



TECHNICAL OVERVIEW & HOW IT WORKS

The Definitive Guide to Rubrik Cloud Data Management



Table of Contents

3 INTRODUCTION

- 3 Infrastructure-as-a-Bottleneck
- 3 Meet Rubrik
- 5 The Principles Behind Our Technology

6 TECHNOLOGY OVERVIEW

- 6 Atlas: Cloud-Scale File System
- 7 Callisto: Distributed Metadata System
- 7 Cerebro: Data Management Layer
- 8 Infinity: Interface & Application Awareness

9 HOW IT WORKS

- 9 Simple Set Up
- 9 Automatic Data Discovery
- 9 Flash-Speed Data Ingestion
- 10 Single SLA Policy Engine
- 12 Continuous Data Protection (CDP)
- 12 API-First Architecture
- 13 Global Predictive Search
- 13 Live Mount
- 14 Security & Compliance

16 CLOUD SOLUTIONS

- 16 CloudOut
- 18 CloudOn
- 19 Cloud Cluster
- 19 AWS Native Data Protection

21 API INTEGRATIONS

23 COMPREHENSIVE PLATFORM SUPPORT

- 24 Virtual Environments
- 25 Physical Operating Systems
- 25 Physical Databases
- 28 Network-Attached Storage (NAS)
- 28 NoSQL Databases
- 29 SaaS Applications

31 BUSINESS USE CASES

- 31 Self-Service Data Management via ServiceNow
- 33 Live Mount for Test/Dev
- 35 Remote and Branch Office Data Protection
- 36 Automated Sensitive Data Discovery and Classification
- 38 Multi-Layered Ransomware Defense
- 40 Centralize Hybrid Cloud Management with Polaris GPS

41 CONCLUSION

Introduction

Infrastructure-as-a-Bottleneck

Enterprise IT today is starkly different from enterprise IT yesterday. The modern IT organization looks to innovate quickly in order to align with critical business objectives and drive growth. To do so, they are adopting initiatives such as transitioning to a software-defined data center, moving to public cloud, or accelerating automation with infrastructure-as-code. However, they are still spending too much time and money just trying to keep the lights on.

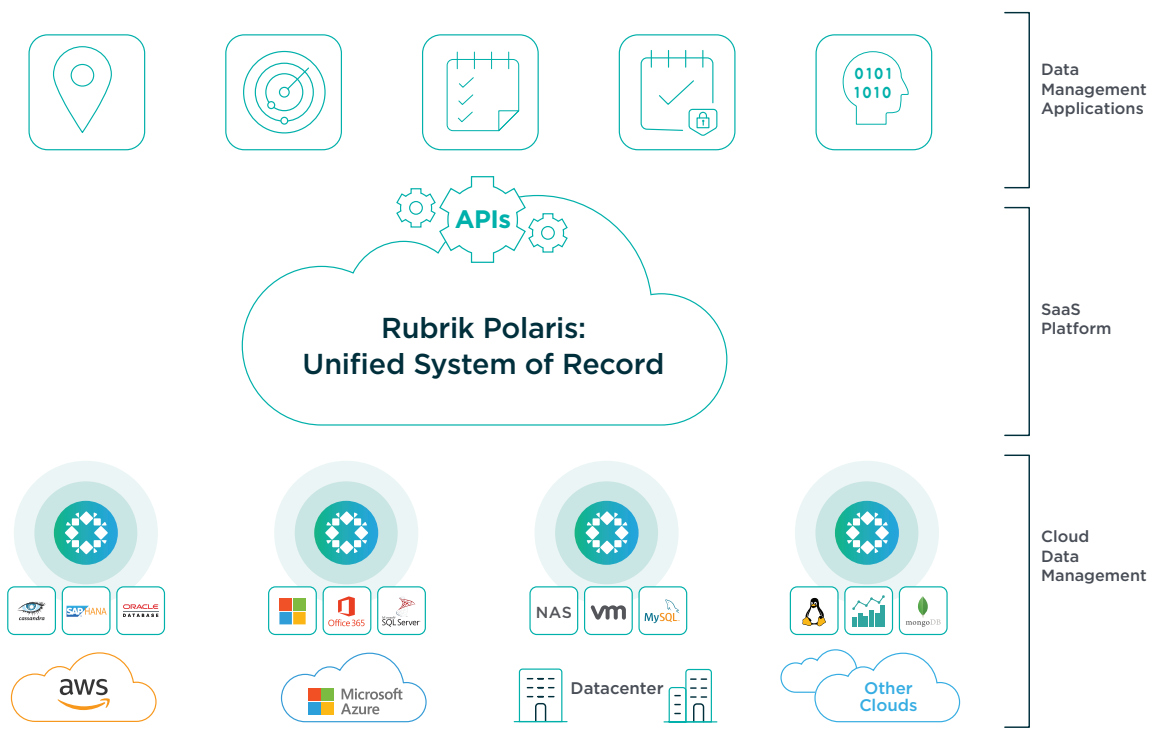
Legacy backup and recovery infrastructure – with its complex, multi-tiered architecture – cannot support forward-looking initiatives. It is time-consuming to manage, stubborn to scale, and expensive to maintain. On top of that, introducing a new cloud provider or application is a massive headache, if possible at all. IT organizations need a solution that eliminates these legacy bottlenecks with an easy-to-use software built to run anywhere, freeing up time for more strategic projects.

Meet Rubrik

Rubrik delivers a single software platform for complete enterprise data management across data center to cloud. It has been designed from the ground-up to simplify and automate core backup and recovery capabilities while also enabling API-first automation, easy cloud mobility, and defense against cyber attacks and data breaches.

Rubrik's product portfolio includes the following:

- **Rubrik Cloud Data Management Platform:** Rubrik's flagship data management solution delivers backup, recovery, data archival, and disaster recovery on-premises, at the edge, and in the cloud.
- **Rubrik Polaris SaaS Platform:** Polaris delivers a unified system of record to organize global business applications and data managed by Rubrik Cloud Data Management. With Polaris, Rubrik delivers new value-add services, such as AI-driven ransomware recovery, governance and compliance, and data intelligence.
- **Rubrik Mosaic:** Rubrik Mosaic simplifies and automates application-consistent backup and recovery for NoSQL databases with a policy-driven approach, enabling digital-minded enterprises to build next-generation customer-facing applications. Rubrik Mosaic can be installed on-premises, in any cloud, and is also available via Microsoft Azure Marketplace.



Rubrik provides Cyso with unprecedented simplicity that allows my team to focus on our customers and continue to innovate. We now have the scalability and flexibility to grow with our customers and all their data management needs.

Tjebbe de Winter
Co-founder and CTO, Cyso

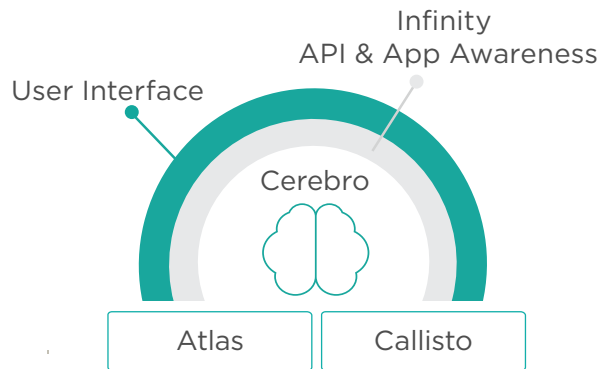


The Principles Behind Our Technology

The Rubrik Cloud Data Management platform incorporates the following design principles:

- 1. Software-defined:** Rubrik consolidates disparate hardware and software components into a single software fabric. Enterprises can run Rubrik anywhere via plug-and-play appliances on-premises, as software on third-party hardware, or as software in the cloud.
- 2. Simplicity at scale:** Rubrik wipes out management complexity with just a few clicks. By adopting the same web-scale technologies used by Google, Facebook, and Amazon, users can easily handle rapidly increasing volumes of information with a linear-scale architecture.
- 3. Cloud mobility:** Rubrik enables cloud and application mobility for all enterprises. Whether deployed on-premises or in the public cloud, Rubrik can protect cloud-native applications, search across applications and files, and quickly spin up instances for test/dev or DR.
- 4. API-first:** Rubrik is the industry's only data management solution built on an API-first architecture. To help users leverage the APIs, Rubrik offers industry-standard OpenAPI documentation, sample code, and pre-built integrations with a variety of automation tools.
- 5. End-to-end security:** Rubrik offers a comprehensive, multi-faceted approach to security. Throughout the entire lifecycle, data is encrypted in-flight and at-rest. Enterprises can recover from ransomware attacks quickly by reverting to the most recent clean snapshot. All applications and data are stored in an immutable format.
- 6. Broad ecosystem support:** Rubrik's software platform is designed to be vendor-agnostic, supporting industry-leading applications, operating systems, databases, hypervisors, clouds, and SaaS applications. Rubrik provides freedom of choice, allowing enterprises to operate both on-premises and in any cloud while avoiding vendor lock-in.

