Configuration Guide

VAST Data Platform with Rubrik® NAS Cloud Direct

Solution & Best Practices Guide

03/25 | Version 1.0

Table of Contents

Introduction 4

Audience 4

Solution 5

Key Takeaways 5

Best Practices 6

Summary 8

General Configuration 9

Rubrik NCD with VAST as a Source 9

Rubrik NCD with VAST as a Target 14

Creating a Policy 17

Running a Backup/Archive 20

Running a Restore 23

VAST Configuration 26

Revision History

Revision Number	Description	Date
1.0	Initial Release	Mar 2025

Introduction

Whether migrating data from older primary storage to VAST or backing up data on VAST to a secondary or disaster recovery site, the VAST Data Platform provides extreme performance with a highly scalable architecture. The Disaggregated Shared Everything (DASE) architecture provides the ability to scale compute and capacity independently depending upon the performance requirements and growth needs.

VAST Data's multiprotocol support allows for pluggable integration as a primary copy of data or a backup target without compromising on storage scale capabilities, redundant appliances, or isolated silos of backup targets. With the flexibility of NFS, SMB and S3 supported protocols Rubrik's NAS Cloud Direct (NCD) integrates seamlessly as a performant mover of data over any protocol.

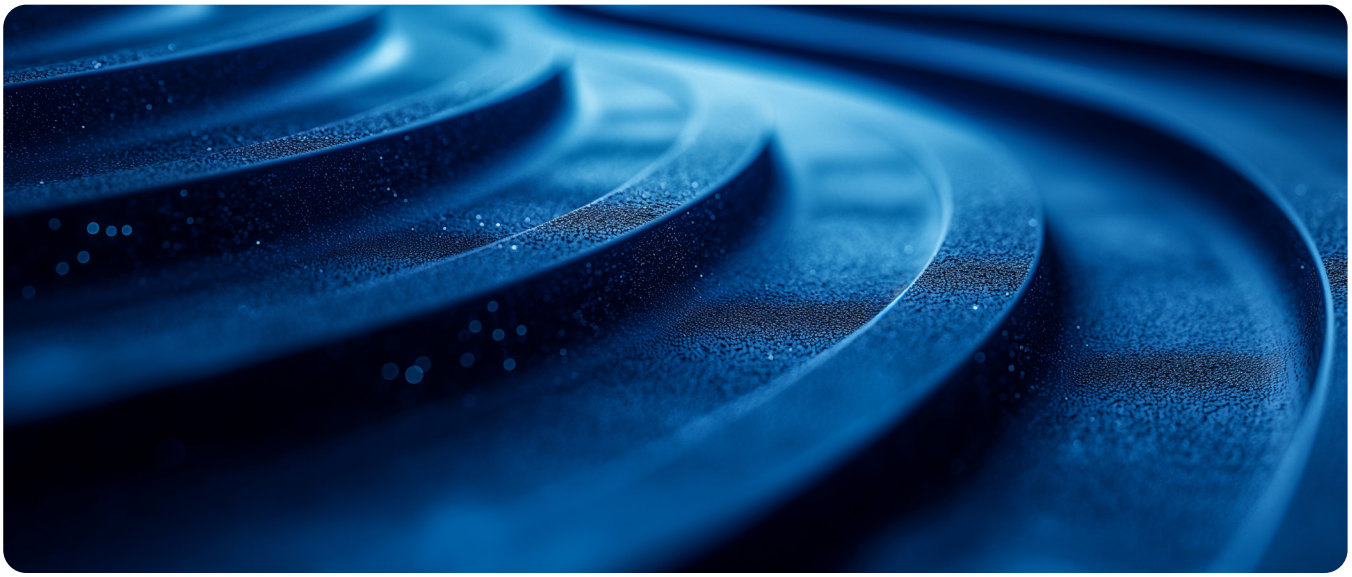
NAS Cloud Direct offers secure data protection that delivers high-performance, scalable enterprise backup and archive services for unstructured file and object data. With Rubrik and VAST you can balance security and performance to protect and manage petabytes of NAS data at scale. Secure all file data stored on any file system under any NAS protocol and write data to VAST over NFS or to an S3 target on-premises or backup VAST from NFS, SMB, or S3 offsite or in the cloud.

NAS Cloud Direct is deployed as a virtual machine and is engineered to move data at line speed without impacting production workloads. Compatible with every NAS architecture and capable of managing billions of files, NAS Cloud Direct is built for the scale of the modern enterprise. VAST's unified global namespace and patented deduplication and similarity technology can achieve enormous capacity savings. Together the performance and capacity needs can scale to meet demand and the combined solution simplifies data protection and reduces costs.

Audience

This document is a high-level integration guide to assist administrators and others interested in understanding the workflow for configuring the VAST Data Platform with Rubrik® NAS Cloud Direct. There are two major use cases with VAST and Rubrik® together and that is VAST as the source of the data and VAST as the target. This document will cover both use cases and will attempt to address the benefits of each.

This paper will describe how Rubrik NCD users should configure their NCD data movers to best operate with VAST. It outlines technical details, architectural considerations, and recommended best practices when combining the VAST Data Platform with a NAS Cloud Direct cluster.



Solution

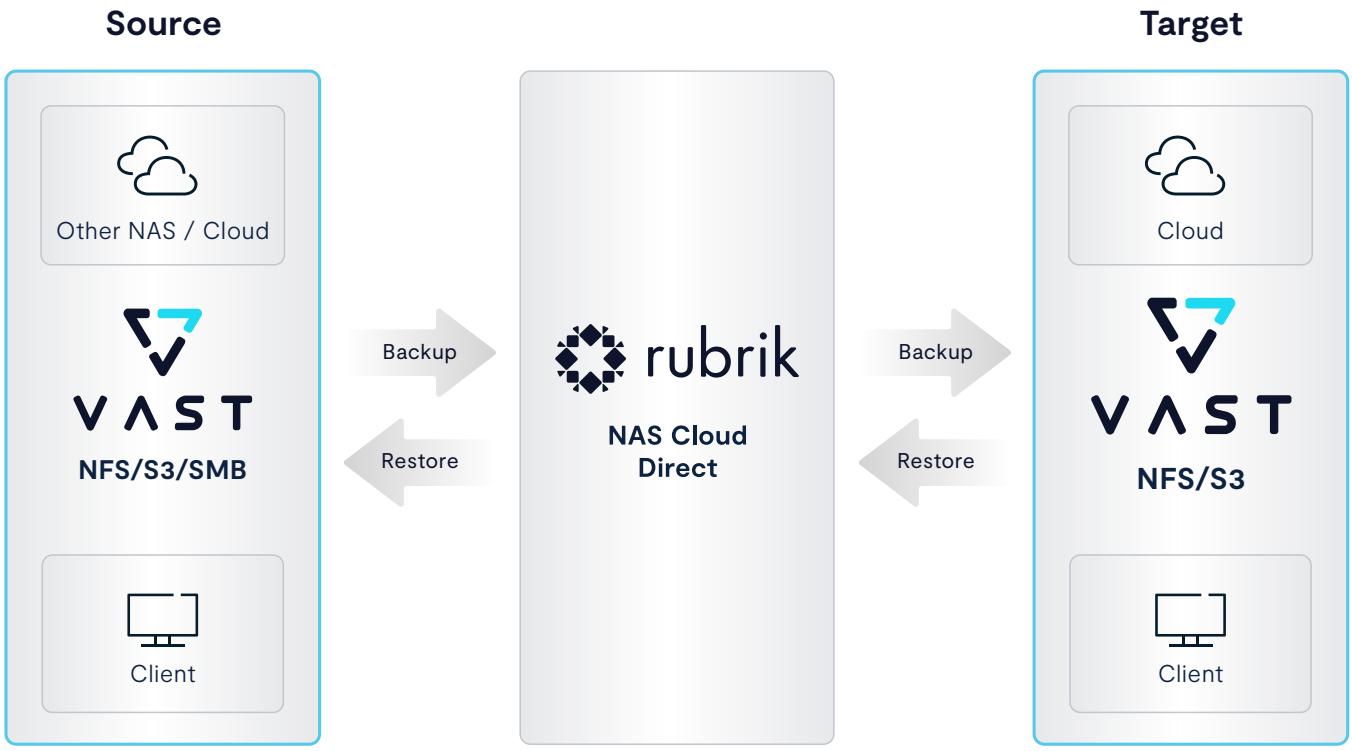


Figure 1 – VAST Data Platform and Rubrik NAS Cloud Direct Solution

Key Takeaways

VAST Data as the source –

- Effective backup, archive and copy of billions of files and objects with minimal setup and management overhead.
- Multiprotocol Support: NAS Cloud Direct can protect data from VAST with multi-protocol support including NFS, SMB and S3.

VAST Data as a target –

- Backing up to a VAST cluster, the target location can be over NFS or S3.
- VAST's single global namespace and patented Similarity data reduction algorithms will deduplicate and compress siloed data to the byte level and maximize capacity savings.
- Extremely performant restore times with a unique ability to allow for the restores right onto the same target system with zero capacity cost

Scalability – Both the VAST Data Platform and Rubrik NCD can scale according to capacity and performance requirements. For additional compute and capacity, CNodes and DNodes, respectively, can easily be added to a VAST cluster. Also, adding capacity to a VAST cluster (DNode) does not silo the data, it's seamlessly added to the global namespace ensuring maximum deduplication and capacity savings. For NAS Cloud Direct additional NCD stateless VM's can be added to a cluster of data movers with minimal clicks ensuring SLAs and backup windows are easily met.

Complementary data reduction – Rubrik NCD's inline compression helps to minimize traffic to accelerate data operations and minimizes data movement across the network, while VAST as a target can further reduce the backup footprint without disabling the NCD reduction. There is no need to compromise storage efficiency when both technologies can work in tandem.

Best Practices

The ease of use with Rubrik NCD and its "from anywhere, to anywhere design" coupled with the multiprotocol VAST Data Platform allows for automated, policy-based backup and archive operations from any NAS source (VAST NFS/SMB/S3), directly to any on-premises (VAST NFS/S3) or cloud target, based on definable policy settings.

The VAST Data Platform supports NFS, S3, and SMB making it extremely versatile as primary storage. With its ability to expose data through one or many protocols it easily integrates with Rubrik NCD. To help further the implementation of this solution the following topics will discuss ways to maximize performance and security.

Encryption

When storing data on a VAST cluster the ideal scenario to maximize deduplication and storage capacity efficiency is to ensure the data is not encrypted, but rather the transport of it is secured across the wire.

With Rubrik NCD's implementation for encryption they do not explicitly encrypt the data, rather they pass on the source data to the Rubrik NCD stateless VM's based on the configured protocol in this manner:

Protocol	Encryption Behavior
S3	Encryption is used if the S3 is configured as an HTTPS endpoint
SMB	Encryption is used if enforced by the source NAS
NFSv4.1	Encryption is used if Kerberos is configured

Table 1 – Rubrik NCD Encryption Behavior

> **Note**
>
> Data only passes through the NAS Cloud Direct (NCD) VM and is not actually stored on the VM.

Taking this behavior into account it is possible to ensure that the data is secure from source to target without encrypting the data itself.

For a secure backup scenario, the simplest method to secure data on the network between Rubrik NCD and VAST is to configure the VAST cluster as a 'cloud' library and to use the SSL protocol for communication which uses TLS encryption on all data in-flight to the VAST cluster. (**Note:** The VAST SSL security certificate may have already been installed during the NCD installation but if not refer to the section, SSL Certificates, for more information). No matter what secure protocol is used, once the data lands on the VAST Data platform all data at rest are encrypted.

Using a secure protocol to transport the data and coupled with VAST Data's encryption at rest ensures secure end to end solutions are possible while allowing for maximum capacity savings.

Compression and Deduplication

With a proper understanding of how encryption plays a role let's address compression and deduplication.

By default, Rubrik NCD compresses the source data into larger blob files and is very effective at getting the most out of compressing data. It uses inline compression with the LZ4 algorithm which achieves an excellent balance of capacity savings and performance. In the end it significantly drives down the volume of data transferred and stored during transfer. This minimizes the network traffic from the NCD stateless VM's to the VAST Data Platform which in turn accelerates both backup and restores.

