application secret key are required when you register the Azure source with the Cohesity DataProtect as a Service.

- The ports listed in the Azure section in the [Firewall Ports](#) topic are open to allow communication between the Cohesity SaaS Connector(s) and Azure environment.
- SaaS Connectors are able to resolve the following URLs by name:
 - Login.windows.net
 - management.azure.com
 - *.blob.core.windows.net
- To whitelist *.blob.storage.azure.net.
- Cohesity DataProtect as a Service supports the [regions](#) where the Azure VMs you want to backup is located.
- To recover files and folders of Azure VMs:
 - Firewall port 50051 must be open on the target VM.
 - The target VM must be reachable via a private IP from the SaaS connector.

Required Permissions

Resource Provider	Operation Name
Microsoft.Resources	Microsoft.Resources/subscriptions/resourceGroups/read Microsoft.Resources/subscriptions/resourceGroups/write
Microsoft.Storage	Microsoft.Storage/storageAccounts/blobServices/containers/read Microsoft.Storage/storageAccounts/blobServices/containers/write Microsoft.Storage/storageAccounts/listkeys/action Microsoft.Storage/storageAccounts/read Microsoft.Storage/storageAccounts/write

Resource Provider	Operation Name
Microsoft.Network	Microsoft.Network/dnszones/A/read
	Microsoft.Network/dnszones/A/write
	Microsoft.Network/dnszones/A/delete
	Microsoft.Network/networkInterfaces/write
	Microsoft.Network/networkInterfaces/read
	Microsoft.Network/networkInterfaces/join/action
	Microsoft.Network/networkInterfaces/delete
	Microsoft.Network/networkInterfaces/ipconfigurations/read
	Microsoft.Network/networkSecurityGroups/read
	Microsoft.Network/networkSecurityGroups/join/action
	Microsoft.Network/networkSecurityGroups/securityRules/read
	Microsoft.Network/privateEndpoints/read
	Microsoft.Network/privateEndpoints/write
	Microsoft.Network/privateEndpoints/delete
	Microsoft.Network/virtualNetworks/read
	Microsoft.Network/virtualNetworks/subnets/read
	Microsoft.Network/virtualNetworks/subnets/join/action
	Microsoft.Network/virtualNetworks/subnets/joinViaServiceEndpoint/action

Resource Provider	Operation Name
Microsoft.Compute	Microsoft.Compute/disks/beginGetAccess/action
	Microsoft.Compute/disks/delete
	Microsoft.Compute/disks/endGetAccess/action
	Microsoft.Compute/disks/read
	Microsoft.Compute/disks/write
	Microsoft.Compute/virtualMachines/start/action
	Microsoft.Compute/virtualMachines/powerOff/action
	Microsoft.Compute/virtualMachines/write
	Microsoft.Compute/virtualMachines/read
	Microsoft.Compute/virtualMachines/delete
	Microsoft.Compute/virtualMachines/runCommand/action
	Microsoft.Compute/virtualMachines/deallocate/action
	Microsoft.Compute/snapshots/write
	Microsoft.Compute/snapshots/read
	Microsoft.Compute/snapshots/beginGetAccess/action
	Microsoft.Compute/snapshots/endGetAccess/action
	Microsoft.Compute/snapshots/delete
	Microsoft.Compute/diskAccesses/write
	Microsoft.Compute/diskAccesses/read
	Microsoft.Compute/diskAccesses/delete
	Microsoft.Compute/diskAccesses/privateEndpointConnectionsApproval/action
Microsoft.KeyVault	Microsoft.KeyVault/vaults/deploy/action

Considerations

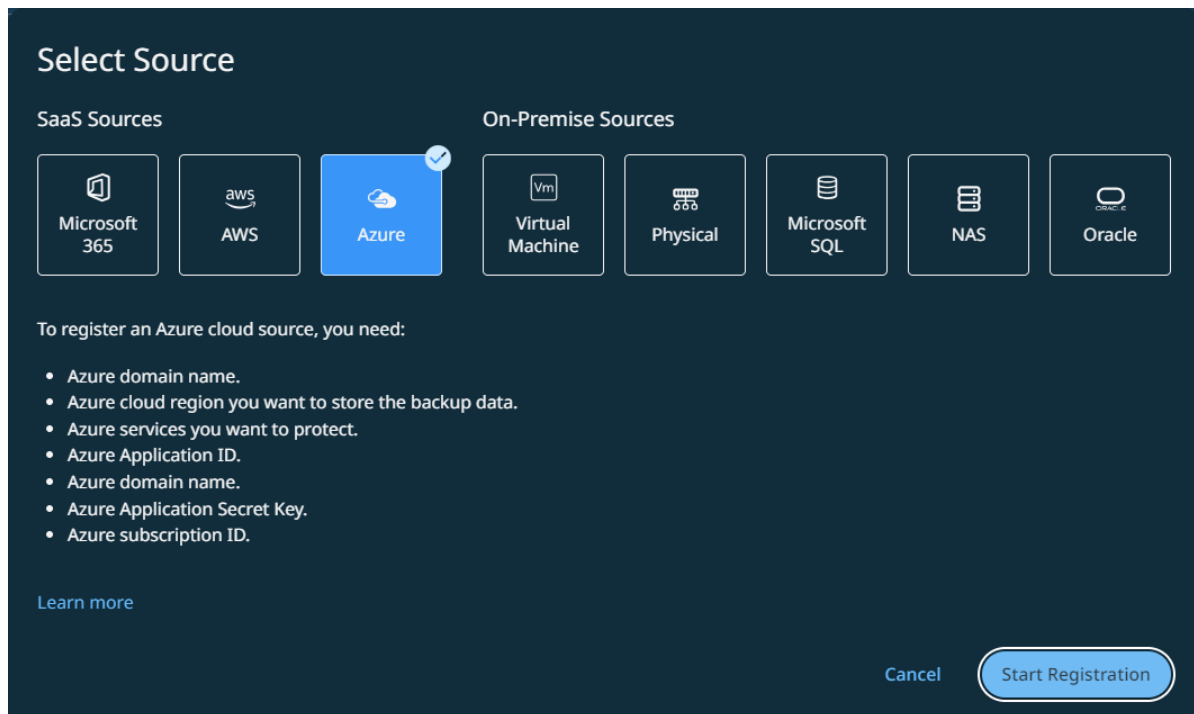
- Cohesity DataProtect as a Service supports the protection of UEFI boot mode-enabled Azure VMs running on the following guest operating systems:
 - Windows 2012, 2016, 2019
 - Ubuntu 14, 16, 18, 20
 - RHEL 6.x, 7.x, 8.x
 - Centos 6.x, 7.x

- All the disk sizes on the Azure VM must be a multiple of 1MB because Azure only allows creating disks whose size is a multiple of 1MB.
- Cohesity DataProtect as a Service supports the protection of Azure VMs with the following configurations:
 - Managed disks - Standard_LRS, Premium_LRS, StandardSSD_LRS, Premium_ZRS, or StandardSSD_ZRS.
 - Unmanaged disks
- VMs encrypted through ADE cannot be restored to a different location unless the user replicates the keys used to encrypt the VM to the new location. VMs encrypted using Azure SSE do not have this issue.
- Managed disk VMs in turn-off state are shown as 0 bytes in size in the entity hierarchy of Azure Source. However, backup and recovery of the VM is supported in a turned-off state.
- Recovery of the unmanaged disk with different SKU types will depend on the storage container where the recovery is performed.
- Recovery of unmanaged disk VM to the original location, scans for the same resource group, storage account, storage container & blobs created during backup. If these resources are deleted, the restore will fail.
- Recovery of the VMs from Availability set to a different location will not have the Availability set parameters.
- VMs with static IP will not be recovered back with static IP.
- You can perform files and folder recovery to an Azure VM Linux instance only if the Azure VM Linux instance is of x86_64 architecture.
- Cohesity DataProtect as a Service does not support:
 - File download.
 - Cloning.
 - App consistent backup.
 - LDM for OS disk.
 - Backup of shared managed disks.
 - Backup of Managed Ultra disks.
 - Backup of the Azure VM having ephemeral volumes.
 - Backup of Azure disks of more than 8TB using a private endpoint.
 - Backup of Azure VMs of new disk of type, Premium SSD v2 LRS.
 - Azure Stack Hub for Azure VM backups.

Register Azure Source

To start protecting your Microsoft Azure services, check the Azure [requirements](#) and then register Azure as a data source in Cohesity DataProtect as a Service.

1. In **DataProtect as a Service**, navigate to **Sources > + Register Source**, and then select **Azure**.



2. Click **Start Registration**.

The Register Azure Source form appears.

3. In the **Domain Name** field, enter the unique tenant ID assigned by Azure, and then select the **Destination cloud region**.
4. Select the Azure services you want to register:
 - **Azure VM**
 - **Azure SQL**Azure VM is selected by default.
5. Click **Continue**.

The screenshot shows the 'Register Azure Source' dialog with three steps: 1. Source Details, 2. Azure Apps, and 3. (unlabeled). Step 1 is active. It contains a 'Domain Name' text field with a red underline and a note 'This is typically the Azure tenant id.' below it. To the right is a 'Destination Cloud Region' dropdown menu. Under the 'Azure Services' section, there is a heading 'Choose Azure Platform to discover' and two options: 'Azure VM' with a checked checkbox and a VM icon, and 'Azure SQL' with an unchecked checkbox and a database icon. At the bottom right are 'Cancel' and 'Continue' buttons.

6. In the **Azure App ID** field, enter the application ID assigned by Azure during the service principal creation process.
7. In the **Azure App Secret Key** field, enter the application secret key for setting up the authentication for the service principal.
8. Click **Continue**.

The screenshot shows the 'Register Azure Source' dialog at Step 2: Azure Apps. Step 1 is marked with a checkmark. A blue information banner at the top states: 'Configuring Azure apps manually requires you to perform certain steps in Microsoft Azure as well. [Show manual registration to-do list.](#)'. Below this, the 'Azure Apps' section has two fields: 'Azure App ID' with the value '2eewff55-3aew-41cdw-9315-fwdwf66a95' and 'Azure App Secret Key' with a masked value '.....'. At the bottom are 'Back', 'Cancel', and 'Continue' buttons.

9. In the **Azure Subscription ID** field, enter the subscription IDs of the subscriptions where the VMs you want to protect belong.
10. Click **Register**.

Next > Once you register the Azure source, you must set up a [SaaS Connection](#) for each region under each Azure subscription in your Azure source.

Edit Azure Source

You can edit the registered Azure Source to add or remove the Azure services, application details, and subscriptions protected by the Cohesity DataProtect as a Service from your Azure source.

To edit an Azure source:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the Actions menu (:) next to the Azure sources and select **Edit**.
3. In the **Edit Azure Source** form, select or unselect the Azure services you need and click **Continue** to update Azure Application ID, Azure secret key, or to add or remove subscriptions.
4. Once you edit the Azure Source, click **Update**.

Unregister Azure Source

If you plan to stop backing up your Azure VMs or Azure SQL, you can unregister the Azure source from Cohesity DataProtect as a Service.

Note: Before you unregister an Azure source from Cohesity DataProtect as a Service, you must unprotect all the protected objects in that Azure source and delete the SaaS Connections.

To unregister the Azure Source:

1. In **DataProtect as a Service**, navigate to Sources.
2. Click the Actions menu (:) next to the Azure source and select **Unregister**.
3. In the **Unregister Source** form, click **Unregister**.

Protect Azure Virtual Machines

Before you protect the Azure VMs in your Azure source, you must set up a [SaaS Connection](#) for each region under each Azure subscription in your Azure source. A SaaS Connection consists of one or more SaaS Connectors, which are VMs that act as data movers between your data sources and the Cohesity DataProtect as a Service.

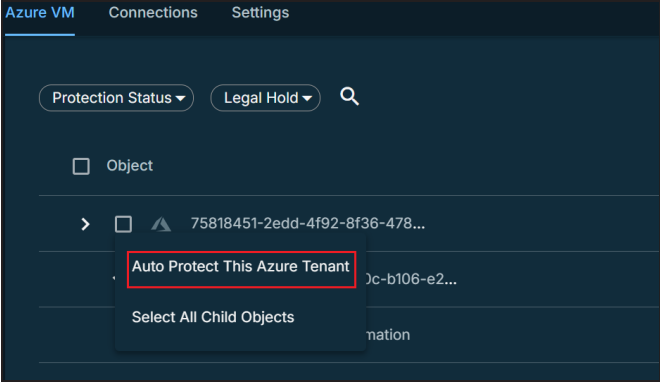
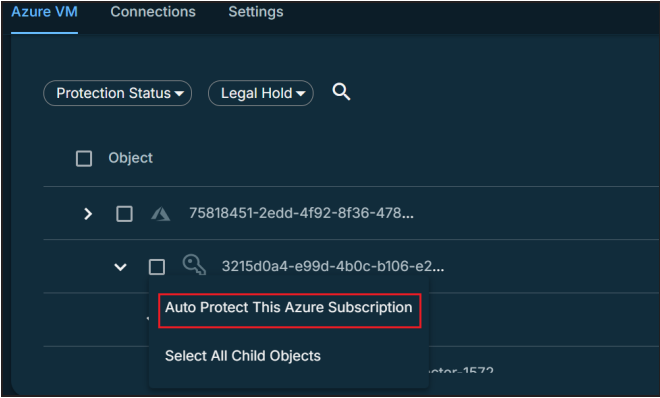
Once you set up a SaaS Connection, you are ready to protect the Azure VMs in the Azure source.

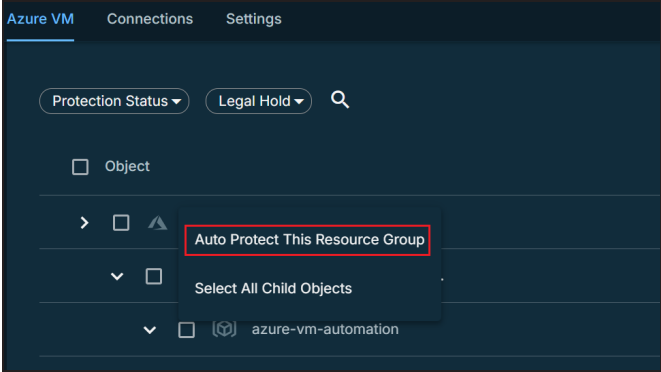
To protect your Azure virtual machines:

1. In **DataProtect as a Service**, navigate to **Sources**, find the registered Azure source and click into it.
2. Click the **Azure VM** tab.
3. Use the checkboxes to select the objects (VMs) for protection.

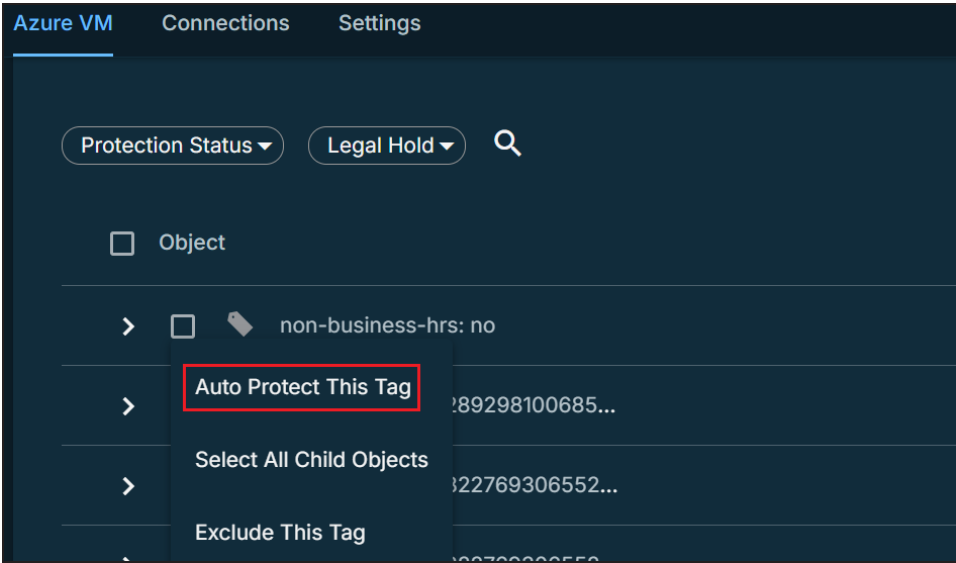
Optionally, you can configure auto-protect at the Azure tenant, subscription, or resource group levels. When this option is enabled at a particular level, all the VMs that are added to that level in the future are automatically protected from the next protection run. Additionally, you can also perform tag-based auto-protection of Azure VMs.



- To auto-protect the VMs based on the hierarchy level, click the **Hierarchy View** icon located at the right corner of the page, and perform one of the following steps listed in the table below:

To auto-protect the Azure VMs..	Action
At the tenant level	<div>Select the checkbox of the Azure source, and then select Auto Protect This Azure Tenant.</div> <div>A screenshot of the Azure VM interface. At the top, there are tabs for 'Azure VM', 'Connections', and 'Settings'. Below these are filters for 'Protection Status' and 'Legal Hold', followed by a search icon. A table lists objects. The first object has a checkbox that is selected. A context menu is open over this object, showing two options: 'Auto Protect This Azure Tenant' (which is highlighted with a red rectangular box) and 'Select All Child Objects'. The second object in the table is partially visible.</div>
At the subscription level	<div>Select the checkbox of the subscription, and then select Auto Protect This Subscription.</div> <div>A screenshot of the Azure VM interface, similar to the one above. It shows the same tabs and filters. In the table, the second object has a selected checkbox. A context menu is open over this object, showing two options: 'Auto Protect This Azure Subscription' (highlighted with a red rectangular box) and 'Select All Child Objects'. The first object in the table is partially visible.</div>

To auto-protect the Azure VMs..	Action
At the resource group level	<div>Select the checkbox of the resource group, and then select Auto Protect This Resource Group.</div> <div></div>

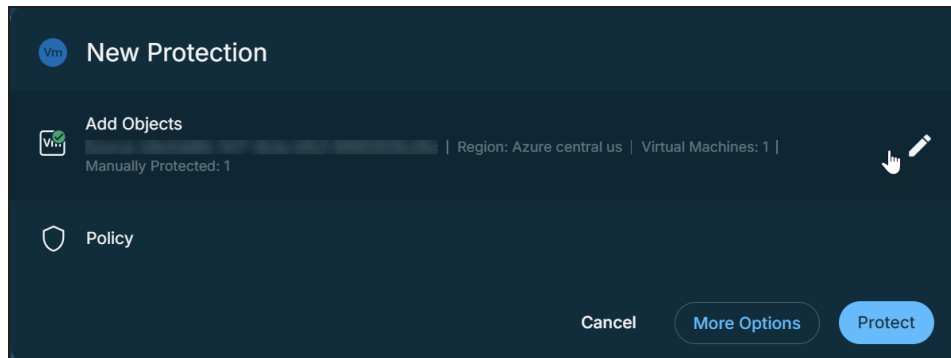
- To auto-protect the VMs based on tag, click the **Tag** icon at the right corner of the page. Select the checkbox of a tag and then select **Auto Protect this Tag** to auto-protect the VMs with this tag.



You can exclude individual VMs from auto-protection by clicking the auto-protect icon () next to the VM. VMs excluded from auto-protection are displayed with the exclude () icon.

4. Click the **Protect** icon above the checkboxes.
5. To exclude any disks of the VM from protection:

- a. Click the edit icon displayed next to **Add Objects**.



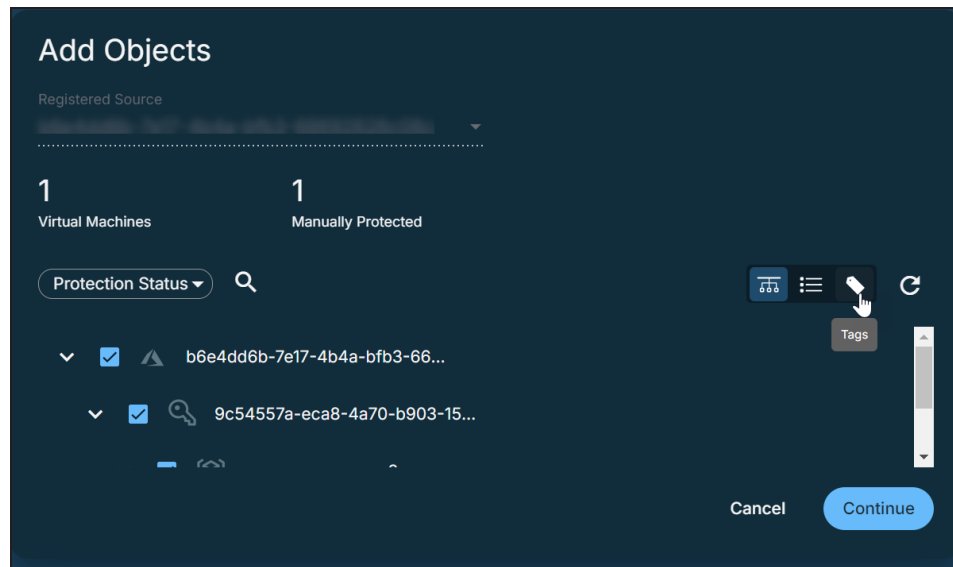
- b. Click the **All Volumes** edit icon displayed next to the VM of which you want to exclude the disk.
- c. Unselect the disks that you want to exclude from protection.

Note: You cannot exclude the root disk from protection.

- d. Click **Save > Continue**.

Note: You cannot exclude disk at the object level for auto-protected VMs. To exclude disk for auto-protect VMs, use the **Volume Exclusion Settings** option in the **Additional Settings**.

6. To exclude auto-protection of VMs based on Tags:
 - a. Click the edit icon displayed next to **Add Objects**.
 - b. Click the **Tags** icon located at the right corner of the page.



The tags associated with the VMs are displayed.

- c. Click the exclude icon next to the tag to exclude the auto-protection of VMs associated with that tag.
- d. Click **Continue**.
7. Choose a policy to specify backup frequency and retention. If you don't have a policy, you can easily [create one](#).
8. To change or configure any of the additional settings, select **More Options** and perform the below steps, or else, click **Protect**.
9. Under **Settings**, edit the **Start Time** if necessary.
10. In the **SLA** field, define how long the administrator expects a protection run to take. Enter:
 - a. **Full**. The number of minutes you expect a full protection run, which captures all the blocks in an object, to take.
 - b. **Incremental**. The number of minutes you expect an incremental protection run, which captures only the changed blocks in an object, to take.

The screenshot shows the 'Virtual Machines' configuration page in Cohesity. It is divided into several sections: 'Source' with a 'Registered Source' dropdown; 'Objects' showing '1 Virtual Machines' and '1 Manually Protected'; 'Policy' with a dropdown set to 'azure_policy' and a 'Backup' button labeled 'Every day Retain 2 weeks'; 'Settings' with 'Start Time' at '10:57pm | Asia/Calcutta', 'SLA' details ('Full: 1 day', 'Incremental: 1 day'), and an informational message: 'SLA will be met if Full Backups complete within 1 day and Incremental Backups complete within 1 day'; and 'Additional Settings' with a dropdown arrow. At the bottom are 'Protect' and 'Cancel' buttons.

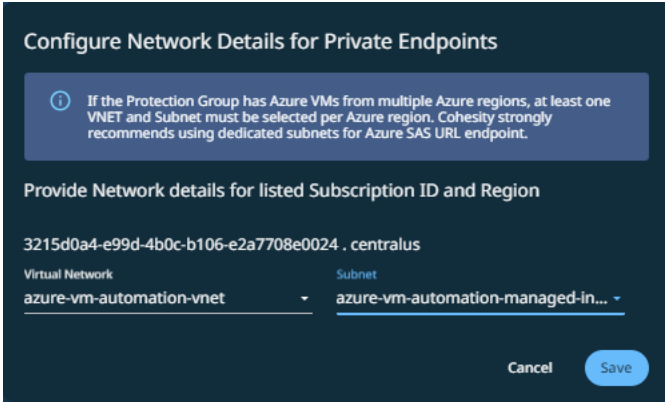
11. If you need to change any of the additional settings, click the down arrow icon next to **Additional Settings** and click **Edit**.

12. Click **Protect**.

Cohesity DataProtect as a Service starts backing up the Azure VMs you selected. You can monitor the status of the backup in the **Activity** page.

Also, the **Activity** tab of a specific Azure VM shows the history of all protection runs, including the one in progress.

Additional Settings

Advance Settings	Description
Pause Future Runs	Toggle on this option to stop protection runs from executing. Once you enable this option, no protection runs will be scheduled.
End Date	If you need to end protection on a specific date, enable this to select the date.
Azure SAS URL Type	<p>For the backup of Azure VMs, Cohesity uses the snapshots of Azure managed disks that are accessible through SAS URLs. Select one of the following endpoint options to access data from the SAS URL:</p> <ul style="list-style-type: none"> • Public Endpoint: Select this option to access snapshots over a public network. • Private Endpoint: Select this option to access snapshots using a private IP address from your virtual network. Click Configure Network Details for Private Endpoints and provide the following details for configuring the private endpoint: <ul style="list-style-type: none"> • Virtual Network: Select a virtual network in Azure. Ensure the virtual network of the SaaS connector and the private endpoint is the same. • Subnet: Select a subnet of the virtual network. 

Advance Settings	Description
Volume Exclusion Settings	<p>Use the following parameters to form a boolean expression for excluding disks:</p> <ul style="list-style-type: none"> • Name • automation • cohesity::helios::uuid • email • environment • expiry • location • non-business-hrs • purpose • rigel • squad • test • user • weekend <p>For example, the expression, <code>(name = net_volume AND environment IN (qa, dev)) OR squad = eng</code>, excludes all volumes from protection for which the volume tags meet the above condition.</p> <div> <p>Note: Exclusion of disks enabled at the object level will take precedence over the exclusion configured here.</p> </div>
Quiet Times	<p><i>(Available only if the selected policy has at least one Quiet Time.)</i></p> <p>When enabled, all the protection runs that are currently executing will cancel when the Quiet Time period starts. By default, this setting is disabled, meaning that after a protection run starts, it continues to execute even when a Quiet Time period starts. However, new protection runs will not start during a Quiet Time.</p>

Recover Azure Virtual Machines

After you protect your Azure VMs, you can recover them to their original location or a new location using Cohesity DataProtect as a Service.

Recover Azure VMs to Original Location

To recover your protected Azure VMs its original location:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the Source name.
3. Select **Protection Status > Protected**.
4. Use the filters, search box, and views to locate and select the Azure VMs you want to recover.

You can also use Global Search to locate, filter, and select the objects you need. Click the Global Search box at the top or type slash (/) anywhere to start your search.

5. Click the **Recover** icon at the top to open the **New Recovery** form. By default, the latest snapshot is pre-selected for recovery. If you need to recover from an earlier snapshot, click the **Edit** (pencil) icon to choose the desired snapshot, and then click **Select Recovery Point**.
6. Under **Recover To**, select **Original Location**.
7. Select the availability set to which the VM has to be recovered. For more information, see [Azure Availability Sets](#).
8. Select your **Recovery Options**:
 - a. **Rename**: Add a **Prefix** and/or **Suffix** to the recovered Azure VMs.
 - b. **Azure SAS URL Type**: For recovering Azure VMs, Cohesity uses the snapshots of Azure-managed disks that are accessible through SAS URLs. Select one of the following endpoint options to access data from the SAS URL:
 - a. **Public Endpoint**: Select this option to access snapshots over a public network.
 - b. **Private Endpoint**: Select this option to access snapshots using a private IP address from your virtual network. You can either use the same SAS URL details as protection or enable the Use custom defined VNET and Subscription and Region in Private Endpoint option, click **Configure Network Details for Private Endpoints**, and provide the following details for configuring the private endpoint:
 - **Virtual Network**: Select a virtual network in Azure. Ensure the virtual network of the SaaS connector and the private endpoint is the same.
 - **Subnet**: Select a subnet of the virtual network.

Note: If you have selected VMs from multiple regions, then you must select at least one virtual network and subnet per region.

Cohesity recommends using dedicated subnets for the Azure SAS URL endpoint.

- c. **Power State:** Disable **Power On** if you want the recovered VMs to remain powered off after they are created.
 - d. **Continue on Error:** Enable this option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.
9. **Task Name:** Change the default name of the recovery task.

The screenshot shows a 'Recover To' dialog box with the following sections:

- Recover To:** Two radio buttons: 'Original Location' (selected) and 'New Location'.
- Recovery Options:** A table with the following rows:

Rename	Prefix: copy-
Azure SAS URL Type	Public Endpoint
Power State	On
Continue on Error	No
Task Name	Recover_Dec_14_2023_3_44_PM
- Buttons:** 'Recover' (highlighted in blue) and 'Cancel'.

10. Click **Recover**.

Cohesity DataProtect as a Service opens the **Activity** page, showing your file recovery task as it runs, along with the recovery progress on the right.

Recover Azure VMs to New Location

To recover your protected Azure VMs to a new location:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the Source name.
3. Select **Protection Status > Protected**.
4. Use the filters, search box, and views to locate and select the Azure VMs you want to recover.

You can also use Global Search to locate, filter, and select the objects you need. Click the Global Search box at the top or type slash (/) anywhere to start your search.

5. Click the **Recover** icon at the top to open the **New Recovery** form. By default, the latest snapshot is pre-selected for recovery. If you need to recover from an earlier snapshot, click the **Edit** (pencil) icon to choose the desired snapshot, and then click **Select Recovery Point**.

6. Under **Recover To**, select **New Location**, and provide the following information:

Location:

- a. **Source:** Select a registered Azure source to recover the VMs to.
- b. **Subscription:** Select the subscription for the VMs selected for recovery.
- c. **Region:** Select a destination Azure region.
- d. **Resource Group:** Select a resource group to which the restored Azure VM will belong.
- e. **Compute:** Select the VM size type such as `Standard_D1` to use when recovering the VMs.
- f. **Availability Set:** Select the availability set to which the VM has to be recovered. For more information, see [Azure Availability Sets](#).

Storage Settings:

- a. **Storage Resource Group:** Select the storage container in which page blobs will be created.
- b. **Storage Account:** Select the Storage Account to which the recovery has to be done. The storage account you select must be in the same region as the resource group. Any storage account not in the same region as the parent resource group will not be displayed as an option in the recovery workflow.
- c. **Storage Container:** Select the storage container in which page blobs will be created.

Network Settings:

- a. **Network Resource Group:** Select the resource group for the virtual network.
- b. **Virtual Network:** Select the virtual network where the restored VM must be placed.
- c. **Subnet:** Subnets in the virtual network where restored VMs must be placed.

7. Select your **Recovery Options**:

- **Rename:** Add a **Prefix** and/or **Suffix** to the recovered Azure VMs.
- **Azure SAS URL Type:** For recovering Azure VMs, Cohesity uses the snapshots of Azure-managed disks that are accessible through SAS URLs. Select one of the following endpoint options to access data from the SAS URL:
 - **Public Endpoint:** Select this option to access snapshots over a public network.
 - **Private Endpoint:** Select this option to access snapshots using a private IP address from your virtual network. You can either use the same SAS URL details as protection or enable the Use custom defined VNET and

Subscription and Region in Private Endpoint option, click **Configure Network Details for Private Endpoints** link, and provide the following details for configuring the private endpoint:

- **Virtual Network:** Select a virtual network in Azure. Ensure the virtual network of the SaaS connector and the private endpoint is the same.
- **Subnet:** Select a subnet of the virtual network.

Note: If you have selected VMs from multiple regions, then you must select at least one virtual network and subnet per region.

Cohesity recommends using dedicated subnets for the Azure SAS URL endpoint.

- **Power State:** Disable **Power On** if you want the recovered VMs to remain powered off after they are created.
- **Continue on Error:** Enable this option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.
- **Task Name:** Change the default name of the recovery task.

Virtual Machines

Recover To

☐ Original Location ☒ New Location

Location

Source Subscription Region

Resource Group Compute Availability Set

Storage Settings

Storage Resource Group Storage Account Storage Container

Network Settings

Network Resource Group Virtual Network Subnet

Recovered VM will be created under this subnet.

Recovery Options

Rename Prefix: copy-

Azure SAS URL Type

Cohesity uses Azure managed disks snapshots which are accessible via SAS URLs. Select Public or Private endpoint for the SAS URL type.

☐ Public Endpoint ☒ Private Endpoint

If the has Azure VMs from multiple regions, atleast one VNET and subnet must be selected per Azure region.

Virtual Network Subnet

Power State On

Continue on Error No

Task Name Recover_Dec_14_2023_3_52_PM

Recover **Cancel**

8. Click **Recover**.

Cohesity DataProtect as a Service opens the **Activity** page, showing your file recovery task as it runs, along with the recovery progress on the right.

Recover Azure VM Files and Folders

You can download or restore specific files and folders from a protected Azure VM to the original or an alternate Azure VM.

To recover or download your files and folders from your protected Azure VM:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the **Source** name.
3. Select **Protection Status > Protected**.
4. Use the filters, search box, and views to locate and select the Azure VMs you want to recover.

You can also use **Global Search** to locate, filter, and select the objects you need. Click the Global Search box at the top or type **slash (/)** anywhere to start your search.

5. Click the required Azure VM backed up and click the **Recover Files** icon for the VM. The page with the VM details is displayed.
6. Select the timeline drop-down list on the top right corner to select the snapshot and click **Apply**.
7. Click the required volume to browse the file system and select the file or folder to be recovered.

You can click **Download Files** to download the selected files.

8. Click **Next**. The **Files** page is displayed.
9. Under **Recover To**, select **Original Server** or **New Server**.
 - For recovery to the *original Azure VM*, you can provide the new recovery path in the **Recover To** field or use the **Recover To Original Path** option to recover to the original path on the original server.
 - For recovery to an *alternate Azure VM*, you can choose any Azure VM and select a **Target**. Provide the new recovery path in the **Recover To** field.
10. Select your **Recovery Options**:
 - a. **Overwrite Existing File/Folder**: Enable this option to overwrite the existing files and folders. Disable this option to create the files and folders in the specified location. If a file with the same name already exists in the target location, the file is overwritten or skipped based on this selection.

If **Overwrite Existing File/Folder** is enabled, recovering a file to source when the file is in use may cause the open file to be overwritten. Whether overwriting occurs depends on the application using the file.
 - b. **Preserve File/Folder Attributes**: By default, this option is enabled and the ACLs, permissions, and timestamps are preserved for all files and folders. If you disable this option, then ACLs and permissions are not preserved. If both folders and files are recovered, then folders will receive the new timestamps, but files retain their original timestamps. If recovering only files, then files will receive the new timestamps.

- c. **Continue on Error:** Enable this option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.
- d. **Task Name:** Change the default name of the recovery task.

11. Click **Recover**.

Cohesity DataProtect as a Service opens the **Activity** page, showing your file recovery task as it runs, along with the recovery progress on the right.

Microsoft Azure SQL Database

Cohesity DataProtect as a Service provides a simple, fast, and cost-effective backup, recovery, and data management solution for growing Azure SQL Server database environments.

This topic covers the following:

- [Azure SQL Requirements and Considerations](#)
- [Register Azure SQL Source](#)
- [Protect Azure SQL Databases](#)
- [Recover Azure SQL Databases](#)
- [Manage Azure SQL Source](#)

Azure SQL Requirements and Considerations

Before you register an Azure SQL source on Cohesity DataProtect as a Service, ensure the following requirements are met:

1. Register an application with Azure Entra ID and create a service principal. For information, see the [Azure documentation](#).
2. Create an application secret key for setting up authentication for the service principal. For information, see the [Azure documentation](#).
3. Create a custom role at the subscription level with the required [permissions](#) for backup and recovery.
For information about creating a custom role, see the [Azure documentation](#).
4. Assign the custom role to the Azure Entra ID application created in step 1.
The application ID and application secret key are required when you register the Azure source with the Cohesity cluster.
5. The ports listed in the Azure section in the [Firewall Port](#) topic are open to allow communication between the Cohesity SaaS Connector(s) and Azure environment.
6. SaaS Connectors are able to resolve the following URLs by name:

- `login.windows.net`
- `management.azure.com`
- `*.blob.core.windows.net`
- To whitelist `*.blob.storage.azure.net`

7. Ensure the following:

- For regular Azure SQL databases (not Managed databases), the following are the connection requirements:
 - Public Access must be enabled in the **Networking** page of the SQL instance
or
 - The SQL instance must have a Private Endpoint Connection created for the VNET assigned to the Saas Connection associated with this Azure source. For information on the steps to create a Private Endpoint Connection, see [Azure Private Link for Azure SQL Database and Azure Synapse Analytics](#).
- For databases in Managed SQL instances, the following are the connection requirements:
 - The SQL Managed instance must be configured on the same VNET assigned to the Saas Connection associated with this Azure source
or
 - A two-way VNET Peering must be created between the VNET of the Saas Connection and the VNET of the SQL Managed instance with the following option checked for both local and remote peering directions: "Allow XXX to access YYY".
For information on the steps to create a Peering, see [Create, change, or delete a virtual network peering](#).
- The server name of the SQL Server (or Managed instance) must be accessible over port **1433**. You can use the **nping** utility to confirm reachability. The server name will be displayed on the Overview page, for example, *sql-managed-instance.c17f785003dd.database.windows.net* OR *cohesitysqlserver.database.windows.net*
For more information, see [Connect to Azure SQL Database Managed Instance with Virtual Network peering](#) and [Connect your application to Azure SQL Managed Instance](#).

Required Permissions

Resource Provider	Operation Name
Microsoft.ManagedIdentity	Microsoft.ManagedIdentity/userAssignedIdentities/assign/action
Microsoft.Resources	Microsoft.Resources/subscriptions/resourceGroups/read Microsoft.Resources/subscriptions/resourceGroups/write
Microsoft.Storage	Microsoft.Storage/storageAccounts/blobServices/containers/read Microsoft.Storage/storageAccounts/blobServices/containers/write Microsoft.Storage/storageAccounts/listkeys/action Microsoft.Storage/storageAccounts/read Microsoft.Storage/storageAccounts/write
Microsoft.Network	Microsoft.Network/dnszones/A/read Microsoft.Network/dnszones/A/write Microsoft.Network/dnszones/A/delete Microsoft.Network/networkInterfaces/write Microsoft.Network/networkInterfaces/read Microsoft.Network/networkInterfaces/join/action Microsoft.Network/networkInterfaces/delete Microsoft.Network/networkInterfaces/ipconfigurations/read Microsoft.Network/networkSecurityGroups/read Microsoft.Network/networkSecurityGroups/join/action Microsoft.Network/networkSecurityGroups/securityRules/read Microsoft.Network/privateEndpoints/read Microsoft.Network/privateEndpoints/write Microsoft.Network/virtualNetworks/read Microsoft.Network/privateEndpoints/delete Microsoft.Network/virtualNetworks/subnets/read Microsoft.Network/virtualNetworks/subnets/join/action Microsoft.Network/virtualNetworks/subnets/joinViaServiceEndpoint/action

Resource Provider	Operation Name
Microsoft.Compute	<p>Microsoft.Compute/disks/beginGetAccess/action</p> <p>Microsoft.Compute/disks/delete</p> <p>Microsoft.Compute/disks/endGetAccess/action</p> <p>Microsoft.Compute/disks/read</p> <p>Microsoft.Compute/disks/write</p> <p>Microsoft.Compute/virtualMachines/start/action</p> <p>Microsoft.Compute/virtualMachines/powerOff/action</p> <p>Microsoft.Compute/virtualMachines/write</p> <p>Microsoft.Compute/virtualMachines/read</p> <p>Microsoft.Compute/virtualMachines/delete</p> <p>Microsoft.Compute/virtualMachines/runCommand/action</p> <p>Microsoft.Compute/virtualMachines/deallocate/action</p> <p>Microsoft.Compute/snapshots/write</p> <p>Microsoft.Compute/snapshots/read</p> <p>Microsoft.Compute/snapshots/beginGetAccess/action</p> <p>Microsoft.Compute/snapshots/endGetAccess/action</p> <p>Microsoft.Compute/snapshots/delete</p> <p>Microsoft.Compute/diskAccesses/write</p> <p>Microsoft.Compute/diskAccesses/read</p> <p>Microsoft.Compute/diskAccesses/delete</p> <p>Microsoft.Compute/diskAccesses/privateEndpointConnectionsApproval/action</p>
Microsoft.KeyVault	<p>Microsoft.KeyVault/vaults/deploy/action</p>

Resource Provider	Operation Name
Microsoft.Sql	Microsoft.Sql/servers/read
	Microsoft.Sql/servers/databases/read
	Microsoft.Sql/servers/databases/write
	Microsoft.Sql/servers/databases/delete
	Microsoft.Sql/servers/databases/usages/read
	Microsoft.Sql/managedInstances/read
	Microsoft.Sql/managedInstances/databases/read
	Microsoft.Sql/managedInstances/databases/write
	Microsoft.Sql/managedInstances/databases/delete

Required Roles

Assign the following roles to the application:

- SQL Managed Instance Contributor - to discover/backup/restore SQL Managed Instance databases.
- SQL DB Contributor - to discover/backup/restore Logical SQL Server databases.

For more information on the permissions granted by these roles, see [Azure built-in roles for Databases](#).

Firewall Ports

For firewall rules, see [Azure SQL Database and Azure Synapse IP Firewall Rules](#).

Considerations

- Transaction Log and Differential Backups are not supported for Azure SQL databases. This is due to a Microsoft limitation.
- Azure SQL Pool/Synapse is not supported. This is due to a Microsoft limitation.
- Databases with external elements are not supported. This is due to a Microsoft limitation.

Register Azure SQL Source

To start protecting your Microsoft Azure services, check the [Azure requirements](#) and then register Azure as a data source in Cohesity DataProtect as a Service.