Technology Overview



Rubrik Cloud Data Management is a single software fabric comprised of four main components: Atlas (Cloud-Scale File System), Callisto (Distributed Metadata System), Cerebro (Data Management Layer), and Infinity (Interface and Application Awareness).

Atlas: Cloud-Scale File System

Rubrik's Cloud-Scale File System is built from scratch to store and manage versioned data and serves as the foundation of the distributed file system. Key properties include:

- **Fault tolerance:** Atlas is designed with no single point of failure. Rubrik uses erasure coding to intelligently encode and partition data such that in the event of a node or two-disk failure, the original data can still be reconstructed. Erasure coding increases usable capacity while maintaining high performance.
- **Self-learning and self-healing:** In the event of a failure, Rubrik automatically self-adjusts and re-balances to ensure resiliency.
- **Linearly scalable:** Rubrik is architected to be an infinitely expandable, web-scale system. Increasing capacity and I/O performance is as simple as adding more nodes to the cluster.
- **Zero-byte clones:** The system utilizes zero-byte clones to make multiple copies from a single "golden image." This accelerates application test and development without a storage penalty. Atlas is also designed for zero-copy recovery, which enables instant recovery and business continuity.
- **Efficient data reduction:** The system delivers intelligent data reduction to optimize storage and bandwidth consumption both on-premises and in the cloud.
- **Flash-optimized:** Rubrik maximizes I/O throughput with a hybrid flash/disk architecture.

Callisto: Distributed Metadata System

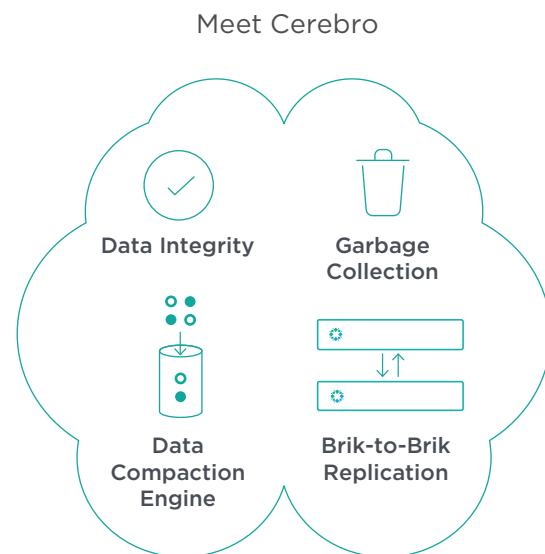
Callisto, Rubrik's Distributed Metadata System, operates alongside the Atlas Cloud-Scale File System to provide a global index and catalogue that can be accessed at high speeds for file-level search. It is built to deliver continuous availability, linear scalability, and operational simplicity with no single point of failure. In the event of multiple node failures, the system maintains continuous access to metadata to facilitate low-latency search.

Cerebro: Data Management Layer

Cerebro serves as the “brains” of the Rubrik system. It creates a master control plane to handle data lifecycle management from ingest to expiration. Cerebro is comprised of two key components: the Blob Engine and Distributed Task Framework.

Blob Engine: The Blob Engine is a distributed version control system that abstracts data from underlying infrastructure to enable application mobility across data centers and clouds. It provides core data management services, including immutability, data reduction, retention, replication, and archival. Key properties include:

- **Data integrity:** The Blob Engine captures all applications and data in an immutable format, preventing cyber threats such as ransomware from encrypting underlying backup data. It ensures data integrity by applying multiple integrity checks within the entire data management stack.
- **Metadata indexing:** The Blob Engine indexes relevant metadata like file name, type, and content to build a single metadata repository. This enables global file-level search across applications, data centers, and clouds.
- **Data reduction:** Cerebro intelligently applies global data reduction while enabling fast data reconstruction to maximize storage efficiency throughout all stages of the data lifecycle.
- **SLA automation:** The Blob Engine ensures that all SLAs are met and marks snapshots as ready for secure deletion when needed. Reports are available to end users to track compliance.



Distributed Task Framework: The Distributed Task Framework globally assigns and executes tasks across the Rubrik cluster in a fault-tolerant and efficient manner. It applies intelligent algorithms to load balance and optimize resource utilization through two methods: Task Scheduling and Task Maintenance. Task Scheduling ensures that tasks are evenly distributed across the cluster while Task Maintenance enforces SLA policies on a daily and long-term basis. Once an SLA policy is set, Task Maintenance strategizes to meet these set goals for data retention, replication, and archival based on priority.

Infinity: Interface & Application Awareness

Infinity is where Rubrik interacts with the external ecosystem in order to bring context and application awareness to the platform. This layer communicates with different data providers via APIs, connectors, and other protocols to ensure that Rubrik can discover and efficiently protect any supported object. Any command that leverages Rubrik's APIs connects to the Infinity interface and leverages the policy engine.

How It Works

Simple Set Up

The system invokes multicast DNS protocols to automatically discover and self-configure each of the nodes within the Rubrik cluster. Users simply assign IP addresses to each of the nodes and provide login credentials for the environments to be managed by Rubrik. To expand or reduce cluster size, users can simply assign new IP addresses or remove nodes directly from the Rubrik UI.



<1 hour to get up and running



Automatic Data Discovery

Discovery is the process by which Rubrik detects the various data sources across physical, virtual, and cloud environments. Auto-discovery happens in a variety of ways depending on the user environment. For example, for physical Oracle databases, Rubrik discovers workloads via Rubrik Backup Service: a lightweight connector. Rubrik Backup Service automates new deployments and updates across all connectors through Rubrik's API, eliminating manual agent management. For cloud environments, such as AWS, users can simply provide login credentials directly from the Rubrik UI.



Rubrik was up and running very quickly. It auto-discovered our entire virtual environment, and we were impressed by how simple it was to deploy.


Jörg Voosen

IT Systems Team Manager, UTSCH

Flash-Speed Data Ingestion

Rubrik is designed as a high-speed data ingestion engine that can easily handle large volumes of data for all environments. Since Rubrik is architected to be a web-scale system, performance for ingestion and disk throughput increases predictably and linearly as more nodes are added to the cluster. As a scale-out system,

Rubrik's distributed workflow management system maximizes the number of parallel data streams processed. Rubrik also delivers an incremental-forever approach for faster performance. For example, for physical Windows or SQL Server environments, Rubrik automates the discovery of changed blocks and the transaction logs that they are associated with. For virtual environments, such as VMware vSphere, Rubrik utilizes VMware's Changed Block Tracking to deliver incremental-forever snapshots.



Single SLA Policy Engine

A Rubrik SLA domain is a declarative policy that captures the core objectives for data protection and lifecycle management. The basic components of an SLA domain are as follows:

- Backup frequency:** This is how often snapshots are taken. Users can specify backup frequency in terms of constant intervals (e.g. take a snapshot every 4 hours) or in terms of calendar elements (e.g. take a snapshot on the last day of each month) for granular control over RPOs.
- Continuous Data Protection (CDP):** To enable CDP, users simply toggle a switch when creating an SLA domain.
- Snapshot window:** This is the specified time in which Rubrik takes snapshots. For example, users can choose to backup from 12pm to 4am. If a snapshot occurs outside of an allowed window, it will be rescheduled for a time that falls within that window.

Create SLA Domain

SLA Domain Name
MISSION CRITICAL

Continuous Data Protection Keep Recovery Points For (Hours)
4

Advanced Configuration

Service Level Agreement
Choose how often we take snapshots and the length of time we keep them.

| Take Snapshots: | Keep Snapshots: |
|----------------------------|---------------------------|
| Every (Hours) 4 | For (Days) 7 |
| Every (Days) 1 | For (Days) 30 |
| Every (Months) 1 | For (Months) 12 |
| Every (Years) 1 | For (Years) 7 |

Local retention set to 7 years.

Snapshot Window

Take snapshots from: _____ to _____

Take first full between: First Opportunity at _____

[Remote Settings](#)

- **Archive location:** Users can automate archival to public or private cloud for cost-effective long-term retention. Rubrik supports archival to public cloud storage, S3-compatible object stores, NFS, and tape.
- **Retention period:** This is how long users keep snapshots before they expire. Users can split the retention period into a first interval where the data resides locally in Rubrik and a second interval where the data resides in the archive location. Users specify the retention period with a simple sliding bar (e.g. 30 days). To define WORM-compliant retention periods in accordance with regulations SEC Rule 17a-4(f) and FINRA Rule 4511(c), users can apply Retention Locked SLA domains, which ensure that retention periods cannot be lowered or removed.
- **Replication:** Users can replicate to other data centers and clouds in the same SLA domain used for backup and archival. Retention policies for replicated data are defined with a simple sliding bar.

