REFERENCE ARCHITECTURE

Rubrik and VMware vSphere

TABLE OF CONTENTS

AN INTRODUCTION TO RUBRIK CDM

Rubrik CDM platform unifies backup, instant recovery, replication, global indexed search, archival, compliance, and copy data management into a single scale-out fabric across the data center and public cloud. Rubrik is used by enterprise organizations to securely manage all data, physical or virtual, across all locations – on-premises, edge of the data center, and cloud.

Instant search delivers near-zero RTOs with predictive search. Easily locate VMs, databases, applications, or files regardless of whether they reside in the cloud or on premises.

Policy-driven engine and programmatic interface eliminates daily operational management by automating how data services are created, consumed, and retired from across the data center and cloud.

Orchestration is the core of Rubrik, providing a suite of APIs that can be used to orchestrate data from data center to cloud. Rubrik provides the freedom to provision data management services with configuration management tools and via custom portals.

Data is **secure** in transit and at rest throughout the entire lifecycle, regardless of location. Granular role-based access can be leveraged while automating compliance reporting in order to successfully meet and complete various industry audits.

Analytics and reporting are provided by Rubrik Envision, which unlocks actionable insight across all environments with customizable reports. Leverage platform analytics that detail operational efficiency, compliance, and capacity utilization across your infrastructure.

Centralized management for a global, distributed Rubrik environment delivered by Rubrik. Designed for a seamless user experience, Rubrik Zero Trust Data Protection provides a comprehensive view of your physical, virtual, and cloud topologies while making management tasks elegantly simple and intuitive.

For more information, watch the on-demand Product Demo or visit the Rubrik website.

AN INTRODUCTION TO VMWARE VSPHERE

VMware vSphere, the industry leading virtualization and cloud platform, is the efficient and secure platform for your hybrid clouds. It provides a powerful, flexible, and secure foundation for business agility that accelerates your digital transformation to hybrid cloud and success in the digital economy. vSphere supports both existing and next-gen workloads through simple and efficient management at scale, comprehensive built-in security that starts at the core, a universal application platform, and a seamless hybrid cloud experience. Applications can be run, managed, connected, and secured in a common operating environment, across a hybrid cloud.

VMware vSphere uses virtualization to transform individual data centers into aggregated computing infrastructures that include CPU, storage, and networking resources. VMware vSphere manages these infrastructures as a unified operating environment and provides you with the tools to administer the data centers that participate in that environment.

The VMware vSphere stack comprises virtualization, management, and interface layers. The two core components of vSphere are ESXi and vCenter Server. ESXi is the virtualization platform where you create and run virtual machines and virtual appliances. vCenter Server is the service through which you manage multiple hosts connected in a network and pool host resources.

AUDIENCE

This reference architecture is intended to provide CTOs, solutions architects, and administrators with information about the architecture, implementation, and benefits of an integrated Rubrik and VMware vSphere solution.

For the remainder of this document, "virtual machines" will be referred to as "VMs" and "disaster recovery" as "DR."

SOLUTION OVERVIEW

Consolidation of workloads onto a plethora of disparate hardware devices was one of the most tangible benefits gained by mainstream, x86 virtualization. Instead of requiring weeks, months, or even years to provision a workload, only a few clicks are required. Many pre-virtualization servers were fragile in construction due to hosting a variety of applications that conflicted with one another. Those applications are now freed to live within their own operating system environment while safely partitioned by the VM construct. At the same time, overall compute utilization rose thanks to the introduction of transparent memory sharing, hyper-threading, smart CPU scheduling, and the holy grail of live migration technologies: VMware vMotion.

One unintended side effect of this phenomenon is workload sprawl. The sprawl only increased the rate of change witnessed by data centers and other compute hosting facilities. One of the primary yet often ignored challenges created by sprawl is ensuring that all workloads are properly identified, protected, and available for quick recovery. Workloads are now being provisioned both quickly and, in many cases, outside of the scope of the Operations group. This makes ensuring that proper guarantees around data protection are met within the business-demanded Recovery Point Objective (RPO) and Recovery Time Objective (RTO) is frustrating at best and disastrous at worst.

The combination of Rubrik and VMware vSphere is a simple and elegant method to capitalize on the extensible and API-focused nature of a virtualized data center to provide policy-driven data protection, test and development workflows, orchestrated data migrations, and much more. This is accomplished by leveraging VMware's vSphere APIs for Data Protection for hypervisor integration, providing Live Mount (Clone) and Instant Recovery (DR) features directly into a vSphere environment, and handing over control to the Operations team by way of declarative policies and SLA Domains.

Let's continue to look at how the solutions positively reinforce one another across these dimensions.

VMWARE VSPHERE STORAGE APIS — DATA PROTECTION

As a technical architecture unfolds, it's important to expand ever deeper across the conceptual, logical, and physical design elements. This ensures that all of the requirements, constraints, and risks are addressed or assuaged before moving towards implementation and operation. It can be a challenge, however, to ensure that integration across an enterprise suite of physical solutions are properly pondered, especially in terms of storage arrays used to fuel datastores for VMs.

In the case of Rubrik and VMware vSphere, leveraging the vSphere APIs for Data Protection ensures that architectural decisions around storage are largely made moot. Regardless of the underlying storage decisions made - such as Fibre Channel SANs, NAS file systems, VMware vSAN (See: Joint VMware and Rubrik White Paper), Storage Virtual Appliances (SVAs), or Hyperconverged Storage endpoints - the usage of an API layer abstracts the nuances present with the selected datastore storage provider.

When data is requested by the Rubrik cluster, the vSphere API layer is used to negotiate with the ESXi hypervisor that is hosting the protected workload to retrieve the needed data from the datastore and transmit it back to the Rubrik cluster. Additionally, the individual node choice within the Rubrik cluster is handled in an automated fashion where the CDM software determines the primary candidate(s) for establishing a session and receiving data.

At a high level, the successful request to begin a backup of a VM kicks off a workflow requesting that the ESXi host running that particular workload to initiate a relatively "no frills" VMware Snapshot in which the quiesce and memory dump options are not used. This snapshot action is intended to redirect storage IO to a snapshot disk, thus freeing the hypervisor's lock on the underlying base disk(s). Subsequent actions are focused on retrieving the data to be parsed and stored by the requesting Rubrik cluster. Once finished, the VMware Snapshot is consolidated (removed), and storage IO returns to the base disk(s). While there are other decision points available in the workflow, such as pre- and post-scripts, the basic flow has been represented.

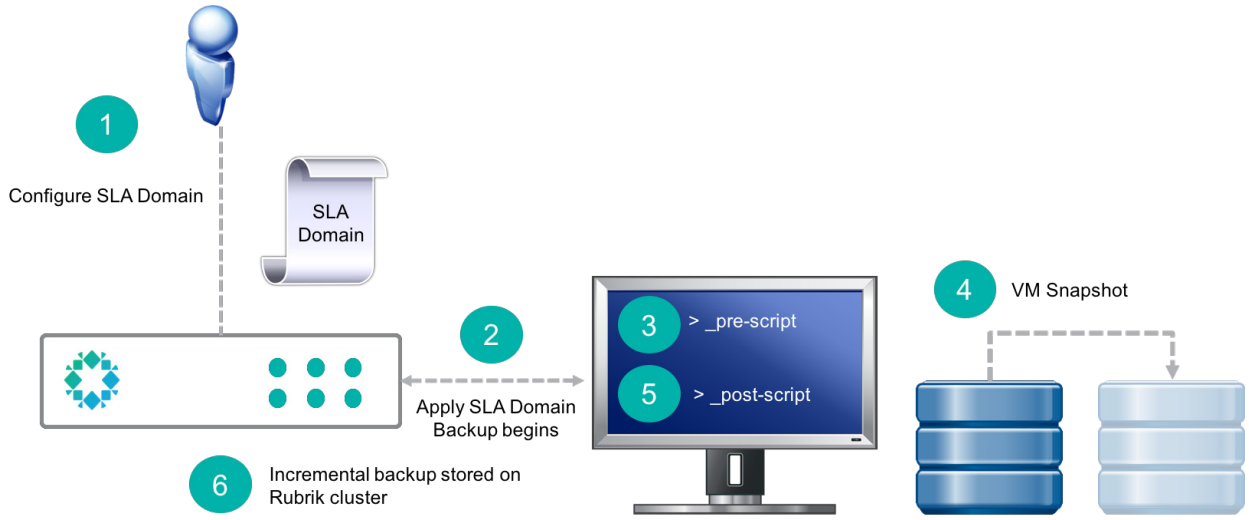

It is worth noting that the process of negotiation, session instantiation, and data transmission is secured using an SSL encrypted Network Block Device protocol and SSL encryption (NBDSSL) transport mode. For applications that require a hardware assisted backup - such as physical workloads with highly-transactional operations - there are options within Rubrik CDM to collaborate on the data retrieval process using hardware offloads and storage snapshots. However, this should be considered out of scope for all but the most exotic VMware vSphere configurations based on customer data on deployed architecture.

RECOVERY METHODS

A Rubrik cluster provides a variety of methods to recover VMs and restore protected data by using snapshots, replicas, and archived snapshots.

When snapshot data exists in a local snapshot and in an archived snapshot, the Rubrik cluster always uses the local snapshot to recover a VM or to restore data. By using the local snapshot, the Rubrik cluster reduces network impact and eliminates any archival data recovery charges associated with a recovery operation or a restore operation.

INSTANT RECOVERY

Rubrik's Instant Recovery can be used to recover VMs that are no longer functioning properly because of:

- Corruption or malware
- Accidental deletion
- Any other service disruption

This functionality allows VMs needing to be restored to be mounted directly off the Rubrik system, thus reducing the recovery time.

Let's visualize the Instant Recovery workflow:

The process first begins by the selecting the VM, snapshot date, and recovery host. You may select to remove a virtual network device if any networking changes or issues would prevent the VM from successfully powering on. This methodology also enables validation of certain services after recovery but before restoring the service.

> Additionally, you may select to preserve the VM-managed object ID (MoRef). This is a managed object ID, which is applicable to vSphere VMs. It will ensure that the VM is recovered using the same MoRef as a part of VM linking, rather than it being recovered as a new object. This method can be important for preserving workflows built around this VM. See the VM Linking section for more information.

At this point, the Rubrik system presents itself as an NFS v3 datastore to ESXi. The original VM is deprecated (renamed); however, keep in mind that the original VM may have been deleted, so this is would not be necessary.

Rubrik coordinates the addition to the VM inventory in vCenter Server. A new copy of the VM running on Rubrik is presented and powered on and services resume.

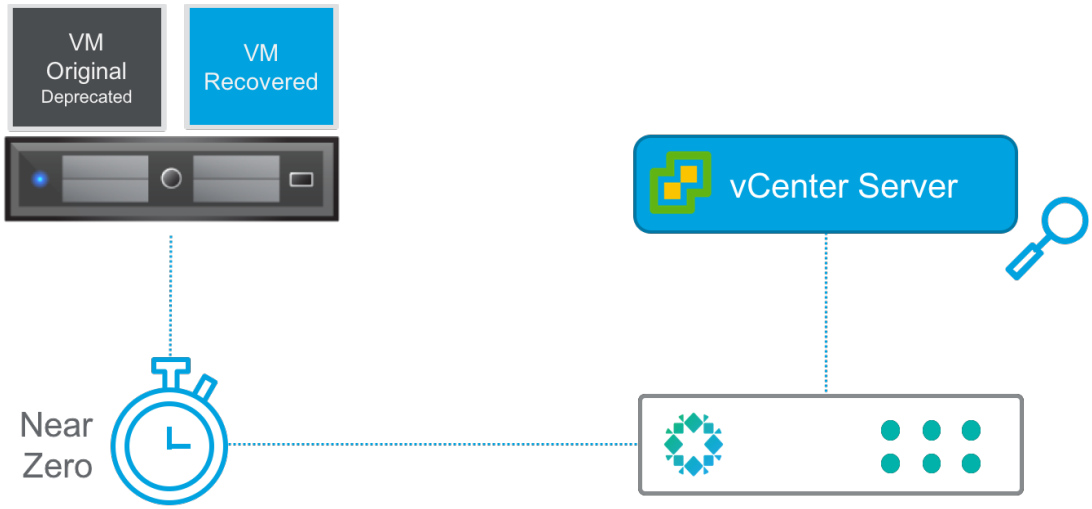

Post-recovery, users can Storage vMotion to the primary storage array.

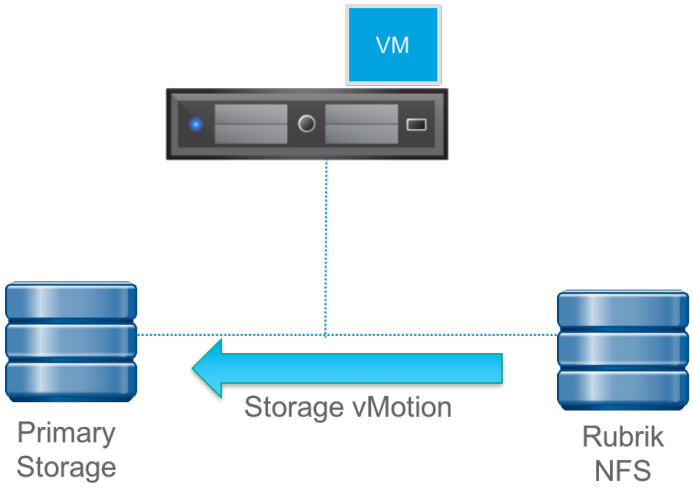

Ultimately, Rubrik serves as a storage endpoint to recover as many vSphere VMs as needed, thus eliminating the complexity and time wasted in transferring data back into the production system. This functionality provides a near-zero recovery time and restores user access near instantly.