1. In **DataProtect as a Service**, navigate to **Sources** and select **+ Register Source > Azure**.

Select Source

SaaS Sources

Microsoft 365

AWS

Azure

On-Premise Sources

Virtual Machine

Physical

Microsoft SQL

NAS

Oracle

To register an Azure cloud source, you need:

- Azure domain name.
- Azure cloud region you want to store the backup data.
- Azure services you want to protect.
- Azure Application ID.
- Azure domain name.
- Azure Application Secret Key.
- Azure subscription ID.

[Learn more](#)

Cancel **Start Registration**

The **Register Azure Source** form appears.

- In the **Domain Name** field, enter the unique tenant ID assigned by Azure and then select the **Destination Cloud region**. For information on getting the tenant ID, see [Get tenant ID](#).
- Select the Azure services as **Azure SQL**.

Register Azure Source

1 Source Details — 2 — 3

Source Details

Domain Name
xxxxxxx
This is typically the Azure tenant id.

Destination Cloud Region
Azure Central US

Azure Services

Choose Azure Platform to discover

☐ **vm** Azure VM
☒ **SQL** Azure SQL

Cancel **Continue**

- Click **Continue**.

5. In the **Azure App ID** field, enter the application ID assigned by Azure during the service principal creation process.
6. In the **Azure App Secret Key** field, enter the application secret key for setting up the authentication for the service principal.

Register Azure Source ✓ — 2 Azure Apps — 3

Info Configuring Azure apps manually requires you to perform certain steps in Microsoft Azure as well. [Show manual registration to-do list.](#)

Azure Apps

Azure App ID: xxxxxx

Azure App Secret Key:

Back Cancel Continue

7. Click **Continue**.
8. In the **Azure Subscription ID** field, enter the subscription IDs of the subscriptions where the Azure SQL you want to protect belongs.

Register Azure Source ✓ — ✓ — 3 Azure Subscriptions

Info Note that the listed subscriptions are added only after an administrator has manually linked them in Microsoft Azure. [Show to-do list.](#)

Add subscriptions for the domain xxxxxx

Azure Subscription ID.

1. XXXXXX +

Back Cancel Register

9. Click **Register**.

Next > Once you register the Azure source, you must set up a [SaaS Connection](#) for each region under each Azure subscription in your Azure source.

Protect Azure SQL Databases

Before you protect the Azure SQL databases in your Azure source, you must set up a [SaaS Connection](#) for each region under each Azure subscription in your Azure source. A SaaS Connection consists of one or more SaaS Connectors, which are VMs that act as data movers between your data sources and the Cohesity DataProtect as a Service.

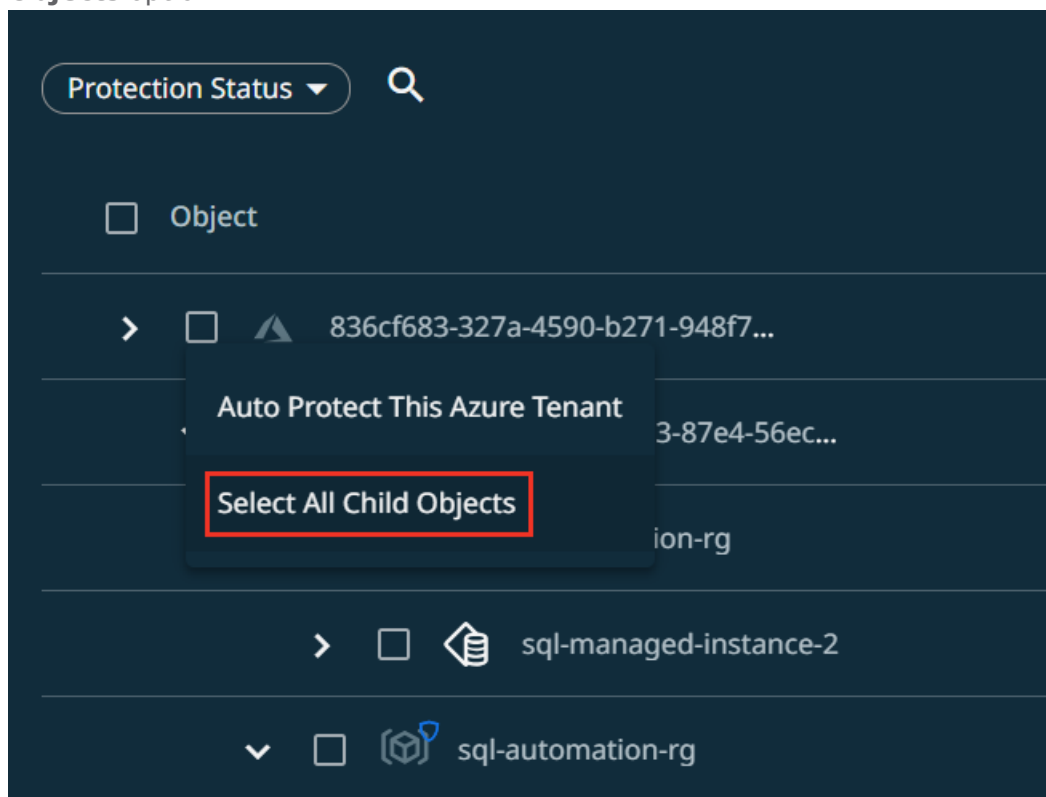
Once you set up a SaaS Connection, you are ready to protect the Azure SQL databases in the Azure source.

Set Azure SQL Server Credentials

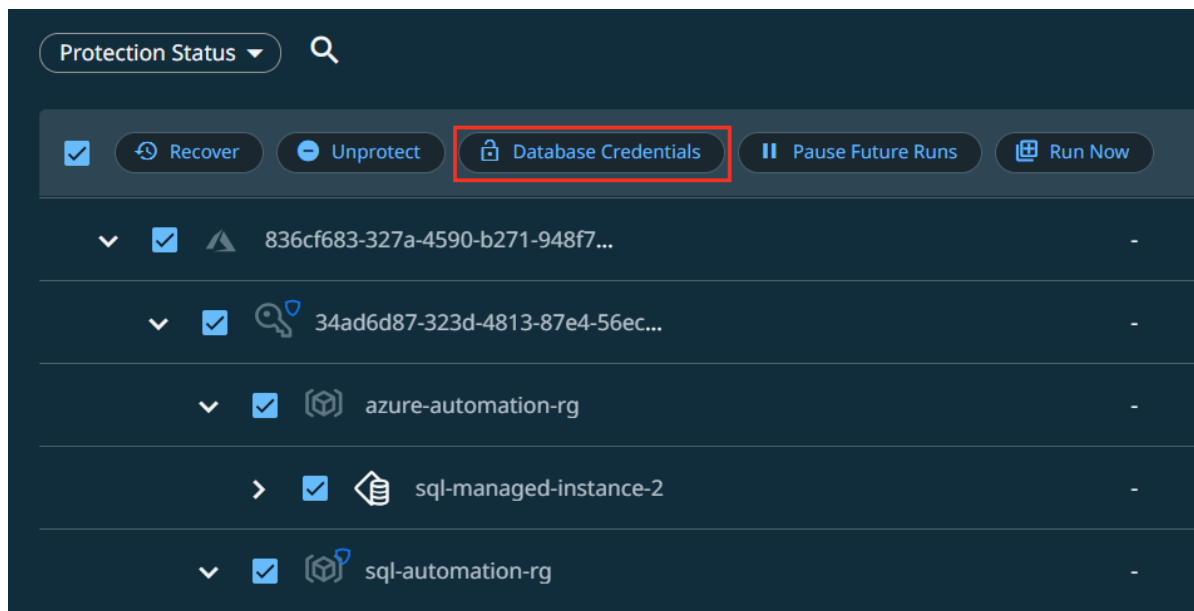
For Azure SQL, you can set credentials at the Azure SQL Server level. This information is used to connect to the database during backup and recovery.

To set the database credentials:

1. In **DataProtect as a Service**, navigate to **Sources**, find the registered Azure source, and click on it.
2. Click the **Azure SQL** tab.
3. Click the checkbox next to the Azure SQL Server and click the **Select All Child Objects** option.



4. Click the **Database Credentials** button.



5. You can select one of the following options:

1. **Username and Password** - Provide the database admin username and password. These credentials are required for leaf-level databases. The account must have access to all databases within the SQL server.
2. **Managed Identity** - Uses the authentication automatically managed by Azure.

Ensure the following:

- Managed Identity is provided during SaaS connector setup. If not, the managed identity must be added to the relevant SaaS Connectors before performing protection using the managed identity method.
- Managed Identity is configured on the SQL Server Microsoft Entra admin field, either directly or indirectly as part of a User Group (this can be helpful if other admins are also required).


Database Credentials

Administrative database credentials are needed to discover leaf level databases and to take RMAN backups

Selected Instances: 7

☒ Username and Password ☐ Managed Identity

Database Admin Username: azure_sql

Enter Password: 

Cancel Save





Note:

- When configuring the SaaS Connector that will serve the SQL Server, it has to be assigned a User Assigned Managed Identity. For more details, see [Manage user-assigned managed identities](#).
- If the SaaS connection has multiple connectors, all of them must be assigned the same Managed Identity.
- The same Managed Identity must also be set as an Entra Admin for the SQL Server. For more details, see [Microsoft Entra admin with a server in SQL Database](#).
- For backup to be transactionally consistent, Cohesity backup creates a copy of the database and uses the copied database to perform the backup. For managed instances, instead of a copy, a new database is created by Azure recovery (native). The new database is created under the same managed instance so that the backup uses the storage and compute resources from the production-managed instance.

6. Click **Save**.

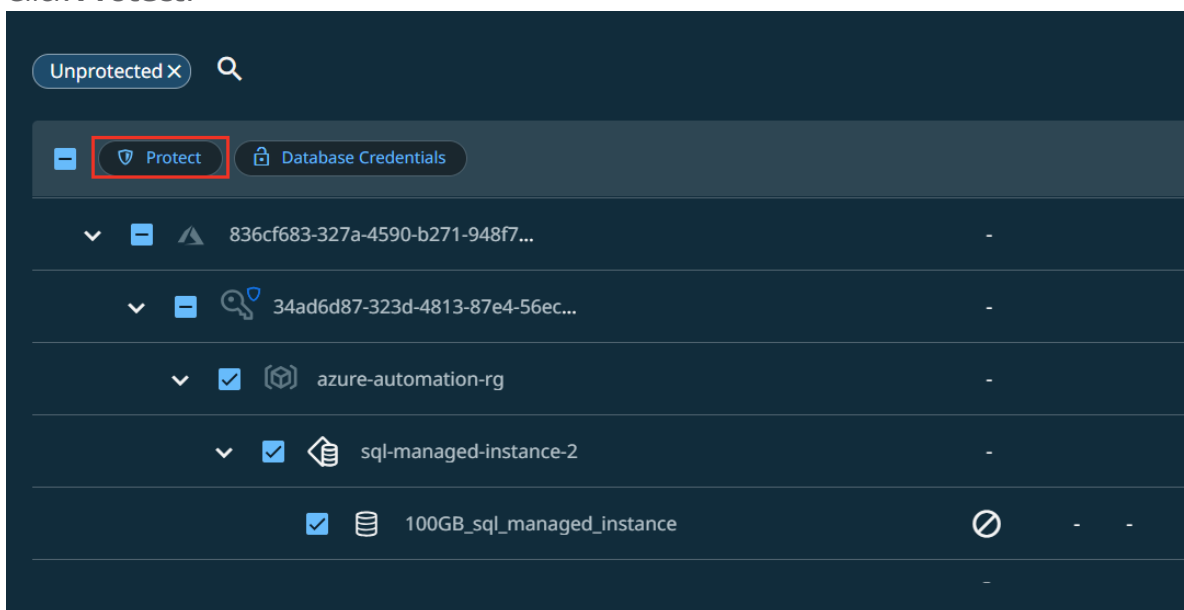
Add Protection to Your Azure SQL Databases

To protect your Azure SQL Server databases:

1. In **DataProtect as a Service**, under **Sources**, find the registered Azure source and click on it.
2. Click the **Azure SQL** tab.
3. Use the checkboxes to select the objects for protection. You can select the tenants , subscriptions , resource groups , or databases  under them.
 - To protect all objects in the source, click the checkbox next to **Object**.
 - To protect all the child objects under the Azure SQL Server, click the checkbox next to the Azure SQL Server and click the **Select All Child Objects** option.
 - To auto-protect the Azure SQL Server, click the checkbox next to the Azure SQL Server object and select the **Auto-Protect This Azure SQL Server** option.

Note: Object-level exclusion is not supported during Azure SQL protection.

4. Click **Protect**.



5. Choose a policy to specify backup frequency and retention. If you do not have a policy, you can easily [create one](#).

Note: Periodic full is not recommended for Azure SQL since all incremental backups are converted to full backups. If periodic full is part of the policy, two full backups will run sequentially when new databases are onboarded.

6. To change or configure any of the additional settings, select **More Options** and perform the below steps, or else, click **Protect**.

7. Under **Settings**, edit the **Start Time** if necessary.
8. In the **SLA** field, define how long the administrator expects a protection run to take. Enter:
 1. **Full**. The number of minutes you expect a full protection run, which captures all the blocks in an object, to take.
 2. **Incremental**. The number of minutes you expect an incremental protection run, which captures only the changed blocks in an object, to take.
9. If you need to change any of the additional settings, click the down arrow icon next to **Additional Settings** and click **Edit**.
10. Click **Protect**.

You can monitor the status of the backup on the **Activity** page. Also, the **Activity** tab of a specific Azure SQL shows the history of all protection runs, including the one in progress.

Cohesity DataProtect as a Service starts backing up the databases you selected.

Additional Settings

Settings	Description
Copy Database	<p>Enable the Copy Database toggle button to make a copy of the database.</p> <p>Backing up a database in use may be transactionally inconsistent. To ensure transactional consistency, you must stop all transactions on the database before backup, or make a copy. Database recovery may not work if the backup is not transactionally consistent.</p> <p>For more details, see Copy a transactionally consistent copy of a database in Azure SQL Database and Microsoft backup limitations.</p>
Copy Database SKU	<p>Specify the Stock Keeping Unit (SKU) details for the copy database. Select the SKU Name, SKU Tier Type, and Capacity. To get a list of valid SKUs for your subscription and region, use the following API: Location Capabilities - List By Location.</p> <p>Backup may be slower when the default Copy Database SKU is used. Cohesity recommends changing to hyperscale SKU with 4 cores for faster backup; this may incur additional charges.</p>
Temporary Database Disk Size	<p>Provide the disk size (in GB) for the temporary disk used during backup. This field is mandatory for databases from Managed Instances (backups may fail if not set) and not required for databases in unmanaged SQL instances.</p>
Pause Future Runs	<p>Toggle on this option to stop protection runs from executing. Once you enable this option, no protection runs will be scheduled.</p>
End Date	<p>If you need to end protection on a specific date, enable this option to select the date.</p>
Quiet Times	<p>Available only if the selected policy has at least one quiet time period. Toggle it ON to specify that all currently executing protection runs should abort if a quiet time period specified for the protection starts. By default this toggle is OFF, indicating that after a protection run starts, it continues to execute even when a quiet time period specified for this protection run starts. However, a new protection run will not start during a quiet time period.</p>

Protect Azure SQL Databases in Bulk

To protect Azure SQL databases in bulk:

1. Click the **Global Search** box at the top or type an asterisk (*).
2. In the **Filter by** section, select **Status** as *Unprotected*, **Type** as *Azure*, and under **Azure**, select *Azure SQL Database* and click **Protect**.
3. You can also select the **Status** as *Protected* and **Pause Future Runs**, **Edit Protection**, **Cancel Run**, and **Unprotect** the protected databases.

Next > When the first protection run completes, you will be ready to [recover your protected databases](#) when and if you need to.

Recover Azure SQL Databases

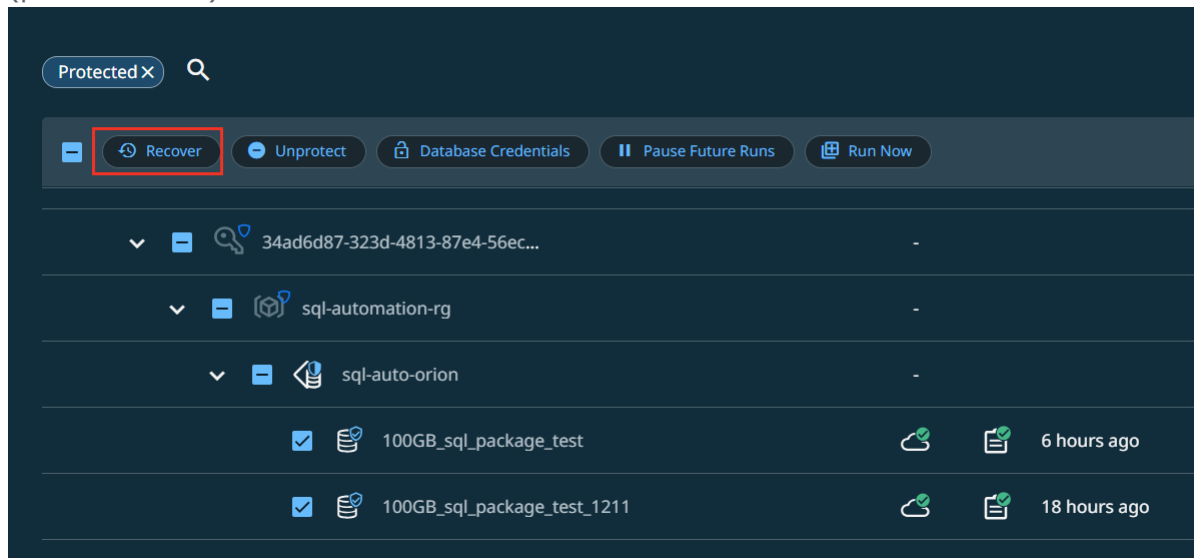
After you [protect your Azure SQL databases](#), you can recover them from Cohesity DataProtect as a Service to their original or a new location.

To recover protected Azure SQL Server databases:

1. In **DataProtect as a Service**, under **Sources**, click the required Azure SQL source and click the **Azure SQL** tab.
2. Select the **Protection Status as Protected**.
3. Use the filters, search box, and views to locate and select the Azure SQL databases you need.

You can also use **Global Search** to locate, filter, and select the objects you need. Click the Global Search box at the top or type **slash (/)** anywhere to start your search.

4. Click **Recover** at the top to open the **New Recovery** form with the **Latest** snapshot (protection run).



If you need to recover from an earlier snapshot, click the **Edit** icon to open the **Recovery Point** page. Click **List** to view the available recovery points by timestamp and click **Select Recovery Point**.

5. Under **Recover To**, select either of the following:
 1. **Original PaaS Instance** - To recover the Azure SQL database to the Azure SQL Server instance or Managed Instance from which the database was backed up.
 2. **Alternate PaaS Instance**. Select a target Azure SQL Server instance or Managed Instance to which the Azure SQL database will be recovered.
6. Select your **Recovery Options**:

1. **Rename.** Add a Prefix and/or Suffix to the full name of the recovered Azure SQL database.
2. **Task Name.** Change the default name of the recovery task.

Azure SQL

2 Azure SQL | Latest Snapshot | Azure Central US Location

836cf683-327a-4590-b271-948f7cf9f4a7
Subscription

Recover To

☒ Original PaaS Instance ☐ Alternate PaaS Instance ⓘ

Recovery Options

Rename	Original Name	New Name
	100GB_sql_package_test	100GB_sql_package_test_rec
	100GB_sql_package_test_1211	100GB_sql_package_test_1211_rec

Task Name: Recover_Dec_14_2023_4_39_PM

Recover **Cancel**

7. Click **Recover**.

You can monitor the status of the recovery on the **Activity** page.

Cohesity DataProtect as a Service starts recovering the selected Azure SQL databases.

Manage Azure SQL Source

Edit Azure Source

You can edit the registered Azure Source to add or remove the Azure services, and subscriptions protected by the Cohesity DataProtect as a Service as a Service from your Azure source.

To edit an Azure source:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the **Actions** menu next to the Azure sources and select **Edit**.
3. In the **Edit Azure Source** form, select or unselect the Azure services you need and click **Continue** to update Azure Application ID, Azure secret key, or to add or remove

subscriptions.

4. Once you edit the Azure Source, click **Update**.

Unregister Azure Source

If you plan to stop backing up your Azure SQL, you can unregister the Azure source from Cohesity DataProtect as a Service.

Before you unregister an Azure source from Cohesity DataProtect as a Service, you must unprotect all the protected objects in that Azure source.

To unregister the Azure Source

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the **Actions** menu next to the Azure source and select **Unregister**.
3. In the **Unregister Source** dialog, click **Unregister**.

Amazon Web Services

Cohesity DataProtect as a Service provides a simple, fast, and cost-effective backup, recovery, and data management solution for Amazon Web Services.

To protect the Microsoft SQL Servers running in AWS EC2 instances, see [Protect Microsoft SQL Server Running in AWS EC2 Instances](#).

AWS Ports and Account Requirements and Considerations

Before you protect your Amazon services using Cohesity DataProtect as a Service, ensure you have met the prerequisites and reviewed the considerations.

Note: For information on the supported cloud regions where you can back up this source, see [Supported Workloads and Cloud Regions](#).

Check Firewall Ports

Ensure that the ports listed in the Amazon Web Services (AWS) section in the [Firewall Ports for User-Deployed SaaS Connectors](#) topic are open to allow communication between the Cohesity SaaS Connector(s) and AWS environment.

Supported AWS S3 Storage Class

Cohesity supports the data protection of the following S3 storage class:

- Amazon S3 Standard
- Amazon S3 Intelligent-Tiering
- Amazon S3 Standard-IA
- Amazon S3 One Zone-IA

Account Requirements

To register your AWS account, run the CloudFormation Template (CFT) and add permissions to the IAM user.

The tables below list the permissions used by Cohesity in your AWS account. You do not need to add these permissions manually (except the IAM User Permissions to Execute CFT), as they are automatically added when you run the CFT provided by Cohesity during your AWS account registration with the Cohesity DataProtect as a Service and SiteContinuity services.

IAM User Permissions to Execute CFT

To register an AWS account with the Cohesity DataProtect as a Service, you need to run the CloudFormation Template on the AWS console. Ensure the IAM user you use has the following permissions to run the CloudFormation Template and to create and view the stack:

Note: Ensure to add these permissions manually.

- cloudformation:CreateChangeSet
- cloudformation:CreateStack
- cloudformation:CreateUploadBucket
- cloudformation>DeleteStack
- cloudformation:DescribeStackEvents
- cloudformation:DescribeStackResources
- cloudformation:DescribeStacks
- cloudformation:GetTemplate
- cloudformation:GetTemplateSummary
- cloudformation:ListStackResources
- cloudformation:ListStacks
- cloudformation:UpdateStack
- iam:AddRoleToInstanceProfile
- iam:AttachRolePolicy
- iam:CreateInstanceProfile
- iam:CreateRole
- iam>DeleteInstanceProfile
- iam>DeleteRole
- iam>DeleteRolePolicy
- iam:DetachRolePolicy
- iam:GetInstanceProfile
- iam:GetRole
- iam:GetRolePolicy
- iam:PassRole
- iam:PutRolePolicy
- iam:RemoveRoleFromInstanceProfile

- iam:TagRole
- lambda:AddPermission
- lambda:CreateFunction
- lambda>DeleteFunction
- lambda:InvokeFunction
- lambda:RemovePermission
- s3:CreateBucket
- s3:GetObject
- s3:ListBucket
- s3:PutObject
- s3: PutBucketPublicAccessBlock

Permissions for Amazon EC2 Data Protection

Note: You do not need to add these permissions manually, as they are automatically added when you run the CFT.

Resource	Permissions	Reason
ebs	ebs:CompleteSnapshot ebs:GetSnapshotBlock ebs:ListChangedBlocks ebs:ListSnapshotBlocks ebs:PutSnapshotBlock ebs:StartSnapshot	These permissions are required for EBS direct APIs to read & write data from/to EBS snapshots.

Resource	Permissions	Reason
ec2	ec2:AssociateIamInstanceProfile ec2:AttachVolume ec2:CopySnapshot ec2:CreateSnapshot ec2:CreateTags ec2:CreateVolume ec2>DeleteSnapshot ec2>DeleteVolume ec2:DeregisterImage ec2:DescribeAccountAttributes ec2:DescribeAddresses ec2:DescribeAvailabilityZones ec2:DescribeInstanceStatus ec2:DescribeInstanceTypes ec2:DescribeInstances ec2:DescribeKeyPairs ec2:DescribeRegions ec2:DescribeReservedInstances ec2:DescribeReservedInstancesOfferings ec2:DescribeSecurityGroups ec2:DescribeSnapshots ec2:DescribeSubnets ec2:DescribeTags ec2:DescribeVolumeAttribute ec2:DescribeVolumes ec2:DescribeVpcEndpointServiceConfigurations ec2:DescribeVpcs ec2:DetachVolume ec2:ModifyInstanceAttribute	<p>These permissions are required to register the AWS account on Cohesity BaaS with the IAM role which got created by the Cloud Formation template. Once the source is registered on BaaS, describe permissions are needed so Cohesity can identify resources present in the account, which will be used for backups as well as at the time of recovery we use this information to provide a list of options(VPC, subnet, KeyPair, etc) to choose from.</p> <p>For Cohesity snapshots we create SaaS Connector instances for doing backup and recovery of AWS EC2 instances. Cohesity creates snapshots of the EC2 volumes while backing up and storing the different instance attributes and tags. While recovering the AWS EC2 instance, Cohesity creates volumes of original disk size. It also attaches the original tags and corresponding network and security groups as part of the recovery, along with IAM Instance Profile if it exists. Cohesity requires the delete snapshots permissions to delete the expired/old snapshots it creates during the backup. Cohesity requires the delete volume and instance termination permissions to tear down the SaaS Connectors.</p>

Resource	Permissions	Reason
	ec2:RegisterImage ec2:RunInstances ec2:StartInstances ec2:StopInstances ec2:TerminateInstances	
iam	iam:PassRole iam:SimulatePrincipalPolicy iam:GetInstanceProfile iam:AmazonSSMManagedInstanceCore	PassRole permission is needed so that we can attach the created role to SaaS Connectors, as well as the original roles on the recovered EC2 instances. SimulatePrincipalPolicy is needed so we can ensure required actions are allowed on the IAM role we created as part of the Cloud Formation template. GetInstanceProfile is needed to check if the required Instance profile is present at the time of recovery in the target location. AmazonSSMManagedInstanceCore is needed to access the AWS Systems Manager Agent (SSM) on the target VM.
kms*	kms:CreateGrant kms:Decrypt kms:DescribeKey kms:Encrypt kms:GenerateDataKey kms:GenerateDataKeyWithoutPlaintext kms:GetKeyPolicy kms:ListAliases kms:ListKeys kms:ReEncryptFrom kms:ReEncryptTo	KMS permissions are needed to read data of encrypted volumes at the time of backup, as well as write encrypted data to the recovered EBS volumes. Describe permissions are needed so we can list & identifies keys associated with EBS volumes.
ssm	ssm:GetCommandInvocation ssm:SendCommand	SSM permissions are needed at the time of claiming (adding) SaaS Connections to Cohesity BaaS.

*If you want to use a KMS key belonging to a different AWS account, then perform the steps described in the [AWS documentation](#).

Permissions for Amazon RDS Data Protection

Note: You do not need to add these permissions manually, as they are automatically added when you run the CFT.

Resource	Permissions	Reason
ec2	ec2:DescribeAvailabilityZones ec2:DescribeInstances ec2:DescribeKeyPairs ec2:DescribeRegions ec2:DescribeReservedInstancesOfferings ec2:DescribeSecurityGroups ec2:DescribeSubnets ec2:DescribeVolumes ec2:DescribeVpcs	Required for AWS source registration, and discover the resources present in the account, which will be used for backups. Also needed for recovery to provide list of options to choose from.
	ec2:CreateVolume ec2:CreateTags ec2:DescribeTag ec2:DescribeVolumeAttribut ec2:DescribeVolumes ec2:DescribeInstances ec2:AttachVolume ec2>DeleteVolume ec2:DetachVolume ec2:ModifyVolume	Required for attaching and detaching volumes of RDS database to SaaS Connectors during the RDS ingest backup.
iam	iam:SimulatePrincipalPolicy	SimulatePricipalPolicy is needed to ensure that the required actions are allowed on the IAM role we created as part of the Cloud Formation template.

Resource	Permissions	Reason
kms*	kms:CreateGrant kms:DescribeKey kms:ListAliases	KMS permissions are needed to read data of an encrypted database at the time of backup, as well as write encrypted data to the recovered database. Describe permissions are needed so we can list & identify keys associated with database instances.
rds	rds:AddTagsToResource rds:CopyDBClusterSnapshot rds:CopyDBSnapshot rds>CreateDBClusterSnapshot rds>CreateDBInstance rds>CreateDBSnapshot rds>DeleteDBClusterSnapshot rds>DeleteDBSnapshot rds:DescribeDBClusterSnapshots rds:DescribeDBClusters rds:DescribeDBInstances rds:DescribeDBParameterGroups rds:DescribeDBSnapshots rds:DescribeDBSubnetGroups rds:DescribeOptionGroups rds:ModifyDBClusterSnapshotAttribute rds:ModifyDBSnapshotAttribute rds:RestoreDBClusterFromSnapshot rds:RestoreDBClusterToPointInTime rds:RestoreDBInstanceFromDBSnapshot rds:RestoreDBInstanceToPointInTime	<p>These permissions are required to register the AWS account on Cohesity BaaS with the IAM role which got created by the Cloud Formation template. Once the source is registered on BaaS, describe permissions are needed so Cohesity can identify resources present in the account, which will be used for backups as well as at the time of recovery we use this information to provide a list of options to choose from.</p> <p>Cohesity creates snapshots of the RDS & Aurora instances while backing up and storing the different database instance attributes and tags. While recovering the database instance, Cohesity creates DB instance/cluster, it also attaches the original tags. Cohesity requires the delete snapshots permissions to delete the expired/old snapshots it creates during the backup. We need to modify snapshot attributes permission so that we can share the snapshot across accounts if cross-account recovery is attempted.</p> <p>RestoreDBInstanceToPointInTime and RestoreDBClusterToPointInTime is needed to do the point in time recoveries.</p>

*If you want to use a KMS key belonging to a different AWS account, then perform the steps described in the [AWS documentation](#).

Permissions for Amazon S3 Data Protection

Note: You do not need to add these permissions manually, as they are automatically added when you run the CFT.

Resource	Permissions	Reason
S3	s3:GetBucketLocation s3:GetBucketNotification s3:GetBucketOwnershipControls s3:GetBucketTagging s3:GetBucketVersioning s3:GetInventoryConfiguration s3:GetObject s3:GetObjectAcl s3:GetObjectTagging s3:GetObjectVersion s3:GetObjectVersionAcl s3:GetObjectVersionTagging s3:ListAllMyBuckets s3:ListBucket s3:PutBucketNotification s3:PutInventoryConfiguration s3:PutObject s3:PutObjectAcl s3:PutObjectTagging s3:PutObjectVersionAcl	These permissions are required for the backup and recovery of S3 objects.
iam	iam:SimulatePrincipalPolicy	SimulatePrincipalPolicy is needed to ensure that the required actions are allowed on the IAM role we created as part of the Cloud Formation template.

Resource	Permissions	Reason
kms*	kms:CreateGrant kms:DescribeKey kms:ListAliases kms:GenerateDataKey	KMS permissions are needed to read data of an encrypted database at the time of backup, as well as write encrypted data to the recovered database. Describe permissions are needed so we can list & identify keys associated with database instances.
Events	events:DeleteRule events:PutTargets events:RemoveTargets	These permissions are required for capturing the incremental changes on the S3 buckets.
Glue	glue:DeleteJob glue:GetJobRun glue:StartJobRun glue:UpdateJob	These permissions are required for sorting the inventory report. The sorted inventory report is then used to back up the S3 objects to the Cohesity DataProtect as a Service.
SQS	sqs:CreateQueue sqs:TagQueue sqs:DeleteMessage sqs:DeleteQueue sqs:GetQueueUrl sqs:PurgeQueue sqs:ReceiveMessage sqs:SetQueueAttributes	These permissions are required for capturing the incremental changes on the S3 buckets.

Permission for AWS S3 Inventory Report

To write objects to the Amazon S3 bucket, you must add the `s3:PutObject` permission to the S3 bucket policy attached to the AWS S3 bucket where you want to create the inventory report.

The following is an example of an S3 bucket policy that allows **s3.amazonaws.com** to write (Put) objects in the S3 bucket:

```
{
  "Version": "2012-10-17",
  "Id": "S3-Console-Auto-Gen-Policy-1698064515475",
  "Statement": [
    {
      "Sid": "InventoryAndAnalyticsExamplePolicy",
      "Effect": "Allow",
      "Principal": {
        "Service": "s3.amazonaws.com"
      }
    }
  ]
}
```

```
    },
    "Action": [
      "s3:PutObject"
    ],
    "Resource": "arn:aws:s3:::inventory-report-bucket/*",
    "Condition": {
      "StringEquals": {
        "aws:SourceAccount": "<account id>",
        "s3:x-amz-acl": "bucket-owner-full-control"
      }
    }
  },
}
```

In the above example, <account id> is the AWS account ID of the Amazon S3 bucket you want to protect.

Permissions for Cohesity SiteContinuity (Disaster Recovery)

Note: You do not need to add these permissions manually, as they are automatically added when you run the CFT.

Resource	Permissions	Reason
ebs	ebs:CompleteSnapshot ebs:GetSnapshotBlock ebs:ListChangedBlocks ebs:ListSnapshotBlocks ebs:PutSnapshotBlock ebs:StartSnapshot	These permissions are required for EBS direct APIs to read & write data from/to EBS snapshots. Reading EBS data is done during failback preparation, and writing to EBS is done at failover.

Resource	Permissions	Reason
ec2	ec2:AssociateIamInstanceProfile ec2:AttachVolume ec2:CancelExportTask ec2:CancelImportTask ec2:CopySnapshot ec2:CreateImage ec2:CreateInstanceExportTask ec2:CreateSnapshot ec2:CreateTags ec2:CreateVolume ec2>DeleteSnapshot ec2>DeleteTags ec2>DeleteVolume ec2:DeregisterImage ec2:DescribeAccountAttributes ec2:DescribeAddresses ec2:DescribeAvailabilityZones ec2:DescribeExportTasks ec2:DescribeImages ec2:DescribeImportImageTasks ec2:DescribeInstanceAttribute ec2:DescribeInstanceStatus ec2:DescribeInstances ec2:DescribeKeyPairs ec2:DescribeRegions ec2:DescribeReservedInstancesOfferings ec2:DescribeSecurityGroups ec2:DescribeSnapshots ec2:DescribeSubnets ec2:DescribeTags	<p>These permissions are required to register the AWS account on Cohesity Helios with the IAM role created by the Cloud Formation template. Once the source is registered, describe permissions are needed so Cohesity can identify resources present in the account like EC2 instance, VPC, subnet, etc. These describe permissions are also used at the time of failover and fallback.</p> <p>The import/export permissions are required because we use AWS Import/Export as our fallback mechanism if Cohesity Import/Export does not work. Cohesity requires all the instance-related permissions to run instances and terminate them if some error occurs.</p> <p>Delete permissions are required so that Cohesity can delete the temporary resources like volumes or snapshots it has created in the process of failover or fallback so that we do not leave any garbage behind.</p>

Resource	Permissions	Reason
	ec2:DescribeVolumeAttribute ec2:DescribeVolumes ec2:DescribeVpcs ec2:DetachVolume ec2:ImportImage ec2:ModifyInstanceAttribute ec2:ModifyNetworkInterfaceAttribute ec2:ModifySnapshotAttribute ec2:RegisterImage ec2:RunInstances ec2:StartInstances ec2:StopInstances ec2:TerminateInstances	
iam	iam:AddRoleToInstanceProfile iam:AttachRolePolicy iam:CreateInstanceProfile iam:CreateRole iam:GetInstanceProfile iam:GetRole iam:GetRolePolicy iam:PassRole iam:PutRolePolicy iam:SimulatePrincipalPolicy	These IAM permissions are needed because we have to SSM into the converter instance, and for that to work, an instance profile should be attached to the converter instance. So to create that instance profile for the role, these permissions are needed.
kms	kms:ListAliases	KMS permission is needed to list the aliases attached to an EC2 instance at the time of source register.

Resource	Permissions	Reason
s3	s3:CreateBucket s3:DeleteObject s3:GetBucketAcl s3:GetObject s3:HeadObject s3:PutBucketAcl s3:PutBucketPublicAccessBlock	These S3 permissions are needed in case of the vmimport role we use in case of failover.
ssm	ssm:GetCommandInvocation ssm:ListCommandInvocations ssm:SendCommand	SSM permissions are needed at the time of failover, where we launch the SaaS Connector and temporary converter instance for creating EC2 instances.

Create a Lifecycle Rule on Amazon S3

To delete the older inventory reports from the Amazon S3 bucket, you must create a lifecycle rule on the Amazon S3 bucket. You can delete all the inventory reports older than 30 days. For information on creating a lifecycle rule, see [Amazon documentation](#).

Permission for AWS Key Management Service (KMS)

If the S3 bucket you want to protect is encrypted with Server-side encryption with AWS Key Management Service keys (SSE-KMS) or Dual-layer server-side encryption with AWS Key Management Service keys (DSSE-KMS), then for Cohesity to access the S3 bucket, you must perform one of the following actions:

- Add the IAM role created by the Cloud Formation template to the AWS KMS user.
- Add the following permission to the Key policy attached to the AWS KMS:
 - kms:Encrypt
 - kms:Decrypt
 - kms:ReEncrypt*
 - kms:GenerateDataKey*
 - kms:DescribeKey

For example:

```
{
  "Version": "2012-10-17",
```

```

    "Id": "AccessKeyId",
    "Statement": [
      {
        "Sid": "Allow use of the key to cohesity role",
        "Effect": "Allow",
        "Principal": {
          "AWS": "arn:aws:iam::<AWS-ACCOUNT>:role/<ROLE-NAME>"
        },
        "Action": [
          "kms:Encrypt",
          "kms:Decrypt",
          "kms:ReEncrypt*",
          "kms:GenerateDataKey*",
          "kms:DescribeKey"
        ],
        "Resource": "*"
      }
    ]
  }
}

```

Considerations

Considerations for Amazon EC2 Cohesity Snapshots

- When using [Cohesity snapshots](#) to back up & recover EC2 instances within the same AWS region, if your [AWS SaaS Connectors](#) are deployed in a:
 - **Public subnet**, configure the [Internet Gateway](#) and [S3 Gateway VPC endpoint](#).
 - **Private subnet**, configure the [EBS VPC Interface Endpoint](#) and [S3 Gateway VPC endpoints](#).
- When using [Cohesity snapshots](#) to back up & recover EC2 across different AWS regions, if your SaaS Connectors are deployed in a:
 - **Public subnet**, configure the [Internet Gateway](#) and [S3 Gateway VPC endpoint](#).
 - **Private subnet**, configure the [EBS VPC Interface Endpoint](#) and the [S3 Interface VPC endpoints](#).

Note: Cross-region data transfer charges apply if Cohesity snapshots are ingested to or recovered from a different AWS region. Using a public subnet for your SaaS Connectors provides cost efficiency compared to a private subnet.

- To prepare your AWS account for Cohesity SaaS Connector deployment in a Public or Private subnet, see [AWS SaaS Connector Deployment](#).
- Backing up NFS mount points mounted on EC2 instance is not supported.
- Cohesity does not support the backup and recovery of AWS EC2 instances with UEFI Preferred boot mode.

Considerations for Amazon RDS

- AWS Aurora cluster is recovered with at most one reader.
- Cohesity does not support the auto-protect of RDS instances with different database types. Auto-protect of an RDS instance is supported only if the databases on the RDS instance are of the same type.
- Cohesity does not support cross region-cross account recovery of Amazon RDS instance encrypted with default KMS key.
- If you are performing cross region - cross account recovery of an Amazon RDS instance encrypted with a customer-managed key, then ensure:
 - The source account's KMS key alias is available in both the source and destination regions of the destination account.

For example, if A1 is the source account and R1 is the source region of the RDS instance you want to recover, and A2 is the destination account with R2 as the destination region, make sure that the key alias from A1 is available in both R1 and R2 of A2.
 - The source account's KMS key alias can be accessed by the destination account. For this, you must grant the necessary permissions to the KMS key alias.

The option group from the source account is available in both the source and destination regions of the destination account.

For example, if A1 is the source account and R1 is the source region of the RDS instance, and A2 is the destination account with R2 as the destination region, ensure the option group from A1 is available in both R1 and R2 of the A2.

Considerations for Amazon S3

- Cohesity does not support:
 - Browse and recover an object in an Amazon S3 bucket. However, you can recover multiple objects by specifying the object prefix in the recovery task under the **S3 Prefixes to Recover** option.
 - The backup of older versions of the AWS S3 versioned bucket. Only the latest version of the versioned Amazon S3 bucket is backed up.
 - The backup and recovery of Amazon S3 buckets that are not in the same cloud region where your data is backed up (Cohesity-managed SaaS platform).

