COHESITY

Version 1.4

July 2024

Cohesity Data Protection for Microsoft 365—Best Practices and Recommendations

ABSTRACT

This guide provides the best practices and recommendations for Cohesity Data Protection for Microsoft 365 (Customer-managed DataProtect and Cohesity-managed DataProtect delivered as a Service).

Table of Contents

Tables

IMPORTANT:

- Sections without superscripts apply to both DataProtect and DataProtect delivered as a Service.
- Sections with superscript [DataProtect] are applicable only to customer managed DataProtect.

ABBREVIATIONS USED:

- M365: Microsoft 365
- AD: Active Directory

Overview

Cohesity Data Protection for Microsoft 365 provides a simple, fast, and cost-effective data protection solution to protect mission-critical data hosted on Microsoft 365 (Exchange Online, OneDrive for Business, SharePoint Online, and Microsoft Teams).

Objective and Intended Audience

This document is intended to be used by IT Admins, M365 Admins, Professional Service Engineers, and System Engineers who want to use Cohesity DataProtect for M365 to protect their or their customer's M365 environment. This document assumes that the reader has a basic understanding of:

- Microsoft 365 and its applications
- Microsoft Azure
- Cohesity DataProtect

Scope of this Document

This document provides you with the best practices and recommendations for using Cohesity DataProtect for M365. This scope of this document does not include giving you expertise on either M365 or Cohesity DataProtect for M365.

LEARN MORE:

- Microsoft 365 Documentation
- Microsoft Azure Documentation
- Cohesity DataProtect Documentation
- Cohesity M365 Protection Documentation

Protecting Your Microsoft 365 Environment

Choose Your Deployment Model

Cohesity provides two flavors of DataProtect.

Table 1: Deployment Model

#	SOLUTION	DESCRIPTION	RECOMMENDATION
1.	Customer-managed Cohesity DataProtect	Customer is responsible for: • CapEx—compute, storage, and network. • OpEx—data center, power, cooling, admin, and maintenance.	Choose this option: • If you want absolute control of your data protection environment. • If you have strict data sovereignty requirements to keep your data on private cloud.
2.	Cohesity-managed DataProtect delivered as a Service	• Cohesity is responsible for all the CapEx and OpEx expenses. • Customer is responsible only for the subscription fees.	Choose this option: • If you are looking for simplicity and do not want to maintain the infrastructure and require an as-a-service model. • If you are already a Cohesity DataProtect user and are looking for a hybrid model.

LEARN MORE:

- Cohesity DataProtect for M365
- Cohesity DataProtect delivered as a Service

M365 Source Registration

Table 2: M365 Source Registration

#	METHOD	RECOMMENDATION
1.	Cohesity-assisted Express Registration	• Cohesity always recommends choosing this method. • It will automatically create Azure app on your behalf, thereby saving time and reducing manual errors.
2.	Manual Registration	Choose this method only if: • You want absolute control of the Azure App such as its name. • You already have created an Azure app or want to utilize an existing app.

LEARN MORE:

COHESITY DATAPROTECT	COHESITY DATAPROTECT DELIVERED AS A SERVICE
• Express Registration • Manual Registration	• Express Registration • Manual Registration

Azure App

NOTE:

- M365 Global Administrator is the only user in M365 who can create an app in Azure.
- The M365 user account with minimum permissions suggested by Cohesity is only used for the backup/restore workflows and not for the Azure app creation. Even with the express registration, you must login to Microsoft via M365 Global Admin credentials for Cohesity to successfully create the Azure apps on your behalf.

Note: Cohesity does not have access to your M365 Global Admin credentials. Also, Cohesity does not store or cache your M365 Global Admin credentials.

- By default, express registration creates one Azure application. If one is not sufficient to meet your SLAs, Cohesity can help you configure a higher number of Azure applications. Contact Cohesity support to learn about adding more Azure apps.

  Note: Microsoft does not recommend adding more than one Azure app. So be mindful when choosing multiple Azure apps, and do not increase it without contacting Cohesity Support.

BEST PRACTICES:

- Azure app should be dedicated only to Cohesity.
- Do not use an Azure App which is already used for other purposes.

LEARN MORE:

- Azure App permissions required by Cohesity
- SharePoint Online permissions required by Cohesity

User Account

NOTE:

- The user account needed to register your M365 source in Cohesity does not require an M365 license unless you want to protect the Public Folders.
- Ensure that multifactor authentication is disabled for the user account.

RECOMMENDATION:

- Do not use your Global Admin account for M365 source registration in Cohesity. Use a user with minimum permissions specified by Cohesity for registering your M365 source.
- Use a dedicated user account with minimum specified permissions only for Cohesity and do not use this user account for any other operations or to use any other product.

