TECHNICAL WHITE PAPER

Protecting Epic with Rubrik Security Cloud

Mike Preston, Technical Marketing Architect
Marcus Henderson, Senior Platform Solutions Architect
July 2023
RWP-0621

Table of Contents

ABSTRACT

The significance of Information Technology (IT) within healthcare today cannot be understated. Across the world, healthcare providers have been adopting Epic's Electronic Medical Record (EMR) system to streamline their IT services and ensure efficient and high-quality patient care. Given this growing dependence on Epic, it's of the utmost importance that healthcare providers ensure patient data is protected and adheres to industry regulations.

This paper walks through how Rubrik Security Cloud protects Epic workloads, provides an in-depth overview of backup and restoration processes, and highlights the results of internal performance tests executed.

AUDIENCE

This paper is for anyone who wants to understand how Rubrik Security Cloud protects the InterSystems Caché, IRIS databases, and associated Epic workloads with secure and immutable backups.

WHY HEALTHCARE BACKUPS ARE MORE IMPORTANT THAN EVER

Ransomware attacks are becoming more sophisticated and cybercriminals are specifically targeting healthcare providers. According to the Wall Street Journal, the Ryuk ransomware gang has targeted at least 235 hospitals and collected over $100 million in ransoms since 2018. However, the true cost of ransomware isn't just the ransom paid but the downtime associated with an attack that affects critical patient care and presents concerns with patient safety. Without proper functioning and secure Electronic Medical Records (EMR) and Electronic Health Record (EHR) systems, healthcare organizations risk reduced access to medical records, which could cause surgical delays or cancellations, inaccurate medical decisions, or even the possibility of breaching regulatory requirements.

Since backups are the last line of defense and the first line of recovery, it's more important than ever for healthcare providers to leverage a backup solution with fast and efficient recoverability in the event of an attack. Rubrik Security Cloud is designed with a zero-trust security model to minimize the surface area of any attack through a variety of features such as:

- An immutable file system - where backups are stored and secured, was designed into the product since inception providing:
  - No requirement for external storage with complex configurations and dependencies.
  - A custom-built filesystem, built from scratch to store and manage versioned data with built-in fault tolerance, linear scalability, zero-byte cloning, and self-healing capabilities.
  - Native Immutability - Once backups are written to the Rubrik they cannot be tampered with or deleted. Even incremental backups are written in an append-only fashion, meaning original data integrity is always maintained.
  - Backups are not exposed over the network through common protocols and cannot be accessed by unauthorized external applications, providing a logical air gap between production and backup environments.

- A hardened and secure Linux-based OS with
  - No Windows-based components
  - Unnecessary and unused services have been removed
  - Firewalls are enabled and ports are restricted to only necessary services
  - Monotonic clock prevents NTP poisoning
- No third-party applications are allowed to run on the Rubrik cluster
  - All features running on the cluster certified by Rubrik
  - Third-party applications dramatically increase the threat vulnerability of the backup platform
- Encryption
  - Data at Rest Encryption
  - Intra-cluster communication is fully encrypted
  - Archive and replication data is fully encrypted

Ultimately, users will need to manage the Rubrik cluster. Rubrik has several features to harden security against both internal and external attacks attempting to compromise credentials:

- Single Sign On (SSO) with SAML 2.0 Support
- Native Multifactor Authentication (MFA) support with Time-based One-Time Password (TOTP) or RSA SecureID
- Role-based Access Control (RBAC) allowing for granular and custom roles to be assigned to end-users, ensuring only the required functionality required to manage their duties is granted.
- Retention Locked SLAs requiring third-party verification through Rubrik for any changes to policies which would reduce the retention of backups.
- Legal Hold provides the ability to lock point-in-time snapshots so they cannot be modified, expired, or deleted.

PROTECTING EPIC WITH RUBRIK SECURITY CLOUD

Rubrik Security Cloud provides robust data protection solutions for the workloads contained within the Epic software suite including:

- **Virtual Machines –** RSC is able to backup the hundreds of VMs that are normally required to run various Epic applications through integrated support for VMware vSphere, Microsoft Hyper-V and Nutanix AHV. RSC also supports cloud-native workloads running within Amazon AWS, Microsoft Azure and Google Cloud.
- **Intersystems Cache or IRIS databases –** Utilizing the fileset technology of Rubrik Security Cloud, healthcare providers can easily protect their central transactional databases hosting patient records.
- **Clarity and Caboodle (CDW) databases on SQL or Oracle –** Epic's data warehousing solutions running on either MSSQL or Oracle can be protected with Rubrik Security Cloud

- **Blob Tier or Objects –** Epic's blob tier hosting on NAS systems can be protected by Rubrik Security Cloud by either backing up directories to the Rubrik platform, directly archiving to another storage system using NAS-DA, or directly to cloud utilizing NAS-CD.

The remainder of this paper focuses on how to protect the Intersystem Cache or IRIS database hosting patient records.

EPIC COMPONENTS

The following lists the required Epic components needed to allow Rubrik Security Cloud to provide data protection capabilities.

Backup Proxy

Epic recommends an off-production host, often called a backup proxy to protect the Cache or IRIS database. This ensures minimal impact on the production environment during the backup process. The backup proxy can be a physical host or a virtual machine when running Cache/IRIS on Red Hat Enterprise Linux (RHEL) and AIX.

Freeze/Thaw Script

Scripts are often used as an automated way of ensuring databases are properly flushed to achieve consistency during the backup process and mount data to be processed to the backup proxy.

The typical script to freeze and thaw the Cache/IRIS databases is as follows:

1. The file system hosting the Cache/IRIS database is unmounted from the backup proxy.
2. The Cache/IRIS database is frozen, flushing all cached content to disk and restricting permitted operations to read commands. In the case of AIX, the JFS2 filesystem is also frozen.
3. A snapshot or clone is executed on the LUNs on the storage array hosting the Cache/IRIS database.
4. The Cache/IRIS database is released or thawed, allowing normal operations to resume to the production database. In the case of AIX, the JFS2 filesystem will be released beforehand as well.
5. The storage array snapshots of the Cache/IRIS database are remounted to the backup proxies filesystem.
6. Once the data has been remounted to the backup proxy, Rubrik can begin the backup process.

RUBRIK COMPONENTS

Rubrik Backup Service

The Rubrik Backup Service (RBS) must be installed on the backup proxy host before backups can occur. RBS is a lightweight connector that runs as a service and is installed as a .rpm on both RHEL and AIX. Some features of RBS include:

- **No required reboots –** RBS can be installed on the backup proxy without the need for any reboots.
- **Small production footprint –** RBS runs as a lightweight service requiring roughly 10 MB of disk space and 100MB of memory.
- **Automatic Upgrades –** As upgrades are performed on the Rubrik environment, RBS will automatically upgrade, eliminating the need for lengthy and complex agent management.

- **Secure Communication –** When installing RBS, security certificates ensure secure communications between the host running RBS and the Rubrik cluster it was downloaded from.

After RBS is installed, the host is added via the Rubrik UI to manage data protection and restores.

SLA Domains

Rubrik orchestrates the movement of data throughout the data lifecycle. From initial ingest to propagation of that data to other locations, such as replicating to remote clusters or Rubrik Cloud Cluster, as well as data archival. A single SLA policy dictates all data lifecycle specifications, and the data control plane of Rubrik does the rest. Furthermore, customers can create as many SLA Domains as desired.