With Auto Protect, objects automatically inherit policies assigned to parent folders and hosts. The most granular policies automatically take priority. Since Rubrik is vendor-agnostic, a single policy can be applied across various applications or locations. For organizations that leverage vSphere tags to organize and group VMs, Rubrik also allows users to assign policies at the vSphere tag-level. SLA domains can be paused or resumed at any moment, allowing users to perform maintenance activities on a subset of protected objects without pausing protection for the entire cluster.



Rubrik's interface is easy enough for anyone on our team to use and requires no special training. We've gone from spending 15 hours a week monitoring backups to 1 hour a week. We're now using this time to invest in automation and strategic initiatives that drive greater business value.

Adam Monnery

Head of Information and Communications Technology, Museum of London



Continuous Data Protection (CDP)

For vSphere VMs, Rubrik provides users with two options for automated protection. Users can deploy a snapshot-based approach as specified by the backup frequency and retention policies, or users can leverage Continuous Data Protection (CDP), which enables near-zero RPOs. To deliver CDP, Rubrik integrates with VMware APIs to capture IO on VMs. Rubrik journals these write operations in a log-file to capture granular changes between snapshots and create a continuous stream of recovery points.

In a snapshot-based approach, the nearest recovery point can be up to multiple hours in the past. However, many enterprises require RPOs on the order of minutes or seconds for more critical applications. With Rubrik's CDP, users can recover to any point-in-time, up to seconds before an application failure or disaster. Activating CDP is as easy as toggling a single switch when creating a Rubrik SLA domain.

API-First Architecture

Rubrik is the industry's only data management solution built on an API-first architecture. This means that every feature available through the Rubrik UI uses Rubrik's own REST APIs. Users can leverage the APIs to easily integrate data protection with existing tools and build new custom services.

Enterprises commonly use Rubrik APIs to integrate backup and recovery into IT service catalogues (e.g. ServiceNow, VMware vRealize Automation, or VMware vCloud Director), simplify management of large, distributed environments via configuration management (e.g. Puppet, Chef, SaltStack, and Ansible), automate lifecycle data management workflows, and centralize monitoring and reporting.

For more information on Rubrik's API integrations, please see [API Integrations](#).



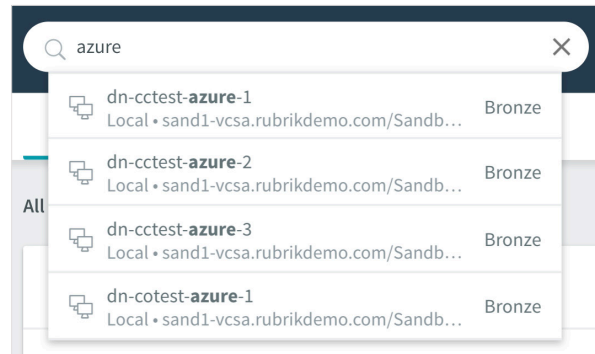
By using Rubrik's integration with vRO, we've automated the process of creating and deleting VMs before archiving it to AWS. This automated workflow was previously impossible without Rubrik.

Daniel Jenkins
IT Network Engineer, Omnitrac



Global Predictive Search

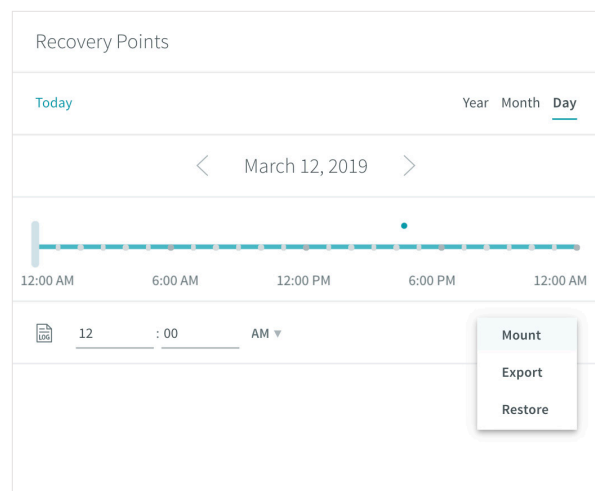
Rubrik offers intuitive search to quickly locate VMs, physical applications, and files regardless of location – on-premises or in the cloud. As users type queries, suggested search results are displayed with auto-complete functionality. For all data sources, Rubrik indexes metadata on key attributes, such as file name, type, and content. A record of all metadata is accessible anytime, even if data is stored in the cloud.



Live Mount

Rubrik's Live Mount delivers near-zero RTOs to radically accelerate data access for instant recovery and application test/dev. Users can browse and recover granular files, objects, and tables from VMs as well as physical SQL Server and Oracle databases.

Rubrik's Cloud-Scale File System has built-in zero-byte cloning capabilities that minimize storage consumption while allowing any number of mounts to be created. As developers alter the mounted instances for test/dev, Rubrik intelligently stores the deltas by forking to new branches. All writes and hot reads utilized during Live Mount are automatically allocated to the flash tier to optimize I/O performance.



For VMware vSphere VMs, post-recovery, users can choose to Storage vMotion the VMDK to the primary storage environment or continue using Rubrik as a storage endpoint. For Oracle databases, Rubrik's Live Migration enables Oracle data files to be migrated back to production while Oracle database files are still being actively served by Rubrik cluster, thus eliminating additional downtime.

For further information on how Live Mount is utilized for application development, please see [Live Mount for Test/Dev](#).



90%

faster restores and minutes to provision Oracle database clones



Security & Compliance

Rubrik offers a comprehensive approach to security regardless of where data is located. To help facilitate the privacy and safety of enterprises' data, Rubrik deploys a multi-layered security framework that consists of the following components:

- **Data-at-rest encryption:** Rubrik encrypts all data at-rest to protect against physical breaches. Data is still secure even if a drive is stolen from a data center. Rubrik delivers both software (FIPS 140-2 Compliant AES-256) and hardware (FIPS 140-2 Level 2 HDD and SSD) encryption.
- **Data-in-flight encryption:** Rubrik encrypts all data before leaving the system, ensuring secure data archival to public or private cloud environments. Rubrik leverages client-side encryption libraries supported by public cloud providers and all archived data undergoes envelope encryption.
- **Flexible key management:** Rubrik offers the flexibility to manage keys with an internal key manager (via the Trusted Platform Module chip) or an external key manager (via Key Management Interoperability Protocol). In both cases, Rubrik facilitates security best practices by allowing users to easily execute a one-time key rotation or automate recurring rotations. Rubrik key management also enables secure cluster erasure to provide government agencies with added security.
- **User authentication:** Rubrik reduces the risk of data breaches and cyber attacks by assigning granular permissions for data access. The Rubrik platform integrates with Active Directory (AD) and supports granting authorizations and groups from the AD. With support for SAM 2.0, users can also securely access the cluster with Single Sign-on from their SAML 2.0 compatible IdP (Identity Provider). This enables users to access multiple applications with a single set of credentials without the interoperability issues associated with vendor-specific designs. Rubrik clusters use Role-Based Access Control (RBAC) to define the capabilities of authenticated users, and also employ multi-factor authentication and API tokens for added layers of security.
- **Data integrity:** Rubrik provides data immutability against cyber attacks such as ransomware. No external or internal operation can modify the data since the underlying backups are read-only.
- **Centralized compliance reporting:** All user activity in the Rubrik platform is logged and made available to administrators through the Rubrik UI and APIs. The activity logs centralize compliance reporting for system operations such as backup and recovery.



We operate in a heavily regulated industry that requires our data management solution to be highly secure and reliable. Rubrik employs a comprehensive approach to security and ensures our data privacy and safety. This includes end-to-end encryption across our entire on-prem and cloud environment as well as audit capabilities for meeting compliance.

Neill Smith

Head of IT Infrastructure, Scottish Government



Rubrik also delivers a comprehensive portfolio of government certifications and accreditations across the hybrid cloud. Rubrik offerings are certified for the United States Department of Defense Information Network Approved Products List (DoDIN APL) and EAL2+ of Common Criteria for Information Technology Security Evaluation. Rubrik supports all major government infrastructure offerings, including Google Cloud Platform (GCP), Microsoft Azure GovCloud, AWS GovCloud, and Commercial Cloud Services (C2S).



Cloud Solutions

Whether deployed on-premises or in the public cloud, Rubrik can automate data protection, deliver global file-level search, and quickly spin up test instances for application development. With a breadth of solutions expanding to all major public and private cloud providers, Rubrik delivers all of the solutions and tools necessary to accelerate the cloud journey.

CloudOut

Rubrik's CloudOut automates archival to public and private cloud storage tiers. Users leverage CloudOut to retain data off-premises at lower costs and with higher availability. For public cloud archival, Rubrik delivers support for Amazon Web Services (AWS), Microsoft Azure, Google Cloud, and Oracle Cloud. Users archiving to private cloud can leverage any S3-compatible object stores, NFS, or tape. For users required to manage tape due to legal or regulatory requirements, Rubrik delivers certified integrations with QStar Archive Manager to support a breadth of tape vendors and configurations. Utilizing CloudOut is a simple process:

- 1. Policy-based management:** Users create and assign an SLA policy with specified backup frequency, archive location, and retention period.
- 2. Automated archival:** Once an SLA policy is assigned, Rubrik auto-archives applications to the desired archive location. All data remains instantly accessible for fast recovery with global predictive search.
- 3. Cloud archive consolidation:** Rubrik's Cloud Archive Consolidation automatically consolidates snapshot chains and performs garbage collection, delivering an incremental-forever approach for data-resilient and storage-efficient cloud archival.

