



TECHNICAL WHITE PAPER

How It Works: End-to-End Microsoft 365 Resilience

Mike Preston, Alpika Singh
RWP-0639

Table of Contents

INTRODUCTION	3	Data Integrity and Availability	16
Audience.	3	Data Storage	16
Objectives.	3	Multi-Geo Support	16
Challenges	3	Data Durability	17
The Rubrik Approach.	5	Operational Security	17
Automate Protection and Manage Sensitive Data Risk	5	Soft Delete	17
Withstand Cyber Attacks	5	Compliance – 3rd Party Validation and Certifications	17
Recover Rapidly	6	Working around API Throttling and Failures	17
ARCHITECTURE AND COMPONENTS	6	Handling M365 Throttling	18
Components.	7	Graph to Application Failover/Failback	18
Entra ID, Service Principals, and Enterprise Applications	7	Data Encryption	18
Azure Kubernetes Service (Exocompute)	8	In-Flight Encryption	18
Customer-Specific Storage Account (Azure Blob)	8	At Rest Encryption	18
Azure Key Vault	8	Key Rotation	21
Microsoft 365 Security	8	ReKeying 22	
Authentication, Authorization, and Accounting (AAA) in Rubrik	9	Key Backups	23
Authentication	9	Recovery.	24
Authorization	9	Exchange Recovery Options	25
Accounting	10	OneDrive Recovery Options	25
Data Encryption	10	SharePoint Recovery Options	26
Data Immutability	10	Teams Recovery Options	26
SLA Retention Lock	10	Mass Recovery	26
Quorum Authorization	10	Prioritized Data Recovery	28
Network Security	11	Cross Subscription Restoration	29
HOW IT WORKS	11	Self-Service Recovery.	30
Initial Configuration of M365 within RSC.	11	Data Threat Analytics	31
Configuration	12	Anomaly Detection	31
Service Principals and Enterprise Applications	12	Threat Monitoring	33
SLA Domains	13	Turbo Threat Hunting	34
Application Assignment	14	Data Security Posture Management	35
Group Assignment	14	Data Discovery and Classification	36
User Assignment	14	Data Access Governance	37
Site Collection Assignment	14	MIP Labeling	39
Teams Assignment	14	SUMMARY	40
Protection	14	APPENDICES	41
Initial Full and Subsequent Backups	16	Appendix A: Required Microsoft 365 API Permissions. .	41
		VERSION HISTORY	42

INTRODUCTION

Welcome to *How It Works: Microsoft 365 (M365) Protection*. The purpose of this document is to aid the reader in familiarizing themselves with the features, architecture, security, and workflows of Rubrik Security Cloud (RSC) Microsoft 365 protection. Such information will prove valuable while evaluating, designing, or implementing the technologies described herein.

AUDIENCE

This guide is for anyone who wants to better understand the capabilities of Microsoft 365 protection on the Rubrik Security Cloud platform and the technical architectures and security that underpin those capabilities. This includes architects, engineers, and administrators responsible for the Microsoft 365 environment and data protection operations, as well as individuals with a vested interest in security, compliance, or governance.

OBJECTIVES

This guide aims to provide the reader with a clear and concise technical reference regarding architecture and workflows utilized by Rubrik Security Cloud's Microsoft 365 protection. After reading this document the reader should understand:

- Why M365 data requires robust data protection
- The problems that Rubrik M365 protection solves
- The architecture of Rubrik M365 protection
- How Rubrik secures and protects M365 data
- How Rubrik M365 protection works

CHALLENGES

Microsoft 365, a widely used productivity suite in the corporate sphere has become an enticing target for cybercriminals due to the abundance of valuable and sensitive information it contains. Various malicious activities, such as ransomware attacks and account takeovers, can lead to the encryption or deletion of crucial data and pose a significant risk to organizations.

While Microsoft offers several native security options to its M365 customers, organizations are ultimately responsible for their data. As illustrated in *Figure 1*, the sole responsibility of information and data is always left up to the customer, whether delivered through SaaS, PaaS, IaaS, or on-premises.

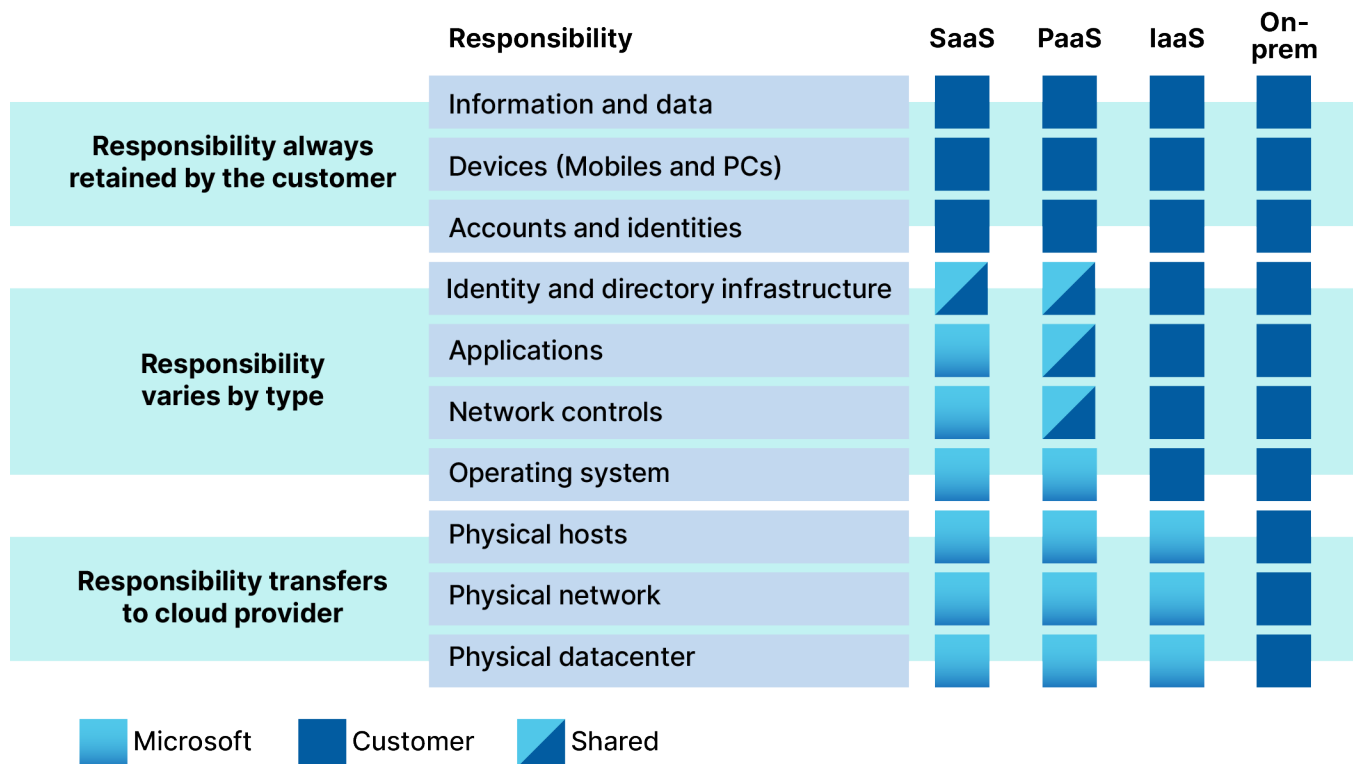


Figure 1 – Microsoft Shared Responsibility Model

Microsoft’s primary responsibilities hinge on tasks related to delivering and managing backend infrastructure and providing resilience in their service delivery. On the flip side, customers, or the data owners, are responsible for the data contained within M365 services. This includes identifying and securing sensitive data, ensuring proper access and control of your data, and ensuring mission-critical data is backed up and secure. For instance, Microsoft is of little to no help if a customer accidentally or purposely deletes all M365 data through a rogue retention policy. The customer would be responsible for that data and its restoration.

Numerous organizations find themselves ill-equipped to handle a potential breach of their M365 environment, resulting in an increased risk of losing access to critical information if compromised. For this reason, organizations must employ a resilient solution that integrates the principles of a zero-trust model. This includes all components required to protect M365, including the platform itself, as well as the underlying Entra ID environment within Azure. By doing so, they can ensure that if their environment is compromised, they can not only recover but recover in the most efficient means possible.

THE RUBRIK APPROACH

Rubrik approaches Microsoft 365 protection by focusing on three key areas: Automating Protection and Managing Sensitive Data Risk, Withstanding Cyber Attacks, and Recovering Rapidly.



Figure 2 – Rubrik's Microsoft 365 Value Props

Let's dive into each area in a bit more detail.

Automate Protection and Manage Sensitive Data Risk

Microsoft 365 is a fluid environment that is constantly changing—new employees getting onboarded to organizations, which results in mailboxes and OneDrive accounts being created and deleted daily. We see new SharePoint sites constantly being instantiated, and Teams contains more and more teams, messages, key attachments, and files daily. With the heavy adoption of M365, we also see a huge rise in sensitive data and intellectual property being housed within the platform, posing a constant risk of infiltration, causing reputational damage and lost business. Legacy approaches that involve creating backup jobs containing specific users and objects simply can't keep up with the constant change of M365. Organizations need to be able to automate the protection, management, and retirement of data within their M365 backup solution. Rubrik minimizes management overhead and configuration delivering data protection for M365 by both the automatic assignment of SLA Domains to objects and intelligent and fully orchestrated task scheduling for large M365 environments. On top of this, Rubrik's robust centralized management and reporting capabilities help organizations ensure their data complies with any regulations they must adhere to. Rubrik's automated data protection allows organizations to avoid critical security gaps in their IT strategy.

Furthermore, Rubrik provides automated Sensitive Data Monitoring by scanning and classifying backup data against many pre-built and custom analyzers. This allows organizations to quickly pinpoint exactly where their sensitive data is located, what types of sensitive data they have, and how much of it exists—key initiatives in mitigating the impact of cyber-attacks should they occur.

Withstand Cyber Attacks

Cybercriminals have long moved past simply targeting production data—instead, they now employ holistic solutions that target both production and backup data to help increase the likelihood that they will receive payment. Quite often, if an organization can't recover from a cyber attack, they have no option but to comply with the bad actors and pay the ransom. The Rubrik M365 solution provides a fully hosted platform to secure your M365 backups in an air-gapped, immutable storage location, completely separated from your internal Microsoft 365 trust boundaries. Data is secured and stored leveraging Azure's Write-Once, Read Many (WORM) time-based retention policies within a Rubrik-managed Azure subscription, with industry standard security layers integrated into the core platform such as Multi-Factor Authentication, least-privilege Role Based Access Control, two-layer envelope encryption, Retention Locked SLA Domains, and Zone redundancy storage principals. Rubrik Security Cloud also provides the ability to protect the underlying Users and Groups within a

customer's Entra ID environment, again, layering on security services to immutable, air-gapped backups. Rubrik provides true data resilience with secure backups outside of your organization's attack radius, ensuring your backups can always be utilized as your last line of defense.

Recover Rapidly

When data is not available and recovery is too slow or nearly impossible, business continuity is at risk. For every minute an organization is down, the business runs the risk of not only lost revenue and wasted operational costs but also increases its chances of reputational and brand damage. It's imperative that organizations are not only able to recover their M365 data quickly but also have options as it pertains to recovery, as every restoration scenario contains different requirements.

Rubrik provides a multitude of efficient options to recover M365 data. Organizations can leverage Rubrik's powerful metadata engine to instantly search across hundreds of restore points to find individual emails, folders, and files and then granularly perform item-level restoration back to the same or a different account. Administrators can confidently perform entire mailbox recoveries or simply export PST files directly from the point-in-time backups of Exchange. With Rubrik's Mass Recovery functionality, organizations can accelerate the recovery of hundreds or even thousands of M365 Exchange and OneDrive Users and Groups, and SharePoint Sites.

ARCHITECTURE AND COMPONENTS

Rubrik's Microsoft 365 protection is delivered as a fully-hosted SaaS solution, allowing customers to easily onboard and protect their M365 data with Rubrik Security Cloud. At a high level, Rubrik's M365 protection involves three main layers of software: Control Plane, Data Plane, and Source Data.

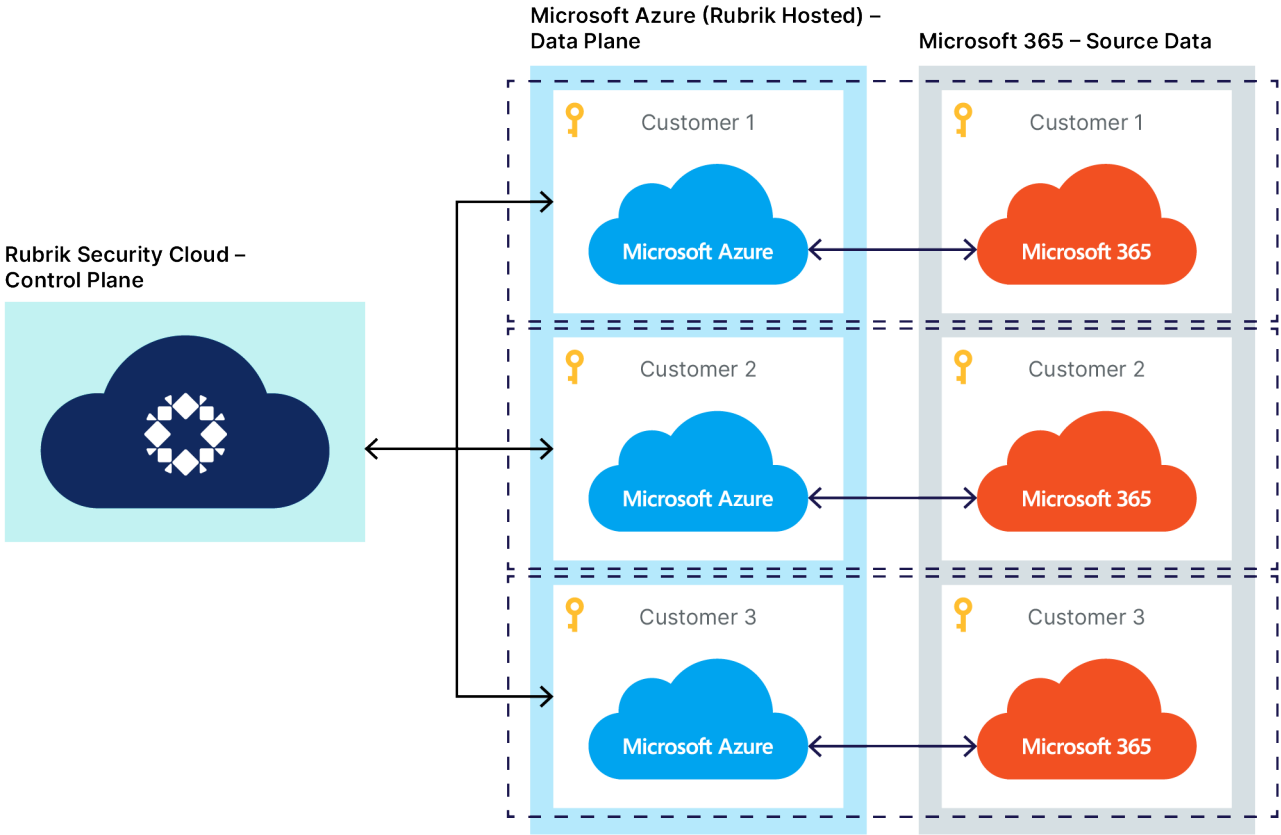


Figure 3 – Breaking down Rubrik's M365 Protection

1. **The Control and Management Plane** – Rubrik Security Cloud acts as the control plane and is the “brains” behind Rubrik’s M365 offering. M365 protection is orchestrated by RSC, a data management engine. It runs as SaaS and is responsible for the entire control and management of the Rubrik M365 software stack. Examples of control and management include access control, audit, backup/recovery tasks, and reporting.
2. **The Data Plane (Exocompute)** – Exocompute is a software layer that performs the heavy lifting in terms of manipulating and moving data. Exocompute is delivered through the means of an Azure AKS Cluster contained within a Rubrik-owned and managed Azure account. The Kubernetes clusters hosting the exocompute are customer-specific, and no data/compute is shared amongst customers. Exocompute is automatically scaled up and down depending on the work to be performed and requires no customer interaction to maintain. Examples of processes performed by the data plane are the processing of API requests for backup and recovery tasks, indexing, and encryption.
3. **The Source Data** – The source data for M365 resides in Microsoft 365, which includes Exchange, OneDrive, SharePoint, and Teams. This data is pulled by the Exocompute layer into the fully hosted customer-specific storage account managed by Rubrik.