During the process, messages about the status appear in Notifications. The Rubrik cluster records the final result of the task in the Activities Log.

The instantly recovered VM derives protection from parent objects. When the recovered VM does not derive protection from any parent objects, add it to an SLA Domain. To protect it using the same SLA rules and policies as the source VM, add the recovered VM to the original SLA Domain or to another SLA Domain. With VM linking, the new VM is linked with the old VM, which preserves the entire snapshot history.

LIVE MOUNT

Live Mount is similar to Instant Recovery, but the distinct difference is that the integrity of the VM snapshot is not altered during Live Mount functions.

Like Instant Recovery, Rubrik becomes an NFS v3 datastore from vSphere ESXi hypervisor perspective. Instead of deprecating the original VM, a VM similar to the original is created but with a trailing date timestamp appending the VM name. The original VM is not altered. Additionally, in order to avoid IP or MAC address conflicts, the Live Mount VM has its NIC disabled by default.

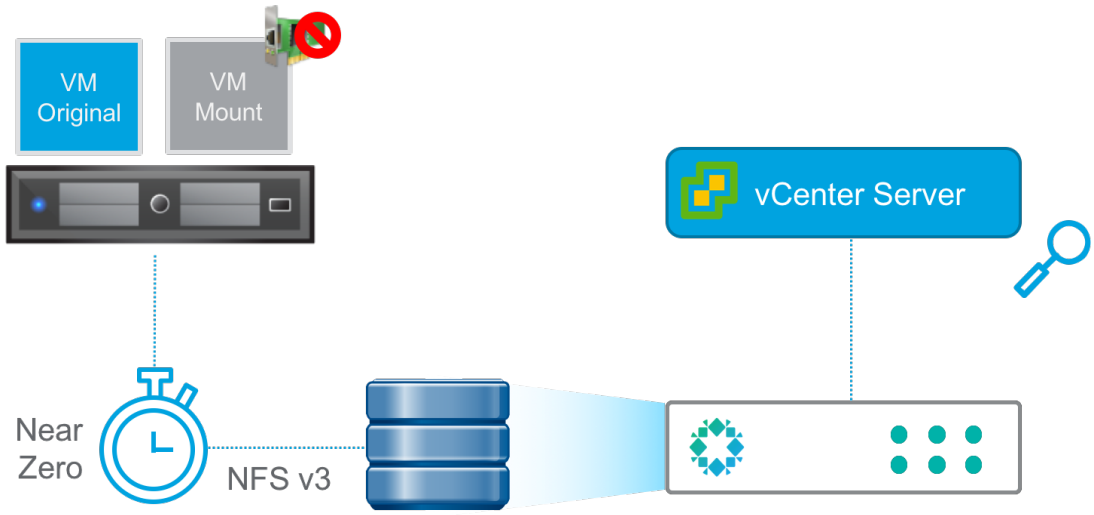

This functionality appeals to application owners and operations teams in order to conduct:

- Functional or regression testing
- Application development
- Software release testing (upgrade the actual applications)

Build isolated environments and leverage the Live Mount feature to instantiate an identical environment in moments. Test VMware Tools or hardware version upgrade, failure scenario, or other use cases using your backup storage. When done, simply throw it away.

No additional configurations are needed on the hypervisor side for Live Mount functionality to work. Simply provide a service account with the documented permissions required for the type of virtualization environment. Rubrik automates the entire process. VMs of any size can be recovered in the amount of time it takes for the OS to boot. Imagine having the ability to spin up an 8 TB VM in under 2 minutes so that a recovery point can be validated by an administrator or application owner.

EXPORT

An Export creates a new VM from a point-in-time copy of the source VM. The chosen ESXi host allows the section of the datastore for the recovered VM.

The Rubrik cluster assigns a new name to the recovered VM and powers it up. The recovered VM is not connected to a network.

FILE-LEVEL RECOVERY

The Rubrik cluster provides file-level restore (FLR) of files and folders from any local snapshot, replica, or archival snapshot that was successfully indexed.

To restore a file or folder, search for the file or folder by name across all local snapshots. You can also browse for the file or folder on a selected snapshot.

> Note: The Rubrik cluster must download an archival snapshot before it can be browsed. Searching by name for a file or folder on an archival snapshot does not require that the archival snapshot be downloaded first.

Files and folders may be restored directly to the source system or by download.

DIRECT RESTORE

For supported Windows and Linux guest operating systems, the Rubrik cluster can restore files and folders directly to the source file system.

When restoring from a snapshot of a supported guest operating system, the web UI provides the option to restore a file or folder directly to the source file system. When this option is selected, the web UI provides a choice to overwrite the source file or folder, or to restore the file or folder to another location.

A restored file or folder inherits the access control of the parent folder and the same owner as the parent folder. The restored file or folder retains the modification time (mtime) of the source file or folder at the time of the snapshot.

To successfully restore directly to the source file system, the Rubrik cluster must be provided the following information:

- Resolvable hostname or IP address of the authentication server
- Username of an account with Administrator privileges for the target
- Password for the account

When the Rubrik cluster has previously accepted the service credentials of a guest operating system, the restore job does not require additional credential information. This feature requires that the Rubrik cluster has successfully used the service credentials for at least one backup prior to the restore task. Otherwise, the credentials can be provided through the Restore File dialog during the restore task.

RESTORE BY DOWNLOAD

The Rubrik cluster generates download links to use for file-level restore (FLR) of files and folders from any local snapshot, replica, or archival snapshot that was successfully indexed. The guest OS of the source virtual machine must have a current version of VMware Tools running to enable successful indexing.

Restore a file from a data protection object through the Rubrik cluster web UI. Once the file is selected, the Rubrik cluster processes the request and provides a link for download of the file.

When restoring a folder, the Rubrik cluster generates a `.ZIP` file containing the folder and all its contents. The `.ZIP` file retains the hierarchy of the selected folder. The Rubrik cluster provides a link for downloading the `.ZIP` file.

DECLARATIVE POLICIES AND SLA DOMAINS

Traditional architecture has long been ruled by the imperative operational model. Historically, administrators have taken some piece of infrastructure and then told it exactly what to do to meet the desired end state. In terms of data lifecycle management, this translates to defining what objects to protect, target destinations, creation and expiration schedules, storage requirements, and so on. Each job requires a non-trivial amount of daily management to function. If there are issues with the job, an administrator must triage the job to determine where the failure occurred (along with re-running the job at a later date).

One of the most positive and impactful shifts in enterprise architecture has been the move towards the *declarative model*. This refers to the ability to express business needs directly to the systems that run applications with the intent of allowing an intelligent fabric of components to make real-time decisions on your behalf.

The declarative model allows technical professionals to plug in their desired state for an object – in this case, the data protection policy for VM workloads – into a policy engine. This engine is elegantly simple because all of the imperative details are abstracted away and handled by an incredibly smart, scale-out system. The resulting input fields are reduced to:

- The Recovery Point Objective (RPO) requirement
- Retention periods for the aforementioned RPOs
- Any archive targets, if desired
- Any replication targets for near-zero RTO requirements, if desired

Policy is logically assigned to any vSphere object: VMs, folders, data centers, clusters, or entire vCenter Servers, as well as constructs outside of vSphere, such as physical workloads, SQL databases, etc. Any of the "jobs," per se, are completely abstracted away by the system. The declarative policy engine funnels your RPO, RTO, availability, and replication requirements into system-level activities. This is where the true value of the system resides – the ability to control end-to-end ingest, placement, and archive for all protected pieces of data. Just set a policy and allow the system to do all of the heavy lifting. This is how the technology industry as a whole is going to tackle the ever-increasing demands for doing more with less, faster and more efficiently.

As an example, imagine you have invited someone over to your house. In order for the person to arrive at your home, you must give exact directions — "start by going straight down Main Street, then right at the In-N-Out Burger, ensure to follow the stop light instructions at the intersection of 1st Ave and A Street. My house is the ninth house on the left past that intersection." This is the imperative model of thinking. Alternatively using a declarative model, I could say "my address is 16 National Ave; input it into a Zero Trust Data Protection app — it will navigate you using the best route."

Rubrik is firmly rooted in the declarative approach; as an administrator, you simply define the desired end state (RPO, retention, replication, archival, etc.) and allow the intelligent software to make it reality. In essence, govern infrastructure and applications using declarative policy rather than imperative jobs.

SLA DOMAINS

Rubrik orchestrates the movement of data from initial ingest and propagation of that data to other data locations, such as replicating to remote clusters or Rubrik Cloud Cluster, as well as data archival. A single SLA policy is used to dictate all data lifecycle specifications, and the data control plane does the rest.

In the SLA Policies section, an example SLA policy was given:

- Take a backup:
  - Run a snapshot every 4 hours and retain hourly backups for a day
  - Run a snapshot every month and retain monthly backups for 7 years
- Archive to Amazon S3 after 30 days
- Replicate data to another Rubrik cluster and retain for 45 days.

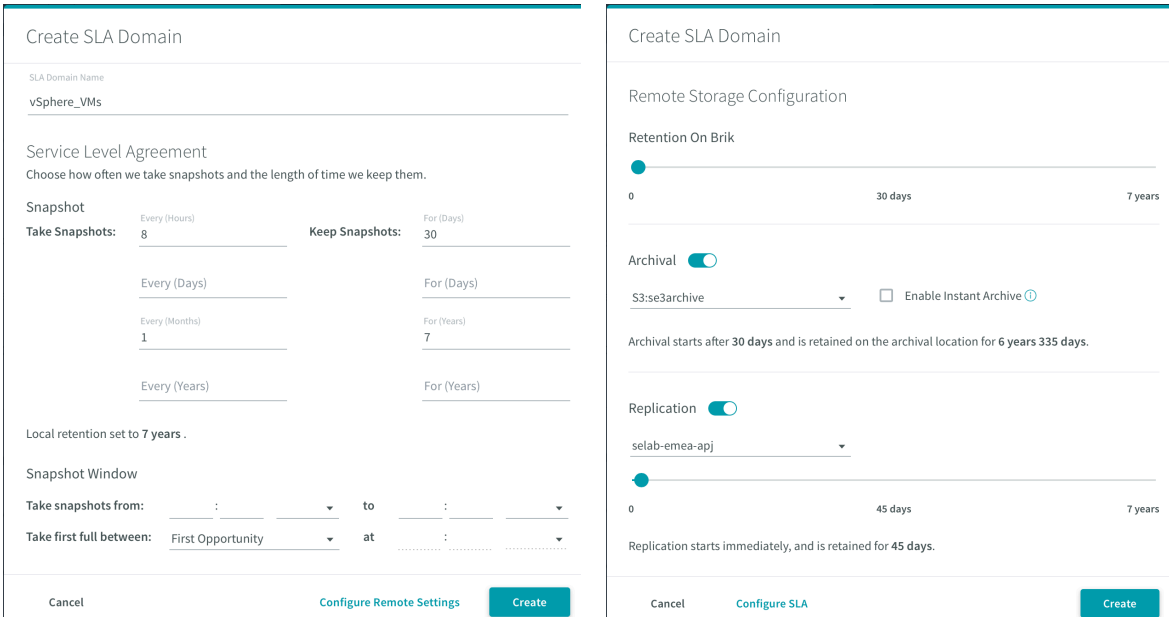

Data is ingested and retained according to the frequency specified in the SLA policy. The example policy is configured to store 30 days of data within the Rubrik cluster. Once that period has elapsed, data is archived to another location for long-term retention. In this case, data is archived to Amazon S3 for another 6 years and 335 days. There is no need for an administrator to manage, prune, or validate that data has been archived; these activities are all handled natively by Rubrik to reflect how they were expressed in the SLA.

The policy also specifies to replicate data from one Rubrik instance to another. For example, a remote office/branch office (ROBO) may replicate workloads into the main data center using Rubrik or a primary site may replicate to a DR site. Eliminate configuring and managing this functionality at the storage layer. Apply policy-based management to workloads and stop babysitting data residing across multiple data centers.

Regardless of where the data is archived, Rubrik ensures instant accessibility of data with real-time predictive search. Metadata is included in the archive to ensure the most cost-efficient way to recover data by removing the need for recovering full backups from archive before restoring. This provides the ability to recover archived data at a snapshot or file-level selectively without having to download the entire workload to restore a single file and reduces egress charges.

ASSIGNING SLA DOMAINS

Once the policy has been created, provide protection for a VM by assigning an SLA Domain.

A VM can be protected by assigning an SLA Domain setting individually to the VM. A VM can also be protected by deriving an SLA Domain setting through automatic protection.

Automatic protection occurs in one of the following ways:

- An administrator assigns an SLA Domain to an object that contains the VM.
- An administrator moves the VM into the hierarchy of an object that is assigned to an SLA Domain.

This means that VMs will be protected through inheritance of the SLA policy assigned to a parent object. If the vCenter Server or a folder has an SLA assigned to it, the VM underneath will automatically inherit the policy. The data control plane detects the newly added VM and automatically applies a protection policy, eliminating the need for any manual administrator interaction. This resolves the common issue of new workloads being brought online and going days or weeks without being protected.

In the event that an SLA policy has been assigned to an individual VM that auto-inherits the policy from a high-level object, conflict resolution occurs. When a conflict is detected, the Rubrik cluster opens the SLA Conflicts dialog box to permit the conflict to be resolved.

In addition to overriding SLA policies, if desired, inheritance may also be blocked by applying a "Do Not Protect" policy at the object level.