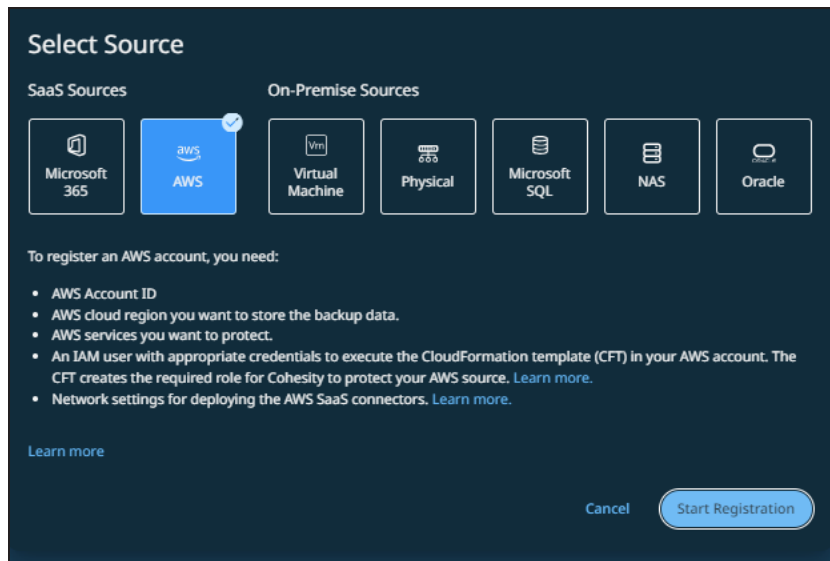
- The Amazon S3 buckets where you want to create the inventory report and the Amazon S3 bucket you want to protect must be in the same region.
- If the SQS is deleted between backups, all the changes between these backups will be skipped in the next incremental backups.
- Cohesity DataProtect as a Service will skip the backup of Amazon S3 objects that are present in the following access tiers of the Amazon S3 Intelligent Tiering during the protection:
 - Archive Access Tier
 - Deep Archive Access Tier
- You do not need to deploy a SaaS connection to protect Amazon S3 buckets.
- Cohesity does not support restoring only metadata. The metadata of Amazon S3 objects will be restored only if the object itself is also restored.
- Cohesity does not remove older objects from the S3 bucket where the inventory report is created. Therefore, you must create a lifecycle rule to remove these objects from the S3 bucket to avoid storage issues.

Register Your AWS Account

To start protecting your AWS account, check the [AWS account requirements](#) and then register the AWS account as a data source in Cohesity DataProtect as a Service. (You can also [unregister an AWS account](#) if and when necessary.)

Register AWS Account

1. In **DataProtect as a Service**, navigate to **Sources > + Register Source**, and then select **AWS**.



2. Click **Start Registration**.

The Register AWS Source form appears.

3. Enter the AWS **Account ID** and select the **Destination cloud region**.

Note: If you decide to create air-gapped [Cohesity snapshots](#) of your Amazon EC2 instances, this is the AWS region where the Cohesity DataProtect as a Service will store them.

4. Enable the option to **Use this account as a backup source in DataProtect** (enabled by default).
5. Select the **AWS Services** you want to register:
 - EC2
 - RDS
 - S3
6. (Applies only for AWS S3 service) Cohesity utilizes the Amazon S3 inventory report to protect the Amazon S3 bucket. Under Inventory Report Location, provide the following information to create the inventory report:
 - **S3 buckets ARN:** Enter the ARN of the Amazon S3 bucket where you want to create the inventory report. The ARN you provide must be of only those Amazon S3 buckets that are in the same AWS account and cloud region as the Amazon S3 you want to protect.
 - **Prefix:** Add a prefix value to the name of the inventory report that will be created.

- The prefix can be of any character and can also include white spaces.
For example, `Report-Source Bucket`.
- The prefix should not begin or end with a forward slash (/) .
- The prefix should not contain consecutive forward slashes.
- You must not upload any files in the prefix of the inventory report.

The inventory report will be created at `<Prefix>/<Path_to_inventoy_report>` when you initiate the protection. Cohesity creates the path to the inventory report on AWS based on the cluster, Amazon S3 bucket, and protection.

Note: SaaS connection is not required for Amazon S3 bucket protection.

7. Disable the **Use this account as a DR target in SiteContinuity** option. Enable this option only if you are planning to use this AWS account as a disaster recovery (DR) target in [Cohesity SiteContinuity](#).
8. Click **Next** to generate a CloudFormation Template, which you will use to complete the AWS source registration.

Register AWS Source

1 Account Information — 2

AWS Account ID
498211347717

Destination Cloud Region
AWS US East (Ohio)

☒ Use this account as a backup source in DataProtect

Enable this option to use Cohesity DataProtect to protect your Amazon EC2 instances and/or RDS databases

AWS Services

☒ EC2 ☒ RDS ☒ S3

Inventory Report Location

S3 Buckets ARN
arn:aws:s3:::auto-dmaas-inventory-re

Prefix
Report-Source Bucket

NOTE: If you plan to create air-gapped backups (Cohesity snapshots) of your Amazon EC2 instances, or to back up SQL and Oracle databases that are running on EC2, you will need to deploy a SaaS Connection for data ingestion.

Once you register this AWS account, click the Actions menu (⋮) and select Setup SaaS Connection. See [Create a SaaS Connection](#).

☐ Use this account as a DR target in SiteContinuity

Enable this option if you plan to use this AWS account as a disaster recovery (DR) target in Cohesity SiteContinuity. This ensures that the Cohesity role will have the necessary permissions to protect, failover, and fallback objects from this account.

Cancel Next

9. Click **Download CloudFormation Template**.
10. [Run the CloudFormation Template](#) in your AWS account to create the IAM roles and policies that the Cohesity DataProtect as a Service needs. On running the CFT, IAM roles and policies are created depending on the AWS services (EC2 instances, S3 buckets, and RDS databases) you selected for registration.

Optionally, you can restrict the granted permissions to a set of resources when creating the CloudFormation stack.

11. Once the roles and policies are created successfully, the **Register AWS Source** form will indicate the account authentication status.
12. Once account authentication is successful, click **Register**. (If authentication fails, contact [Cohesity Support](#).)

If you plan to protect Amazon RDS and EC2 instances using [Cohesity snapshots](#), make sure you deploy one or more [SaaS Connectors](#) in your AWS account by going to **Sources** and editing your AWS source. From there, you can enter the SaaS Connector configuration details.

Unregister AWS Account

If you plan to stop backing up your Amazon EC2 instances, Amazon S3, or Amazon RDS, you can unregister the AWS account from Cohesity DataProtect as a Service.

Note: Before you unregister an AWS account from Cohesity DataProtect as a Service, you must unprotect all the protected objects in that AWS account.

To unregister the AWS account:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the **Actions** menu (⋮) next to the AWS account and select **Unregister**.
3. In the **Unregister Source** dialog, click **Unregister**.

Edit AWS Account

You can edit the registered AWS account to add or remove the AWS services protected by the Cohesity DataProtect as a Service from your AWS account.

To edit an AWS Account:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the **Actions** menu (⋮) next to the AWS account and select **Edit**.
3. In the **Edit AWS Source** form, select or unselect the AWS services (EC2, S3, and/or RDS) you need and click **Update**.

Important: If you add or remove any AWS services, then you must update the [CloudFormation Template](#) and execute it in your AWS account again to update the existing CloudFormation stack.

Next > You are now ready to protect the [Amazon EC2 instances](#), [RDS databases](#), or [Amazon S3](#) in your AWS account!

Amazon EC2 Instances

Cohesity DataProtect as a Service provides a simple, fast, and cost-effective backup, recovery, and data management solution for Amazon EC2 instances in your AWS account.

Protect Your Amazon EC2 Instances

Once you have [registered your AWS account](#), you are ready to protect the EC2 instances in that account.

Note: If you have already registered your AWS account to protect AWS RDS or AWS S3 workloads, then you must [Update the Existing CloudFormation Template](#) to update the Cohesity permissions in your AWS account.

Cohesity's Options for EC2 Backup: AWS or Cohesity Snapshot

Cohesity DataProtect as a Service provides two options for Amazon EC2 backup:

- **AWS snapshot:** Cohesity DataProtect as a Service protects the EC2 instances using the native AWS snapshots and stores them in the same AWS account and region as the source EC2 instances.
- **Cohesity snapshot:** Cohesity DataProtect as a Service protects the EC2 instances by ingesting the backup data to an AWS region supported by the Cohesity DataProtect as a Service. The target AWS region is the region that is selected during [AWS source registration](#). Cohesity snapshots provide an air-gapped backup and granular file & folder level recoveries. With air-gapped backup approach, the backed up data is isolated from any network connectivity, ensuring that your data remains safe. Network connectivity is resumed only during the recovery process, minimizing the risk of ransomware attacks.

When selecting a protection policy below, you can choose to back up your EC2 instances using either approach, or both.

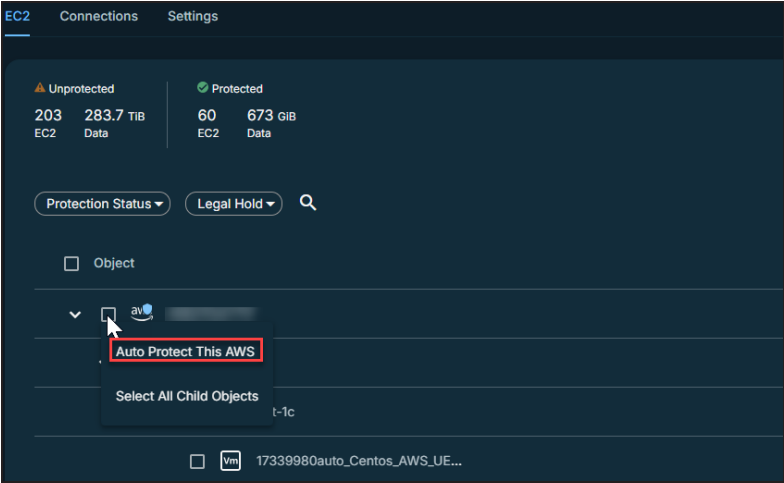
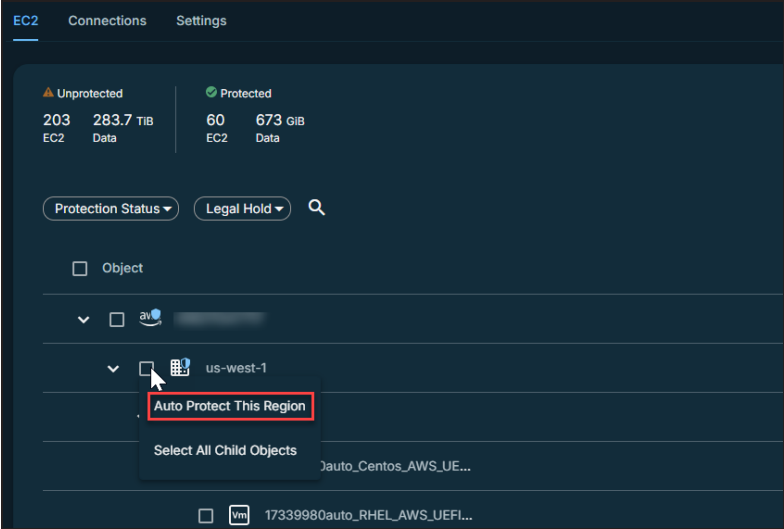
Add Protection to Your Registered Amazon EC2 Instances

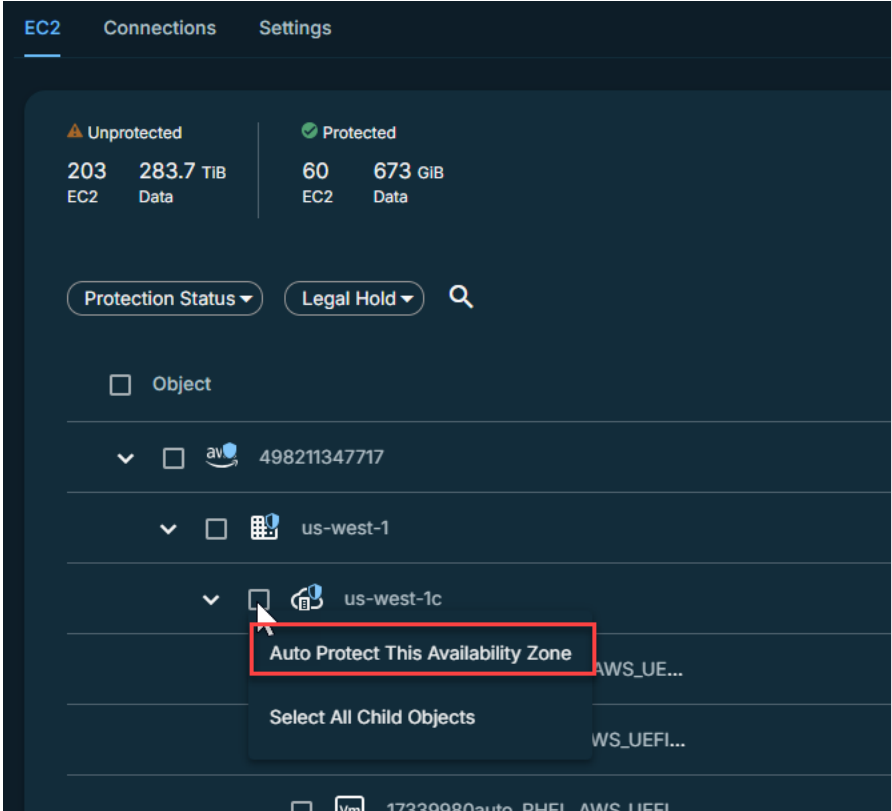
To protect your Amazon EC2 instances:



1. In **DataProtect as a Service**, navigate to **Sources**.
2. Find the registered AWS account and click into it.
3. Click the **EC2** tab.
4. Use the checkboxes to select the EC2 instances for protection.

Optionally, you can configure auto-protect at the AWS account, region, or zone level. When this option is enabled at a particular level, all the EC2 instances that are added to that level in the future are automatically protected from the next protection run. Additionally, you can also perform tag-based auto-protection of EC2 instances.

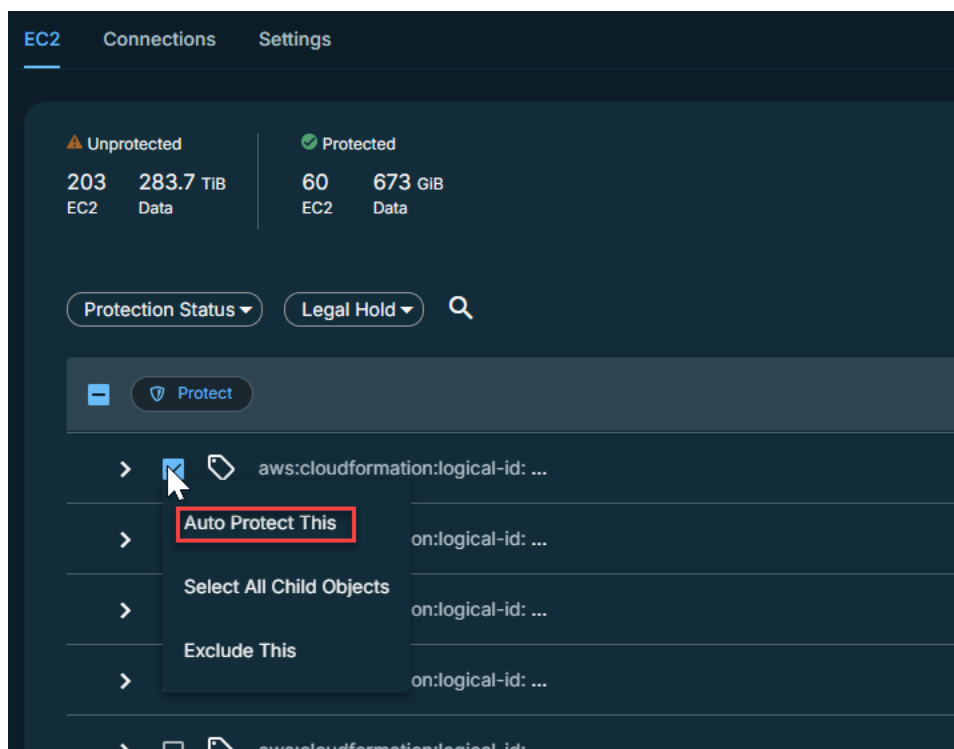
- To auto-protect the EC2 instances based on the hierarchy level, click the **Hierarchy View** icon located at the right corner of the page, and perform one of the following steps listed in the table below:

To auto-protect the AWS EC2 instances..	Action
<p>At the account level</p>	<p>Select the checkbox of the AWS account, and then select Auto Protect This AWS.</p>  <p>The screenshot shows the AWS console interface for EC2. At the top, there are tabs for 'EC2', 'Connections', and 'Settings'. Below these, there are two summary boxes: 'Unprotected' (203 EC2, 283.7 TiB Data) and 'Protected' (60 EC2, 673 GiB Data). Below the summary boxes are filters for 'Protection Status' and 'Legal Hold'. A table lists various objects. The first object is an AWS account, represented by the AWS logo. A red box highlights the 'Auto Protect This AWS' option in the context menu that appears when the account is selected. Below this option is the text 'Select All Child Objects'. The second object in the table is a VM instance named '17339980auto_Centos_AWS_UE...'.</p>
<p>At the region level</p>	<p>Select the checkbox of the subscription, and then select Auto Protect This Region.</p>  <p>The screenshot shows the AWS console interface for EC2, similar to the one above. In this view, the 'us-west-1' region is selected. A red box highlights the 'Auto Protect This Region' option in the context menu that appears when the region is selected. Below this option is the text 'Select All Child Objects'. The table lists various objects, including a VM instance named '17339980auto_RHEL_AWS_UEFI...'.</p>

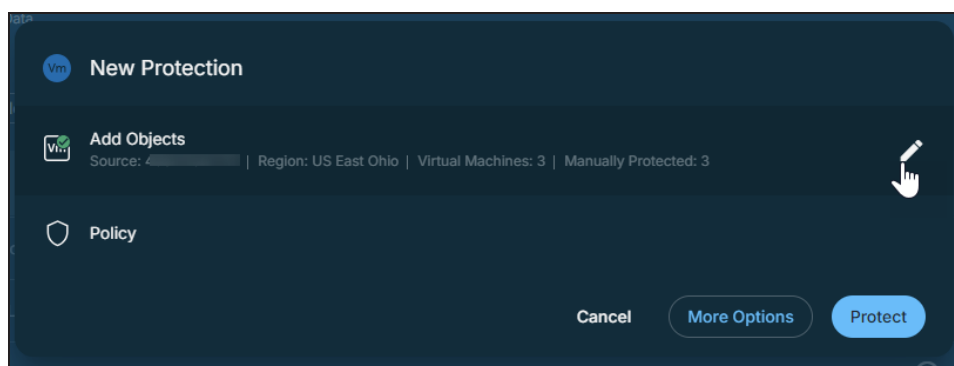
To auto-protect the AWS EC2 instances..	Action
At the zone level	<div>Select the checkbox of the resource group, and then select Auto Protect This Availability Zone.</div> <div></div>

You can exclude individual EC2 instance from auto-protection by clicking the auto-protect icon () next to the instance. EC2 instances excluded from auto-protection are displayed with the exclude () icon.

- To auto-protect the EC2 instances based on tag, click the **Tag** icon at the right corner of the page. Select the checkbox of a tag and then select **Auto Protect This** to auto-protect the EC2 instance with this tag.

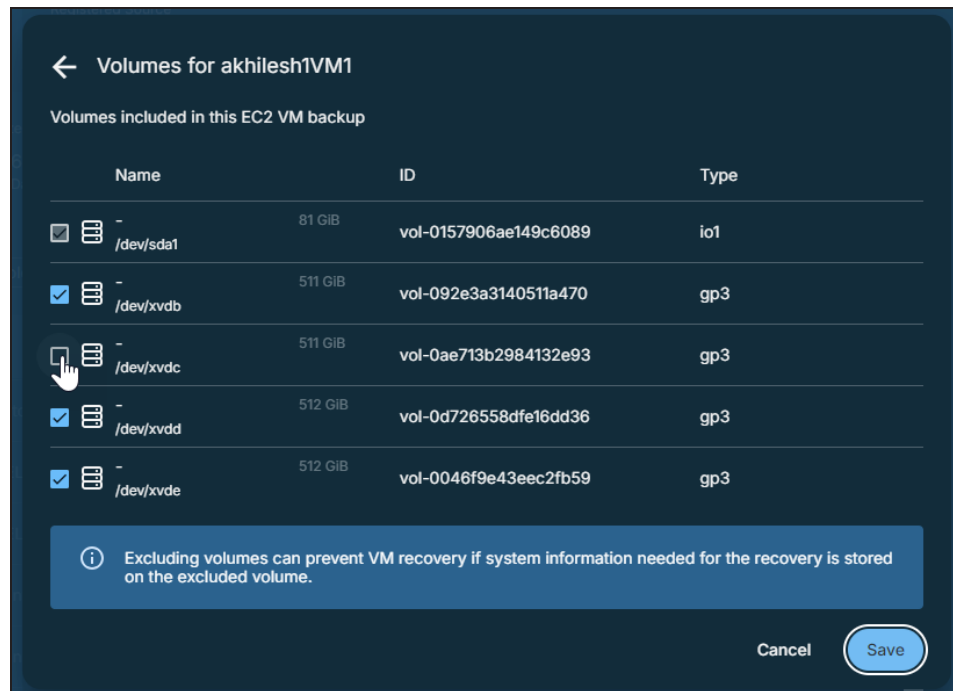


5. Click the **Protect** icon above the checkboxes.
6. To exclude any disks of the EC2 instance from protection:
 - a. Click the edit icon displayed next to **Add Objects**.

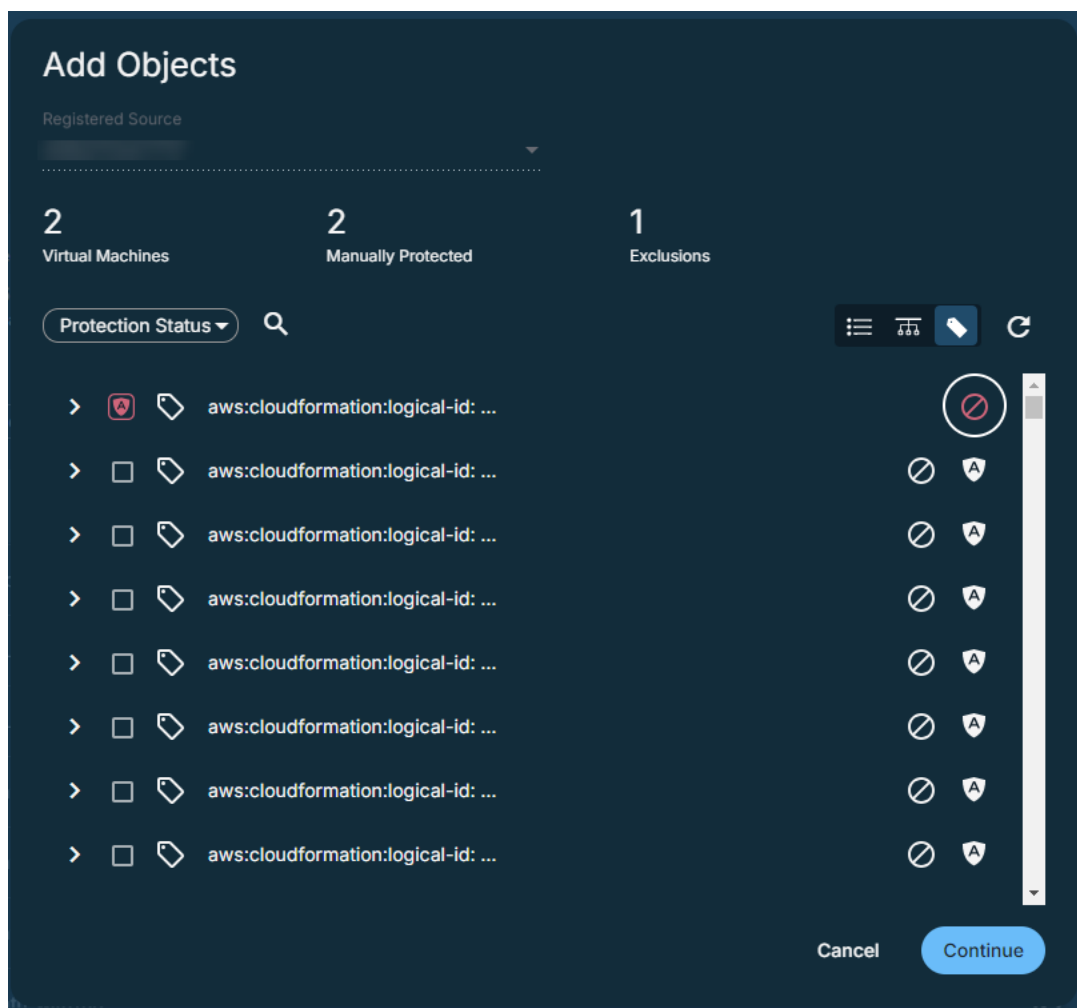


- b. Click the **All Volumes** edit icon displayed next to the EC2 instance of which you want to exclude the disk.
 - c. Unselect the disks that you want to exclude from protection.

Note: You cannot exclude the root disk from protection.



- d. Click **Save > Continue**.
7. To exclude auto-protection of EC2 instances based on Tags:
 - a. Click the edit icon displayed next to **Add Objects**.
 - b. Click the **Tags** icon located at the right corner of the page.
The tags associated with the VMs are displayed.
 - c. Click the exclude icon next to the tag to exclude the auto-protection of VMs associated with that tag.




If an EC2 instance has multiple tags with Auto Protect and Exclude applied, exclusion takes precedence. These excluded EC2 instances are also excluded even if a parent object is auto protected.

- d. Click **Continue**.
8. In the **New Protection** dialog, select a **Policy** from the following [snapshot options](#):
 - **Policy (AWS snapshot)**
 - **Policy (Cohesity snapshot)**

You can create AWS snapshots, Cohesity snapshots, or both. If you choose to create both snapshot types, you can use either the same policy or different policies to specify the backup frequency and retention.

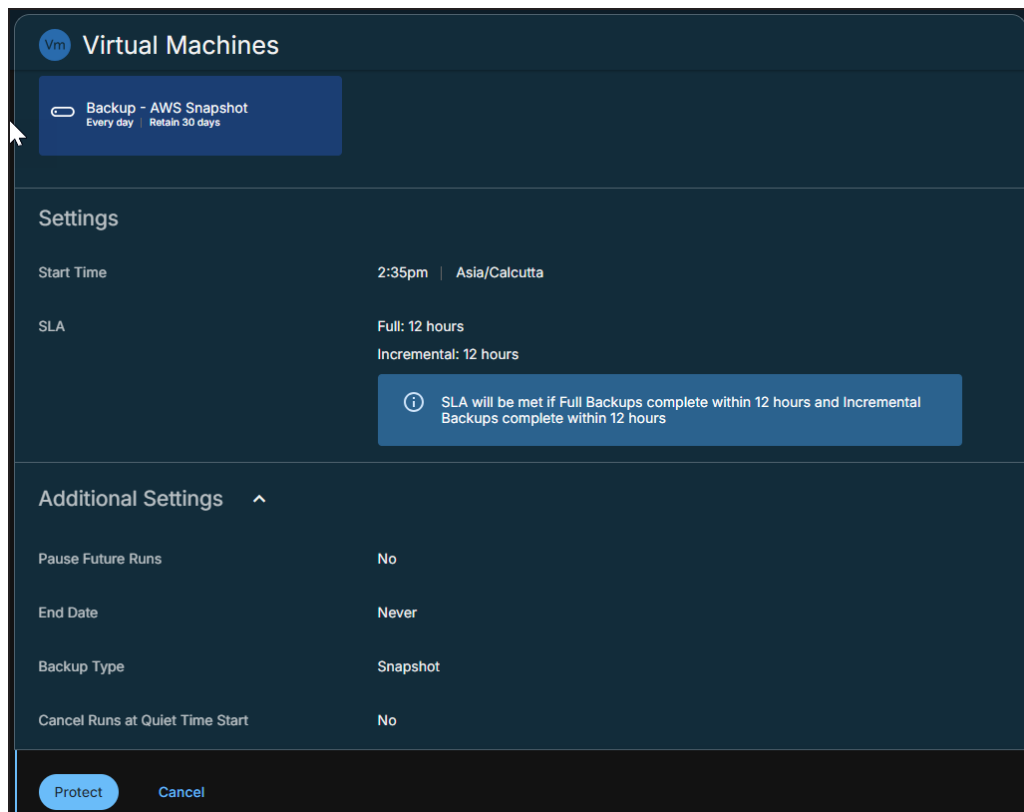
If the existing policies do not meet your needs, you can [create a new policy](#) with the backup frequency and retention settings as desired.

Note: If you have selected [Policy \(Cohesity snapshot\)](#), ensure that an [AWS SaaS Connection](#) is deployed for all the AWS regions where you have instances to protect. If a region in your AWS account does not have a SaaS Connection deployed, protecting the Amazon EC2 instances in that region will fail.

To view the SaaS Connections that are already configured, click the **Actions** menu () next to the registered AWS source and select **Setup SaaS Connection**.

9. If you wish to change or configure any of the additional settings, select **More Options** and perform the below steps or else, click **Protect**.
10. Under **Settings**, edit the **Start Time** if necessary.
11. In the **SLA** field, define how long the administrator expects a protection run to take. Enter:
 - **Full.** The number of minutes you expect a full protection run, which captures all the blocks in an object, to take.
 - **Incremental.** The number of minutes you expect an incremental protection run, which captures only the changed blocks in an object, to take.

12. If you need to change any of the additional settings, click the down arrow icon next to [Additional Settings](#) and click **Edit**.
13. Click **Protect**.



Cohesity DataProtect as a Service starts backing up the Amazon EC2 instances you selected. You can monitor the status of the backup in the **Activity** page.

Also, the **Activity** tab of a specific Amazon EC2 instance shows the history of all protection runs, including the one in progress.

If you have selected both **AWS snapshot** and **Cohesity snapshot** policies, then the **Activity** page will display two protection runs for the objects that are being backed up:

- **Backup**. The protection run created for Cohesity snapshot-based protection.
- **Backup (AWS Snapshot)**. The protection run created for AWS snapshot-based protection.

To learn about managing the existing protection, see [Manage Existing Protection](#).

Additional Settings

Advance Settings	Description
End Date	If you need to end protection on a specific date, enable this to select the date.
Backup Type	Available only if you have selected AWS snapshot policy. Enable Create AMI and specify how often AMI should be created. For example, for the protection, you have configured an AWS snapshot policy with backup frequency set as daily . Now if you specify to create AMI for Every 5 runs , then in a month, AMI will be created for 6 protection runs.
Quiet Times	Available only if the selected policy has at least one quiet time period. Toggle it ON to specify that all currently executing protection runs should abort if a quiet time period specified for the Protection Group starts. By default this toggle is OFF, which means after a protection run starts, it continues to execute even when a quiet time period specified for this protection run starts. However, a new protection run will not start during a quiet time period.

Next > When the first protection run completes, you will be ready to [recover your protected Amazon EC2 instances](#) if and when you need to.

Manage Existing Protection


Edit protection settings, change the policy, and start, stop, & pause protection.

Once you have [applied protection](#) to the objects in your sources, Cohesity DataProtect as a Service makes it easy to make changes to that protection quickly. You can:

- Edit additional settings.
- Apply a different policy.
- Start an on-demand protection run, pause and resume it, or even remove protection.

Edit Protection Settings

To edit protection settings:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click into the **Source** name.
3. Select **Show All > Protected** and use the other filters, search box, and views at the top to narrow your search.
4. Click the **Actions** menu () next to the object and select **Edit Protection** to open the protection settings for that object.

Apply a New Protection Policy

To change the **Policy**, click the drop-down and select a different policy. To help you choose, each policy in the list shows the **Backup** frequency and the **Retain** period for each backup.

If you don't have a policy that meets your needs, scroll to the bottom of the list and click **Create Policy** to [create your own policy](#).

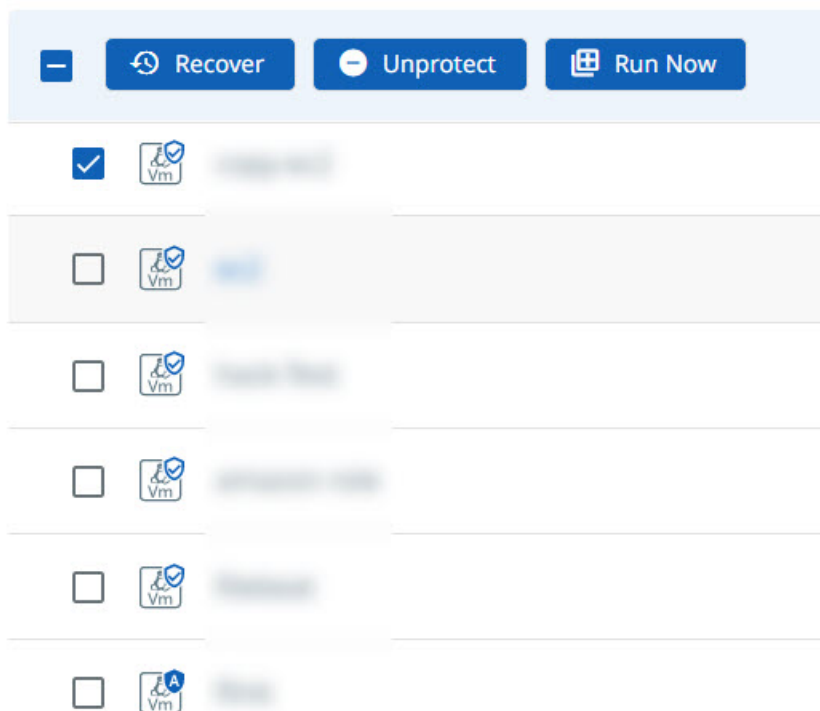
Edit Additional Protection Settings

Under **Settings**, you can change the protection **Start Time** (and select the **Time Zone**).

Click the drop-down next to **Additional Settings** to change more options. See [Additional Protection Settings](#) for details.

Start, Stop, or Remove Protection

When you select protected objects in one of your sources, Cohesity DataProtect as a Service presents buttons for the actions that are possible for those objects.



With the protected objects selected, you can click:

- **Recover** to recover the object or file.
- **Unprotect** to remove protection from the object.

Tip: If a protected object is deleted from the source, you can search the object using Global Search and unprotect it.

- **Run Now** to start an on-demand protection run immediately.

Additional Settings

Advance Settings	Description
End Date	If you need to end protection on a specific date, enable this to select the date.
Quiet Times	Available only if the selected policy has at least one quiet time period. Toggle it ON to specify that all currently executing protection runs should abort if a quiet time period specified for the Protection Group starts. By default this toggle is OFF, which means after a protection run starts, it continues to execute even when a quiet time period specified for this protection run starts. However, a new protection run will not start during a quiet time period.

Recover Your Amazon EC2 Instances

After you [protect your Amazon EC2 instances](#), you can recover them to their [original location](#) or a [new location](#) using Cohesity DataProtect as a Service.

We recommend that you also review the [Amazon EC2 Recovery Support Matrix](#) and [Important Considerations](#) at the end of this article.

Recover EC2s to Original Location

To recover your protected Amazon EC2 instances to their original location:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the **Source** name.
3. Select **Show All > Protected**.
4. Use the filters, search box, and views to locate and select the EC2 instances you want to recover.

Tip: You can also use Global Search to locate, filter, and select the objects you need. Click the **Global Search** box at the top or type **slash (/)** anywhere to start your search.

5. Click the **Recover** icon at the top to open the **New Recovery** form. By default, the **Latest** snapshot is pre-selected for recovery. If you need to recover from an earlier snapshot, click the **Edit** (pencil) icon to choose the desired snapshot. The icon(s) displayed under **Location** indicates the snapshot type(s) available (**AWS snapshot** and/or **Cohesity snapshot**) for recovery. Choose a snapshot type and click its icon to proceed with the recovery task:
 - Click **Select Recovery Point**.
 - Click **Next: Recover Options** to return to the form.

Note: To recover from a [Cohesity snapshot](#), you need an [AWS SaaS Connection](#) deployed in the target AWS region.

6. Under **Recover To**, select **Original Location**.
7. Select your Recovery Options:
 - **Rename:** Add a **Prefix** and/or **Suffix** to the recovered Amazon EC2 instances.
 - **Power State:** Disable **Power On** if you want the recovered **EC2 instances** to remain powered off after they are created.
 - **Custom AWS Resource Tags:** Add your custom AWS tag for the EC2 instance that will be recovered by specifying the **Tag Key** and **Tag Value** for the new custom tag. You can add multiple tags by clicking the **add** icon. These new tags will be attached to the recovered EC2 instance along with the existing tags.

If you provide a new **Tag Value** for an existing tag, the recovered EC2 instance will have this tag attached to the new **Tag Value**.

AWS allows you to add up to 50 tags to an AWS resource. When Cohesity recovers data, it adds 2 tags to the resource. If the number of custom and existing tags exceeds 48, the custom tags will override the existing tags. In this case, Cohesity will randomly discard some of the existing tags to accommodate the custom tags.
 - **Continue on Error:** Enable this option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.

- **Task Name:** Change the default name of the recovery task.

The screenshot shows the 'Virtual Machines' recovery configuration interface. At the top, it displays 'Virtual Machines', a snapshot from 'Oct 23, 2023 6:23pm', and the location 'AWS US East (Ohio)'. Below this, the 'Recover To' section has 'Original Location' selected. The 'Recovery Options' section includes a 'Rename' field with 'cpal-' as a prefix and an 'Add Suffix' field. The 'Power State' is set to 'On'. The 'Custom AWS Resource Tags' section shows a tag with key 'custom_tag_1' and value 'custom_value_1'. The 'Continue on Error' option is set to 'No'. Finally, the 'Task Name' is 'Recover_Oct_25_2023_10_46_AM'.

8. Click **Recover**.

Cohesity DataProtect as a Service begins to restore the selected Amazon EC2 instances.

Recover EC2s to New Location

To recover your protected Amazon EC2 instances to a new location:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the **Source** name.
3. Select **Show All > Protected**.
4. Use the filters, search box, and views to locate and select the EC2 instances you want to recover.

Tip: You can also use Global Search to locate, filter, and select the objects you need. Click the **Global Search** box at the top or type **slash (/)** anywhere to start your search.

5. Click the **Recover** icon at the top to open the **New Recovery** form. By default, the **Latest** snapshot is pre-selected for recovery. If you need to recover from an earlier snapshot, click the **Edit** (pencil) icon to choose the desired snapshot. The icon(s) displayed under **Location** indicates the snapshot type(s) available (**AWS snapshot** and/or **Cohesity snapshot**) for recovery. Choose a snapshot type and click its icon to proceed with the recovery task:

- Click **Select Recovery Point**.
- Click **Next: Recover Options** to return to the form.

Note: To recover from a [Cohesity snapshot](#), you need an [AWS SaaS Connection](#) deployed in the target AWS region.

6. Under **Recover To**, select **New Location** and provide:

- **Source.** Select a registered AWS account as the new recovery destination.
- **Region.** Select a destination AWS region.
- **Key Pair.** Select an AWS key pair to be associated with the recovered EC2 Instance(s).
- **Encryption Settings:** Toggle on and provide the following details to encrypt the EC2 instance(s) to be recovered.
 - **KMS Key Location:** From the drop-down, select whether you want to use the KMS key from the **Same AWS Account** where you are recovering to or from a **Different AWS Account**.
 - **KMS Key:** If you are using the KMS key from the **Same AWS Account**, then from the drop-down, select the KMS Key. If you are using the KMS key from a **Different AWS Account**, then enter the KMS Key ARN in the following format:

```
arn:<partition>:kms:<region>:<account-id>:key/<key-id>
```

Note: If **Encryption Settings** is enabled, all disks of EC2 will be encrypted with the provided key.

By default, this option is disabled.

- For Cohesity snapshot recovery with **Encryption Settings** disabled, the recovered EC2 instances will remain unencrypted irrespective of its encryption status before recovery.
- For AWS snapshot recovery with **Encryption Settings** disabled, the EC2 instance will be recovered with the *default* AWS KMS encryption key of the destination region.
- **Subnet.** Select the subnet where the recovered EC2 Instance(s) will be launched.
- **Network Security Group.** Select the network security group to be associated with the recovered EC2 Instance(s).

7. Select your **Recovery Options**:

- **Rename:** Add a **Prefix** and/or **Suffix** to the recovered Amazon EC2 instances.
- **Power State:** Disable **Power On** if you want the recovered **EC2 instances** to remain powered off after they are created.
- **Custom AWS Resource Tags:** Add your custom AWS tag for the EC2 instance that will be recovered by specifying the **Tag Key** and **Tag Value** for the new custom tag. You can add multiple tags by clicking the **add** icon. These new tags will be attached to the recovered EC2 instance along with the existing tags.

If you provide a new **Tag Value** for an existing tag, the recovered EC2 instance will have this tag attached to the new **Tag Value**.

AWS allows you to add up to 50 tags to an AWS resource. When Cohesity recovers data, it adds 2 tags to the resource. If the number of custom and existing tags exceeds 48, the custom tags will override the existing tags. In this case, Cohesity will randomly discard some of the existing tags to accommodate the custom tags.

- **Continue on Error:** Enable this option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.
- **Task Name:** Change the default name of the recovery task.

8. Click **Recover**.

Cohesity DataProtect as a Service begins to restore the selected Amazon EC2 instances.