Since Rubrik NCD does not explicitly encrypt the data, VAST Deduplication and Similarity can work its magic and provide additional capacity savings making it complimentary to Rubrik's compression. With VAST as the target for multiple NAS silos or vendors it can significantly reduce similar datasets across all the silos, minimizing storage usage.

Even if data is in different native formats (vmdk, vhdx, Oracle dbf, MS SQL mdf, various proprietary backup formats, etc.), VAST's similarity can provide significant data reduction between these formats storing them all on the same VAST cluster deduplicating across its global namespace.

Rubrik NCD and VAST data reduction techniques are complementary to each other and having them both simultaneously provides an overall best practice for both network efficiency and storage capacity considerations. By utilizing both technologies it maximizes space savings while optimizing network efficiency.

Performance

Both VAST and NCD can scale performance as needed in any solution. To ensure that the appropriate configuration is designed a few subtleties need to be addressed.

When reading off a VAST, either VAST contains the source data for a backup/archive/copy or has the backup data for a restore, the VAST platform will usually not be the bottleneck as its current restore speeds are almost 6x faster than writes. With read speeds up to 64GB/s (1 x 1 – entry level configuration) this makes for an extremely performant system whether the source or target. Write speeds on an entry level system can achieve up to 10GB/s.

NCD moves data in parallel streams and can leverage up to the full amount of network bandwidth, however due to a limitation of the VMXNET3 driver, a single NCD VM is bottlenecked at 5Gb/s (671MB/s). To fully saturate, for example, a 10Gb/s network Rubrik recommends 3 NAS Cloud Direct VMs, and this is easily done in just a few clicks.

Scaling the NCD VMs to meet performance needs is only part of the equation. There are 4 tasks that NCD performs – Discover, Backup/Archive (Full), Backup/Archive (Incremental) and restore and each VM instance is allocated a maximum of each according to the following:

Task		Limit
Discover		1
Backup / Archive	Full	1
Backup / Archive	Incremental	3
Restore		8

Table 2 – NCD Task Limits Per VM

Note

The maximum number of tasks allowed to run by task type are configurable by Rubrik Support.

These values scale linearly by adding NAS Cloud Direct VMs to the NCD cluster. For example, if there are 2 NCD VMs in the cluster, then the total maximum full backups running at any time (by default) is 2.

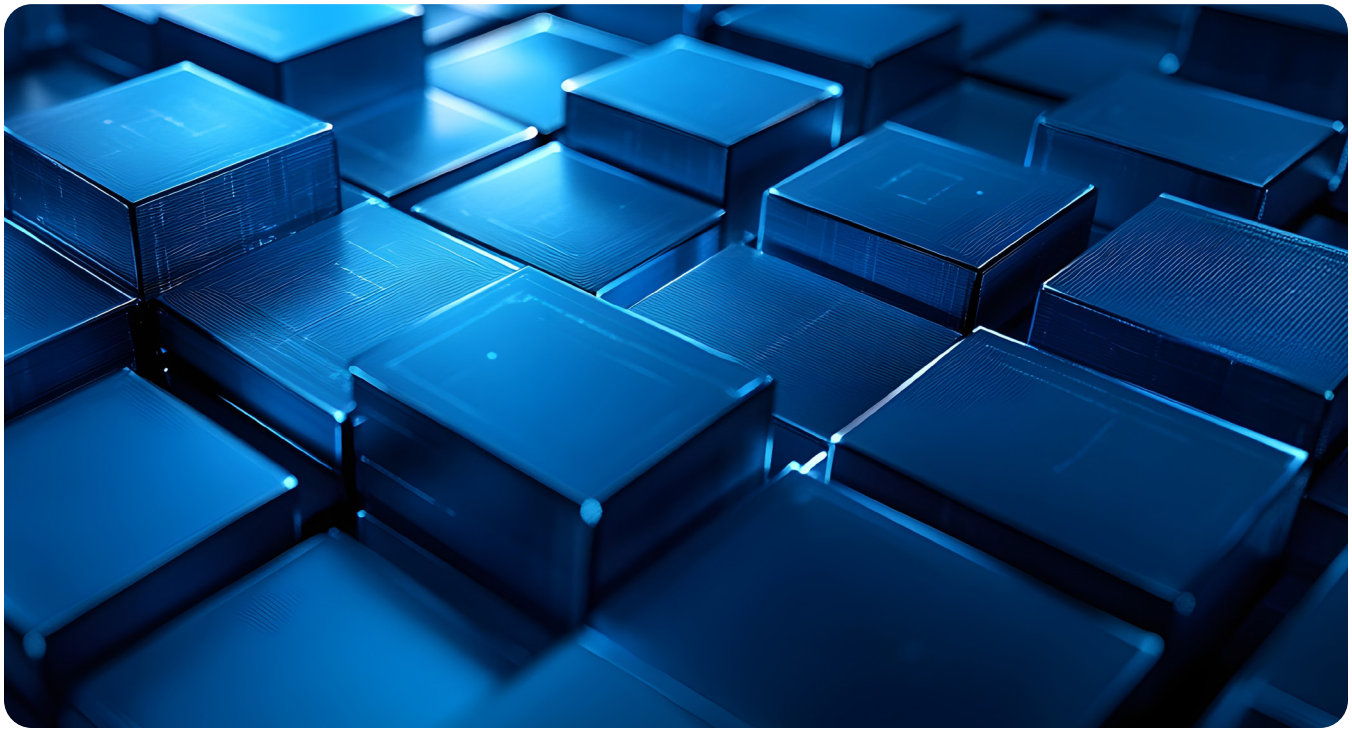
One additional task to ensure maximum performance when backing up data from VAST (source) is to ensure proper load balancing across all the VAST cluster CNodes. This is easily achieved with NCD's load balancing mechanism implemented with their connection pool and discussed more in the section NCD Connection Pool.

Summary

The combination of Rubrik software and the VAST all-flash platform creates a secure, cyber-resilient data protection solution that harnesses the technology of both vendors synergistically.

The data reduction capabilities of VAST and Rubrik are complementary. Rubrik will compress data minimizing local network traffic and speed up backups and then VAST will further reduce Rubrik's compressed data by up to an additional 2x. Almost cost-free data reduction is also achieved between backup and native formats and opens additional use cases in disaster recovery, CDP, high availability, and test/dev operations.

The solution is best in class, offering flash-speed backups and unrivaled performance for restores. Additionally, it offers economic advantages by reducing infrastructure requirements, thereby improving the total cost of ownership. It is particularly well-suited for large-scale environments, thanks to Rubrik's scalable infrastructure and VAST's disaggregated shared everything (DASE) architecture. This joint solution is fully integrated and operational without the need for complex configurations, providing robust cyber-resilient data management and protection capabilities.



General Configuration

This section covers the configuration tasks that are required to properly setup the VAST and Rubrik solution. The goal is to cover as many of the critical configuration steps as possible to ensure the implementation can be up and running quickly.

> **Note**
>
> Both VAST and Rubrik NCD UIs may have changed slightly from the release of this document. However, the concepts should remain the same.

Rubrik NCD with VAST as a Source

One of the use cases with VAST and Rubrik NCD is when VAST contains the source data and NCD will be used to move data from it to another location. When adding VAST as a new system (Figure 3) there are 3 supported protocols that NCD can connect to VAST with – NFS, SMB and S3. All mapped systems with source data will appear in the Rubrik NCD UI under Systems (Figure 2).

To begin the process of adding VAST click on the **Systems** icon in the left windowpane and then select the **Add system** button (Figure 2).

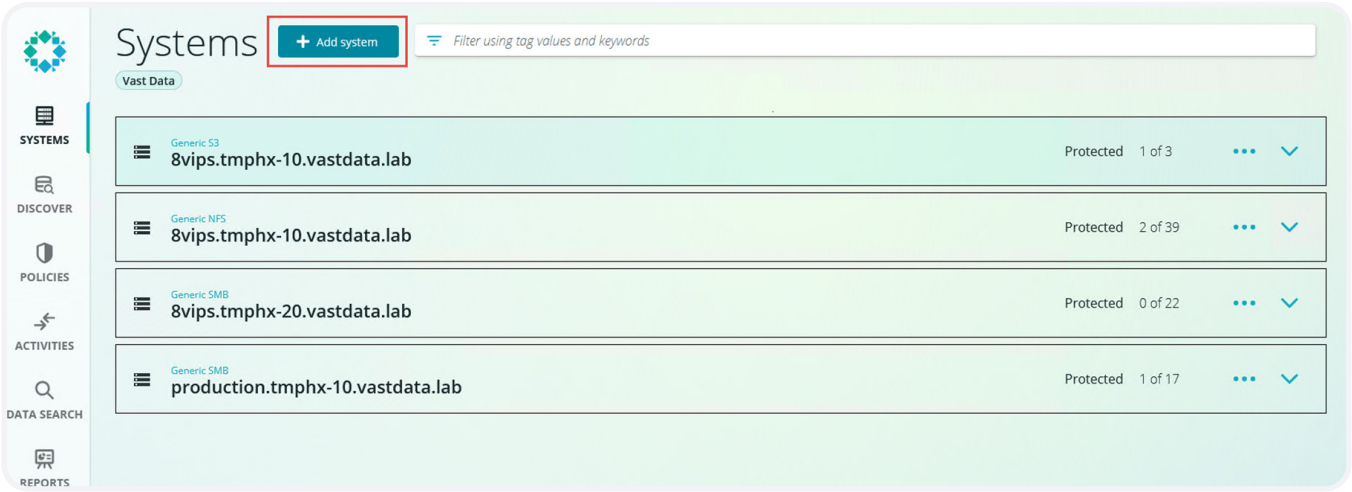


Figure 2 – Add System

This will pop up the Add a new system window as shown in Figure 3. We will address each of the protocols one at a time starting with NFS. Note that in the following examples a hostname (VAST VIP Pool) was used in the creation of each system type. An IP address can be used as well but in either case the configured system should be modified afterwards for load balancing. In the end the entry that's used to create the system will be its 'name' and what is seen within NCD. To assist administrators with recognizing systems more easily it's recommended that the hostname be used. (See NCD Connection Pool for more information)

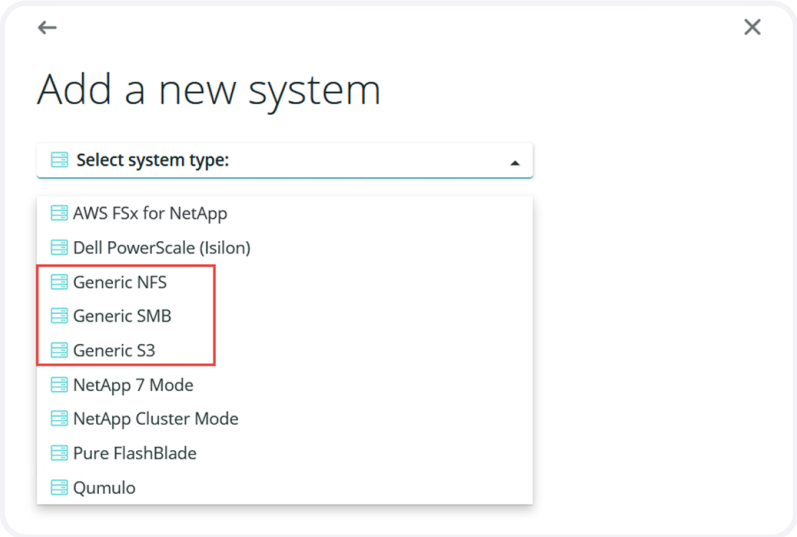


Figure 3 – VAST Supports NFS, SMB, S3 as a Source

Adding VAST NFS

From the menu in Figure 3 select **Generic NFS** and this will pop up the **Add a new system window** as shown in Figure 4.