Edit Archival Location

Archival Type
Amazon S3

Region
US West (N. California)

Storage Class
Standard

AWS Access Key
AKIABJQQZ4D6GDRFSONQ

AWS Secret Key

AWS Bucket Name
gaia-sand1-archive-bucket

Archival Location Name
AWS S3:gaia-sand1

Encryption Type
 KMS Master Key ID RSA Key

KMS Master Key ID

Cancel Update

Advanced Settings

Cloud Compute Settings
Cloud compute settings will be used to launch a temporary Rubrik instance on cloud.

Virtual Network ID
vpc-0d35d62f38fd64914

Subnet ID
subnet-02d06d07fb6e4e8bf

Security Group ID
sg-0e28917658d322ca6

Enable Archive Consolidation

Cancel Save

For Azure, Rubrik also delivers Azure Instant Tiering and Smart Tiering, which are two modes of archiving to the Azure Archive Access tier. This provides users with the ability to optimize cost for archive data that is rarely accessed. Instant Tiering and Smart Tiering are available in addition to Rubrik's functionality for the Azure Default Access tier:

- **Default Access Tier:** Rubrik archives to Azure Hot or Cool Access tiers, which offer faster recovery times.
- **Instant Tiering:** Archived snapshots are immediately uploaded to Azure's Archive Access tier, which offers the most cost-effective storage for snapshots that are rarely accessed. These snapshots have a minimum retention requirement of 180 days.
- **Smart Tiering:** Archived snapshots are initially uploaded to the default Hot or Cool Access tiers before transitioning to the Archive Access tier to optimize cost. To enable Smart Tiering, users simply specify the minimum number of days or minutes that a snapshot should be kept in the Default Access (Hot or Cool) tier before transitioning to the Archive tier. Users can design SLAs such that only snapshots with long retentions will be Smart Tiered (e.g. monthly snapshots with retention of at least a year), ensuring that more recent snapshots are available with faster RTOs while older snapshots are stored more cost-effectively.

The screenshot shows the configuration for Azure Archival. At the top, the 'Archival' toggle is turned on. Below it, the container name is 'Azure:gaia-emea2-archive-container' and the 'Enable Instant Archive' checkbox is unchecked. Under the 'Azure Tiering' section, three radio buttons are present: 'Default Access Tier Only', 'Archive Access Tier Only', and 'Smart Tiering', with 'Smart Tiering' selected. A 'Duration (days)' field is set to '30'. An information icon is visible next to the duration field. At the bottom, a note states: 'Archival starts after 3 years and is retained on the archival location for 4 years. Snapshots will be transitioned to Azure archive at minimum of 3 days.'

CloudOn

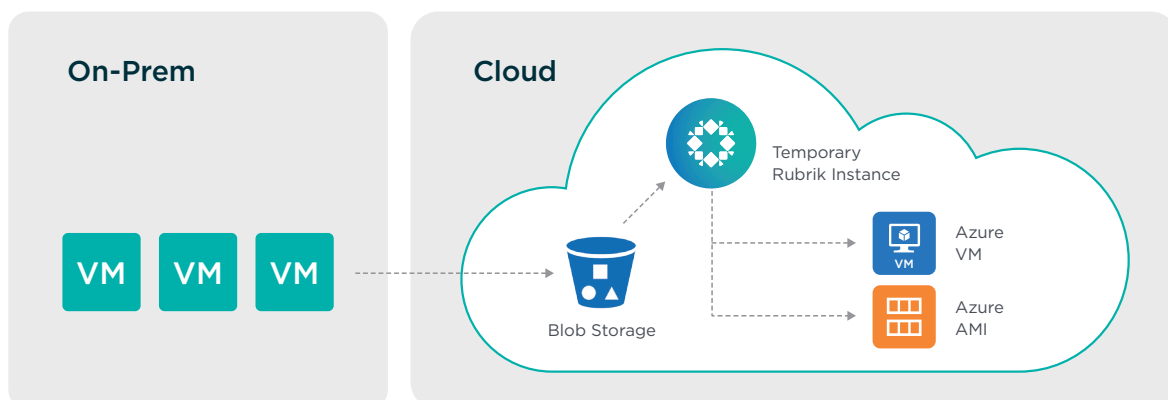
Rubrik's CloudOn automates the conversion of on-premises VMs to cloud instances in AWS and Microsoft Azure for cost-effective disaster recovery, application migration, and accelerated test/dev.

To power on applications in the cloud, users configure the desired security group and AWS virtual private cloud (VPC) or Azure Virtual Network details. Rubrik scans VM configuration files to understand characteristics such as compute, memory, and storage before recommending a compatible cloud instance type. Thereafter, Rubrik converts VM snapshots through the following process:

1. Rubrik spins up temporary compute in the user's cloud environment. This prevents the need for any data to exit the cloud region, saving both bandwidth and egress costs.
2. The temporary Rubrik instance pulls the VMs from the archive location and converts them into cloud-native images.
3. Rubrik automatically powers down the temporary cloud instance after the conversion is complete in order to minimize cloud compute costs. This approach is unlike traditional solutions, which require continuously running compute resources in the cloud. Since the Rubrik compute is launched within the user's cloud environment, all cloud resource consumption is transparent to the user.

Users can leverage two approaches for CloudOn:

- **On-demand conversion:** Users can convert point-in-time VM snapshots into cloud instances on-demand at lower costs. This is commonly used for migrating applications to the cloud or for ad-hoc application test and development.
- **Policy-based conversion:** To deliver faster RTOs, users can set up SLA domains to automatically convert VMs after archival to cloud storage. To minimize cloud storage consumption, users can choose to only retain the most recent converted VM image. In the event of a disaster, users can leverage global search to quickly locate the pre-converted cloud instances and spin them up in minutes.



Cloud Cluster

Users can deploy Rubrik as a software instance in a public cloud provider to orchestrate all critical data management functions, including policy-based backup, recovery, replication, disaster recovery, archival, and search. Backups are automatically stored in elastic cloud storage so that users can efficiently scale protection in-line with cloud service consumption. All cloud-native data is automatically indexed and stored in a single, scale-out repository, delivering instant access and fast recovery with global predictive search. Users simplify management across hybrid cloud with a single Rubrik UI for cloud-native and on-premises applications.

Users can also automate cloud-native protection of large-scale AWS and Azure deployments via Polaris GPS. As enterprises create hundreds of AWS accounts or Azure subscriptions across different regions, each becomes an individual silo to manage. With Polaris GPS, enterprises exercise global control over their cloud instances via global SLA policies. For more information on GPS, please see [Global Hybrid Cloud Management](#).

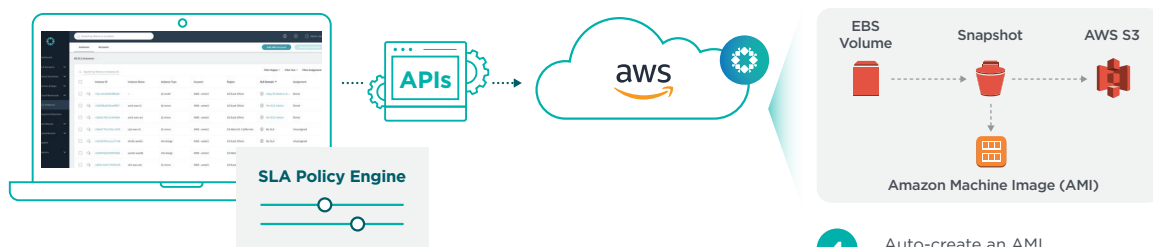
AWS Native Data Protection

Users face several challenges with native AWS tools to protect EC2 instances, such as error-prone, manual scripting and fragmented management. Rubrik integrates with native AWS APIs to deliver policy-driven and cloud-native data protection for AWS EC2. Users can manage applications on-premises, at the edge, and in the cloud with the same easy-to-use interface. Benefits include:

- **Auto-discovery:** Users enter AWS account credentials directly from the Rubrik UI. Rubrik then auto-discovers all EC2 instances in configured AWS accounts, allowing users to specify which accounts and regions to manage.
- **Policy-based management:** Users assign SLA policies to EC2 instances with the same simple process as for other applications. Rubrik automatically spins up a single lightweight Rubrik node in the cloud to index EBS snapshots. Once complete, the Rubrik node is automatically powered down, minimizing consumption of cloud compute resources.
- **Fast recovery:** For each snapshot, Rubrik auto-creates an AMI with associated permissions and billing information. This allows users to perform file-level restores in the event of a corrupt or missing config file without rolling back the entire system. With global predictive search, users can locate and recover files directly from EC2 instances. Users can also choose to export a replica of any EC2 instance to another AWS Availability Zone or Region.