An example SLA policy is as follows:

- 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

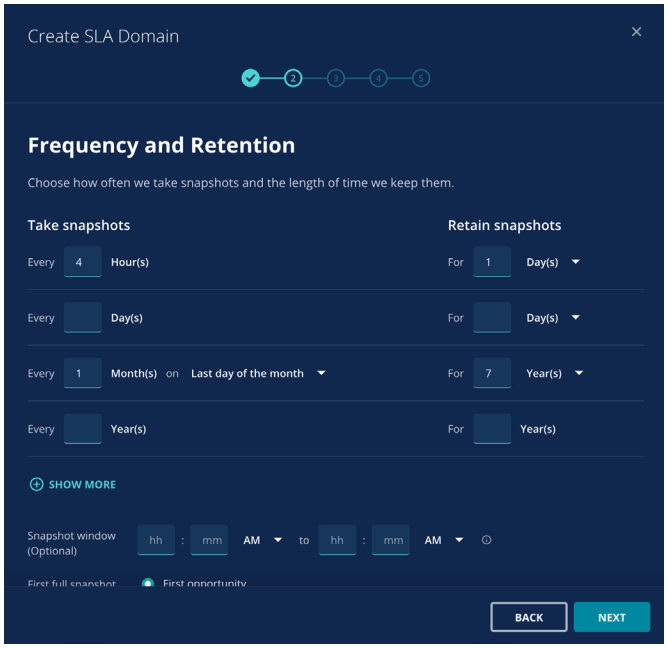

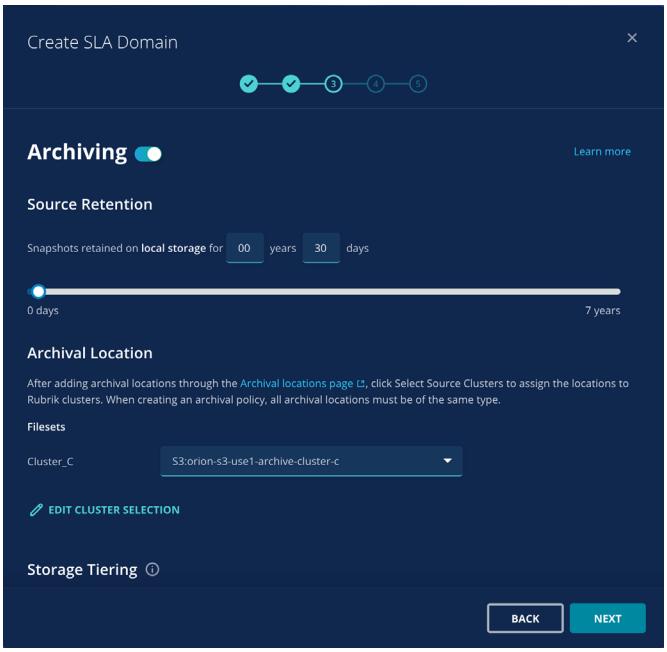

Data is ingested and retained according to the frequency specified in the SLA policy. The example policy is configured to store 30 days of data within the Rubrik cluster. Once that period has elapsed, data is archived to another location for long-term retention. In this case, data is archived to Amazon 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 can also specify replicating 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. This capability mitigates the need to configure and manage this functionality at the storage layer. Apply policy-based management to workloads and stop babysitting data residing across multiple data centers.

Note: Rubrik provides three built-in SLA Domains by default -- each representing a set level of protection:

- Gold (highest protection)
- Silver (medium protection)
- Bronze (lowest protection)

Administrators may choose to use the built-in SLA Domains or to create additional SLA Domains.

From the data archive locations, Rubrik also provides simplified data retrieval with predictive search capabilities. 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 the archive before restoring. This provides the ability to recover archived data at a snapshot or file-level selectively without downloading the entire workload to restore a single file, or racking up significant egress charges.

Filesets

Rubrik protects files and folders on host systems by using filesets. A fileset defines a set of files and folders on a host system or NAS share. Rubrik uses these defined filesets assigned to the hosts or NAS shares to determine what data to manage and protect.

When defining filesets within Rubrik, administrators can provide a listing of values that represent the files to process. This is provided by configuring files/folders to Include, files/folders to exclude, as well as specific overriding rules for certain files to never exclude. These values can be defined by providing a full path to the file(s) and folder(s), segments of the paths utilizing wildcards, as well as portions of the filenames themselves utilizing wildcards. Values are provided through the use of comma-separated segments. Filesets are not host-specific, meaning the same fileset can be attached to multiple hosts along with multiple filesets attached to the same host.