Amazon EC2 Recovery Support Matrix

Backup Type	Data Source	Recovery to Same AWS account, same AWS region	Recovery to same AWS account, different AWS region	Recovery to a diff AWS account, same AWS Region	Recovery to a diff AWS account, different AWS Region
AWS Snapshot	Encrypted	Supported	Supported	Supported	Supported
AWS Snapshot	Non-Encrypted	Supported	Supported	Supported	Supported
Cohesity Snapshot	Encrypted	Supported	Supported	Supported	Supported
Cohesity Snapshot	Non-Encrypted	Supported	Supported	Supported	Supported

Important Considerations

For recovery to:

- **Same AWS Account, same AWS Region:** No prerequisites.
- **Different AWS Account, same AWS Region:**
 - a. Target AWS account should be registered as a data source in the same destination cloud region.
 - b. If you have enabled **Encryption Settings** for an EC2 instance that is already encrypted, then the EC2 instance will be recovered with the encryption provided in the **Encryption Settings**.
 - c. Additional limitations for AWS snapshot recovery:
 - AWS KMS encryption key should be shared from the source AWS account to the target AWS account before the recovery is attempted.
 - If the source EC2 instances were encrypted with the *default* AWS KMS encryption key, their recovery to a different AWS account will fail. (AWS limitation).
 - You cannot unencrypt an EC2 instance that is already encrypted. That is, if you have configured to recover an encrypted EC2 instance with the **Encryption Settings** option disabled, then the EC2 instance will be recovered with the default AWS KMS encryption key of the destination region.
- **Different AWS Account, different AWS Region:**
 - a. Target AWS account should be registered as a data source in the same destination cloud region.
 - b. If you have enabled **Encryption Settings** for an EC2 instance that is already encrypted, then the EC2 instance will be recovered with the encryption provided in the **Encryption Settings**.
 - c. Additional limitations for AWS snapshot recovery:
 - AWS KMS encryption key should be shared from the source AWS account to the target AWS account before the recovery is attempted.
 - If the source EC2 instances were encrypted with the *default* AWS KMS encryption key, their recovery to a different AWS account will fail. (AWS limitation).
 - You cannot unencrypt an EC2 instance that is already encrypted. That is, if you have configured to recover an encrypted EC2 instance with the **Encryption Settings** option disabled, then the EC2 instance will be recovered with the *default* AWS KMS encryption key of the destination region.

Recover Amazon EC2 Files and Folders

You can download or restore specific files and folders from a protected EC2 instance to either the original or an alternate EC2 instance.

Prerequisites

- The SaaS Connector must be able to reach the target VM on port 50051 so that the SaaS Connector can push the files being recovered to the target VM using the Cohesity agent.
- If the Cohesity Agent is to be installed as part of the recovery task in Cohesity, ensure that:
 - AWS Systems Manager Agent (SSM) access is available on the target VM. For more information, see [AWS documentation](#).
 - The target VM is able to reach the SaaS Connector on port 443 so that the target VM can pull the agent installer from the SaaS Connector.

Note: For enhanced security, when installing the agent on the target EC2, Cohesity automatically deploys an X.509 certificate.

Considerations

When recovering files and folders from protected Amazon EC2 instances, remember:

- Files and folders download is only available for EC2 Cohesity snapshots and not for AWS snapshots.
- The maximum number of files that can be recovered is up to 100k.
- Download of symlinks is not available.
- Recovery of Windows symlinks is not supported.
- Recovery of files and folders from a combination of different volumes is not supported.

Recover Amazon EC2 Files and Folders

Important: To restore files from a [Cohesity snapshot](#) to an Amazon EC2 instance, you need an [AWS SaaS Connection](#) deployed in the target AWS region.

To recover or download your files and folders from your protected Amazon EC2 instances:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the **Source** name.
3. Select **Protection Status > Protected**.

4. Use the filters, search box, and views to locate and select the EC2 instances you want to recover.

You can also use **Global Search** to locate, filter, and select the objects you need. Click the Global Search box at the top or type **slash (/)** anywhere to start your search.

5. Click the required EC2 backed up as a Cohesity Snapshot and click the **Recover Files** icon for the EC2. The page with the EC2 details is displayed.
6. Select the timeline drop-down list on the top right corner to select the snapshot and click **Apply**.
7. Click the required volume to browse the file system and select the file or folder to be recovered.

You can click **Download Files** to download the selected files.

8. Click **Next**. The **Files** page is displayed.
9. Under **Recover To**, select **Original Server** or **New Server**.
 - For recovery to the *original* EC2, you can provide the new recovery path in the **Recover To** field or use the **Recover To Original Path** option to recover to the original path on the original server.
 - For recovery to an *alternate* EC2, you can choose any AWS server and select a **Target**. Provide the new recovery path in the **Recover To** field.

Note: The recovery process will attempt to install the Cohesity Agent on the target EC2 instance using AWS SSM. If the SSM agent is not running on the target EC2 instance or if the Cohesity IAM role does not have access to send SSM commands to the target EC2 instance, then you can download the agent using the **Download Cohesity Agent** link and install it on the target EC2 instance before starting the recovery. For more details, see [Download and Install the Cohesity Agent](#).

10. Select your Recovery Options:
 - **Overwrite Existing File/Folder:** Enable this option to overwrite the existing files and folders. Disable this option to create the files and folders in the specified location. If a file with the same name already exists in the target location, the file is overwritten or skipped based on this selection. If **Overwrite Existing File/Folder** is enabled, recovering a file to source when the file is in use may cause the open file to be overwritten. Whether overwriting occurs depends on the application using the file.
 - **Preserve File/Folder Attributes:** By default, this option is enabled and the ACLs, permissions, and timestamps are preserved for all files and folders. If you

disable this option, then ACLs and permissions are not preserved. If both folders and files are recovered, then folders will receive the new timestamps, but files retain their original timestamps. If recovering only files, then files will receive the new timestamps.

- **Continue on Error:** Enable this option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.
- **Task Name:** Change the default name of the recovery task.

11. Click **Recover**.

Cohesity DataProtect as a Service begins to restore the selected Amazon EC2 files and folders.

Download and Install the Cohesity Agent

Install the Cohesity Agent on each Windows and Linux Amazon EC2 instance that you want to recover to.

Install the Cohesity Windows Agent

To download and install the Cohesity Windows Agent:

1. Navigate to the **Files** page to recover the Amazon EC2 instance. To access the **Files** page, follow steps 1-8 in [Recover Amazon EC2 Files and Folders](#) above.
2. Click **Download Cohesity Agent** and download it to the appropriate server.
3. As an administrator with local system privileges on that server, run the executable and complete the installation wizard.

Install the Cohesity Linux Agent

The Cohesity Linux Agent is available with different installer packages, providing support on multiple Linux distributions. You'll need to [install the appropriate package](#) (RPM, Debian, or SUSE RPM) for your Linux distribution or [install the script installer package](#).

The installer packages and Linux distributions on which the installer package is supported are:

Installer Package	Linux Distribution
(Default) RPM	RHEL and its click derivative
Suse RPM	SUSE
Debian	Ubuntu
Script Installer	All supported Linux Operating Systems

The Cohesity Linux Agent has dependencies on the following packages, which must be installed on the Linux server:

Command/Package	RHEL	SUSE	CentOS	Ubuntu	Debian
rsync	rsync	rsync	rsync	rsync	rsync
mount	util-linux	util-linux	util-linux	mount	mount
lvm2	lvm2	lvm2	lvm2	lvm2	lvm2
sudo	sudo	sudo	sudo	sudo	sudo
coreutils	coreutils	coreutils	coreutils	coreutils	coreutils
util-linux	util-linux	util-linux	util-linux	util-linux	util-linux
nfs client	nfs-utils	nfs client	nfs-utils	nfs-common	nfs-common
lsof	lsof	lsof	lsof	lsof	lsof
wget	wget	wget	wget	wget	wget

Install RPM, Debian, or SUSE RPM Installer Package

To install the RPM, Debian, or SUSE installer package:

1. Navigate to the **Files** page to recover the Amazon EC2 instance. To access the **Files** page, follow steps 1-8 in [Recover Amazon EC2 Files and Folders](#) above.
2. Click **Download Cohesity Agent**. Based on your Linux distribution, from the **Download Agents** window, select RPM, Debian, or SUSE RPM and download it to the server you want to protect.
3. As the root user with local system privileges on that server, change the directory to the location of the installer package.
4. Run the following command depending on the installer package:

Installer Package	Command
RPM	rpm -i el-cohesity-agent-6.5.1-1.x86_64.rpm or yum localinstall ./el-cohesityagent-6.5.1-1.x86_64.rpm
Debian	dpkg -i cohesity-agent_6.5.1-1_amd64.deb
Suse RPM	rpm -i cohesity-agent-6.5.1-1.x86_64.rpm

Note:

By default, the installation uses the root user permission for all the files, and the service is started as root. Therefore, it is necessary to add non-root users to the sudoers list by making the following changes in the /etc/sudoers file:

```
<username> ALL=(ALL) NOPASSWD:ALL
Defaults:<username> !requiretty
```

5. To start the service as a non-root user, create a new user or use an existing user with sudo permission and run the following command:

Installer Package	Command
RPM	export COHESITYUSER= <username>; rpm -i el-cohesity-agent- 6.5.1-1.x86_64
Debian	COHESITYUSER= <username> dpkg -i cohesity-agent_6.5.1- 1_amd64

Installer Package	Command
Suse RPM	<pre>export COHESITYUSER= <username> rpm -i cohesity-agent-6.5.1- 1.x86_64</pre>

6. Provide the location details for:

- **Installation directory:** /opt/cohesity
- **Log file:** /var/log/cohesity

Install Script Installer Package

To install the script installer package:

1. Navigate to the **Files** page to recover the Amazon EC2 instance. To access the **Files** page, follow steps 1-8 in [Recover Amazon EC2 Files and Folders](#) above.
2. Click **Download Cohesity Agent**. In the **Download Agents** window, select **Script Installer** based on your Linux distribution, and download it to the server you want to protect.
3. As the root user with local system privileges on that server, change the directory to the location of the installer package.

Note: For SLES 11 SP4, you are required to install the Agent as the root user.

4. Make the installer executable. For example:

```
chmod +x cohesity_agent_6.5.1-master_linux_x64_installer
```

5. Run the executable:

```
sudo cohesity_agent_6.5.1-master_linux_x64_installer -- --install
```

6. Provide the location details for:

- **Installation directory:** /home/<username>/cohesityagent or /root/cohesityagent
- **Log file:** /home/cohesityagent/cohesityagent/logs

The Agent starts after installation completes, as follows:

- **CentOS and RedHat** (distributions with the "systemd" init system): The Agent starts automatically.
- **Ubuntu** (distributions with the "upstart" init system): The Agent starts automatically.

If a Linux server's `/etc/sudoers` file is managed by a deployment engine such as Chef, Puppet, or others, this might affect Cohesity's interaction with servers that have the Linux Agent installed. Take the corresponding actions depending on user type:

Agent Installation by User Type	Action Required
As the default cohesityagent user	<p>The Cohesity Linux Agent is installed using the cohesityagent user by default.</p> <p>For default installations, the cohesityagent user is created by the installer. During installation, the installer updates the <code>/etc/sudoers</code> file to allow cohesityagent sudo and no-tty sudo access.</p> <p>Ensure the following settings in the <code>/etc/sudoers</code> file for the cohesityagent user are preserved:</p> <pre>cohesityagent ALL= (ALL) NOPASSWD:ALL Defaults:cohesityagent !requiretty</pre> <p>For example:</p> <pre>#includedir /etc/sudoers.d dgoble ALL= (ALL) NOPASSWD:ALL cohbackup ALL= (ALL) NOPASSWD:ALL Defaults:cohbackup !requiretty</pre>

Agent Installation by User Type	Action Required
As a non-default user, for example, foo	Ensure the above settings in the <code>/etc/sudoers</code> file for the foo user are preserved by replacing the occurrences of 'cohesityagent' with 'foo'.
As root user	No changes required.

Amazon RDS Instance

Cohesity DataProtect as a Service provides a simple, fast, and cost-effective backup, recovery, and data management solution for Amazon RDS instances in your AWS account.

Protection Type Options

Cohesity DataProtect as a Service provides two **Protection Type** options for the protection of AWS RDS:

- **AWS Snapshot:** Cohesity DataProtect as a Service uses this option to protect RDS instances. For protecting the RDS instances, Cohesity DataProtect as a Service leverages the AWS native snapshot and stores them in the same AWS account and region as the source RDS instances.

When creating protection for the RDS instance, by default, **AWS Snapshot** is selected as the Protection Type and you cannot change this value.

- **Amazon RDS Ingest:** Cohesity DataProtect as a Service uses this option to protect the databases on the RDS instance by ingesting the backup data to an AWS region supported by the Cohesity DataProtect as a Service. The target AWS region is the region that is selected during [AWS source registration](#). Amazon RDS Ingest snapshots provide an air-gapped backup and recovery. With air-gapped backup approach, the backed up data is isolated from any network connectivity, ensuring that your data remains safe. Network connectivity is resumed only during the recovery process, minimizing the risk of ransomware attacks.

When creating protection for databases on the RDS instance, by default, **Amazon RDS Ingest** is selected as the Protection Type and you cannot change this value. For information on protecting RDS database, see [Amazon RDS Database](#).

Note: Cohesity performs only full backups of RDS using the Amazon RDS Ingest method. Incremental backup is not supported in the Amazon RDS Ingest method.

Supported RDS Databases Using Ingest-Based Protection

Using the Amazon RDS Ingest method, Cohesity supports the protection of the following individual database of the RDS instance:

- PostgreSQL: 11.22-R2- 15.6.
- Aurora (PostgreSQL Compatible): 11.9 - 15.6

Protect Your Amazon RDS

Once you have registered your AWS account, you are ready to protect the RDS of your AWS account. You can protect the RDS using either the AWS Snapshot protection method or the Cohesity Ingest protection method.

- With the **AWS Snapshot** protection method, Cohesity protects the RDS of your AWS account at the instance level.
- With the **Amazon RDS Ingest** protection method, Cohesity protects the RDS of your AWS account at the database level.

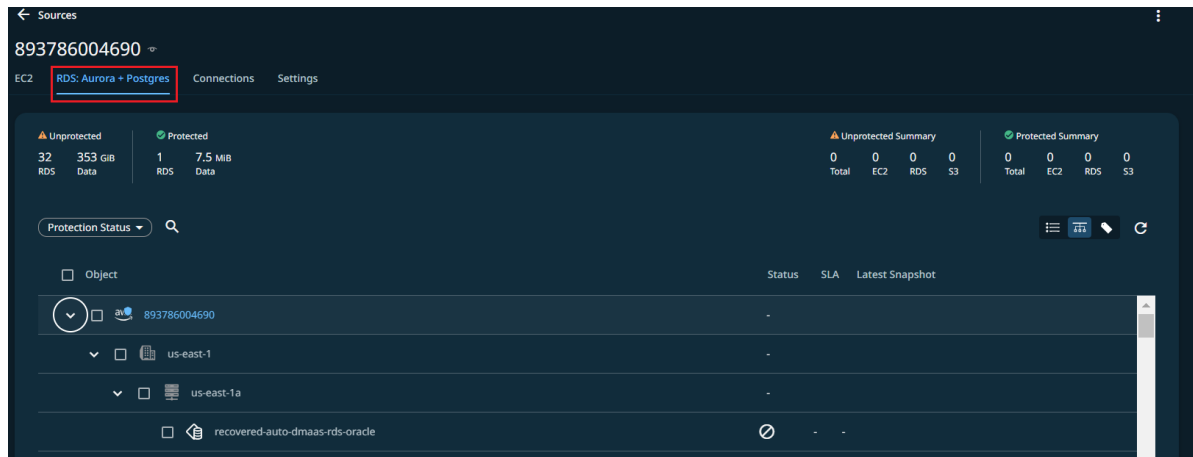
To protect the RDS of your AWS account at the database level using the Amazon RDS Ingest method, you must provide the credentials of the RDS instance while creating the protection. By providing the credentials, Cohesity discovers the databases as objects for protection. For the list of supported RDS databases for protection, see [Supported RDS Databases Using Ingest-Based Protection](#) .

Note: Cohesity performs only full backups of RDS using the Amazon RDS Ingest method. Incremental backup is not supported in the Amazon RDS Ingest method.

Note: If you have already registered your AWS account to protect AWS S3 or AWS EC2 workloads, then you must [Update the Existing CloudFormation Template](#) to update the Cohesity permissions in your AWS account.

To protect your Amazon RDS:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Find the registered AWS account and click into it.
3. Click the **RDS: Aurora + Postgres** tab.




4. Click the **Hierarchy View** icon located at the right corner of the page.
5. Perform one of the following steps based on the protection type method you prefer:

- **AWS Snapshot method:**


- i. Use the checkboxes to select the RDS instances for protection. You can also select the RDS instances at various hierarchy levels (AWS account level, region, and availability zone) by selecting **Select All Child Objects**, provided the instances are of the same type at the level you select for protection.
- ii. Optionally, you can auto-protect the RDS instances at various hierarchy levels- **account level**, **AWS region level**, or **Availability zone level**. The auto-protect option enables you to automatically protect the new RDS instances that are added.

- To auto-protect the RDS instances at AWS account-level, select the checkbox of the AWS account, and then select **Auto Protect This AWS**.




- To auto-protect the RDS instances at the AWS region level, select the checkbox of the region, and then select **Auto Protect This**

Region.



- To auto-protect the RDS instances at the AWS availability zone level, select the checkbox of the availability zone, and then select **Auto Protect This Availability Zone**.



- Click the **Protect** icon above the checkboxes.

In the **New Protection** dialog, the AWS Snapshot-based policy is enabled by default. Create or select an existing policy from the **Policy (AWS Snapshot)** drop-down.

- Amazon RDS Ingest method:**
 - Select the RDS instance having the databases to protect. You can also select the RDS instances at various hierarchy levels- AWS account level, region, and availability zone by selecting **Select All Child**.
 - Click **Database Credentials** above the checkboxes.
The Backup Database Credentials page appears.
 - Click **+ Add Credential**.
 - Select the **Type** of database you want to discover for protection:
 - PostgreSQL
 - Aurora PostgreSQL
 - Select one of the following authentication methods:
 - Credentials:** Enter the Username, and Enter Password for the RDS instance.
 - IAM Authentication:** Enter the Username for IAM authentication.
Ensure the username you provide has the following permissions:

- `GRANT rds_iam TO <username>`
- `GRANT pg_read_all_data TO <username>`
- `GRANT pg_write_all_data TO <username>`
- `ALTER USER <username> CREATEDB`

- **Kerberos Authentication:** Enter the **Username**, **Enter Password**, **Real Name**, **Active Directory DNS Address**.

Ensure the username you provide has the following permissions:

- `GRANT rds_ad TO <username>`
- `GRANT pg_read_all_data TO <username>`
- `GRANT pg_write_all_data TO <username>`
- `ALTER USER <username> CREATEDB`

- vi. Click **Add Credentials** and follow the two steps above to discover a database of a different type.
- vii. Click **Save**.
- viii. Optionally, you can auto-protect databases of the RDS instances. The auto-protect option enables you to automatically protect the new RDS databases that are added. To auto-protect the database, click the checkbox of the RDS instance. Depending on the databases on the RDS instance, one of the following options is displayed.
 - **Auto Protect This RDS:** If the databases on the RDS instances are PostgreSQL.
 - **Auto Protect This Aurora Cluster:** If the databases on the RDS instances are Aurora.
- ix. Click the **Protect** icon above the checkboxes.

In the **New Protection** dialog, the Cohesity ingest-based policy is enabled by default. Create or select an existing policy from the **Policy (Cohesity Ingest)** drop-down.

Note: If you are protecting RDS instances with discovered databases, both the AWS snapshot and Cohesity Ingest policy options will be available. Depending on the policy you choose, the RDS instances will be protected either at the instance level or at the database level. If you select both policies, the RDS instances will be protected at both the instance level and the database level.

6. To change or configure any of the additional settings, select **More Options** and perform the below steps or else, click **Protect**.
7. Under **Settings**, edit the **Start Time** if necessary.
8. Under **Additional Settings**, configure the following option:
 - **Cancel Runs at Quiet Time Start:** (Available only if the selected policy has at least one [Quiet Time](#)) When enabled, all the protection runs that are currently executing will cancel when the Quiet Time period starts. By default, this setting is disabled, meaning that after a protection run starts, it continues to execute even when a Quiet Time period starts. However, new protection runs will not start during a Quiet Time.
9. Click **Protect**.

Manage Existing Protection


Edit protection settings, change the policy, and start, stop, & pause protection.

Once you have [applied protection](#) to the objects in your sources, Cohesity DataProtect as a Service makes it easy to make changes to that protection quickly. You can:

- Edit additional settings.
- Apply a different policy.
- Start an on-demand protection run, pause and resume it, or even remove protection.

Edit Protection Settings

To edit protection settings:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click into the **Source** name.
3. Select **Show All > Protected** and use the other filters, search box, and views at the top to narrow your search.
4. Click the **Actions** menu () next to the object and select **Edit Protection** to open the protection settings for that object.

Apply a New Protection Policy

To change the **Policy**, click the drop-down and select a different policy. To help you choose, each policy in the list shows the **Backup** frequency and the **Retain** period for each backup.

If you don't have a policy that meets your needs, scroll to the bottom of the list and click **Create Policy** to [create your own policy](#).

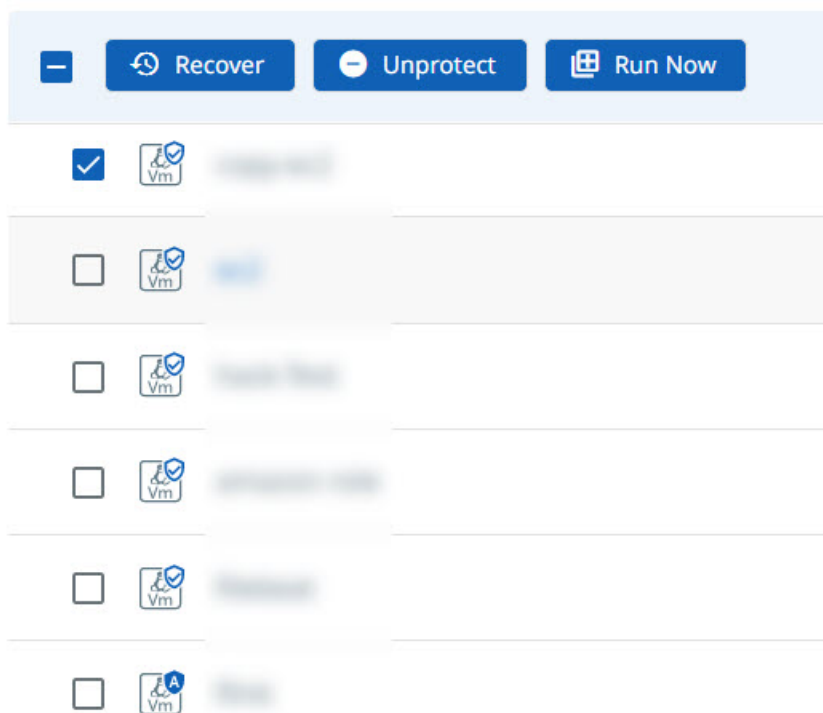
Edit Additional Protection Settings

Under **Settings**, you can change the protection **Start Time** (and select the **Time Zone**).

Click the drop-down next to **Additional Settings** to change more options. See [Additional Protection Settings](#) for details.

Start, Stop, or Remove Protection

When you select protected objects in one of your sources, Cohesity DataProtect as a Service presents buttons for the actions that are possible for those objects.



With the protected objects selected, you can click:

- **Recover** to recover the object or file.
- **Unprotect** to remove protection from the object.

Tip: If a protected object is deleted from the source, you can search the object using Global Search and unprotect it.

- **Run Now** to start an on-demand protection run immediately.

Additional Settings

Advance Settings	Description
End Date	If you need to end protection on a specific date, enable this to select the date.
Quiet Times	Available only if the selected policy has at least one quiet time period. Toggle it ON to specify that all currently executing protection runs should abort if a quiet time period specified for the Protection Group starts. By default this toggle is OFF, which means after a protection run starts, it continues to execute even when a quiet time period specified for this protection run starts. However, a new protection run will not start during a quiet time period.

Recover Your Amazon RDS Instances

After you [protect your Amazon RDS DB instances](#), you can recover them to their [original location](#) or a [new location](#) using Cohesity DataProtect as a Service.

We recommend that you also review the [Amazon RDS Recovery Support Matrix](#) and [Important Considerations](#) at the end of this article.

Recover Amazon RDS Instances to Original Location

To recover your protected Amazon RDS DB instances to their original location:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the **Source** name.
3. Select **Show All > Protected**.
4. Use the filters, search box, and views to locate and select the DB instances you want to recover.

Tip: You can also use Global Search to locate, filter, and select the objects you need. Click the **Global Search** box at the top or type **slash (/)** anywhere to start your search.

5. Click the **Recover** icon at the top to open the **New Recovery** form. By default, the **Latest** snapshot is pre-selected for recovery. If you need to recover from an earlier snapshot, click the **Edit** (pencil) icon to choose the desired snapshot. Selecting an invalid time from the slider automatically selects the closest available snapshot.
6. Under **Recover To**, select **Original Location**.
7. Enable the **Multi A-Z Deployment** option if you want the database instances to be recovered to have a standby instance deployed in another availability zone. This

option is disabled by default.

8. Configure the following **Additional Settings**:

- a. **Database Instance Identifier**: Specify the unique key that identifies the database instance that will be recovered.
- b. **DB Port**: Specify the TCP/IP port that the DB instance will use for application connections. The connection string of any application connecting to the DB instance must specify the port number of the DB instance. Both the security group applied to the instance and your company's firewalls must allow connections to this port.
- c. **IAM DB Authentication**: Enable this option if you want to manage your database user credentials through AWS IAM users and roles. This option is disabled by default.
- d. **Public Accessibility**: Enable this option if you want the DB instance to also have a public IP address in addition to the private IP address. This option is disabled by default.
- e. **Copy Tags To Snapshots**: Enable this option for copying tags to snapshots. This option is disabled by default.
- f. **Auto Minor Version Upgrade**: Enable this option if you want the DB instance to automatically upgrade when a new minor database engine version is available. This option is disabled by default.

9. Optional. Change the default name of the recovery task in the **Task Name** field.

10. Click **Recover**.

Cohesity DataProtect as a Service begins to restore the selected Amazon RDS databases.

Recover Amazon RDS Instances to New Location

To recover your protected Amazon RDS DB instances to a new location:

1. Go to **Sources**.
2. Click the **Source** name.
3. Select **Show All > Protected**.
4. Use the filters, search box, and views to locate and select the Amazon RDS instance you want to recover.

Tip: You can also use Global Search to locate, filter, and select the objects you need. Click the **Global Search** box at the top or type **slash (/)** anywhere to start your search.

5. Click the **Recover** icon at the top to open the **New Recovery** form. By default, the **Latest** snapshot is pre-selected for recovery. If you need to recover from an earlier snapshot, click the **Edit** (pencil) icon to choose the desired snapshot. Selecting an invalid time from the slider automatically selects the closest available snapshot.
6. Under **Recover To**, select **New Location**.
7. Enable the **Multi A-Z Deployment** option if you want the Amazon RDS database instance to be recovered to have a standby instance deployed in another availability zone. This option is disabled by default.
8. Under **Location**, provide the following information:
 1. **Source**: Select a registered AWS account as the new recovery destination.
 2. **Region**: Select a destination AWS region.
9. Under **Network Settings**, configure the following settings:
 1. **Subnet**: Select a subnet in the Amazon VPC to store the recovered RDS instance.
 2. **Network Security Groups**: Select the security group that should be applied to the DB instance.
 3. **Availability Zone**: Select an availability zone in AWS to recover the RDS instance.
10. Configure the following **Additional Settings**:
 1. **Database Instance Identifier**: Specify the unique key that identifies the database instance that will be recovered.
 2. **DB Port**: Specify the TCP/IP port that the DB instance will use for application connections. The connection string of any application connecting to the DB instance must specify the port number of the DB instance. Both the security group applied to the instance and your company's firewalls must allow connections to this port.
 3. **DB Option Group**: Select an option group that contains the option you want to attach to the DB instance that will be recovered. If there are not any option groups compatible with the selected engine, a default group will be created at launch.
 4. **DB Parameter Group**: Select the database parameter group to associate with the DB instance.
 5. **IAM DB Authentication**: Enable this option if you want to manage your database user credentials through AWS IAM users and roles. This option is disabled by default.
 6. **Public Accessibility**: Enable this option if you want the DB instance to also have a public IP address in addition to the private IP address. This option is

disabled by default.

7. **Copy Tags To Snapshots:** Enable this option for copying tags to snapshots. This option is disabled by default.
8. **Auto Minor Version Upgrade:** Enable this option if you want the DB instance to automatically upgrade when a new minor database engine version is available. This option is disabled by default.
11. Optional. Change the default name of the recovery task in the **Task Name** field.
12. Click **Recover**.

Cohesity DataProtect as a Service begins to restore the selected Amazon RDS databases.

Amazon RDS Recovery Support Matrix

Backup Type	Data Source	Recovery to Same AWS account, Same AWS region	Recovery to same AWS account, different AWS region	Recovery to a diff AWS account, same AWS region	Recovery to a diff AWS account, different AWS region
AWS Snapshot	Encrypted	Supported	Supported	Supported	Supported
AWS Snapshot	Non-Encrypted	Supported	Supported	Supported	Supported

Important Considerations

For recovery to:

- **Same AWS Account, same AWS Region:** No prerequisites.
- **Same AWS Account, Different AWS Region:** To recover encrypted RDS instance(s), you must create a KMS encryption key in the target AWS account & region with the same alias name as the KMS encryption key used to encrypt the source RDS instance(s).

Recover Your Amazon RDS Databases

After you protect your RDS databases, you can recover them to their original location or a new known location.

Note: You can recover the database on RDS only if you have protected the RDS at the database level (**Amazon RDS Ingest**) and not at the instance level.

Recover Database on RDS to Original Location

To recover your protected RDS database to its original location:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the **Source** name.
3. Select **Show All > Protected**.
4. Use the filters, search box, and views to locate and select the DB instances you want to recover.

Tip: You can also use Global Search to locate, filter, and select the objects you need. Click the **Global Search** box at the top or type **slash (/)** anywhere to start your search.

5. Click the **Recover** icon at the top to open the **New Recovery** form. By default, the **Latest** snapshot is pre-selected for recovery. If you need to recover from an earlier snapshot, click the **Edit** (pencil) icon to choose the desired snapshot. Selecting an invalid time from the slider automatically selects the closest available snapshot.
6. Under **Recover To**, select **Original Aurora PostgreSQL Cluster**.
Since you are recovering to the original location, the options to select Source, Region, and Instance are disabled.
7. Enable **Overwrite Existing Objects** if you want to recover the database by overwriting the original RDS database. This option is disabled by default.
8. In the **Rename** field, add **Prefix** and/or **Suffix** strings to the name of the database that will be recovered.
9. In the **Task Name** field, change the default name of the recovery task.

10. Click **Recover**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Recover Database on RDS to Another Location

You can recover the PostgreSQL or Aurora databases to an alternate RDS database. The alternate location can be:

- **Known RDS/Aurora PostgreSQL Cluster:** You are recovering the database to an alternate RDS instance which already exist.
- **Non-RDS Server:** You are recovering the database to a PostgreSQL running on an AWS account but not managed by AWS.

Important: When recovering the RDS database to an alternate location, the AWS region you select must have an AWS SaaS Connection deployed. The recovery will fail if there is no SaaS connection on the region you select.

Recover RDS Database to Known Aurora PostgreSQL Cluster

To recover the database to an alternate RDS instance which already exist, perform the following steps:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the **Source** name.
3. Select **Show All > Protected**.
4. Use the filters, search box, and views to locate and select the Amazon RDS instance you want to recover.

Tip: You can also use Global Search to locate, filter, and select the objects you need. Click the **Global Search** box at the top or type **slash (/)** anywhere to start your search.

5. Click the **Recover** icon at the top to open the **New Recovery** form. By default, the **Latest** snapshot is pre-selected for recovery. If you need to recover from an earlier snapshot, click the **Edit** (pencil) icon to choose the desired snapshot. Selecting an invalid time from the slider automatically selects the closest available snapshot.
6. Under **Recover To**, select **Known Aurora PostgreSQL Cluster**.
7. From the **Source** drop-down, select the AWS source having the RDS instance to which you want to recover the database. You can also register a new AWS account by selecting **Register Source**.
8. From the **Region** drop-down list, select the region of the RDS instance to which you want to recover the database.

Note: Ensure the region you select has an AWS SaaS Connection deployed. The recovery will fail if there is no SaaS connection in the region you select.

9. From the **Instance** drop-down list, select the RDS instance to which you want to recover the database.
10. If there is already a database on the target RDS instance with the same name as the database selected for recovery, then you can enable **Overwrite Existing Objects** if you want to recover the database by overwriting the RDS database on the target RDS instance. This option is disabled by default.
11. In the **Rename** field, add Prefix and/or **Suffix** strings to the name of the database that will be recovered.
12. Optional. Change the default name of the recovery task in the **Task Name** field.
13. Click **Recover**.

Amazon RDS

postgres RDS Instances Latest Snapshot AWS US East (Ohio) Location

Recover To

☐ Original RDS PostgreSQL Instance
 ☒ Known RDS PostgreSQL Instance
 ☐ Custom PostgreSQL Server

Source: 893786004690
 Region: AWS US East (Ohio)
 Instance: auto-dmaas-aurora-encrypted
Select RDS PostgreSQL Instance

☒ Overwrite Existing Objects

Rename: Prefix: copy_

Task Name: Recover_Dec_14_2023_12_49_PM

Recover Cancel

The recovery task is initiated. You can monitor the recovery task from the Recoveries page

Recover RDS Database to Non-RDS Server

To recover the database to an alternate RDS instance which already exist, perform the following steps:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the **Source** name.
3. Select **Show All > Protected**.
4. Use the filters, search box, and views to locate and select the Amazon RDS instance you want to recover.

Tip: You can also use Global Search to locate, filter, and select the objects you need. Click the **Global Search** box at the top or type **slash (/)** anywhere to start your search.

5. Click the **Recover** icon at the top to open the **New Recovery** form. By default, the **Latest** snapshot is pre-selected for recovery. If you need to recover from an earlier snapshot, click the **Edit** (pencil) icon to choose the desired snapshot. Selecting an invalid time from the slider automatically selects the closest available snapshot.
6. Under **Recover To**, select **Custom PostgreSQL Server**.
7. Provide the following information:
 1. **Region:** Select the region of the AWS account to which you want to recover the database.

Note: Ensure the region you select has an AWS SaaS Connection deployed. The recovery will fail if there is no SaaS connection on the region you select

2. **IP Address:** Enter the IP address of the server (target) to which you want to recover the RDS database.
3. **Port:** Enter the port number that can be used for connecting to the target server.
4. **Username:** Enter the user name of the target server.
5. **Enter Password:** Enter the password of the target server.
8. If there is already a database on the target RDS instance with the same name as the database selected for recovery, then you can enable **Overwrite Existing Objects** if you want to recover the database by overwriting the RDS database on the target RDS instance. This option is disabled by default.
9. In the **Rename** field, add **Prefix** and/or **Suffix** strings to the name of the database that will be recovered.
10. In the **Task Name** field, change the default name of the recovery task.
11. Click **Recover**.

The screenshot shows the 'Recover To' dialog in the Amazon RDS console. At the top, it displays 'postgres RDS Instances', 'Latest Snapshot', and 'AWS US East (Ohio) Location'. The 'Recover To' section has three radio buttons: 'Original RDS PostgreSQL Instance', 'Known RDS PostgreSQL Instance', and 'Custom PostgreSQL Server' (which is selected). Below this, there are fields for 'Region' (set to 'AWS US East (Ohio)'), 'IP Address' (set to '19.3.44.33'), and 'Port' (set to '50052'). There is also a 'Username' field set to 'admin' and a password field labeled 'Enter Password' with masked characters and a toggle icon. A checkbox labeled 'Overwrite Existing Objects' is checked. At the bottom, there is a 'Rename' section with a 'Prefix' set to 'copy_' and a 'Task Name' field set to 'Recover_Dec_14_2023_12_49_PM'. At the very bottom, there are 'Recover' and 'Cancel' buttons.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Amazon S3 Buckets

Cohesity leverages the Amazon S3 inventory report to protect the Amazon S3 bucket. The inventory report contains the list of all the objects available on the Amazon S3 bucket you selected for protection. Cohesity uses this report to perform the first full backup. This report is also used to periodically reconcile the list of objects. The subsequent incremental backups are performed by using the AWS EventBridge capability.

With this protection approach, Cohesity can back up multi-billion objects at a faster rate.

Protect Your Amazon S3 Buckets

Cohesity utilizes the Amazon S3 inventory report to protect the Amazon S3 bucket. When you [register the AWS account](#), you can specify the Amazon S3 bucket where you want to create the inventory report. The S3 bucket you specify must be within the same AWS account and cloud region as the Amazon S3 bucket selected for protection.

When you initiate the Amazon S3 bucket protection, AWS will create an inventory report. It may take up to 48 to 72 hours for AWS to create an inventory report. This inventory report will contain the list of all the objects available on the Amazon S3 bucket you selected for protection. Cohesity uses this inventory report to perform the first full backup of the Amazon S3 bucket. Once Cohesity performs the first full backup, Cohesity utilizes AWS EventBridge and SQS queues to perform incremental backups of the Amazon S3 bucket.

With this protection approach, Cohesity can back up multi-billion S3 objects at a faster rate.

Note:

- If you have already registered your AWS account to protect AWS RDS or AWS EC2 workloads, then you must [Update the Existing CloudFormation Template](#) to update the Cohesity permissions in your AWS account.
- You do not need to deploy a SaaS connection to protect Amazon S3 buckets.

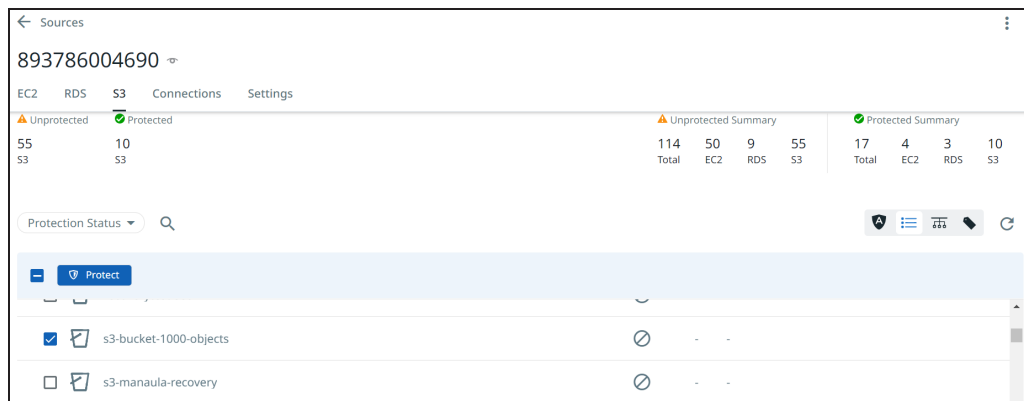
Add Protection to Your Registered Amazon S3 Buckets

Note: Before protecting your Amazon S3 bucket, ensure you have met the [prerequisites and understood the considerations](#).

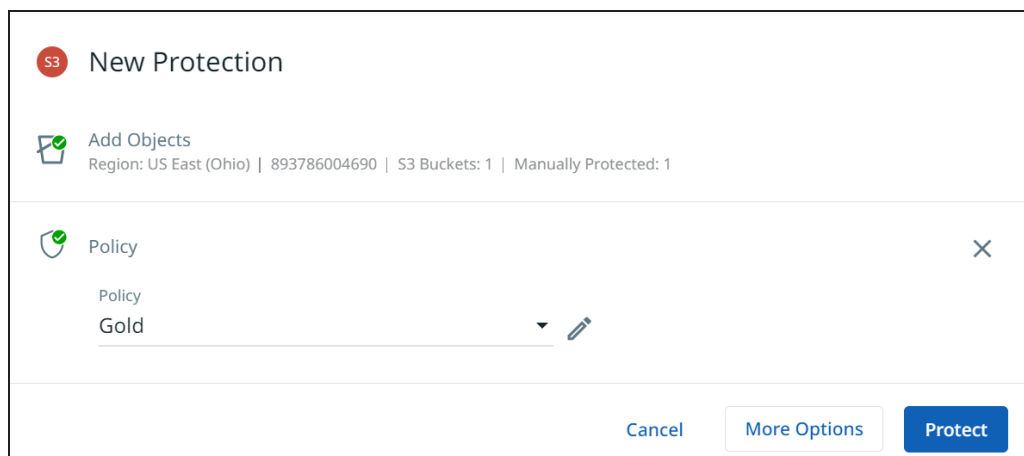
To protect your Amazon S3 buckets:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Find the registered AWS account and click into it.
3. Click the **S3** tab.

- Use the checkboxes to select the objects for protection. To protect the whole source, click the checkbox above the column.
- Click the **Protect** icon above the checkboxes.



- Choose a policy to specify backup frequency and retention. If you don't have a policy, you can easily create one. Ensure that the backup frequency set in the policy for S3 bucket is at most weekly. Backups may fail if the frequency is set for more than a week.
- If you wish to change or configure any of the additional settings, select **More Options** and perform the below steps or else, click **Protect**.



- Under **Settings**, edit the following options if necessary:
 - Start time:** Indicates what time the protection run should start. Enter the **Start Time** and select **AM** or **PM**. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone here.
 - SLA:** Defines how long the administrator expects a protection run to take. Enter:

- **Full.** The number of minutes you expect a full protection run, which captures all the blocks in an object, to take.
- **Incremental.** The number of minutes you expect an incremental protection run, which captures only the changed blocks in an object, to take.

9. Under **Additional Settings**, configure the following option:

- **Skip File on Errors:** Enable this option to continue the protection run even if any error is encountered when backing an S3 object. By default, this option is enabled. By disabling this option, the protection run will fail if one of the objects in S3 encounters an error.
- **Enable ACL Backups:** Enable this option to backup ACL. You can backup ACLs only if ACLs are configured in the S3 bucket you chose to protect. By default, this option is disabled.

10. Click **Protect**.

Cohesity starts backing up the Amazon S3 buckets you selected. You can monitor the status of the backup in the **Activity** page.