Rubrik’s Microsoft 365 protection is a fully hosted, air-gapped solution. This allows customers to take advantage of the power of Rubrik’s SLA policy engine within RSC without the need to run compute instances within their Azure subscriptions. Instead, Rubrik leverages a scalable AKS cluster within a customer-specific Rubrik-managed Azure account, resulting in no compute costs to the customer. Rubrik intelligently consumes both the Microsoft Graph API as well as individual M365 application APIs to provide efficient backup of the customers’ M365 data. Data is written to a customer-specific storage account within the Rubrik-managed Azure account, leveraging both encryption and immutability.

COMPONENTS

In this section we will discuss the deployed resources in more detail and describe their roles within M365 protection before delving deeper into the architectures and workflows.

Entra ID, Service Principals, and Enterprise Applications

Rubrik creates an Entra ID Enterprise Application (and associated Service Principal) for each customer who enables Rubrik’s M365 Backup and Recovery software. These resources are controlled by Rubrik but authorized by the customer through Modern Authentication (OAuth 2.0) to provide Rubrik access to their M365 subscription. Once Rubrik creates the associated Service Principal and Enterprise Application, the global administrator credentials are no longer needed and purged from memory. At no time are they ever stored or written to storage.

Rubrik’s Enterprise Application architecture for M365 provides scalable performance while performing backup and recovery tasks. By default, Microsoft rate limits the number of calls an entity can make against the M365 APIs. This rate limiting is done both at the Enterprise Application level and at the subscription level. Rubrik has designed the solution around this to make it throttling-aware and to ensure sustainable performance as the protected environment grows. To learn more about the API permissions Rubrik requires, see [Appendix A: Required Microsoft 365 API Permissions](#).

Azure Kubernetes Service (Exocompute)

Azure Kubernetes Service (AKS) is a managed container orchestration service based on the open-source Kubernetes system and available on Microsoft's Azure public cloud.

AKS resides in a private Azure Virtual Network (VNet), allowing only inbound access from Rubrik. The connection between AKS and Rubrik is also secure using TLS 1.3 for transport.

Worker Node images are managed by Microsoft and are running stripped-down, hardened Linux operating systems. Security patches are applied regularly by Microsoft.

Lastly, AKS is managed by Rubrik – upgrades and scale are all handled automatically.

Rubrik's software for M365 is run through a set of containerized applications within the AKS cluster. These containers are utilized to provide critical M365 protection services such as performing backups and restores, generating metadata, and applying encryption to data at rest. Container definitions are stored in a private Azure Container Registry, and updates are pushed via secure and automated deployment pipelines.

All containers are scanned for vulnerabilities as part of the deployment pipeline and then checksummed to ensure data integrity before being promoted to production. This ensures that only authorized containers are ever run and have passed appropriate validation checks.

Customer-Specific Storage Account (Azure Blob)

Rubrik's Exocompute data plane writes data to both Azure Table and Blob Storage to store metadata and data, respectively. Data is stored encrypted-at-rest using both key and data encryption keys (KeK, DeK). Keys are stored securely in an Azure Key Vault KMS and accessed via secured credentials.

Storage accounts in Azure are set up to accept TLS 1.3 connections with stringent firewall rules to allow only connections from trusted sources (e.g. Rubrik). Storage accounts use zone-redundant storage (ZRS) where data is replicated across three different availability zones where available.

For customers with a large amount of users, Rubrik will leverage multiple storage containers and load balance the data in order to provide increased read/write performance. This is one of the ways that Rubrik is able to be performant at scale, especially during onboarding and mass recovery scenarios. This storage sharding happens automatically without requiring any customer interaction and provides more aggregate IOPs and throughout.

Azure Key Vault

Azure Key Vault is a cloud service for securely storing and accessing secrets. A secret is anything that you want to tightly control access to, such as API keys, passwords, certificates, or cryptographic keys.

Rubrik leverages Azure Key Vault to manage customer-specific Intermediate Key Encryption Keys (KeKs) to ensure data is secure and encrypted when at-rest while ensuring keys are securely stowed.

MICROSOFT 365 SECURITY

As with any fully hosted solution, the security of customer data is of the utmost importance. Rubrik Security Cloud provides a true defense-in-depth approach as it pertains to securing customers' M365 data through the usage of multiple security layers.

Authentication, Authorization, and Accounting (AAA) in Rubrik

Rubrik Zero Trust Data Management approach of trust nothing, always verify, incorporates the “Triple A” (AAA) framework.

AUTHENTICATION

Authentication restricts access to Rubrik to a specified set of users.

Rubrik requires and enforces the use of strong passwords for all user authentication by default (or as configured). This helps detect brute-force attacks and to block credentials identified in security breaches. Customer data is protected from certain common SaaS security issues, such as session hijacking, script insertion, and cross-site-request-forgery. Beyond this, IP allowlisting also enables Rubrik to restrict login access to a specified list of IP addresses, address ranges, or subnets.

Single-Sign-On (SSO)

Rubrik supports single sign-on (SSO) using the Security Assertion Markup Language (SAML) 2.0 standard. SSO allows login to Rubrik using credentials associated with an identity provider configured by the customer.

SAML 2.0 uses metadata files to exchange information between an identity provider (IdP) and a Service Provider (SP), such as Rubrik. The information in these files establishes a trust relationship between the two entities. The files also specify where authentication requests and responses should be sent, along with formatting details.

Rubrik can be integrated with any SAML 2.0-enabled IdP that supports SP-initiated SSO, such as ADFS, Azure, Okta, and OneLogin.

Two-factor authentication (2FA)

Time-based one-time passwords (TOTP) enable Rubrik Two-step Verification for Rubrik.

Rubrik Two-step Verification is an implementation of two-factor authentication (2FA) for RSC. Users can enable Rubrik Two-step Verification to use TOTP-mediated 2FA, which provides an additional layer of authentication security. In addition to the username and password, TOTP uses an app to provide a single-use numeric code that serves as the second authentication factor. Administrators can enforce Rubrik Two-step Verification for users. When enforced, each user must configure Rubrik Two-step Verification on an individual basis.

Rubrik Two-step Verification supports authenticator apps from Microsoft, Google, and Okta.

AUTHORIZATION

Rubrik offers role-based access control (RBAC) that restricts access based on the roles of individuals within an organization. Access rights are restricted to the relevant associated operations and resources depending on the assigned role. Access rights are also based on the least privilege principle, a key tenant of the Zero Trust framework with regard to M365.

Custom roles in Rubrik can be created to provide access to specific M365 resources at the workload, AD group, or individual user/site level.

In addition, roles can be assigned granular privileges, for example limiting a role to only read-only access or specific recovery options. This is particularly useful in cases when helpdesk personnel have the permissions to perform the recovery of most employees, except for those in C-Suite, HR, or Legal departments. The recovery for these departments is reserved for trusted senior administrators.

For access to the Microsoft 365 APIs, Rubrik requests the minimum set of permissions needed to achieve the task for each Enterprise Application. Rubrik also uses multiple Enterprise Applications, each one scoped to a specific Microsoft 365 Application (i.e. Exchange Mailbox, OneDrive, etc.).

ACCOUNTING

The Rubrik Events feature identifies, isolates, and prioritizes incidents with a unified view of global Rubrik events.

The Events feature makes it possible to find point-in-time events (by event and object type) with easy-to-use filters and real-time search. M365 Events in Rubrik can also be forwarded to SIEMs or log management systems.

Audit log functionality shows log messages for Rubrik domain system events.

Data Encryption

Data in transit between Rubrik, Exocompute (Azure), and M365 can utilize TLS 1.2+ to communicate. Communication internal to Rubrik or Exocompute can also be encrypted with TLS 1.2+ enforced.

Rubrik stores customers' backup data in Azure Blob Storage. To ensure strict isolation, each customer has their own Storage Account to host their storage. Data at rest is encrypted using the AES 256-bit cipher. Keys for data encryption are stored in Azure Key Vault with features like purge protection enabled to prevent accidental key deletion. For more details about data encryption, see the [How it Works: Data Encryption](#) section of this paper.

Data Immutability

Rubrik provides fully hosted immutable storage for M365 backups by storing data in a Write Once, Read Many (WORM) state within Azure Blob. While in a WORM state, data cannot be modified or deleted, protecting it from intentional or accidental overwrites or deletes. Rubrik leverages Azure's Time-Based Retention Policies to adhere to the constructs configured within the Rubrik SLA Domain to ensure that written objects can be created (backed up) and read (restored), but not modified or deleted until the desired retention time has been reached.

SLA Retention Lock

A retention lock can be applied to a Rubrik Security Cloud SLA Domain to prevent premature deletion of backups belonging to Microsoft 365 workloads. With Retention Lock enabled, customers will be unable to perform actions such as deleting the SLA Domain, changing retention within the SLA that results in a decreased retention for backups, re-assigning an SLA with a shorter retention period to an object, or disabling retention lock on the SLA itself. Retention Lock and Data Immutability go to great lengths to mitigate the risk of your backups being affected by ransomware, internal/external bad actors, and/or accidental deletion.

Quorum Authorization

Quorum Authorization (Q-Auth) is a security feature that enforces the requirement of getting additional approvals before a requester can perform data-modifying actions in Rubrik. Q-Auth ensures that no single user has the authority to perform important actions on critical data secured by RSC. With Q-Auth enabled, a user requires additional approvals to perform data-modifying actions. Q-Auth approvers should be different people than normal Rubrik administrators.

With Q-Auth configured, a retention lock can be applied to a SLA Domain policy in Governance mode. Any restricted action to objects protected with a retention locked SLA will require Q-Auth approval in order to be executed. Restricted actions include prematurely deleting backups, weakening the SLA Domain policy such as decreasing the retention for backups, re-assigning an SLA with a shorter retention period to an object, or disabling retention lock on the SLA itself.

Retention Lock and Data Immutability go to great lengths to mitigate the risk of your backups being affected by ransomware, internal/external bad actors, and/or accidental deletion.

Network Security

Rubrik's Exocompute data plane, along with Azure Storage, resides in a private VNet, which has very restricted firewall rules, only allowing inbound access from trusted sources (e.g. Rubrik Security Cloud IPs). All connectivity into the data plane occurs through secure tunnels.

Each tunnel is authenticated when established and can be secured using TLS 1.2+ with mutual authentication by digital certificates. The certificates used for the tunnel are rotated periodically.

HOW IT WORKS

The following section provides a deeper technical understanding of how Rubrik Security Cloud protects Microsoft 365 data, from the onboarding of Microsoft 365 to RSC, initial configuration operations, backup and recovery processes, and the discovery of Sensitive Data.

INITIAL CONFIGURATION OF M365 WITHIN RSC

Before protecting M365 data, the customer's M365 subscription must be initially configured within Rubrik Security Cloud. Performing the initial configuration requires customers to input the following information into Rubrik Security Cloud:

- The region in which to host backups
- Whether or not Multi-Geo should be leveraged
- The Azure Key vault information if leveraging "Bring your own key" functionality
- One-Time authentication with global admin

By default, Rubrik will generate a Key Encryption Key within the Rubrik Azure environment to utilize for data encryption. If the customer desires to bring their own key (BYOK), then the Tenant ID, Key Vault ID, and Key Name must also be provided during the onboarding process. For more information on exactly how Rubrik's encryption works and the options available, see the [How it works: Data Encryption](#) section of this paper.

A simple wizard walks the customer through the remainder of the authorization and configuration process. First, the customer will be prompted to log into their M365 tenant using global administrative credentials. The provided credentials are only utilized once to provide the authentication for an OAuth2 token, are promptly discarded afterward, and are never stored by Rubrik Security Cloud. Service Principals and Enterprise Applications are then created within the customers' tenant and are leveraged, along with the token, for any subsequent authentication and authorization requests coming from Rubrik. For those customers who wish to not leverage OAuth, the Enterprise Applications can be created manually and the subsequent IDs manually entered into Rubrik Security Cloud. Also, during the configuration process, dedicated customer specific resources are deployed within the Rubrik Managed Azure Account including, Azure storage account(s), Azure Key Vault instances, and Azure Kubernetes clusters (Exocompute). *Figure 4* illustrates the configuration process.