1 Auto-discover AWS EC2 instances. User specifies AWS accounts and regions for management.

3 Spin up a single lightweight Rubrik node in AWS to index snapshots in S3. Once complete, the Rubrik node is automatically powered down.



2 Automate scheduling and retention via one SLA policy engine.

4 Auto-create an AMI with permissions and billing information.

API Integrations



Rubrik's REST APIs enable a breadth of automation capabilities through pre-built SDKs and integrations with automation, self-service, and monitoring tools.

Rubrik Build: All of Rubrik's SDKs, tooling integrations, and API-driven use cases are documented extensively on Rubrik Build: Rubrik's developer community for building custom applications with the Rubrik API. For more information, please visit build.rubrik.com.

Pre-built SDKs: Rubrik delivers pre-built SDKs for Python, Golang, and PowerShell to enable developers to access Rubrik APIs and build custom applications. The languages are easy to use and are supplemented with extensive documentation on Rubrik's public GitHub repositories.

Automation tools: Rubrik integrates with a variety of third-party automation and configuration management tools (such as Puppet, Chef, SaltStack, Ansible, and more) to simplify deployments across hundreds of servers or VMs. Users automate compliance testing, SLA assignments, and lifecycle management at speed and scale in minimal lines of code.

Self-service portals: Rubrik's pre-built integrations with popular service catalogs simplify orchestration of data management tasks through self-service access:

- **ServiceNow and vRealize Automation:** Rubrik's API integrates with ServiceNow and vRealize Automation to deliver self-service access for data protection, instant recovery, application test/dev, and customized analytics. Users eliminate lengthy wait times at the help desk by accessing Rubrik's integration directly from their service catalogs. For more information on Rubrik's support for ServiceNow, please see [Self-Service Data Management Through ServiceNow](#).
- **VMware vCloud Director:** Enterprises and managed service providers running on vCloud Director (vCD) integrate with Rubrik to deliver backup-as-a-service in multi-tenant environments. Tenants gain self-service access for SLA policy assignments, instant recoveries, application test/dev, and customized reporting. Users can access Rubrik's integration directly from the vCD UI.

Monitoring tools: Rubrik integrates with monitoring tools like Splunk and Nagios to provide information on compliance, infrastructure health, and application consumption. Users can stream metrics from Rubrik Cloud Data Management and Polaris into Splunk, providing centralized consumption of dashboards, alerts, and analytics in an easy-to-use web interface.

Comprehensive Platform Support

Rubrik is designed to be vendor-agnostic from day one, delivering broad support for both traditional and modern enterprise applications.

Virtual



Physical



NoSQL



Cloud + SaaS



Virtual Environments

VMware vSphere

Rubrik automatically discovers vSphere VMs by connecting to VMware vCenter through the native vSphere API. Once the VMs are discovered, users can apply Rubrik's SLA policy engine to automate protection at the VM, folder, cluster, or host-level. Users can also apply policies to vSphere tags to organize protection by applications and VM groupings. Rubrik mitigates VM stun effects by delivering rapid parallel ingest and integrating with VMware's vStorage API for Data Protection (vADP). This delivers faster incremental-forever backups and minimizes impact to production environments. For VMs running on Pure Storage FlashArray, Rubrik delivers additional integrations with the native FlashArray API. This integration significantly reduces backup windows by immediately offloading VM snapshots to FlashArray storage.

For mission-critical applications running on VMware vSphere, Rubrik also offers Continuous Data Protection (CDP) for near-zero RPOs. For more information on CDP and how it works, please see [Continuous Data Protection](#).

For users managing vSphere through native VMware automation and self-service tools, Rubrik delivers extensive API integrations with services such as vRealize Automation and vCloud Director. For more information, please see [API Integrations](#).

Microsoft Hyper-V

Rubrik can discover Hyper-V environments by connecting directly to Microsoft System Center Virtual Machine Manager (SCVMM) or to physical hosts running Hyper-V. Users can automate data protection at the VM, folder, cluster, or host-level.

For both VMware vSphere and Hyper-V VMs, Rubrik delivers Live Mount for instant recoveries and clones, as well as CloudOn for disaster recovery or test/dev.

Nutanix AHV

Rubrik protects Nutanix AHV VMs through native integration with Nutanix APIs. To connect Rubrik to a Nutanix cluster, users simply enter the connection information and administration credentials in Rubrik. Once complete, Rubrik automatically discovers the VMs so that users can begin assigning SLA policies immediately. Rubrik's integration leverages Nutanix's native REST APIs to deliver incremental-forever snapshots at the VM, folder, cluster, or host-level.

Physical Operating Systems

Rubrik delivers comprehensive support for physical servers running Windows, Linux, and Unix (IBM AIX and Oracle Solaris) operating systems. Rubrik connects to physical servers through Rubrik Backup Service, allowing users to simply specify the files and directories to be protected. SLA policies can be assigned with granularity down to the file or folder level, where users can specify inclusion or exclusion for each object. With incremental-forever backups, users can drive network and storage savings both on-premises and in the cloud.

For physical Windows servers, users can also perform volume-based backup and recovery of entire drives. Rubrik Backup Service invokes calls to Microsoft's Volume Shadow Copy Service (VSS) to capture crash-consistent snapshots of all drives, filesets, and relevant metadata. When recovering, users can choose to restore to compatible bare metal Windows Servers or virtualize to a hypervisor or cloud of their choice.

Physical Databases

Elastic App Service

For Oracle, SQL Server, SAP HANA, MySQL, PostgreSQL, MongoDB, and Cassandra databases, Rubrik delivers Elastic App Service (EAS): intelligent application-aware data reduction for DBAs utilizing database-native backup tools. EAS delivers high-performance data reduction while minimizing use of compute resources, leading to better network and storage efficiency on-premises and in the cloud.

With EAS, DBAs can use their preferred management tools while benefiting from Rubrik's automated data lifecycle management and immutability from ransomware. This allows DBAs to maintain control over operational recovery objectives while giving backup administrators control over long-term retention and compliance objectives. To enable Elastic App Service, users simply add a new managed volume, specify the application type, and assign an SLA domain.

Microsoft SQL Server

Rubrik auto-discovers all SQL Server hosts and databases through Rubrik Backup Service, enabling users to assign granular SLA policies at the host, instance, or database-level. Rubrik delivers incremental-forever and application-consistent snapshots. Crash-consistency is delivered through automated calls to Microsoft's Volume Shadow Copy Service (VSS), which briefly quiesces I/O on the target database to ensure that snapshots are in a consistent state.

In addition to protecting SQL Server databases, users can choose to apply SLA policies to the database transaction logs. Protecting transaction logs enables users to recover databases with extremely granular RPOs. When delivering point-in-time recoveries for a particular database, Rubrik automatically applies the protected transaction logs to the nearest full snapshot, rolling the database's state forward to the date and time requested.

For SQL Server databases that are replicated through Microsoft's AlwaysOn Availability Groups, Rubrik automatically detects the parent database and only applies SLA policies to the parent copy. This delivers greater network and storage efficiencies while accommodating native tools.

Oracle Database

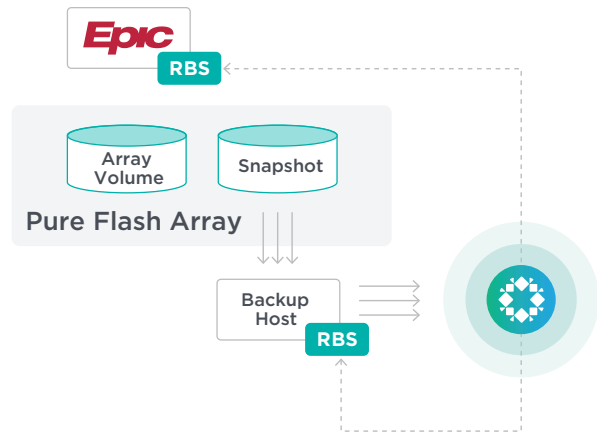
Rubrik drastically improves data management for Oracle databases with instant recoveries and clones, automated lifecycle management, and streamlined collaboration between DBAs and Backup Admins. Rubrik connects directly to Oracle database hosts or RAC nodes via Rubrik Backup Service, allowing automatic discovery of all Oracle clusters, hosts, databases, and tablespaces. Once the Oracle Database environment has been discovered, users can apply SLA policies at the host, instance, or database level directly from the Rubrik UI or API. Rubrik also drives capacity and network savings through Oracle RMAN's Incremental Merge, which provides incremental-forever backups and high-performance data reduction. All metadata is indexed for granular search and recovery at the database or tablespace level.

For DBAs that want to continue using Oracle RMAN scripts while also benefiting from Rubrik's app-aware data reduction, immutability from ransomware, and Live Mount capabilities, Rubrik offers [Elastic App Service](#).