SLA policies may be hierarchically assigned to:

- vCenter Server
- Clusters
- Folders
- ESXi hosts
- VMs
- Tags, using PowerShell

Once the policy is assigned, Rubrik will ensure adherence to user-defined policies such as frequency, retention, archival, etc. as described above. All manual configuration is eliminated by the data control plane, which applies intelligent algorithms to ensure efficiency and performance for the entire backup workload. These intelligent algorithms assist with balancing the workload as more VMs are created and added into the system. The automatic scheduling of tasks ensures that all workloads are evenly distributed across the Rubrik cluster, preventing cluster resource contention.

PROTECTION OVERVIEW

Rubrik provides backup protection for VMs by combining native snapshot technology with the fast and scalable converged data management platform of the Rubrik cluster.

PERFORMANCE AND SCALABILITY

The Rubrik cluster provides a high-performance, highly-scalable integration with vSphere APIs for Data Protection to backup VMs hosted on ESXi hypervisors.

By efficient use of vSphere APIs for Data Protection calls and by providing very fast data ingestion, the Rubrik cluster minimizes the time that a VM is quiescent during a backup. This reduces and, in most cases, eliminates the application time-outs caused by many other backup products. The time that a VM is quiescent, sometimes referred to as VM stun or application stun, is the time between the following:

- The point where execution of the VM is paused, at an instruction boundary, and all in-flight disk input/output operations are completed
- The point where execution resumes

The period a VM is quiescent, is very brief, just long enough to create a snapshot. The VM does not remain quiescent during the processing and ingestion of the snapshot data.

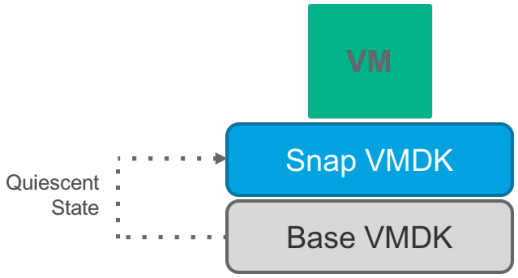

To help minimize the time that a VM is quiescent, the Rubrik cluster maintains multiple concurrent connections with a vSphere environment and opens five threads for each ESXi host in that environment.

The Rubrik cluster also efficiently uses the 10 Gigabit Ethernet connection to the vSphere environment. It provides a very high rate of data ingestion to the flash-based write cache that is the initial storage of the Rubrik cluster.

The result is an extremely short time that a VM is quiesced.

For best performance, use a 10 Gigabit Ethernet connection between the Rubrik cluster and the vSphere environment. Also, for replication, it is recommended to provide a 10 Gigabit Ethernet connection between the source Rubrik cluster and the target Rubrik cluster.

The Rubrik cluster uses a distributed task scheduler that permits the Rubrik cluster to schedule tasks to run on any node and on multiple nodes, as needed. Since the distributed task scheduler can seamlessly schedule tasks on all available nodes and across multiple nodes, adding nodes to a Rubrik cluster further increases ingestion and processing efficiency.

BACKUP PROCESSES

A Rubrik cluster backs up a VM by using vSphere APIs for Data Protection to create a snapshot of the VM. When a Rubrik cluster begins protecting a VM, the Rubrik cluster starts by creating a first full snapshot of it. This first full snapshot is a complete backup of the VM.

After the first full snapshot, the Rubrik cluster continues protection by creating incremental snapshots based on the change information provided by change block tracking (CBT). The Rubrik cluster creates each incremental snapshot very quickly because the snapshot only includes the data blocks that have changed since the last snapshot.

The vSphere environment transmits the snapshot data to the Rubrik cluster using the most efficient available transport mode. Normally, the vSphere environment uses the NBD/NBDSSL transport mode. The high efficiency of the Rubrik cluster eliminates data bottlenecks, allowing the NBD/NBDSSL transport mode to provide data transmission rates that minimize the time that a VM is quiescent.

For VMDKs that are stored on a SAN, the Rubrik cluster can use the SAN transport mode. In this mode, the Rubrik cluster uses the iSCSI protocol to obtain snapshot data over a direct connection to the storage array resulting in very fast data transmission.

ARCHITECTURAL OVERVIEW

Rubrik is a vendor-agnostic platform that is built on an API-first architecture. The SLA policy is the heart of every Rubrik configuration. Rubrik reduces daily operational management, providing a step-function change in simplicity by enabling a single-policy engine to orchestrate SLAs across the entire data lifecycle. SLA policies can be applied anywhere in the vSphere hierarchy stack: vCenter Server, the cluster, host, folder, or VM levels. The Rubrik cluster provides a variety of methods to recover VMs and restore protected data. Rubrik recovers VMs and restore data by using snapshots, replicas, and archival snapshots.

End-to-end data management is provided by Rubrik for all applications running on vSphere. Users can securely access data instantly, automate protection policies, and orchestrate data across their VMware environments.

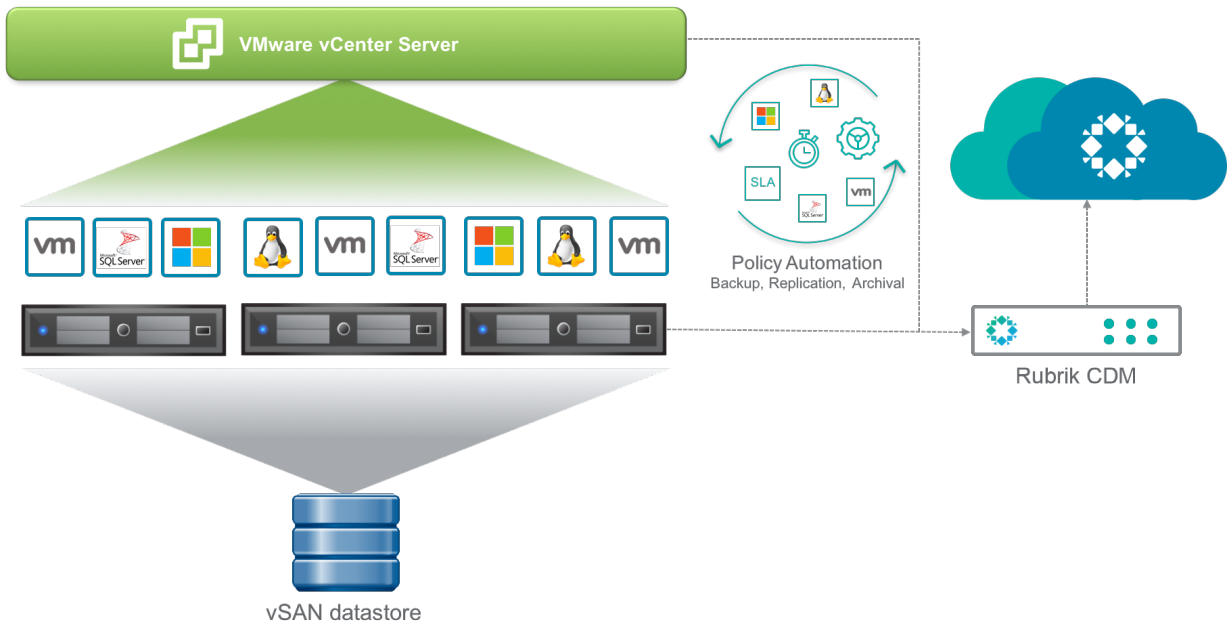

The following details a few requirements for this integration detailed in the reference architecture:

- VMware vSphere 6.5 or later
  - VMware ESXi 6.5 or later
  - VMware vCenter Server 6.5 or later
  - Minimum of virtual hardware version 7 for CBT
- Primary Storage
  - Storage agnostic
  - Specific storage integration support matrix may be found at https://support.rubrik.com.
- Ethernet (IP) Network
  - A VMkernel port is required for NFS if using Live Mount or Instant Recovery functionality.
  - Additionally, a separate VMkernel network may be configured for Rubrik data traffic.
  - Details surrounding required ports may be found in Appendix A.
- Rubrik CDM 4.x
- vCenter Server privileges
  - It is recommended that that a custom vCenter role be created with only the minimum required privileges. Appendix B details all required privileges.
- VMware Tools
  - The Rubrik cluster requires the current version of VMware Tools to perform administrative operations and enable application-consistent snapshots.
  - If VMware Tools is out of date, the backup will proceed but will not be application consistent.

Additionally, the following assumptions have been made in the writing of this document:

- Workloads are supported by VMware vSphere.
- Workloads are using supported versions of their operating system and application release(s).
- 10 GbE network connectivity exists between the ESXi Host(s) and Rubrik Cluster.
- 4 Node Rubrik Cluster

Lastly, a few constraints since Rubrik cannot protect data that exists on any of the following using native vSphere integration through APIs:

- VMDKs that are set to Independent-Persistent mode or to Independent-Nonpersistent mode
- Network drives that are mounted on the file system of a protected virtual machine
- Any VM for which the Rubrik cluster does not have snapshot creation permission because of settings on the VM or on a vSphere folder that contains the VM
- Any VM data that resides on raw disk mappings (RDMs), where the compatibility mode of the RDMs is set to Physical

That being said, these constraints apply to protection using native vSphere APIs but can be backed up using RBS.

These requirements and assumptions should be taken into account for the remainder of the document.

CHANGED BLOCK TRACKING

Changed Block Tracking (CBT) is a VMware feature that helps perform incremental backups. This is leveraged by Rubrik for efficient data protection.

VMs running on ESXi hosts can track disk sectors that have changed. On many file systems, CBT identifies the disk sectors altered between two change set IDs. On VMFS partitions, CBT can also identify all the disk sectors that are in use.

Virtual disk block changes are tracked outside the VMs in the virtualization layer. When Rubrik performs a backup, it requests transmission of only the blocks that changed since the last backup or the blocks in use. The CBT feature is accessed as part of the VMware vSphere Storage APIs – Data Protection. Rubrik calls vSphere APIs for Data Protection to request that the VMkernel return blocks of data that have changed on a virtual disk since the last backup snapshot.

Rubrik's distributed file system, Atlas, is built from the ground up by the Rubrik engineering team. It marries the design principles of modern, web-scale file systems with the ability to handle random writes and intelligent version awareness. Conceptually, Atlas stores a set of versioned files - all protected data is immutable. As new data enters the system, Atlas preserves older versions by writing to new blocks and always utilizing full stripe writes. Combined with using CBT for incremental changes, this makes later access those that data extremely efficient.

ADAPTIVE DATA CONSISTENCY

Data consistency in recovery points is broadly classified into three categories: ***inconsistent***, ***crash-consistent***, and ***app-consistent***.

An *inconsistent recovery point* is taken with zero pre-work. It is not suitable for data sets that contain complex and interdependent relationships, as only data captured on disk is backed up. The in-memory changes are not captured, so it won't be a true representation of the data in the system for that point in time.

A *crash-consistent recovery point*, when restored, gives you the state of the data from the time of the backup. All the data is captured at the same time, but I/O operations and transactions in process may not be captured. For most modern applications, crash-consistent recovery points may be sufficient.

An *app-consistent recovery point* captures all of the data simultaneously, just like a crash-consistent recovery point. But it also waits for the applications to flush I/O operations and transactions in process. For apps running in Microsoft Windows-based operating systems, the Volume Shadow Copy (VSS) service helps with creating app-consistent snapshots. For Linux-based systems, the applications may have native tools and/or file system sync tools to help with creating app-consistent recovery points.

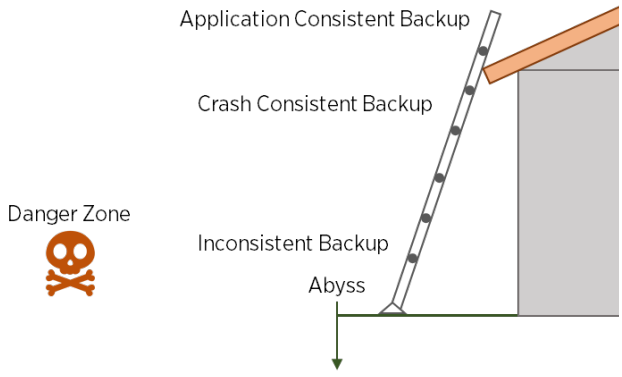

When talking to our customers, we discovered that backup administrators had been dealing with two key pain points in this area.

Pain Point #1: Too many knobs to turn and tweak

Traditional backup solutions require administrators to make decisions on consistency of recovery points in advance. Several 'knobs' are provided in backup job definitions, application processing settings, guest operating system settings, and so on, which allows the administrator to make decisions on a per-job or per-VM basis. While this method is flexible, it can get overwhelming, especially when you need to protect hundreds of VMs. Furthermore, the human errors in setting up backup job or agent attribute can be costly because there is no reporting for data consistency, as the user is supposed to make decisions deliberately.

Pain Point #2: Secure or easy to use?

Traditional backups also force the administrators to decide the security profile of the guest operating systems in advance. Most of today's VM backup solutions provide agentless operation if User Access Control (UAC) is disabled in the guest operating system. If UAC is enabled, the backups fail or the backup solution blindly performs a crash-consistent backup of the VM. In order to perform application-consistent backup of UAC-enabled VMs, you must use separate backup jobs with agent-centric workflows. Or you have to expose elevated credentials of the VMs to the backup system, which defeats the very purpose of securing the VM using UAC! Now, imagine making this decision for hundreds, or even thousands, of VMs.

To address these two problems, Rubrik designed *adaptive data consistency*.

Rubrik features a custom VSS provider for Windows VMs running in VMware vSphere. There are several enhancements in this custom VSS provider designed to leverage Rubrik's scale-out architecture. Furthermore, this custom provider handles SQL Server and Exchange differently; it does not break application-level backups of SQL Server and performs application-consistent log truncation for Exchange.

Using Rubrik eliminates the two pain points above. No additional configurations are required to make use of Rubrik's adaptive data consistency methodology! The system adaptively determines the best possible data consistency path for the VM. If the chosen path is not the best, you are notified so that you can perform corrective actions.