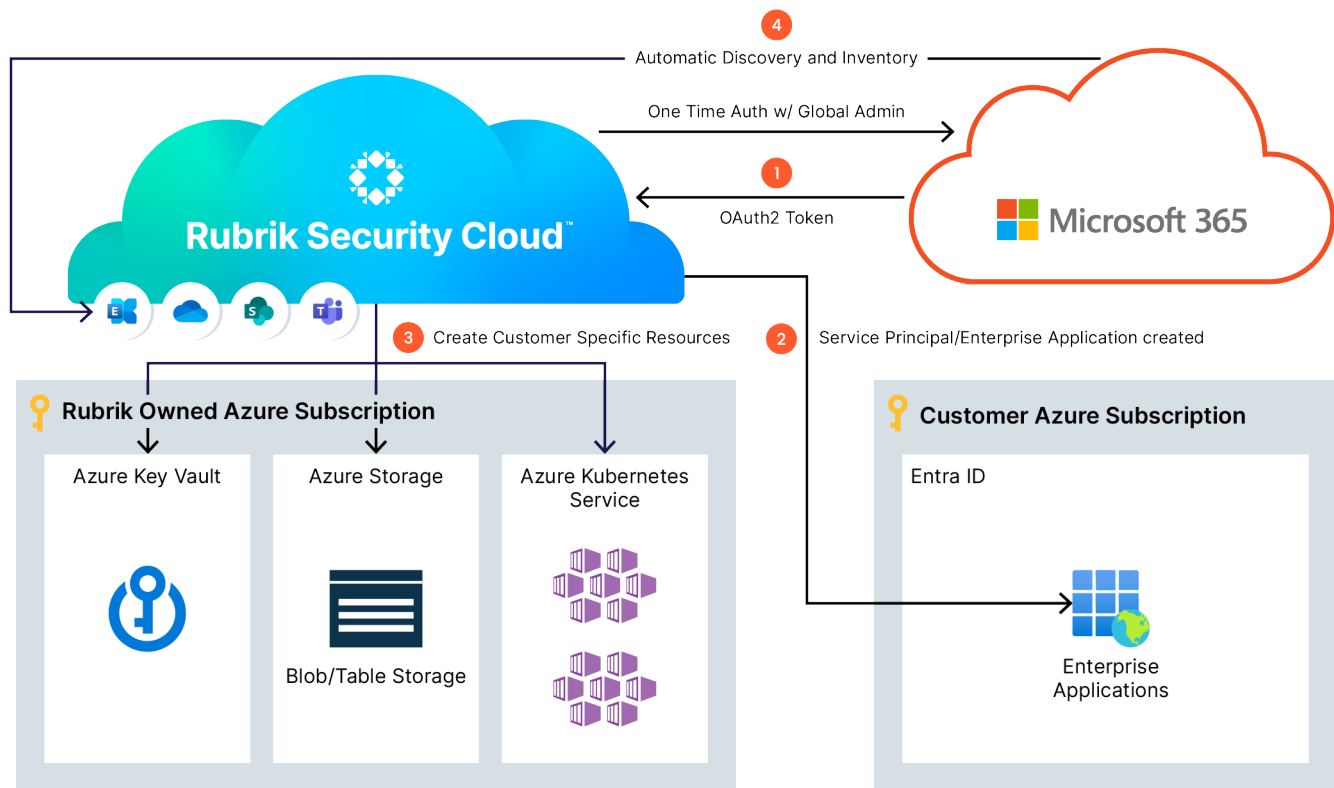


Figure 4 – Initial Configuration Microsoft 365 to Rubrik Security Cloud

After the initial configuration has been completed, Rubrik Security Cloud will leverage Exocompute to begin automatically discovering and performing an inventory of all of the Users, Groups, Mailboxes, SharePoint sites, Teams Channels, etc within the Microsoft 365 environment. Once this process is complete, protection of discovered objects is possible.

CONFIGURATION

This section will walk through the various technical options in terms of configuration and day-to-day management of Rubrik's M365 protection.

Service Principals and Enterprise Applications

Rubrik creates an Entra ID Enterprise Application (and associated Service Principal) for each customer who enables Rubrik's M365 Backup and Recovery software. The app is controlled by Rubrik but authorized by the customer through Modern Authentication (OAuth 2.0) to provide Rubrik access to their M365 subscription. Because Modern Authentication is used to perform this authentication, Rubrik never receives (and therefore, never stores) the customers' credentials for their M365 subscription, as its access is revoked once Service Principals are created.

Many data protection solutions will leverage multiple Enterprise Applications in order to provide scale. While this solution has worked in the past, Microsoft has since strongly stated that this process is no longer a recommended solution. Increasing the number of Enterprise Applications only increases an organizations

attack surface and presents a security risk to the organization. In addition to this, leveraging many Enterprise Applications has the ability to lead to severe tenant wide throttling. This has the potential to not only result in your data protections solution being throttled, but all other enterprise applications within the tenant as well.

Rubrik, along with Microsoft have developed a solution that follows Microsoft's best practices of employing only one Enterprise Application per protected service. For example, customers protecting Exchange, OneDrive, SharePoint and Teams will only require a total of 4 enterprise applications. Based upon the number of licensed users, Rubrik has tuned a single enterprise application to perform a specific number of concurrent jobs as recommended by Microsoft. This not only provides scale and performance, but also reduces an organizations overall security risk as well as management overhead.

SLA Domains

Once onboarding has been completed the next step is to assign SLA Domains to the individual applications or individual objects and/or groups discovered within Microsoft 365, allowing Rubrik to begin performing backup processes against the protected objects. SLA domains are a powerful replacement to the job scheduling approach used by many traditional data protection solutions largely due to their declarative nature, which generally maps very nicely to the RPOs and RTOs required by businesses.

Building an SLA for use with M365 protection is a straightforward process that is outside this document's scope. Please reference the [Rubrik Security Cloud User Guide](#) for details on SLA creation within RSC. Once SLA Domains have been created, they can be assigned to M365 workloads on a variety of hierarchies as illustrated in *Figure 5*.

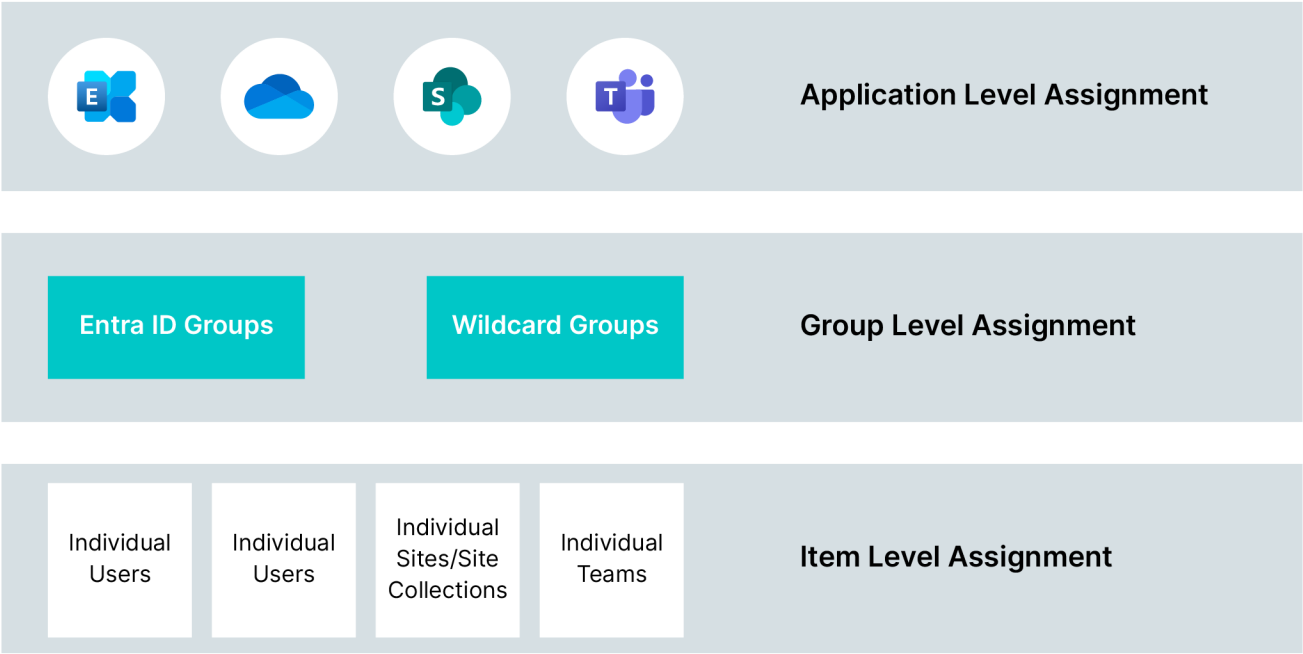


Figure 5 – SLA Domain Assignment Hierarchies

APPLICATION ASSIGNMENT

SLA Domains can be applied directly to the individual Microsoft 365 applications at the top level and, by default, will be inherited by the objects within that application. For instance, if an SLA Domain is assigned to Exchange, all of the existing mailboxes and any new mailboxes will automatically inherit the SLA assigned to the parent. This allows administrators to ensure that data protection is never an afterthought by deploying automated protection to all objects. Inheritance can be broken by assigning a different SLA at one of the other levels explained below or by setting the object to 'Do Not Protect,' which won't be processed during backup processes.

Application assignment is available for Exchange, OneDrive, SharePoint, and Teams.

GROUP ASSIGNMENT

Rubrik can import AD Group membership information that can be leveraged for Exchange and OneDrive SLA assignment. For SharePoint and Teams, you can create Group Wildcards within Rubrik based on 'regex' patterns that match on the SharePoint or Teams name in order to assign a SLA policy at scale based on the matched name.

When a SLA is assigned to an AD Group or Group Wildcard, all users or matched names will automatically inherit the group-assigned SLA, including any new memberships that occur after the initial assignment. A common use-case for group-level assignments is around providing a more aggressive SLA to a group of individuals vs the rest of the business—for instance, backing up executive mailboxes every 4 hours while only providing daily backups for the remainder of the employees.

USER ASSIGNMENT

SLA Domains can also be assigned directly to individual users. A user-level assignment takes precedence over any other type of assignment, meaning any user-level assigned SLA will override the inheritance of application or group-level assignments.

User Assignment applies to only the Exchange and OneDrive applications.

SITE COLLECTION ASSIGNMENT

Site Collection Assignment is specific to SharePoint Site Collections and will override any inheritance of application-level or wildcard group assignments applied directly to the SharePoint application. All components of the site collection, such as document libraries, lists, etc, will be processed, adhering to the constructs set forth within the SLA assigned to the Site Collection.

TEAMS ASSIGNMENT

Assigning an SLA directly to a Team will override any inheritance of SLA Domains assigned to the Teams application or wildcard group. As with SharePoint, any components within the team, such as messages, attachments, files, etc will be processed adhering to the constructs defined within the SLA assigned to the Team.

PROTECTION

Once an SLA Domain has been assigned to an object through any hierarchy, the Rubrik job framework will begin automatically scheduling and taking backups of the M365 objects following the SLA constructs defined. There is no need for customer interaction and no concept of managing jobs. *Figure 6* provides a high-level overview of how Rubrik Security Cloud processes backups of M365.

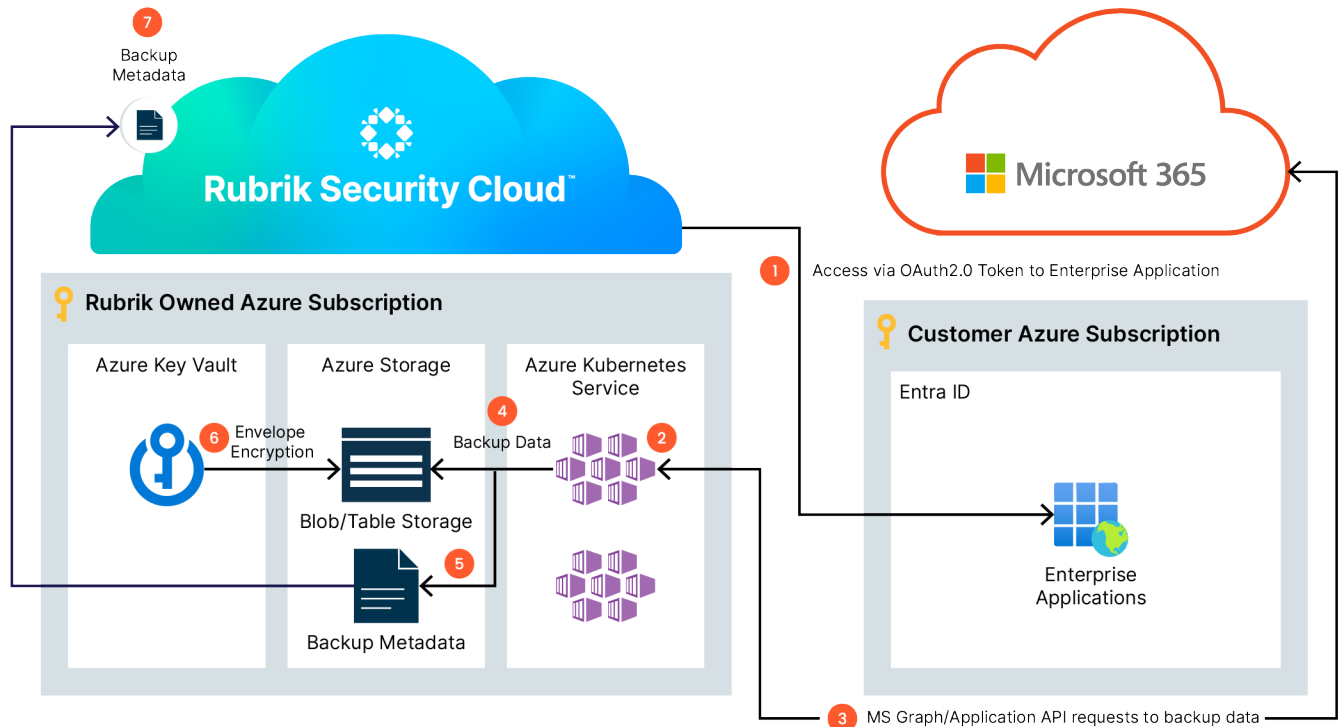


Figure 6 – High-Level Overview of Backup Process

1. RSC leverages existing Service Principals and Enterprise Applications to gain access to the M365 tenant through an OAuth2 token.
2. If required, RSC will instruct AKS to instantiate a specific number of pods containing Rubrik's Exocompute.
3. Excompute will consume the MS Graph API to retrieve the respective data from the M365 tenant. If the MS Graph APIs are throttled and/or failing, RSC will intelligently provide time based retries in order to maintain compliance.
4. Excompute writes backup data to provisioned blob and table storage within a dedicated customer storage account located in the Rubrik hosted environment.
5. The backup data is attached to an excompute instance and read, generating index metadata around the source data within the backup.
6. Excompute leverages the defined encryption parameters retrieves keys from an Azure Key Vault, and encrypts the newly written data.
7. Excompute uploads the generated metadata to Rubrik Security Cloud to be used for granular search capabilities.
8. RSC instructs AKS to scale down the number of instantiated nodes in the event no more jobs are scheduled.

The above represents only a high-level overview of the backup process, and many other functions and processes are executed within. Let's dive into these, starting with data integrity and availability.

Initial Full and Subsequent Backups

Rubrik employs a forever incremental approach to backing up Microsoft 365 data; meaning the initial backup of any object (Mailbox, OneDrive, SharePoint, Teams) is a full backup, while subsequent backups only contain incremental copies of the data that has changed since the last backups.

As discussed, Rubrik leverages multiple techniques to deliver the most performant solution on the market. Sharding backup storage allows for more IOPs to be delivered, while batching API requests helps avoid API throttling mechanisms put in place by Microsoft. The first initial full backup however has the potential to take a significant amount of time. Contributing factors to the first full backup window include not only capacity, but focus more so around the number of users within the environment, along with the specific number of items contained within the users mailbox and OneDrive account, the number of items contained within SharePoint sites, along with the number of Teams channels and files.

For those customers who employ a large number of users (>20,000), the Rubrik team will work with the organizations to tune the initial first full backup process to ensure they meet any compliance requirements they may have.

Subsequent incremental backups are generally much faster and generally do not require support from the Rubrik team.

Data Integrity and Availability

Rubrik employs several measures to ensure customers' backup data is both integral and available if a restoration process needs to be performed. While the customer does not control these characteristics, it's important to understand them.

DATA STORAGE

M365 backup data and its respective metadata are stored in Azure Blob and table stores and encrypted using the AES 256-bit cipher. To maintain isolation between customers, each customer is assigned their own dedicated Azure Storage account, which resides in an Azure account managed by Rubrik. During onboarding, the customer selects the region to store the data, along with whether or not M365 Multi-Geo is leveraged.

Root Key Encryption Keys are stored in Azure Key Vault. Access credentials to storage and Key Vaults are unique per customer and are either provided by Rubrik or the customer.