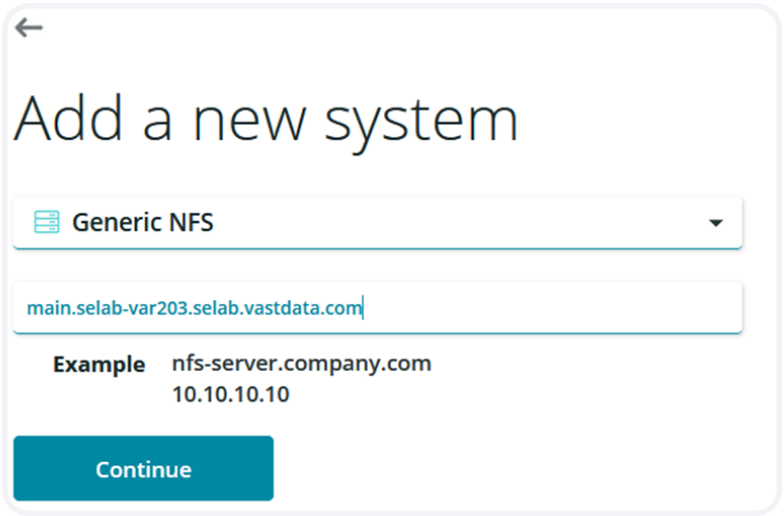


Figure 4 – Adding VAST NFS

When entering the NFS system using the hostname will make it easier to manage. Click **Continue** when ready.

Rubrik NCD will now query the VAST system via the hostname and tally all the NFS exports discovered (Figure 5). To finish adding the VAST NFS system click on **Complete adding system.**

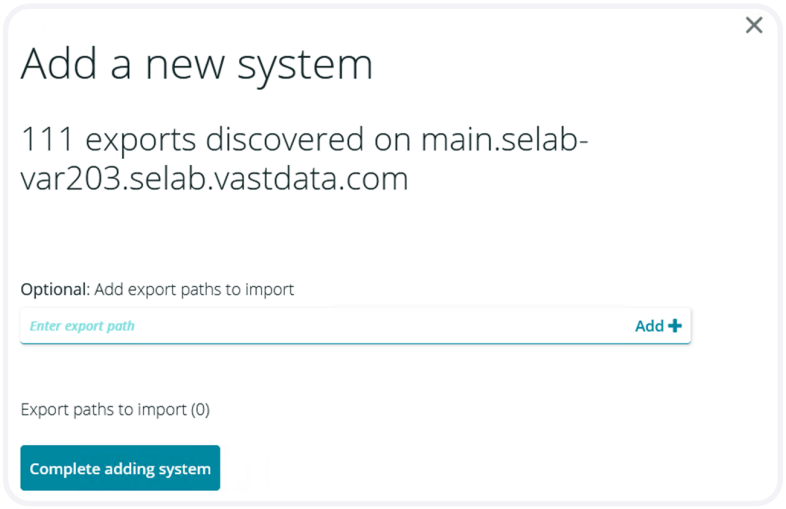


Figure 5 – Complete Adding System

Adding VAST S3

Starting from Figure 3 select **Generic S3**. This will bring up the window in Figure 6. Again, in this example the same FQDN that was used in adding the NFS system is used here. An S3 user, already created on the VAST side, is used and its secret and access keys are added (See Creating an S3 User). For security across the wire check the Verify SSL box. With VAST's encryption at rest all data from Rubrik is secure with today's most rigorous encryption methods.

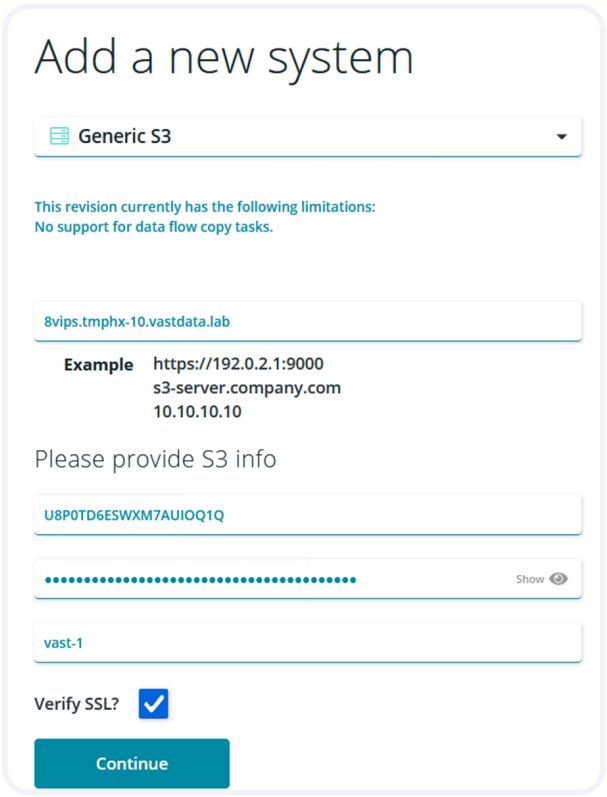


Figure 6 – Adding VAST as an S3 Source

After clicking **Continue** Rubrik will browse the VAST endpoint with the S3 user's credentials and discover all available buckets. The summary is shown in Figure 7. Simply click **Complete adding system** to finish the process.

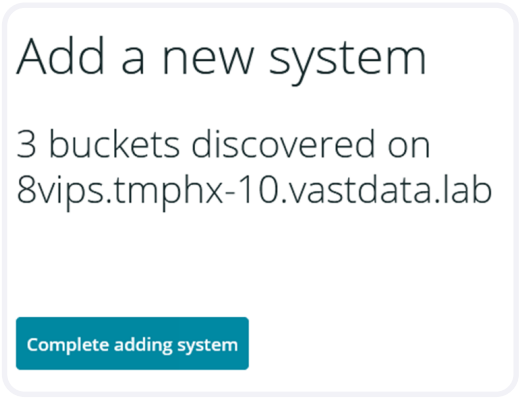


Figure 7 – Complete Adding VAST S3 as a Source

Adding VAST SMB

Again, starting from Figure 3 select **Generic SMB**. This will bring up the window in Figure 8. One thing to note here is that in the figure the same hostname is being used. The same SMB hostname cannot be used if there is already an existing NFS system with the same hostname. A similarly named S3 system can coexist with NFS or SMB but NFS and SMB systems of the same FQDN name cannot be created. Using two different VIP pools would be an acceptable solution.

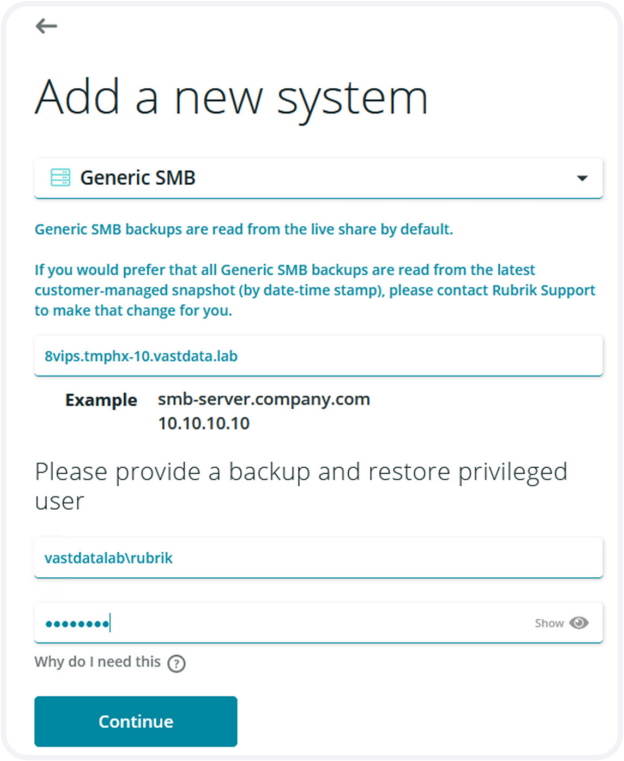


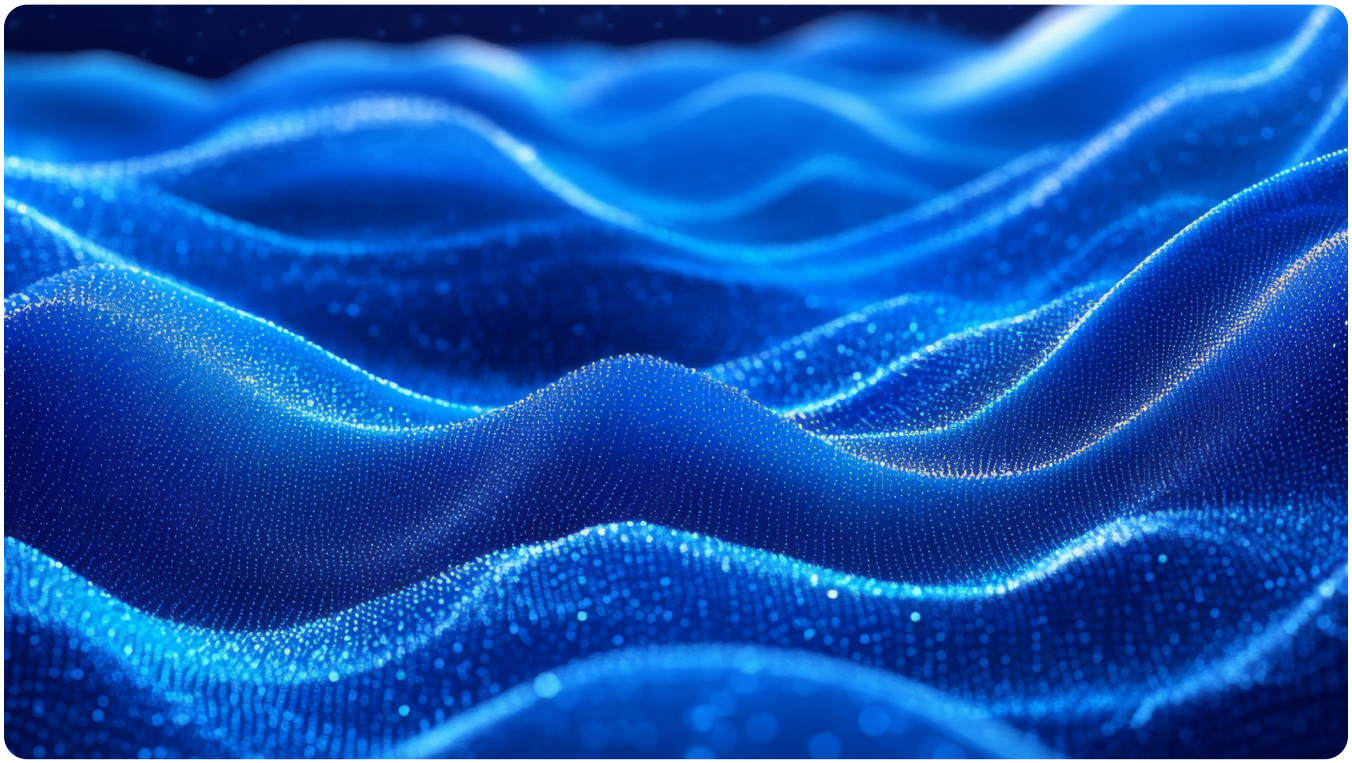
Figure 8 – Adding VAST SMB Source

The user being used is part of a domain that the VAST cluster itself has joined allowing for simple user privilege management. Simply click **Continue** when done.



Figure 9 – Complete Adding VAST SMB Source

Figure 9 shows the summary screen after Rubrik has scanned the filesystem and discovered all the available shares. Click **Complete adding system** when ready.



NCD Connection Pool

After adding a VAST cluster as a source through any of the supported protocols, it is a best practice to then modify that system and configure it within NCD for load balancing across all the CNodes in a VAST cluster.

Since the system was successfully added within NCD, it means that there is an existing Virtual IP (VIP) pool configured on the VAST cluster. The VIP pool is a collection of IP addresses that connect to the front end CNodes and allows client access to any exposed protocol. The general best practice for a VIP pool is 1 IP per CNode in a cluster. This ensures that each CNode participates in IO transfer. For NCD to properly load balance across all the IPs in the VIP pool, they will need to be entered into the configuration section of the system.