How does it all work under the hood? Rubrik attempts to contact the VM over the network to see if Rubrik Backup Service (RBS) is listening to the VM. If RBS is listening, our platform orchestrates app-consistent backup by making use of custom a VSS agent that is already installed as part of RBS. This agent-assisted method has the following benefits:

- Unlike traditional backup solutions, the administrator does not have to make complex decisions on data consistency, security (UAC), and job setup.
- The method works in both UAC enabled and disabled environments.
- The method does not have the adverse delays generally seen when using vSphere VIX APIs used solely for the purpose of creating app-consistent snapshots
- The actual data movement occurs via NBDSSL transport through ESXi server, hence still an agentless ingest method.

Additionally, if RBS is not listening, Rubrik has the ability to auto-install RBS on the VM. This capability is turned off by default, as it may be considered intrusive to install persistent binaries on a production system. If needed, this capability may be turned on temporarily to streamline the installation process. As this process requires UAC to be temporarily disabled, it may not be viable in some environments.

If RBS is not active or cannot be installed, Rubrik pushes an ephemeral agent into the guest OS using vSphere VIX APIs. This ephemeral agent has the custom VSS provider. Rubrik attempts app-consistent snapshots of the VM using its custom VSS provider. In the event this relatively slow VIX based interaction makes it impossible to finish VSS snapshot with the (Microsoft mandated) 10 second VSS window, Rubrik will notify the situation and proceed with a crash consistent backup of VM. The administrator can perform corrective actions based on notification received

Providing an adaptive capability to provide best effort data consistency is important. Manually tweaking job settings, guest and application processing settings, client attributes, consistency choices, and more is doable with a handful of VMs. But in most cases, you also need to work with application admins/owners of the VMs to choose an approach that meets their recovery requirements. As your environment scales to hundreds or thousands of VMs, this is not only hard to do, but the cost of making a mistake can be very high.

The time consuming but extremely critical decisions on VM and app data consistency are now in self-learning mode because of Rubrik's adaptive approach. Combined with auto-discovery and the adaptive throttling already baked into the product, Rubrik's self-driving (declarative) capabilities ensure simplicity and manageability for protecting VMs at scale.

RUBRIK BACKUP SERVICE (RBS)

Linux and Windows continue to be the most pervasive operating systems in today's data centers. IBM AIX has been utilized by many large organizations for decades and continues to be deployed to support mission-critical applications. Rubrik has innovated a method used to protect all of these different system types with Rubrik Backup Service. Specifically, the RBS provides the Rubrik cluster with the ability to granularly protect data on Linux, AIX, and Windows file systems, whether physical servers or VMs.

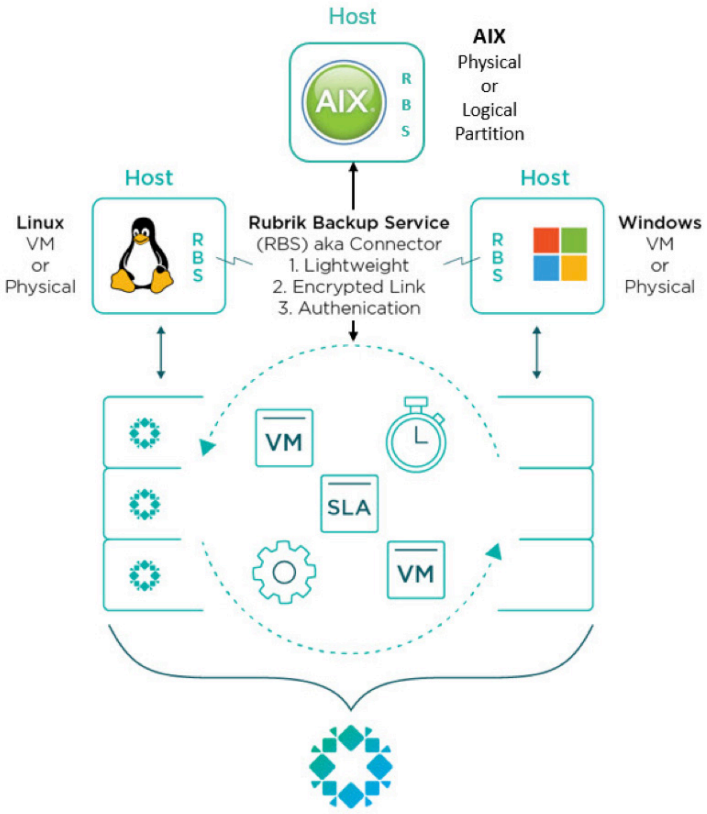

The RBS software can only be used with the Rubrik cluster from which the software is obtained. Each Rubrik cluster generates a copy of the RBS software that includes authentication information specific to that Rubrik cluster. This method ensures that the Rubrik cluster and a hosted deployment of the RBS can reliably authenticate each other. As an additional security measure, all communication between Rubrik and RBS is encrypted.

> It is recommended to use Group Policies or automation tools such as Chef, Puppet, or something similar to automate RBS installation.

Rubrik provides automatic upgrade of the RBS software as part of a general upgrade of the Rubrik cluster software. After upgrading the Rubrik cluster software, the Rubrik cluster automatically upgrades the RBS software on all protected Windows Server hosts. This increases manageability in a large site that is using RBS across thousands of objects.

The RBS must run as an account that is a member of the Administrators group of the Windows Server host.

When first installed, the RBS runs as a LocalSystem account. A LocalSystem account includes the permissions that are provided by the local Administrators group.

Instead of running the RBS as a LocalSystem account, it can be configured to run as a member of the local Administrators group.

To run as a member of the local Administrators group, run the RBS as a user account that is one of the following:

- Local user account that is a member of the local Administrators group
- Domain user account that is a member of the local Administrators group

VM LINKING

Each time a new vSphere VM is discovered, the Rubrik cluster assigns it a UUID. The Rubrik UUID is a combination of the UUID assigned to a vCenter Server when added to the Rubrik system + the MoRef ID (Managed Object Reference ID) of the VM within the vCenter Server.

A MoRef ID is guaranteed to be unique within a vCenter Server instance. It is used to track all entities such as VMs, hosts, data centers, resource pools, etc. and is stored within the vCenter Database. An entity's MoRef ID will remain the same throughout its lifecycle unless it (e.g. VM or ESXi host) is destroyed or unregistered from vCenter Server.

A VM's Rubrik-assigned UUID will change when either the vCenter Server or its moRef ID is changed. This could occur when:

- **vCenter UUID Change –** a vCenter Server spans data centers and multiple Rubrik clusters with each Rubrik cluster residing in one of the data centers and both protecting the same vCenter Server while replicating to the other Rubrik cluster. In this scenario, the MoRef ID of the VMs may stay the same, but the UUID assigned to the vCenter Server will be unique to each Rubrik cluster, avoiding the conflicts among the objects during replication.

- **VM MoRef ID Change –** the moRef ID changes whenever a VM moves from one vCenter Server to another. The moRef ID also changes whenever a VM is unregistered and re-registered to the same vCenter Server. Rubrik creates a brand new VM object and archives the old one, turning it into a relic VM.

This behavior has two side effects.

1. **Backup History –** the history of snapshots associated with a given virtual machine is reset, which makes compliance enforcement complicated from a customer perspective.

2. **New Objects –** whenever the VM UUID changes within Rubrik, it is treated as a brand new object that, when protected, results in a full snapshot that may or may not be deduplicated.

A linked VM is when Rubrik joins two VMs into the same snapshot chain / history, such as when customers migrate vCenter Servers without migrating the database. This results in new MoRefs on all VMs. Without VM linking, this would result in all new VM full backups, which wastes space from unmanaged objects.

Rubrik provides two options to the user: 1) let Rubrik to resolve conflicts automatically or 2) discard conflicts results by creating an independent VM in Rubrik (existing behavior). In the UI, when a user adds or edits a vCenter Server, a prompt will appear to select either automatic or discard for conflict resolution. Rubrik will follow this instruction for a vCenter Server while discovering the VMs as part of the vCenter Server refresh.

This new option addresses multiple use cases, which include:

- **vCenter Server failure or migration:** Large organizations with multiple data centers commonly conduct failover tests across data centers to validate disaster recovery preparedness. In such cases, the DR protected VMs should be linked with the primary VM in Rubrik for tracking and automation of failovers/failbacks.
- **Automated recovery with vRA:** VM linking is also an important component in automating a VM restore workflow through vRealize Automation (vRA). This is a common use case for customer using vRA for vMotion or VM restores. vRA must restore a VM in place (same storage and same VM UUID), otherwise the VM becomes orphaned from vRA.
- **Instant Recovery:** When a virtual machine is instantly recovered, the recovered VM comes with a new moRef ID that generates a new Rubrik assigned UUID. Previously, this new VM is independent and loses snapshot history. With VM linking, the new VM is linked with the old VM, which preserves the entire snapshot history. (Note that in this case, VM linking will occur regardless of the vCenter Server settings for conflict resolution.)

OPERATIONAL OVERVIEW

Rubrik and VMware vSphere product lines leverage a complementary progressive hybrid cloud enterprise architecture, with the goal of accelerating applications and business requirements. This section aims to highlight how lightweight the Rubrik and VMware vSphere joint solution is.

VCENTER SERVER

The Rubrik cluster accesses VM data through a connection to the VMware vCenter Server that manages the hypervisor where the VM is running. To successfully connect with a vCenter Server, the Rubrik cluster requires connection information for that vCenter Server. This information includes:

- vCenter Server FQDN or IP Address
- Username (with proper permissions, details can found at Appendix B)
- Password

The Rubrik cluster provides access to vCenter Server information on the vCenter Servers page. That page provides the FQDN or IP address and the connection status for every vCenter Server that is added to the Rubrik cluster.

After connection information for a vCenter Server is added, the Rubrik cluster requests relevant metadata from the vCenter Server, such as folder, cluster, and host information. The Rubrik cluster uses the metadata to display and work with the VMs on the vCenter Server.

The Rubrik cluster automatically refreshes the metadata from a vCenter Server every 30 minutes. This is referred to as a light refresh. The Rubrik Edge appliance performs a light refresh of a vCenter Server every six hours.

The Rubrik cluster automatically refreshes the metadata and rescans the VMDK files of a vCenter Server every two hours. This is referred to as a full refresh. The Rubrik Edge appliance performs a full refresh of a vCenter Server every 24 hours.

VMDK files are also automatically scanned as part of every create snapshot job. A full refresh can be manually initiated at any time.

PROTECTING VCENTER SERVER AND PLATFORM SERVICES CONTROLLER

vCenter Server comes in two forms, a Linux-based virtual appliance or Windows. VMware has announced the deprecation of vCenter Server for Windows, vSphere 6.7 will be the final release for vCenter Server for Windows. The vCenter Server Appliance has surpassed its Windows counterpart in functionality as of vSphere 6.5 and is now considered the default deployment for vCenter Server.

vCenter Server can be configured in the following ways:

- As an embedded deployment model with an internal Platform Services Controller
- As an external deployment model with an external Platform Services Controller. To support enhanced linked mode in vSphere 6.0 and 6.5 an external deployment was required.

As of vSphere 6.5 Update 2 and vSphere 6.7 or later, the vCenter Server Appliance now supports Enhanced Linked Mode regardless of whether configured as embedded or distributed deployment model. Embedded enhanced linked mode support is only for the vCenter Server Appliance.

The role of the Platform Services Controller is to manage authentication, licensing, tags and categories, global permissions, and custom roles. It is also the certificate authority for the vSphere Single Sign-On domain. The Platform Services Controller is multi-master, and automatically replicates the same data to any other Platform Services Controller in the same vSphere Single Sign-On domain.

Important note: The vSphere Distributed Switch environment should be backed up as often as the VCSA. It is easy to script using PowerCLI `Export-VDSwitch` and `Export-VDPortGroup` cmdlets. The script can be executed as part of a pre-action in the backup so there is always a current backup of the vSphere Distributed Switch (vDS). Alternatively, this can accomplished using the vSphere Web Client.

EMBEDDED DEPLOYMENT

This deployment model installs all the vCenter Server services on a single node. The following diagram represents a vCenter Server Appliance embedded deployment:

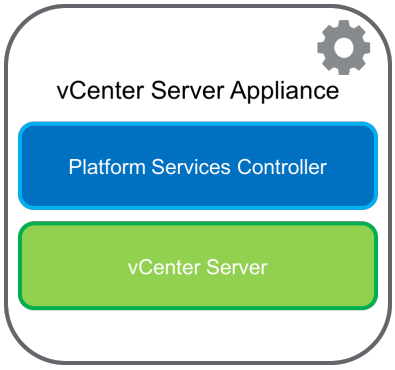

In this example, the entire VM should be protected. If necessary, the entire VM can be restored. It is always best to backup or restore the whole appliance. The database should not be modified.

While Rubrik is capable of file level protection, it is important to note VMware supports only an image-level or file-based backup of the vCenter Server Appliance. File-based backups and restores are supported by using the VMware Appliance Management Interface using port 5480 (https://vCenter Server Appliance FQDN or IP Address:5480) or RESTful APIs found in the API Explorer (https://vCenter Server Appliance FQDN or IP Address/apiexplorer)

EXTERNAL DEPLOYMENT

This deployment model installs the Platform Services Controller on a separate node from where vCenter Server is installed. Installing the Platform Services Controller is a prerequisite for installing vCenter Server. The configuration is demonstrated in the following diagram:

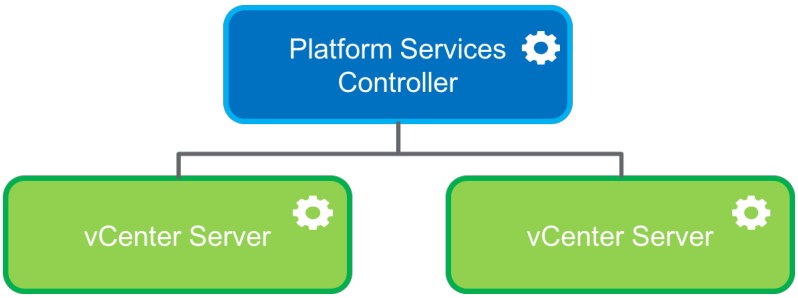

It is recommended to protect each vCenter Server Appliance node and at least one Platform Services Controller (if in the same SSO domain). If there are multiple SSO domains, then a minimum of one Platform Services Controller should be protected in each SSO domain.