Backup data is stored within the Azure storage account inside Azure blob stores. Metadata surrounding that backup data is written to table storage, again, within the Rubrik managed Azure account. Metadata is also copied to Rubrik Security Cloud and stored within a multi-tenant relational database in order to provide efficient search and browse functionality when determining data to restore.

Data is also deduplicated across a customer's entire M365 subscription and compressed.

MULTI-GEO SUPPORT

Microsoft 365's Multi-Geo capabilities allow customers to expand their M365 presence to multiple geographic regions or countries within a single M365 tenant, satisfying any data residency requirements they may have.

Rubrik supports Multi-Geo configurations by deploying the resources to process and store the backups within the same regions as the production data. During initial onboarding, a Preferred Data Location (PDL) is specified. Rubrik reads the PDL tag for each user, SharePoint site, or Teams groups, and ensures that the backups for those applications will reside in the same PDL region. During initial onboarding, a Central Data Location (CDL) is also specified. If an object doesn't have a PDL tag, then data is stored in the CDL region.

DATA DURABILITY

Backup data is written into the customer-specific storage account and replicated via Azure's Zone-Redundant Storage (ZRS), which copies data synchronously three times to physically separated data centers in the same geographic region.

OPERATIONAL SECURITY

Being a fully-hosted solution, Rubrik employs many operational security measures to deliver the utmost privacy as it pertains to customers data. That said, either through support requests or internal processes where Rubrik engineers require access to customer data for troubleshooting purposes. In the event this access needs to be provisioned, a number of checks and balances are required.

First, no internal Rubrik employees have standing access to any production environments. Those that have the ability to access production data are limited, and in order to gain access a request must be made with any subsequent justifications required.

If in the event access is granted, a time limit is placed on the session. In addition to this, Rubrik maintains audit logs of all actions performed against production data within the Rubrik hosted environment.

In addition to all these requirements, Rubrik also conducts regular internal and 3rd party reviews of security policies and access to any production resources.

SOFT DELETE

Soft Delete is enabled on storage hosting customer M365 backups. When a backup is deleted, either through scheduled retention policies or manually deleted, the recovery point will be removed from the user view with RSC. The data however, remains on the Rubrik hosted storage for a set number of days following the deletion event. If a customer runs into a situation where data was deleted accidentally or maliciously, they can engage with Rubrik support to attempt to retrieve it. Once this time period has expired, the data is permanently deleted and is not retrievable by any means.

COMPLIANCE – 3RD PARTY VALIDATION AND CERTIFICATIONS

Rubrik's security policies are reviewed regularly by internal and 3rd party companies to ensure they are to best practices. Continual security scans and penetration testing are also part of Rubrik's security practice.

Rubrik also has achieved several certifications and attestations of compliance against global standards. For a complete list, visit the [Rubrik Compliance Program](#).

Working around API Throttling and Failures

Due to the nature of Microsoft 365 being a production application, Microsoft implements throttling limits to the number of concurrent calls to a specific service to prevent the overuse of resources and mitigate latency and downtime.

When exceeding the throttling threshold, Microsoft will temporarily limit future requests from that specific client. While most users may not notice throttling during day-to-day usage of M365, you can imagine its impact on Rubrik when trying to perform backups of hundreds of thousands of objects.

To alleviate both throttling and failures and increase performance, Rubrik leverages several techniques to ensure successful backups and maintain compliance. Let's explore these techniques in detail.

HANDLING M365 THROTTLING

Rubrik, along with Microsoft have developed a solution that follows Microsoft's best practices to alleviate situations that might cause throttling of M365 APIs. Global throttling has the potential to not only affect your data protection solution, but can lead to the throttling of other applications and services consuming M365 APIs as well. Rubrik adheres to Microsoft best practice by employing the use of only one Enterprise Application per protected service. For example, customers protecting Exchange, OneDrive, SharePoint and Teams will only require a single Enterprise Application per protected service, for a total of four. Based upon the number of licensed users, Rubrik provides the ability to tune a single enterprise application to perform a specific number of concurrent jobs. This not only provides scale and performance, but also reduces an organization's overall security risk as well as management overhead. Furthermore, Rubrik will batch multiple queries into single API requests to optimize each round trip when performing backup and recovery of M365.

The specific number of API calls allocated per application within a Microsoft 365 tenant depends on the amount of licenses purchased. The following table outlines the resource unit limits for SharePoint (OneDrive + Teams) at the time of writing this paper.

License Count	0-1K	1K-5K	5K-15K	15K-50K	50K+
App per minute	1,200	2,400	3,600	4,800	6,000
App per day	1,200,000	2,400,000	3600000	4,800,000	6,000,000

It should be noted that these API request limits are shared across all applications calling the Microsoft Graph APIs, not just Rubrik Security Cloud.

GRAPH TO APPLICATION FAILOVER/FAILBACK

To interface with the Microsoft environment, Rubrik communicates with M365 via AKS, through multiple Microsoft APIs. The recommended M365 API is known as Microsoft Graph, which provides a single endpoint that allows Rubrik to connect to each of the underlying M365 applications. Also available are the traditional APIs that are specific to each M365 product—such as [EWS](#) for Exchange. While the MS Graph API is recommended and more efficient, it is prone to failures. Thus, to provide both performance and reliability, Rubrik takes a “hybrid” approach to consuming APIs with the ultimate goal of creating the most reliable and performant experience possible. For example, if Microsoft Graph fails during an operation for M365 Exchange, Rubrik can switch to the traditional Exchange Web Services API to complete the task with minimal interruptions in the data protection process while automatically failing back afterward.

Data Encryption

Encryption is pivotal in safeguarding our sensitive information if a breach occurs. Rubrik employs in-flight and at-rest encryption, ensuring that data is encoded so that only authorized parties can access and decipher it.

IN-FLIGHT ENCRYPTION

All data in transit between Rubrik Security Cloud, Exocompute (Azure), and M365 uses TLS 1.3 to communicate. Absolutely no unencrypted communication is permitted. Furthermore, even internal communication within the Rubrik components, including RSC and Exocompute, is encrypted with TLS 1.3 enforced.

AT REST ENCRYPTION

Rubrik leverages AES-256 ciphers to encrypt data at rest within the customer-specific storage account.

The at-rest data is encrypted utilizing an envelope encryption scheme, whereas each file is encrypted with a unique key known as the “Data Encryption Key” or DEK.

The DEK is then encrypted with an additional encryption key, known as an Intermediate “Key Encryption Key” or KEK. This Intermediate KEK is also stored alongside the backup data. A new Intermediate KEK is generated every 30 days for new data and completely rotated once a year automatically.

Finally, a Root KEK is used to encrypt the Intermediate KEK. This Root KEK is stored in Azure Key Vault, either in Rubrik’s environment or if a customer opts to bring their own encryption key, in an Azure Key Vault that the customer owns. *Figure 7* illustrates how Rubrik achieves envelope encryption at a high level.




Figure 7 – Envelope Encryption

Rubrik provides two options for managing Key Encryption Keys: Customers can depend on Rubrik generated keys or they can provide the key using Bring Your Own Key (BYOK). These encryption options must be decided when initially onboarding the M365 subscription and cannot be changed afterward. While both options are similar, there are differences between available functionality between them—let’s have a closer look.

Rubrik Generated Keys

By utilizing Rubrik generated keys, Rubrik will manage and host the entire encryption process from end to end. This includes hosting the Azure Key Vault within the Rubrik environment and creating all DEK and KEK keys. The Azure Key Vault within the Rubrik environment is configured with the Soft Delete and Purge Protection features enabled and is a customer-specific resource. When using Rubrik generated keys, on-demand and automatic key rotation, along with annual rekeying, is also available. Below, *Figure 8* outlines the process of utilizing Rubrik Generated Keys.




Figure 8 – Rubrik Generated Keys

1. An Intermediate Key Encryption Key (KEK) is generated and stored within the customer-specific storage account managed by Rubrik.
2. A Data Encryption Keys (DEK1, DEK2, and DEK3) mapping to the respective data to be encrypted.
3. The Data Encryption Keys are used to encrypt the respective blob stores.
4. The Intermediate Key Encryption Key is utilized to encrypt the newly created Data Encryption Keys and stored alongside the encrypted data within Azure blob storage.
5. The Intermediate KEK is encrypted using Key Vault's Root KEK and also stored within blob storage.

Bring your own Key (BYOK)

For customers who wish to maintain more control over the encryption process, Rubrik offers a BYOK option. When leveraging BYOK, customers provide their own Root KEK and can revoke access at any time, leaving the data within the Rubrik environment unreadable. The customer provided Root KEK is stored within an Azure Key Vault running in the customers' environment and shared with the Rubrik environment. Rubrik accesses the Root KEK by leveraging an application registration configured within the customer's account. Key Rotation and ReKeying are still supported within the BYOK model; however, this is a manual process that involves the Rubrik support team. *Figure 9* outlines the process of utilizing BYOK for encryption.




Figure 9 – Bring your own key

1. Rubrik generates a Data Encryption Key for encryption purposes.
2. The Data Encryption Keys are utilized to encrypt the data on the Azure blob storage.
3. Rubrik generates an Intermediate Key Encryption Key (KEK).
4. The Intermediate KEK is utilized to encrypt the Data Encryption Keys, which are then stored alongside the encrypted data.
5. The Intermediate KEK is encrypted using the customer provided Root Key Encryption Key.

Note: In the event of key loss, all encrypted data becomes irretrievable and unreadable. This irreversible consequence underscores the need for customers to exercise extreme caution in managing and preserving their encryption keys.


Note: Rubrik never stores the customers Root encryption keys on persistent storage. Instead, keys are constantly queried for when needed, at times, multiple times per minute. At any given time the customers can revoke access to the keys, leaving the data unreadable by Rubrik.

KEY ROTATION

Rubrik performs key rotation every 30 days for those customers leveraging Rubrik Generated Keys. The process of key rotation creates a new Intermediate KEK that will be utilized to encrypt future DEKs that are generated for any newly written data after the key rotation process has completed. Existing data remains encrypted utilizing the original DEK and KEKs. Key Rotation limits the amount of data encrypted with any given key, essentially reducing the risk of exposure should any single key be exfiltrated. It should also be noted that before performing key rotation, Rubrik performs backups of existing keys in the event they are needed. *Figure 10* illustrates how key rotation is processed within Rubrik.

🔑 Rubrik Owned Azure Subscription

Azure Key Vault



Azure Storage

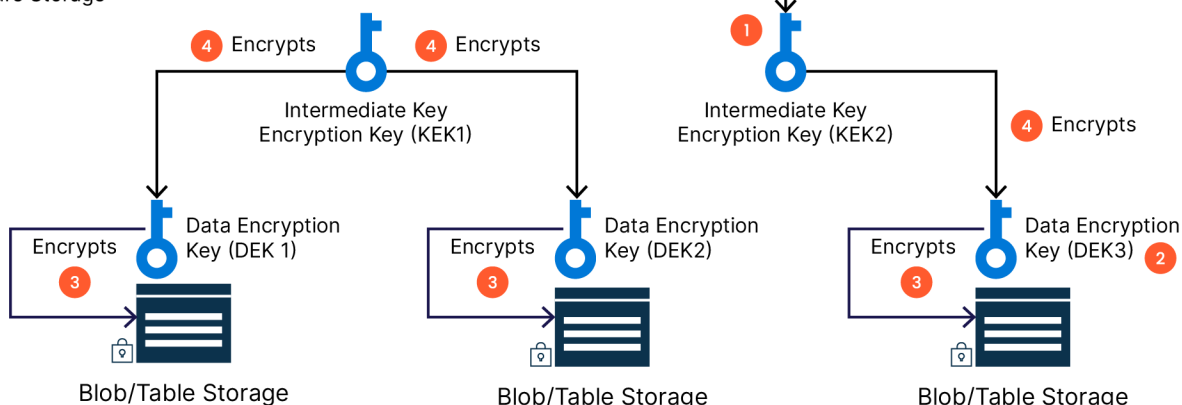


Figure 10 – Key Rotation

1. A new Intermediate Key Encryption Key is generated and stored within the Azure Storage.
2. A new Data Encryption Key is generated by Rubrik.
3. Blob data is encrypted utilizing the newly generated Data Encryption Key.
4. The newly generated Data Encryption Key is encrypted using the newly generated Intermediate Key Encryption Key.
5. The newly generated Intermediate Key Encryption Key is encrypted using the Root KEK.

Note: Existing data, Key Encryption Keys, and Data Encryption Keys remain untouched. Key Rotation only affects new data written after the rotation process

REKEYING

Like Key Rotation, ReKeying creates a new set of Intermediate KEKs. However, instead of just utilizing the newly created Intermediate KEKs for newly written data, ReKeying will go back and re-encrypt existing DEKs with the newly generated KEKs. The DEK themselves do not change; therefore the existing data does not have to be re-encrypted—the DEK is only re-encrypted using the new KEK generated. By default, ReKeying occurs once every year and is supported only for those customers leveraging Rubrik Generated Keys. ReKeying will generate as many new Intermediate KEKs as existing Intermediate KEKs that have been created. *Figure 11* outlines the process of ReKeying within Rubrik Security Cloud.

Rubrik Owned Azure Subscription

Azure Key Vault

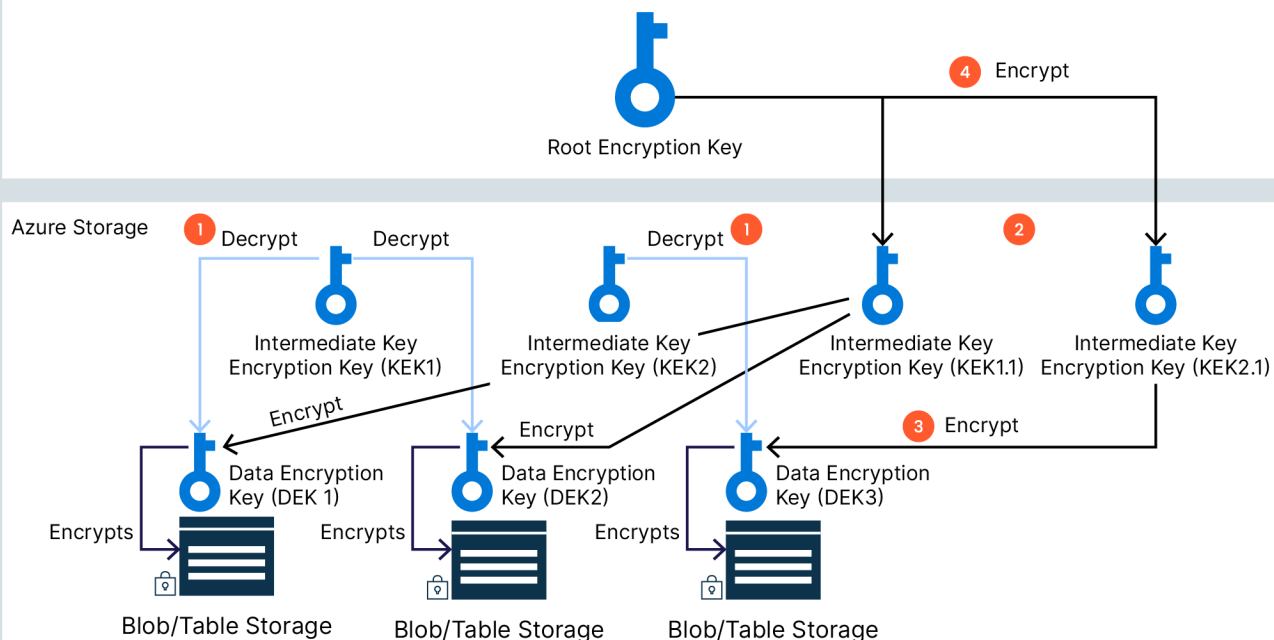


Figure 11 – ReKeying

1. Existing Data Encryption Keys mapped to rotated Intermediate KEKs are all decrypted using the available Intermediate Key Encryption Keys.
2. New Intermediate Key Encryption Keys are generated and stored within the Azure Storage. New Intermediate KEKs are mapped on a one-to-one basis to existing KEKs.
3. The DEKs are then re-encrypted using the newly generated Intermediate Key Encryption Keys.
4. The new Intermediate KEKs are encrypted using the Root KEK.