To do this click on the three dots at the far right of any system (Figure 10).

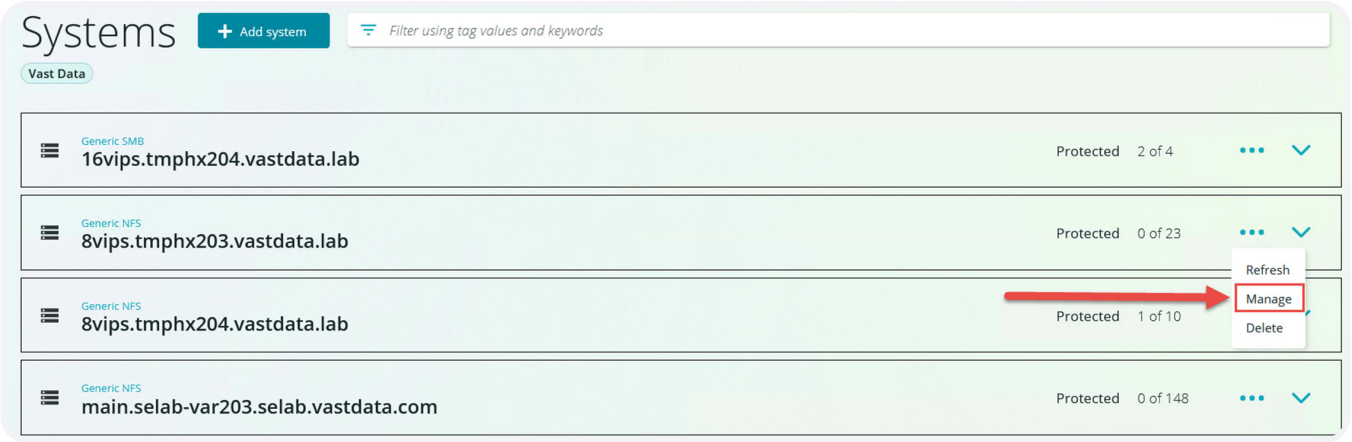


Figure 10 – Select Manage to Configure System

This brings up the settings of the system. Enter the IPs of the VIP pool. In Figure 11 the IPs were listed using a range separated by a hyphen.

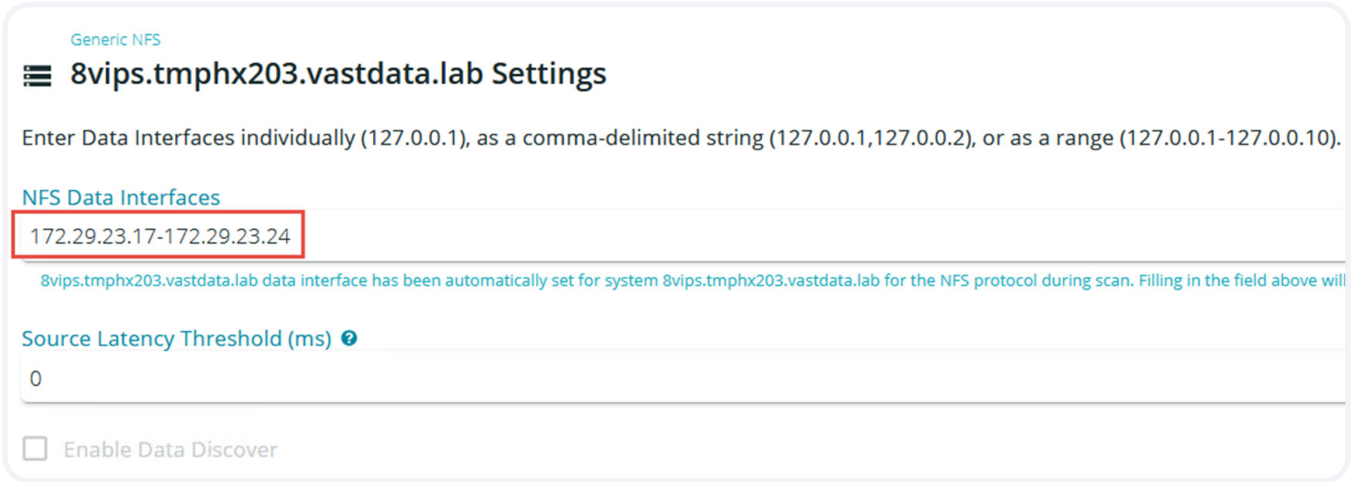


Figure 11 - Enter VIP pool IPs

NCD will automatically instantiate the IP range so that the IPs are listed as shown in Figure 12. When finished, click Save.



Figure 12 - VIP pool IPs Fully Instantiated

Rubrik NCD with VAST as a Target

Configuring a VAST cluster as a target can be done in a couple of different ways, from the Targets window or during the process of creating a backup policy. Both methods land the user on the same wizard, so the process is roughly identical. This section will highlight Targets window approach of adding VAST as a new target for both NFS and S3.

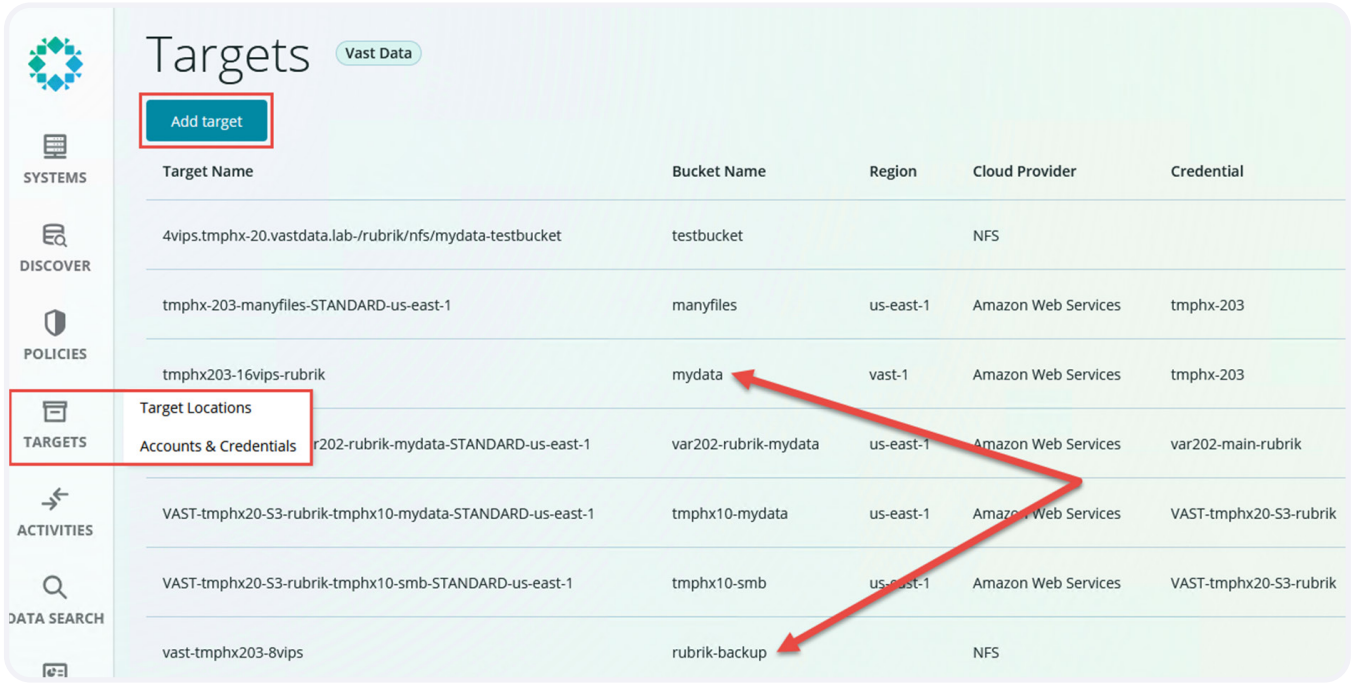


Figure 13 – Adding a VAST Target / Credential

To begin, click on the Targets icon in the left panel and select Target Locations (Figure 13). Then click on the Add target button at the top and selection of cloud providers will be shown as in Figure 14.

VAST supports both NFS and S3 as a target. VAST supports SMB but from an NCD perspective SMB is only supported as a source currently and not a target.

The next couple of subsections address the creation of a target with these protocols.

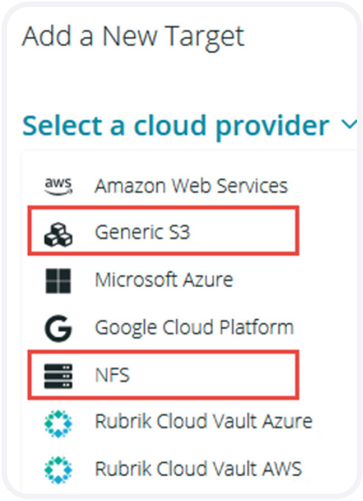


Figure 14 – Selecting VAST Source Protocol

Adding a VAST NFS Target

From Figure 14 click on NFS, and it will bring up the new target window shown in Figure 15. Enter an appropriate target name and then the VAST FQDN (a single IP is sufficient but will limit throughput) that will be used to connect to VAST.

For the mount path enter an NFS path (or alias) that is available on the VAST cluster (See - VAST NFS Configuration). The Data Bucket entry in this NFS context is a subdirectory inside the mount path and it needs to be created ahead of time. Once finished, click **Save Credentials**. The destination will now appear as shown in Figure 13.

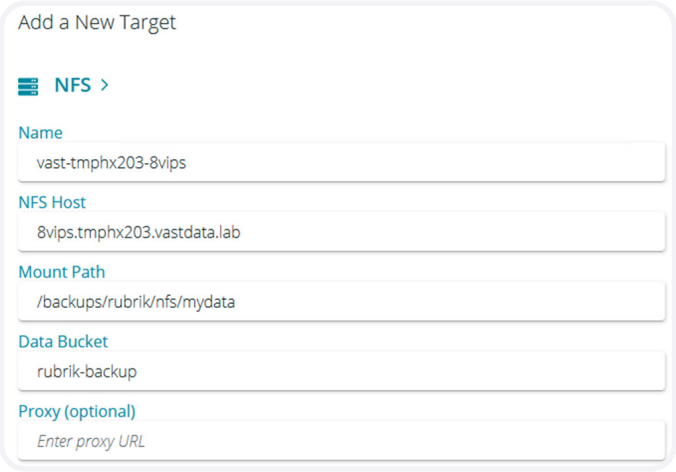


Figure 15 – Configuring VAST NFS Target

Adding a VAST S3 Target

Selecting **Generic S3** back in Figure 14 brings up the new target window for **Generic S3** as shown in Figure 16.

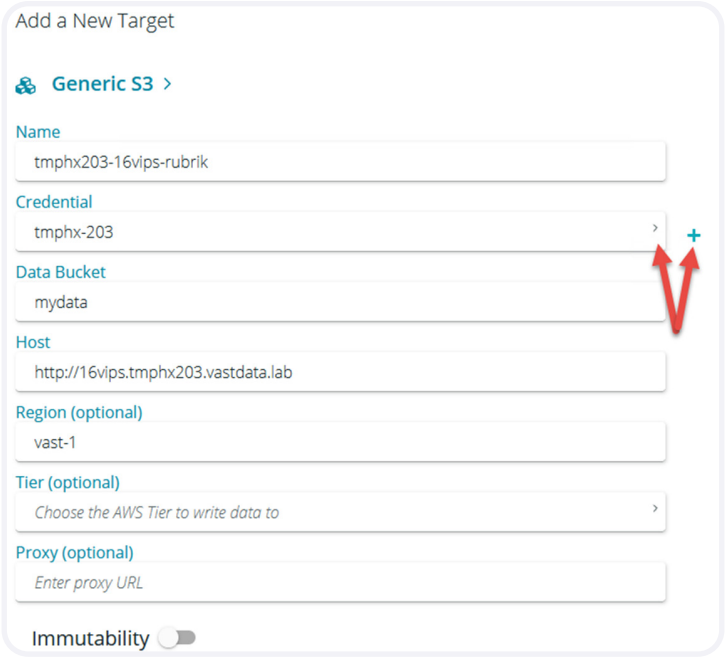


Figure 16 – Configuring VAST S3 Endpoint

To be able to fill out this form properly, a VAST user with S3 credentials will be required. Also, for greater management ease, the VAST best practice of using an endpoint should be considered. Both are addressed in the section – VAST S3 Configuration.