RESTORING VCENTER SERVER AND PLATFORM SERVICES CONTROLLER

This section is intended to provide several scenarios to assist with architecting an appropriate backup solution for vSphere management components.

EMBEDDED DEPLOYMENT

In this scenario, vCenter Server is installed on the same node with the Platform Services Controller. If the entire VM is down, the embedded vCenter Server deployment should be recovered from backup, at which time the entire VM should be restored.

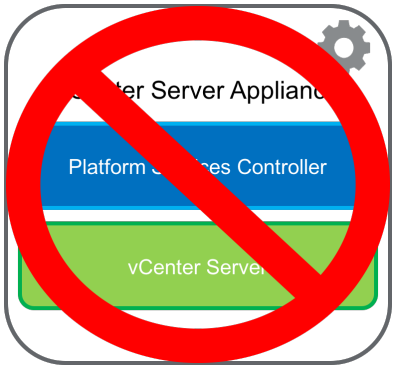

Rubrik allows restores directly to an ESXi host in the event that vCenter Server is unavailable.

EXTERNAL DEPLOYMENT

If vCenter Server and Platform Services Controller nodes fail at the same time, you must first restore the Platform Services Controller and then the vCenter Server nodes.

Scenario 1 — One Platform Services Controller Unavailable

This scenario has multiple Platform Services Controllers in the same vSphere Single Sign-On Domain being accessed by multiple vCenter Server Appliances. The following diagram visualizes this architecture:

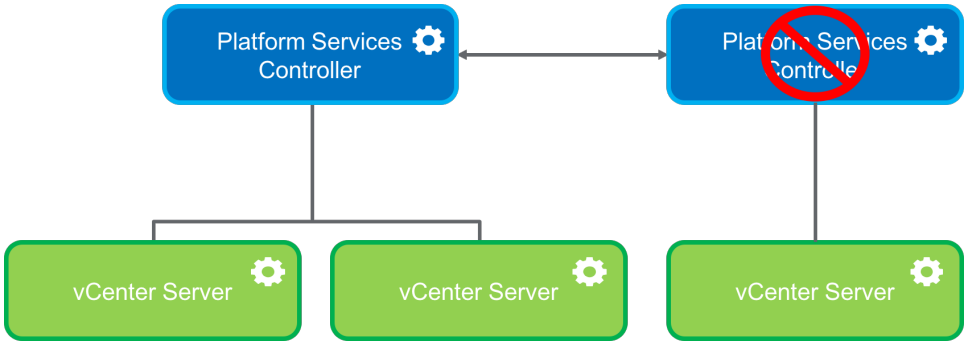

In the situation that a single Platform Services Controller is unavailable but there is at least one surviving Platform Services Controller in the same vSphere Single Sign-On domain (vCenter Servers are unaffected by this failure). When a Platform Services Controller is unavailable any vCenter Server registered to it cannot authenticate, thus preventing successful login. The following procedures should be followed to repoint to another Platform Services Controller in same vSphere Single Sign-On domain:

1. Repoint the vCenter Server Appliance to available surviving PSC. In vSphere 6.5 repointing is only allowed within the same site. vSphere 6.7 allows repointing within and across sites. Keep latency in mind when repoint across sites in vSphere 6.7.

    - The command on the vCenter Server Appliance to repoint is: `cmsso-util repoint --repoint -psc`
    - On average this process can take 5-10 minutes.
    - Authenticating with the vCenter Server Appliance is now available after the repointing process has completed to the other available Platform Services Controller. To test, simply try to log in.

The aforementioned test may fail in the event that not all Platform Services Controllers are configured for Active Directory. In this event, use the administrator account for the local SSO domain and configure Active Directory.

It is recommended that in a distributed deployment all PSCs are configured for Active Directory.

2. Decommission failed Platform Services Controller.

    - The command on the PSC to decommission is: `cmsso-util unregister --node-pnid --username --passwd`

Decommissioning allows the use of the same FQDN and IP Address when the Platform Services Controller is re-deployed in a later step.

    - Validate the Platform Services Controller has been removed from the vSphere Single Sign-On domain using the `vdcrepadmin` command to check all the available Platform Services Controllers.

        i. Example command: `vdcrepadmin -f whoservers -h PSC_FQDN -u administrator -w Administrator_Password`

        ii. This may be found in `/usr/lib/vmware-vmdir/bin`

3. Deploy a new PSC in the same vSphere Single Sign-On domain

4. Repoint the vCenter Server Appliance back the newly deployed Platform Services Controller, if necessary

This scenario does not require a restore of the Platform Services Controller from backup. Keep in mind that Platform Services Controllers are multi-master and the same data is replicated between Platform Services Controller nodes.

Scenario 2 — All PSCs Unavailable

This scenario has multiple Platform Services Controllers in the same Single Sign-On domain being accessed by multiple vCenter Server Appliances. The following diagram visualizes this architecture:

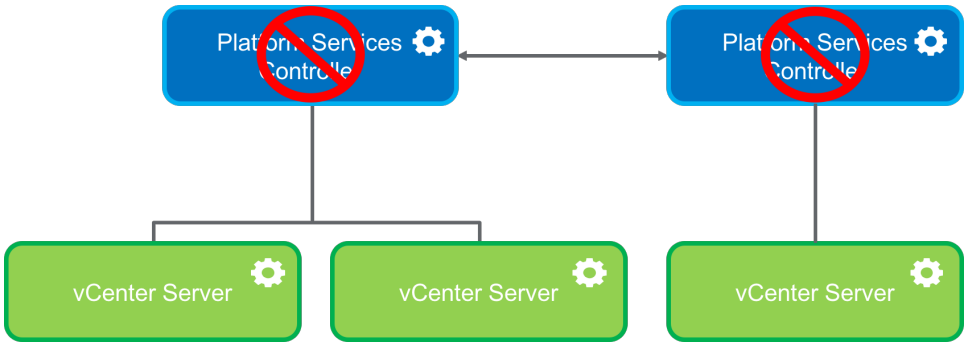

In the situation that all Platform Services Controllers are unavailable, the following procedures should be followed:

1. Restore one PSC.
2. Decommission the other failed Platform Services Controller.
    - The command on the PSC to decommission is: `cmsso-util unregister --node-pnid --username --passwd`

Decommissioning allows the use of the same FQDN and IP Address when the Platform Services Controller is re-deployed in a later step.

    - Validate the other failed Platform Services Controller has been removed from the vSphere Single Sign-On domain using the `vdcrepadmin` command to check all the available Platform Services Controllers
        i. Example command: `vdcrepadmin -f whoservers -h PSC_FQDN -u administrator -w Administrator_Password`
        ii. This may be found in `/usr/lib/vmware-vmdir/bin`

    Again, attempt to authenticate to confirm.

3. Deploy a new PSC in the same vSphere Single Sign-On domain.
4. Run the vCenter-restore script on the vCenter Server node registered with the restored Platform Services Controller (only for vCenter Server 6.5).
    - Example command run from vCenter Server Appliance shell: `vcenter-restore -u psc_administrator_username -p psc_administrator_password`
    - For vCenter Server 6.7 reconciliation is not required, stop and start all services.
        i. `service-control --stop --all`
        ii. `service-control --start --all`

Because all Platform Services Controllers are multi-master and the same data is replicated between all the Platform Services Controllers instances, it is only necessary to restore one PSC per SSO domain. Any additional Platform Services Controllers required can be newly deployed and synced with the restored Platform Services Controller.

Scenario 3 — All vCenter Server Appliances and Platform Services Controllers Unavailable

This scenario has multiple Platform Services Controllers in the same Sign Sign-On domain being accessed by multiple vCenter Server Appliances. The following diagram visualizes this architecture:

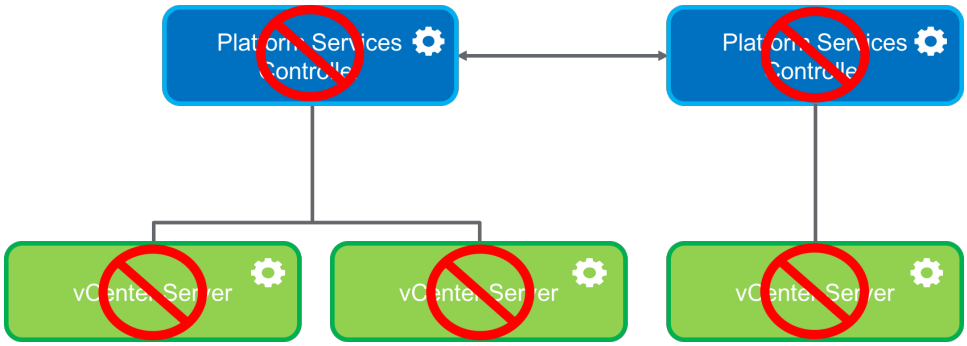

In the situation that all PSCs and VCSAs are unavailable, the following procedures should be followed:

1. Restore a PSC in the SSO domain.
    - Note: The PSC must be restored directly to an ESXi host since the vCenter Server Appliance that Rubrik is configured to use is unavailable.
2. Run the command on the PSC to decommission other unavailable Platform Services Controllers:
    - `cmsso-util unregister --node-pnid --username --passwd`

    Decommissioning allows the use of the same FQDN and IP Address when the Platform Services Controller is re-deployed in a later step.
3. Validate the other unavailable Platform Services Controller has been removed from the vSphere Single Sign-On domain using the `vdcrepadmin` command to check all the available Platform Services Controllers
    - Example command: `vdcrepadmin -f whoservers -h PSC_FQDN -u administrator -w Administrator_Password`
    - This may be found in `/usr/lib/vmware-vmdir/bin`
4. Restore the vCenter Server Appliances.
5. Repoint, as needed.
6. Deploy additional PSCs, as required.
7. Run the vCenter-restore script on the vCenter Server node registered with the restored Platform Services Controller (only for vCenter Server 6.5). For vCenter Server 6.7 reconciliation is not required, stop and start all services.
    - `service-control --stop --all`
    - `service-control --start --all`

Though it is only necessary to restore one PSC per SSO domain, each vCenter Server Appliance must be restored individually because each vCenter Server Appliance inventory is unique, whereas Platform Services Controllers within an Single Sign-On domain are multi-master.

Scenario 4 - Use of Load Balancer

Rather than recoverability, the Platform Services Controller can be made highly available by using a load balancer. A load balancer can ensure automatic failover without downtime. An example of this architecture is seen in the following diagram:

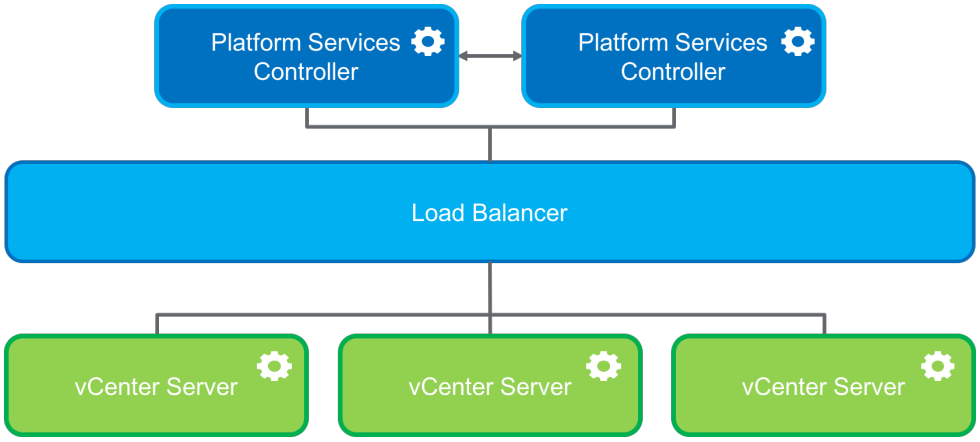

The vCenter Server nodes are connected to the load balancer. When a Platform Services Controller instance stops responding, the load balancer automatically repoints the vCenter Servers registered the unavailable Platform Services Controller to other functional Platform Services Controller node without downtime.

This methodology increases complexity while delivering a higher level of availability. If the SLA does not require this level of availability, the aforementioned scenarios may suffice.

Note: It is still recommended to backup at least 1 Platform Service Controller per SSO domain. Restore is not required unless all Platform Service Controller nodes are unavailable. See Scenario 2 — All PSCs Unavailable for more details.

Scenario 5 - vCenter High Availability

vCenter High Availability (vCenter HA) protects vCenter Server Appliance against host and hardware failures. The active-passive architecture of the solution can also help you reduce downtime significantly when you patch vCenter Server Appliance.

A vCenter HA cluster consists of three vCenter Server Appliance instances. The first instance, initially used as the Active node, is cloned twice to a Passive node and to a Witness node. Together, the three nodes provide an active-passive failover solution.

Deploying each of the nodes on a different ESXi hosts protects against hardware failure. Adding the three ESXi hosts to a DRS cluster can further protect the environment.