Epic EHR

For Epic Caché databases running on Pure Storage FlashArray, Rubrik automates daily management, minimizes backup windows, and delivers native encryption and immutability to defend against ransomware. For SQL or Oracle databases running in Epic Clarity environments, Rubrik offers additional support through Live Mount for instant recovery and test/dev.

To protect Epic Caché with Rubrik, users simply deploy Rubrik Backup Service on the corresponding FlashArray production hosts. Rubrik automatically discovers all LUNs so that users can begin assigning SLA policies immediately. Through Rubrik Backup Service, Rubrik integrates with native FlashArray APIs to deliver volume-based backups of the Caché databases. All snapshots are application-consistent and incremental-forever, enabling users to deliver seamless recoveries and drive capacity and network savings.



As with all other applications, Rubrik protects Epic EHR with encryption both in-flight and at-rest, regardless of where the data is stored. With encryption and immutability built into the software fabric, potential intruders can neither see nor alter backups of sensitive patient data. In the event of a potential ransomware attack, administrators can quickly restore data back to the pre-infected state.

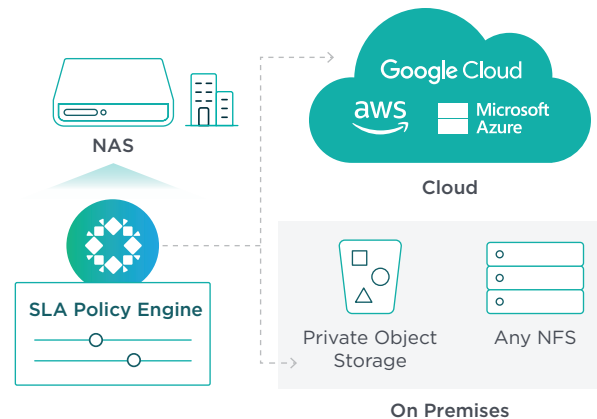
For Microsoft SQL Server, Oracle, and Epic Clarity databases, Rubrik also delivers Live Mount for near-zero RTOs and self-service access. For more information, please see [Live Mount](#).

SAP HANA

Rubrik integrates with SAP HANA through a Backint API integration, enabling users to backup entire databases to Rubrik directly from SAP HANA Studio or Cockpit. The first layer of data protection is provided by the native SAP HANA backup scheduling engine, which delivers full and incremental backups directly to Rubrik. Thereafter, Rubrik's SLA policy engine automates replication and archival across on-premises and cloud. Users can leverage the elasticity of the cloud for long-term retention or archive elsewhere on-premises. Rubrik also protects the database log files in a separate managed volume from the database backups, enabling users to leverage Studio and Cockpit for point-in-time recoveries as well.

Network-Attached Storage (NAS)

Rubrik provides vendor-agnostic support for all NAS vendors that speak NFS or SMB protocols. For NetApp, Isilon, and Pure Storage FlashBlade, Rubrik integrates with native snapshot APIs in Isilon OneFS, NetApp ONTAP, and Purity//FB to enhance snapshot consistency and accelerate file-scanning. With Rubrik's vendor-native API integrations, users can take consistent, point-in-time snapshots of NAS volumes with up to 10x faster file system scans than for generic NAS solutions without integrations.



Users can mount NAS shares directly on Rubrik without proxies. New shares are added directly from the Rubrik UI by specifying the host and NFS or SMB path. Once a NAS share has been registered, users can assign SLA policies with granularity down to the fileset level.

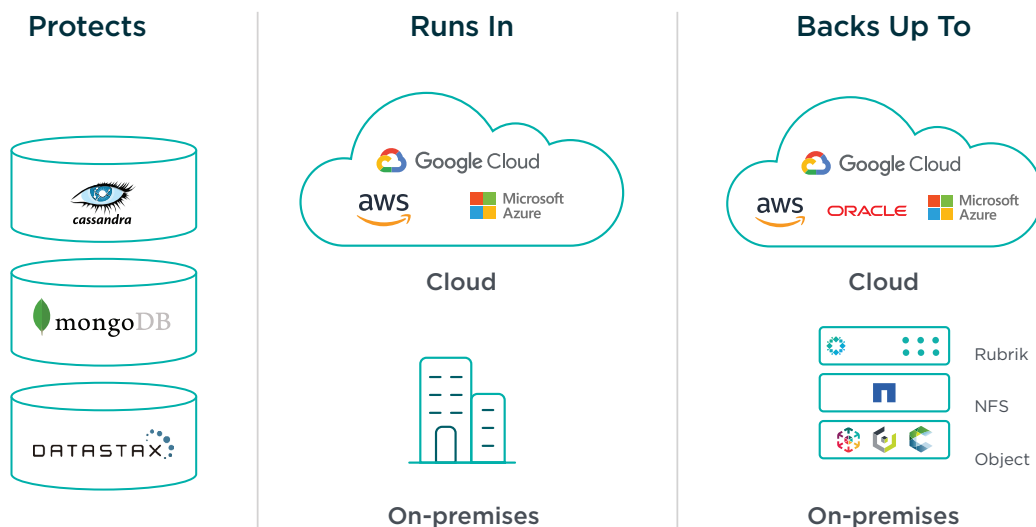
Rubrik provides NAS Direct Archive to enable enterprises to efficiently manage massive unstructured datasets on-premises or in the cloud. Rubrik captures and indexes all metadata before instantly archiving NAS data to the desired third-party NFS or object storage location. This enables users to maintain instant access with file-level search and recovery. To enable NAS Direct Archive, users simply select a NAS share or fileset, confirm that they would like to apply Direct Archive, and apply any policy with a specified archive location.

NoSQL Databases

Rubrik Cloud Data Management works with Rubrik Mosaic to protect and manage NoSQL databases such as MongoDB and Apache Cassandra (DataStax), allowing DevOps and database administrators to accelerate innovation for customer-facing applications without worrying about data loss. Users can backup NoSQL databases protected by Rubrik Mosaic directly to a local Rubrik Cloud Data Management environment, or to any NFS, object storage or public cloud. With Rubrik Mosaic, organizations can benefit from the following key features:

- **Application-consistent backups:** Application-consistency enables repair-free recoveries and dissimilar topology restores, enabling rapid-RTOs for customer-facing applications.
- **Incremental-forever backups:** Rubrik Mosaic automatically tracks changes in NoSQL databases to deliver incremental-forever snapshots. This approach minimizes data transfer and impact on production both on-premises and in the cloud. Together with database compaction operations, this delivers greater space efficiency.

- **Industry-first semantic deduplication:** The software platform provides a deep understanding of the data to comprehend exactly what is stored and how it is updated. Because it is data aware, Mosaic condenses multiple NoSQL replicas into a single backup to deliver a +70% reduction in secondary storage requirements.
- **Streaming recovery:** Users don't have to wait until a recovery is complete in order to access NoSQL data. With streaming recoveries, data is made available in real-time as each recovery is in process for the entire data set.
- **Advanced recovery operations:** In addition to recovering at the database or object-level (database or table), users can restore data that meets a specific query, time range, or column family. Furthermore, Rubrik Mosaic delivers advanced data automation features such as data masking, under-sampling and others for test/dev and test data management use cases.
- **Topology Independent restores:** Restores can also be across dissimilar topologies to simplify application development in large deployments. Users simply query the relevant data for test/dev and restore to their desired test environment, regardless of the number of nodes in the production/source database environments or non-prod/test database environments.



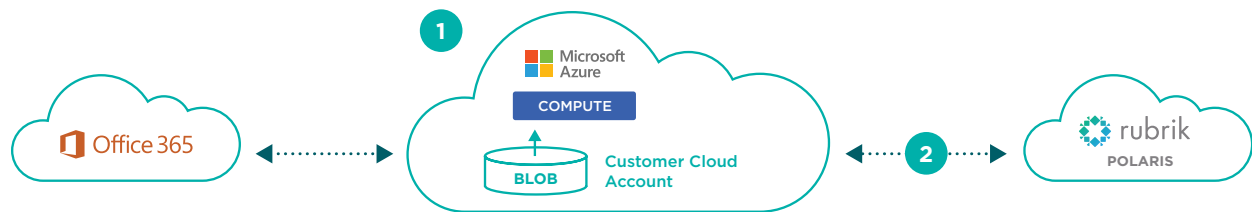
SaaS Applications

IT Leaders are increasingly making the move to Software-as-a-Service (SaaS) to leverage the agility of the cloud, reduce capital expenditures, and increase employee productivity. With SaaS-based offerings, IT organizations can eliminate routine, manual tasks and focus on delivering IT-as-a-Service for end users.