Note: The data housed within the blob storage itself is not decrypted and encrypted. Only the encryption of the DEK is performed, and since the unencrypted DEKs remain the same, there is no need to re-encrypt data.

KEY BACKUPS

To avoid accidental corruption of keys during rotation and rekeying operations, a backup of rotated keys is taken before the operation begins. Backups are stored within a dedicated metadata backup storage account that is created during the initial encryption setup. In the event a key needs to be restored, Rubrik Support must be contacted

RECOVERY

Without recovery, backup is useless. Rubrik provides a multitude of recovery options as it pertains to M365 data, providing solutions to nearly every restoration scenario. Rubrik's M365 recovery employs many of the same features as its backup process, leveraging AKS to scale required infrastructure resources, batching multiple API requests to avoid throttling, and performing API failover between MS Graph and individual application APIs to increase recovery success. *Figure 12* outlines the high-level process for recovering all M365 applications with Rubrik Security Cloud, while subsequent sections explore recovery options within each individual M365 application.

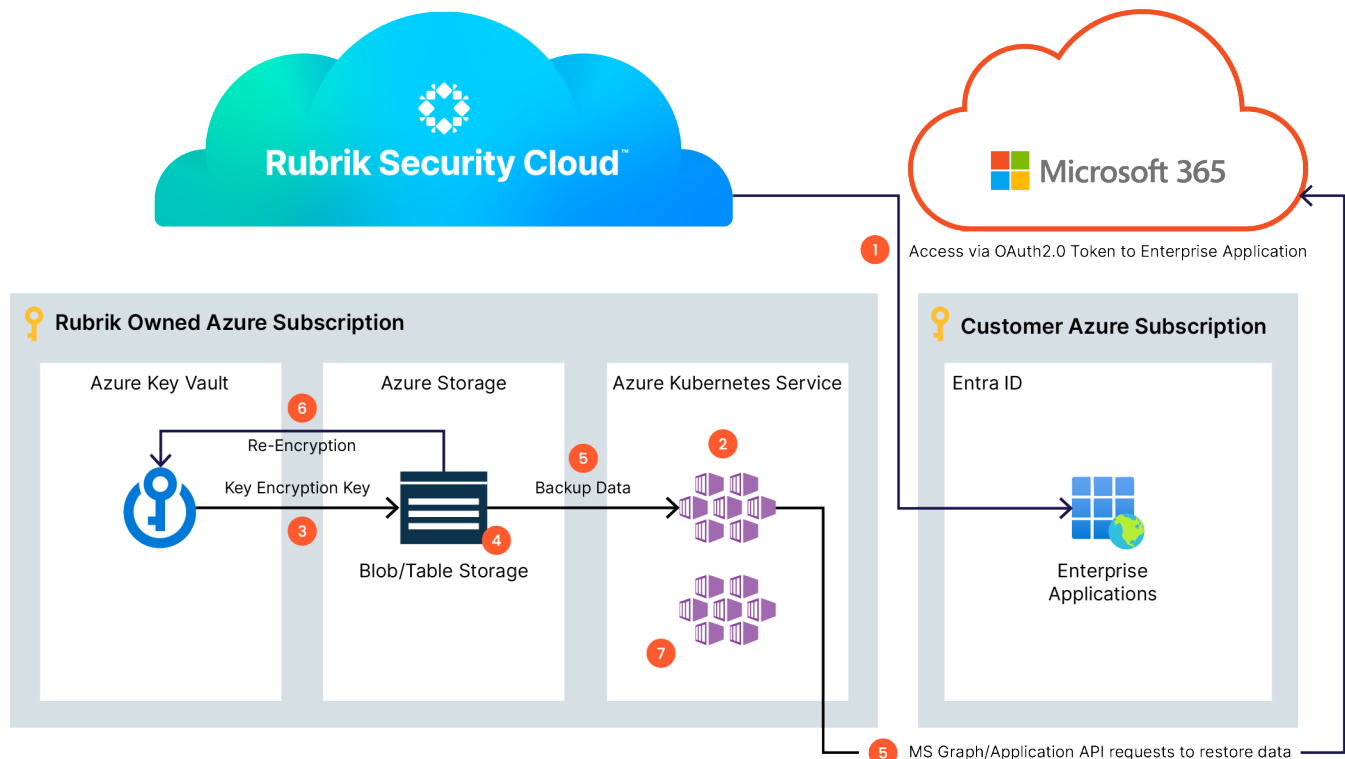


Figure 12 – High-Level Recovery Process

1. RSC leverages existing Service Principals and Enterprise Applications to gain access to the M365 tenant through an OAuth2 token.
2. If required, RSC will instruct AKS to instantiate a specific number of pods containing Rubrik's Exocompute.
3. Exocompute retrieves Key Encryption Key from Azure Key Vault and decrypts the Data Encryption Key that was used to encrypt the data to be restored.
4. The unencrypted Data Encryption Key is utilized to decrypt the backup data.
5. Excompute will combine multiple API requests into one and consume the MS Graph API to restore the data to the M365 tenant. If the MS Graph APIs are throttled and/or failing, RSC will intelligently provide failover/failback to the application-specific APIs to maintain compliance.
6. Data is re-encrypted with its respective Data Encryption Key, Data Encryption Key is re-encrypted with Key Encryption Key.
7. RSC instructs AKS to scale down the number of instantiated nodes in the event no more jobs are scheduled.

Exchange Recovery Options

Leveraging Rubrik Security Cloud's powerful metadata engine, organizations can easily search across multiple point-in-time backups of Exchange, easily locating emails based on content within the subject, date, sender, or recipient fields. Browsing individual point-in-time backups is also available, allowing organizations to drill into mailboxes and folders to discover data to be restored. Aside from email, Rubrik also supports the recovery of Exchange calendars and contacts.

In terms of recovery options for Exchange, Rubrik provides the following:

- **In-place recovery of the entire mailbox** – This allows organizations to recover a user's entire mailbox within Exchange, including their calendar and contacts, overwriting the existing production mailbox.
- **Recovery of individual items** – This provides the ability to search and recover individual mailbox items, such as emails and folders, back to the user's original mailbox.
- **Export to PST** – A common recovery task for any M365 administrator. Rubrik can extract a PST file containing the entire mailbox directly from the backup data, allowing administrators to perform manual restorations.
- **Restoration to a different user** – Whether it be for legal reasons or since an employee is no longer with the company, often the requirement to restore email items to a user other than the source user is required. Rubrik Security Cloud provides an easy workflow, allowing the restoration of messages and folders to different accounts.
- **Restoration to a different user in an alternate subscription** – Rubrik also provides the option to restore mailbox items from a user within one subscription to an alternate user located within an alternate subscription.

OneDrive Recovery Options

As with Exchange, Rubrik's OneDrive protection leverages RSC's metadata to enable users to search across multiple point-in-time OneDrive backups based on the files, folders, and dates within the backup. Whether searching across all point-in-time backups or drilling through an individual point in time, Rubrik provides an efficient way to discover data within the platform.

In terms of recovery of OneDrive, Rubrik enables the following:

- **In-Place Recovery of OneDrive Account data** – Easily recover entire OneDrive account data back to production accounts, essentially rewinding a user's OneDrive instance.
- **Recovery of individual items** – Simply select an individual file or folder to restore and Rubrik restores only the individual files and folders to the user's account.
- **Recovery to a different user** – Like Exchange, Rubrik offers the ability to restore individual files and folders from one OneDrive account to another.
- **Recovery to a different user within an alternate subscription** – Rubrik offers the ability to restore OneDrive data across different subscriptions.

SharePoint Recovery Options

SharePoint is quickly becoming a mission-critical application and organizations must be able to efficiently recover data should the need arise. Rubrik supports the following recovery options as they pertain to SharePoint within M365:

- **Full Site/Site Collection Recovery** – Recover entire SharePoint sites and site collections back to the original production environment.
- **Item-Level Recovery** – Rubrik provides a means to granularly recover individual Lists, Files, and Document Libraries within a SharePoint Site.
- **Recovery to an alternate subscription** – Rubrik allows M365 administrators to recover SharePoint data from one subscription to an alternate subscription.

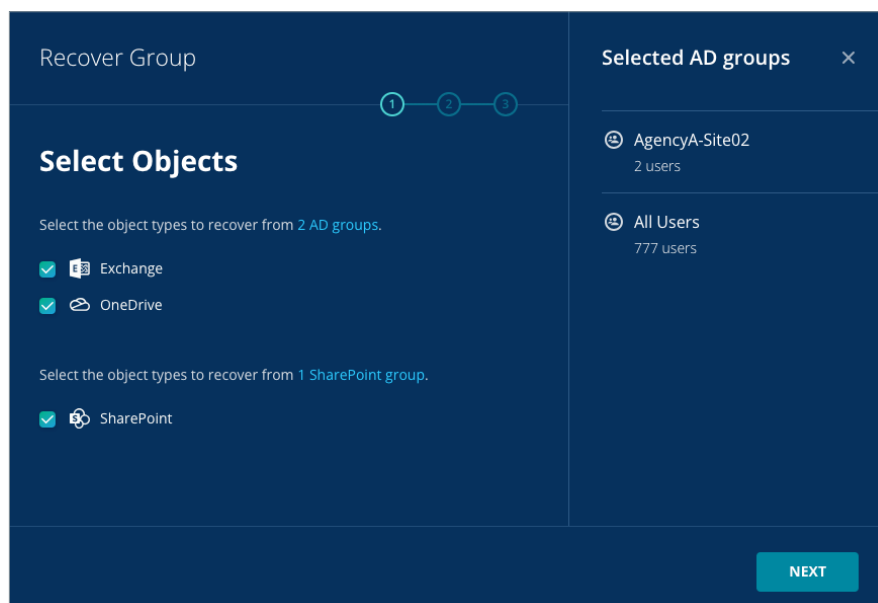
Teams Recovery Options

Microsoft Teams is the backbone of an organization's day-to-day collaboration environment. Efficiently recovering data back to Teams is a crucial business continuity requirement for any organization leveraging the application. Rubrik provides the following options when restoring Microsoft Teams:

- **Recovery Entire Teams** – Rubrik can easily recover entire Microsoft Teams teams in the event of a breach or accidental deletion.
- **Support for private teams** – Rubrik fully supports the recovery of both public and private teams created within the application.
- **Granular Recovery** – Granular recovery of individual posts (including attachments, code snippets, links, etc) as well as individual files uploaded into the Teams environment.

Mass Recovery

As organizations drive to deliver business continuity, data resilience, and compliance with regulatory requirements, M365 mass recovery becomes a paramount priority within their cyber recovery efforts. Whether it be the result of a cyber event or rogue script, Mass Recovery is a must-have, allowing customers to recover hundreds or even thousands of accounts efficiently without the need to select individual objects.



Rubrik's M365 Mass Recovery works by allowing organizations to select groups of M365 objects to recover, rather than performing individual recovery of each object. To select multiple M365 Exchange and OneDrive accounts, users can be added to Entra ID groups, which are then selected for recovery within Rubrik Security Cloud. To select multiple M365 SharePoint and Teams objects, Rubrik uses a concept called wildcard groups. Wildcard groups are populated using simple regex patterns to match against SharePoint site and Team names. If a match is found, the object is deemed a member of the wildcard group.

Mass Recovery of a group of objects is performed much in the same way as recovering individual objects, with the exception of multiple objects being recovered in parallel.

As shown below, *Figure 13* provides a high-level overview of how to perform Mass Recovery of M365 within Rubrik Security Cloud.

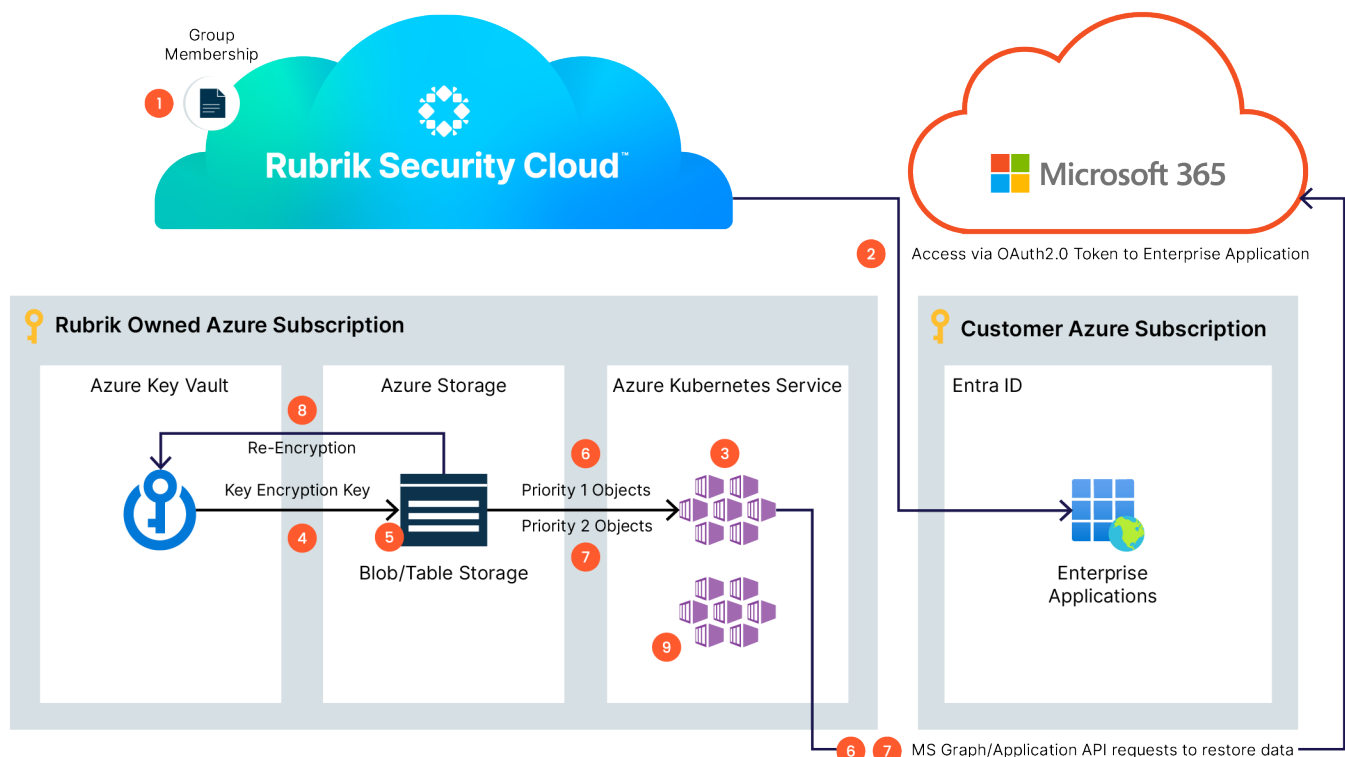


Figure 13 – Rubrik Mass Recovery

1. Information around the Users, SharePoint Sites, and Teams Data included within the group to recover is gathered to prepare for the recovery.
2. RSC leverages existing Service Principals and Enterprise Applications to gain access to the M365 tenant through an OAuth2 token.
3. If required, RSC will instruct AKS to instantiate a specific number of pods containing Rubrik's Exocompute.
4. Exocompute retrieves Key Encryption Key from Azure Key Vault and decrypts the Data Encryption Key that was used to encrypt the data to be restored.
5. The unencrypted Data Encryption Key is utilized to decrypt the backup data.