LEARN MORE:

- User permissions required by Cohesity

Cluster Setup [DataProtect]

RECOMMENDATION:

For optimal performance in a large-scale environment the following best practices should be considered:

- Have a dedicated cluster for M365 workload. i.e., avoid having other workloads (such as VMs, Database, NAS, SAN etc.) on the cluster which is being used for M365 protection.
- Get your cluster sized for M365 workload by the Cohesity team for optimum results.

Windows Connector [DataProtect]

NOTE:

- Configure Windows Connector only if you want to use the PST export feature for Exchange Online.
- For Cohesity versions below 6.8.1, Windows Connector is required for SharePoint protection as well.
- A Windows Connector must be registered as a physical source on your Cohesity cluster.
- A single Windows Connector can serve a maximum of five M365 sources.

RECOMMENDATION:

- Windows Connector should be dedicated to M365 data protection and should not be a part of any other workflows provided by the Cohesity Data Cloud.
- If your Cohesity cluster is part of an AD, then make sure the Windows Connector is also a part of the same AD.

LEARN MORE:

- Firewall and Port requirements
- Setting up Windows Connector

Protection Policy

NOTE:

- There are four predefined policies available by default. In addition to these, you can create your own customized policies.

RECOMMENDATION:

- If you have enabled archival, then you can reduce the retention of the primary backup for space savings. [DataProtect]
- Always use daily incremental backups for space savings and increased performance.
- For compliance and regulatory requirements enable the data lock. Data lock ensures your backup data cannot be tampered (modified/deleted) by any user.

  Note: Only a user with Data Security role can add, modify, or remove a data lock.
- Set Retry **Options** to **0**, as M365 protection in Cohesity already has the retry mechanism to handle failed requests. You need not increase the value for this setting in the Cohesity Protection Policy.
- Do not schedule multiple backup jobs to run at the same time. Schedule different backup jobs at different time of the day for increased efficiency.
- For better performance avoid having backup frequency less than 24 hours. [DataProtect]

  Note: Workload here means data protection source which you add in Cohesity cluster, such as VMs, Database, NAS, SAN, M365, etc**.**

LEARN MORE:

- Creating Policies
- Managing Policies
- Predefined Policies
- Extended Retention

Protection Group [DataProtect]

NOTE:

- For backup schedule, Cohesity allows you to set the day for weekly backup and any day of the week for monthly backup.
- If you specify multiple extended retention rules for the same set of snapshots with different retention periods, the snapshots are retained for the longest specified time period.

  Example: If you define two extended retention rules that retain the same set of snapshots, but one rule specifies a retention period of 90 days, and another specifies a retention period of 180 days, the set of snapshots are retained for 180 days.

RECOMMENDATION:

- When you have a large number of objects, make sure to create multiple jobs with a maximum of up to 8000 objects per job, instead of creating 1 job with all in one. Contact Cohesity Support to get help with the job splitting tool, which can assist you in splitting the jobs efficiently.
- Schedule the backup time during the non-working hours of your organization, when your M365 is not busy.
- Stagger the job times for different jobs. i.e., schedule different backup time for each job, so that all the jobs won't begin at the same time. This can prevent the SLA violations during the incremental runs.
- Always use the default QoS Policy—**Backup HDD**, which writes data directly to HDD. Use the policy **Backup SSD** if you need fast ingest speed for a small number of Protection Groups. Use policy **Backup Auto** if you want to reduce SSD wear out (in comparison to policy Backup SSD).
- Do not add the same objects to multiple Protection Groups.
- Use Exclusion List to exclude unwanted folders from the backup. For example, in Exchange Online, if you do not want to back up folders such as junk emails or any custom created folder where you have saved unwanted emails, then you can add that folder/s under Exclusion List to exclude them from the backup.

  Note: Exclusion list is not supported for Public Folders.
- Exclude unwanted folders from indexing for faster indexing, resulting in faster search results.

  Note: If you exclude any folder from indexing, then you cannot search the contents within these folders for granular recovery.
- Ensure to add alerts on failures for the Protection Group, so that you receive notifications if any backups fail.

LEARN MORE:

QoS Policies, Exclusion List, Indexing, Index Exclusion, Alerts

Backup

NOTE:

- The initial backup will be a full backup and would take a significant amount of time compared to subsequent incremental backups. The initial backup will also be throttled more by Microsoft due to the significant amount of data in the backup. Due to these reasons, the backup SLA can be violated for this first full backup.

  Note: For Exchange Online workload, consider raising a Microsoft ticket for temporary relaxation of throttling.