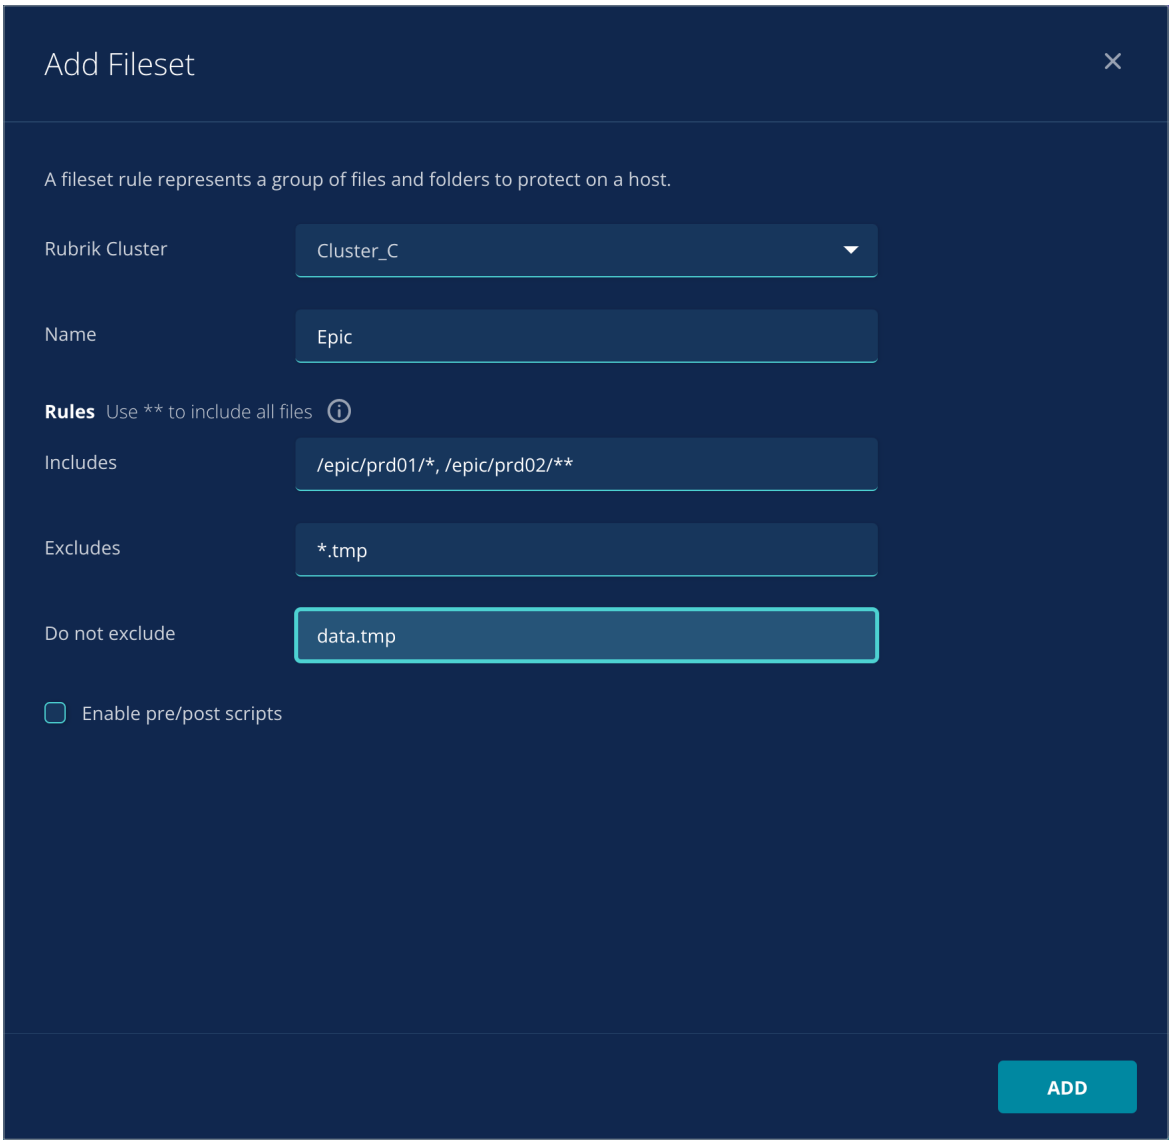

For example, the image above dictates a fileset which abides by the following rules:

- /epic/prd01/ and all files within the directory will be included within the backup, however any subdirectories within the directory will not
- /epic/prd02 and all files, including subdirectories, will be included within the backup
- Any file with the extension of .tmp will be excluded from the backup.
- The data.tmp file is the only exception as it is explicitly configured within the Do not exclude section.

HIGH LEVEL OVERVIEW OF BACKING UP CACHE/IRIS WITH RUBRIK SECURITY CLOUD

Backing up Cache/IRIS begins with creating an SLA domain and a fileset. The SLA Domain dictates how often to backup, along with how long to retain those backups on the Rubrik platform, and several other data protection constructs such as archival and replication settings.

From there, a fileset is created, specifying the path to the Cache/IRIS database files on the backup proxy. These files should be marked as "include" when creating the fileset. After installing RBS on the backup proxy and adding the backup proxy to RSC, the SLA Domain can be assigned to the host containing the target data.

The next step is to ensure that the freeze/thaw scripts on the backup proxy are executed before the backup begins. Administrators have a few options for executing these scripts:

1. The scripts can be executed by the backup proxy itself, ensuring that the refreshed Cache/IRIS database files are mounted to the backup proxy at a specified time every day. The SLA Domain can then be configured to run at a time after that has occurred

    a. For example, if the freeze/thaw scripts run every night at 8:30 PM, the SLA domain can be configured to kick off at 9:00 PM daily automatically. Compliance-based reporting is also available to ensure backups adhere to the data protection constructs set forth within the policies. This is the simplest way to schedule backups.

2. Freeze/Thaw scripts can be integrated into the SLA Domains rather than scheduled outside of Rubrik Security Cloud. During the creation of the SLA/Fileset, administrators can configure the execution of pre and post-backup scripts that will be executed before and after the backup begins/ends. The scripts remain on the backup host; however, before Rubrik begins a backup, they will be executed to ensure that the Cache/IRIS data is refreshed and available on the backup proxy before Rubrik moves any data.

3. On-Demand backups can be triggered within freeze/thaw scripts directly from the backup proxy through API calls. Similar to option 2, backups will occur immediately after the database files have been remounted, however, this option allows administrators to centralize all backup tasks within the scripts themselves. Rubrik Security Cloud provides a robust, easy-to-use API implementation, as well as various Software Development Kits (SDKs), allowing for automation tasks such as this. Administrators can easily integrate data protection tasks within Rubrik into their freeze/thaw process to trigger on-demand backups of their Cache/IRIS databases. For a more detailed overview of utilizing the GraphQL API see Appendix A.

Once the backup completes, Rubrik will then perform several additional processes depending on the constructs set forth within the SLA Domain, such as:

- Indexing of the backup data to enable efficient search across backups
- Archival of older point-in-time backups to a different location
- Replication of most recent backups to a different location
- Various scans of data to support Anomaly Detection and Sensitive Data Monitoring

A more detailed description of configuring Rubrik to backup Cache/IRIS can be found in Appendix B.

HIGH LEVEL OVERVIEW OF RESTORING CACHE/IRIS WITH RUBRIK SECURITY CLOUD

Rubrik supports flexible restore options, either to the same host or to an alternate host of the same OS type that has been added to the Rubrik cluster. This allows you to restore to a test or development system to perform periodic Epic restore tests or to seed a new Epic environment. After an environment is restored, you can run an integrity check to verify the database files.

To perform a restore, navigate to the host and fileset that you want to restore from. The recovery calendar will show a dot for the days that there are valid recovery points. Simply, select the date that you want to restore, and then click “Recover Files” for the time that you want to perform the restore from:

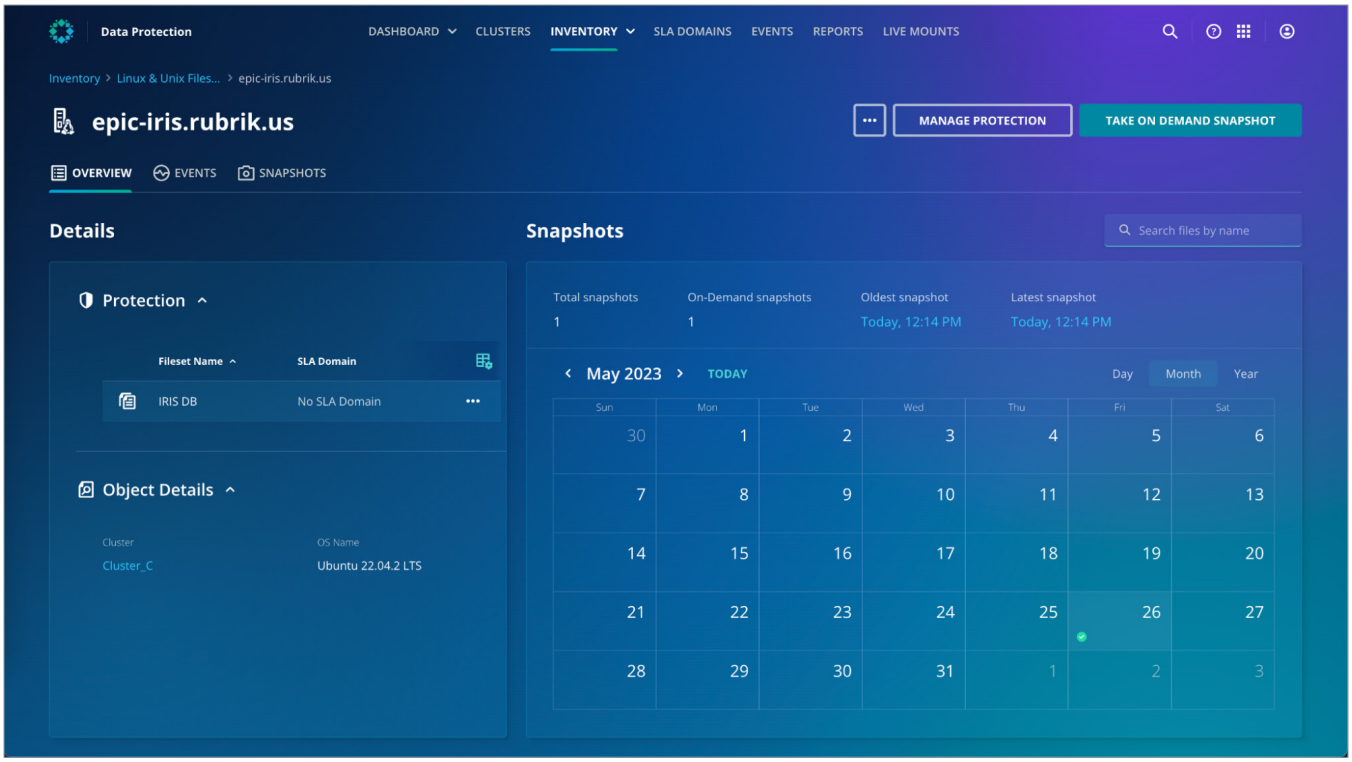

After selecting the desired files to be restored, the following recovery options are available:

- **Download –** The requested files are bundled into a zip file, and a download link is generated, allowing end users to download the data to any location. This option is not often utilized when restoring Epic.
- **Overwrite original –** The requested files are restored in-place to the original location on the host from which they were backed up. This option will overwrite any existing files located on the original host.
- **Recover to a separate folder –** The requested files are restored to a user-specified folder on the original host from which they were backed up.
- **Export to a different host –** The requested files are restored to the same, or a different host added to the Rubrik Security Cloud inventory.

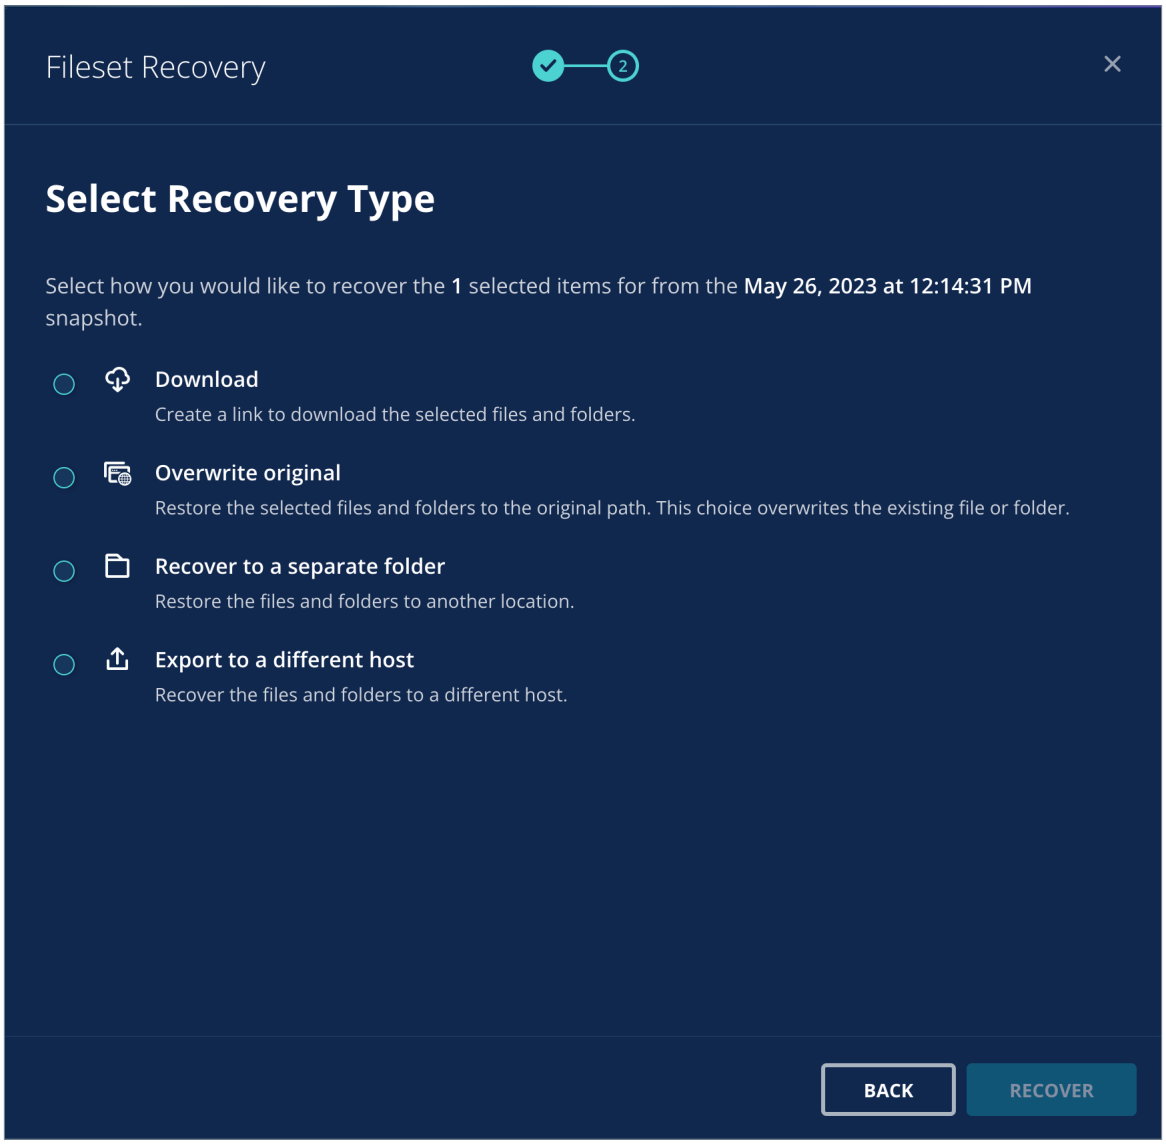

THE RUBRIK ADVANTAGE

Legacy backup technologies typically have challenges backing up and restoring Caché/IRIS databases in a reasonable amount of time due to the large total size of the database and the large sizes of single `CACHE.DAT` or `IRIS.DAT` files. It is now common to see Caché/IRIS databases in the 10's or even 100's of TB with single `CACHE.DAT` or `IRIS.DAT` files being multiple TB in size. Healthcare data is roughly 30% of the world's data volume being generated, leading customers to wonder if their existing legacy solutions can keep up with daily backups.

For legacy solutions, a lot of time is then spent tuning, breaking the backups into multiple sets, juggling different jobs to get concurrency, making config changes for the number of threads, and other modifications which may not offer much improvement. Furthermore, these processes only add to the management complexity required to perform backup and restore operations. Legacy backups might also take occasional full backups, which exacerbate and can lead to long backup times and storage inefficiencies.

Rubrik has the speed and technology to protect large, challenging Caché/IRIS database environments. In addition, Rubrik reduces operational complexity and improves the security of your backups.

During each backup, Rubrik performs the following:

1. **Scan** – RBS traverses the file system to identify what files have been added or modified since the last backup. A list of incremental files that need to be backed up is built and automatically grouped into partitions that enable parallel backups. The scan phase typically finishes in less than a minute since there are only 200-400 files in the Caché/IRIS directory.