Protect an S3 Bucket Located in a Different AWS Region

If you want to protect an Amazon S3 bucket located in a different AWS region from where the inventory report's S3 bucket is located, perform the following steps:

1. Add the region of the Amazon S3 bucket you want to protect as a new region to store your backup. For more information, see [Select Regions and Encryption Key Management System](#).
2. [Re-register](#) the AWS Account with the following details:
 - a. Specify the region of the Amazon S3 bucket you want to protect as the **Destination cloud region**.
 - b. The S3 bucket you specify for creating the inventory report must be in the same region as the S3 bucket you want to protect.

Next > When the first protection run completes, you will be ready to [recover your protected Amazon S3 buckets](#) if and when you need to.

Manage Existing Protection


Edit protection settings, change the policy, and start, stop, & pause protection.

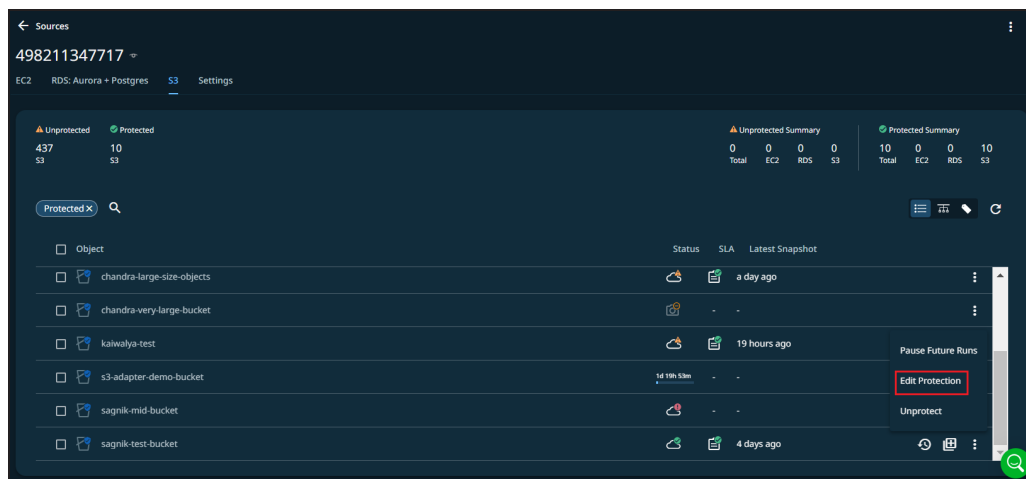
Once you have [applied protection](#) to the objects in your sources, Cohesity DataProtect as a Service makes it easy to make changes to that protection quickly. You can:

- Edit additional settings like **Start Time**, **SLA**, and more.
- Apply a different policy.
- Start an on-demand protection run, pause and resume it, or even remove protection.

Edit Protection Settings

To edit protection settings:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click into the **Source** name.
3. Select **Show All > Protected** and use the other filters, search box, and views at the top to narrow your search.
4. Click the **Actions** menu () next to the object and select **Edit Protection** to open the protection settings for that object.



Apply a New Protection Policy

To change the **Policy**, click the drop-down and select a different policy. To help you choose, each policy in the list shows the **Backup** frequency and the **Retain** period for each backup.

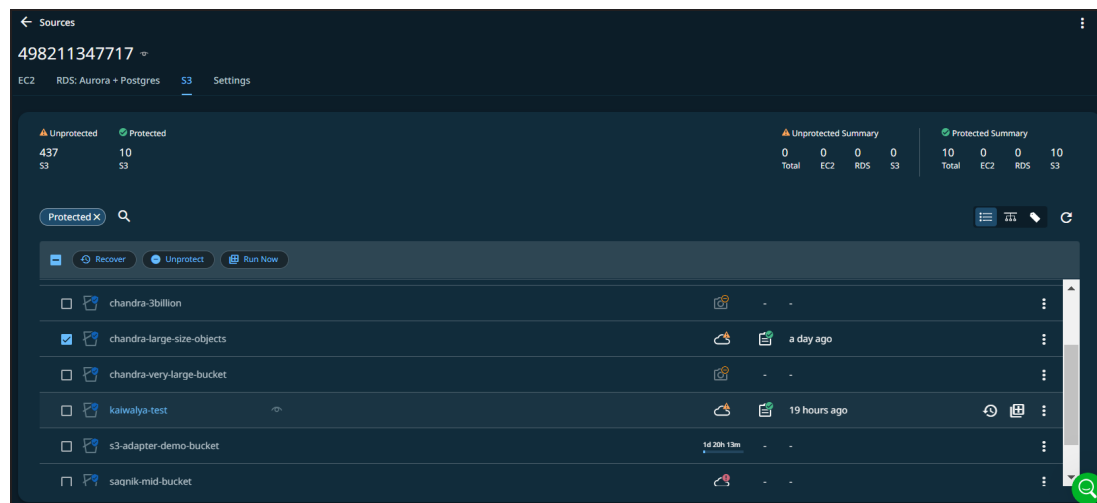
If you don't have a policy that meets your needs, scroll to the bottom of the list and click **Create Policy** to [create your own policy](#).

Edit Additional Protection Settings

1. Under **Settings**, you can change the protection **Start Time** (and select the **Time Zone**).
2. Click the drop-down next to Additional Settings to change more options such as **Skip File on Errors** and **ACL Backups**.

Start, Stop, or Remove Protection

When you select protected objects in one of your sources, Cohesity DataProtect as a Service presents buttons for the actions that are possible for those objects.



With the protected objects selected, you can click:

- **Recover** to recover the Amazon S3 bucket.
- **Unprotect** to remove Amazon S3 bucket from protection.

Note: The notifications enabled for the protected S3 bucket will remain active even after deleting the protection. If required, you can disable these notifications.

Tip: If a protected object is deleted from the source, you can search the object using Global Search and unprotect it.

- **Run Now** to start an on-demand protection run immediately.

Recover Your Amazon S3 Buckets

After you [protect your Amazon S3 buckets](#), you can recover them to their [original location](#) or a [new location](#) using Cohesity DataProtect as a Service.

Recover Amazon S3 to Original Location

To recover your protected Amazon S3 buckets to their original location:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the **Source** name.
3. Select **Protection Status > Protected**.

4. Use the filters, search box, and views to locate and select the S3 buckets you want to recover.

Tip: You can also use Global Search to locate, filter, and select the objects you need. Click the **Global Search** box at the top or type **slash (/)** anywhere to start your search.

5. Click the **Recover** icon at the top to open the **New Recovery** form. By default, the **Latest** snapshot is pre-selected for recovery. If you need to recover from an earlier snapshot, click the **Edit** (pencil) icon to choose the desired snapshot. You can also select any point from the green solid line on the slider if you want to restore to a specific point in time. Selecting an invalid time from the slider automatically selects the closest available snapshot.

Edit recovery point for auto-dmaas-s3-ga-bucket-1612

Timeline **List**

Dec 18, 2023 – Jan 17, 2024 ✕

Snapshot Time	Backup Type	Location
<input checked="" type="radio"/> Jan 6, 2024 3:23pm	Full	☁
<input type="radio"/> Jan 5, 2024 3:23pm	Full	☁
<input type="radio"/> Jan 4, 2024 3:23pm	Full	☁

Items per page 50 1 - 3 of 3 < >

Cancel **Select Recovery Point**

Task Number: Recover 2x3 Jan 12, 2024 11:12 AM

6. Under **Recover To**, select **Original Location**.
7. Enable **Overwrite Existing Bucket** if you want to recover the S3 bucket by overwriting the original S3 bucket. By default, this option is disabled.

If the data in the source S3 object and target S3 are the same, then the target S3 will not be updated even if the **Overwrite Existing Bucket** option is enabled. If the target bucket is versioned and the S3 object name of the target and the source match, then a new version of the S3 object will be created.
8. In the **Object Prefix** field, add prefix value to the names of the S3 objects that are recovered by this task. For information about the S3 object naming convention, see [AWS documentation](#).
9. Select your **Recovery Options**:
 - **S3 Prefixes to Recover:** To perform granular recovery the AWS S3 bucket, enable the **Turn ON to recover specific prefixes** option and provide the prefix of the objects to be recovered.

For example, `Finance/statement/document`

You can click the **add** icon to provide multiple prefixes.

The maximum number of prefixes that you can provide per recovery task is 100.

- **Continue on Error:** Enable this option if you want to continue the recovery even if one of the S3 objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.
- **Task Name:** Change the default name of the recovery task.

10. Click **Recover**.

The screenshot shows the 'Amazon S3' configuration window. At the top, it displays 'kaiwalya-test-tagged-bucket' (S3 Bucket), 'Latest' (Snapshot), and 'AWS US East (Ohio)' (Location). Below this, the 'Recover To' section has 'Original Location' selected, 'Overwrite Existing Bucket' is disabled, and 'Object Prefix' is 'new_'. The 'Recovery Options' section shows 'S3 Prefixes to Recover' with a toggle 'Turn ON to recover specific prefixes' which is turned on. Below the toggle, 'Finance/statement/document' is entered as a prefix. The 'Continue on Error' option is set to 'Yes'. The 'Task Name' is 'Recover_S3_bucket'. At the bottom, there are 'Recover' and 'Cancel' buttons.

Cohesity DataProtect as a Service begins to restore the selected Amazon S3 bucket.

Recover Amazon S3 to New Location

To recover your protected Amazon S3 buckets to a new location:

1. In **DataProtect as a Service**, navigate to **Sources**.
2. Click the **Source** name.
3. Select **Protection Status > Protected**.
4. Use the filters, search box, and views to locate and select the Amazon S3 you want to recover.

Tip: You can also use Global Search to locate, filter, and select the objects you need. Click the **Global Search** box at the top or type **slash (/)** anywhere to start your search.

5. Click the **Recover** icon at the top to open the **New Recovery** form. By default, the **Latest** snapshot is pre-selected for recovery. If you need to recover from an earlier snapshot, click the **Edit** (pencil) icon to choose the desired snapshot. You can also select any point from the green solid line on the slider if you want to restore to a specific point in time. Selecting an invalid time from the slider automatically selects the closest available snapshot.
6. Under **Recover To**, select **New Location** and provide the following information:
 - **AWS Account:** Select a registered AWS account as the new recovery destination.
 - **Region:** Select a destination AWS region.

Note: The target bucket must be in the same cloud region where your data is backed up (Cohesity-managed SaaS platform).

- **S3 Bucket:** Select the S3 bucket on to which you want to recover.
7. Enable **Overwrite Existing Bucket** if you want to recover the S3 bucket by overwriting the original S3 bucket. By default, this option is disabled.

If the data in the source S3 object and target S3 are the same, then the target S3 will not be updated even if the **Overwrite Existing Bucket** option is enabled. If the target bucket is versioned and the S3 object name of the target and the source match, then a new version of the S3 object will be created.
 8. In the **Object Prefix** field, add prefix value to the names of the S3 objects that are recovered by this task. For information about the S3 object naming convention, see [AWS documentation](#).
 9. Select your **Recovery Options**:
 - **S3 Prefixes to Recover:** To perform granular recovery the AWS S3 bucket, enable the **Turn ON to recover specific prefixes** option and provide the prefix of the objects to be recovered.

For example, `Finance/statement/document`

You can click the **add** icon to provide multiple prefixes.

The maximum number of prefixes that you can provide per recovery task is 100.

- **Continue on Error:** Enable this option if you want to continue the recovery even if one of the S3 objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.
- **Task Name:** Change the default name of the recovery task.

10. Click **Recover**.

The screenshot shows the Amazon S3 console interface for a bucket named 's3-adapter-demo-bucket'. The 'Recover To' section is active, showing options to recover to the 'Original Location' or 'New Location' (selected). The 'Location' section displays the AWS Account (498211347717), Region (AWS US East (Ohio)), and S3 Buckets (auto-dmaas-inventory-report-bucket-dev4-2201). The 'Overwrite Existing Bucket' toggle is off, and the 'Object Prefix' is 'new_'. The 'Recovery Options' section shows 'S3 Prefixes to Recover' with a toggle 'Turn ON to recover specific prefixes' (checked) and a list containing 'Finance/statement/document'. The 'Continue on Error' option is set to 'Yes', and the 'Task Name' is 'Recover_S3'. At the bottom, there are 'Recover' and 'Cancel' buttons.

Cohesity DataProtect as a Service begins to restore the selected Amazon S3.

Databases

Cohesity DataProtect as a Service unifies fragmented data protection solutions for databases.

Microsoft SQL Server

Cohesity DataProtect as a Service provides a simple, fast, cost-effective backup, recovery, and data management solution for growing Microsoft SQL Server database environments.

Requirements for Microsoft SQL Server

To register Microsoft SQL Server sources, ensure you meet the version and permission requirements, then download and install the Cohesity Agent.

Before you [register your Microsoft SQL Server source](#), ensure that you have the supported Microsoft SQL Server deployments. For more information, see [Supported Software for Cohesity DataProtect as a Service](#).

Also, make sure you meet the minimum permissions below and then [install the Cohesity Agent](#) on each SQL server you wish to protect.

Note: For information on the supported cloud regions where you can back up this source, see [Supported Workloads and Cloud Regions](#).

Minimum Permissions

To be able to register a Microsoft SQL Server source, you need to first install the Cohesity Agent on that source. To install the Cohesity Agent, you can use either the LOCAL SYSTEM account or an account that:

- Is a member of the local Windows Administrators group. For example, if qa01\tme-backup is an Active Directory user account in the data center that the backup admin plans to use, qa01\tme-backup must be part of the local Windows Administrators group on the SQL server.
- Has [Log on as a service](#) in the **User Rights Assignment** on the Microsoft SQL Server to install the CohesityAgent.
- Has the **sysadmin** role in the Microsoft SQL Server instance for transaction (T-log) log backup requirements. The **sysadmin** role is a Microsoft requirement that allows third-party solutions to back up transaction logs (T-logs) for full and bulk-logged recovery model databases.

Check Firewall Ports

Ensure that the ports listed in the Microsoft SQL Servers section in the [Firewall Ports for User-Deployed SaaS Connectors](#) topic are open to allow communication between the Cohesity SaaS Connector(s) and Microsoft SQL Server.

Considerations

- Windows 2008 R2 Operating System is not supported for SQL Server protection on Cohesity DataProtect as a Service.

Download and Install the Agent

Install the Cohesity Agent on each SQL server that you want to protect.

To install the Cohesity Agent:

1. In **DataProtect as a Service**, navigate to **Sources** and select **Register Source > Microsoft SQL Server**.
2. At the bottom of the **Register a Microsoft SQL Server** dialog, click **Download Cohesity Agent**. Make sure you download the Agent on the server you plan to protect.
3. As an administrator with local system privileges on that server, run the executable and complete the installation wizard.
4. **Service Account Credentials**. Enter either the LOCAL SYSTEM account credentials or an account that meets the [minimum permissions](#) above.
5. Wait for the Agent installation to finish. In SQL Server Management Studio (SSMS), validate that the account used to install the Cohesity Agent has **SQL Server Role: sysadmin** in the SQL server instances.
6. The Agent starts automatically.

Repeat the Agent installation process on each SQL server you want to protect. This includes any standalone Microsoft SQL Servers and Microsoft SQL Server nodes with AGs.

Next > [Register your Microsoft SQL Server source](#) to protect your databases!

Microsoft SQL Server on Linux OS

Cohesity provides backup and recovery solutions for the Microsoft SQL Server databases running on the Linux Operating System.

Requirements

Operating System	Database	Notes
Linux OS RHEL 8 and higher versions	Microsoft SQL Server 2019	<ul style="list-style-type: none"> Only VDI-based backups are currently supported on Linux OS. High Availability Microsoft SQL Server configurations like FCI and AG are not supported.

Note: This is an Early Access feature. Contact [Cohesity Support](#) to enable the feature.

Considerations

- Cross-platform database restores between Windows and Linux are currently not supported.
- For VDI-based backups to work, Microsoft requires the backup application account (Cohesity Agent's service account) and the Microsoft SQL Server service account to be added to each other's primary groups using the usermod command.

For example:

```
sudo usermod -a -G mssql cohesityagent
```

```
sudo usermod -a -G cohesityagent mssql
```

where, mssql is the user account of the mssql-server service and cohesityagent is the user account of the Cohesity Agent service.

Following this, the Microsoft SQL Server and the Cohesity Agent services must be restarted for the changes to take effect.

Note: The usermod -a -G command may not work for domain accounts.

- Currently, only password-based SQL Server authentication is supported and Active Directory-based authentication is not supported for Linux.

For supported features of SQL Server 2019 on Linux, see [Editions and supported features of SQL Server 2019 on Linux](#).

Register SQL Server on Linux OS

- In **DataProtect as a Service**, select **Sources**.
- Click **+Register Source > Microsoft SQL > Start Registration**.

- In the **Register Microsoft SQL Server** page, select the **SaaS Connection** and click **Continue**.
- In the **Hostname or IP Address** field, enter the IP address of the server, FQDN of the server, or VIP of the SQL FCI. Cohesity recommends you to provide the FQDN of the server.
- Select the **Server Type** as Linux.
- Provide the username and password for the **Database Authentication**.
- Click **Complete**.

Register Microsoft SQL Server Sources

To start protecting a Microsoft SQL Server database, once you meet the [Requirements for Microsoft SQL Server](#), you need to register the SQL Server as a source.

Note: To connect with sources in your data center, you'll need to use a SaaS Connection (or [create one](#)) to establish connectivity between the sources and the Cohesity DataProtect as a Service.

To register a Microsoft SQL Server, check that it meets the prerequisites below and then [add it as a source in DataProtect](#).

Prerequisites

- Verify Microsoft SQL Server services are running.
- On the server's Windows system, set the **Power Plan** to **High performance**.
- On the SQL Server where you have installed the Cohesity Agent, open the following ports:
 - **50051**, for backup operations (incoming).
 - **11113** and **11117**, for VDI-based backup and restore (outgoing).

Note:

For SQL running in an Amazon EC2 instance, add **inbound rules** to the EC2 and SaaS connector(s) security groups, to allow the backup and recovery of SQL Server.

- For **SaaS Connector(s)** (source) to **EC2 instance** (destination), use TCP and Port **50051**.
- For **EC2 instance** (source) to **SaaS Connector(s)** (destination), use TCP and Ports **11113** and **11117**.

- If you're using the Windows Firewall, set:
 - **Inbound rules:**
 - Add a rule to accept SQL Server traffic and TCP connections on local port 1433.
 - Set **Remote Port** to **All Ports**.
 - **Outbound rules** (for Microsoft SQL Server 2016 running on Windows 2016):
 - Update the 'Block network access for R local user accounts in SQL server instance MSSQLSERVER' rule by navigating to **General > Action window** and selecting **Allow the connection**.

Register a Microsoft SQL Server Source

To add a Microsoft SQL Server as a Cohesity DataProtect as a Service source:

1. Confirm that you meet the Microsoft SQL Server requirements for software version and user account minimum permissions.
2. In **DataProtect as a Service**, navigate to **Sources** and click **+ Register Source**.
3. In the **Select Source** dialog box, select **Microsoft SQL** and click **Start Registration**.
4. In the **Register Microsoft SQL Server** dialog box, select an existing healthy SaaS connection marked *Unused*, or click **Create SaaS Connection** and follow the instructions in [Create a SaaS Connection](#), and then click **Continue**.

Register Microsoft SQL Server

1 SaaS Connection — 2

Cohesity DataProtect service uses SaaS connections to protect your on-premise or cloud-based data sources. A SaaS Connection consists of one or more SaaS Connectors, which are data mover VMs deployed in your environment. [Learn more.](#)

SaaS Connection
Saas_Connection-1

Cancel Continue

Note: Ensure that the agent is installed on all the nodes that are part of the Always on Availability Group (AG) and Failover Cluster Instance (FCI). For more details, [Download and Install the Agent.](#)

5. Enter the Microsoft SQL Server **Hostname** or **IP Address**, the FQDN of the server, or the VIP of the SQL FCI.
6. Click **Complete**. Cohesity auto-discovers the entire Microsoft SQL Server topology on the Windows cluster.
7. From the topology list, select **Register all MSSQL Nodes** to register the Microsoft SQL Server nodes as individual Microsoft SQL Server sources. For AG and FCI, ensure to register all the nodes that are part of the AG and FCI.
8. Click **Complete Registration**.

Next > You are now ready to [protect your SQL databases](#).

Protect Microsoft SQL Server Databases

Once you have [Registered a Microsoft SQL Server](#) as a source, you're ready to protect the Microsoft SQL Server databases on that server.

To protect your Microsoft SQL Server databases:

1. In **DataProtect as a Service**, under **Sources**, find the Microsoft SQL Server source, click the **Actions** menu (:), and select **Protect**.
2. Click **Add Objects**. Browse through the SQL Server instances and select the databases that you want to protect including MS SQL standalone, FCI, and AG. Click **Continue**.

Note: AG and FCI are represented using the  icon.



For AG, the AG Replica details like Hostname, Server Name, Primary Role Allow Connections, Secondary Role Allow Connections, and Role are available in the **Settings** tab.



If appropriate for your organization, enable Auto Protect.

Click  to turn on **Auto Protect**. By enabling this option:

- All the SQL databases that are added to the SQL server in the future are automatically added to the Protection Group and are protected from the next protection run.
- All the SQL databases that are removed from the SQL server in the future are automatically removed from the Protection Group and are not protected from the next protection run. The existing backup for the user is preserved until the snapshot expires.

For example, you create a Protection Group that backs up a SQL Server. The Protection Group runs for two weeks and then a new database is added to the SQL Server. The next time the Protection Group runs, if the object hierarchy has been refreshed on the Cohesity cluster, the new database is also backed up even though the new database has not been explicitly selected to be included in the Protection Group. The object hierarchy is automatically refreshed every four hours.

To manually refresh the object hierarchy, select **Data Protect as a Service > Sources**, find the source in the list and click . To exclude an object from Auto Protect, click .

-  Indicates an Auto Protected object.
-  Use this option to exclude SQL databases from the protection group. Excluded objects are not protected. The icon indicates that an ancestor of the object is Auto Protected but this object is explicitly excluded.

To exclude a database from the protection group, specify the absolute path of the database to be excluded from the backup at a protection job level. You can optionally use:

- **Wildcard Exclusion** - You can optionally use the Wildcard * in any location in the search string to exclude all the matching databases from the backup at a protection job level.

- **Regex Exclusion** - You can use the RegEx string name match and manual user selection to exclude DBs from backup at a protection job level.

Note: You can use the exclude SQL database feature only if Auto Protect is enabled for a backup job.

Protection of dismounted databases, recovery database types, and deleted databases are automatically skipped by the Protection Group.

After defining the filter click **Apply** to save the selection.

The following table lists the examples for excluding objects in a SQL database:

Sample Configuration

10.2.157.20 (Host)

* SQL2012	* SQL2014	* MSSQLSERVER
(SQL Server Instance)	(SQL Server Instance)	(SQL Server Instance)
- UserDB1	- UserDB3	- ProdDB1
- UserDB2	- UserDB4	- ProdDB2
- TestDB1	- TestDB3	
- TestDB2	- TestDB4	

Example of Absolute Path

- **10.2.157.20/SQL2012/UserDB1** - Excludes the database 'UserDB1' from the SQL server instance 'SQL2012' on the host 10.2.157.20.
- **10.2.157.20/SQL2012/** - Excludes the 'SQL2012' from the SQL server instance and its databases.

Note: The absolute instance path must end with '/' and the plain text instance path must contain two or more slashes.

Example of Wildcard Exclusion

- **10.2.157.20/SQL2012/User*** - Excludes the databases 'UserDB1' and 'UserDB2' from the SQL server instance 'SQL2012' on the host 10.2.157.20.
- **10.2.157.20/SQL201?/*** - Excludes the SQL server instances 'SQL2012' and 'SQL2014' from the host 10.2.157.20.

Example of Regex Exclusion

10.2.157.20/SQL.*/* - Excludes the SQL server instances 'SQL2012' and 'SQL2014' from the host 10.2.157.20.

Considerations

Consider the following when using the **Auto Protect** option on the backup job

- If new databases are added on the host, the existing filters are applied on the newly added databases. For example, if the exclude filter matches the newly added databases then these databases will be excluded in the backup job.
 - If any system databases are excluded by the filters in a SQL server, it will exclude all the system databases belonging to that SQL server.
 - When defining filters to exclude AG databases, ensure that the filters are configured on all the AG hosts. For example, to exclude an AG database and if the AG database has 4 hosts, then ensure that the filters are defined on all the four AG hosts.
3. Choose a policy to specify backup frequency and retention.* If you don't have a policy, you can easily [create one](#).

If you select a policy that includes Log Backup (Database), the Log Backup Settings option will be enabled.

4. Click **More Options** and review the following settings:

Microsoft SQL Settings

- **Make Full Backups Copy-only.** Enable if you want full backups to be copy-only backups so they do not affect the differential base. Note that copy-only full backups do not take log backups even if the policy schedules them.
- **WITH Clause.** Define the WITH clause that you want to use to customize the backup. For more information, see [BACKUP \(Transact-SQL\)](#) in the Microsoft documentation.

The WITH clause text box allows you to invoke the SQL Server general options during backup and recovery. The text field accepts any SQL Server general backup or recovery option like "ENCRYPTION". Note that specifying an option could add to the overall duration of the backup or restore process.

This is the full/incremental WITH clause with the default value *"WITH MAXTRANSFERSIZE = 4194304, BUFFERCOUNT = 64"*.

- **Number of Streams.** Define the number of **.bak** files you want to create for better backup performance. By default, Cohesity DataProtect as a Service creates three **.bak** files for each database backup for better backup performance.
- **User Databases.** Select how to handle AG databases during backup.
- **AG Backup Preferences.** Select this option if AG databases will be backed up.

- **Use Server Preferences** uses Microsoft SQL preferences.
- **Override Preferences** enables you to override Microsoft SQL preferences with your selection.

AG replica preference for Microsoft SQL backups

Cohesity uses "replica priority" to select the best replica when more than one qualified replica matches the backup preference. For the following AG Backup Preference settings, Cohesity uses replicas in the indicated order of preference to back up Microsoft SQL databases:

Backup Preference Setting	Replica Used for MS SQL Backup
Prefer Secondary or Any	1. Sync Secondary Replica 2. Async Secondary Replica 3. Primary Replica
Secondary Only*	1. Sync Secondary Replica 2. Async Secondary Replica
Primary Only	1. Primary Replica

* If the AG uses the **Secondary Only** Backup Preference setting, ensure the AG replicas are set to "**Readable secondary=Yes**" or "**Readable secondary=Read-Intent**".

You can set **Readable Secondary** field to **No** if the Backup Preference Setting is set to **Prefer Secondary or Any** or **Primary Only**.

Differential backup can be performed only when the Backup Preference is set to **Primary Only**.

Note: You must apply the exclusion filter on each AG replica to ensure the database is excluded when a SQL Server failover occurs.

- **System Databases.** Select whether to back up or skip system databases.
- **Databases to Backup.** Select the User Databases and System Databases. For AG, select the **AG Backup Preferences**. You can **Use Server Preferences** or **Override Preferences** (the options include Primary Only, Secondary Only, Preferred Secondary, and Any).

Log Backup Settings

Cohesity enhances the SQL Server native T-log backup and restoration performance by splitting the T-log backup into multiple backup files and ingesting them in parallel through the Cohesity nodes.

- **WITH Clause.** Define the WITH clause that you want to use to customize the backup. For more information, see BACKUP (Transact-SQL) in the Microsoft documentation.

The WITH clause text box allows you to invoke the SQL Server general options during backups and restores. The text field accepts any SQL Server general backup or restore option like "ENCRYPTION". Note that specifying an option could add to the overall duration of the backup or restore process.

The default WITH clause is "WITH MAXTRANSFERSIZE = 4194304, BUFFERCOUNT = 16".

- **Number of Streams.** Define the number of **.bak** files you want to create for better backup performance. The default number of streams for log backup is 1.

5. Click **Protect**.

Note: The backups start immediately after you protect the objects, regardless of the time you set for the protection run.

Cohesity DataProtect as a Service starts backing up the databases you selected.

When choosing or configuring your policy, ensure the full, incremental (SQL Differential), and T-Log backup retention periods are properly configured. The retention period requirements for SQL VDI are identical to those for SQL native backups. For example, we recommend aligning your retention periods for each backup type along these lines:

- **Full Backups.** Daily at 1 AM with a 7-day retention.
- **Incremental Backups** (equivalent to SQL Differential backups). Every 12 hours with a 3-day retention.
- **T-Log Backups.** Every 15 minutes with a 1-day retention.

Note: The following error message is displayed when the Physical SQL Server source includes manually protected AG databases: "*Protected objects are missing from the source. This might lead to backup failures.*"

You can click **View and unprotect the objects** or use the search box to filter the objects with status, **Unavailable**, and unprotect them.

Next > When the first protection run completes, you will be ready to [recover your protected databases](#) when and if you need to.

Troubleshoot

- The log back up of AG MS SQL Server fails with the following error if there is a break in the log chain.

```
Log chain break error: Discovered a break in the log chain for  
<Database Names>
```

Resolution: To resolve this issue, ensure that no other third-party applications are running a log backup, and then perform a full backup run to reset the log chain.

- The log back up of the AG MS SQL Server fails with the following error if a database is added or removed from the AG MS SQL Server.

```
AG relationship error: Discovered a AG relationship error for database  
<Database Names>
```

Resolution: To resolve this error, perform a full backup run.

For more troubleshooting information, see the following KB article:

- [Collecting troubleshooting information for MS SQL issues](#)

Log in to the [Cohesity Support Portal](#) to see more KB articles.

Protect Microsoft SQL Server Running in AWS and Azure

Cohesity supports protecting the Microsoft SQL Servers (Standalone, Failover Clusters, and Availability Groups) running in the following:

- [AWS EC2 Instances](#)
- [Azure VMs](#)

Protect Microsoft SQL Server Running in AWS EC2 Instances

The following is the procedure to protect the Microsoft SQL Servers running in the AWS EC2 instances:

- [Register the AWS Account.](#)
- [Deploy AWS SaaS Connectors.](#)
- [Register the Microsoft SQL Server Using the SaaS Connectors.](#)
- [Protect Microsoft SQL Server Databases.](#)
- [Recover Microsoft SQL Server Databases.](#)

Protect Microsoft SQL Server Running in Azure VM

The following is the procedure to protect the Microsoft SQL Servers running in the Azure VMs:

- Register the Azure Subscription.
- Deploy Azure SaaS Connectors.
- Register the Microsoft SQL Server Using the SaaS Connectors.
- Protect Microsoft SQL Server Databases.
- Recover Microsoft SQL Server Databases.

Recover Microsoft SQL Server Databases

After you [protect your Microsoft SQL Server databases](#), you can recover them from Cohesity DataProtect as a Service, to their original or a new location.

To recover protected Microsoft SQL Server databases:

1. In **DataProtect as a Service**, go to **Sources** to set up your recovery task.
2. Click on the **Source** name.
3. Above the tree, select **Show All > Protected**.
4. Use the filters, search box, and views to locate and select the SQL databases you need including MS MSQL standalone, FCI, and AG.

Tip: You can also use Global Search to locate, filter, and select the objects you need. Click the **Global Search** box at the top or type **slash (/)** anywhere to start your search.

5. Click **Recover** at the top to open the **New Recovery** form with the **Latest** snapshot (protection run).

If you need to recover from an earlier snapshot, click the **Edit** icon to open the **Recovery Point** calendar. Click **List** to view the available recovery points by timestamp and click one.

- Click **Select Recovery Point**.
 - Click **Next: Recover Options** to return to the form.
6. Under **Targets**, select **Recover as a new Database** or **Overwrite Original Database**. If you choose:
 - **Recover as a new Database**, select a registered **Microsoft SQL Server Instance** or **Restore to Original SQL Server Instance**.
 - **Overwrite Original Database**, Cohesity will overwrite the original SQL Server instance. Note that this is a destructive action that cannot be undone.
 7. If necessary, under **Database File Paths**, you can:

- Update the **Database Files** and **Log Files** paths.
 - Enter additional **File Path Rules**.
8. Select your **Recovery Options**:
- **Rename**. Choose whether to **Bulk Rename** with a **Suffix** or **Rename Individual Objects**.
 - **WITH RECOVERY**: By default, a Microsoft SQL Server restore WITH RECOVERY is performed. You can optionally toggle this off to perform a restore WITH NORECOVERY.
 - **Keep CDC**: Use this option to restore a backed-up database with the change data capture (CDC) enabled. By default, the Keep CDC switch is **ON**. If the backed-up database is not CDC enabled and the user tries to restore it with Keep CDC, the database will be restored without CDC.
 - **Overwrite Alternate Database**: Enable this option to recover the database by overwriting the existing database.
 - **WITH Clause for Full/Incremental Restore**: Specify the WITH clause that you want to use for the full/ incremental recovery.
 - **WITH Clause for Log Restore**: Specify the WITH clause that you want to use for the log recovery.
 - **Capture Tail Logs**: You can optionally choose to **Capture tail logs**. Tail logs capture records that have not yet been backed up. They are captured to ensure all transactions are backed up before restoring the database.
 - **Task Name**. Change the default name of the recovery task.
9. Click **Start Recovery**.

Oracle Database

Cohesity DataProtect as a Service provides a simple, fast, cost-effective backup, recovery, and data management solution for growing Oracle database environments.

Oracle Requirements

To register your Oracle servers and protect your databases, be sure you meet the requirements and install the Cohesity Agent on each server.

Before you register your Oracle servers to protect your Oracle Databases, confirm that you meet the software version, [prerequisites](#), [credentials](#), [choose an authentication method](#), and set [sudoers permissions](#) below, then [download and install the Cohesity Linux Agent for Oracle](#) on the servers you wish to protect.

Also, be sure to review the [limitations](#) at the end.

Note: For information on the supported cloud regions where you can back up this source, see [Supported Workloads and Cloud Regions](#).

Support Matrix

Cohesity DataProtect as a Service supports Oracle Database protection. For information on the supported Oracle versions, see [Supported Software for Cohesity DataProtect as a Service](#).

Check Firewall Ports

Ensure that the ports listed in the Oracle Servers section in the [Firewall Ports for User-Deployed SaaS Connectors](#) topic are open to allow communication between the Cohesity SaaS Connector(s) and Oracle Server.

Prerequisites

Make sure the following prerequisites are met before you proceed with Oracle source registration:

- **UUIDs.** All the Oracle Databases that are protected using Cohesity DataProtect as a Service should have a unique UUID on the Oracle source where the databases reside.
- **Archive Log Mode.** Archive Log mode must be enabled for databases to be opened in Read-Write mode.
- **Read Only Mode:** The Oracle Databases should be opened in Read-Write mode.
- **Version.** The recovery source and target database must be the same Oracle database version. For example, snapshots of an 11g Oracle Database cannot be recovered to a 12c Oracle Database.
- **Oracle Single Instance Deployment.** For an Oracle single-instance database, the database must be entered into the `/etc/oratab` file. Otherwise, Cohesity DataProtect as a Service will not be able to discover this database.
- **Authentication.** If you choose DB authentication, all the databases on the system should have the same username and password or OS Authentication. At the backup level, they can have individual passwords for the databases.
- **Ports.** On the Oracle Server where you [install the Cohesity Linux Agent](#) (below), open the 50051 port for backup operations (incoming) and 59999 port for self-monitoring and debug pages.

Credentials and Privileges

Once you register your physical servers with Cohesity DataProtect as a Service as Oracle servers, Cohesity DataProtect as a Service will discover your Oracle databases automatically. For Cohesity DataProtect as a Service to successfully discover your Oracle databases, the user account running the Cohesity Linux Agent must have the appropriate credentials and privileges.

You can install the Cohesity Linux Agent to [run with the ROOT user](#) or [with a separate OS user](#) (also known as the 'OS Service Account user').

When connecting to Oracle databases, Cohesity DataProtect as a Service can use either the Oracle OS Authentication or [Oracle DB Authentication](#) method. These two types of Oracle authentication are available whether the Agent is run with the ROOT user or a separate OS Service Account user.

Note: While most Oracle operations are available using either OS or DB authentication, some specific operations specifically require one or the other. For details, see [Oracle Authentication Method Requirement](#) below.

Running Agent with ROOT User

You can install Cohesity's Linux Agent to run with the ROOT user. When you take this approach, the agent runs every command using the ROOT user, except for Oracle commands and utilities like RMAN or SQLPLUS. To run Oracle commands and utilities, the Agent will 'su' to the user who is the owner of the Oracle binary in the current Oracle Home. If an Oracle operation is run against a source database that has DB Authentication configured (where the user has previously configured DB credentials for this Oracle source database), DB Authentication will be used to run Oracle commands and utilities. Otherwise, OS Authentication via the Oracle binary owner will be used.

When you install the Cohesity Agent to run with the ROOT user, there is no need to configure additional SUDOERS privileges.

To start the service as a ROOT user, add the following permission to the sudoers file:

```
Defaults:<oracle_binary_user> !requiretty.
```

Running Agent with OS Service Account user

You can install Cohesity's Linux Agent to run with a specific OS Service Account user account, as long as it meets the following requirements:

- The OS user is automatically granted the required sudo privileges. This allows the Cohesity Agent to execute specific privileged commands. For details, see [Oracle Sudoers Permissions for Linux Databases](#) below.
- The OS user should be part of the OS group with SYSDBA or SYSBACKUP privileges (for example, dba).

You can run the Cohesity Agent as a different service user, the *cohesityagent* user, if this user is part of the OSDBA group in Oracle.

If you choose DB authentication, then all the databases on the system should have the same username and password.

If you wish to add the OS user to the Oracle Database as an OS-authenticated user, use the IDENTIFIED EXTERNALLY clause.

Oracle Authentication Method Requirement

You can either use either OS user or DB user authentication to connect to your Oracle Databases, but for recovery to *alternate* servers, you must use OS authentication.

Table: Available Oracle Operations by Authentication Method.

Oracle Operation	Authentication Method	Notes
Backup	OS Authentication or DB Authentication	None
Restore to Original Server (a.k.a. Overwrite Restore)	OS Authentication or DB Authentication	Restoring data to the same server overwrites the original database.
Restore to Alternate Server	OS Authentication	DB Recovery or Restore into a different server is available, assuming the Oracle binaries already exist and the target Oracle server has free space to store the newly created database files.

Oracle Sudoers Permissions for Linux Databases

The following tables list the sudoers permissions required for the Cohesity Linux Agent for Oracle.

Note: When you install the Cohesity Agent to run with the ROOT user, there is no need to configure additional SUDOERS privileges.

Operating System	Sudoers Permissions	Sudoers Permissions
	Cohesity Linux Agent Commands for both Oracle sources & Linux servers	Additional commands only for Linux servers

Operating System	Sudoers Permissions	Sudoers Permissions
Linux	<ul style="list-style-type: none"> • cp • chown • chmod • mkdir • rm • tee • hostname • stat • timeout • ls • rsync 	<ul style="list-style-type: none"> • blkid • lsof • losetup • dmsetup • lvs • vgs • lvcreate • lvremove • lvchange

Download and Install the Cohesity Agent

The Cohesity Linux Agent can be installed to [run as a ROOT user](#) or [as an OS Service Account user](#). Install the Cohesity Linux Agent on each Oracle server that you want to protect.

Cohesity Linux Agent Best Practices

We recommend you follow these best practices when you plan to deploy the Cohesity Linux Agent on Oracle servers and hosts:

- If you choose DB authentication, then all the databases on the system should have the same username and password.
- Create a database user for your Cohesity Oracle backup and restore workflows. *(Optional)*
- Both the Oracle host and the Cohesity Linux Agent should have permission to write to the `adump` and `diag` directories, control file, and the database restores locations.
- Enable Block Change Tracking (BCT) to improve the incremental backup performance of the Oracle server. *(Optional)*
- Assign sudoers to the user running the Cohesity Linux Agent.
- Make the Cohesity Linux Agent user part of the Oracle dba group.
- Given that Oracle Secure Backup (SBT)-based incremental backups are not fully hydrated (unlike imagecopy-based backups), we recommend you take a full database backup regularly.

Install the Cohesity Linux Agent to Run with ROOT User

To install the Cohesity Linux Agent to run as the ROOT user on your Oracle server:

1. In **DataProtect as a Service**, navigate to the **Sources** page and click **+ Register Source** in the upper-right corner of the page and then click **Oracle**.
2. Click **Start Registration**.
3. In the Register Physical dialog box, select an existing SaaS connection marked Unused or click Create SaaS Connection and follow the instructions in [Create a SaaS Connection](#), and then click **Continue**.
4. Click **Download Cohesity Agent**. Ensure the agent has been downloaded to the appropriate server.

Register Oracle Source

SaaS Connection

The source you are registering requires a SaaS Connection to communicate with Cohesity DataProtect. [Learn about setting up a SaaS Connection](#).

SaaS Connection

Source Details

i A Cohesity Agent needs to be pre-installed on all servers being registered. If this has not been completed, download the agent using the link below and install it before continuing the registration.

[Download Cohesity Agent](#)

Enter Host Address

This field is required

Authentication Type

☒ OS Authentication ☐ Database Authentication

Cancel Register

5. Run the executable file with sudo using the following command syntax:
`sudo /<path_to_installer_file> -- --install -c 0 -S root -G root`

The command options are:

- **-S:** The user to run the Agent. Specify 'root'.
- **-G:** The group permission the Agent will use for files and directories installed by the agent. Specify 'root'.
- **-c:** The boolean switch that controls whether the OS user and group should be created. '0' means do not create the OS user and group, and '1' means the Agent installation will create the specified OS user and group. (If you choose to run with the root user, specify '-c 0' as 'root' already exists.)

The Agent starts automatically after the installation, as well as on a subsequent Oracle host reboot.

At the end of the installation, the commands used to start, stop, or get Agent status are displayed for future reference.