- Optional practices during initial full backup:
  - Perform stage-wise protection of the services. Do not protect all the services together. Finish the protection of one service completely before starting it for another service. [DataProtect]

    Example: Begin the protection of Exchange Online for all the mailboxes. Once it completes, then begin the protection of other services (OneDrive, SharePoint, Teams) sequentially one after the other (after each service finishes the protection completely).

    This is because if you protect all the services in one go, then they will wait in a queue and might give the impression that the backup is taking too long, whereas the backup has not even started for that service as it is waiting for its turn in the queue. If you are not concerned about it, then you can choose to protect all the services together.
  - While performing stage-wise protection of the services as described above, consider pausing the incremental runs for the finished jobs until all other jobs are completed. [DataProtect]

    This is just to make sure that the full backup jobs get priority and completes quickly.
  - For Exchange Online, exclude the “In-place archive” mailboxes from the Protection Job, as access to data in in-place archives is much slower than the regular mailboxes. Once the backups of regular mailboxes are completed successfully, you can remove the exclusion and include the “In-place archive” mailbox on the Protection Job, so that they will be protected in the next run.
  - You can follow the same process as mentioned above for Shared Mailboxes during the initial full backup.
  - Consider excluding SharePoint templates. Enable templates after initial full backup is complete.
- Initial full backups are likely to be slow due to throttling from Microsoft end. Do not cancel the jobs as it will start from scratch when you restart.
- Complete any M365 migrations before running backups. i.e., If you are migrating your data from On-prem Exchange Server to M365 Exchange Online, then start the backups after the migration is complete. Running backups during migration can lead to increased throttling from Microsoft, which in-turn can lead to slower backups and violated SLAs.
- Do not run two backup software (e.g., Cohesity and a competitor) in parallel for the same M365 source.
- There are no egress charges in M365. So, any backups done to your on-prem environment or to Cohesity managed service does not cost any egress charge.
- If a backup job completes partially with certain objects skipped, then the skipped objects will be backed up in the next backup job.

RECOMMENDATION:

- Always follow the 3-2-1 backup rule. Configure your replication and archival [DataProtect].
- Use auto-protect objects to avoid manual intervention. Auto-protect will make sure to protect all your future objects added to your M365 automatically.

LEARN MORE:

- Auto-protect
- Protect Exchange Online Mailboxes
- Protect Public Folders
- Protect OneDrives
- Protect SharePoint Online Sites
- Protect Microsoft 365 Teams
- Protect Microsoft 365 Groups

Recovering Microsoft 365 data

Restore Type

Table 3: Restore Type

#	RESTORE TYPE	THINGS TO NOTE
1.	Complete restore—Restore entire application data (Restore Mailbox, OneDrive, SharePoint Site, or Team).	If the size of the application data (such as mailbox) is large, then restoring takes a longer duration.
2.	Granular restore—Restore individual items within an application. • Emails, folders, or subfolders in Exchange Online • Files, folders, or subfolders in OneDrive, SharePoint Online, and Teams	• If the item already exists at the destination, then the item will be overwritten in case of OneDrive, SharePoint, and Teams. • If the item already exists at the destination, then the item will be skipped in case of Exchange Online.
3.	Download offline data of backup (pst for Exchange Online, zip for OneDrive, and SharePoint Online).	• If the size of the mailbox is large, then exporting the mailbox to pst format takes longer duration. • The pst export process might take a few minutes or hours depending on the size of the mailbox. • The exported pst of the mailbox or the email is valid for 72 hours. Ensure that you download the pst file within 72 hours of the recovery task completion. • pst export can be protected with password for more secure access.

Restore Location

Table 4: Restore Location

#	Restore Location	Things to note
1.	Restore to the original location on the same M365 tenant	• If the item already exists at the destination, then the item will be overwritten in case of OneDrive, SharePoint, and Teams

#	Restore Location	Things to note
		• If the item already exists at the destination, then the item will be skipped in case of Exchange Online.
2.	Restore to an alternate location on the same M365 tenant	• If a folder with the specified name does not exist at the destination, Cohesity creates the folder and recovers the data to it. • Permission and metadata attributes are not restored.
3.	Restore to a different M365 tenant	• Both the source and destination M365 tenants must be registered as sources in Cohesity Data Cloud. • If a folder with the specified name does not exist at the destination, Cohesity creates the folder and recovers the data to it. • Permission and metadata attributes are not restored.

Exchange Online

NOTE:

- A Windows Connector is required to export mailboxes to pst. See section Windows Connector for more information. [DataProtect]
- Restoration at the Protection Group level is not supported for M365.
- Restoring items in the same mailbox from different recovery points is not supported in the same recovery task. You need to create different restore tasks for different recovery points.
- Schedule the restores during non-working hours of your organization, when your M365 is not busy.