6. Excompute will combine multiple API requests into one and consume the MS Graph API to restore the data to the M365 tenant. If the MS Graph APIs are throttled and/or failing, RSC will intelligently provide failover/failback to the application-specific APIs to maintain compliance. Objects within the group are recovered in parallel.
7. Data is re-encrypted with its respective Data Encryption Key, Data Encryption Key is re-encrypted with Key Encryption Key.
8. RSC instructs AKS to scale down the number of instantiated nodes in the event no more jobs are scheduled.

Prioritized Data Recovery

Depending on the size of your organization and its M365 data landscape, performing a Mass Recovery of thousands of users has the potential to take some time to complete. Rubrik's Prioritized Data Recovery augments the Mass Recovery process by allowing administrators to instruct Rubrik to prioritize the restore of the most recent data first. Often, the most recent data is the most critical in terms of returning to operations, while the historical data recovery can be completed at a later time.

Select Recovery Type for Exchange

Select a recovery type below.

☒ **Prioritized Data Recovery**

Recover the following items immediately:

- All contacts
- Upcoming calendar events and events from the past 2 weeks

☒ Emails from the last days

☐ Recover all mailboxes except archive mailbox.

☒ Automatically recover the remaining data after prioritized data recovery is complete

☐ **All Data Recovery**

Recover all emails, calendars, and contacts.

Rubrik's Prioritized Data Recovery is currently available for Exchange only (at this time), is applicable to only mass recovery of groups and follows the below restore path in terms of operations:

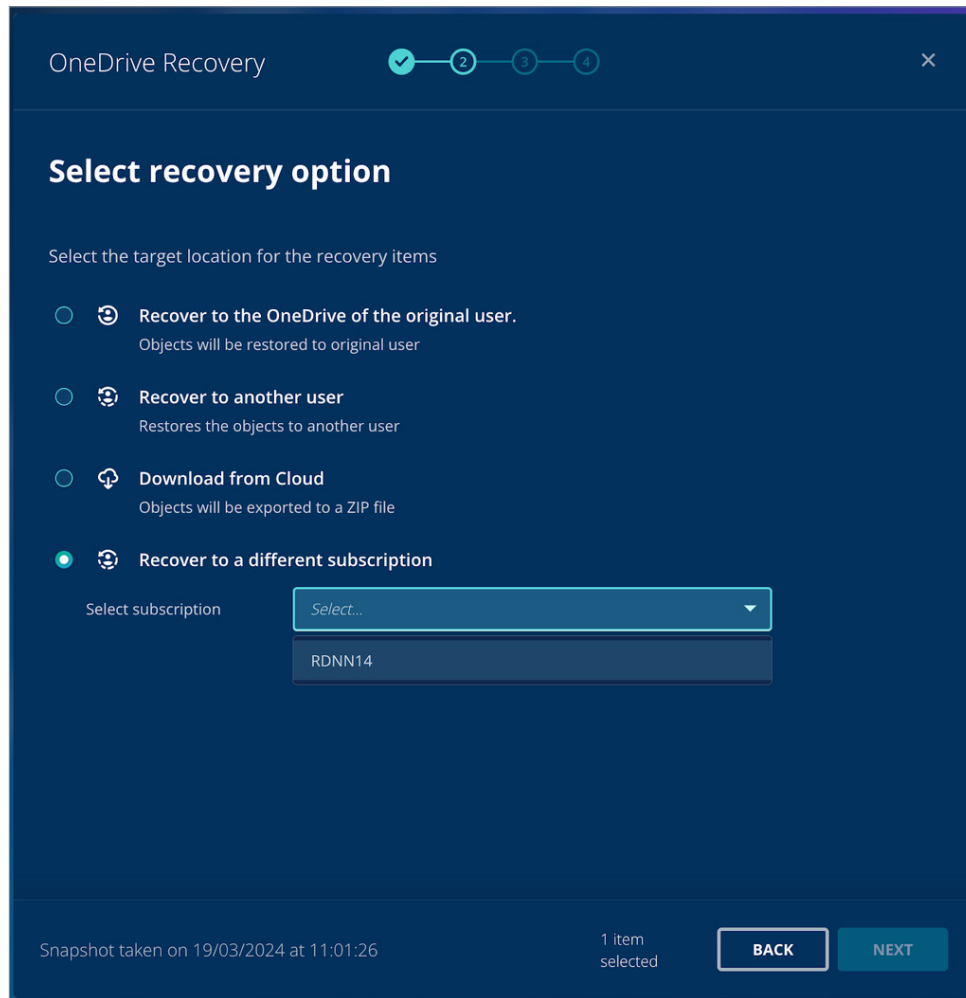
1. All contacts are recovered for all users
2. Any upcoming calendar events and past calendar events for the last 2 weeks are recovered
3. User specifies the date threshold as to which recovery emails. IE Emails from the last two weeks are recovered.
4. Administrators can choose to automatically begin restoration of historical data after prioritized data has been restored, or can pause the process and trigger the remaining recovery manually.

By leveraging Rubrik's Prioritized Data Recovery organizations are able to access their most recent M365 Exchange data without having to wait for entire restore processes to complete. In fact, Rubrik has estimated with 30 days of recovery, using prioritized data recovery the restore time takes less than 1% of the overall total mailbox restoration time, drastically reducing an organization's recovery times.

Cross Subscription Restoration

In today's dynamic business environments, organizations often find themselves managing multiple Microsoft 365 subscriptions across different departments, subsidiaries, or even geographical locations. This fragmentation can present significant challenges, especially when the need arises to restore data lost or compromised due to accidental deletions, cyber incidents, employee movement, or compliance requirements. The ability to efficiently and securely restore M365 data across various subscriptions is crucial for minimizing downtime and maintaining business continuity.

Enter Rubrik's sophisticated solution, embodying a seamless approach to Cross Subscription Restoration. This capability is especially vital for businesses undergoing restructuring, mergers, or acquisitions, where the consolidation or realignment of resources is a strategic necessity. Additionally, enterprises looking to enhance their disaster recovery strategies or ensure strict adherence to regional compliance and data sovereignty laws will find this functionality indispensable. In addition to this, often after a cyber incident, organizations need the ability to restore to a new "clean" M365 subscription to avoid any further issues of reinfection. Rubrik's technology not only bridges the gap between disparate M365 subscriptions but also ensures that key data remains accessible, secure, and restorable on demand, underpinning agile business operations and robust data governance frameworks.



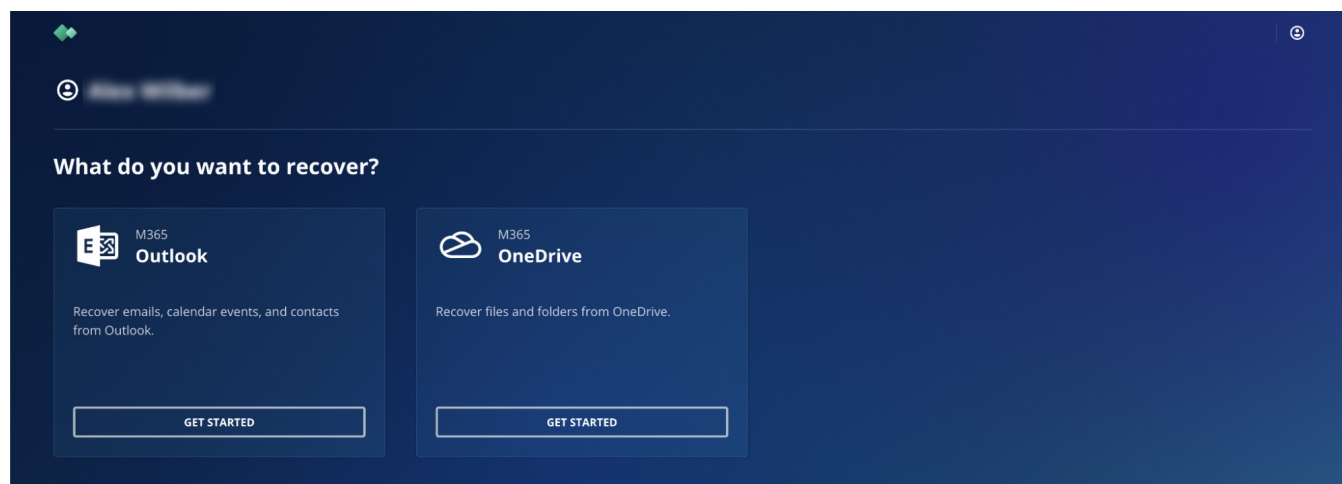
Cross subscription restoration contains the following caveats:

- Support for Exchange, OneDrive and SharePoint (Teams currently on the roadmap)
- Both the source and target M365 subscriptions must be onboarded to Rubrik Security Cloud
- Users must already exist within the target subscription to facilitate recovery of Exchange and OneDrive data
- When recovering an entire SharePoint site collection to another subscription you must specify a new owner for the restored site collection.
- While SharePoint permissions are recovered when restoring to the original subscription, they are not recovered when specifying an alternate subscription

SELF-SERVICE RECOVERY

Rubrik's Self-Service Recovery for Microsoft 365 empowers our customers to quickly restore their own critical data, significantly reducing the burden on IT teams and accelerating business continuity. The process is remarkably simple:

- IT administrators initially configure a dedicated Self-Service Recovery role within Rubrik Security Cloud and assign it to relevant user groups.
- Once set up, end-users can seamlessly log into the Rubrik Security Cloud using their existing organizational credentials or Single Sign-On (SSO).
- From there, they can intuitively navigate to their Exchange or OneDrive data, efficiently search for and select the specific items they need to recover, and then initiate an in-place restore.
- A key feature is that recovered items always appear in a new folder, ensuring the original data remains untouched and preserving data integrity.



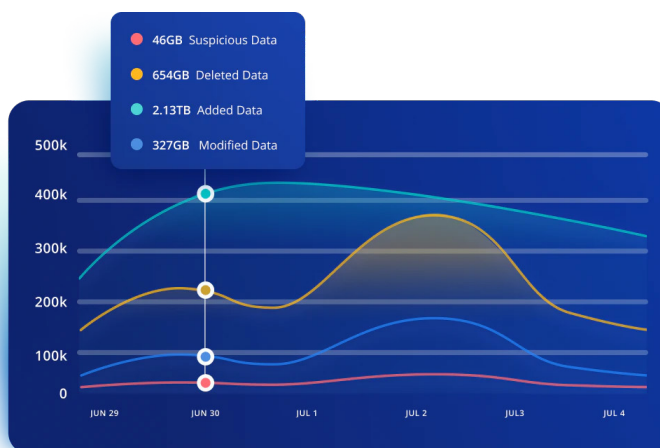
This streamlined approach minimizes downtime and allows IT to focus on more strategic initiatives, while users regain access to their vital information without delay.

DATA THREAT ANALYTICS

Rubrik proactively defends your data with anomaly detection, rapid threat hunting, and real-time threat monitoring. Rubrik Data Threat Analytics helps you uncover hidden risks, investigate faster, and stay ahead of cyberattacks. Let's break down each of the security applications.

Anomaly Detection

Anomaly Detection can proactively identify unusual activity, particularly encryption, within your M365 backup data by continuously analyzing your data, helping you catch ransomware and insider threats before they spread, so customers can act fast and minimize damage. It provides clear visibility into potential threats, ensuring that only genuine incidents are flagged.



Here's how it works:

1. **Difference Identification:** Metadata from M365 snapshots is generated and compared to previous versions to identify differences across the snapshot. This is referred to as **DIFF File Meta Data** or **DIFF FMD**.
2. **ML Anomaly Detection:** This DIFF FMD data is fed into an **ML model** to pinpoint potential anomalies.
3. **Entropy Scan:** Concurrently, an **Entropy Scan** is performed on the DIFF FMD data, with encryption statistics saved for further analysis.
4. **UI Presentation:** All validated anomaly and encryption detection results are then presented in the **Anomaly Detection UI**, offering clear insights into potential threats.

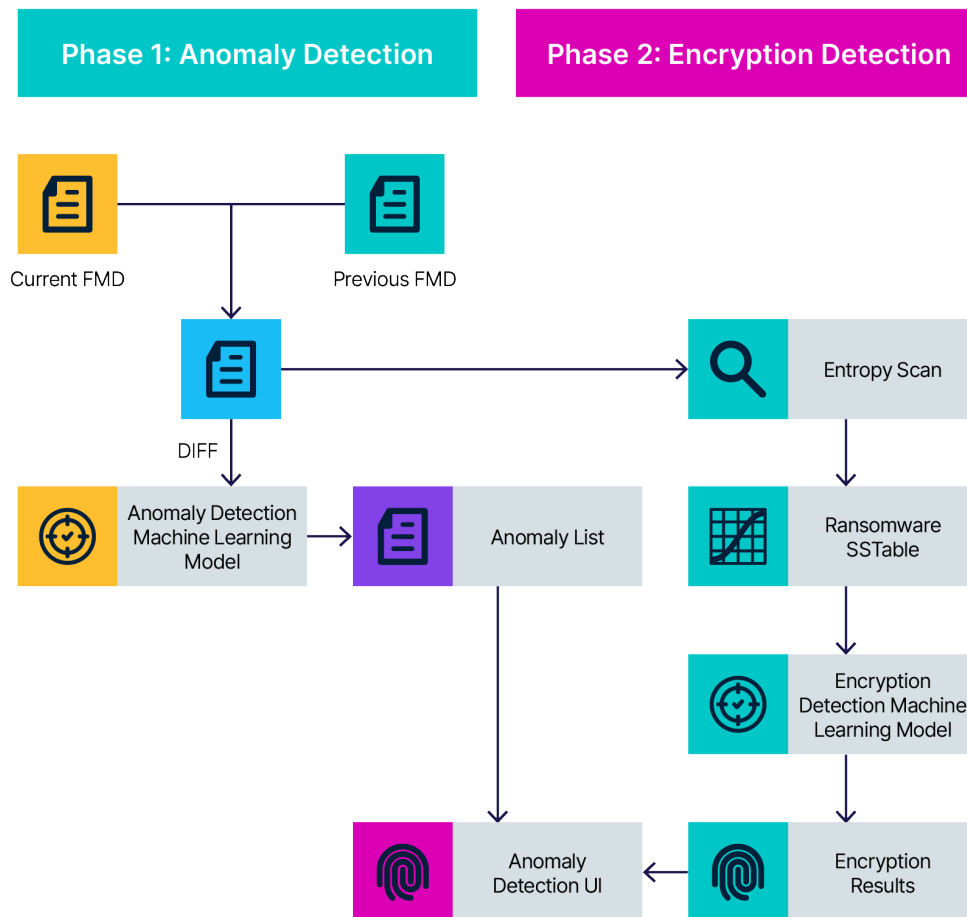
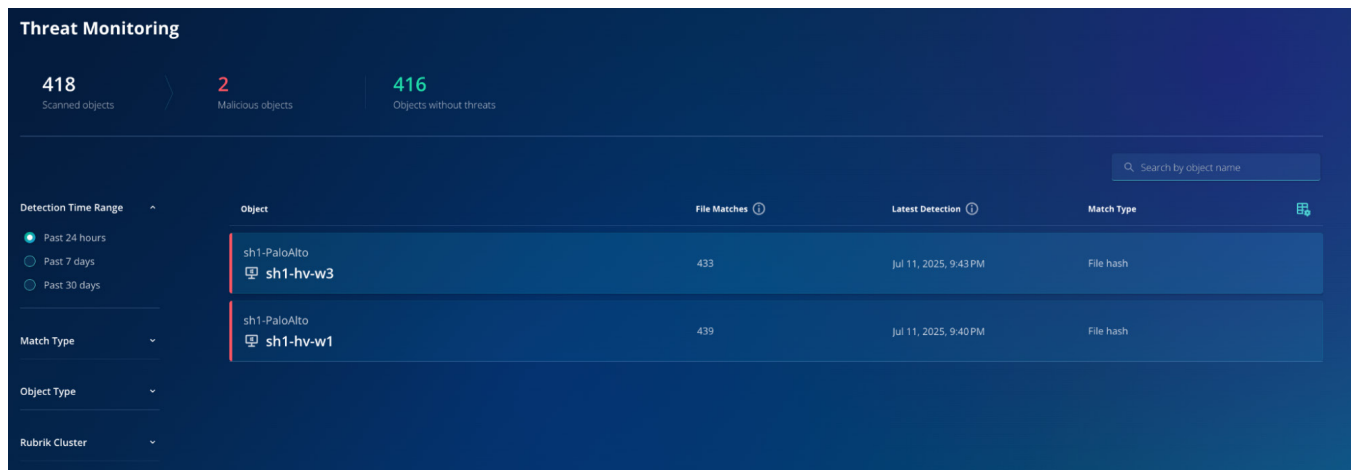