Office 365

Rubrik delivers cloud-native data protection for Office 365 through the Polaris SaaS platform. Rubrik eliminates costly, manual infrastructure adjustments by spinning up compute resources on-demand that dynamically scale in-line with Office 365 consumption. Additionally, Rubrik ensures enterprises can maintain control of where data is stored. All data remains in enterprises' Azure subscription. Benefits include:

- 1. Set up in minutes:** Rubrik leverages OAuth to authenticate users' Azure and Office 365 subscriptions. Users simply enter their login credentials, allowing Rubrik to automatically discover all users and organizations in the Office 365 subscriptions.
- 2. Automate management via SLA policies:** Users can assign SLA policies at the subscription or user-level for granular control.
- 3. Search and recover quickly:** To restore email data, administrators can leverage the Polaris UI to drill down into a single snapshot or search across all snapshots by date range, keyword, sender, or recipient. Users can choose to restore to the original location or to export to another user's mailbox.



1 All data remains in your Azure subscription, including emails, calendar items, and metadata, so you maintain control of where data is stored.

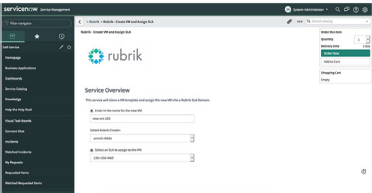
2 Polaris orchestrates Azure resources to complete data management tasks – backup, search, recovery – as needed.

Business Use Cases

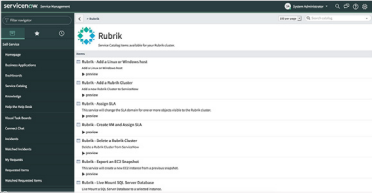
Self-Service Data Management via ServiceNow

Enterprises are constantly looking for ways to increase automation. By automating mundane daily tasks, administrators can save management time and focus on more strategic projects. But often, backup and recovery sits isolated from centralized IT automation suites, such as ServiceNow, resulting in hours to days at the IT help desk to just recover a file or assign an SLA.

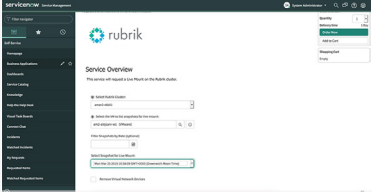
With Rubrik's API integration with ServiceNow, Rubrik allows enterprises to eliminate help desk tickets and provide self-service backups and restores to IT administrators – all through the ServiceNow catalog. This reduces time to resolution from days to minutes, frees up the help desk for more critical requests, and helps enterprises recover quickly from unplanned disruptions. Capabilities include:



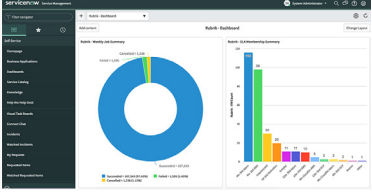
Automated data protection: Users can assign or modify SLAs directly from the ServiceNow UI. Rubrik automates protection for VMs as soon as they are provisioned without additional steps or help desk tickets are required.



Self-service file recovery: Rubrik's API integration delivers predictive file-level search so that users can locate and recover files on a self-service basis. This capability reduces recovery times from days to minutes.



Live Mount for test/dev: With Live Mount, users can instantly provision clones for test/dev without rehydrating production environments. Developers can instantiate VM and database clones via self-service, accelerating development and eliminating wait times at the help desk.



Customized reporting and analytics: Users can deliver rich data visualizations from within ServiceNow's reporting dashboard. Rubrik's API integration enables users to gain insights on data management, compliance, and capacity utilization across hybrid and multi-cloud environments.

Customer Story



Before Rubrik

Secure-24 was challenged by complex, lengthy recoveries with their legacy solution. Simply recovering a file required multiple teams to be involved with different tools and procedures for their virtual and physical environments. It could often take hours to recover a workload.

After Rubrik

With Rubrik's ServiceNow integration, Secure-24 consolidated multiple tools and services into a single software platform. By enabling self-service file-level recovery and creation of SLA policies within the ServiceNow catalog, Rubrik significantly reduces management complexity and slashes recovery times to minutes. They can also now track platform analytics and compliance for more efficient reporting,

Results include:



**Management time savings
with self-service**



**Reduced RTOs from
hours to minutes**



**Enhanced reporting
capabilities to
ensure SLA compliance**

Read more about how Secure-24 uses Rubrik [here](#).

Live Mount for Test/Dev

The ability for developers to quickly access data is essential for accelerating application development. Legacy solutions deliver lengthy RTOs for VMs and databases, complicating daily tasks such as recovering data, testing updates, or building new tools.

Rubrik's Live Mount delivers near-zero RTOs to accelerate test/dev for VMs and databases. Administrators can grant granular permissions for self-service access, enabling users to leverage a single "golden image" snapshot to mount multiple clones directly on Rubrik. Users can even access Live Mount through preferred IT service catalogs, such as VMware vRealize Automation and ServiceNow via Rubrik's APIs. For databases, use cases include:

- 1. Ad hoc queries and restores:** Users can quickly perform restores of specific rows or tables via a simple query, export, and import. By mounting multiple recovery points, users can easily examine a database to track when specific data changed without provisioning extra disk space.
- 2. Health checks:** Users can quickly spin up database clones to validate backups without impacting production. With Rubrik's APIs, users can create a script to bring a Live Mount online, perform a database validation, and create an alert upon completion. Users can also perform Oracle recovery fire-drill without needing a full restore of data files.
- 3. Test/dev workloads:** Users can spin up database clones in real-time without storage penalty. If Rubrik Cloud Cluster is in place, users can instantiate a cloud instance of SQL Server for many test/dev scenarios, such as data loads or schema test changing.

For more information on how Live Mount works, please see [Live Mount](#).

Customer Story



Before Rubrik

Rubrik helped Compeer Financial consolidate and simplify three disparate IT systems following a merger with Agstar financial services. Prior to Rubrik, Compeer struggled with poor RPO and RTOs, especially for SQL databases. Additionally, SQL database administrators had to send requests to the backup team in order to find and restore databases, which was time-consuming and inefficient.

After Rubrik

With Rubrik's Live Mount, Compeer's SQL database administrators can instantly restore any of its approximately 500 SQL databases to quickly access records, perform backup health checks, or use for testing and development. Additionally, IT organizations can grant self-service access for SQL database administrators through role-based access control.

Results include:



Near-zero RTOs for SQL



Reduced daily management time from hours to minutes



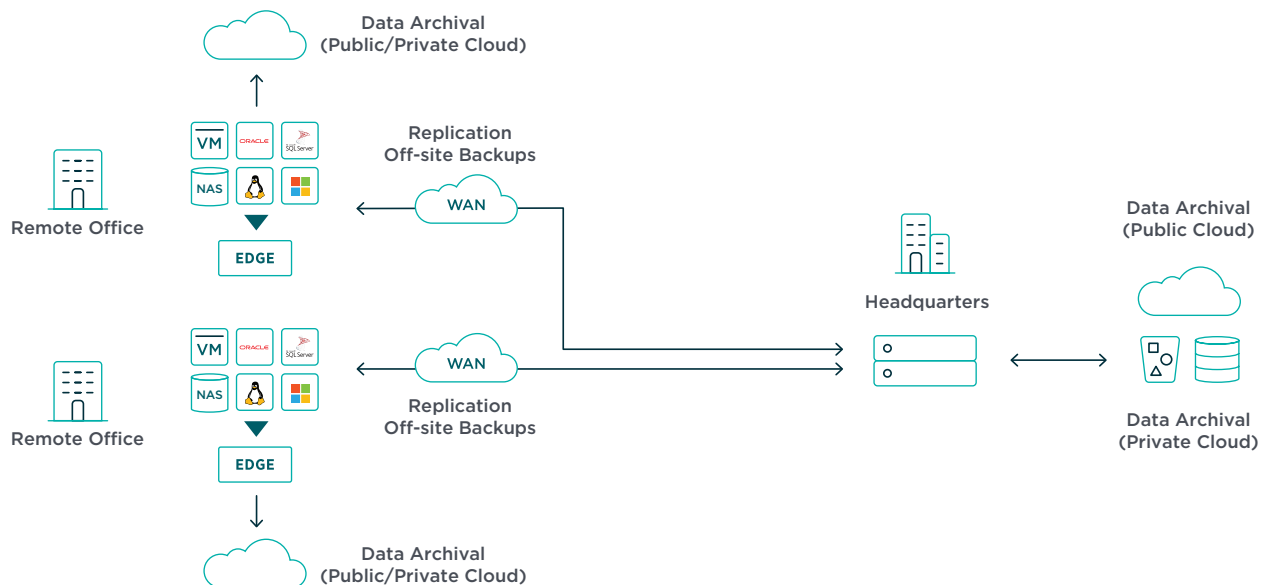
Self-service access for SQL DBAs

Read more about how Compeer Financial uses Rubrik [here](#).

Remote and Branch Office Data Protection

As enterprises expand to new sites, administrators are looking for a single solution to protect all their applications and data at their primary and remote data centers while minimizing costs and complexity. Rubrik Edge is a full-featured software appliance that extends all of Rubrik's data management capabilities to virtual and physical environments in remote and branch offices. Users can bring their own hardware, with all hardware requirements determined by capabilities needed from the software. With a single interface, users can deliver policy-based management, global search, fast recovery, and compliance reporting on-premises, at the edge, and in the cloud.