2. **Fetch** – RBS performs incremental forever backups by reading data from multiple partitions concurrently and using fingerprints to identify changed chunks within a file since the last backup. Only the changed chunks are transferred to the Rubrik cluster, and data is compressed and encrypted before being sent. This makes for efficient network utilization between the backup proxy host and the Rubrik cluster. RBS transfers data to multiple Rubrik nodes for fast, parallel, multi-node ingest.

3. **Copy** – As partitions finish ingesting on the Rubrik cluster, the data is de-staged into the immutable Rubrik filesystem, purpose-built to maintain data availability and integrity. Data is encrypted at rest for security and erasure coded for resiliency. CRCs and fingerprints are calculated and stored with the data to ensure data integrity. Continuous background scans are performed against the data, and any chunks failing CRC or fingerprint verification will automatically be rebuilt. Any operation that reads the data, such as a restore task, will also verify the CRCs and fingerprints to ensure that the data being read is the same as when it was written.

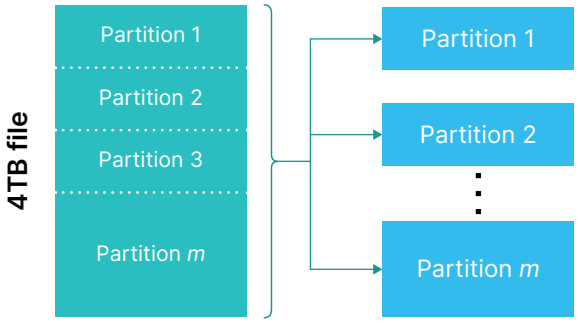

In addition, Rubrik shards large single files into multiple partitions, as illustrated to the left.

The key to Rubrik's fast backups and restores is that the partitions are spread across all Rubrik nodes to be processed. This provides a high degree of parallelism and concurrency automatically without the need for time-intensive, manual tuning.

In production environments, Rubrik typically sees 5:1 or 80% data reduction for Caché/IRIS databases and 10-12% daily change rates. This means that a 50 TB Caché/IRIS database can be reduced to only 10 TB when stored on the Rubrik cluster. Since RBS can figure out the incremental changes, it will typically only need to transfer 5-6 TB daily that is compressed down to 1-1.2 TB, making for very fast backups that are proved out in our testing and also seen in the real world.

When it comes to recovery, Rubrik provides a multitude of options to support nearly every restore scenario. Organizations can leverage in-place recovery to quickly restore their Cache/IRIS databases to a specified point in time, essentially rewinding their production Epic environment. In addition, Export functionality gives organizations an easy way to restore point-in-time copies of their production Epic environment into an isolated recovery environment to be used for both recovery validation and development and test purposes.

INTERNAL RUBRIK PERFORMANCE TESTING

The following section outlines the hardware, software, and technologies utilized for Rubrik's internal performance testing of backing up and recovering an Epic Cache database.

OVERVIEW OF VENDORS UTILIZED

Throughout the internal performance testing, many components were utilized to achieve the results. The following is an outline of the hardware and software vendors included within the testing process

Epic

Epic Systems Corporation, or Epic, is the leading provider of electronic health record (EHR) software. More than 250 million patients have a current electronic record in Epic. Their technology supports the entire end-to-end patient care experience: registration and scheduling, clinical systems for medical personnel, systems for pharmacists and radiologists, and billing systems for insurers.

Epic's suite of healthcare software is centered on the InterSystems Caché and IRIS databases. Caché/IRIS is a proprietary database architecture developed and maintained by Intersystems Corp that runs on AIX or Red Hat Enterprise Linux (RHEL) OS. The Epic EHR also makes use of a file blob store it uses as a storage location for various images/files, a 'Clarity' database running on Microsoft SQL Server or Oracle used for clinical reporting, and other applications which are usually virtualized.

Rubrik Security Cloud

Rubrik helps enterprises achieve cost-effective data protection and cyber recovery, cyber resiliency, and accelerate cloud adoption. Rubrik Security Cloud unifies backup, instant recovery, replication, globally indexed search, compliance, and copy data management into a single scale-out fabric across the data center, SaaS, and 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, SaaS, and cloud.

Pure Storage

Pure Storage is a leading provider of data storage solutions that are revolutionizing the way businesses manage and utilize their data. With a relentless focus on innovation, Pure Storage offers a comprehensive portfolio of all-flash storage arrays and software-defined storage solutions that deliver exceptional performance, efficiency, and simplicity.

Supermicro

Supermicro is a global leader in high-performance, high-efficiency server technology and solutions. With a strong commitment to innovation and quality, Supermicro designs and manufactures a wide range of server and storage solutions that cater to the unique needs of businesses across various industries. Their products are renowned for their superior performance, energy efficiency, and scalability, allowing organizations to optimize their computing infrastructure and drive operational efficiency.

TESTING ENVIRONMENT

The following outlines the specification of the hardware and software components leveraged during this performance testing.

Hardware Components

The test environment included the following hardware components:

- Rubrik CDM Cluster r6000
  - 12 node cluster running CDM 8.1.2-24509
  - 25Gbs networking in active/passive bonds
- Compute for ESXi
  - Supermicro 1029U-TR25M
  - 2x Intel Gold 5218 CPU
  - 512 GB Memory
- Storage
  - Pure Storage FlashArray X70
  - 16 × 8TB Volumes
  - Presented over 32Gpbs Fibre Channel

Software Components

The test environment included the following software components:

- VMware ESXi v7.0.3.0
- RHEL v7.9
- Rubrik Security Cloud (management plane)
- Rubrik CDM 8.1.2-24509
- Pure Storage Purity v6.4.1
- GenerateIO v1.17.4

Test VMs

The test environment consisted of the following VMware virtual machines:

- Primary VM (RHEL):
  - 24 vCPU, 64GB RAM
  - 16 × 8TB Volumes attached as RDM
  - GenIO data files were created on this VM
  - I/O was run against the data files to generate changes between backups
- Backup Proxy VM (RHEL):
  - 24 vCPU, 64GB RAM
  - Snapshots of Primary VM Luns were mounted to this VM for processing
  - RBS was installed on this VM to facilitate backups

The following architecture diagram illustrates the above information:

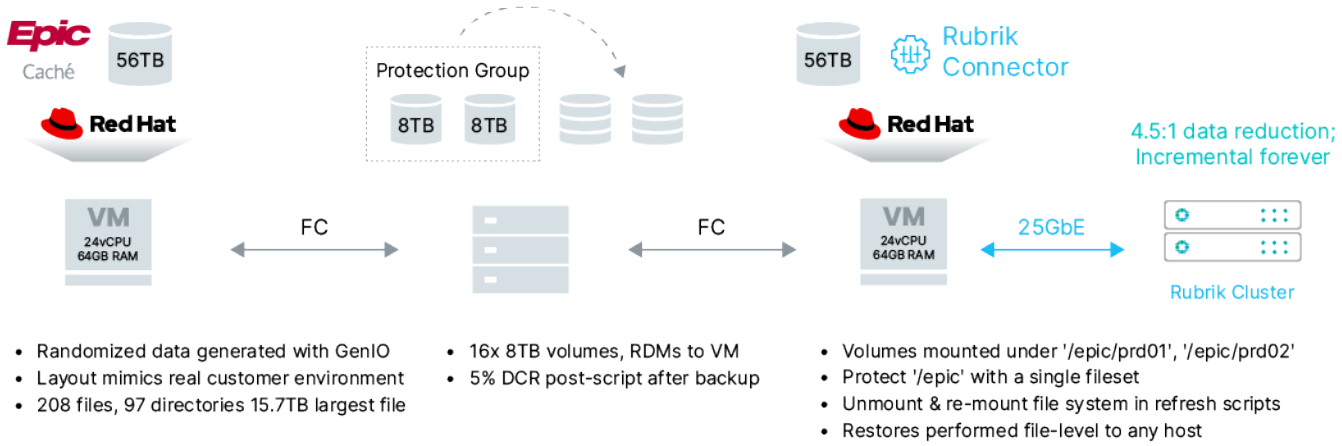

GENERATING DATA FILES AND CHANGES

GenerateIO (GenIO) is a tool used to test storage systems with synthetic workloads simulating a Caché/IRIS database load. For this testing, GenIO was used to create data files that mirrored a real customer environment that would be backed up. A total of 56 TB of data was created. I/O was generated against the data files between each backup to simulate daily changes. These daily changes amounted to roughly ~5% of changes between each test backup.