Figure 14

Anomaly Detection allows Rubrik customers to pinpoint the extent of cyber attacks by identifying all affected workloads, folders, and files. This functionality helps our customers reduce the time spent on discovering the impact of attacks and minimizes data loss by providing granular visibility. As a result, they can accelerate their response and recovery efforts.

Threat Monitoring

Threat Monitoring, a core component of Rubrik's Data Threat Analytics, provides continuous and automated scanning of your M365 backup data for Indicators of Compromise (IOCs). This “set it and forget it” feature ensures proactive and effortless threat detection without impacting production systems.



The screenshot displays the Threat Monitoring dashboard. At the top, there are three summary cards: '418 Scanned objects', '2 Malicious objects', and '416 Objects without threats'. Below these is a search bar labeled 'Search by object name'. On the left, there are filters for 'Detection Time Range' (Past 24 hours, Past 7 days, Past 30 days), 'Match Type', 'Object Type', and 'Rubrik Cluster'. The main table has the following data:

Object	File Matches	Latest Detection	Match Type
sh1-PaloAlto sh1-hv-w3	433	Jul 11, 2025, 9:43 PM	File hash
sh1-PaloAlto sh1-hv-w1	439	Jul 11, 2025, 9:40 PM	File hash

Here's how it works:

- 1. Pre-computed Hashes:** As part of Threat Monitoring, Rubrik calculates file hashes for executable files and stores the details in a hash table in RSC.
- 2. Automated Scanning:** Rubrik automatically initiates scans for IOCs in your time-series backup data. Threat Monitoring periodically checks for new snapshots after every backup, ensuring continuous active monitoring
- 3. Threat Intelligence Feed:** The most up-to-date and actively relevant IOCs associated with cyberattacks are integrated into the process, leveraging verified intelligence M365 hashes, sources from Google Threat Intelligence. Customers have the option to bring their own IoCs to RSC as well.
- 4. Match Reporting:** If any matches are found, the object name, match type, and location are immediately displayed on your dashboard.
- 5. Effortless Operation:** Once enabled, it automatically ingests new threat intelligence and scans for IOCs in your backups, requiring no further manual intervention.
- 6. Future Threat Hunting:** All executable files have a hash value calculated as part of threat detection, and these values, along with metadata about the file and snapshot, are stored in a hash catalog within RSC, which we'll see used in the next section.

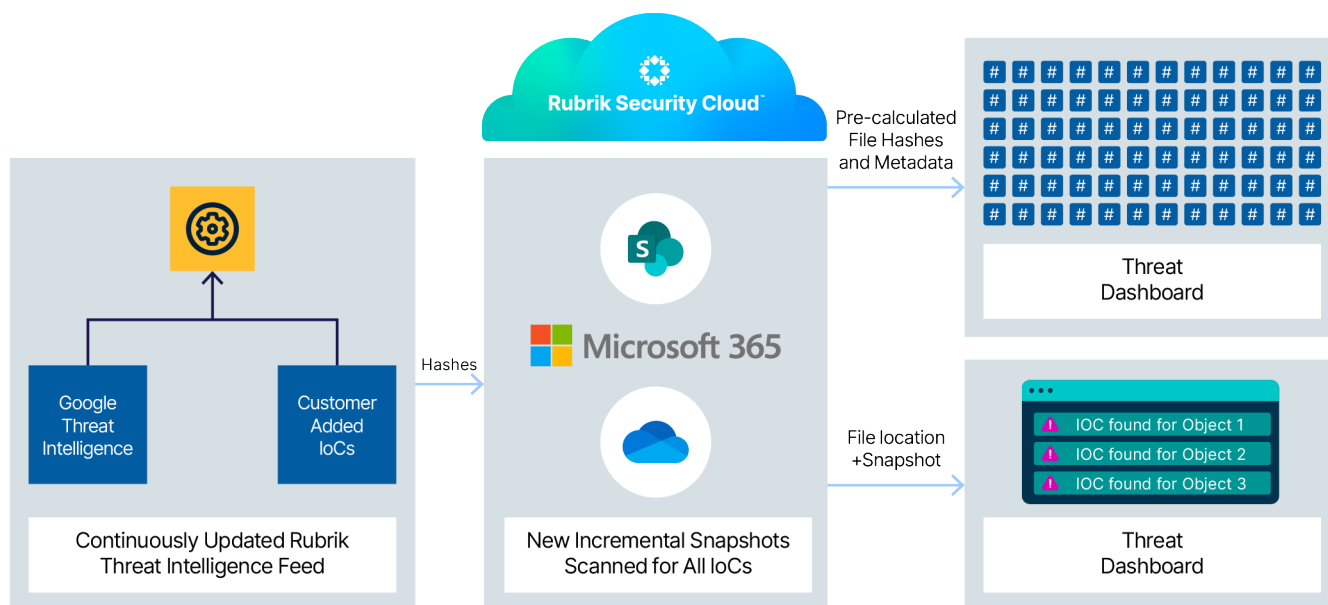


Figure 15

Rubrik Threat Monitoring offers an innovative solution for swift threat detection and response. It leverages automated, proactive threat detection on every backup using up-to-date, validated threat intelligence to help you stay ahead of evolving threats, reduce risk exposure, and mitigate incident impact. This proactive approach lets you tap into Rubrik’s expertise, avoiding costly disruptions and reputational damage.

Turbo Threat Hunting

Rubrik Threat Hunting drastically speeds up incident response by reducing the time to find malware-free recovery points from hours to seconds. This powerful capability allows organizations to quickly select secure recovery points, prevent reinfection, and significantly accelerate incident response and forensic analysis.

Threat Hunts > Lockbit IOC Threat Hunt

Lockbit IOC Threat Hunt

MATCHES **PARAMETERS**

All Objects ... **DOWNLOAD CSV** **START CYBER RECOVERY** **QUARANTINE**

Hunt Details	Name	Location	Files Matched	IOC Matches	Match Type	Earliest Matched Snapshot	Latest Matched Snapshot	Matched Sna...	Latest Snapshot without Mat...
Unique File Matches 6	<input type="checkbox"/> sh2-zaffre-fs-01	sh2-cork-vcsa.rubr...	1	10	YARA rule	July 10, 2025 at 12:40 AM	July 11, 2025 at 12:40 PM	10 / 10	--
Scan most recent snapshot before July 11, 2025 at 4:30 PM	<input type="checkbox"/> sh2-fs-01	sh2-cork-vcsa.rubr...	5	10	YARA rule	July 7, 2025 at 12:51 AM	July 11, 2025 at 1:12 PM	10 / 10	--

Here's how it works:

1. **Pre-computed Hashes:** This is calculated as part of Threat Monitoring as outlined in the previous section. Rubrik only needs to look in the hash table when a new IOC is discovered (e.g., a new zero-day). This allows for near instant identification of clean recovery points without needing to mount and scan individual files.
2. **Rapidly Hunt Across All Snapshots:** During an incident, security teams can simply input up to 100 specific IOC hashes to initiate a threat hunt. Rubrik [Turbo Threat Hunting](#) then swiftly scans the hash table rather than the actual backup environment to pinpoint unaffected recovery points.

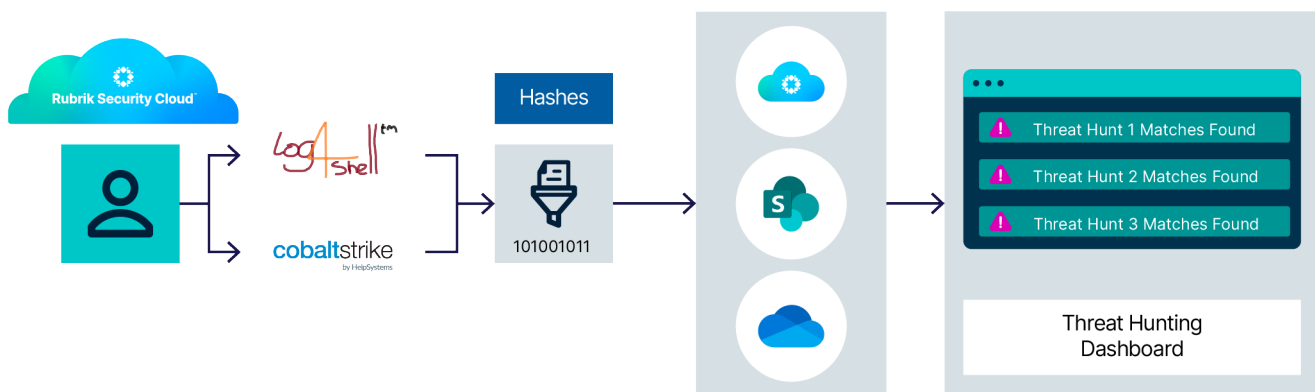


Figure 16

Turbo Threat Hunting empowers Rubrik customers' security teams to trace attack paths, identify impacted data, pinpoint precisely when the specific indicator of compromise was introduced, and accelerate response with actionable intelligence.

Now, based on the insights provided by the security application, Rubrik customers can proactively monitor their M365 data for the most common indicators of compromise, identify the blast radius for all affected data during any cyber or ransomware attack, identify clean recovery points, and perform recovery operations without reinfesting the production data.

DATA SECURITY POSTURE MANAGEMENT

In today's digital landscape, protecting sensitive data is more critical than ever. Proactive Data Security Posture Management (DSPM) by Rubrik offers a comprehensive solution to help organizations navigate the complexities of data privacy and security. By moving beyond traditional reactive approaches, Rubrik DSPM enables enterprises to proactively manage data risks in real-time environments. With powerful capabilities to discover, classify, and remediate data exposure, as well as ensure compliance with rigorous regulations, organizations can safeguard their valuable information while promoting secure innovation. Embrace a strategic approach to data security with Rubrik DSPM and empower your team to operate confidently in an evolving technology landscape.

Data Discovery and Classification

Data Discovery and Classification is a data classification tool that actively scans the contents of backups, looking for specific sensitive data (as outlined below). Data Discovery and Classification leverages the systems and application data within a Rubrik backup environment and uses that data to determine where this sensitive data exists and who has access. It can also classify specific sensitive data without the arduous deployment of individual agents or interfering with production systems.

Data Discovery and Classification fundamentally uses the concept of an analyzer and a policy. It uses an analyzer to define what should be identified in the contents of the data, and policies enable the bundling of multiple types of analyzers into a single report. Built-in analyzers are available out-of-the-box for common classifications such as social security numbers, email addresses, passport numbers, and credit card numbers.

At the time of writing, there are over 200 out-of-the-box analyzers, with new analyzers constantly added. As every customer’s needs are different, if you find that the shipped analyzers don’t meet your needs, then you can create custom analyzers to ensure the discovery of the data that is important to you. They can be tailored with customized dictionary terms or regular expressions to meet an organization’s particular needs.

Rubrik provides Data Discovery and Classification services on top of the two main applications within M365 that are prone to show the most risk: OneDrive and SharePoint. Organizations can easily discover exactly what sensitive data they have within M365 and take proactive steps to ensure it’s either deleted, moved, or properly secured, essentially lowering the impact of an attack. The following provides a detailed description of how Rubrik’s Data Discovery and Classification processes work with Microsoft 365 OneDrive and SharePoint data.

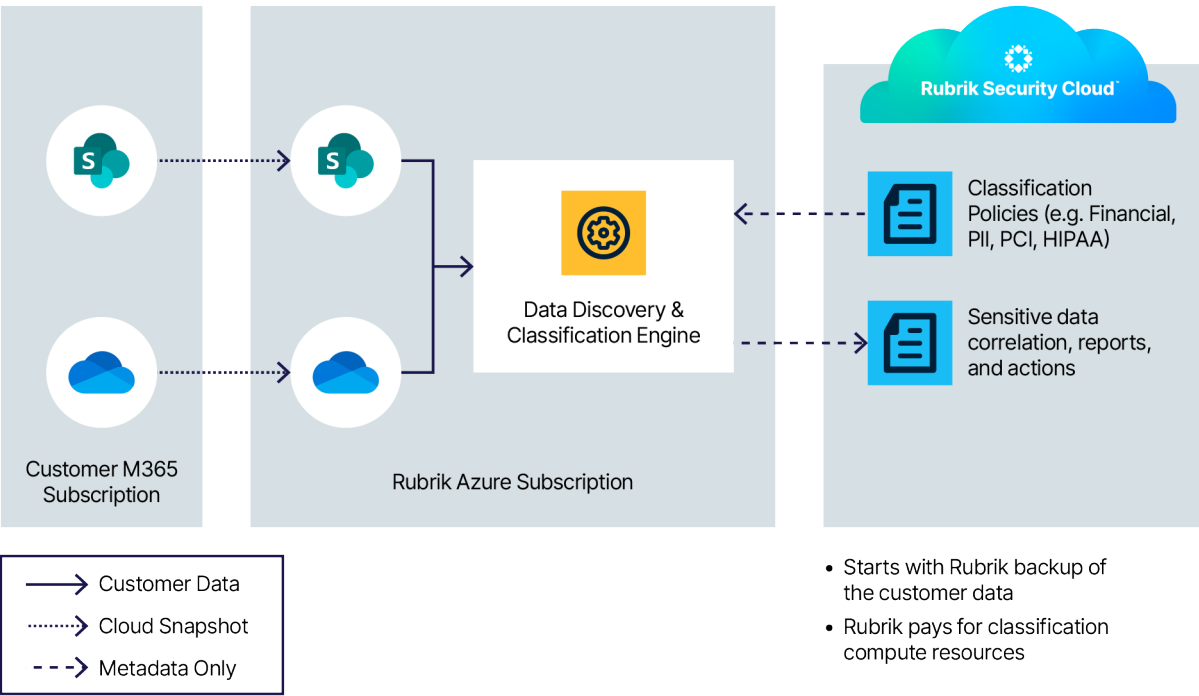


Figure 17 – High-Level Architecture of Data Discovery and Classification

1. Rubrik first creates backups for OneDrive and SharePoint as described earlier within a customer-specific Rubrik Azure subscription. This subscription is dedicated to the customer and won't be shared with any other customer's data.
2. The Data Discovery and Classification Engine also operates within this subscription and will ingest the backed-up data and process it based on classification policies that are defined in RSC.
3. Metadata about the sensitive data that was discovered, any correlation with user access, and reports will then be sent to Rubrik Security Cloud and displayed within the UI. Actions can also then be taken to remediate any identified issues. No actual customer data is sent outside the customer's environment.
4. All the compute resources required for data classification are instantiated in the Rubrik Azure subscription.