Enter an appropriate name in the first field. Since it's possible that there could be multiple entries to the same VAST cluster through varying FQDNs, it may be prudent to name the destination with appropriate nomenclature to help identify the target such as the FQDN or the S3 user. The example in the figure shows a name that highlights the VAST cluster (tmphx203), the VIP Pool (16vips) and the user (rubrik).

When entering a credential in the form, the pull-down menu allows for choosing already configured users. If a new user needs to be created click on the plus symbol and a wizard will start (Figure 17). When creating an S3 user for VAST always select the Amazon Web Services category.

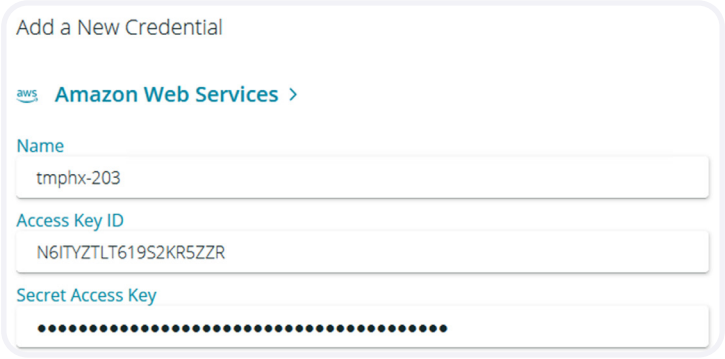


Figure 17 – Adding VAST S3 Credentials

The bucket field in previous versions allowed NCD to create the bucket automatically. This functionality was not available while crafting this document, so a bucket may need to be created ahead of time.

The **Host** field is the VAST FQDN or again a single IP would work. The **Region** field is listed as '**optional**'. In testing, the region was never required but, in some environments, not having this set can lead to noisy logs. If this is an issue, then use the region – **VAST-1** or **vast-1**.

When using TLS over HTTPS for secure data transfer over the wire, select the **Verify SSL** check box. See the SSL Certificates section for information on setting up a proper CA certificate. After filling out all the fields click **on Save Credentials.** The VAST S3 target will now appear in the available list of destinations as shown back in Figure 13.

Modifying VAST Targets and Credentials

Adding or modifying users can be done by clicking on the **Targets** button and selecting **Accounts & Credentials** (Figure 13). This will bring up the **Credentials** windows (Figure 18) where all the existing users will be displayed. For any VAST user that is created, select Amazon Web Services as the cloud provider. For organizational best practice assign the user as a ***bucket creator*** in a VAST Endpoint (See – Creating an S3 Endpoint on VAST).

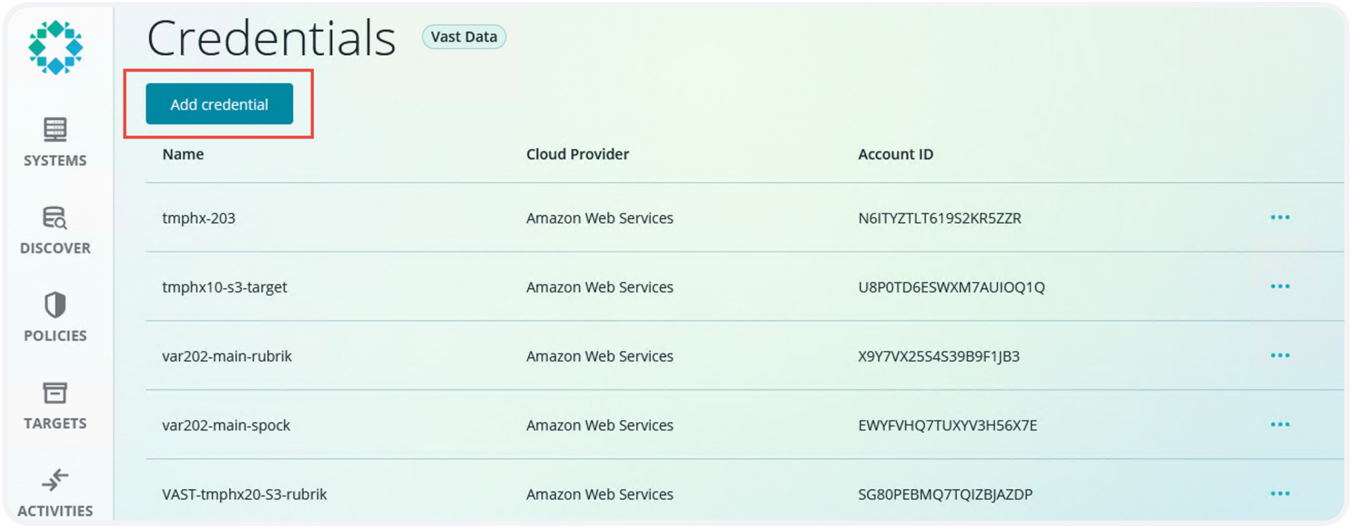


Figure 18 – Cloud Settings for Adding / Editing S3 Targets

Creating a Policy

With either NFS or S3 used to create a VAST target it can now be assigned to a policy which will then be applied to a source system for backup. To begin creating a policy click on **Policies** in the left windowpane and then select **Create Policy** (Figure 19).

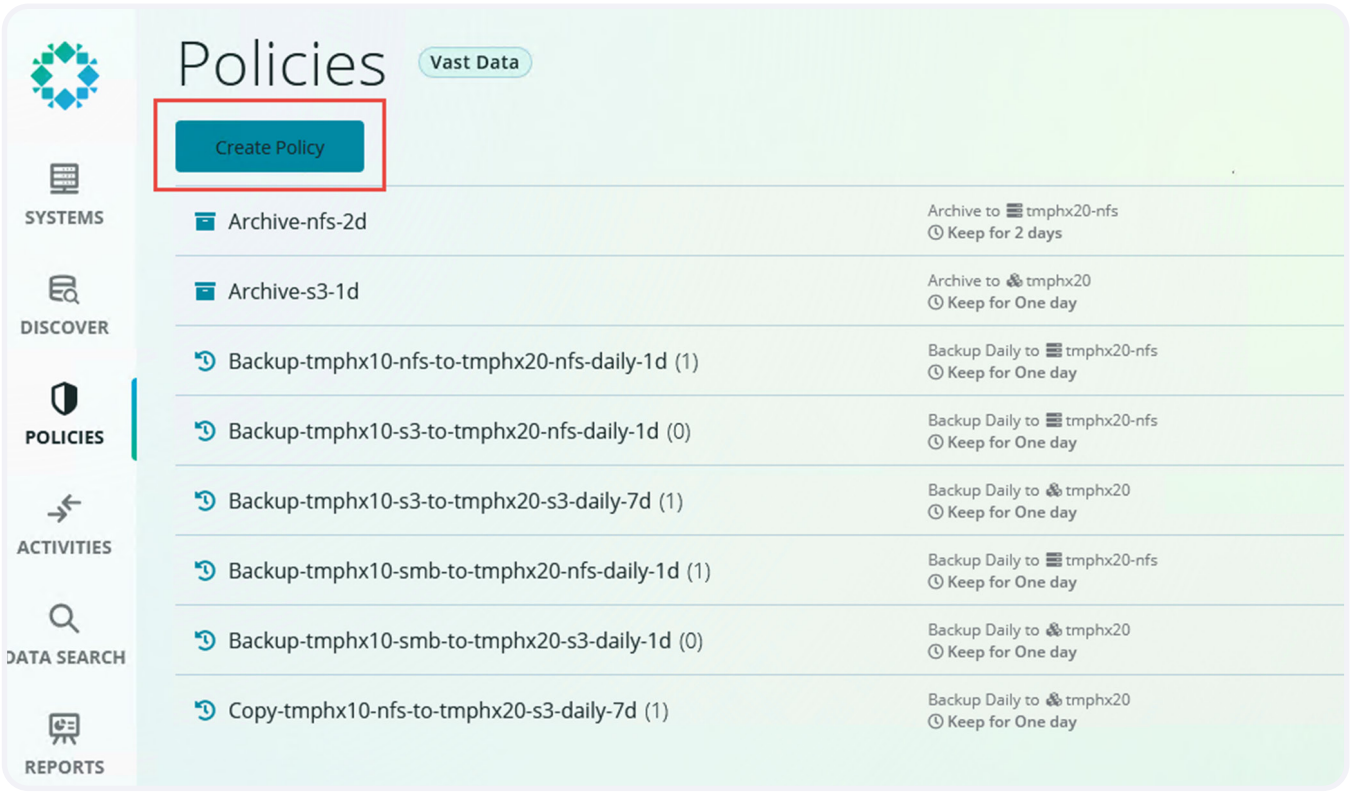


Figure 19 – Create Policy

This will pop up a **Create a policy** window (Figure 20). In this example **Backup** will be used as the policy type.

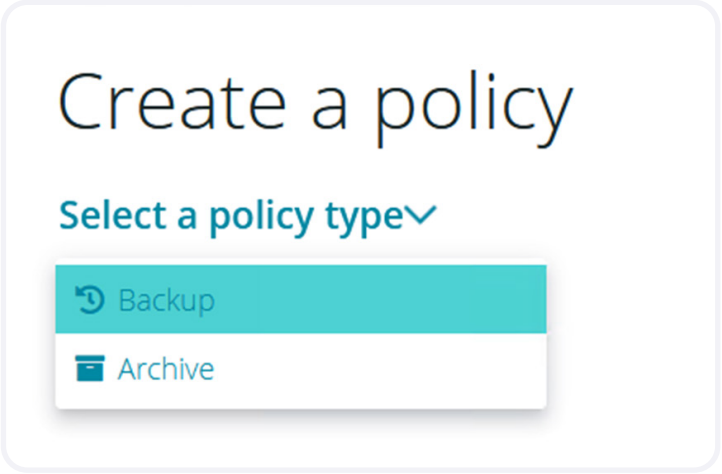


Figure 20 – Selecting Policy Type

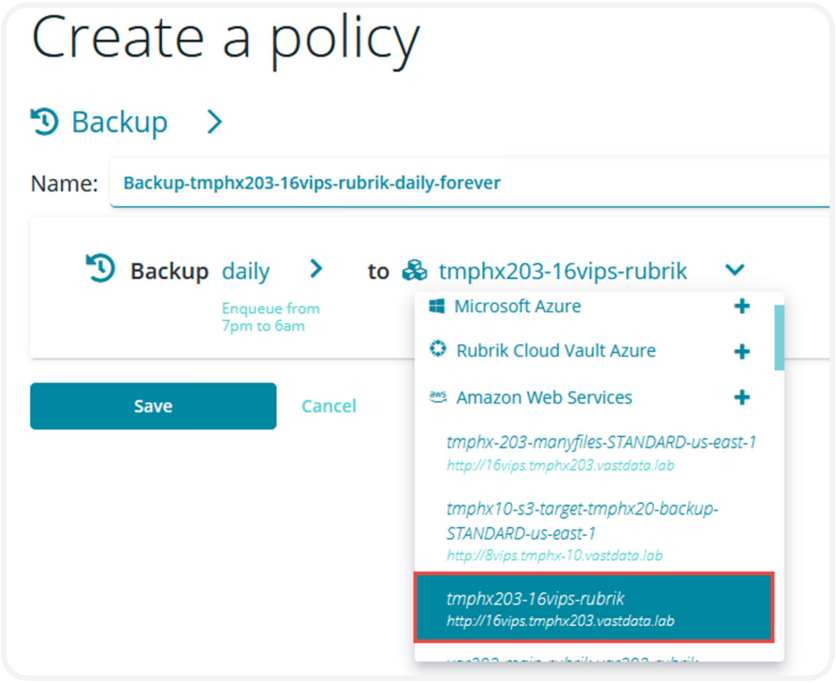


Figure 21 – Selecting Configured VAST S3 Destination

The **Backup** configuration window will now appear as in Figure 21. A default name is displayed, and changes based on the input of the pull-down menus. The default name, however, is just a suggestion, and was changed in this example. To add VAST, click on **Select a destination** and a list of compatible destinations appears. Select one of the configured VAST targets and then choose the retention time from the pull-down menu on the right (Figure 22) and when finished click **Save**.