Install the Cohesity Linux Agent to Run with OS Service Account User

To install the Cohesity Linux Agent to run as the OS Service Account user on the Oracle server:

1. In **DataProtect as a Service**, navigate to the **Sources** page and click **+ Register Source** in the upper-right corner of the page and then click **Oracle**.
2. Click **Start Registration**.
3. Click **Download Cohesity Agent**. Ensure the agent has been downloaded to the appropriate server.
4. Grant sudo permission to the user who will install the agent. This user must be part of the OS DBA group. For details, see [Credentials and Privileges](#) above.
 - If you plan to run the Oracle SQL commands as OS authenticated user, we recommend you perform the installation as the Oracle OS user. Even if the Cohesity Agent user is part of the DBA group, you can run the Oracle SQL commands.
 - Because restoring to alternate locations requires OS authentication, we recommend you use OS instead of DB authentication. The restore to alternate locations will succeed only if the Cohesity Agent is installed with **dba** or **oinstall** as the user group.
 - The Cohesity Agent installer grants sudo permission for the following commands:

```
/usr/bin/cp, /usr/bin/chown, /usr/bin/chmod, /usr/bin/mkdir,
/usr/bin/rm, /usr/bin/tee, /usr/bin/hostname, /usr/bin/stat,
/usr/sbin/blkid, /usr/sbin/lsof, /usr/bin/ls, /usr/sbin/losetup,
/usr/sbin/dmsetup, /usr/bin/rsync, /usr/bin/timeout,
/usr/sbin/lvs, /usr/sbin/vgs,
/usr/sbin/lvcreate, /usr/sbin/lvremove, /usr/sbin/lvchange
```

5. Copy the downloaded file to the target Oracle host and run the executable file as a sudo user using the following command syntax:

For script-based installer:

```
sudo /<path_to_installer_file> -- --install
```

For RPM-based installer:

```
sudo rpm -i path_to_install_file
```

The installer creates the user group, 'cohesity agent,' and installs the Agent.

The Agent starts automatically after the installation or on reboot.

Considerations

- **Oratab.** Only standalone databases listed in the `oratab` file on the Oracle server can be registered and protected. Cohesity DataProtect as a Service cannot discover databases that are not in `oratab`.
- **Auto Protect.** Auto Protect is not supported for Oracle databases.
- **Point-in-Time Restore.** During a point-in-time restore to a time near the end of a full backup, the restore might fail due to this [Oracle issue](#).

Next > [Register your Oracle servers!](#)

Register Oracle Sources

To start protecting your Oracle Databases, you need to register your Oracle servers and hosts as Cohesity DataProtect as a Service sources. Confirm you've met the Oracle requirements and then register your Oracle sources.

Note: To connect with sources in your data center, you'll need to use a SaaS Connection (or [create one](#)) to establish connectivity between the sources and the Cohesity DataProtect as a Service.

To register an Oracle Server as a Cohesity DataProtect as a Service source:

1. Confirm that you meet the [Oracle requirements](#) for software version and the required credentials and privileges.
2. In **DataProtect as a Service**, navigate to the **Sources** page and click **+ Register Source** in the upper-right corner of the page and then select **Oracle**.
3. From the SaaS selection drop-down, choose the **Existing Connection** and select one that is marked **Healthy**, or click **Create SaaS Connection** and follow the instructions in [Create a SaaS Connection](#).

4. Enter the **Hostname (FQDN)** or **IP address** of the Oracle server you're registering. We recommend that you use the FQDN.
5. Choose your [Oracle authentication method](#): **OS Authentication** (the default) or **DB Authentication**.

Note: If you choose DB authentication, then all the databases on the system should have the same username and password.

6. Click **Register**.

Your Oracle server appears under **Sources** in Cohesity DataProtect as a Service.

Next > You're ready to [protect your Oracle Databases](#)!

Protect Oracle Databases

Once you have [registered an Oracle server](#) as a source, you're ready to protect the Oracle Databases on that server.

To protect your Oracle Databases:

1. In **DataProtect as a Service**, navigate to **Sources**, find the Oracle source name and then click **Protect**.
2. Click **Add Objects**. Browse through the Oracle server instances and select the databases that you want to protect. Click **Continue**.

3. Click the **Edit** (pencil) icon next to the selected object and select one of the following options:
 - **System selects active node.** Cohesity auto-selects an active single-instance Oracle node and configures the number of RMAN channels for the database object. (*Default*)
 - **Select specific node(s).** If you select this option, you can choose the number of RMAN channels and the SBT library path to be used for the database object.
 - **Delete Archive Log.** Toggle on and specify the days after which the archived logs on the source database should be deleted. If you enter a value of "0" days, source archived logs will be deleted immediately after each successful protection run.

Important: If you do *not* enable this option, Cohesity will not delete the archived logs after each protection run and you are responsible for deleting the archived logs from the source server.

4. Click **Save**.
5. In the **New Protection** dialog, select a **Policy** that matches the schedule and retention period you need. If the existing policies do not meet your needs, you can [create a new policy](#) with the settings you need.

Note: The Oracle adapter for the CCS SaaS solution relies on the SBT library to back up and recover. This requires occasional full backups. Please consider this while configuring the backup policies for the Oracle database.

6. If you wish to configure a specific **End Date, Alerts**, and other additional settings, click [Additional Settings](#).
7. Click **Protect**.

Note: The backups start immediately after you protect the objects, regardless of the time you set for the protection run.

Additional Settings

Advance Settings	Description
Start Time	Available only if the selected policy is set to Backup Daily . Indicates what time the protection run should start. Enter the Start Time and select AM or PM . The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone here.
End Date	If you need to end protection on a specific date, enable this to select the date.
Priority	Select a priority for the protection task execution. Cohesity supports concurrent backups, but if the number of tasks exceeds the ability to process them, they are executed in this priority order: <ol style="list-style-type: none"> 1. High-priority tasks 2. Medium-priority tasks 3. Low-priority tasks
Alerts	Click to enable one or more of these alert types to trigger alerts for the following events and click Add to enter email addresses. <ul style="list-style-type: none"> • SLA Violation. Creates warning alert when a protection run exceeds the configured SLA. Sends email. • Failure. Creates critical alert when object protection fails to complete. Sends email. • Success. Creates information alert when object protection completes. Does not send email.
SLA	The service-level agreement (SLA) defines how long the administrator expects a protection run to take. Enter: <ul style="list-style-type: none"> • Full. The number of minutes you expect a full protection run, which captures all the blocks in an object, to take. • Incremental. The number of minutes you expect an incremental protection run, which captures only the changed blocks in an object, to take.
Cancel Runs at Quiet Time Start	<i>(Available only if the selected policy has at least one Quiet Time.)</i> When enabled, all the protection runs that are currently executing will cancel when the Quiet Time period starts. By default, this setting is disabled, meaning that after a protection run starts, it continues to execute even when a Quiet Time period starts. However, new protection runs will not start during a Quiet Time.

Next > When the first protection run completes, you will be ready to [recover your protected Oracle Databases](#) if and when you need to.

Recover Oracle Databases

Once you have [protected your Oracle Databases](#), you can recover them from Cohesity DataProtect as a Service to their original or a new location.

To recover protected Oracle Databases:

1. Go to **Sources** to set up your recovery task.
2. Click into the **Source** name.
3. Click the **Recover** icon.
4. Select the **Recovery Type**:
 - **Databases** to recover an entire database.
 - **Archive Logs** to recover just the archive logs.
5. Under **Targets**, select **Alternate Database** or **Overwrite Original Database**.
 1. For **Alternate Database**:
 - **Oracle Host**. Provide the hostname to which you want to restore the database.
 - **Configure Channels**:
 - **System selects active node**. Selecting this option auto-selects an active single-instance Oracle node and configures the number of RMAN channels for the database object. (*Default*)
 - **Select specific node(s)**. If you select this option, you can choose the number of RMAN channels and the SBT library path to be used for the database object.
 - **Recovery Options**. Enter:
 - **Restore Database Files to**. Specify the path to an existing empty directory. The newly created database files will reside in this path.

You can restore to an ASM path using this option. If you do, enter just the ASM volume name instead of the entire path. For example, if the restore path for ASM volume **data1** is "+data1" and you enter the entire path, +data1/dbname, the restore task will fail.
 - **Oracle Home**. The ORACLE_HOME value for the host where the database is restored.
 - **Base Directory**. The directory for the restored database.
 - **Target Database Name**. The name for the target database to recover the database components and data files.
 - **Enable Archive Log mode for the Database**. Select the checkbox to enable redo log archiving on the recovered database.

- **BCT File Path.** The BCT file path specifies the location where the block change tracking file will be created. If not provided, BCT is not enabled for the restored database.
- **Leave database in Recovery Mode.** Select the checkbox if you do not want the recovered database to be opened.
- **Shell Environment.** Configure your shell environment that executes the Cohesity DataProtect as a Service restore tasks. For example, define your TNS_ADMIN shell variable here to point to a different `sqlnet.ora` file for use as the restore target database. For TDE support, the wallet location for your restore target database might depend on a shell variable. Use this to specify your wallet location for restoring a backup taken from a TDE database.
- **Task Name.** Change the default name of the recovery task.

2. For **Overwrite Original Database:**

- **Leave database in Recovery Mode.** Select the checkbox if you do not want the recovered database to be opened.
- **Shell Environment.** Configure your shell environment that executes the restore tasks. For example, define your TNS_ADMIN shell variable here to point to a different `sqlnet.ora` file for use as the restore target database. For TDE support, the wallet location for your restore target database might depend on a shell variable. Use this to specify your wallet location for restoring a backup taken from a TDE database.
- **Task Name.** Change the default name of the recovery task.

6. Click **Recover**.

SAP HANA

Cohesity DataProtect as a Service natively integrates with SAP HANA to provide backup and recovery for SAP HANA deployments. It provides an SAP HANA Connector Agent to back up and restore enterprise SAP HANA deployments. The SAP HANA Connector leverages SAP HANA's [backup](#) and [restore](#) capabilities.

The SAP HANA Connector enables you to back up and restore TenantDB and SystemDB. It provides:

- Backup:
 - Full backup
 - Delta backups (incremental backups and differential backups)
 - Automatic Log backup
- Restore:

- Restore the database to its most recent state.
- Restore the database to the point in time.
- Restore the database to a specific data backup.
- Alternate restore of SAP HANA database
- Alternate restore to a new or active database on the same SAP HANA system.
- Alternate restore to a new or active database on a different SAP HANA system.

Requirements for SAP HANA

To register your SAP HANA source and protect your databases, be sure you meet the requirements.

Prerequisites

Before you register your SAP HANA deployment as a source with Cohesity and protect SAP HANA databases, ensure the following prerequisites:

- Identify the node(s) in your SAP HANA cluster with adequate resources to execute the SAP HANA connector agent.
- Synchronize the SAP HANA Linux machine time with a Network Time Protocol (NTP) server. For detailed instructions, see [How to Sync Linux Server Time with an NTP Server](#).
- Ensure the NTP time is in sync between the client and the SaaS Connector. Perform the following on the client:

- Verify if NTP is enabled and synchronized in the client using the following command:

```
timedatectl
```

- If the NTP time is not synchronized, use the following command to synchronize the NTP time

```
timedatectl set-ntp true
```

- Use "bash" shell to install the SAP HANA Connector and to run backups and restore workflows of SAP HANA databases.
- A database user with database admin user privileges.
- Download and install the Cohesity Linux Agent on the SAP HANA node.
- Download and install the SAP HANA connector agent on the SAP HANA host.

Supported Versions

For more information on supported SAP HANA versions, see [Supported Software for Cohesity DataProtect as a Service](#).

Port Requirements

You must open certain ports in the firewall to allow the Cohesity SaaS Connector(s) to transmit and receive data. The following ports should be opened in the firewall to allow data transfer between the SAP HANA system and the Cohesity SaaS Connector(s):

Source	Destination	Port	Protocol	Purpose
SAP HANA host	SaaS Connector	11117 (secure gRPC)	gRPC	Required for Backup and Recovery operations.

Considerations

Review and understand the following before you protect your SAP HANA database:

- SAP HANA 1.0 is not supported.
- Indexing is not supported.
- Scheduling log backups in the Cohesity DataProtect as a Service is not supported, as SAP HANA automatically performs log backups.
- You cannot cancel a running restore operation.

Create the Secure User Store (hdbuserstore) Key

Create the SAP HANA hdbuserstore key to connect to SYSTEMDB. For detailed instruction, see Secure User Store (hdbuserstore). Use the following command to create the hdbuserstore key:

```
hdbuserstore SET <KEY> <host:port> <USERNAME> <PASSWORD>
```

Note:

Ensure that the user account of the secure user store has the following administrator privileges:

- BACKUP ADMIN
- DATABASE ADMIN
- DATABASE START
- CATALOG READ
- MONITORING

Sample:

```
hdbuserstore SET SYS_KEY test-FI:30013 SYSTEM testpassword
```

After you have created the hdbuserstore key, verify if the key has the required SAP HANA connection information, such as hostname or login credentials. Verify the key using the following command:

```
tecadm@test-FI:/usr/sap/TEC/home> hdbsql -AU SYS_KEY
"select DATABASE_NAME,ACTIVE_STATUS from m_databases"
| DATABASE | ACT |
| ----- | --- |
| SYSTEMDB | YES |
| TEC      | YES |
| SRC      | YES |
| DST      | YES |
4 rows selected (overall time 28.150 msec; server time 659 usec)
tecadm@test-FI:/usr/sap/TEC/home>
```

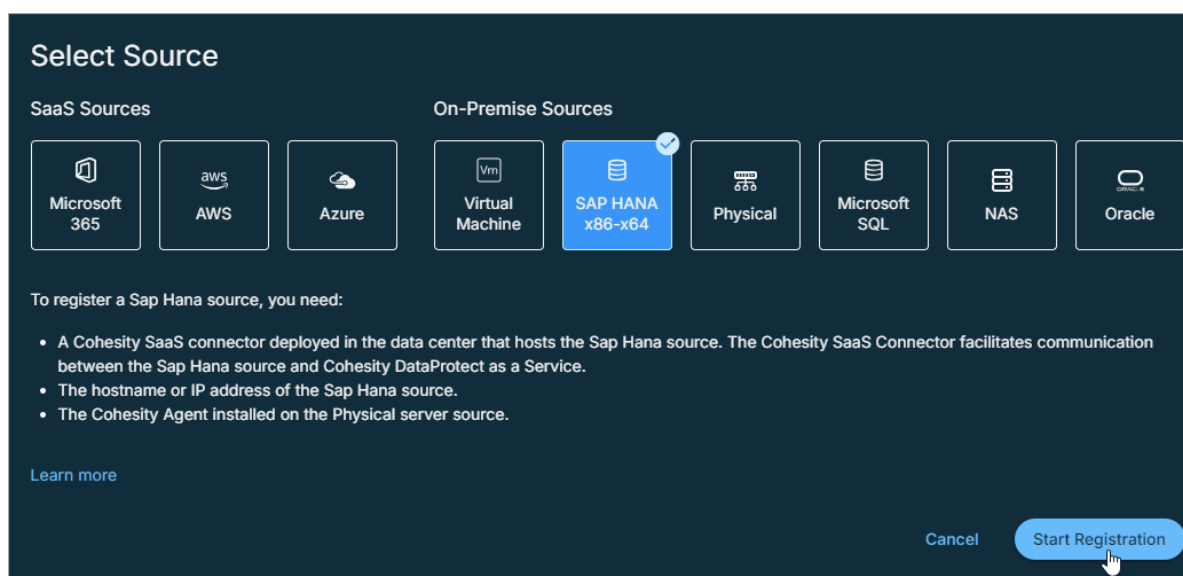
Download and Install the Cohesity Linux Agent

The Cohesity Linux Agent is available as a script-based, RPM, Debian, and SUSE RPM Installer Package. Based on your requirements, download the installer and install it on the SAP HANA node(s).

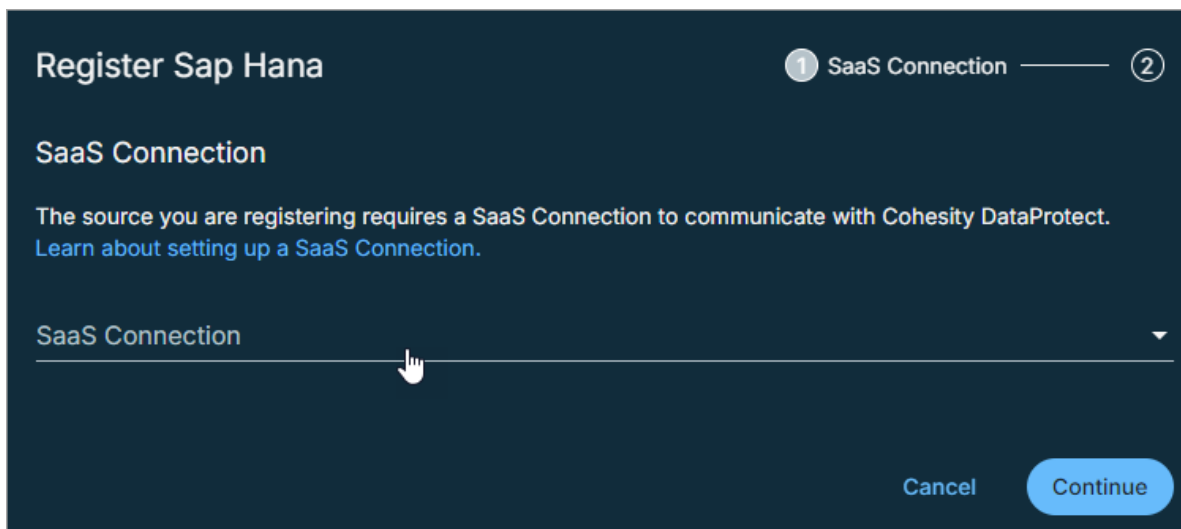
Download the Cohesity Linux Agent

To download the Cohesity Linux Agent:

1. In **DataProtect as a Service**, navigate to the **Sources** page and click **+ Register Source** in the upper-right corner of the page.
2. In the **Select Source** dialog box, select **SAP HANA** and then click **Start Registration**.

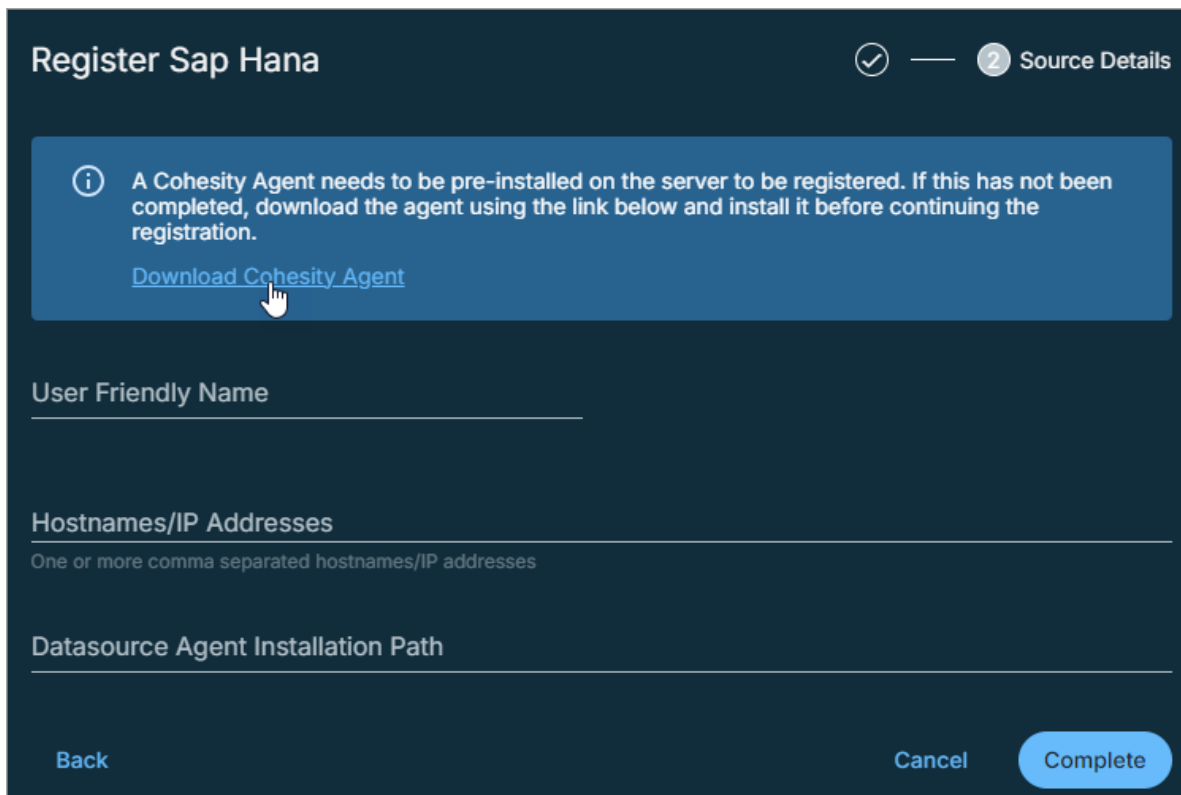


3. In the **Register SAP HANA** dialog box, select an existing healthy SaaS connection marked Unused, or click **Create SaaS Connection**, follow the instructions in [Create a SaaS Connection](#), and then click **Continue**.



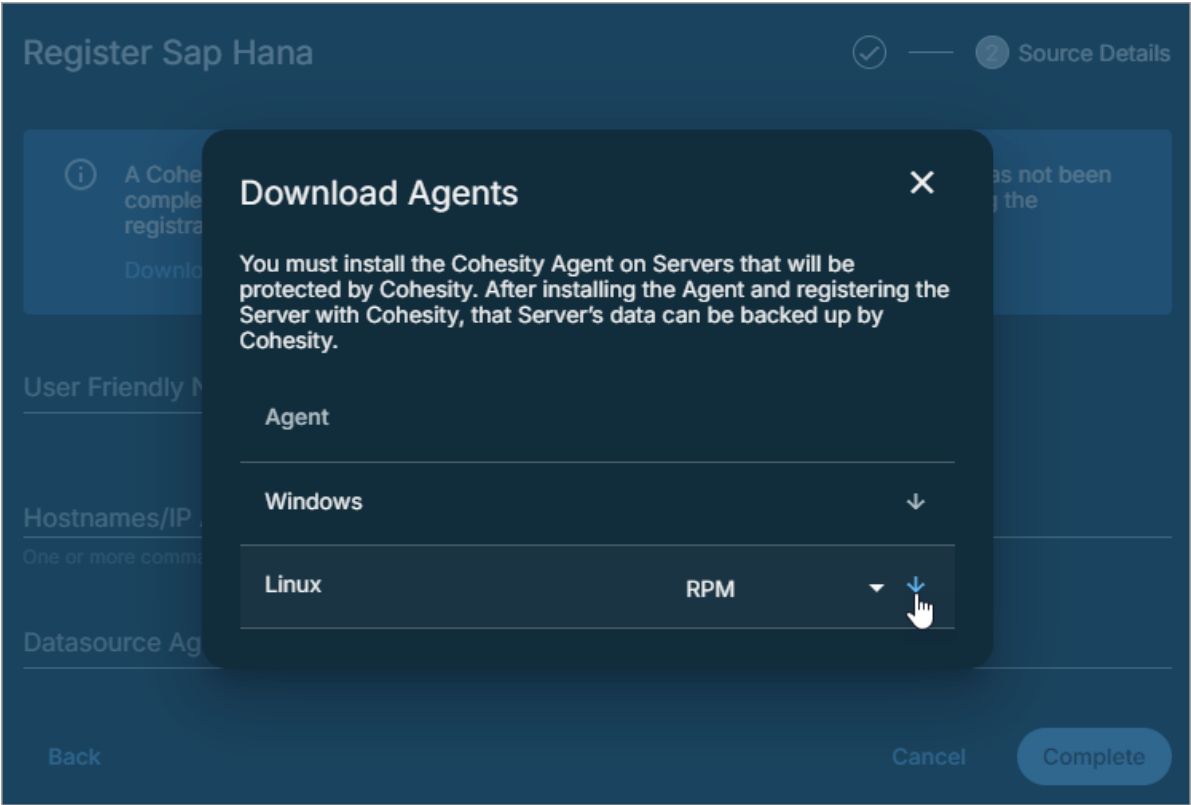
The screenshot shows the 'Register Sap Hana' dialog box with the 'SaaS Connection' step selected. The title bar shows '1 SaaS Connection' and '2'. The main content area has the heading 'SaaS Connection' and a message: 'The source you are registering requires a SaaS Connection to communicate with Cohesity DataProtect. [Learn about setting up a SaaS Connection.](#)' Below this is a dropdown menu labeled 'SaaS Connection' with a hand cursor pointing to it. At the bottom right are 'Cancel' and 'Continue' buttons.

4. Click **Download Cohesity Agent**.



The screenshot shows the 'Register Sap Hana' dialog box with the 'Source Details' step selected. The title bar shows a checkmark, a separator, and '2 Source Details'. A blue information box at the top contains an 'i' icon and the text: 'A Cohesity Agent needs to be pre-installed on the server to be registered. If this has not been completed, download the agent using the link below and install it before continuing the registration.' Below this is a link 'Download Cohesity Agent' with a hand cursor pointing to it. The main content area has three input fields: 'User Friendly Name', 'Hostnames/IP Addresses' (with a hint 'One or more comma separated hostnames/IP addresses'), and 'Datasource Agent Installation Path'. At the bottom left is a 'Back' button, and at the bottom right are 'Cancel' and 'Complete' buttons.

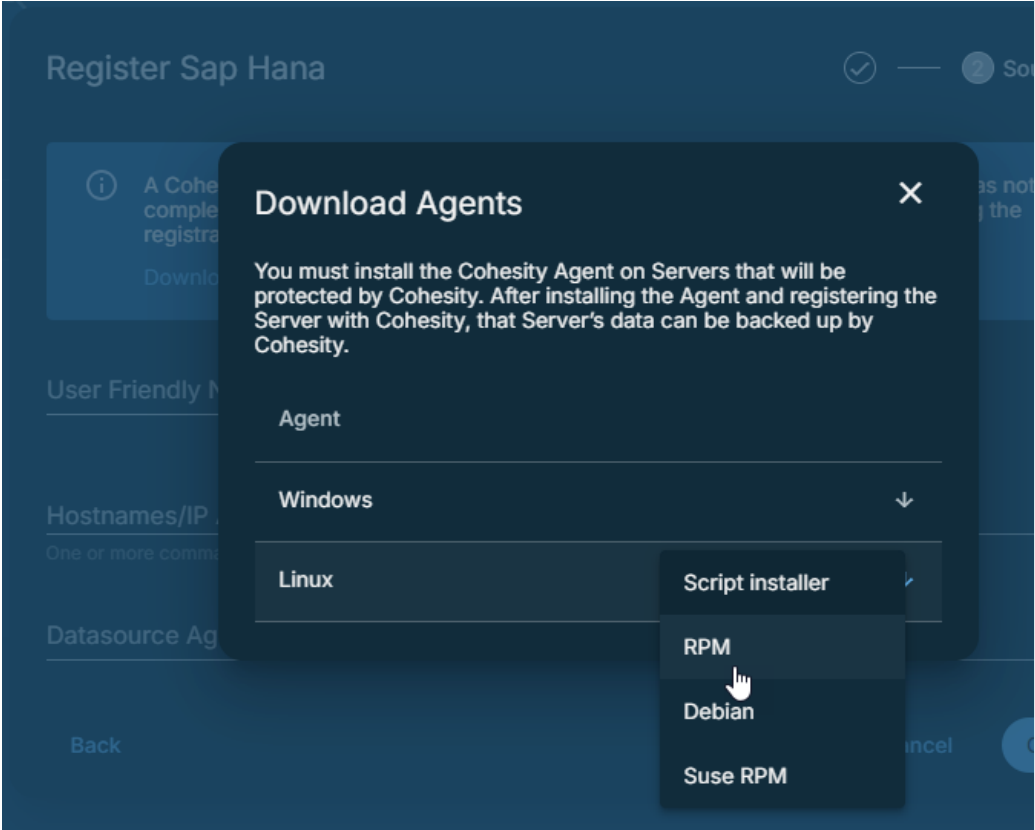
5. In the **Download Agents** dialog box, under **Linux**, select the Cohesity Agent installer package from the drop-down menu and then click the download icon.



Install RPM, Debian or SUSE RPM Installer Package

The following table explains the action to be performed for installing the RPM, Debian or SUSE installer package:

Note: In the following commands, replace the Cohesity agent version with the desired agent version.

Action	How to
Download the Agent installer	<p>From the Download Agents window, based on your Linux distribution, select RPM, Debian, or SUSE RPM and download it to the server you want to protect. For detailed instructions, see Download the Cohesity Linux Agent.</p> 
Navigate to the downloaded directory	<p>As the root user with local system privileges on that server, change the directory to the location of the installer package.</p>

Action	How to
Install the Custom certificates	<p>LINUX</p> <p>If you are using a Cohesity certificate then run the following commands to install the user certificate followed by agent installation:</p> <pre>export ENFORCE_USE_CUSTOM_CERTS=true && export AGENT_CERT_FILE=<path-to-agent-certificate-in-Cohesity-format></pre> <p>If you are using a user certificate then run the following commands to install the user certificate followed by agent installation:</p> <pre>export ENFORCE_USE_CUSTOM_CERTS=true && export USE_THIRD_PARTY_CERTS=true && export ROOT_CA_FILE=<path-to-root-ca-cert> && export PRIVATE_KEY_FILE=<path-to-server-private-key> && export CERT_CHAIN_FILE=<path-to-server-cert></pre>
Install the Agent	<p>Run the following command depending on the installer package:</p> <ul style="list-style-type: none"> • RPM: <code>rpm -i el-cohesity-agent-<version>-1.x86_64.rpm</code> or <code>yum localinstall ./el-cohesityagent-<version>-1.x86_64.rpm</code> • Debian: <code>dpkg -i cohesity-agent-<version>-1_amd64.deb</code> • SUSE RPM: <code>rpm -i cohesity-agent-<version>-1.x86_64.rpm</code>
Install the agent as root user	<p>Run the executable file with the following command syntax:</p> <pre>/<path_to_installer_file> -- --install -c 0 -S root -G root</pre> <p>Followings are the command options in details:</p> <ol style="list-style-type: none"> 1. -S this is the user to run the agent, specify root. 2. -G this is the group permission the Agent will use for files and directories installed by the agent, specify root. 3. -c this boolean switch controls if the OS user and group should be created. 0 means do not create the OS user and group, 1 means the agent installation will create the specified OS user and group. (When choose to run with the root user, specify '-c 0' as 'root' already exists)

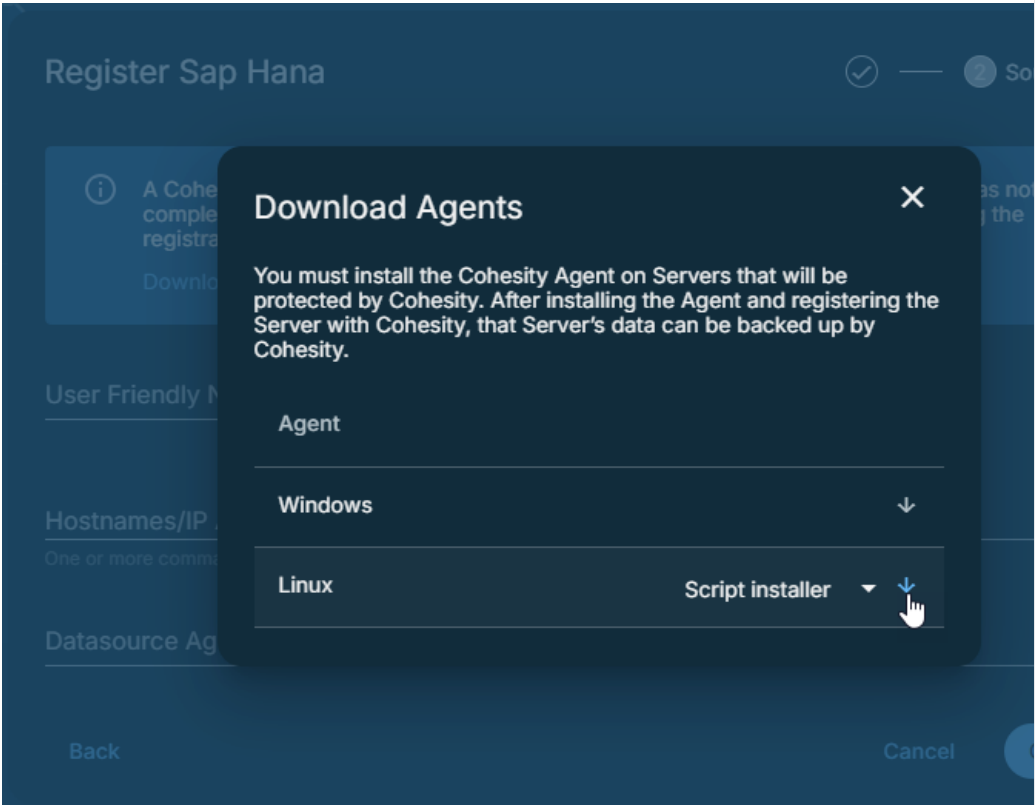
Action	How to
Install the agent as non-root user	<p>To start the service as a non-root user, create a new user or use an existing user with sudo permission and run the following command:</p> <ul style="list-style-type: none"> • RPM: <code>export COHESITYUSER=<username> ; rpm -i el-cohesity-agent-<version>-1.x86_64</code> • Debian: <code>COHESITYUSER=<username> dpkg -i cohesity-agent-<version>-1_amd64</code> • SUSE RPM: <code>export COHESITYUSER=<username> rpm -i cohesity-agent-<version>-1.x86_64</code> <p>Where, <code><username></code> is the user name with sudo permission.</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <p>Note:</p> <p>By default, the installation uses the root user permission for all the files, and the service is started as root. Therefore, it is required to add non-root users to the sudoers list by making the following changes in the <code>/etc/sudoers</code> file:</p> <pre><user_name> ALL=(ALL) NOPASSWD:ALL Defaults:<user_name> !requiretty</pre> </div> <p>An example of the sudoers file is as given below:</p> <pre>dgoble ALL=(ALL) NOPASSWD:ALL Defaults:dgoble !requiretty cohbackup ALL=(ALL) NOPASSWD:ALL Defaults:cohbackup !requiretty</pre>

Action	How to
	<div><p>Note:</p><p>You can add sudo privileges for non-root users using the syntax given below:</p><pre><user_name> ALL=(ALL) NOPASSWD:/path/to/command1, /path/to/command2, /path/to/command3</pre><p>Remove <code>NOPASSWD</code> from the command above if you do not want to authenticate every sudo commands with password.</p><p>The example below shows the sudo commands for user <code>cohesityagent</code>:</p><pre>cohesityagent ALL=NOPASSWD: /bin/mount,/bin/umount,/bin/chmod,/bin/chown,/bin/hostname,/bin/mkdir, /bin/rm,/usr/bin/tee,/usr/bin/stat,/bin/cp,/bin/sed, /usr/bin/touch,/sbin/blkid,/usr/bin/lsof,/bin/ls,/sbin/losetup,/sbin/dmsetup, /usr/bin/rsync,/usr/bin/timeout,/usr/bin/getfacl,/usr/bin/setfacl,/tmp/installer_ch_cohesity_linux_agent, /home/cohesityagent/installer,/home/cohesityagent/cohesityagent/data/agent/tmp/installer_ch_cohesity_linux_agent, /bin/systemctl,/home/cohesityagent/cohesityagent/software/crux/bin/cohesity_linux_agent.sh</pre><p>The actual paths for each command may vary in each system. Cohesity keeps adding or removing commands across releases.</p></div>
Location	<ul style="list-style-type: none">• Installation directory: <code>/opt/cohesity</code>• Log file: <code>/var/log/cohesity</code>

Install Script Installer Package

The following table explains the action to be performed for installing the script installer package:

Note: In the following commands, replace the Cohesity agent version with the desired agent version.

Action	How to
Download the Agent installer	<p>From the Download Agents window, based on your Linux distribution, select Script Installer and download it to the server you want to protect.</p> 
Navigate to the downloaded directory	<p>As the root user with local system privileges on that server, change the directory to the location of the installer package.</p> <p>For SLES 11 SP4, installing the agent as the root user is required.</p>
Make the installer executable	<p>Make the installer executable, for example: <code>chmod +x cohesity_agent_<version>_linux_x64_installer</code></p>
Install the Agent	<p>Run the executable:</p> <pre>sudo /cohesity_agent_<version>_linux_x64_installer - --install</pre>

Action	How to
Location	<ul style="list-style-type: none"> Installation directory: /home/<username>/cohesityagent or /root/cohesityagent Log file: /home/cohesityagent/cohesityagent/logs <p>Note: You can specify the option <code>-log-dir <installtion_path></code> to specify the directory path where the agent logs will be saved. If this option is not provided, the logs will be saved in the default path</p>

Download the SAP HANA Connector Agent Installer

Cohesity provides a SAP HANA connector for SAP HANA to backup and restore SAP HANA databases. The SAP HANA connector for SAP HANA is available as a script-based installer on the [Cohesity Download](#) portal. Log on to the [Cohesity Download](#) portal and download the respective SAP HANA Connector installer to the SAP HANA host.

Install the SAP HANA Connector Using Script Based Installer

To install the SAP HANA connector using script-based installer:

1. Log on to the SAP HANA node's bash CLI using <SID>adm user.
2. Make the SAP HANA connector installer (`cohesity_secure_connector_service_7.2.1c_sap_hana_installer`) executable using the following command:

```
SAP_HANA_SERVER:/home/adminse # chmod +x cohesity_secure_connector_service_<version>_sap_hana_installer
```

Note: During the SAP HANA connector installation, by default, the script installer extracts the files to the `/tmp` directory. If you plan to extract the files to a different directory, you can define the directory path in the `TMPDIR` environment variable as follows: `export TMPDIR=<directory_path>`.

3. Install the SAP HANA connector using the following command:

```
./cohesity_secure_connector_service_<version>_sap_hana_installer -- --install --install-dir=<backint install directory> --sid=<SID> --systemdb-key=<systemdb_key> --workflow=<workflow>
```

Note: You can view the details of SAP HANA Connector installer options using the `--help` option. For example: `./cohesity_secure_connector_service_sap_hana_installer --help`

Where:

Options	Mandatory?	Description
<code>--install-dir</code>	No	<p>The directory on the SAP HANA node on which you want to install the SAP HANA connector.</p> <p>Note: If <code>--install-dir</code> is not specified, the SAP HANA connector is installed in the <code>/usr/sap/\${SAPSYSTEMNAME}/home</code> directory.</p>
<code>--sid</code>	No	SAP System Identification (SID) is a unique three-character code that is available for every R/3 installation (SAP system), which consists of a database server and several application servers.
<code>--systemdb-key</code>	Yes	The SAP HANA HDBUSERSTORE key to connect to the SYSTEM DB. If you plan to backup your SAP HANA databases using the SAP HANA connector, ensure that you use the option <code>--systemdb-key <user_store_key></code> .
<code>--debug</code>	No	Option to enable installer debug.

After the SAP HANA connector is installed, the SAP HANA connector agent are available at: `<backint_install_directory>/cohesity_backint_plugin/uda_scripts`

For example:

```
/usr/sap/SID/home/cohesity/cohesity_backint_plugin/uda_scripts
```

- Log on to the SAP HANA node as root user and create a soft link of `<backint_install_directory>/cohesity_backint_plugin/uda_scripts` folder to `/opt/cohesity/agent/uda_scripts` using the following command:

```
ln -s /usr/sap/SID/home/cohesity/cohesity_backint_plugin/uda_scripts /opt/cohesity/agent/uda_scripts
```

Note: Before creating the soft link, make sure to remove any existing soft link to the `<backint_install_directory>/cohesity_backint_plugin/uda_scripts` folder.

Uninstall the SAP HANA Connector

To uninstall the SAP HANA connector, run the following command with appropriate parameters on SAP HANA node's bash CLI.

Using the Script installer:

```
./cohesity_secure_connector_service_<version>_sap_hana_installer -- --uninstall --install-dir=<backint_install_directory> --sid=<sid> --workflow=<workflow>
```

Note: Specify the `--standby` option, if you are uninstalling the SAP HANA connector on a standby node.

After you have successfully uninstalled the SAP HANA connector, log on to the SAP HANA node as root user and unlink the soft link from `/opt/cohesity/agent/uda_scripts`.

Upgrade the SAP HANA Connector

To upgrade the SAP HANA connector, run the following command on all SAP HANA nodes.