Requirements include:

- Minimum of 3 ESXi hosts
- The network latency between the Active, Passive, and Witness nodes should be less than 10 milliseconds

Constraints include:

- The vCenter HA network IP addresses for the Active, Passive, and Witness nodes must be static
- The vCenter HA network must be on a different subnet than the management network. The three nodes can be on the same subnet or on different subnets
- The default gateway entry must not be added for the cluster network

Use the basic mode, unless the vCenter Server node configured for vCenter HA is managed by another vCenter Server in a different Single Sign-On domain, in that case use the Advanced mode. The Advanced mode requires manual intervention of cloning and node placement.

By default basic mode will auto recommend placement if DRS is available this include anti-infinity rules (placement of the nodes on separate hosts). If DRS is not available or Advance mode is used manual intervention will be required.

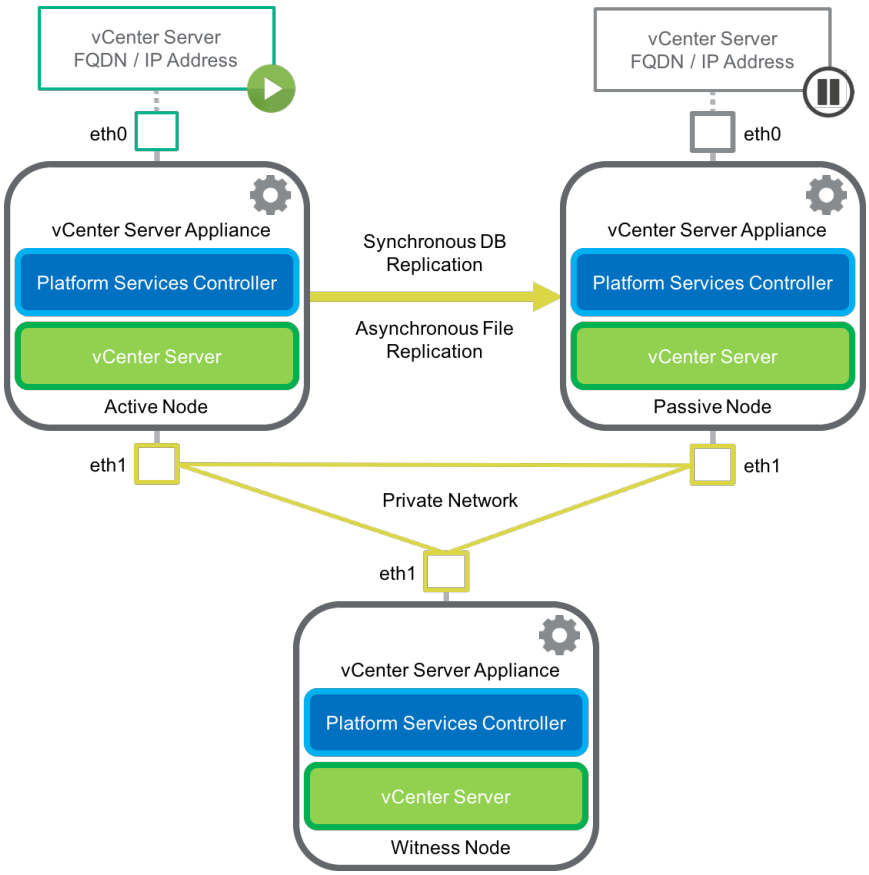

When vCenter HA configuration is complete, only the Active node has an active management interface (public IP). The three nodes communicate over a private network called vCenter HA network that is set up as part of configuration. The Active node and the Passive node are continuously replicating data. This data includes:

- Synchronous database replication
- Asynchronous file replication

In this architecture, only the Active node needs to be backed up. Restore should only occur in the event of catastrophic failure.

In order to restore a vCenter HA cluster:

1. Power off and delete the cluster nodes
2. Restore the active vCenter Server node from backup
3. Reconfigure vCenter HA

STORAGE ARRAY INTEGRATION

A Rubrik cluster can integrate with a storage array to further reduce the time that a VM is quiescent during a snapshot operation. To qualify for storage array integration, all of the datastores that are assigned to the VM must reside on storage arrays.

Normally, a Rubrik cluster ingests the VMDK files of a VM as part of the snapshot process. During this time the VM must be kept quiescent. A Rubrik cluster ingests the VMDK files very quickly, resulting in extremely short periods of quiescence. However, for large VMDK files, the time that is required for ingesting them can impact the VM.

With storage array integration, a Rubrik cluster can use the API of the storage array to move ingestion of the VMDK files out of the vSphere environment and onto the storage array. Using storage array integration, a Rubrik cluster can release a VM for normal operation immediately after a hypervisor snapshot. The Rubrik cluster takes storage array snapshots and uses those for ingestion of the VMDK files.

After releasing the VM, the Rubrik cluster mounts the storage array-level snapshots as temporary datastores on the VM host. The Rubrik cluster then attaches the VMDK files from the temporary datastores to an automatically created temporary proxy VM. The Rubrik cluster completes the data ingestion through the proxy and removes the temporary datastore objects and the proxy.

Storage array integration can employ custom scripts running on the guest operating system to provide application-level quiescence or application consistency. A pre-backup script can prepare an application for the brief quiescence, and a post-snap script can resume the application immediately after the snapshot.

USE CASES

This section is intended to provide a few examples of how Rubrik and VMware vSphere may be used together. It is not intended to be an exhaustive list but merely a set of sample use cases.

AUTOMATED DATA PROTECTION

Legacy backup systems require the creation of backup jobs and manual scheduling of each workload. This approach is time consuming and cumbersome because it requires strategizing when all critical servers should be backed up within the daily backup/archive windows. Additionally, many IT teams are still restricted, only allowing backups during non-business times. This typically occurs between 6:00 pm and 6:00 am. This 12-hour window begins to shrink as more systems that need to be protected are continually added.

Furthermore, this task of manually building backup jobs becomes a juggling act as critical workloads are balanced against workload and business requirements. It is unlikely that an SLA can be met if the workload must be recovered since the data is being protected only during this time period. What if the backup infrastructure is already resource constrained and the backup jobs are exceeding the protection window and going into business operating hours?

Dare to imagine never again having to define backup jobs, selecting media agents, configuring backup storage settings, or creating backup chains. Eliminate spending days or weeks spent mapping out backup jobs, designing and deploying the required infrastructure. Rubrik removes one of the most inefficient and storage wasting tasks in Backup & Recovery — scheduling periodic full backups.

Rubrik eliminates the imperative approach and replaces it with declarative, taming the circus of manually juggling backup tasks with simple data management by defining SLA protection policies. An SLA policy provides the choice of Rubrik snapshots frequency and the length of retention, effectively translating business logic into an automated task. Rather than manually creating a policy and applying per workload, an SLA policy may be applied at a broad level - such as the management server (vCenter Server), folder, host, cluster, etc. - or granularly (per VM) to achieve specific data protection objectives.

The automatic protection mechanism simplifies assigning protection to large numbers of VMs and provides an easy method to uniformly assign specific SLA Domains to groups of functionally similar VMs.

This paper uses a VMware vSphere environment to illustrate Rubrik automating protection, although Rubrik protects far more than just VMware vSphere workloads (including other hypervisors such as Nutanix AHV and Microsoft Hyper-V, physical servers such as Windows and Linux, NAS, and databases such as Microsoft SQL and Oracle). Rubrik's CDM platform accesses VM data through an API connection with the VMware vCenter Server that manages the hypervisor running the VM.

Rubrik CDM provides Auto Protect of VMs through inheritance of the SLA policy assigned to a parent object. If the vCenter Server or a folder has an SLA assigned to it, the VM underneath will automatically inherit the policy. Rubrik detects the newly-added VM and automatically applies a protection policy, eliminating the need for any manual administrator interaction. This resolves the common issue of new workloads being brought online and going days or weeks without being protected.

Once the policy is assigned, Rubrik adheres to user-defined policies such as frequency, retention, archival, etc. as described above. All manual configuration is eliminated by leveraging the Distributed Task Framework and the Blob Engine, which apply intelligent algorithms to ensure efficiency and performance for the entire backup workload. These intelligent algorithms assist with balancing the workload as more VMs are created and added into the system. The automatic scheduling of tasks ensures that all workloads are evenly distributed across the Rubrik cluster, preventing cluster resource contention.

RECOVERY

For a Rubrik cluster, recovery of a source VM means to mount a point-in-time copy of it.

A VM can be recovered by using any of the Rubrik data protection objects: snapshots, replicas, and archival snapshots. Recover a VM by using one of the available recovery actions. The Rubrik cluster provides the following recovery actions for VMs:

- Instant Recovery
- Live Mount
- Export

Action	Name of recovered VM	Datastore	Power state	Network	Source VM
Instant Recovery	Assigned the name of the source virtual machine	Local Rubrik cluster	On	Connected (Optional)	Powered off and renamed
Live Mount	Composite[1]	Local Rubrik cluster	On	Disconnected	No impact
Export	Composite	Datastore of hypervisor	On	Disconnected	No impact

[1] The name of the recovered virtual machine is constructed as follows: name of source virtual machine + timestamp of snapshot + incremented integer. For example, the first mount of the snapshot of the virtual machine "Rebecca01" that was created at "08-04 06:48" is named "Rebecca01 08-04 06:48 1".

FILE OR FOLDER-LEVEL RECOVERY

The Rubrik cluster provides file-level restore (FLR) of files and folders from any local snapshot, replica, or archival snapshot that was successfully indexed.

To restore a file or folder, search for it by name across all local snapshots or browse for it on a selected snapshot.

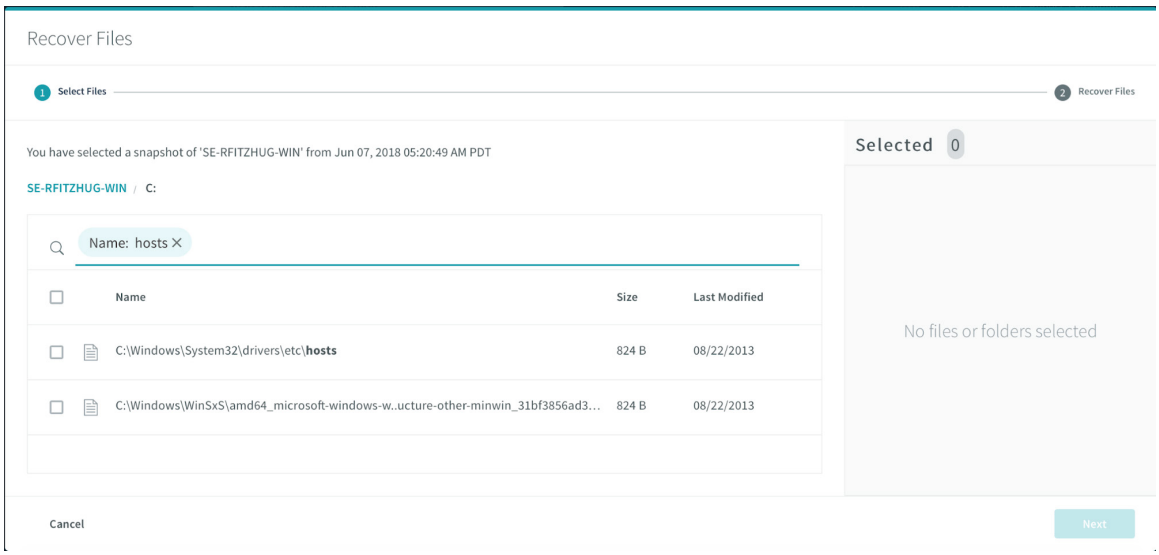

For supported Windows and Linux guest operating systems, the Rubrik cluster can restore files and folders directly to the source file system.

When restoring from a snapshot of a supported guest operating system, the UI provides the option to restore a file or folder directly to the source file system. When this option is selected, the UI provides a choice to overwrite the source file or folder, or to restore the file or folder to another location.

A restored file or folder inherits the ACL of the parent folder and the same owner as the parent folder.

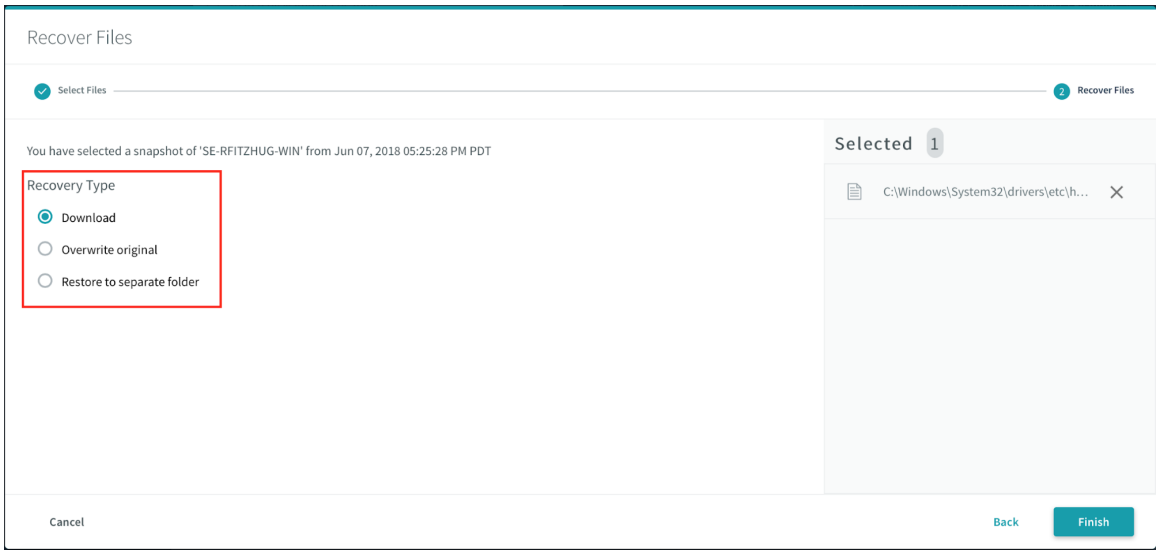

Files and folders may also be downloaded through the UI. The Rubrik cluster processes the request and provides a link for download of the file. If a folder is selected, then a `.ZIP` file containing the folder and all that the folder contains. The `.ZIP` file retains the hierarchy of the selected folder. This functionality is useful for administrators who want to locally view a file or folder at a specific point in time rather than restoring back to the original VM.