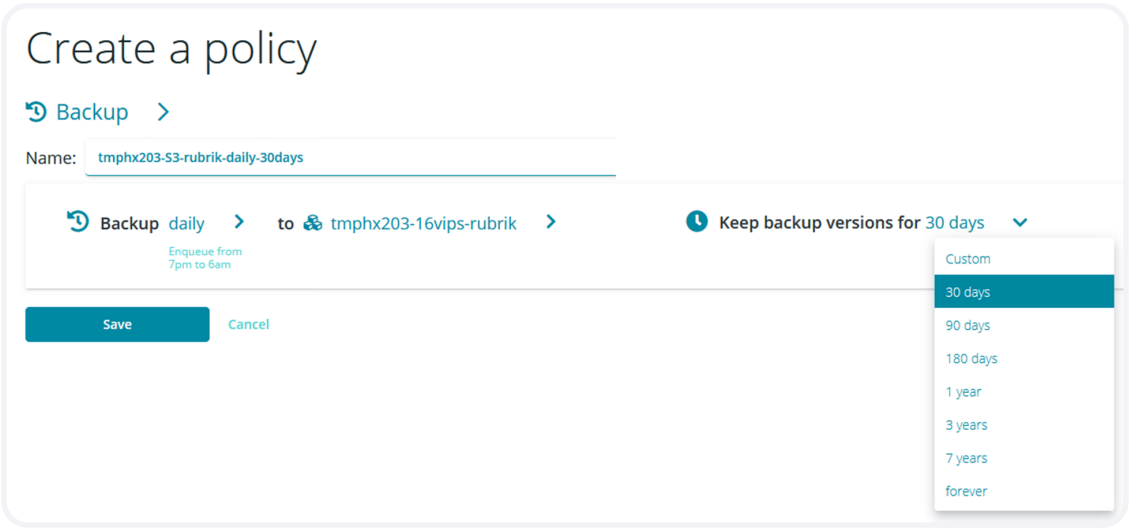


Figure 22 – Selecting Retention Time

This policy (Figure 23) is now available to assign as a backup location for any NFS directory, S3 bucket or SMB folder that has been added in the **Systems** section. As a reminder Rubrik NCD can move data between any supported protocol – From NFS, SMB, S3 as a source to NFS and S3 as the target.

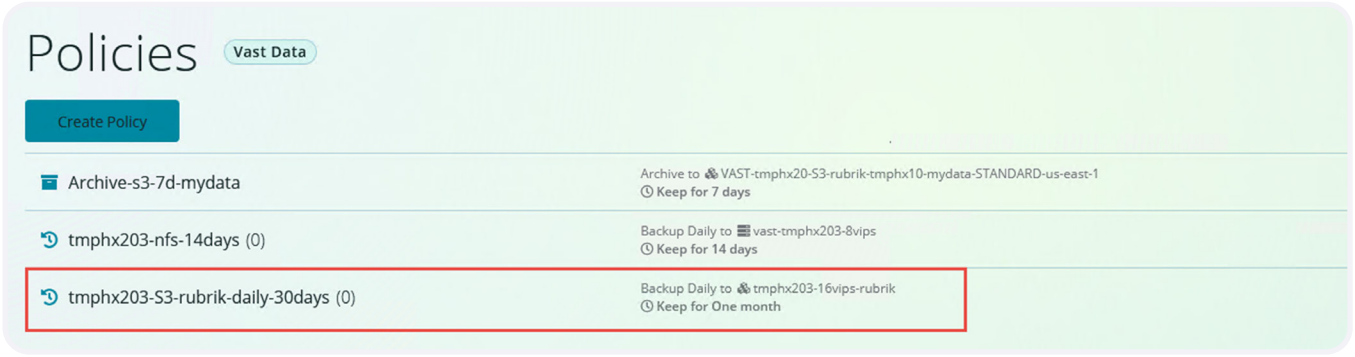


Figure 23 – New Policy is Listed

Running a Backup/Archive

We'll tackle both assigning the policy and running a backup in this section. First, to assign a policy, click on **Systems** in the left panel and this will show the list of available systems (Figure 24). Again, this is the list of systems with source data to be backed up. In this example the source is a VAST NFS export - click on the upside-down caret as shown in Figure 24.

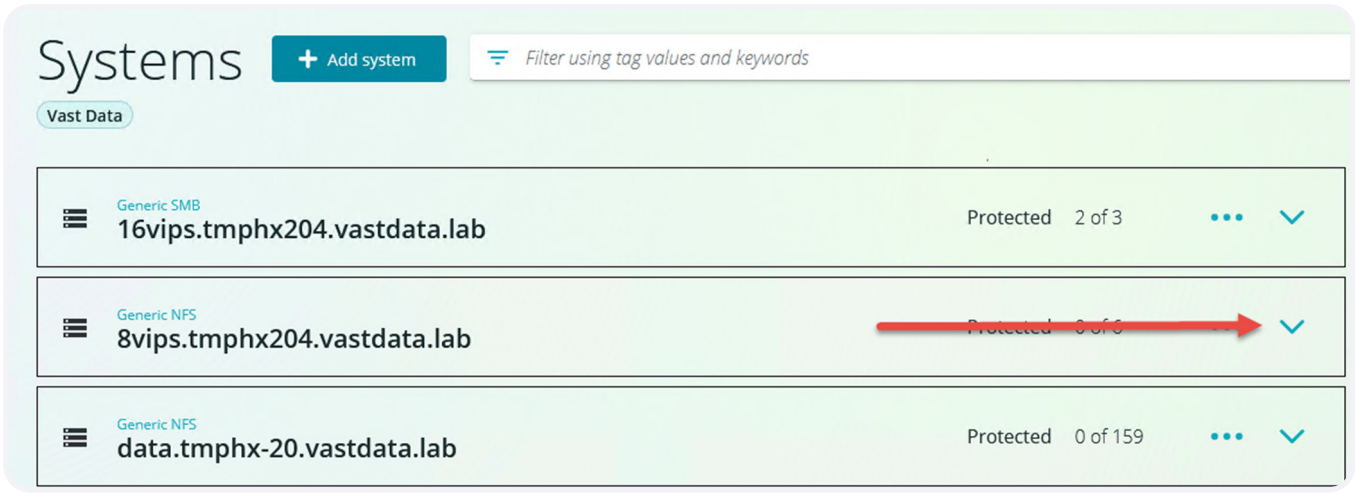


Figure 24 – Selecting Source System for Assigning Policy

This will now list all the available directories, folders or buckets within that service point depending on the protocol (Figure 25). Select the directory for backup and this will highlight its current protection policy. As highlighted in Figure 25 the directory has no policy assigned to it.

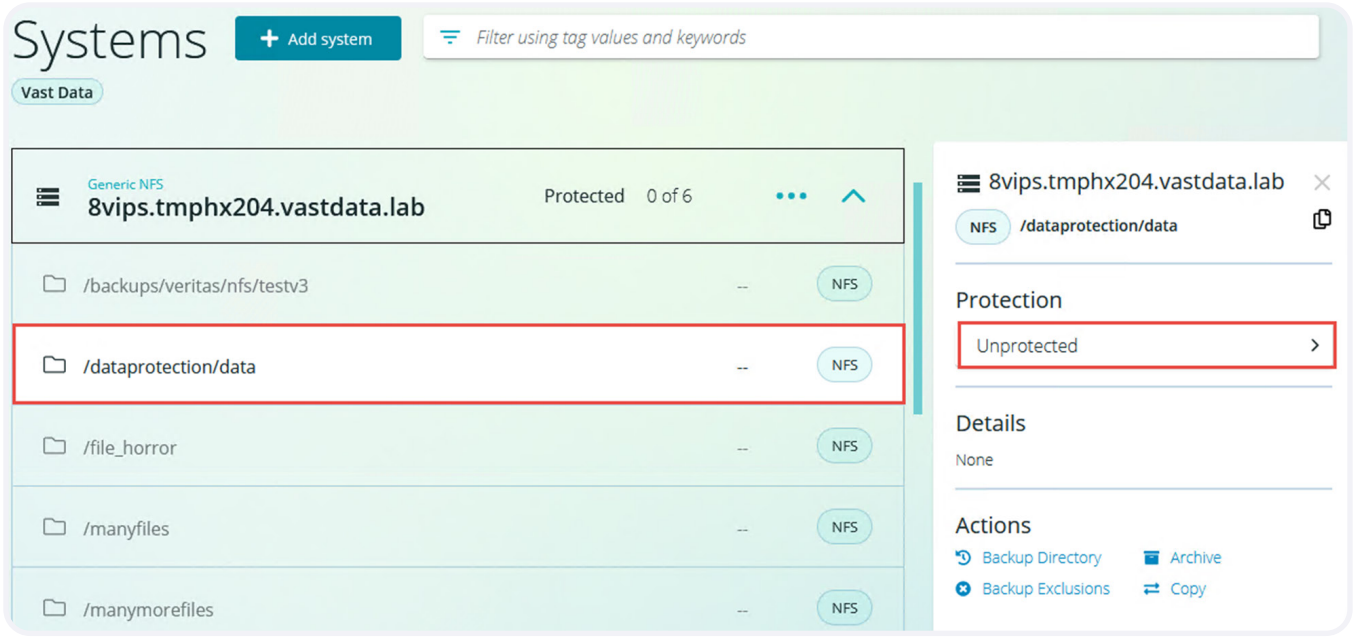


Figure 25 – Selecting Directory to Backup

Clicking on the pull-down menu will show a list of all the existing policies. Figure 26 shows that the policy that was created previously is now assigned to the source system location. In the policy configuration there is a setting to **Enqueue** the policy to a particular backup window. This can be modified to ensure that the backup happens when desired.

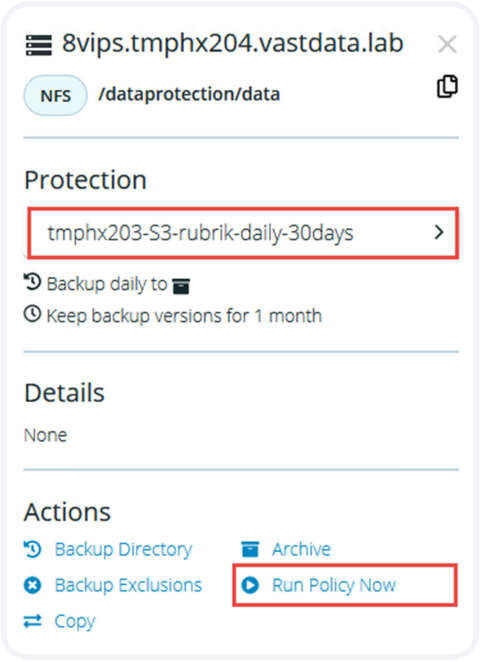


Figure 26 – Policy Assigned to System Export

To back up the source data immediately simply click on **Run Policy Now** as shown in Figure 26.

There are myriad ways to view the current (and previous) activity of the backup policy or system. The most obvious way is by clicking the main Activities link in the left panel. However, now that this policy has now been run at least once against this system, a couple of other options are now available, **View Activity** or **View Status**, as shown in Figure 27.