Using the Script installer:

```
./cohesity_secure_connector_service_<version>_sap_hana_installer -- --upgrade --install-dir=<backint_install_directory> --sid=<sid> --systemdb-key=<user_store_key> --workflow=<workflow>
```

After you have upgraded the SAP HANA connector, log on to the SAP HANA node as root user and create a soft link of `<backint_install_directory>/cohesity_backint_plugin/uda_scripts` folder to `/opt/cohesity/agent/uda_scripts` using the following command:

```
ln -s /usr/sap/SID/home/cohesity/cohesity_backint_plugin/uda_scripts /opt/cohesity/agent/uda_scripts
```

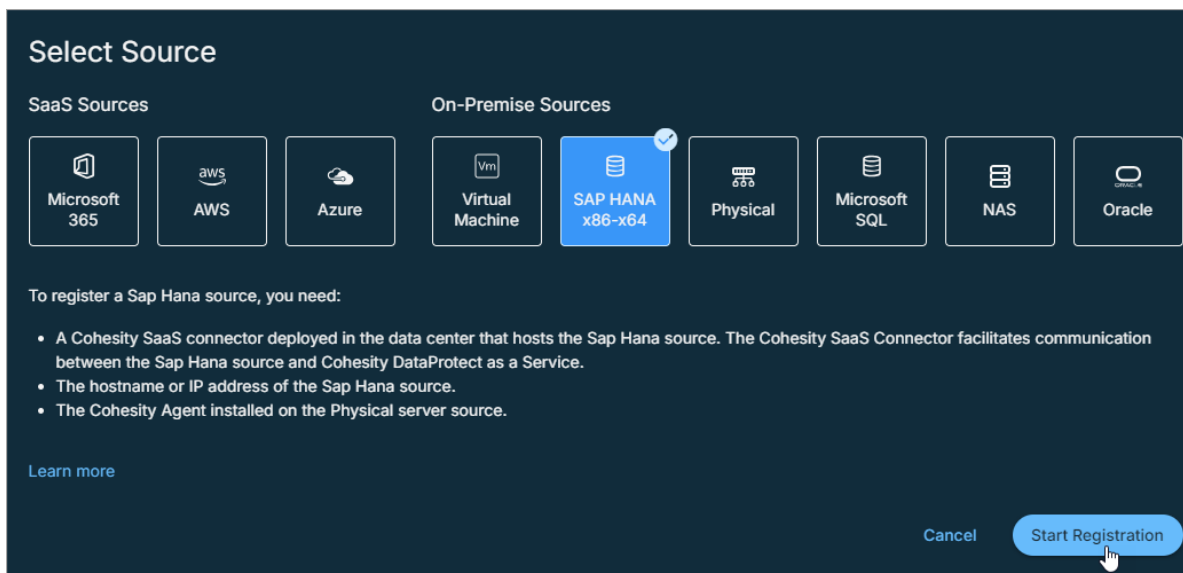
Register SAP HANA Sources

To start protecting a SAP HANA database, once you meet the [requirements](#), you need to register the SAP HANA deployment as a source.

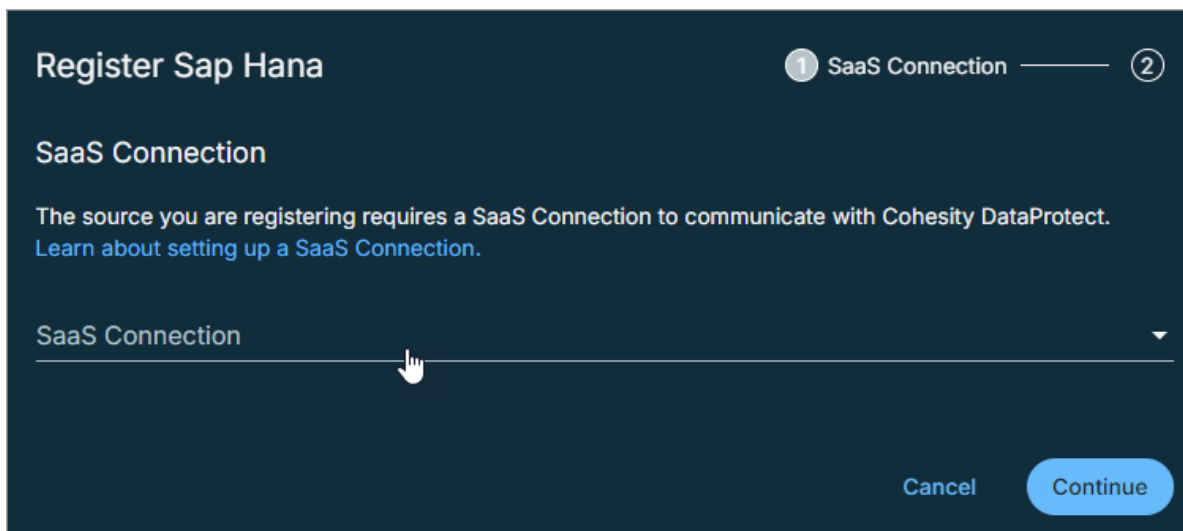
Note: To connect with sources in your data center, you'll need to use a SaaS Connection (or [create one](#)) to establish connectivity between the sources and the Cohesity DataProtect as a Service.

To register your SAP HANA deployment as a source:

1. In **DataProtect as a Service**, navigate to the **Sources** page and click **+ Register Source** in the upper-right corner of the page.
2. In the **Select Source** dialog box, select **SAP HANA** and then click **Start Registration**.

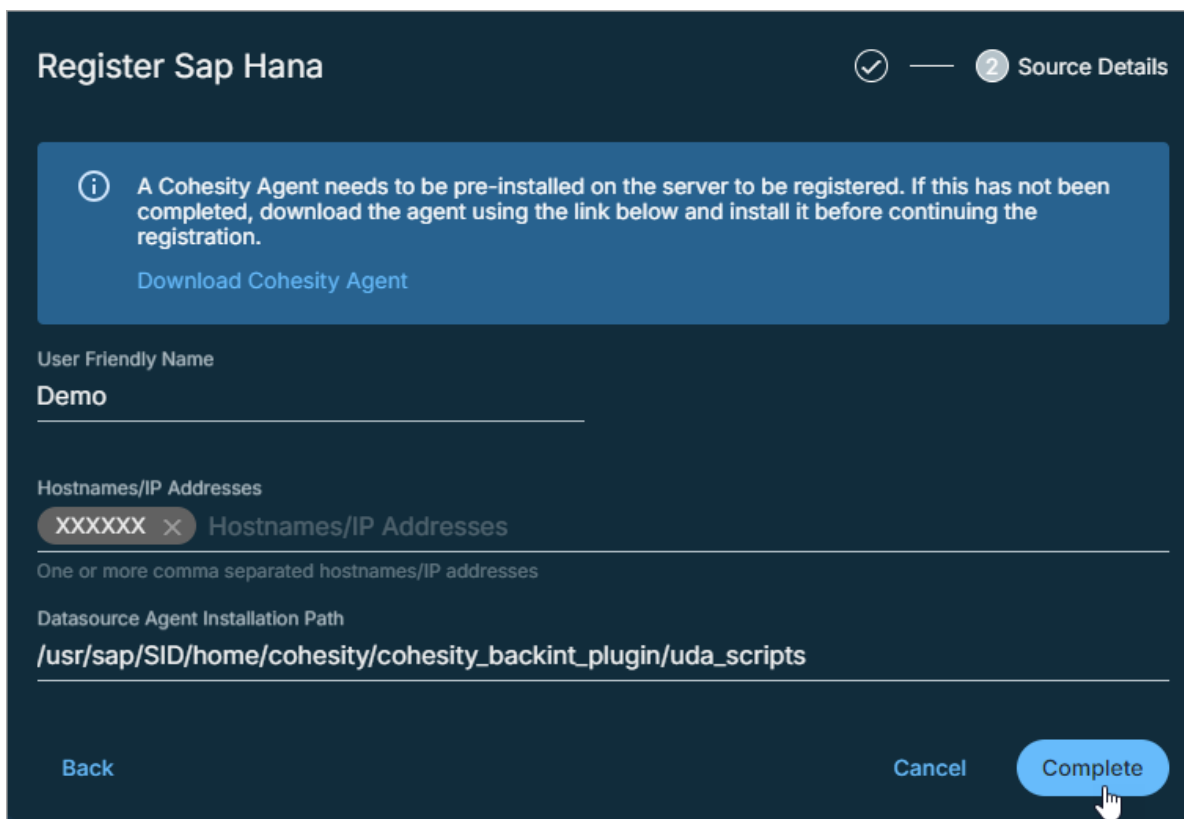


3. In the **Register SAP HANA** dialog box, select an existing healthy SaaS connection marked *Unused*, or click **Create SaaS Connection**, follow the instructions in [Create a SaaS Connection](#), and then click **Continue**.



4. Enter a unique name to identify the SAP HANA source.
5. Enter the **Hostnames or IP Addresses** of the SAP HANA node(s) you have identified to run the SAP HANA connector agent.

6. In the **Datasource Agent Installation Path** field, enter the directory path on the SAP HANA node where you have installed the SAP HANA connector agent.
7. Click **Complete**.



Register Sap Hana ✓ — 2 Source Details

i A Cohesity Agent needs to be pre-installed on the server to be registered. If this has not been completed, download the agent using the link below and install it before continuing the registration.

[Download Cohesity Agent](#)

User Friendly Name
Demo

Hostnames/IP Addresses
XXXXXX × Hostnames/IP Addresses

One or more comma separated hostnames/IP addresses

Datasource Agent Installation Path
/usr/sap/SID/home/cohesity/cohesity_backint_plugin/uda_scripts

[Back](#) [Cancel](#) [Complete](#)

Protect SAP HANA Databases

Once you have [registered your SAP HANA deployment](#) as a source, you're ready to protect the SAP HANA database in your deployment.

To protect your SAP HANA databases:

1. In **DataProtect as a Service**, find the **SAP HANA** source under **Sources** and click on it.
2. Use the checkboxes to select the databases for protection. To protect the whole source, click the checkbox above the column.
3. Click the **Protect** icon above the checkboxes.
4. In the **New Protection** dialog, select a **Policy** that matches the schedule and retention period you need. If the existing policies do not meet your needs, you can [create a new policy](#) with the settings you need.
5. To change or configure any of the additional settings, select **More Options** and perform the below steps or else, click **Protect**.

6. In the **Start Time** field, enter the time the protection run should start. The default time zone is the browser's time zone. You can change the time zone of the job by selecting a different time zone.
7. If you need to change additional settings, click the down arrow icon next to [Additional Settings](#) and click **Edit**.
8. Click **Protect**.

Cohesity DataProtect as a Service starts backing up the database(s) you selected. You can monitor the status of the backup on the **Activity** page.

Note: The backups start immediately after you protect the objects, regardless of the time you set for the protection run.

Also, the **Activity** tab of a specific database shows the history of all protection runs, including the one in progress.

Additional Settings

Additional Settings	Description
End Date	If you need to end protection on a specific date, enable this to select the date.
Concurrency	The maximum number of backup streams that are created to exchange data.
SLA	<p>The service-level agreement (SLA) defines how long the administrator expects a protection run to take. Enter:</p> <ul style="list-style-type: none"> • Full. The number of minutes you expect a full protection run, which captures all the blocks in an object, to take. • Incremental. The number of minutes you expect an incremental protection run, which captures only the changed blocks in an object, to take.
Delta Backup	The type of backup: Incremental or Differential

Recover SAP HANA Databases

After you [protect your SAP HANA databases](#), you can recover them from Cohesity DataProtect as a Service, to their original or new location.

To recover protected SAP HANA databases:

1. In **DataProtect as a Service**, go to **Sources** to set up your recovery task.
2. Click on the **Source** name.
3. Above the tree, select **Show All > Protected**.
4. Use the filters, search box, and views to locate and select the SAP HANA database.

You can also use **Global Search** to locate, filter, and select the objects you need. Click the **Global Search** box at the top or type **slash (/)** anywhere to start your search.

5. Click **Recover** at the top to open the **New Recovery** form with the **Latest** snapshot (protection run).

If you need to recover from an earlier snapshot, click the **Edit** icon to open the **Recovery Point** calendar. Click **List** to view the available recovery points by timestamp and click one.

- Click **Select Recovery Point**.
 - Click **Next: Recover Options** to return to the form.
6. Under **Recover To**, select **Original Location** or **New Location**. If you choose:
 - **Original Location**, Cohesity will overwrite the original SAP HANA database on the same SAP HANA host. Note that this is a destructive action that cannot be undone. Type **YES**.
 - **New Location**, Select the Target SAP HANA host to which the data needs to be restored. Enable **Overwrite existing object with the same name** and Type **YES**.
 7. Select your **Recovery Options**:
 - **Concurrency**. The maximum number of restore streams that is created to exchange data. By default, maximum scaling is set to 16.
 - **Start Database**. Set it to ON if you want to activate the database after completing the recovery.
 - **Rename**. Provide the new name for the database if you want to restore the database with a new name.
 - **Task Name**. Change the default name of the recovery task.
 8. Click **Recover**.

Cohesity DataProtect as a Service begins to restore the selected database to the selected location.

Monitoring

Reports

Cohesity provides one-stop-shop reporting on Cohesity DataProtect as a Service. You have an aggregated view of your Cohesity deployment regardless of the use case, workload, or deployment type (on-premises, consumed as a Cohesity-hosted service, or a combination).

The built-in reports are designed to address your top use cases out-of-the-box. You can view an overall summary of your data protection jobs and storage systems, or analyze data at the granular level using powerful filtering options. You can filter, schedule, email, and download reports.

Note: A user logging in to Cohesity DataProtect as a Service through SSO cannot schedule reports if its user account is not available on the **Access Management** page.

The report that you schedule or download inherits the filters that you have applied.

Tip: You can also watch the [Helios Next Generation Reporting](#) video to know more about Reporting.

View Reports

To view a report:

- 1. [Log in to Helios](#).
- 2. In **DataProtect as a Service**, navigate to **Reports**.
By default, the **Library** tab is displayed.
- 3. Click a report card. For more information, see [Choose a Report Type](#).

Each report helps you view, visualize, and analyze data. The following table describes the key features of reports:

Filters	Each report provides various filters that help you pare down the report until it only shows the data that you want in the report. The filter options change depending on the type of report. For more information, see Filter Report Data .
---------	---

Glance bar	The glance bar provides a summary of the report for the time period you set in the filter.
Charts	Each report includes chart(s) that provide a graphical representation of data.
Data table	The Data table in the report provides deeper insights to help you analyze the data. You can customize the columns in the table. For more information, see Customize Table Columns .
Common tasks	<p>You can perform the following tasks:</p> <ul style="list-style-type: none"> • Download Reports • Schedule Reports • Manage Scheduled Reports • Reset to Default View

Choose a Report Type

Each different report type can help you identify the information you need. Currently, 16 built-in reports are available:

- [Failures](#)
- [Protected / Unprotected Objects](#)
- [Protected Objects](#)
- [Protection Runs](#)
- [Recovery](#)
- [Service Consumption](#)

Filter Report Data

Reporting in Cohesity DataProtect as a Service provides a comprehensive view of the data under management. You have full control over what data you want to include and view in your reports. Use the filters to pare down your report until it only shows the data that you want in the report. The filter options change depending on the type of report.

For more information about the filtering options available in each report, refer to the help page for the respective report.

Customize Table Columns

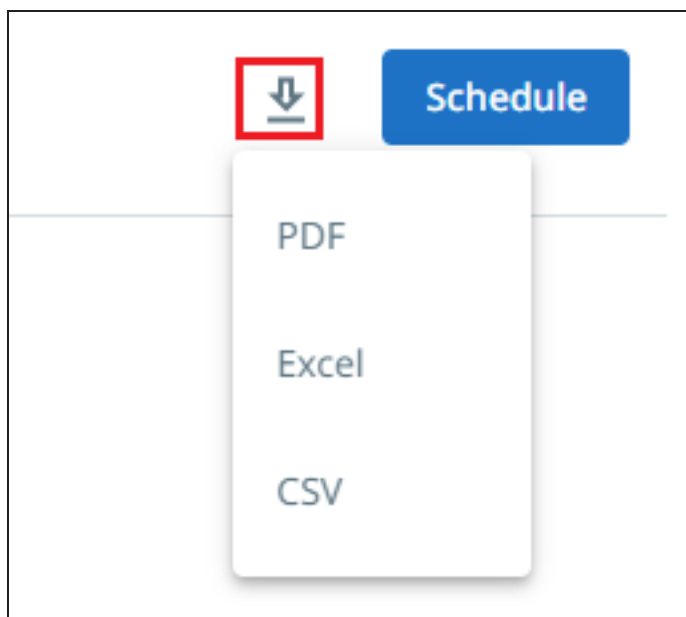
Each report in Cohesity DataProtect as a Service provides comprehensive data. In each report, data is displayed in a tabular format. You can add and remove columns from the **Data** table. The changes you make to columns in a table persist until you change them again or restore the report to the default view.

To customize table columns:

1. In **DataProtect as a Service**, navigate to **Reports**.
2. Click a report card.
3. In the upper-right corner of the table, click the **Settings** (⚙️) icon:
 - Enable the toggle to add a column
 - Disable the toggle to remove a column

Download Reports

You can download reports in different file formats from the reports page. On any report, click the **Download** icon and select one of the file formats:



The report in the selected file format gets downloaded to your system.

Note: The time taken to generate a report depends on multiple factors such as the number of clusters selected, other filters applied on the report, amount of data, and so on. If the report is very large, it may take a few moments to download the report.

Schedule Reports

You can schedule reports to run at periodic intervals. Once you select a report and filter the scope, you can schedule the report to run and send an email to recipients at specified times.

Important Points to Note

- SSO users can view and download reports. To schedule reports, SSO users must be explicitly added in Cohesity DataProtect as a Service. For more information about explicitly adding users, see [Add SSO Users & Groups](#).
- If the report is too large, the email will contain a download link instead of an attachment.
- Columns included in the scheduled report are the columns available in the default view. If you have customized the table, those changes are not reflected in the scheduled report.

To schedule reports:

1. In **DataProtect as a Service**, navigate to **Reports**.
2. Click a report card. For more information, see [Choose a Report Type](#).
3. Click **Create Report**.
4. Configure the following details:
 - **Report Name**—Enter a name for your report.
 - **Report Type**—Choose **Schedule** to create an schedule.
 - **Schedule**—Choose the frequency and the time at which to run the report.
 - **Recipients**—Enter the email address of the recipient. You can enter multiple email addresses.
 - **Email Subject**—Enter a subject line for the email.
 - **Format**—Select the format(s). The recipients receive the report in the format that you select.
 - **Filters**—Apply the necessary filters based on your requirements.
5. Click **Create**.

The recipients receive a new email with the updated report on the schedule you selected. See your scheduled reports under the **Scheduled** tab on the **Reporting** page.


Manage Scheduled Reports

You can perform the following tasks from the **Scheduled** tab:

- Instantly run a report
- Pause a report
- Modify the settings of a report
- Delete a report

Note: Users with the **Super Admin** role can view and manage all scheduled reports in the same Cohesity DataProtect as a Service account.

To manage scheduled reports:

1. In **DataProtect as a Service**, navigate to **Reports**.
2. Click the **Scheduled** tab.
3. Hover over a report and click the **Actions** menu ():
 - Select **Run Now** to instantly run and email the report.
 - Select **Pause** to pause the schedule.
 - Select **Edit** to modify the settings of a scheduled report. Update the settings as necessary and click **Schedule**.
 - Select **Delete** to delete a scheduled report. You must click **Delete** to confirm the deletion.

On-Demand Reports

Generating reports on data from an extended time range can occasionally lead to an API timeout, which prevents access to the report. Cohesity Helios provides on-demand reporting that enables asynchronous report generation, allowing you to create reports for a broader time range without encountering these issues.

Note: This is a Private Preview feature. Private Preview was termed as Early Access in the earlier releases. Contact your Cohesity account team to enable the feature.

After selecting a report and applying the necessary filters, you can create an on-demand report and choose to send the report through an email to designated recipients or manually download it for further analysis.

Important: Generating an on-demand report may take a few minutes to several hours, depending on the amount of data being scanned.

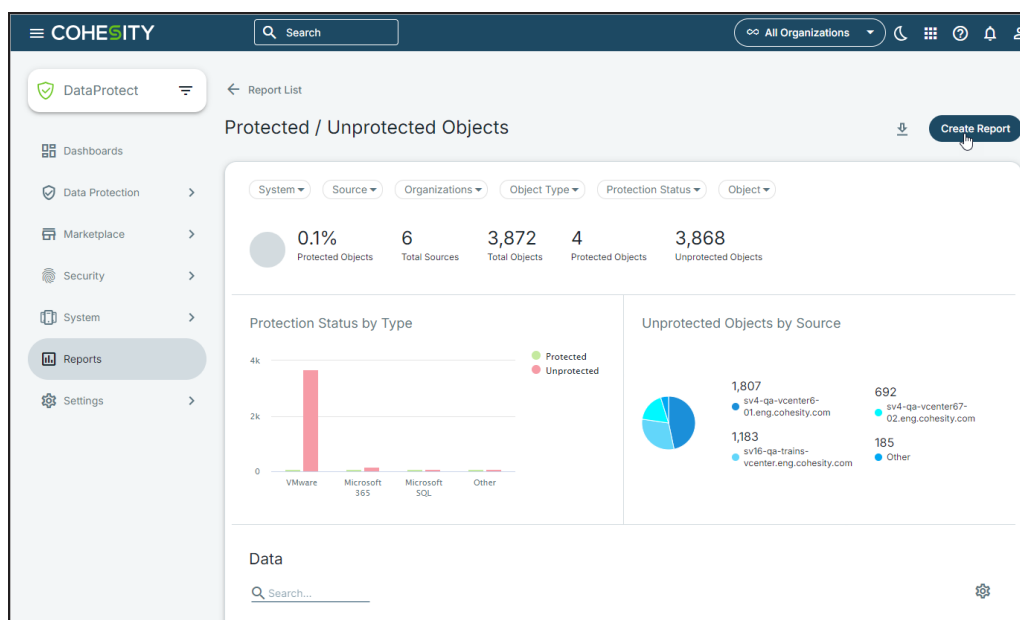
Important Points to Note

- SSO users can view and download reports. To schedule reports, SSO users must be explicitly added in Cohesity DataProtect as a Service. For more information about explicitly adding users, see [Add SSO Users & Groups](#).

- On-demand reports are created once.
- The columns included in the on-demand report are those available in the default view. If you have customized the table, the on-demand report does not reflect those changes.
- For reports, the maximum supported lookback duration is 4 months, and the maximum query duration is 1 month.
- For on-demand reports, the maximum supported lookback duration is 18 months, and the maximum query duration is 1 year.
- On-demand reports cannot be viewed directly in the user interface; they can only be downloaded and viewed.

To create on-demand reports:

1. In **DataProtect as a Service**, navigate to **Reports**.
2. Click a report card. For more information, see [Choose a Report Type](#).
3. Click **Create Report**.



4. In the **Create Report** dialog box, configure the following details:
 - **Report Name**—Enter a name for your report.
 - **Report Type**—Choose **On Demand** to create an on-demand report.
 - **Recipients**—Enter the recipient's email address. You can enter multiple email addresses.
 - **Email Subject**—Enter a subject line for the email.

- **Format**—Select the format(s). The report is generated in the format that you select.
- **Filters**—Apply the necessary filters based on your requirements.


5. Click **Create**.

Once the on-demand report creation is initiated, you can check the report generation status from the On Demand tab on the **Report** page. After the on-demand report is created, the recipients will receive an email with the report attached. Additionally, you can download the on-demand report from the **On Demand** tab on the **Report** page.

Report	Recipients	Filters	Duration
10:30 AM test doc Protection Activities	-	- Date Range: Sep 10, 2024 - Sep 20, 2024	5s
4:05 AM	-	- Date Range: Sep 1, 2024 - Sep 20, 2024	4s
	-	- Date Range: Sep 1, 2024 - Sep 20, 2024	5s
	-	- Date Range: Sep 1, 2024 - Sep 20, 2024	4s

Reset to Default View

Once you filter a report or customize table columns, you can reset the report page's view to the default view. To switch to the default reports page view, click the **Restore to default display** button:



The page refreshes and reverts to the default view.

Helios Reporting APIs

The Helios architecture is API driven. You can programmatically interface with the Helios Reporting service. For more information about using Helios Reporting APIs, see [Helios Reporting Service APIs](#).

Failures

The **Failures** report provides a summary and list of objects that had one or more backup run failures. It also helps you identify consecutive failures in the last three backups, and breaks down the failed objects by object type.

Example use case: Which object do I have no successful backup of in the last week?

Filter Report Data

The report supports multiple filters to pare down the data that you want to view in the report:

- **System**—Select all cluster(s) to include.
- **Source**—Select all the sources to include.
- **Type**—Choose the types of objects to include — Generic NAS, Isilon, NetApp, Physical, Pure, VMware, and so on.
- **Time Range**—Set the time period for your report.
- **Object**—Enter an object name to filter by the name of the object.
- **Organization**—Choose one or more organizations to see the report data specific to the selected organizations.

Glance Bar

The glance bar provides a summary of the report for the specified period:


- **Total Sources**—The total number of sources.
- **Total Objects**—The total number of objects.
- **Failed Objects**—The total number of objects that experienced one or more backup run failures during the specified date range.
- **Without Snapshots**—The total number of objects without any snapshots.

15	77	31	11
Total Sources	Total Objects	Failed Objects	Without Snapshots


Charts

The report includes the following two charts:

• Failures in Last 3 Backups



• Success and Failed Objects by Object Type



Report Data

The following table describes the data displayed in the **Data** table. Use the search bar to filter the data by object name, source, system name, or policy.

Note: You can add or remove columns. For more information, see [Customize Table Columns](#).

The data displayed in the **Policy** and **System** columns are from the last backup run of the object in the specified time period.

Column Name	Description
Object Name	The name of the object.
Source	The hostname or IP address of the registered source.
System	The name of the cluster on which the protection job was run.
Policy	The protection policy associated with the Protection Group.
Last Failed Run	The date and time at which the last backup run failed.
Failed Backups	The total number of backup runs that failed.
Failures in Last 3 Backups	The total number of failures in the last three backups.
Last Fail Reason	The reason for the failure of the last backup.

Related Topics

- [View Reports](#)
- [Filter Report Data](#)
- [Download Reports](#)
- [Schedule Reports](#)
- [Manage Scheduled Reports](#)
- [Reset to Default View](#)

Protected Objects

The **Protected Objects** report provides a summary and list of all protected objects that had a backup run. You can view the backup status and the objects with an active snapshot.

Example use case: Do I have a good backup of my VM in the last month?

Filter Report Data

The report supports multiple filters to pare down the data that you want to view in the report:

- **System**—Select all cluster(s) to include.
- **Source**—Select all the sources to include.
- **Type**—Choose the types of objects to include — Generic NAS, Isilon, NetApp, Physical, Pure, VMware, and so on.

- **Backup Status**—Filter by objects with successful backups or unsuccessful backups.
- **Last Run Status**—Filter by the status of the most recent protection run — Canceled, Failed, Running, Success, and/or Warning.
- **Time Range**—Set the time period for your report.

Note: If you set a time period, the report displays all objects that had a backup run during the selected time period. If an object is no longer protected, the report would still display data if the object had a backup run during the selected time period. If an object is protected and if it did not have a backup run during the selected time period, the report does not display the data specific to this object.

- **Object**—Enter an object name to filter by the name of the object.
- **Organization**—Choose one or more organizations to see the report data specific to the selected organizations.

Glance Bar

The glance bar provides a summary of the report for the specified period:


- **Success Rate—Without Successful Backup / Total Objects.**
- **Total Objects**—The total number of objects.
- **With Successful Backup**—The total number of objects that have one or more successful backups.
- **Without Successful Backup**—The total number of objects that did not have any successful protection runs.
- **With Snapshots**—The total number of objects with snapshots retained. This number can differ from the earlier “With Successful Backups”, for example, all backups fail for an object during the selected date range but the object still has actively retained snapshots from earlier backups (that occurred before the selected date range).
- **Without Snapshots**—The total number of objects without snapshots.




Charts

The report includes the following two charts:

• Protected Objects by Type



• Object Protection by Type



Report Data

The following table describes the data displayed in the **Data** table. Use the search bar to filter the data by object name, system name, source, or policy.

Note: You can add or remove columns. For more information, see [Customize Table Columns](#).

Column Name	Description
Object Name	The name of the protected object.
Source	The hostname or IP address of the registered source.
Policy	The protection policy associated with the latest run of the object.

Column Name	Description
Last Run	The date and time at which the last backup for the object ran.
Last Successful Backup	The date and time at which the last successful backup for the object ran.
Active Snapshots	The total number of active snapshots for the object.
Successful Backups	The total number of successful backups for the object.
Unsuccessful Backups	The total number of unsuccessful backups for the object.
System	The name of the cluster on which the object had the latest run.

Related Topics

- [View Reports](#)
- [Filter Report Data](#)
- [Download Reports](#)
- [Schedule Reports](#)
- [Manage Scheduled Reports](#)
- [Reset to Default View](#)

Protected / Unprotected Objects

The **Protected / Unprotected Objects** report provides a summary and list of objects along with their protection status. You can identify objects that are not associated with a Protection Group. The report does not contain data about Cohesity views.

Example use case: Are all the objects in my vCenter protected?

Filter Report Data

The report supports multiple filters to pare down the data that you want to view in the report:


- **System**—Select all cluster(s) to include.
- **Source**—Select all the sources to include.
- **Type**—Choose the types of objects to include — Generic NAS, Isilon, NetApp, Physical, Pure, VMware, and so on.
- **Protection Status**—Filter by object protection status — Protected or Unprotected.
- **Object**—Enter an object name to filter by the name of the object.

- **Organization**—Choose one or more organizations to see the report data specific to the selected organizations.

Glance Bar

The glance bar provides a summary of the report for the specified period:


- **Protected Objects**—The percentage of **Protected Objects** to **Total Objects**.
- **Total Sources**—The total number of sources.
- **Total Objects**—The total number of objects.
- **Protected Objects**—The total number of protected objects.
- **Unprotected Objects**—The total number of unprotected objects.




Charts

The report includes the following two charts:

- **Protection Status by Type**



- **Unprotected Objects by Source**



Report Data

The following table describes the data displayed in the **Data** table. Use the search bar to filter the data by object name, protection status, source, or system name.

Note: You can add or remove columns. For more information, see [Customize Table Columns](#).

Column Name	Description
Object Name	The name of the object.
Protection Status	The protection status of the object.
Source	The name of the registered source.
System	The name of the cluster on which the object is registered.
Logical Data	<p>The combined total of data in the objects that are protected by Cohesity. These metrics are different depending on workload type.</p> <ul style="list-style-type: none"> • VMs—The data size reported by VMware is the provisioned amount, not the actual data residing in the VM. For example, if a VM is provisioned for 1 TB but contains only 100 GB of data, VMware reports it as 1 TB. • All Other Workloads—The data size reported is the actual front end data residing on the server. If a server with 1 TB capacity contains 100 GB of data, the server reports 100 GB. <p>Note: Cohesity does not include unprotected objects in these metrics.</p>
Organization	The name specified for the organization when added to the cluster.

Related Topics

- [View Reports](#)
- [Filter Report Data](#)
- [Download Reports](#)
- [Schedule Reports](#)
- [Manage Scheduled Reports](#)
- [Reset to Default View](#)

Protection Runs

The **Protection Runs** report provides a summary and list of all backup activities per object per run. You can view the summary and success rate of protection runs. You can also view the snapshot status of the protection run.

Example use case: How many failed protection runs did I have in the last week?

Filter Report Data

The report supports multiple filters to pare down the data that you want to view in the report:

- **System**—Select all cluster(s) to include.
- **Source**—Select all the sources to include.
- **Type**—Choose the types of objects to include — Generic NAS, Isilon, NetApp, Physical, Pure, VMware, and so on.
- **Run Status**—Filter by the status of the protection run — Canceled, Failed, Running, Success, and/or Warning.
- **Snapshot Status**—Filter by the status of the snapshot — Active or Expired.
- **Time Range**—Set the time period for your report.
- **Organization**—Choose one or more organizations to see the report data specific to the selected organizations.

Glance Bar

The glance bar provides a summary of the report for the specified period:


- **Success Rate—Total Successful / Total Runs.**
- **Total Runs**—The total number of protection runs.
- **Total Successful**—The total number of successful runs.
- **Success**—The total number of protection runs with status Success.
- **Warning**—The total number of protection runs with status Warning.
- **Failed**—The total number of protection runs with status Failed.
- **Canceled**—The total number of protection runs with status Canceled.
- **Running**—The total number of protection runs with status Running.
- **SLA Met**—The total number of protection runs that met SLA.
- **SLA Missed**—The total number of protection runs that missed SLA.




Charts

The report includes the following two charts:

• **Run Status by Policy**



• **Run Status by Type**



Report Data

The following table describes the data displayed in the **Data** table. Use the search bar to filter the data by object name, source, policy, system name, or snapshot status.

Note: You can add or remove columns. For more information, see [Customize Table Columns](#).

Column Name	Description
Start Time	The date and time at which the protection run started.
End Time	The date and time at which the protection run was completed.
Object Name	The name of the protected object.
Source	The hostname or IP address of the registered source.
Policy	The protection policy associated with the protection run for the corresponding object.
System	The name of the cluster on which the object had a protection run.

Column Name	Description
Snapshot Status	The status of the snapshot.
Duration	The time taken by the protection run.
Logical Data	<p>The combined total of data in the objects that are protected by Cohesity. These metrics are different depending on workload type.</p> <ul style="list-style-type: none"> • VMs—The data size reported by VMware is the provisioned amount, not the actual data residing in the VM. For example, if a VM is provisioned for 1 TB but contains only 100 GB of data, VMware reports it as 1 TB. • All Other Workloads—The data size reported is the actual front end data residing on the server. If a server with 1 TB capacity contains 100GB of data, the server reports 100 GB. <div> <p>Note: Cohesity does not include unprotected objects in these metrics. Currently, the logical data value shown on the Cohesity DataProtect as a Service Dashboard is a sum of the logical data values captured across all the protection runs. For instance, if the source has 100 GB of logical data, and assuming it remains at 100 GB for the first 10 protection runs, Cohesity would report, after 10 runs, the Logical Data to be 1000 GB (1 TB).</p> </div>
Data Read	Size of the set of protected objects as read by Cohesity for a single backup run. This number is a per protection run statistic and is not additive across backup runs.
Data Written	<p>Data written on the Cohesity platform after the unique logical data has been reduced by data deduplication and data compression.</p> <div> <p>Note: This number reflects unique data written, before resiliency operations.</p> </div>
Organization	The name specified for the organization when added to the cluster.

Related Topics

- [View Reports](#)
- [Filter Report Data](#)
- [Download Reports](#)
- [Schedule Reports](#)
- [Manage Scheduled Reports](#)
- [Reset to Default View](#)

Recovery

The **Recovery** report provides a summary and list of all the clone and recovery tasks that were executed. It also provides other details such as the time taken for the operation and status of the operation.

Note: If a Cohesity view is unprotected, the report does not display data about clone view operations.

Example use case: How many recovery tasks failed in the last week?

Filter Report Data

The report supports multiple filters to pare down the data that you want to view in the report:

- **System**—Select all cluster(s) to include.
- **Source**—Select all the sources to include.
- **Organization** – Choose one or more organizations to see the report data specific to the selected organizations.
- **Type**—Choose the types of objects to include — Generic NAS, Isilon, NetApp, Physical, Pure, VMware, and so on.
- **Status**—Filter by the status of the recovery task — Canceled, Failed, Running, Success, and/or Warning.
- **Time Range**—Set the time period for your report.
- **Object**—Enter an object name to filter by the name of the object.

Glance Bar


The glance bar provides a summary of the report for the specified period:

- **Success Rate—Successful / Total Recoveries.**
- **Total Recoveries**—The total number of recovery runs.
- **Successful**—The total number of recoveries with status Success.
- **Failed**—The total number of recoveries with status Failed.
- **Warning**—The total number of recoveries with status Warning.
- **Canceled**—The total number of recoveries with status Canceled.
- **Running**—The total number of recoveries with status Running.



Chart

The report includes the **Recovery Status by Type** chart:



Report Data

The following table describes the data displayed in the **Data** table. Use the search bar to filter the data by object name, source, system name, task name, or username.

Column Name	Description
Start Time	The date and time at which the recovery task started.
Object Name	The name of the object.
Source	The hostname or IP address of the registered source.
System	The name of the cluster on which the recovery task was run.
Recovery Point	The date and time of the backup run from which the object was recovered.
Duration	The time taken by the recovery task.
Task Name	The name of the recovery task.
Username	The name of the user who initiated the recovery.
Organization	The name specified for the organization when added to the cluster.

Related Topics

- [View Reports](#)
- [Filter Report Data](#)
- [Download Reports](#)
- [Schedule Reports](#)
- [Manage Scheduled Reports](#)
- [Reset to Default View](#)

Service Consumption

The **Service Consumption** report provides statistics — like average usage, peak usage, and change rates — about the Cohesity DataProtect as a Service consumed by your protected objects. It also helps break down current usage and monthly peak usage by type.

Detect Ransomware Attacks

Ransomware can take over enterprise data and threaten to publish it or block access to it until a ransom is paid. Cohesity DataProtect as a Service detects potential ransomware attacks in your environment.

We use machine learning algorithms to continuously monitor change rates in the backup data. If the rate is out of the normal range — based on daily and historical rates — Cohesity DataProtect as a Service flags it as a potential ransomware attack.

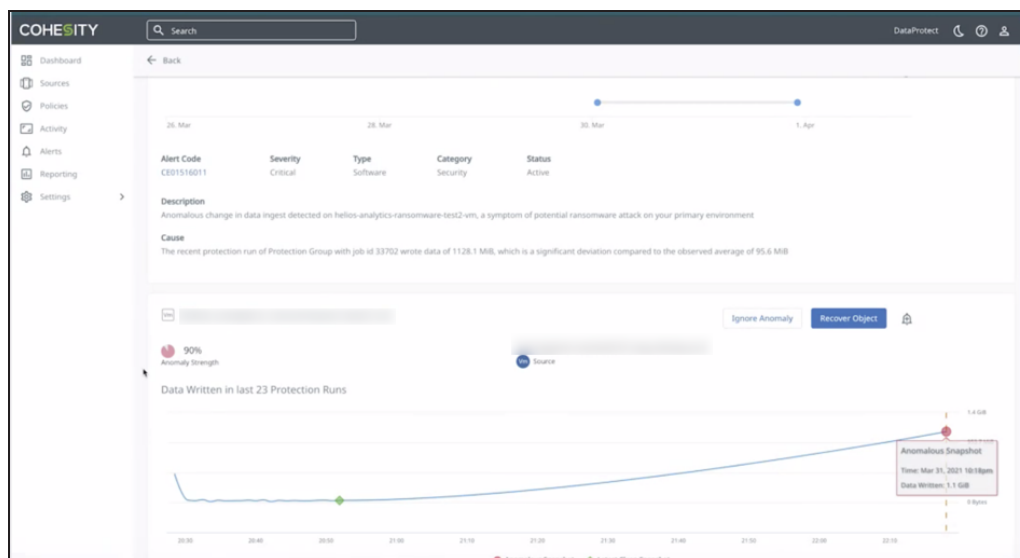
If Cohesity DataProtect as a Service detects an anomaly during a protection run of your data, it triggers the critical alert, **DataIngestAnomalyAlert**. Using the alert information, you can investigate the anomaly and decide on the next course of action.

After reviewing the anomaly, you can either ignore the anomaly or recover the object from the last clean snapshot.

To locate and inspect potential anomalies:

1. In **DataProtect as a Service**, navigate to **Health > Alerts** and then click the **Severity** filter.
2. Select **Critical** and click **Apply**.
3. If you see a **DataIngestAnomalyAlert** alert, click into it.
4. On the **DataIngestAnomalyAlert** page, review the alert details.
5. Once you have thoroughly reviewed the alert, click:
 - **Ignore Anomaly** to dismiss the anomaly.

- **Recover Object** to recover the object from the last clean snapshot.



Alerts

The Cohesity DataProtect as a Service creates an alert for various reasons:

- It finds a potential problem
- Certain criteria exceed the defined threshold
- Informational events which occur in the system
- To indicate the success or failure of the protection run.

Each alert has a severity rating that indicates the seriousness of the problem:

- **Critical**—Immediate action is required because it detects a severe problem that might be imminent or major functionality is not working, such as a missing VM backup.
- **Warning**—Action is required, but the affected functionality is still working, such as the restore task failed due to some external target connectivity and/or credentials issues.
- **Informational**—Immediate action is not required, and the alert provides an informational message.

For a listing of the Alerts created by the Cohesity DataProtect as a Service, see [Alerts References](#).

Analyze the Alert

You can click on an alert from the **Alerts** tab and view the alert details on the **Details for <Alert_Name>** page.

The **Details for <Alert_Name>** page includes a timeline view showing the date and time the alert was triggered. The page also provides the following details of the alert:

Details	Description
Alert Code	The alert code. You can click on the alert code for detailed information about the alert.
Severity	The severity rating of the alert.
Type	The alert type. It defines the Cohesity component that triggered the alert.
Category	The alert category.
Status	The status of the alert. It can be Active, Resolved, or Note.
Description	A brief description of the problem that triggered the alert.
Cause	A brief description of the cause of the problem.

Alert Notification

You can configure general alert email notifications or enable Webhooks for alerts notification in the **Health > Notification** tab. For more information, see [Configure Alert Notification Settings](#).

Resolve Alerts

In case if you are aware of the problem and confirm that the issue has been resolved or if the issue does not require further attention, from the Alerts tab, you can manually resolve those alert(s). You can either create a new resolution of the alert(s) or attach an existing resolution to the alert(s).

Create a New Resolution

To create a new resolution:

1. In the **Alerts** tab, select an alert or multiple alerts that you plan to resolve and click **Resolve Selected Alerts**.

2. In the **Resolution** dialog, do the following:

1. Select **Create new resolution**.

2. In the **Resolution Summary** field, add a resolution summary for the alert.

3. In the **Resolution Description** field, add a brief description of the resolution.

4. Click **Resolve**.

The resolution is added to the selected alerts, and the alert(s) status is marked as **Resolved**.

Attach an Existing Resolution

To attach an existing resolution to the alert(s):

1. In the **Alerts** tab, select an alert or multiple alerts that you plan to resolve and click **Resolve Selected Alerts**.

The screenshot shows the COHESITY Health Alerts tab. The left sidebar contains navigation options: Dashboard, Sources, Policies, Activity, Health (selected), Audit Logs, Reporting, and Settings. The main area displays a table of alerts with columns for ID, Status, Message, Count, Region, Time, Action, and Service. Two alerts are selected, and the 'Resolve Selected Alerts' button is highlighted. A date range filter 'Apr 27, 2023 - May 03, 2023' is applied. The bottom right shows 'Items per page 50' and '1 - 5 of 5'.