APPLICATION FAILURE

Three-tier applications can be found in nearly every data center. It is endlessly frustrating to wait for a full database restore when a DBA only needs to query or recover a few tables.

Given an immediate Live Mount, a DBA could do a selective restore of specific rows and tables via a simple ad hoc query, export, and then import. This would provide far faster and granular database restore capabilities. Similarly, a database could easily be examined to track when specific data changed; there could even be multiple Live Mounts created to iterate back in time and pinpoint the exact time of a change. In a traditional backup infrastructure, this may not be possible whether due to long restores or not having extra disk space available for ad hoc requests.

Ad hoc scenarios are specifically enabled by the immediacy of Live Mount. Rubrik customers have found many variations on this theme based on the ease and speed of bringing a Live Mount online.

Additionally, Instant Recovery functionality may be used to recover a VM or set of VMs, providing a near-zero RTO.

If an entire application stack were to fail, it may be recovered in a similar manner. Simply choose whether to instantly recover the VMs using Rubrik as storage or to fully rehydrate the VMs back to their primary storage using the Export option.

PRIMARY STORAGE FAILURE

In the event that primary storage fails, critical VMs may be instantly recovered or Live Mounted using Rubrik as the temporary storage. Once the primary storage is back online, the VMs may be simply Storage vMotioned back to the desired datastores.

If secondary NAS or SAN is online, the VMs may be recovered via export to designate another storage location.

CLOUDOUT (ARCHIVAL)

CloudOut is a capability within Rubrik CDM used to archive data to the cloud for short and long-term retention. Users may leverage Rubrik to intelligently and cost-effectively store backup data in Amazon S3, Microsoft Azure Blob storage, or Google Cloud Storage. More importantly, Rubrik is optimized to provide rapid and efficient data restores both on-premises and in the public cloud. Data is indexed by Rubrik CDM before it is stored in the cloud archive, enabling customers to quickly browse for, search for, and restore any file. During restores, Rubrik only retrieves the specific files that need to be recovered to minimize bandwidth and egress costs.

Rubrik customers typically leverage CloudOut as a solution to replace their tape storage infrastructure, eliminating the need to copy backup data to tapes which would then need to be manually stored offsite. Rubrik with cloud archive provides a tape-replacement solution that is more durable, available, cost-effective, and agile.

If on-premises archive solutions are preferred, Rubrik also supports NFS, tape, and object storage.

CLOUDON (INSTANTIATION)

CloudOn, or instantiation, allows users to migrate existing on-premises workloads to the cloud for test/development or even disaster recovery purposes. Rubrik's CloudOn feature converts workload data (VM image) sent to the cloud into a compute instance. There is no need to run Rubrik in the cloud to migrate workloads to the cloud for test/dev, increasing overall cloud savings.

Better yet, imagine not needing a separate cluster for test/development workloads or an identical physical infrastructure for disaster recovery. Using Rubrik CloudOn, workloads can be migrated at a VM level from on-premises to AWS or Azure.

Rubrik offers three options that can be applied to on-premises workloads that customers choose to instantiate in AWS or Azure:

- **On-Demand –** The default configuration in which Amazon Machine Images (AMIs) or Azure Virtual Hard Drives (VHDs) are created only at the time of a "power on in the cloud" request.
- **Auto Convert Latest Snapshot - Keep One –** Rubrik will automatically construct an AMI or VHD reflecting the latest snapshot to be archived into S3 or Azure. When a new snapshot is sent to the archive, a new AMI or VHD is constructed with the new archive data. Once completed, the older AMI or VHD is removed.
- **Auto Convert Latest Snapshot - Keep All –** Rubrik will automatically construct an AMI or VHD reflecting the latest snapshot to be archived into S3 or Azure. When a new snapshot is sent to the archive, a new AMI or VHD is constructed with the new archive data. The older AMI or VHD is retained if desired (configurable via policy), creating a series of AMIs or VHDs representing each snapshot.

The following screenshot demonstrates the required information to instantiate a workload in AWS:

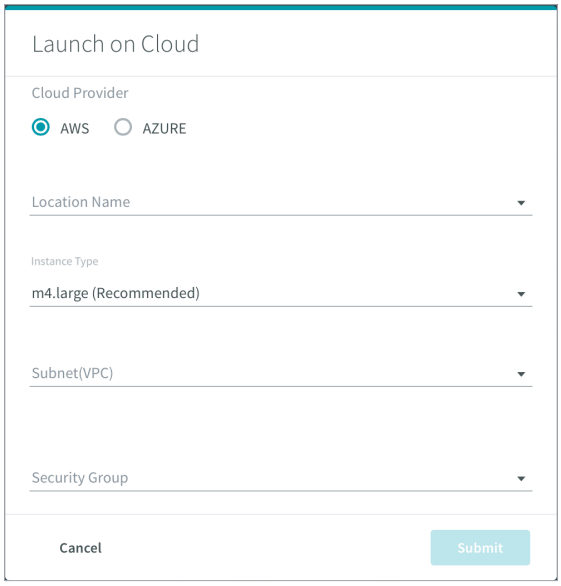

Whether instantiating workloads on-demand or automatically with the latest snapshot, spinning up copies of workloads in the cloud results in faster development cycles as developers are unblocked from the constraints of physical infrastructure. Picture the cost savings garnered when avoiding a dedicated on-premises infrastructure for test/development. Developers can spin up instances when required and shut down when not in use.

TEST AND DEVELOPMENT

Upgrades can be scary! What if it was possible to instantiate an environment that looked exactly like production (at the time of snapshot) automatically through APIs and push code into that environment to test? With Rubrik, that possibility is reality.

Build isolated environments and leverage the Live Mount feature to instantiate an identical environment in moments. Additionally, VMs may be instantiated using CloudOn in AWS or Azure from any point-in-time without modifying the underlying immutable data. Test a VMware Tools or hardware version upgrade, failure scenario, or other use cases using your backup storage. When done, simply throw it away.

Rather than tinkering with a production VM, use Live Mount to spin up a copy from a specific point-in-time to modify. As an example, Active Directory work, such as rehearsing an upgrade to the Forest Level, may be done during the day in a bubble network.

This functionality appeals to application owners and operations teams in order to conduct:

- Functional or regression testing
- Application development
- Software release testing (upgrade the actual applications)

No additional configurations are needed on the hypervisor side for Live Mount functionality to work. Simply provide a service account with the documented permissions required for the type of virtualization environment. Rubrik automates the entire process. VMs of any size can be recovered in the amount of time it takes for the OS to boot. Imagine having the ability to spin up an 8 TB VM in under 2 minutes so that a recovery point can be validated by an administrator or application owner.

VSPHERE UPGRADES

With the new releases of vSphere, along with the ending of support for older vSphere versions approaching, many customers are planning upgrades as part of their data center modernization efforts. Some view this as an opportunity to start with a fresh slate and build a new environment in accordance with modern architectural best practices and enterprise data management needs. In these cases, customers must "swing" over their VMs from the old environment to the new environment with as little disruption as possible.

Rubrik's VM linking capability enables users to seamlessly execute these swing migrations by deploying a new vSphere environment alongside their existing setup. They can then share the primary VM storage to both the existing and new deployments, which provides the new vSphere environment access to the VM data without requiring a data migration to take place. As customers migrate VMs to the new environment and register them with the new vCenter deployment, the Rubrik system automatically links the VMs to their previous identities across vCenter boundaries. This linking lets Rubrik preserve the historical snapshot information and any SLA domains directly assigned to their VM, minimizing operational disruption.

CONCLUSION

Rubrik's support for VMware vSphere protection is robust and full-featured while extending Rubrik's market leading focus on simplicity. Using Rubrik and VMware vSphere together helps accelerate companies on their journey to meet hybrid cloud business requirements by protecting on-premises workloads, providing archival and replication to public cloud, and the ability to instantiate vSphere workloads in AWS or Azure.

ABOUT THE AUTHORS

Chris Wahl is the author of Wahl Network and host of the Datanauts Podcast. In addition to co-authoring "Networking for VMware Administrators" for VMware Press, Chris also travels globally to speak at industry events, provide subject matter expertise, and offer perspectives to startups and investors as a technical adviser. You can find him on Twitter @ChrisWahl.

Rebecca Fitzhugh is a Technical Marketing Engineer at Rubrik. She is VCDX #243, a published author, and blogger. You can find her on Twitter @RebeccaFitzhugh.

Emad Younis is a Staff Technical Marketing Architect and VCIX 6.5-DCV working in the Cloud Platform Business Unit, part of the R&D organization at VMware. He currently focuses on the vCenter Server Appliance, vCenter Server Migrations, and VMware Cloud on AWS. His responsibilities include generating content, evangelism, collecting product feedback, and presenting at events. Emad can be found blogging on emadyounis.com or on Twitter via @emad_younis.