With Edge, enterprise can automate data protection locally, replicate data to a central data center, and archive to public or private cloud directly from either location. Additionally, Edge minimizes data transfer and storage costs by deduplicating data before replication occurs.



Automated Sensitive Data Discovery and Classification

In today's cloud-y world, sensitive data resides in more places and can be accessed by more people across dynamic, distributed architectures, intensifying risk of data breaches and non-compliance penalties. Therefore, businesses need consistent processes for continuous data governance. However, businesses spend too much time and money on complex classification and compliance tasks consisting of manual tagging, native auditing, and annual cleanups. To efficiently safeguard sensitive corporate and customer data adequately, automation is key.

Polaris Sonar automates the discovery, classification, and protection of sensitive data for efficient data governance. Sonar applies machine learning to automate sensitive data classification, allowing enterprises to easily identify and identify data exposure without impacting production environments. By leveraging their existing Rubrik CDM and Polaris deployments, users get up and running in just a few minutes with zero additional infrastructure. For management simplicity, we provide pre-built data classification policy templates to identify common data types from standards and regulations such as GDPR, PCI/DSS, SOX, HIPAA, GLBA, and CCPA. Users can also define their own custom dictionaries, expressions, and policies. Use cases include:

- **Investigative Data Discovery:** Sonar uses machine learning to automatically scan your entire environment and identify what kinds of sensitive data exists and where it is located.
- **Assist with Access Requests:** Searching for personal data with Sonar does not, by itself, allow a customer to comply with access requests.
- **Simplifying Compliance:** Easily gain insight into types of data processed and where sensitive data may be stored when managed with Rubrik CDM.

Customer Story



Before Rubrik

As a state government, City of Sioux Falls must protect citizens' personal information and minimize serious risk of data breaches. In the past, data classification tasks required dedicated teams with multiple full-time engineers. To search for keywords, each query had to be manual scripted one at a time across each location and categorized via spreadsheets, taking 1-2 weeks to complete.

After Rubrik

With Polaris Sonar's on-demand search and ML-driven data classification, they could search for hundreds of keywords in just 1 hour, freeing up employees for higher-value add work. They were also able to identify and lock down servers to minimize sensitive data exposure. They can now continuously monitor sensitive data for high risk incidents on existing backup data without impacting production.

Results include:



90%+ operational savings



FTE focused on higher-value work

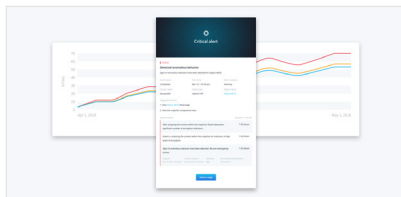


Reduced risk of data breach and negative PR

Multi-Layered Ransomware Defense

In the event of a ransomware attack, it's often difficult to recover fast. Common challenges include days to weeks to discover infection, manually identifying the scope of impact, and hours to restore to a clean state. Not to mention, ransomware attacks often are detrimental to the business with significant revenue, information, and productivity loss.

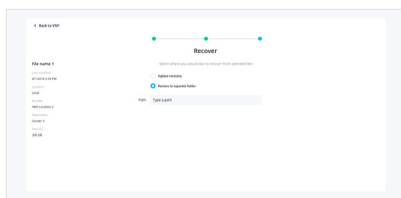
Rubrik complements enterprises' defense-in-depth approach to security with faster ransomware recovery. All data is stored in an immutable format, preventing ransomware from encrypting or overwriting backups. In the event of an attack, users can simply restore applications to the most recent clean snapshot and resume operations quickly with minimal business downtime. Radar, delivered on the Polaris SaaS platform, makes it even faster and easier to recover from a ransomware attack with a multi-layered defense:



AI-driven detection: Radar applies machine learning algorithms against application metadata to alert anomalous activity.

A screenshot of the 'Select Data' interface. It shows a table with columns for 'Application', 'File Path', 'Size', and 'Last Modified'. The table lists various applications and their associated files, with some rows highlighted in red to indicate impacted data. There are search and filter options at the top of the table.

File-level impact analysis: Users can quickly investigate the attack surface with simple visualizations on what applications and files were impacted and where they are located.



One-click recovery: Users can simply select all impacted applications and files, specify the desired location, and restore to the most recent clean versions with just a click. Rubrik automates the rest of the restore process.

Customer Story



Before Rubrik

ASL's previous solution was not built for a strong defense against ransomware. The cargo airline industry is a common target for ransomware, experiencing a minimum of one attack per month. In the past, ASL would have to resort to manual scripting to identify and erase infected files, which was painful and killed productivity for days. On top of that, they are required to maintain 99.9% availability. If ASL's IT system is down for more than 15 minutes, airplanes cannot take off, customers cannot receive their cargo, and the airline is at risk of being hit with massive fines.

After Rubrik

ASL Airlines leverages Radar's AI-driven detection to reduce the time spent monitoring for ransomware from two hours per day to just a few minutes. In the event of an attack, they can replace manual scripting with one-click recovery to minimize downtime. With Radar, ASL Airlines was also approved for a cyber insurance contract, which was not possible before.

Results include:



25% IT admin time savings monitoring for cyber threats



15 to 100+ hours of IT admin time saved in case of an attack



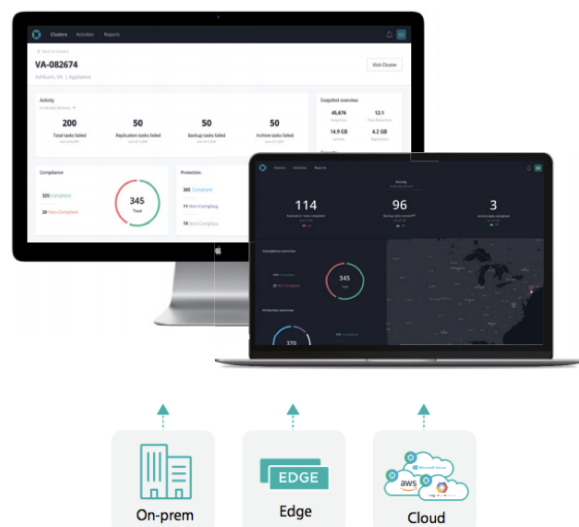
Millions of euros in potential savings in case of an attack

Read more about how ASL uses Rubrik [here](#).

Centralize Hybrid Cloud Management with Polaris GPS

As enterprises adopt cloud initiatives, it is increasingly difficult for administrators to monitor and manage all their applications and data across multiple data centers and clouds. As a result, they spend too much time monitoring backups, troubleshooting errors, and analyzing siloed information, resulting in an inability to scale. Polaris GPS provides a single SaaS-based management console for global management on-premises, at the edge, and in the cloud. This allows businesses to streamline operational costs and free up administrative time for projects more critical to the business. Key capabilities include:

- **Global inventory:** GPS aggregates the discovered environments of each Rubrik Cloud Data Management (CDM) cluster. This provides users with a single global inventory of objects across the enterprise. With Global Object Search, users can search by object name across clusters to easily locate data regardless of location, providing full visibility and control over all applications and data.
- **Global monitoring and reporting:** GPS delivers a single easy-to-use dashboard to view global metrics on SLA compliance, infrastructure health, and performance. Metrics can be dynamically filtered for ad-hoc analyses across a variety of dimensions, like time, location, status, and application. Any chart or table can be converted into a point-in-time analysis or time-series with just a single click.
- **Global SLA policy management:** GPS enables users to create and assign global SLA policies across multiple CDM clusters. Users can effectively centralize management across the entire enterprise with a single pane of glass across on-prem, edge, and cloud applications – including workloads running in AWS and Azure.



Conclusion

As enterprises face explosive data growth and embrace cloud, data management is increasingly critical for driving operational efficiencies and innovation. To meet today's demands and prepare for the future, IT organizations need management simplicity across all their workloads, automation of routine daily tasks, and easy scalability across data center to cloud. Rubrik delivers a radically new approach to data management, providing the foundation necessary to simplify backup and recovery, reclaim valuable time for business-critical projects, and transform IT organizations into business enablers.

Ready to learn more about Rubrik Cloud Data Management? To see it in action, speak to one of our sales representatives at rubrik.com/contact-sales.



Global HQ

1001 Page Mill Rd., Building 2
Palo Alto, CA 94304
United States

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

Rubrik, the Multi-Cloud Data Control™ Company, enables enterprises to maximize value from data that is increasingly fragmented across data centers and clouds. Rubrik delivers a single, policy-driven platform for data recovery, governance, compliance, and cloud mobility. For more information, visit www.rubrik.com and follow [@rubrikinc](https://twitter.com/rubrikinc) on Twitter. © 2019 Rubrik. Rubrik is a registered trademark of Rubrik, Inc. Other marks may be trademarks of their respective owners.

20191204_v1