In production Caché/IRIS environments, the CACHE.DAT or IRIS.DAT database files are stored under the '/epic/prd##' directory with 80 – 130 sub-directories each containing the CACHE.DAT or IRIS.DAT files. Several of these CACHE.DAT or IRIS. DAT files can be substantially larger than others, oftentimes in the multi-TB range. During our testing, 208 total files were processed within 97 directories with the largest single file sitting at 15.7TB in size.

TESTING RESULTS

The following outlines the backup and recovery results from the internal Rubrik performance testing.

Backup Test Results

Rubrik's testing proves that RSC is capable of excellent backup performance as seen in the results below:

Test Performed	Test Details	Results
Epic Cache Full Backup	56TB Cache Database	7.8 hours
Epic Cache Incremental Backup	5% Daily Change Rate	7.1 hours

Rubrik achieved excellent results in processing the backup of an Epic Cache database. Through the use of file sharding Rubrik was able to ingest the entire 56TB database in just under 8 hours (7.8 hours) resulting in data throughput of roughly 7.1 TB per hour. During an incremental backup, Rubrik first reviews changed files, breaking them into 64K segments, transferring only those segments with changed data to Rubrik. While this does add to the incremental backup time, it ensures peak network optimizations in terms of data transfer.

With Rubrik's scale-out architecture, backup performance should scale linearly with the environment size. If the database size is 100TB, we should expect to see the backup complete in just under 16 hours.

Restore Test Results
Rubrik's testing proves that RSC is capable of excellent restore performance as seen in the results below:

Test Performed	Test Details	Results
Epic Cache Full Restore	56TB Cache Database	5.25 hours | 10.6 TB/hour

Rubrik is capable of very fast restore speeds, restoring the entire 56TB of files in 5 hours and 15 minutes. As with backup, this performance should scale linearly with the environment size, meaning if the database was 112TB, restore should take roughly 10 hours and 30 minutes.

TESTING CONSIDERATIONS FOR PRODUCTION ENVIRONMENTS

Backup and restore speeds will depend on the overall environment that Rubrik is deployed in. The test environment was set up to mirror a production environment as closely as possible. Factors that may lead to better or worse performance include:

- **OS:** The testing was done on RHEL and performance may be different on AIX
- **Caché/IRIS file system layout:** Each Caché or IRIS database will have a different file system layout and file size skew.
- **Backup proxy host resources:** The test VM was given 24 virtual cores and 64 GB RAM. Based on our testing we found the following:
  - RBS utilized more CPU during the first full backup and during restores
  - RBS does not use relatively too much CPU during steady-state incremental backups. Although we provided 24 cores, good backup times could be achieved with as little as 16 cores
  - RBS does not use or require that much RAM — 16 – 32 GB of RAM would be sufficient
- **Storage array:** The performance of the storage array will determine how fast RBS can read and restore backup data
- **SAN connectivity:** The available SAN connectivity between the proxy host and storage array will limit how fast RBS can read and restore the backup data
- **IP network connectivity:** The available IP network connectivity between the backup proxy host and the Rubrik cluster will limit how fast RBS can read and write data to the Rubrik cluster

SUMMARY

Rubrik is uniquely positioned to help customers protect and secure their Epic EHR environments. This paper shows how easily Rubrik protects Epic Caché/IRIS databases but there may be other workloads that prove challenging to existing legacy solutions. These might include:

- **Clarity Database –** Rubrik can backup SQL and Oracle databases natively and protect the logs for granular, automated recovery.
- **Blob tier –** Typically 100's of millions of files that legacy NDMP-based solutions cannot backup. Rubrik can backup NAS at a massive scale – billions of files and multi-PB – cost-effectively, to either on-prem or cloud targets.
- **General Apps & VMs –** Protect virtual and physical environments with a single solution that provides operational simplicity at scale across all workloads.

By choosing Rubrik, you can:

- Dramatically improve backup & recovery times compared to existing legacy solutions
- Reduce complexity with an easy-to-use UI and automation with SLA Domain policies
- Improve your security posture with a data protection solution designed around zero trust
- Accelerate public cloud adoption
- Provide a compelling TCO

For more information about how Rubrik is transforming the healthcare data security landscape, visit rubrik.com.

APPENDICES

APPENDIX A - TRIGGERING AN ON-DEMAND BACKUP USING RUBRIK SECURITY CLOUD APIS

Rubrik Security Cloud deploys a full set of GraphQL APIs which can be utilized to automate nearly every process within the platform. To take an on-demand snapshot of a configured fileset to backup your Cache/IRIS database, the desired SLA Domain ID and Fileset ID must first be obtained.

The following query will return a list of all SLA Domains configured within Rubrik Security Cloud:

```
query getAllSLADomains {
  slaDomains {
    edges {
      node {
        id
        name
      }
    }
  }
}
```

Once executed, a response similar to that below should be returned:


```
{
  "data": {
    "slaDomains": {
      "edges": [
        {
          "node": {
            "id": "00000000-0000-0000-0000-000000000002",
            "name": "Bronze"
          }
        },
        {
          "node": {
            "id": "00000000-0000-0000-0000-000000000000",
            "name": "Gold"
          }
        },
        {
          "node": {
            "id": "00000000-0000-0000-0000-000000000001",
            "name": "Silver"
          }
        },
        {
          "node": {
            "id": "8b9f29a9-5cec-47e1-af6a-c73f377565a7",
            "name": "SLA_EPIC_1D-7D"
          }
        },
        {
          "node": {
            "id": "cc2e764f-d666-4469-a580-9de4ce12941a",
            "name": "Tier1-Production-4hr-14d"
          }
        }
      ]
    }
  }
}
```


Make note of the SLA Domain ID that will be applied to the on-demand backup.

The following query will obtain a list of all Linux filesets configured within Rubrik Security Cloud:

```
query getLinuxFilesets {
  filesetTemplates (hostRoot: LINUX_HOST_ROOT ) {
    edges{
      node {
        id
        name
      }
    }
  }
}
```

Once executed, output similar to the following should be shown:

```
{
  "data": {
    "filesetTemplates": {
      "edges": [
        {
          "node": {
            "id": "62ab6527-b374-5ba5-9a39-a91f28e0fabe",
            "name": "IRIS DB"
          }
        }
      ]
    }
  }
}
```

Make note of the desired fileset ID to be utilized within the on-demand backup.