APPENDIX A: PORT DIAGRAM

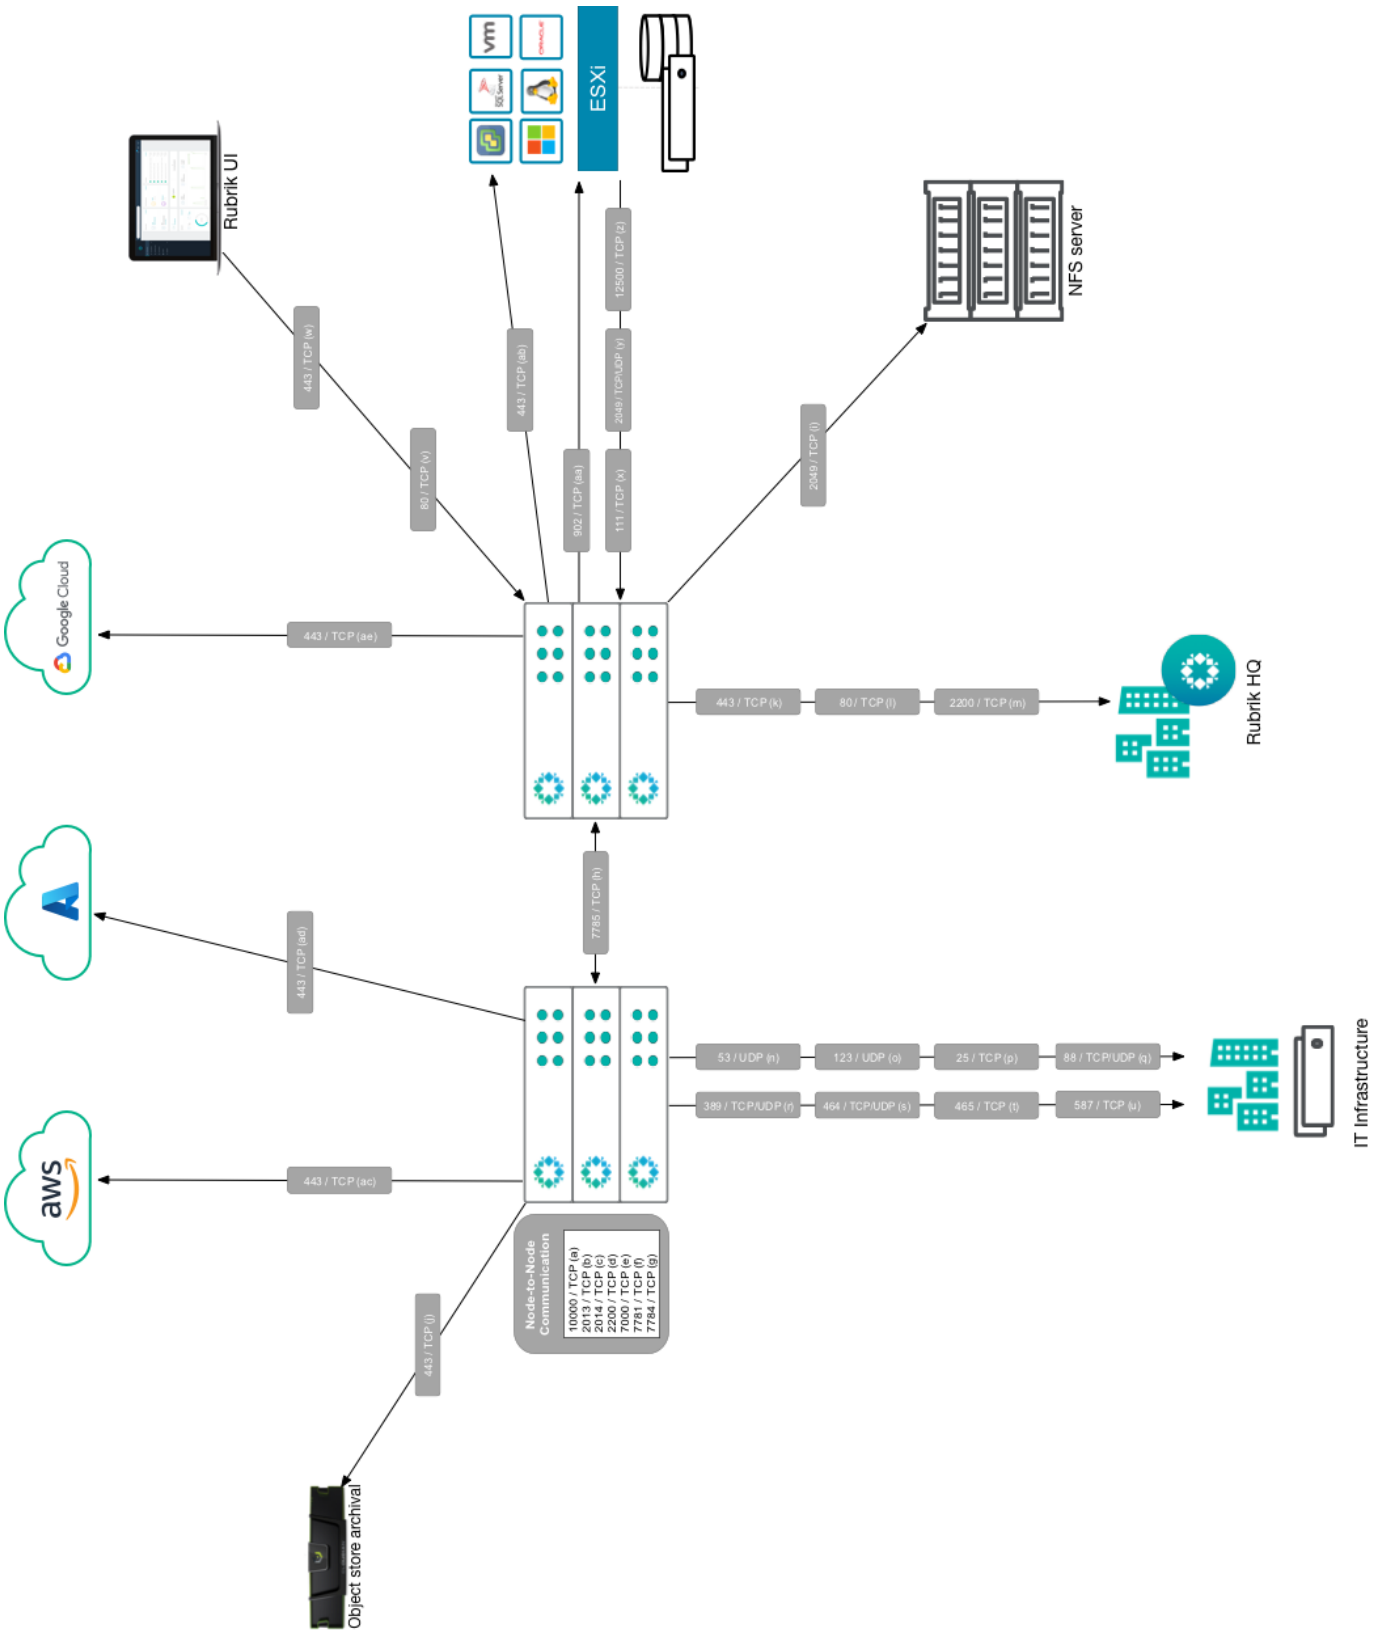

ID	Port	Protocol	Source	Destination	Description
a	10000	TCP	Rubrik cluster	Rubrik cluster	Allows sharing of Rubrik cluster file system (SDFS) data between the nodes of a Rubrik cluster
b	2013	TCP	Rubrik cluster	Rubrik cluster	Allows sharing of statistics between the nodes of a Rubrik cluster
c	2014	TCP	Rubrik cluster	Rubrik cluster	Allows sharing of statistics between the nodes of a Rubrik cluster
d	2200	TCP	Rubrik cluster	Rubrik cluster	Allows node to node SSH communication during upgrade
e	7000	TCP	Rubrik cluster	Rubrik cluster	Allows process arbitration between the nodes of a Rubrik cluster
f	7781	TCP	Rubrik cluster	Rubrik cluster	Permits the Rubrik cluster to load basic software and configuration information (bootstrap) during cluster configuration
g	7784	TCP	Rubrik cluster	Rubrik cluster	TLS over TCP communication between nodes within a Rubrik cluster
h	7785	TCP	a. Replication source b. Replication target	a. Replication target b. Replication source	Replication data transmission
i	2049	TCP	Rubrik cluster	NFS server	Permits communication with a NAS device that is being used as an archival location
j	443	TCP	Rubrik cluster	Object store archive	Transmitting data to the archival location.
k	443	TCP	Rubrik cluster	logs.rubrik.com	Error logs for support
l	80	TCP	Rubrik cluster	live.rubrik.com	Stats upload
m	2200	TCP	Rubrik cluster	supporttunnel. rubrik.com	Support Tunnel
n	53	UDP	Rubrik cluster	DNS server	Permits hostname resolution
o	123	UDP	Rubrik cluster	NTP server	Provides access to NTP servers for time synchronization
p	25	TCP	Rubrik cluster	Email server	Allows the Rubrik cluster to send email alerts to administrators (if email servers support port)
q	88	TCP/UDP	Rubrik cluster	Active Directory server	Permits Kerberos communication
r	389	TCP/UDP	Rubrik cluster	Active Directory server	Permits LDAP communication
s	464	TCP/UDP	Rubrik cluster	Active Directory server	Permits Kerberos password set/change communication
t	465	TCP	Rubrik cluster	Email server	Allows the Rubrik cluster to send email alerts to administrators (if email servers support port)
u	587	TCP	Rubrik cluster	Email server	Allows the Rubrik cluster to send email alerts to administrators (if email servers support port)
v	80	TCP	Web UI clients	Rubrik cluster	Handles redirection of web UI clients to HTTPS
w	443	TCP	Web UI clients	Rubrik cluster	HTTPS interface
x	111	TCP	VMware ESXi hosts	Rubrik cluster	Provides an NFS datastore for ESXi hosts.
y	2049	TCP/UDP	VMware ESXi hosts	Rubrik cluster	Permits contact with the NFS daemon running on the Rubrik cluster for Live Mount operations.
z	12500	TCP	VMware ESXi hosts	Rubrik cluster	For Live Mount, allows an ESXi host to perform an NFS mount on this port to acquire a virtual machine

ID	Port	Protocol	Source	Destination	Description
aa	902	TCP	Rubrik cluster	VMware ESXi hosts	Permits network block device (NBD) data transfers
ab	443	TCP	Rubrik cluster	VMware vCenter Server	Information queries about virtual machines.
ac	443	TCP	Rubrik cluster	AWS	Required for all data protection, instantiation, and archival operations
ak	443	TCP	Rubrik cluster	Azure	Required for all data protection, instantiation, and archival operations
al	443	TCP	Rubrik cluster	Google Cloud	Required for archival operations

APPENDIX B: VCENTER PRIVILEGES

The following table describes the minimum privileges on the vCenter Server that are required by the vCenter Server role that is assigned to the Rubrik cluster. The table uses an asterisk (*) to indicate a privilege that Rubrik does not require in the current release but anticipates requiring in a later release.

Privilege Category	Privilege	Description
Datastore	Allocate space	Used by Rubrik to create virtual machines for export. Also used by Rubrik to provide space for delta files on the datastore when creating a snapshot.
Datastore	Browse datastore	Permits Rubrik to find and download the vmware.log file for a virtual machine after a failed snapshot and to send the vmware.log file out for support.
Datastore	Configure datastore	Allows Rubrik to connect the datastore on a Rubrik cluster to the vCenter Server for Live Mount and Instant Recovery.
Datastore	Low level file operations	Permits Rubrik to ingest and to export the contents of snapshot VMDKs.
Datastore	Move datastore*	Allows Rubrik to place a Live Mount datastore into a vCenter Server folder to enhance manageability.
Datastore	Remove datastore	Used by Rubrik to detach a Live Mount datastore that is no longer in use.
Global	Enable methods	Permits the Rubrik cluster to provide vSphere APIs for Data Protection license information to the vCenter Server. Required when using vSphere APIs for Data Protection to transfer VMDK contents.
Global	Licenses	Permits the Rubrik cluster to provide vSphere APIs for Data Protection license information to the vCenter Server. Required when using vSphere APIs for Data Protection to transfer VMDK contents.
Host	Configuration: A. Storage partition configuration	Configuration privileges: A. Used by Rubrik for storage partition configuration when attaching Live Mount datastores to ESXi hosts.
Network	Assign network	Permits Rubrik to connect Instant Recovery virtual machines to a network when powering on the virtual machines.
Resource	Assign virtual machine to resource pool	Allows Rubrik to allocate resources on an ESXi host for powering on virtual machines that are created through the Export, Live Mount, and Instant Recovery features.
Sessions	Validate session	Used by Rubrik to discover, cache, and reuse previous vCenter Server sessions.
Sessions	View and stop sessions	Used by Rubrik to discover, cache, and reuse previous vCenter Server sessions.

Privilege Category	Privilege	Description
Virtual Machine	Configuration: A. Add existing disk B. Add new disk C. Advanced D. Change resource E. Disk change tracking F. Disklease G. Rename* H. Settings I. Swapfile placement	Configuration privileges: A. Used by Rubrik when creating virtual machines through the Export, Live Mount, and Instant Recovery features. B. Used by Rubrik when creating virtual machines for the Export, Live Mount, and Instant Recovery features. C. Required for Live Mount, Instant Recovery, and Export. Also permits creation of the proxy virtual machine that is required for storage array integration. D. Permits Rubrik to configure virtual machine resources that are created in resource pools. E. Used by Rubrik to enable incremental snapshots, and to reset CBT when required. F. Allows Rubrik to acquire leases to permit using vSphere APIs for Data Protection for transferring VMDK contents. G. Permits Rubrik to rename the Live Mount datastore to enhance manageability. H. Used by Rubrik to configure virtual machines that are created through the Export, Live Mount, and Instant Recovery features. I. Allows Rubrik to power on virtual machines that are created through the Export, Live Mount, and Instant Recovery features.
Virtual Machine	Guest Operations: A. Guest Operation Modifications B. Guest Operation Program Execution C. Guest Operation Queries	Guest Operations privileges: A. Permits Rubrik to deploy the Rubrik VSS agent into guest operating systems when creating application consistent snapshots. B. Permits Rubrik to start the Rubrik VSS agent on guest operating systems when creating application consistent snapshots. C. Allows Rubrik to monitor and manage the Rubrik VSS agent while the agent is running on guest operating systems.
Virtual Machine	Interaction: A. Answer question* B. Backup operation on virtual machine C. Device connection D. Guest operating system management by VIX API E. Power Off F. PowerOn G. Reset* H. Suspend* I. VMware Tools install*	Interaction privileges: A. Permits Rubrik to automatically handle situations where a virtual machine is in a stuck state waiting for a question to be answered. B. Used by Rubrik to perform backup operations on virtual machines. C. Used by Rubrik to connect and disconnect devices which are attached to virtual machines that are created through the Export, Live Mount, and Instant Recovery features. D. Permits Rubrik to manage a guest operating system along with the Rubrik VSS agent when creating application consistent snapshots. E. Allows Rubrik to power off Live Mount virtual machines and Instant Recovery virtual machines before deleting the virtual machine. F. Allows Rubrik to power on Export virtual machines, Live Mount virtual machines and Instant Recovery virtual machines after creating the virtual machine. G. Permits Rubrik to manage Export virtual machines, Live Mount virtual machines and Instant Recovery virtual machines after creating the virtual machine. H. Permits Rubrik to manage Export virtual machines, Live Mount virtual machines and Instant Recovery virtual machines after creating the virtual machine. I. Allows Rubrik to upgrade VMware Tools on a guest OS as needed to prevent the guest OS from hanging or crashing when quiescing for a snapshot.

Privilege Category	Privilege	Description
Virtual Machine	Inventory: A. Create new B. Move C. Register D. Remove E. Unregister	Inventory privileges: A. Used by Rubrik to create Export virtual machines, Live Mount virtual machines and Instant Recovery virtual machines. B. Permits Rubrik to move an original virtual machine into a "deprecated" folder before replacing the original with an Instant Recovery virtual machine. C. Used by Rubrik to create Export virtual machines, Live Mount virtual machines and Instant Recovery virtual machines. D. Allows Rubrik to remove Export virtual machines, Live Mount virtual machines and Instant Recovery virtual machines. E. Allows Rubrik to remove Export virtual machines, Live Mount virtual machines and Instant Recovery virtual machines
Virtual Machine	Provisioning: A. Allow disk access B. Allow read-only disk access C. Allow virtual machine download D. Allow virtual machine files upload	Provisioning privileges: A. Permits Rubrik to write the VMDK contents of Export virtual machines, Live Mount virtual machines and Instant Recovery virtual machines. B. Permits Rubrik to read the VMDK contents of Export virtual machines, Live Mount virtual machines and Instant Recovery virtual machines when backing up the virtual machines. C. Allows Rubrik to download the non-VMDK files of protected source virtual machines, including configuration files and support logs. D. Allows Rubrik to upload non-VMDK files of Export virtual machines, Live Mount virtual machines and Instant Recovery virtual machines, when creating and configuring the virtual machines.
Virtual Machine	Snapshot management: A. Create snapshot B. Remove snapshot C. Rename snapshot* D. Revert to snapshot	Snapshot management privileges: A. Permits Rubrik to create temporary snapshots of virtual machines for ingest into Rubrik cluster storage. B. Permits Rubrik to remove temporary snapshots of virtual machines that were created for ingest into Rubrik cluster storage. C. Allows Rubrik to manage the temporary snapshots of virtual machines that were created for ingest into Rubrik cluster storage. D. Used by Rubrik to prepare an Export virtual machine with data from a Rubrik snapshot.

VERSION HISTORY

Version	Date	Summary of Changes
1.0	July 2018	Initial Release
1.1	January 2022	Update product naming in line with Winter 2021 Release

rubrik

Global HQ
3495 Deer Creek Road
Palo Alto, CA 94304
United States

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

20220119_v1