ID	Status	Message	Count	Region	Time	Action	Service
256643	Failed	RestoreTaskFailed	1	US East (Ohio)	7 days ago	Backup & Restore	Data Service
252469	Failed	RestoreTaskFailed	1	US East (Ohio)	7 days ago	Backup & Restore	Data Service
251329	Failed	RestoreTaskFailed	1	US East (Ohio)	7 days ago	Backup & Restore	Data Service
247050	Failed	RestoreTaskFailed	1	US East (Ohio)	7 days ago	Backup & Restore	Data Service
97974	Failed	RestoreTaskFailed	1	US East (Ohio)	7 days ago	Backup & Restore	Data Service

2. In the **Resolution** dialog, do the following:
 1. Select **Associate with existing resolution**.

The screenshot shows the Resolution dialog box overlaid on the alerts table. The dialog has two radio buttons: 'Create new resolution' and 'Associate with existing resolution'. The second option is selected. Below the radio buttons are fields for 'Resolution Summary *' and 'Resolution Description'. At the bottom are 'Resolve' and 'Cancel' buttons. The background shows the same alerts table as the previous screenshot.

2. From the **Resolution Summary** drop-down, you can search and select the

resolution that you plan to attach to the alert.

3. Click **Resolve**.

The existing resolution is attached to the selected alerts, and the status of the alert(s) are marked as **Resolved**.

Resolve an alert in the Details for <Alert_Name> page

Once you have reviewed the alert, you can resolve the alert using the page's **Resolution** section. You can create a new alert resolution or attach an existing one in the **Resolution** section.

The screenshot displays the COHESITY DataProtect interface. On the left is a sidebar with navigation icons for Dashboard, Sources, Policies, Activity, Health, Audit Logs, Reporting, and Settings. The main panel is titled 'Details for RestoreTaskFailed' and shows 'Showing Latest 1 of 1 Occurrences' with a date range of 'Apr 27, 2023 4:54pm'. Below this is a timeline showing a single occurrence at 04:00 pm on 27 Apr. A table lists the alert details: Alert Code (CE00610034), Severity (Critical), Type (Data Service), Category (Backup & Restore), and Status (Active). The 'Description' section states: 'Restore for object auto-dmaas-ui-test-win in job Recover_Files_Apr_27_2023_11_02_AM failed'. The 'Cause' section explains: 'One of the restore tasks in job Recover_Files_Apr_27_2023_11_02_AM of type kRestoreFiles for object auto-dmaas-ui-test-win of type kAWS with task ID 677 failed with error: [1-10-11] Cohesity service on host 3.144.222.76 cannot be reached. Please confirm Cohesity Agent is installed successfully on the host, check connectivity to the service, including the required ports, firewall rules, the service status, and the permissions for the service logon account..'. The 'Resolution' section is highlighted with a red box and contains two radio buttons: 'Create new resolution' (selected) and 'Associate with existing resolution'. Below these are text input fields for 'Resolution Summary' and 'Resolution Description', and a blue 'Resolve' button at the bottom.

To create a new resolution:

1. In the **Resolution** section, select **Create new resolution**.
2. In the **Resolution Summary** field, add a resolution summary for the alert.
3. In the **Resolution Description** field, add a brief description of the resolution.
4. Click **Resolve**.

To attach an existing resolution:

1. In the **Resolution** section, select **Associate with existing resolution**.
2. From the **Resolution Summary** drop-down, you can search and select the resolution that you plan to attach to the alert.
3. Click **Resolve**.

Configure Alert Notification Settings

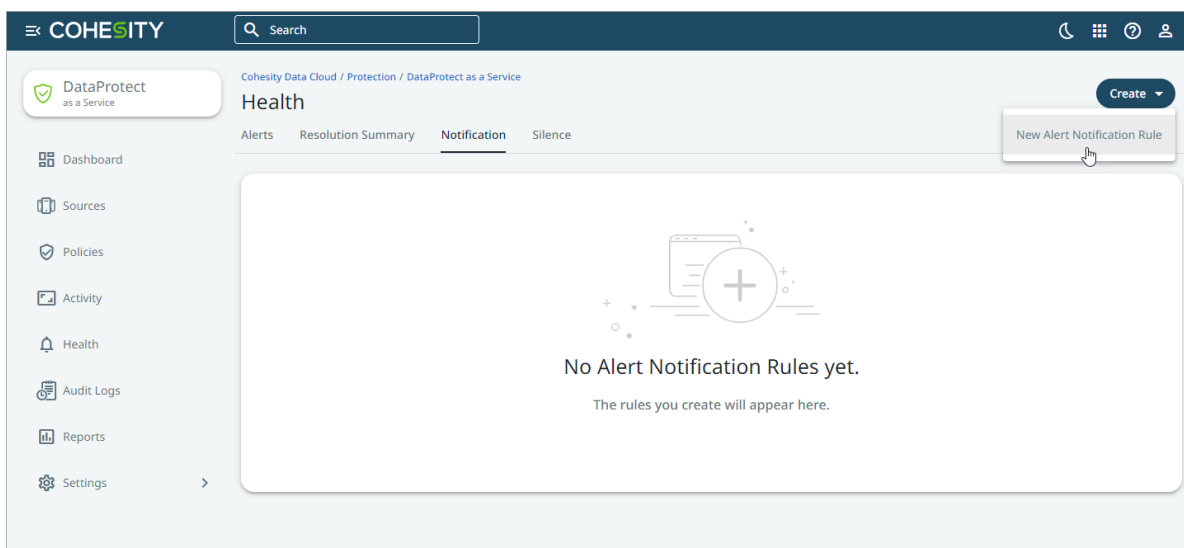
You can configure general alert notification rules from the **Health** page in the **Notification** tab. You can configure email and webhook as the notification output for the alert notification.

Create Alert Notification Rule for Email Notifications

You can add different alert notification rules that send emails based on the alert categories, severities, and names.

To create an alert notification rule for email notifications:

1. In **DataProtect as a Service**, navigate to the **Health > Notification** tab.
2. Click **Create > New Alert Notification Rule**.



3. In the **Create Alert Notification Rule** dialog, perform the following:
 1. Enter a unique **Notification Name** for the alert notification rule.
 2. In the **Notification Filters** section, select the filter based on your requirements:

Note: The alert notification is sent when an alert matches the combination of the filter settings you have configured.

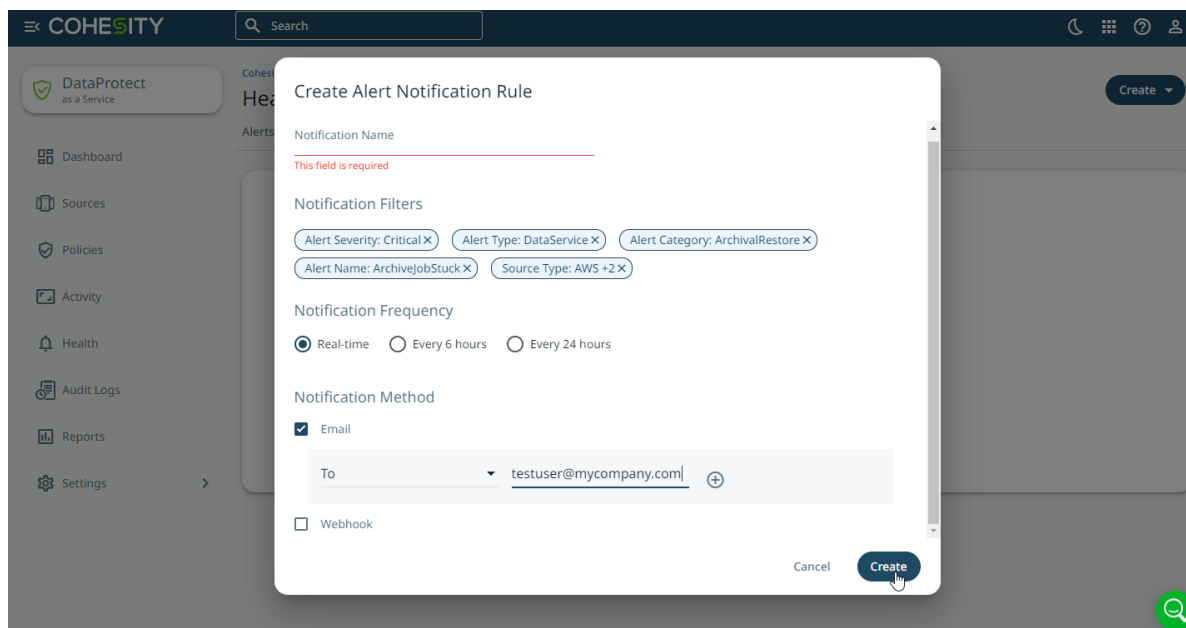
Details	Description
Alert Severity	Select one or more severities from the drop-down. Otherwise, all alerts with any severity will trigger the rule.

Details	Description
Alert Category	Select one or more categories from the drop-down. Otherwise, all alerts in any category will trigger the rule.
Alert Name	Select one or more names from the drop-down. Otherwise, any Alert name will trigger the rule. If you selected any categories, the list includes only alerts in those categories.
Source Type	Select one or more sources from the drop-down. Otherwise, any source will trigger the rule.

3. In the **Notification Method** section, select **Email**. Choose one of the options from the drop-down based on your requirement:

Details	Description
To	Type an email address or distribution list of the recipients to whom you plan to send the email notification.
Cc	Type an email address or distribution list of the recipients to whom you plan to send a copy of the email notification.

Click **+** to add multiple email addresses based on your requirement.



4. Click **Create**.

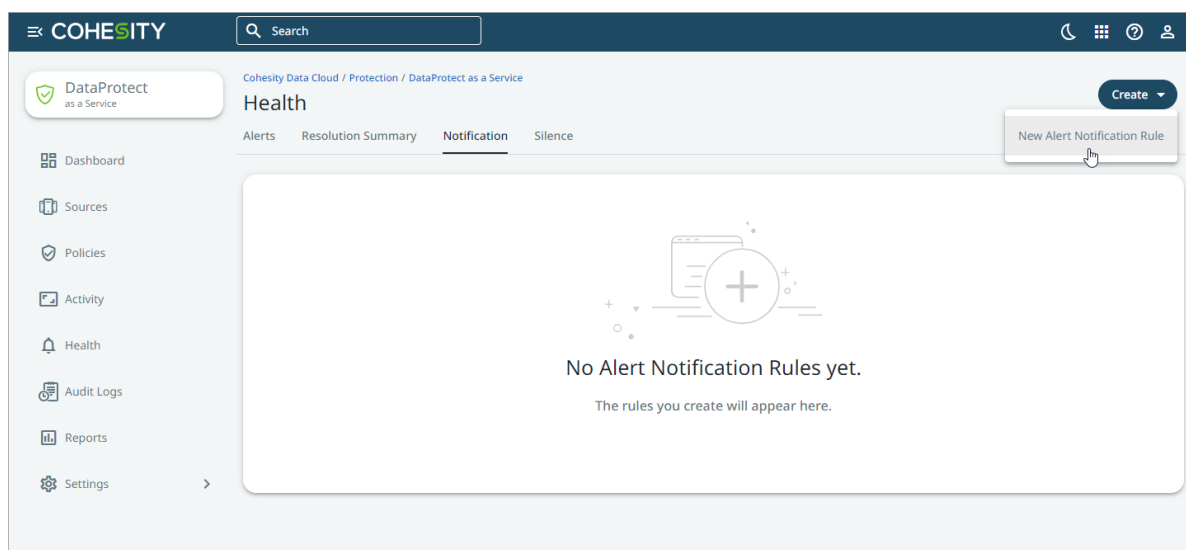
Create Alert Notification Rule for Webhooks Notification

Webhooks are HTTP callbacks that are usually triggered by some event. Webhooks are configured on one website, and when an event occurs on this website, an HTTP request is made to the configured URL, which invokes an action on the other website.

You can enable webhooks for Cohesity DataProtect as a Service alerts by creating an alert notification rule. When the alert is triggered and meets the criteria in the rule, Cohesity DataProtect as a Service sends an HTTP request to the specified website. Your application can interpret the request. For example, the webhook might notify the website about a critical protection run alert, and your application might open a trouble ticket to track the problem.

To create an alert notification rule for Webhook notifications:

1. In **DataProtect as a Service**, navigate to the **Health > Notification** tab.
2. Click **Create > New Alert Notification Rule**.



3. In the **Create Alert Notification Rule** dialog, perform the following:
 1. Enter a unique **Notification Name** for the alert notification rule.
 2. In the **Notification Filters** section, select the filter based on your requirements:

Note: The alert notification is sent when an alert matches the combination of the filter settings you have configured.

Details	Description
Alert Severity	Select one or more severities from the drop-down. Otherwise, all alerts with any severity will trigger the rule.
Alert Category	Select one or more categories from the drop-down. Otherwise, all alerts in any category will trigger the rule.
Alert Name	Select one or more names from the drop-down. Otherwise, any Alert name will trigger the rule. If you selected any categories, the list includes only alerts in those categories.
Source Type	Select one or more sources from the drop-down. Otherwise, any source will trigger the rule.

3. In the **Notification Method** section, select **Webhook**, and provide the URL and cURL options.

COHESITY Search

DataProtect as a Service

Dashboard Sources Policies Activity Health Audit Logs Reports Settings

Create Alert Notification Rule

Notification Name
This field is required

Notification Filters
Alert Severity: Critical X Alert Type: DataService X Alert Category: ArchivalRestore X
Alert Name: ArchivejobStuck X Source Type: AWS +2 X

Notification Frequency
☒ Real-time ☐ Every 6 hours ☐ Every 24 hours

Notification Method
☐ Email
☒ Webhook

e.g. http://example.com/? Options

Cancel Create

4. Click **Create**.

Alert Request

When an alert is triggered, a sample payload, as shown below, will be available at the configured URL:

Request:

```
'https://test-service-now.com/api/x_hesin_cohesity_c/cohesitywebhook'
```

The Payload sent to the above URL:

```

{
  "receiver": "00101000005nBps_test1",
  "status": "firing",
  "alerts": [
    {
      "status": "firing",
      "labels": {
        "account_id": "00101000005nBps",
        "alert_category": "BackupRestore",
        "alert_code": "CE00610005",
        "alert_id": "10534",
        "alert_state": "Open",
        "alert_type_bucket": "DataService",
        "alert_type_id": "10005",
        "alertname": "ProtectedObjectFailed",
        "cluster_id": "1609127048663690",
        "cluster_id_str": "4327092961767844",
        "cluster_name": "DPCluster",
        "failure_reason": "Testing DP alerts raise.",
        "first_occurrence_usecs": "1682699539084721",
        "hidden_from_user": "false",
        "job_id": "18211",
        "job_name": "Test12",
        "job_type": "kOracle",
        "matchedTags": "WorkloadSource_kOracle",
        "object_id": "181",
        "object_name": "obj181",
        "run_id": "182",
        "run_start_time": "2023.02.07 11:21:00 Pacific Time",
        "run_url": "https://test.com",
        "severity": "Critical",
        "tenant_id": "d520840916/",
        "type": "kOracle"
      },
      "annotations": {
        "cause": "Testing DP alerts raise..",
        "description": "Backup of obj181 that is part of protection group Test12 of type kOracle failed with error Testing DP alerts raise",
        "help": "Please refer to KB for details/resolution.",
        "occurrence": "Start at 2023-04-28 16:32:19.084721 +0000 UTC, total 1 time."
      },
      "startsAt": "2023-04-28T16:32:19.084721Z",
      "endsAt": "0001-01-01T00:00:00Z",
    }
  ]
}

```



```

    "generatorURL": "",
    "fingerprint": "bfef9abae71570f0"
  }
],
"groupLabels": {
  "account_id": "00101000005nBps",
  "alertname": "ProtectedObjectFailed",
  "severity": "Critical"
},
"commonLabels": {
  "account_id": "00101000005nBps",
  "alert_category": "BackupRestore",
  "alert_code": "CE00610005",
  "alert_state": "Open",
  "alert_type_bucket": "DataService",
  "alert_type_id": "10005",
  "alertname": "ProtectedObjectFailed",
  "cluster_id": "1609127048663690",
  "cluster_id_str": "4327092961767844",
  "cluster_name": "DPCluster",
  "failed_objects": "obj181",
  "failure_reason": "Testing DP alerts raise.",
  "hidden_from_user": "false",
  "job_id": "18211",
  "job_type": "kOracle",
  "matchedTags": "WorkloadSource_kOracle",
  "run_start_time": "2023.02.07 11:21:00 Pacific Time",
  "run_url": "https://test.com",
  "severity": "Critical",
  "tenant_id": "d520840916/",
  "type": "kOracle"
},
"commonAnnotations": {
  "help": "Please refer to KB for details/resolution."
},
"externalURL": "https://helios-dev3-internal.cohesitycloud.co/alertmanager-d1",
"version": "4",
"groupKey": "{}/{account_id=\"00101000005nBps\",alertname=~\"^(?:ProtectedObjectFailed)$\",hidden_from_user=\"false\",matchedTags=~\"^(?:.*WorkloadSource_kOracle.*)$\",tenant_id=\"d520840916/\"}:{account_id=\"00101000005nBps\", alertname=\"ProtectedObjectFailed\", severity=\"Critical\"}",
"truncatedAlerts": 0

```

}


Silence Alert Notifications

Sometimes, it makes sense to silence alert notifications, such as during maintenance or testing windows.

You can silence alerts that match the rules you define in the Silence tab. Optionally, you can silence alerts for specific periods that you define. Once silenced, alerts are triggered and displayed on the Alerts page, but email or Webhook notifications are not sent.

To create an alert silence rule:

1. In **DataProtect as a Service**, navigate to the **Health > Silence** tab.
2. Click **Create > New Silence Rule**.

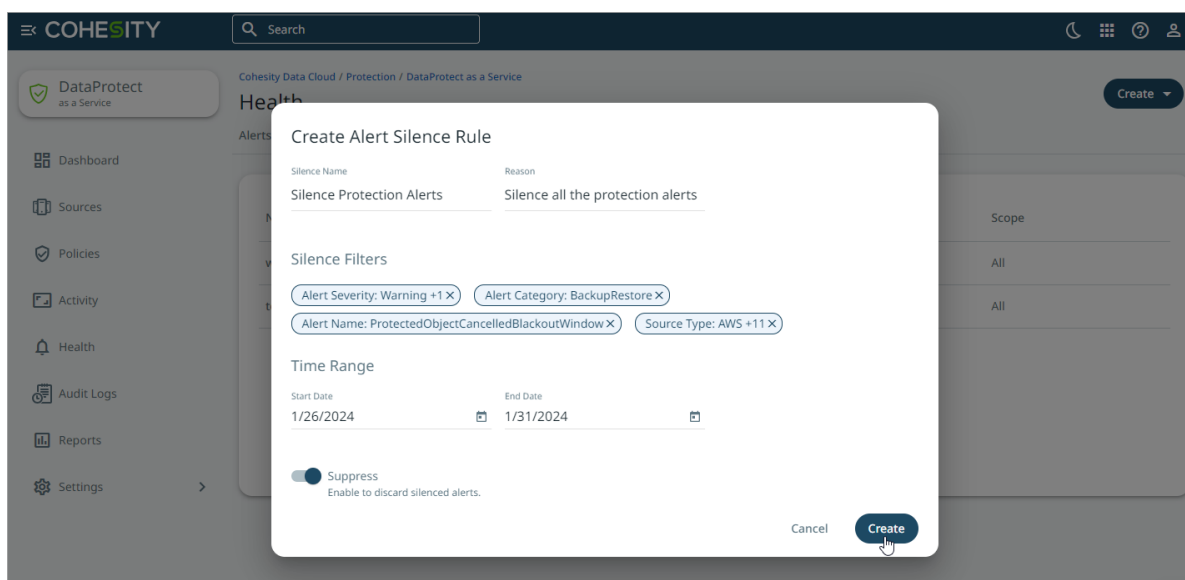


3. In the **Create Alert Silence Rule** dialog, perform the following:
 1. Enter a **Silence Name** for this alert silence rule and provide the **Reason** why you are creating the alert silence rule.
 2. In the **Silence Filters** section, select the filters based on your requirements:

Details	Description
Alert Severity	Select one or more severities from the drop-down you want to silence.
Alert Category	Select one or more categories from the drop-down you want to silence.
Alert Name	Select one or more names from the drop-down you want to silence.

Details	Description
Source Type	Select one or more sources from the drop-down for which you want the alerts silenced.

3. In the **Time Range** section, select a date in the **Start Date** and **End Date** fields to set the period within which the alert notifications must be silenced.
4. Enable **Suppress** if you do not want the alert to persist and appear on the **Alerts** page.



5. Click **Create**.

Alerts References

This topic provides details on all the alerts triggered by Cohesity DataProtect as a Service:

- [Archival and Restore Alerts](#)
- [Backup and Restore Alerts](#)
- [Security Alerts](#)

Archival and Restore Alerts

This topic covers provides details on all the alerts triggered by Cohesity DataProtect as a Service:

CE00820004 ArchiveJobFailed

Alert Description: The archive task for a job failed.

Reason: This alert is triggered when the archive task of a job fails. This alert is triggered only once per archive task. This alert can be caused by one of the following conditions:

- When an archive task fails due to external target connectivity and/or credentials issues.
- When an archive task fails due to an internal Cohesity issue.

Action: Check and fix external target connectivity or credential issues. If the archive task fails after fixing these issues, contact [Cohesity Support](#).

Severity: Warning

Dedup Interval: 604800 seconds

CE00820005 ArchiveJobStuck

Alert Description: The archive task for a job is stuck.

Reason: This alert is triggered when the archive task of a job is stuck and does not make any progress for more than 3 hours. If the archive task continues to be stuck/queued, this alert is triggered once a week per archive task. This alert can be caused by one of the following conditions:

- When an archive task doesn't make progress due to some external target connectivity and/or credential issues.
- When an archive task doesn't progress due to an internal Cohesity issue.

Action: Check and fix the external target connectivity or credential issues.

Severity: Warning

Dedup Interval: 604800 seconds

CE00820008 IceboxDedupCacheFull

Alert Description: The Icebox dedup cache is full.

Reason: The alert is triggered when the number of archived chunks maintained in the Distributed Key Value Store exceeds the default threshold value. When this happens, the effectiveness of deduplication for the archived data is impacted and could result in transferring more data to the external target.

Action: Contact [Cohesity Support](#).

Severity: Warning

Dedup Interval: 604800 seconds

CE00820001 MediaErrorDuringArchival

Alert Description: The archival job is waiting to correct an error.

Reason: This alert is triggered when no more tapes are available for archiving data.

Action: Add new tapes.

Severity: Critical

Dedup Interval: 604800 seconds

CE00820002 MediaErrorDuringRestore

Alert Description: The restore task is waiting to correct the error.

Reason: One or more tapes required to restore data are unavailable. This alert is triggered when the tape required to restore data is unavailable in the tape drive.

Action: Insert the required tapes to continue with the restore. The required tapes are listed in the alert.

Severity: Critical

Dedup Interval: 604800 seconds

CE00820007 RestoreJobFailed

Alert Description: The restore task failed.

Reason: This alert is triggered when the restore task of a job fails. This alert is triggered only once per archive task. This alert can be caused by one of the following conditions:

- When a restore task fails due to some external target connectivity and/or credentials issues.
- When a restore task fails due to an internal Cohesity issue.

Action: Check and fix external target connectivity or credential issues. If the restore task fails after fixing these issues, contact [Cohesity Support](#).

Severity: Warning

Severity: 604800 seconds

CE00820006 RestoreJobStuck

Alert Description: The restore task is stuck.

Reason: This alert is triggered when the restore task of a job is stuck and does not make any progress for more than one day. If the restore task continues to be stuck/queued, this alert is triggered once a week per Restore task. The alert can be caused by one of the following conditions:

- When a restore task doesn't make progress due to some external target connectivity and/or credentials issues.
- When a restore task doesn't progress due to an internal Cohesity issue.

Action: Check and fix external target connectivity or credential issues. If the restore task doesn't progress after fixing these issues, contact [Cohesity Support](#).

Severity: Warning

Dedup Interval: 604800 seconds

Backup and Restore Alerts

CE00608002 MissingVMBBackup

Alert Description: Missing VM backup.

Reason: This alert is triggered when a descriptor VMDK file has been backed up, but the corresponding flat VMDK file is missing from the backup.

Action: No action is required.

Severity: Critical

Dedup Interval: 3600 seconds

CE00610016 ObjectBackupSlaViolated

Alert Description: The SLA for the backup of an object is violated.

Reason: The alert can be triggered when the load on the Cohesity DataProtect as a Service is higher than anticipated, or the primary source is loaded, and the Cohesity DataProtect as a Service cannot back it up fast enough.

Action: Verify if a new workload is recently added to the Cohesity cluster or if the primary source is throttling Cohesity APIs/calls.

Severity: Warning

CE00610006 PolicyFieldsDeprecated

Alert Description: The policy settings in a policy have been deprecated

Reason: The alert is raised after the Cohesity cluster is upgraded from a 4.x release to a 5.x release and the cluster detects that some policy settings used in the current policies on the cluster have been deprecated.

Action: Open the listed policy in the Cohesity Dashboard, verify the current settings, and make any necessary adjustments. See the [ALERT: CE00610006 POLICYFIELDSDEPRECATED](#) KB article.

Severity: Warning

Dedup Interval: 86400 seconds

CE00610005 ProtectedObjectFailed

Alert Description: The backup of an object that is part of a protection run failed with an error.

Reason: The alert is raised when the Cohesity cluster detects that an object (such as a VM) failed to be backed up during a Protection Run. One alert is raised for each object (such as a VM) that failed to be backed up. For instructions on how to enable this alert, contact [Cohesity Support](#). A protection run can fail to back up an object for the following reasons:

- There is an issue with the primary environment, such as a removed VM or a Snapshot failure.
- The primary storage is full. (The primary storage contains the objects backed up by the Cohesity cluster.)
- The Cohesity Agent is unreachable while attempting to back up physical servers.

Action: See the [CE00610005 | BackupRestore - BackupObjectFailed](#) KB article for a resolution.

Severity: Critical

Dedup Interval: 86400 seconds

CE00610009 ProtectedObjectSLaViolated

Alert Description: The service level agreement violation (SLA) of an object in the protection run was violated.

Reason: The alert is triggered when the service level agreement violation (SLA) occurs for an individual object in a Protection run. A Protection run may take longer than the specified SLA for the following reasons:

- If the primary storage is slow.
- The network is slow.
- You specified SLA that is too short.

Action: Investigate why the Protection run took longer than the specified SLA. If appropriate, adjust the time period specified in the SLA.

Severity: Warning

Dedup Interval: 86400 seconds

CE00608003 VMCrackingSkipped

Alert Description: The VM contents are not indexed.

Reason: The alert is triggered when the Cohesity DataProtect as a Service detects 5 consecutive unsuccessful attempts to index a VM. The alert can be caused by the following conditions:

- The Cohesity DataProtect as a Service is not able to mount the VMDK.
- The VM Snapshot has an issue.

Action: No action is required.

Severity: Warning

Dedup Interval: 3600 seconds

CE00610014 VMMigrationIdentified

Alert Description: The VM(s) present in the vCenter have been identified to be migrated from other VCenter(s).

Reason: The alert is triggered when the Cohesity DataProtect as a Service identifies a VM in a vCenter that was earlier part of another vCenter registered on Cohesity DataProtect as a Service.

Action: No action is required if the migrated VMs are mentioned in the alert. If not, contact [Cohesity Support](#).

Severity: Critical

Dedup Interval: 86400 seconds

CE00610021 ProtectionPolicyModified

Alert Description: The Protection Policy was modified by a user.

Reason: This alert is triggered when a Protection Policy is modified. The modification might include any changes apart from DataLock-related changes in the policy.

Action: No action is required.

Severity: Informational

Dedup Interval: 86400 seconds

CE00610019 PolicyDataLockChanged

Alert Description: DataLock settings were changed in the Protection Policy.

Reason: This alert is triggered if you enable or disable DataLock for a Protection Policy.

Action: No action is required. Using this alert, you can validate if the DataLock was enabled or disabled by a valid user.

Severity: Informational

Dedup Interval: 86400 seconds

CE00610020 PolicyDataLockDurationChanged

Alert Description: DataLock retention for the Protection Policy was changed.

Reason: This alert is triggered when you change the DataLock duration for a Protection Policy.

Action: No action is required. Using this alert, you can validate if a valid user modified the DataLock configuration.

Severity: Informational

Dedup Interval: 86400 seconds

CE00610017 ProtectionPolicyDeleted

Alert Description: A Protection Policy was deleted.

Reason: This alert is triggered when you delete a Protection Policy.

Action: No action is required. Using this alert, you can validate if a valid user deleted the Protection Policy.

Severity: Warning

Dedup Interval: 86400 seconds

CE00610023 ProtectionRunModified

Alert Description: A protection run was modified.

Reason: This alert is triggered when a Protection run is modified. The modification might include deleting a Protection run, enabling a legal hold, etc.

Action: No action is required.

Severity: Informational

Dedup Interval: 86400 seconds

CE00610027 ObjectDeletionRejected

Alert Description: A protection run was modified.

Reason: This alert is triggered when the user deletes a specific snapshot for an object instead of deleting the entire view. The deletion is rejected because the view is marked immutable, and therefore individual object deletion can not be performed.

Action: Evaluate if the user can delete the entire view instead of individual snapshots.

Severity: Warning

Dedup Interval: 3600 seconds

Security Alerts

CE01516011 DataIngestAnomalyAlert

Alert Description: Anomalous change in data ingests detected on your Source, which might be a symptom of a potential ransomware attack on your primary environment.

Reason: This alert is triggered when an anomalous change in the data ingest rate for a protected Source is detected and is only generated if the cluster is registered with Cohesity DataProtect as a Service. The change might be a symptom of a ransomware attack on your primary environment.

Action: Consider restoring the Source from a Snapshot. This alert provides a link to begin an Instant Recovery using the latest clean Snapshot. For more information on detecting anomalies and ransomware attacks, see [Detect Anomalies](#) and [Detect Ransomware Attacks](#).

Severity: Warning

Dedup Interval: 3600 seconds

Audit Logs

The **Audit Logs** page records the events that occur in Cohesity DataProtect as a Service. The events are:

- Read or write actions performed by the users on Cohesity DataProtect as a Service.
- Login and logout actions performed by the Helios users.

View Audit Logs

On the **Audit Logs** page in Cohesity DataProtect as a Service, you can find the following details for the events that are logged by the registered regions:

- Date
- Time
- User & action
- System (Cohesity DataProtect as a Service region)

Note: By default, only the write actions performed by the users on Cohesity clusters are displayed on the **Audit Logs** page. To see read actions, select **Read Actions** from the **Actions** filter and click **Apply**. See [Use Filters to Locate Specific Logs](#) next.

Use Filters to Locate Specific Logs

Use the following filters to narrow the listed audit logs and locate the specific logs.

Filter	Purpose
Date Range	Filter the audit logs based on the selected time window.
System	Filter the audit logs based on the Cohesity DataProtect as a Service regions.
Users	View the audit trails of specific users.
Category	Filter the audit logs based on predefined categories. See Review Audit Log Categories next.
Action	Filter the audit logs based on the read or write actions performed by the users in the registered regions. See Logged Actions below

Review Audit Log Categories

Audit logs are logged under predefined categories for you to find the relevant audit logs and analyze the correct logs quickly.

- API Key
- Access Token
- Alert
- Alert Notification Rule
- Group
- Helios Event
- IDP Configuration
- Protection Group
- Protection Policy
- Recovery Task
- Region
- Resolution
- SaaS Connector
- Snapshot
- SNMP Config
- Source
- Tenant
- User

Logged Actions

Along with the read actions, the following write actions are logged:

Write Actions	Descriptions
Accept	A user accepted the license agreement.
Activate	A user activated an entity such as Protection Run.
Add	A user added a Region.
Apply	A user applied a setting or configuration.
Assign	A user assigns a source to a tenant.
Cancel	A user canceled an entity such as a running Protection run or a Recovery task.
Clone	A user cloned an entity such as a Snapshot, VM, or SQL Database.

Write Actions	Descriptions
Close	A user closed an SMB file.
Cloud Spin	A user deployed a VM on the cloud.
Cluster Expand	A user expanded the cluster.
Create	A user created an entity such as a Protection run.
Deactivate	A user deactivated a Protection run.
Delete	A user deleted an entity such as a Protection run, or Protection Policy.
Disjoin	A user disjoined the Cluster from an AD domain.
Download	A user downloaded a VMX file or a file from a VM Snapshot.
Import	A user performed a generic action for any import operations. For example, the user has imported patch binary.
Install	A user performed a generic action for any installation. For example, the user has installed an app.
Join	A user joined the Cluster to an AD domain.
Login	A user logged in to the Cohesity DataProtect as a Service.
Logout	A user logged out of the Cohesity DataProtect as a Service.
Mark	A user marked an entity for removal such as a disk.
Modify	A user modified an entity such as a User, Protection run, or Remote Cluster.
Notification Rule	A user modified the notification rule.
Overwrite	A user performed an overwrite operation.
Pause	A user paused an entity such as a running Protection run.
Recover	A user recovered an entity such as a VM, file, or SQL Database.

Write Actions	Descriptions
Refresh	A user performed a refresh of the entities in the Cohesity DataProtect as a Service. For example, the user refreshed the source configuration.
Register	A user registered an entity such as an External Target (Vault).
Mark Removal	A user marked an entity for removal. For example, the user marked a disk for removal.
Rename	A user renamed an entity.
Restart	A user restarted a service.
Resume	A user performed a resume action on a Protection run.
Revert	A user reverted a setting or action.
Run Diagnostics	A user ran a diagnostics. For example, the user ran diagnostics on the agent to collect logs and other metrics.
Run Now	A user performed a Run Now action on a Protection run.
Schedule	A user scheduled an event.
Schedule Report	A user scheduled an email report.
Search	A user searched for a term.
Start	A user started a service.
Stop	A user stopped a service.
Unassign	A user removes a source from a tenant.
Uninstall	A user uninstalled an app.
Unregister	A user unregistered an entity such as a Source.
Update	A user updated an entity in a Cohesity cluster.
Upgrade	A user upgraded the Cohesity cluster.


Write Actions	Descriptions
Upload	A user uploaded an entity.
Validate	A user validated an entity.

Set Log Retention Period for Cluster Audit Logs

You can set the retention period for cluster audit logs. When you set a retention period, the logs are retained on the cluster until the retention period ends.

Note: The default retention period is 180 days. The minimum retention period is 90 days, and the maximum is 365 days.

To set a retention period for cluster audit logs, follow the steps below:

1. In **DataProtect as a Service**, navigate to **Security > Audit Logs > Settings**.
2. In the **Settings** tab, click the edit icon for **Log Retention Period**.
3. Enter the desired number and choose a type of retention period (Days, Weeks, Months, or Years).
4. Select the  icon to save.

A push notification with the message **Settings Updated** is displayed.


Cohesity converts weeks, months, or years into days and displays it as the **Log Retention Period**.

Download Audit Logs

You can download the Audit Logs in CSV format from Cohesity DataProtect as a Service for analysis and sharing.

Note: The downloaded .CSV file contains more details than what the Helios Dashboard displays. For example, the file contains details about the IP addresses of the systems from which the cluster is accessed, tenants, impersonation, and so on.

To download audit logs:

1. In **DataProtect as a Service**, navigate to **Audit Logs**.
2. In the top right, click the **Download**  icon.



The audit logs CSV file is downloaded.

Subscription Status

Cohesity Helios displays banners on the UI, providing details on your Cohesity DataProtect delivered as a Service subscription status, allowing you to take necessary actions. The banners are of three types:



- **Information**

Sample:


 Your Cohesity DataProtect delivered as a service (1 BETB); AWS data plane paid subscription has expired. Contact your Cohesity account team as soon as possible.
 


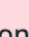
- **Warning**

Sample:


 Your Cohesity DataProtect delivered as a service (1 BETB); AWS data plane subscription will expire in 3 day(s). Renew your subscription now to continue.
 

- **Critical**

Sample:


 Your Cohesity DataProtect delivered as a service (1 BETB); AWS data plane subscription has expired. Contact your Cohesity account team immediately.
 

Banner Messages

Based on your subscription type and status, DataProtect delivered as a Service UI displays different types of banners. The table below shows the various scenarios and the types of banners displayed in each scenario:

Subscription Type	Subscription Status	Description
Free Trial	Expiration	Before the free trial expires, an information banner is displayed 15 days prior, and a warning banner is displayed 7 days prior to the expiry.
	Post Expiration	A day after the free trial expires, the DataProtect service UI displays the following message: "Your Cohesity DataProtect Delivered as a Service - Free Trial (1 FETB) free trial has expired. Contact your Cohesity account team for extension or purchase."
	Grace Period	After the free trial period ends, access to the service will be restricted immediately with no grace period.

Subscription Type	Subscription Status	Description
Paid Subscriptions	Expiration	DataProtect as a Service UI shows a banner 30 days before the subscription expires, a warning at 15 days, and critical after expiry.
	Post Expiration	A day after the paid subscription expiry, the following banner is shown on the DataProtect delivered as a Service UI: "Your Cohesity DataProtect delivered as a service (1 BETB); AWS data plane subscription has expired. Contact your Cohesity account team immediately."
	Grace Period	Once a paid subscription expires, there is a grace period to renew it. During this time, access is unrestricted, but product functionality is limited.

Sample Banner Messages

The following are different banner messages that provide details on the Cohesity DataProtect delivered as a Service subscription status:

3 days left in the grace period:

The screenshot displays the Cohesity DataProtect delivered as a Service UI. At the top, an orange banner with a warning icon states: "Your Cohesity DataProtect delivered as a service (1 BETB); AWS data plane subscription will expire in 3 day(s). Renew your subscription now to continue." Below the banner, the UI shows the Cohesity logo and a search bar. The main content area is titled "Regions" and features a world map with two red dots indicating active regions. To the right of the map, statistics are shown: 2 Regions, 0 Sources, and 0 Bytes Data. Below the map, a table lists the regions and their status.

Region	Encryption Option	Sources	Data Stored	Status
AWS Asia Pacific (Mumbai)	Cohesity KMS	0	0 Bytes	Failed to provision

The subscription has expired:

<1/2>

ⓘ

Your Cohesity DataProtect delivered as a service (1 BETB); AWS data plane paid subscription has expired. Contact your Cohesity account team as soon as possible.

×

☰COHESITY

🔍

Search

🌙

☰

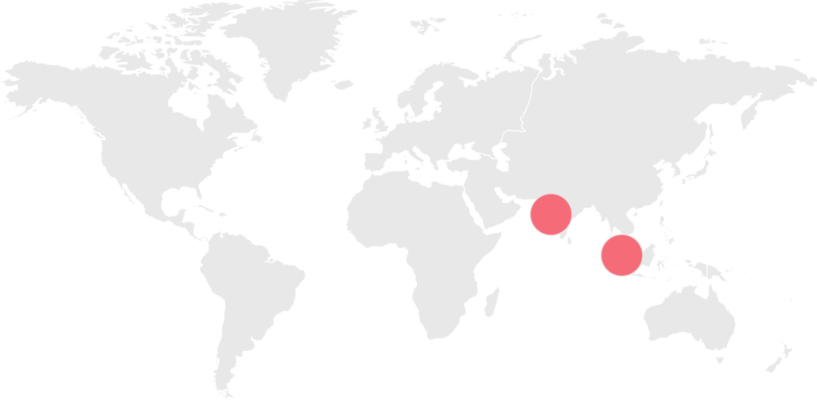
?

👤

Cohesity Data Cloud / Protection / DataProtect as a Service

Add Regions

Regions



2

Regions

0

Sources

0

Bytes

Data

Region	Encryption Option	Sources	Data Stored	Status
AWS Asia Pacific (Mumbai)	Cohesity KMS	0	0 Bytes	<div>⚠️ Failed to provision</div>

🔍

How-To Videos

Use these [videos](#) to learn some of the key tasks you'll be performing in Cohesity DataProtect as a Service in detail.

Cohesity Support

Reach Cohesity Support

There are several ways to create a Cohesity support case.

- Go to [Cohesity Support](#), to search in our knowledge base; or contact us by phone - United States and Canada: 1-855-9CO-HESI (926-4374), option 2.
- Log in to the [Cohesity Support Portal](#) to create a new case.
- Click the (?) icon on the Cohesity UI and select Support Portal.

Creating a customer support case for Cohesity Cloud Services (CCS)

When creating a customer support case for Cohesity Cloud Services (CCS), follow the steps listed below:

1. Mention CCS in the subject and select **CCS** as the **Issue Type**.
2. Provide the case information.
3. Edit the **Case Subject** as per your cloud region. For example, for AWS region, **CCS (AWS_Region): <Input Issue Subject Information>**.
4. Update the **Issue Type** field to **CCS**.

Additionally, provide the **Cluster ID** and the **Support Token** information if a SaaS connector is involved.

Support/Service Assistance

First contact the Service Provider that you have contracted for service and support. If you work directly with Cohesity and have a product warranty/entitlement, repair pricing or technical support related question, see your options below:

- To find solutions to your product issues or for suggestions or best practices, visit [Cohesity Knowledge Base](#).
- Log in to the [Cohesity Support Portal](#) to create a new case.
- To monitor your open cases, log in to the portal, click the **Cases** tab on the home page. This page should have all the case statuses and updates. You can also view individual case status.

Cohesity Software Running on Partner Hardware

For Cohesity software running on qualified third-party hardware, the following support workflow applies:

1. The customer may contact Cohesity Support first if the issue cannot be determined as a hardware issue.

Note: Cohesity cannot process hardware replacement requests for partner hardware.

2. Cohesity Support triages the issue. If it is a software issue, Cohesity Support continues to work on it.
3. If it is a hardware/firmware issue or is suspected to be a hardware/firmware issue, Cohesity provides information about the issue to the customer and requests that the customer open a support ticket with the appropriate partner.
4. If needed, Cohesity Support can join a three-way call with the partner and the customer.
5. The customer informs Cohesity Support on the progress of the partner's case.