Finally, the following mutation can be executed in order to take the on-demand backup, replacing appropriate placeholders with the id's of the SLA Domain and Fileset.

```
mutation EpicOnDemandBackup {
  createFilesetSnapshot (input: {
    config: {
      slaId: <id of sla or null>
    },
    id: "<id of fileset>"
  }) {
    id
    status
  }
}
```

Once executed, the following output should be displayed:

```
{
  "data": {
    "createFilesetSnapshot": {
      "id":
"CREATE_FILESET_SNAPSHOT_b78a796f-727d-4418-b122-74a56f252aa5_7061b1c1-a509
-44dd-a39a-131784196ea1:::0",
      "status": "QUEUED"
    }
  }
}
```

For more information around Rubrik Security Cloud and GraphQL please visit the Rubrik API page.

APPENDIX B - DETAILED RUBRIK BACKUP CONFIGURATION

It is extremely easy to configure Rubrik for Caché/IRIS backups. The steps to do so:

1. Download Rubrik Backup Service (RBS) from the Rubrik cluster
2. Install RBS on the backup proxy host
3. Add the backup proxy host in Rubrik
4. Create a new SLA domain policy that defines backup frequency, retention, backup window, archive, and replication configuration
    a. Once created, a SLA domain policy can used with any object
5. Create a fileset containing the desired directory to protect
    a. One fileset for the CACHE.DAT or IRIS.DAT files, usually including: "/epic/prd01/" or "/epic/prd**"
    b. Optional - one fileset for the OS, including: "**" and excluding: "/epic/prd01/" or "/epic/prd**"
    c. Once created, a fileset can be used for any host of the same OS type
6. Assign the Fileset + SLA to the backup proxy host

Rubrik backups will automatically be sharded with parallel backup ingest across all nodes in the cluster. There is no requirement to manually create additional threads, tuning, etc.

Download Rubrik Backup Service from the Rubrik Cluster

In the UI, navigate to **Settings → Data Sources → Linux & Unix** and select the **Add Hosts** button. In the resulting dialog, select the desired Rubrik cluster to store the backups. Select the "I want to install and register the RBS software now" option and use the associated links to either download or copy the download path, ensuring you have the rpm for RHEL or AIX.

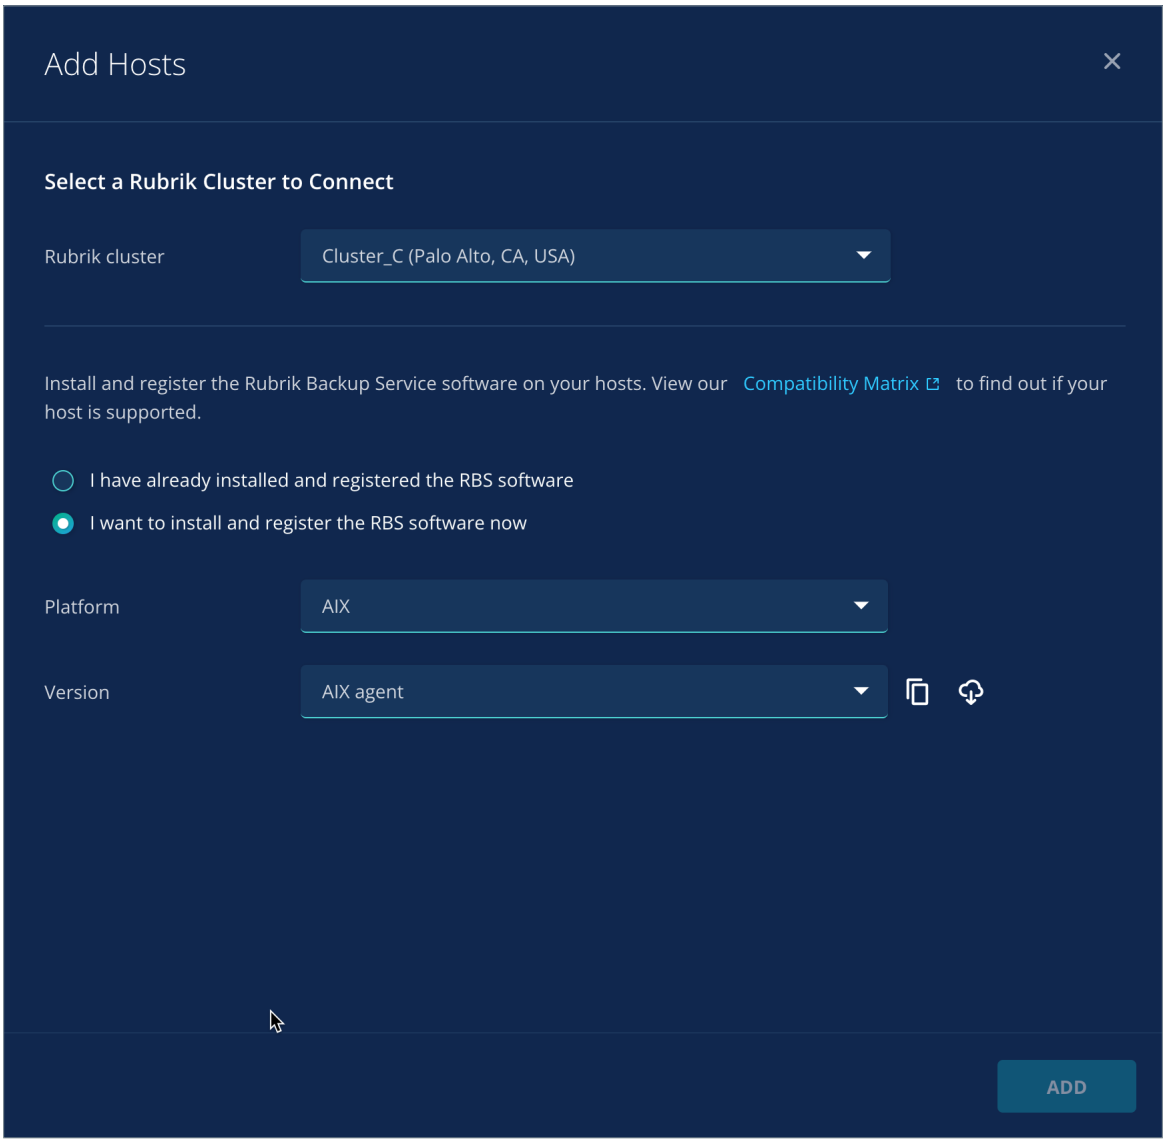

Note: Do not exit this dialog as there are more steps to perform once RBS has been installed

Installation of RBS on the backup proxy

Once the RBS rpm for your OS is downloaded, transfer it to the backup proxy and install it as root:

- sudo rpm -i rubrik-agent.x86_64.rpm
  - Or use 'yum' package manage to install
- sudo rpm -ivh rubrik-agent-aix7.1.pcc.rpm
- sudo rpm -ivh rubrik-agent-aix7.2.pcc.rpm

You can verify that RBS is running using:

- ps -ef | grep rubrik
  - Two Rubrik services will be running, 'bootstrap_agent_main' and 'backup_agent_main'

Addition of backup proxy to Rubrik Security Cloud

Once RBS is installed on the backup proxy host you can add it by hostname or IP address using the "Add Hosts" screen that was seen earlier. If the existing dialog is still available, simply select the "I have already installed and registered the RBS software" option and fill in either the IP Address or FQDN of the backup proxy. Once complete, the host should appear in the UI and be in a Connected state.