LEARN MORE:

- Restore Considerations
- Restore Mailboxes
- Restore Emails and Folders

OneDrive For Business & SharePoint Online

NOTE:

- Use search and restore if you know the filename or other attributes. This would save time in recovery compared to browsing and restoring.
- If you are not sure about the search attributes for the file or folder, use browse and restore, where you can browse the backup contents and locate the file/folder manually.
- Empty folders cannot be downloaded.
- If the file is being restored to original location and if it already exists in the folder, then the file is restored as the latest version.
- The contents of SharePoint Document Library should not be in checked out state.

LEARN MORE:

- Restore OneDrive
- Restore OneDrive Files and Folders
- Restore SharePoint Online Sites
- Restore SharePoint Online Files and Folders

Microsoft Teams

NOTE:

- If a source team is deleted from the M365 tenant after its backup, then during recovery, Cohesity creates a team with the same name and recovers the team's data to the newly created team.
- The **Restore Original** Owner Members option is not applicable if you are restoring the Team's data to a different Microsoft 365 tenant.
- If you are restoring to a different team, the channels are restored with suffix: `<_coh><some_int_value>`. Where _coh is appended to the channel name and some_int_value is a random integer value.
- Only public channels can be restored, not private channels.

LEARN MORE:

- Restore a Team
- Restore Teams Files
- Teams Metadata Restore Scenarios

Further Reading

Table 5: Additional Info

#	TOPIC	DETAILS
1.	Replication [DataProtect]	Cohesity enables you to replicate M365 data to another Cohesity cluster for disaster recovery use cases. For more details, see: • Replicate M365 data • Replication and remote access setup
2.	Archival [DataProtect]	Cohesity enables you to archive the M365 data to a designated external target, such as tape or cloud storage. For more details, see: • Archive Exchange Online data • Archive OneDrive for Business data • Manage External Targets
3.	Multitenancy [DataProtect]	Cohesity cluster supports the multitenancy feature. The multitenancy feature enables you to configure the multitenant environment on your Cohesity Data Cloud platform and securely isolate each tenant in your environment. For more details, see Multitenancy on Cohesity Cluster
5.	Security and Privacy	Cohesity is secure by design and employs industry leading security and privacy features to secure your data. For more details, see: • Cohesity Security Features • Cohesity Security Practices • Cohesity Trust Centre
6.	Access Management	For more details on access management, RBAC users, groups, and roles see: • Access management in Cohesity DataProtect. • Access Management in Cohesity DataProtect delivered as a Service.

#	TOPIC	DETAILS
7.	Monitoring	For more details on monitoring, reporting, and audit logs, see: • Monitoring in Cohesity DataProtect. • Monitoring in Cohesity DataProtect delivered as a Service.
8.	Networking [DataProtect]	For information on networking, see Cohesity Networking.
9.	How to videos	See How-to-videos for detailed product videos.
10.	Support	See Cohesity Support for more details on support.
11.	Cohesity REST APIs	See Cohesity Developer Portal to get more details on Cohesity REST APIs and supported integrations.

Your Feedback

Was this document helpful? Send us your feedback!

About the Author

Shashanka SR is a Technical Solution Engineer at Cohesity. In his role, Shashanka focuses on Microsoft 365 and Salesforce Protection.

Other major contributors include:

- Aditya Vasudevan, VP Customer Success
- Karthick Radhakrishnan, Director, Technical Solution Engineering
- Ravi Luhadiya, Director Engineering
- Kunal Bose, Product Solution Engineer
- Prajakta Ayachit, Staff 2 Engineer

Document Version History

VERSION	DATE	DOCUMENT HISTORY
1.4	July 2024	Republishing
1.3	Apr 2024	Content updates
1.2	Apr 2023	Minor updates
1.1	Mar 2023	Content updates
1.0	Dec 2022	First version

ABOUT COHESITY

Cohesity is a leader in AI-powered data security and management. Aided by an extensive ecosystem of partners, Cohesity makes it easier to protect, manage, and get value from data – across the data center, edge, and cloud. Cohesity helps organizations defend against cybersecurity threats with comprehensive data security and management capabilities, including immutable backup snapshots, AI-based threat detection, monitoring for malicious behavior, and rapid recovery at scale. Cohesity solutions are delivered as a service, self-managed, or provided by a Cohesity-powered partner. Cohesity is headquartered in San Jose, CA, and is trusted by the world's largest enterprises, including six of the Fortune 10 and 42 of the Fortune 100.

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