Note that backups are required for M365 DSPM and that all data processing happens in the isolated Rubrik subscription.

Rubrik's Data Discovery and Classification for M365 uses Exocompute within the customer-specific Rubrik environment to process the scanning, and at no time is any of the organization's sensitive data uploaded to Rubrik Security Cloud - only generated metadata and overall scan results are sent to RSC.

Data Access Governance

Least privilege access is a simple term but a difficult principle to implement. In a typical organization, you'll see a mixture of access directly applied to M365 data and access gained through roles and groups. Multiply this by the number of different platforms and products in use; figuring out where to start can be challenging. Data Access Governance significantly simplifies establishing a baseline and helps you proactively monitor your environment to stay on top of changes as they occur.

For on-premises and M365 workloads, Rubrik starts by taking backups of Active Directory, Entra ID, VMs, SharePoint, and OneDrive. Once you have identified where sensitive data is present across your environments and classified the types of data there, your next step is to map the data to those with access. The Data Access Graph provides a simple visualization of who has access to what, making it easy to ensure least privilege access.

With visibility into which identities have access to which data, Identity Inventory can help identify excessive and potentially misconfigured permissions, breaking down the different types of data and the different levels of access that specific Roles, Groups, and Users have access to. By mapping in this way, it can easily be determined what the potential impact would be if a specific identity were compromised.

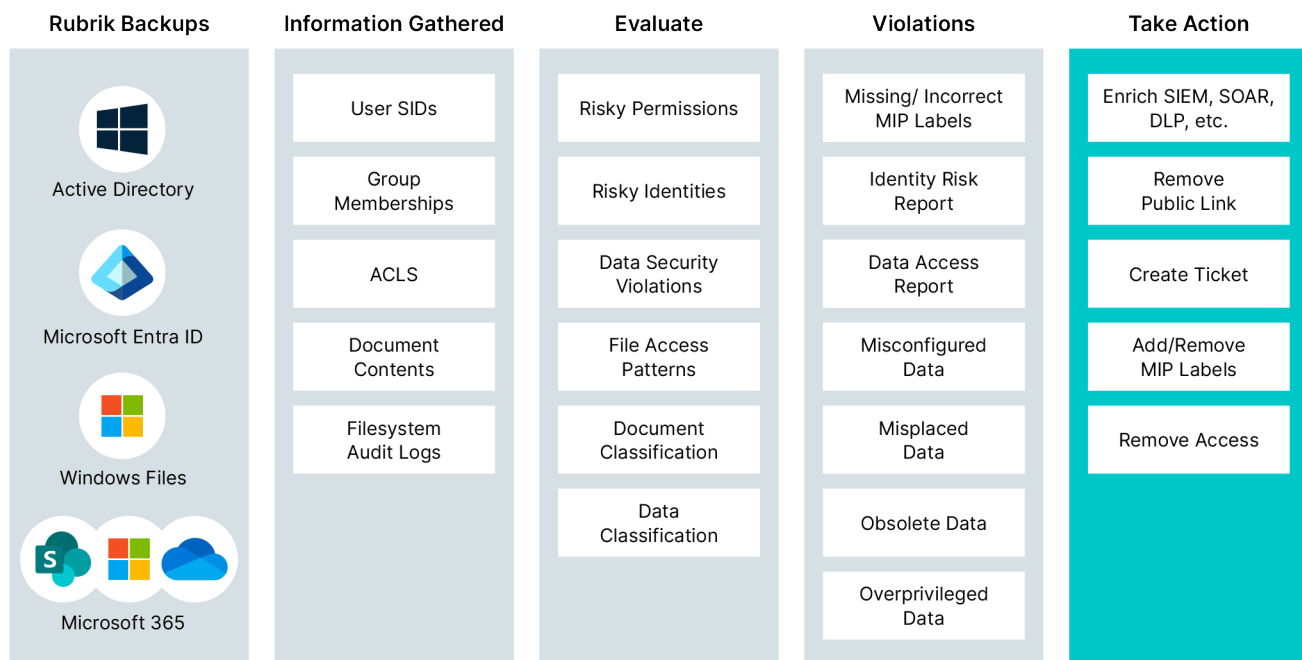


Figure 18 – High-Level Architecture of Data Access Governance

1. For on-premises environments and M365 workloads, Rubrik begins by backing up key components, including Active Directory, Entra ID, VMs, SharePoint, and OneDrive.
2. From these backups, Rubrik collects vital information such as user SIDs, group memberships, file-level permissions, the contents of each file, and filesystem-level audit logs. The Data Discovery and Classification Engine then analyzes this information to identify risky permissions, potentially problematic identities, and improper file access patterns based on the audit logs. It also investigates and classifies sensitive data, assessing individual documents accordingly.
3. Using the policies defined within RSC, violations are detected and reported. These violations may include missing or incorrect MIP labels, identity risks, data access issues, and misconfigured, misplaced, obsolete, or overly permissive data.
4. Identifying these violations is just the beginning; the issues must also be resolved. Rubrik DSPM provides customers with effective options for quickly remediating these violations directly from the violation reports.
5. Information about these findings can be integrated with other applications, such as SIEM, SOAR, DLP, or ITSM ticketing solutions. Additionally, Rubrik can take actions in the live environment to enhance data security, including the removal of public links or risky permissions.

MIP Labeling

Microsoft Information Protection (MIP) labels play a crucial role in this effort by allowing organizations to classify, control access, and enforce protection settings like encryption for their sensitive data. These labels are embedded directly within the files, ensuring that they travel with the data, regardless of where it moves across platforms and devices. With the integration of MIP labels into data loss prevention (DLP) and cloud access security broker (CASB) tools, organizations can enhance their data security posture.

Rubrik's DSPM leverages AI-driven classification to automatically apply MIP labels without requiring user intervention. This capability not only enforces consistent data protection policies, even when users may inadvertently mislabel or skip files, but also prepares organizations for the integration of AI tools like Copilot, ensuring that thousands of documents are accurately labeled and accessible only to the appropriate entities.

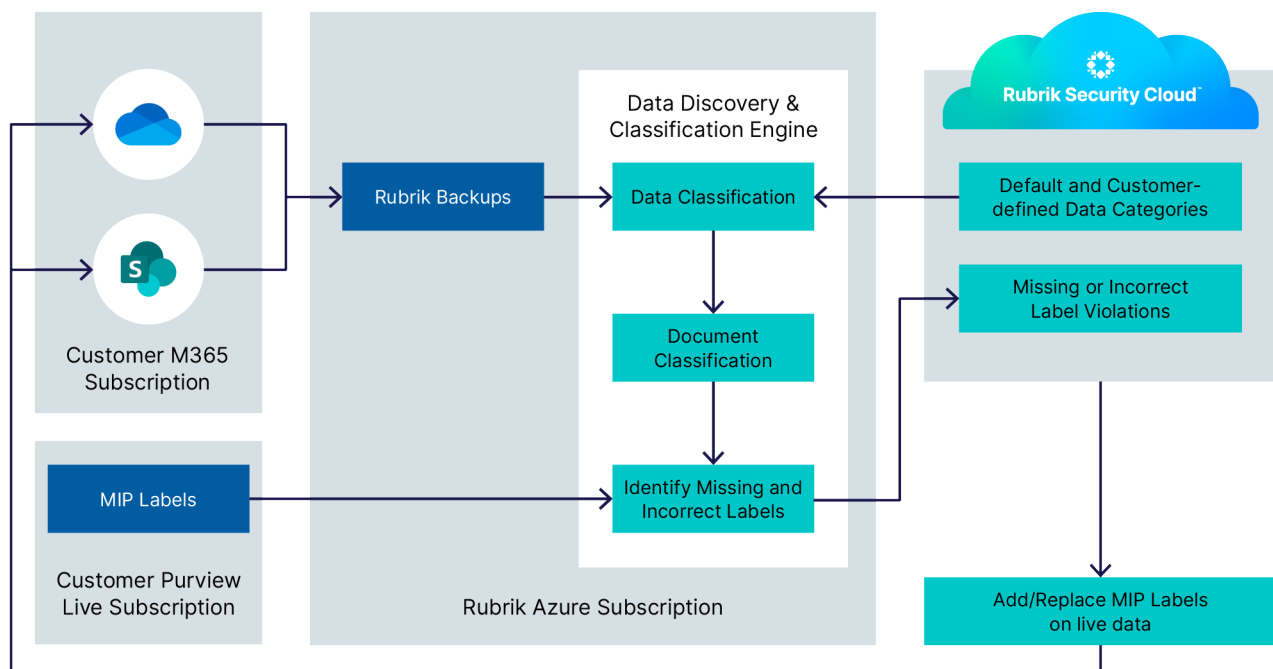


Figure 19 – High-Level Architecture of MIP Labeling

1. Rubrik is equipped to generate insights from data and identity information as described earlier. The platform is particularly well-suited to assist customers in managing their MIP (Microsoft Information Protection) labeling at scale within their M365 environment. This process encompasses the customer's OneDrive and SharePoint environments, their Purview configuration, a Rubrik Azure subscription, and Rubrik Security Cloud.
2. The procedure begins with backing up the data from the OneDrive and SharePoint environments and storing these backups in a customer-specific Azure subscription provided by Rubrik.
3. Following this, the newly backed-up data is classified according to the data categorization policies defined within RSC.
4. Using AI, the platform analyzes the patterns of data categories and the business context of each document to identify its type, whether it is a tax document, bank statement, patient information, or another category.

5. Based on this classification information, the Data Discovery and Classification Engine evaluates whether the current MIP labels associated with each document are appropriate. If there is a mismatch, violations are generated within RSC.
6. This process allows customers to take direct action on these violations, enabling them to add or replace labels in the live environment. In the near future, these actions will be automated, eliminating the need for human intervention.
7. The entire process is triggered whenever a backup captures a file modification or a change to MIP labels.

Rubrik DSPM for Microsoft 365 is designed for the era of AI tools like Copilot, which are becoming essential for collaboration and creation within teams. As the risk of data leakage increases, Rubrik DSPM ensures that your data security is robust. It provides the necessary classification, labeling, and access controls to protect sensitive information. With Rubrik, you don't just back up your Microsoft 365 data; you secure it, classify what is important, and take proactive measures to prevent damage from attackers or AI.

SUMMARY

This concludes the How It Works guide on protecting Microsoft 365 with Rubrik Security Cloud. The guide explained the general architecture and overall value proposition of Rubrik's Microsoft 365 protection as well as educated the reader on the nuances of each major workflow within the product. Additionally, the guide equipped the reader with some common techniques and best practices for using Rubrik Security Cloud efficiently from both an operations and cost perspective, allowing them to fully understand and unlock the potential of protecting Microsoft 365 with Rubrik.

APPENDICES

APPENDIX A: REQUIRED MICROSOFT 365 API PERMISSIONS

OFFICE 365 SHAREPOINT ONLINE permissions are needed for all OneDrive, SharePoint and Teams apps.

EXCHANGE WEB SERVICE

- Calendars.ReadWrite.All
- Contacts.ReadWrite
- full_access_as_app
- Mail.ReadWrite
- Tasks.ReadWrite
- User.Read.All
- Tasks.ReadWrite

MICROSOFT GRAPH—GENERAL

- User.Read.All
- Group.Read.All
- Reports.Read.All

MICROSOFT GRAPH—EXCHANGE

- Calendars.ReadWrite
- Contacts.ReadWrite
- Mail.ReadWrite
- Group.Read.All
- User.Read.All
- Reports.Read.All

MICROSOFT GRAPH—ONEDRIVE

- User.Read.All
- Sites.Read.All
- Files.ReadWrite.All
- Sites.FullControl.All

MICROSOFT GRAPH—AZURE APP

- Directory.AccessAsUser.All
- User.Read
- user_impersonation

MICROSOFT GRAPH—TEAMS

- Group.ReadWrite.All
- Channel.Create
- Teamwork.Migrate.All
- ChannelMessage.Send
- Chat.ReadWrite
- Chat.ReadWrite.All
- ChannelMessage.Read.All
- Sites.Read.All
- Files.ReadWrite.All
- User.Read.All
- ChannelMessage.Read.All
- Sites.FullControl.All

MICROSOFT GRAPH—SHAREPOINT

- Sites.Read.All
- Files.ReadWrite.All
- User.Read.All
- Sites.FullControl.All

OFFICE 365 SHAREPOINT ONLINE

- Sites.FullControl.All

MICROSOFT GRAPH—MANAGEMENT

- ServiceHealth.Read.All
- Group.Read.All
- User.Read.All
- Sites.Read.All
- Reports.Read.All
- MailboxSettings.Read
- Mail.ReadBasic.All
- AdministrativeUnit.Read.All
- Sites.Read.All
- Application.Read.All

VERSION HISTORY

Version	Date	Summary of Changes	Author
1.0	April 2024	Initial Release	Mike Preston
1.1	August 2024	Recovery section updates	Mike Preston
1.2	September 2025	Data Threat Analytics and DSPM sections added	Alpika Singh



Global HQ

3495 Deer Creek Road
Palo Alto, CA 94304
United States

1-844-4RUBRIK
inquiries@rubrik.com
www.rubrik.com

Rubrik (NYSE: RBRK) is on a mission to secure the world's data. With Zero Trust Data Security™, we help organizations achieve business resilience against cyberattacks, malicious insiders, and operational disruptions. Rubrik Security Cloud, powered by machine learning, secures data across enterprise, cloud, and SaaS applications. We help organizations uphold data integrity, deliver data availability that withstands adverse conditions, continuously monitor data risks and threats, and restore businesses with their data when infrastructure is attacked.

For more information please visit www.rubrik.com and follow @rubrikInc on X (formerly Twitter) and Rubrik on LinkedIn. Rubrik is a registered trademark of Rubrik, Inc. All company names, product names, and other such names in this document are registered trademarks or trademarks of the relevant company.