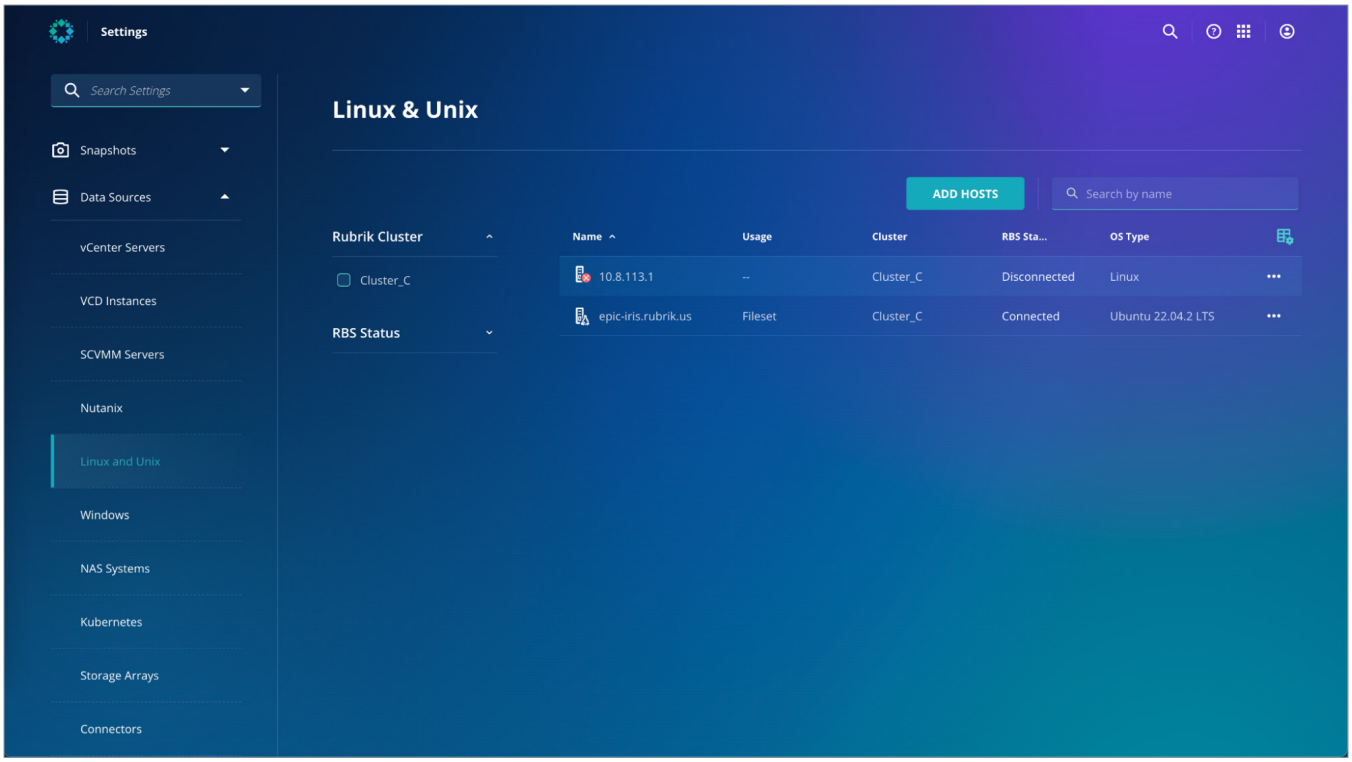

Creation of new SLA Domain

You can create a new SLA Domain Policy by navigating to Data Protection → SLA Domains and selecting the Create SLA Domain button.

The core of Rubrik's simplicity at scale revolves around the SLA Domain policies. A SLA Domain policy defines the following:

- **SLA Name –** Give the SLA a name
- **Backup frequency & retention –** define a policy aligned to business requirements such as: daily backups retained for 30 days, a monthly backup retained for 12 months
- **Backup window –** defines when a backup is allowed to kick-off
- **Full backup window –** defines when full backups are allowed to kick off since they are the most impactful
- **Archiving –** send a copy of the backups to an archive location, which can be NFS, S3-compatible object storage, or a cloud provider
  - All tiers of storage classes for Azure, AWS, and GCP are supported
  - Long term retention – offload backups greater than the archive threshold to the archive location
  - Instant archive – send a copy of the backup as soon as it is taken

- **Replication –** send a copy of the backups to another Rubrik cluster, which can be on-prem or running in the cloud
  - Send the last x days of backups to the remote cluster

Once an SLA Domain policy is created it can be assigned to ANY workload. This allows a single SLA Domain policy to protect VMs, Windows & Linux hosts, SQL, Oracle, and NAS shares, etc. Most customers leverage a handful of policies that are aligned with business requirements.

Creation of new fileset

You can create a fileset by navigating to **Inventory → Data Center → Linux & Unix Filesets → Filesets** and selecting the **Add Fileset** button.

A fileset defines what files and folders to backup on a host being protected using RBS:

- **Fileset name –** give the fileset a name
- **Include –** defines which directories or files to include
  - **/epic/prd01/*** – includes only the files within the directory, but no sub-directories
  - **/epic/prd01/**** – includes all files within the directory and all sub-directories
    - **/epic/prd01/**– this will add an implicit ** and is the same as **/epic/prd/01/****
- **Exclude –** defines which directories or files to exclude
- **Do not Exclude –** defines which directories or files to never exclude
- **Follow network shares –** option of whether to follow network shares, default off
- **Enable pre/post scripts –** option to run a pre-/post- script before/after the backup completes

Once a fileset is created it can be assigned to one or more hosts. A single host can also have multiple filesets assigned.

For example, a fileset to protect the Caché/IRIS database files will look similar to the following:

- **Include – /epic/prd01/** or **/epic/prd**** if you have multiple /epic/prd## directories

A second fileset can be created to protect the OS binaries by using the following rules:

- **Include – **** to backup all files
- **Exclude – /epic/prd01** or **/epic/prd**** to exclude the Caché/IRIS database files in this fileset

Having two filesets, one for the Caché/IRIS database and another for the OS binaries, makes the backup and reporting cleaner.

A fileset can also call pre/post-scripts. In some of the smaller Caché/IRIS environments, for example, the ACE or training environments, a fileset may be created to directly backup those hosts, calling a script to freeze and thaw the Caché/IRIS database as part of the backup process.

Assignment of fileset + SLA to backup proxy

Once you have the SLA and Fileset created you can use them to protect the backup proxy host file system. There are two options:

1. Assign the Fileset + SLA to the backup proxy host. Rubrik will automatically schedule backups to be taken according to the backup window.
    a. You can schedule the backup to be run at a time you know the backup proxy will be ready by
    b. You can have Rubrik trigger pre-scripts to prepare the proxy host when the backup kicks off. Rubrik backups will continue when the script exits with a code 0
2. Trigger an on-demand backup as part of the backup proxy refresh script
    a. Use an API call to the Rubrik cluster to trigger an on-demand backup

To assign a Fileset + SLA to the backup proxy to start protecting it, navigate to the host by either searching for it using the global search toolbar at the top, or navigating to the host under **Inventory → Linux & Unix Hosts**.

Once you are at the host overview screen you can click "**Manage Protection**" to assign the Fileset and SLA to the host.

VERSION HISTORY

Version	Date	Summary of Changes
1.0	August 2023	Initial Release

rubrik

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

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

rwp-protecting-epic-with-rubrik-security-cloud / 20230810