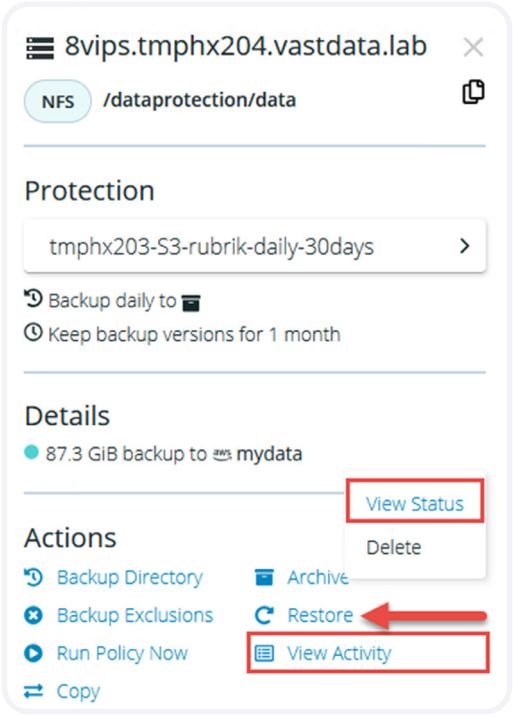


Figure 27 – View Policy Activity

Regardless of method for navigating to the Activities windows, Figure 28 highlights the activity of the specified policy on that source. Notice that this is incremental forever after the first full.

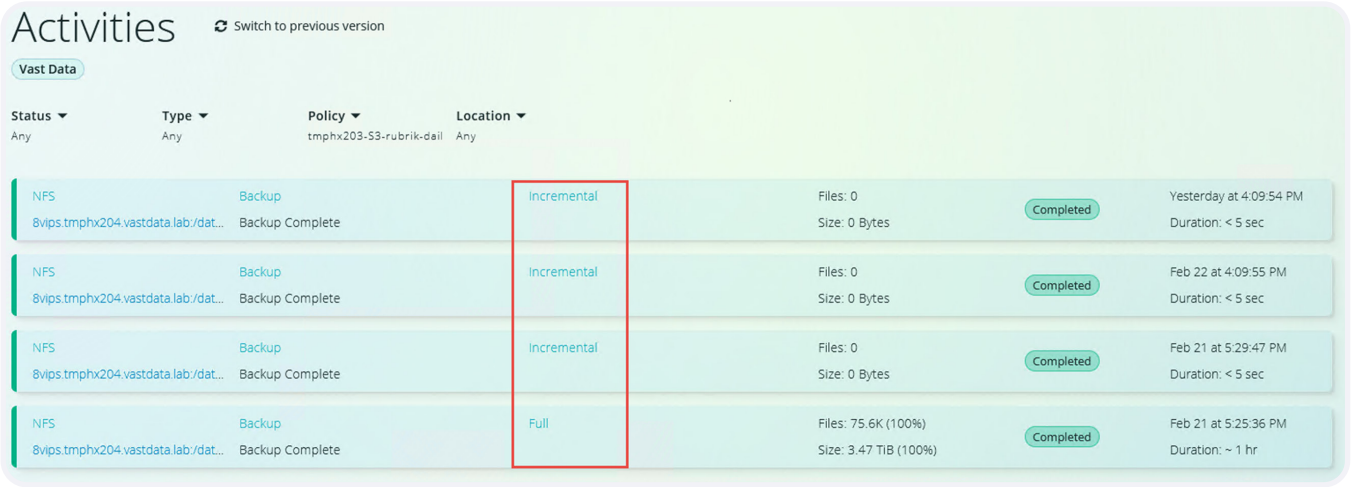


Figure 28 – Activity of Policy

The details of the first full backup can be seen in Figure 29. It shows the source and targets of the data, size and number of files.

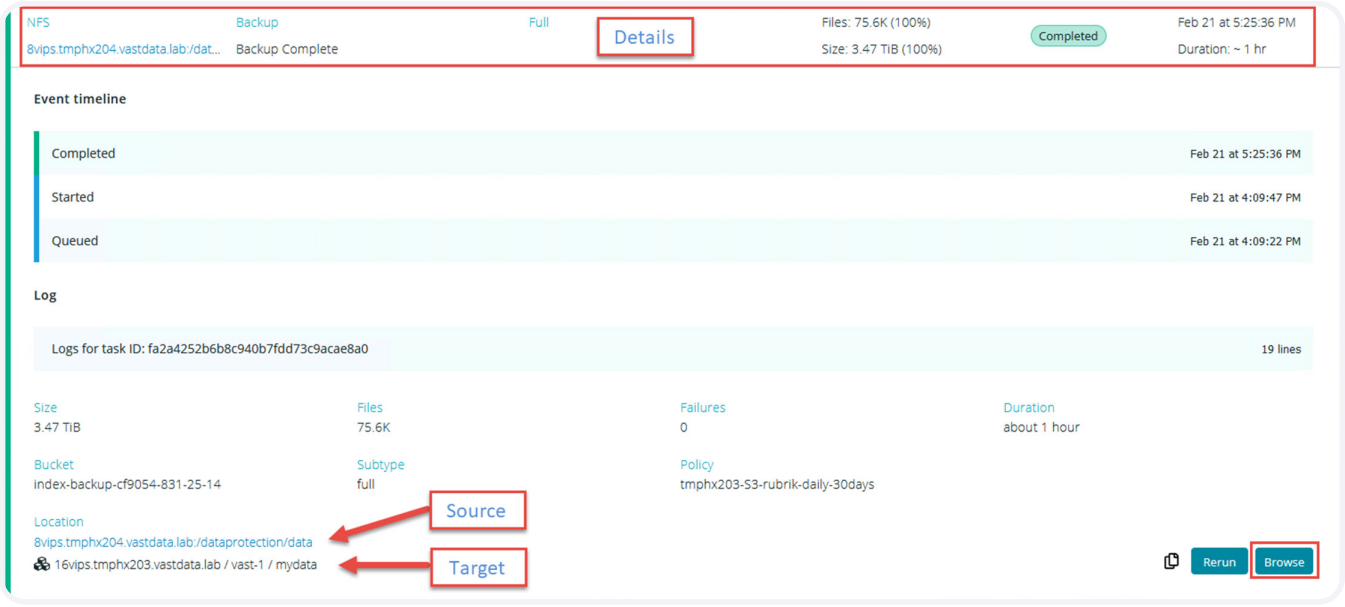


Figure 29 – Status of Backup/Archive in Activities

Running a Restore

Restoring data from any of the restore points is straight forward. There are several ways to start a restore process. Looking back at Figure 27, it shows a restore link within the source system's menu. Figure 29 shows it as a **Browse** button in the lower right of the backup activity window. No matter which method is used, the **Data Search** window will appear as shown in Figure 30. Since this example does not have any additional incremental data, all the restore points contain the same original full backup so picking any point to restore from is possible to recover the data.

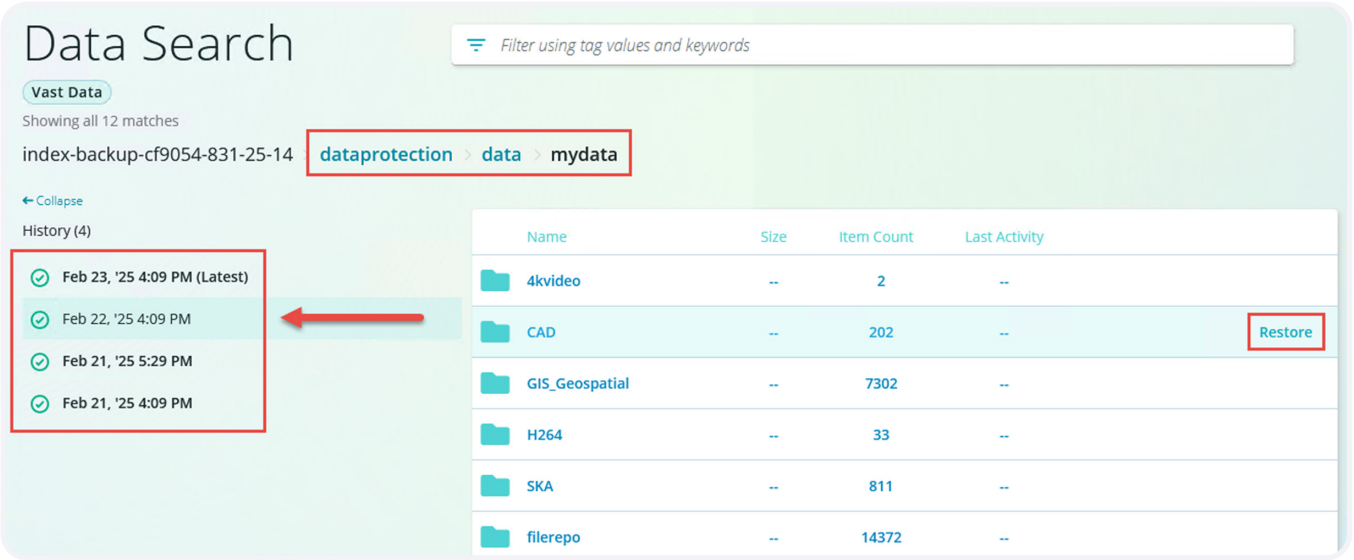


Figure 30 – Selecting Data to Restore

In this example the second most recent incremental is selected. Data can be restored from the top directory of this backup (**/data**) for a full restore but here only a specific subfolder (**/data/mydata/CAD**) is being restored. The directory structure is displayed at the top as shown in Figure 30 so it's easy to know where to find any particular file or folder.

Clicking on the **Restore** button brings up the **Restore Data** window shown in Figure 31. Restores can be written to existing system devices already configured on Rubrik NCD or to a custom location. The example in the figure is restoring back to the original VAST cluster but to a new view that is already configured (**/dataprotection/rubrik/restores**).

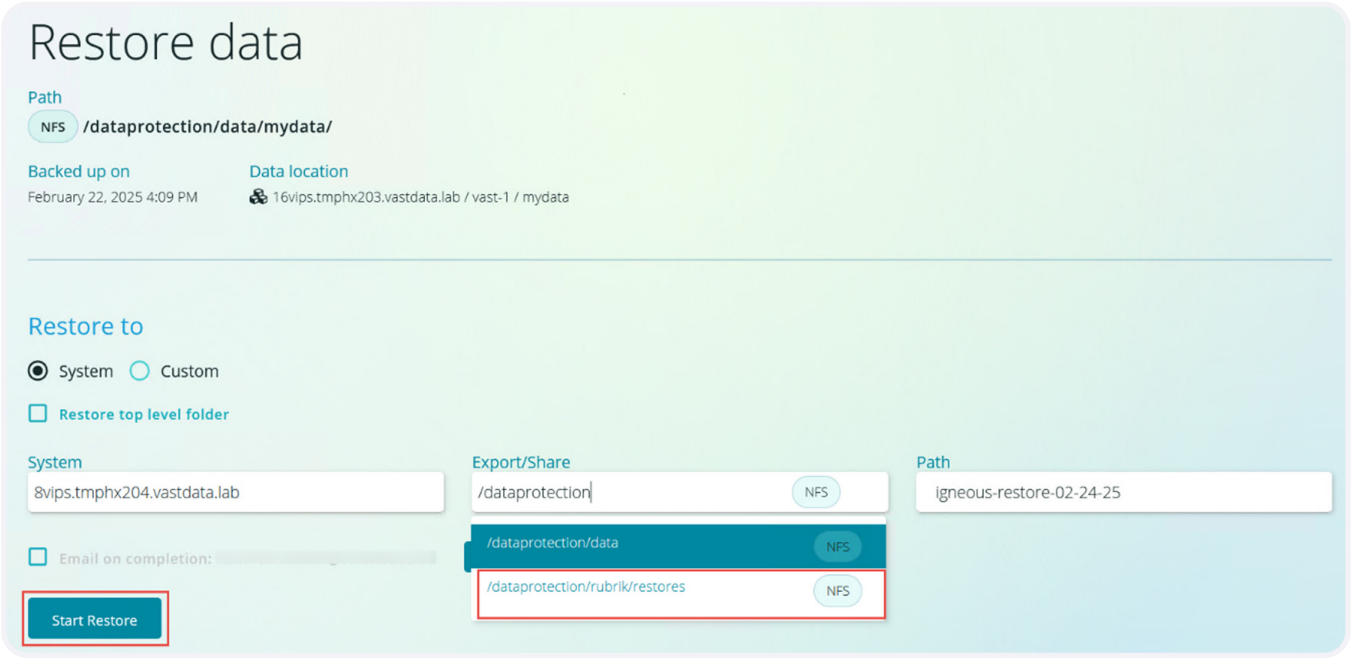


Figure 31 – Restore Data

Once the restore location has been chosen simply click the **Start Restore** button. The restore can be monitored from the Activities windows shown in Figure 32 and Figure 33.

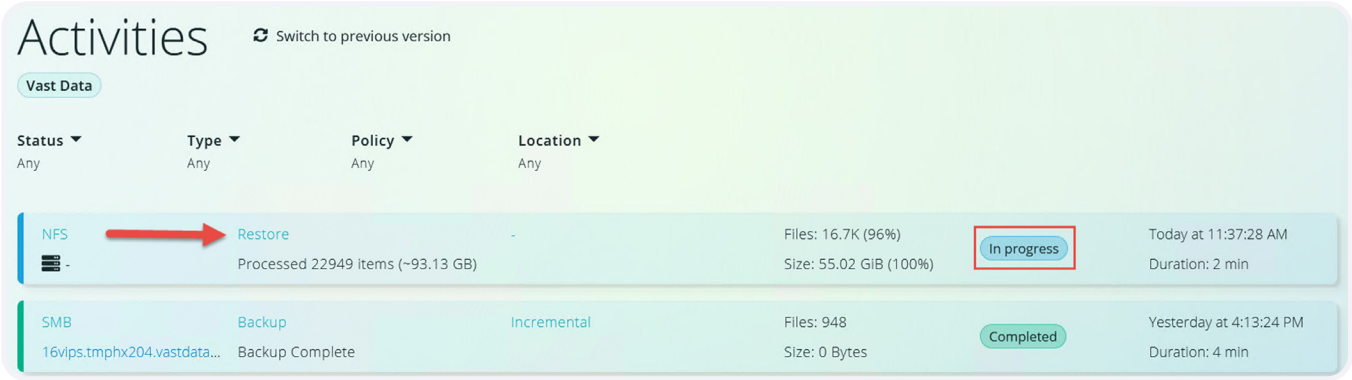


Figure 32 – Restore in Activities Window

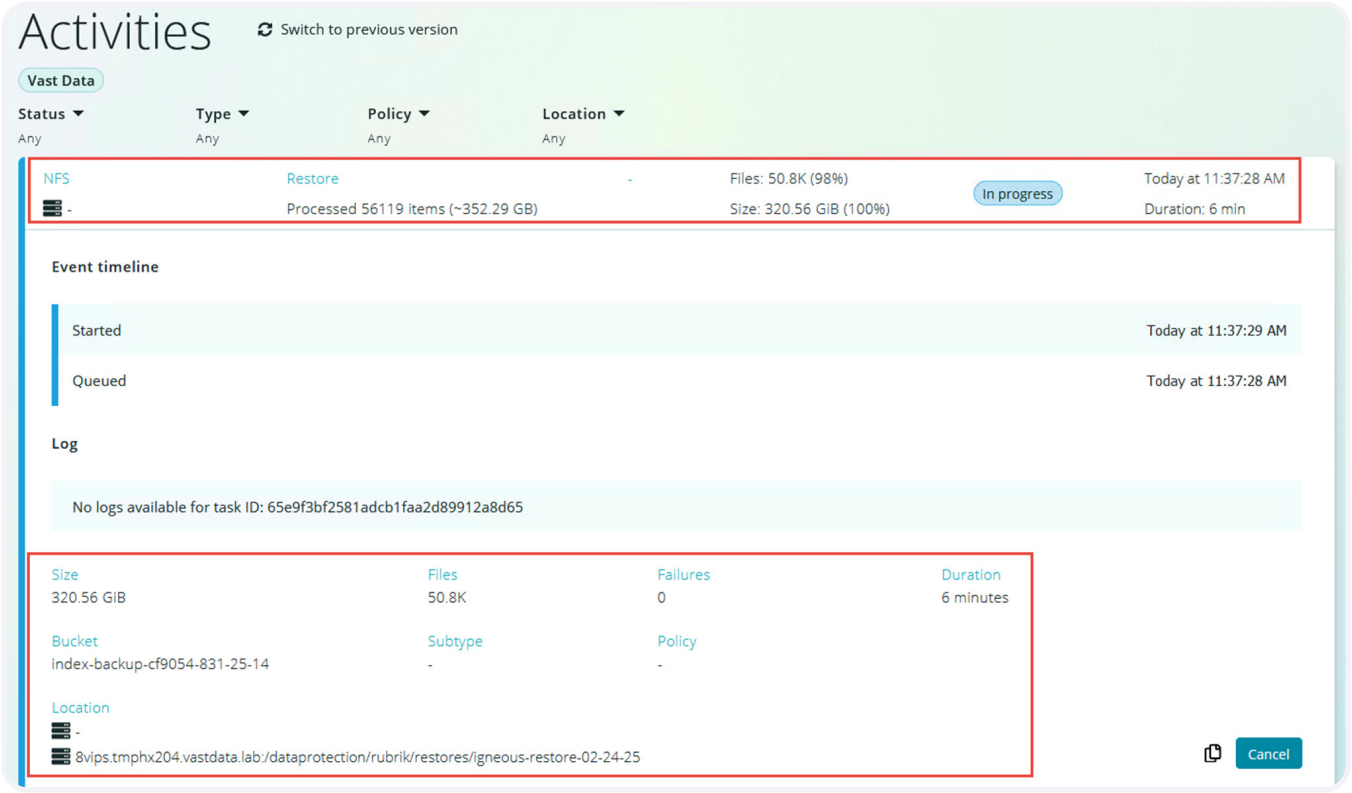
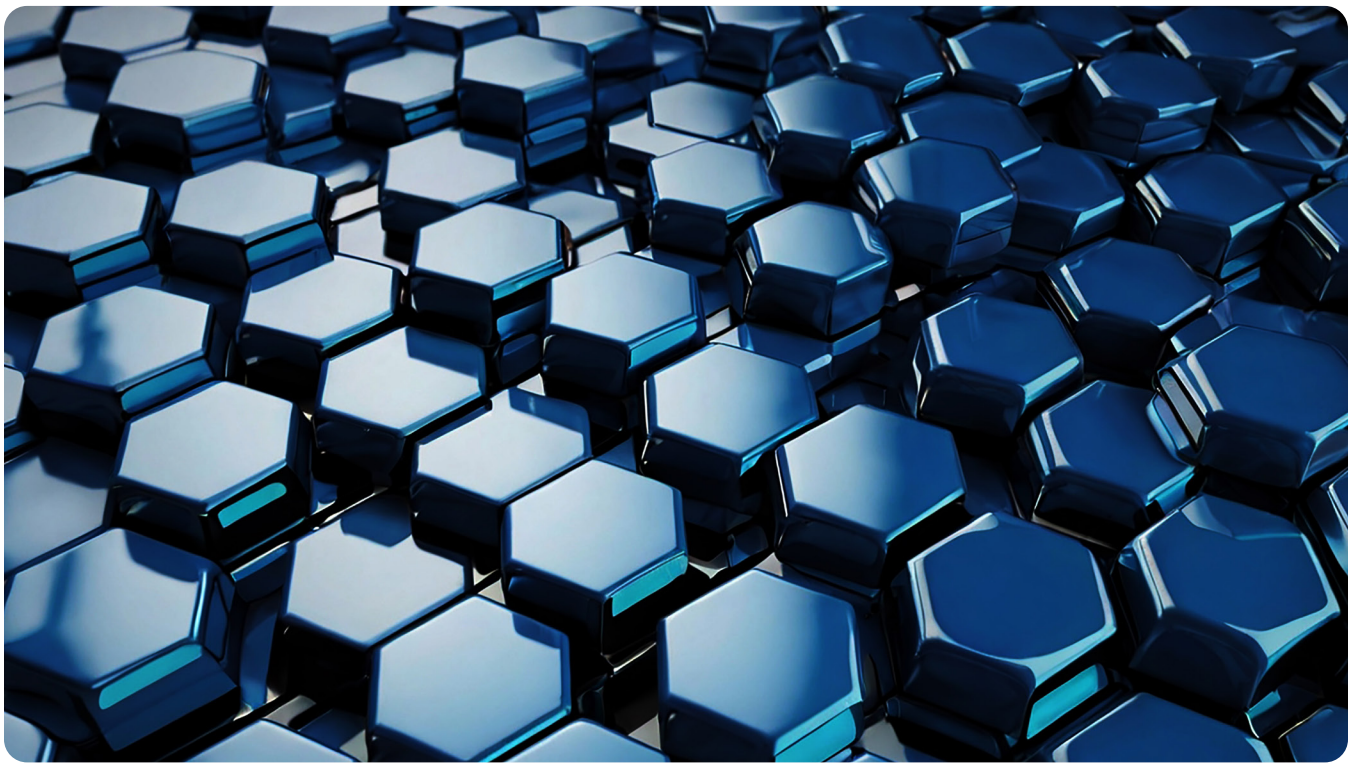


Figure 33 – More Details on Restore



VAST Configuration

During the process of adding a policy and targeting a VAST cluster there are some prerequisites that will need to have been created on the VAST side. If the target is an S3 bucket the S3 user credentials (secret and access keys) will be needed. If the target is an NFS share, then a proper VAST view (and optional alias) will be needed. So, this section covers the VAST side configuration of these components.

VAST S3 Configuration

There are several components to configure with a VAST system for S3. Not all of them are mandatory but we will look at each of them.

Creating an S3 User

To create an S3 user with the VAST UI, go to the Dashboard and then User Management and select Users (Figure 34).

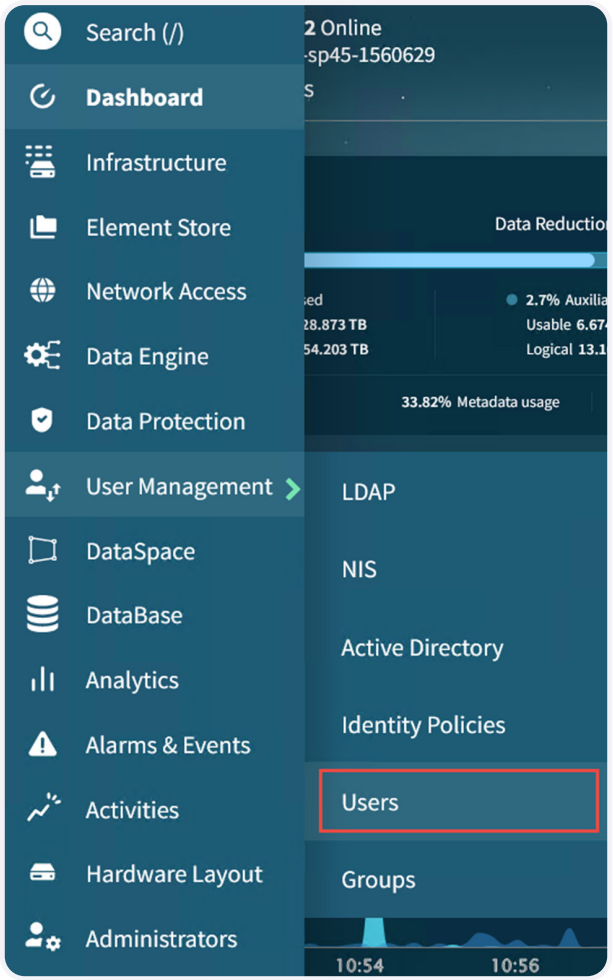


Figure 34 – Selecting Users Category

The Add User window appears (Figure 35) and a local username needs to be entered along with any desired UID or an Active Directory user can be used if the VAST cluster is integrated with a domain. The user in this example was not granted full S3 credentials as it is not a requirement for basic bucket creation which is all that is needed to create an S3 target for a backup policy.

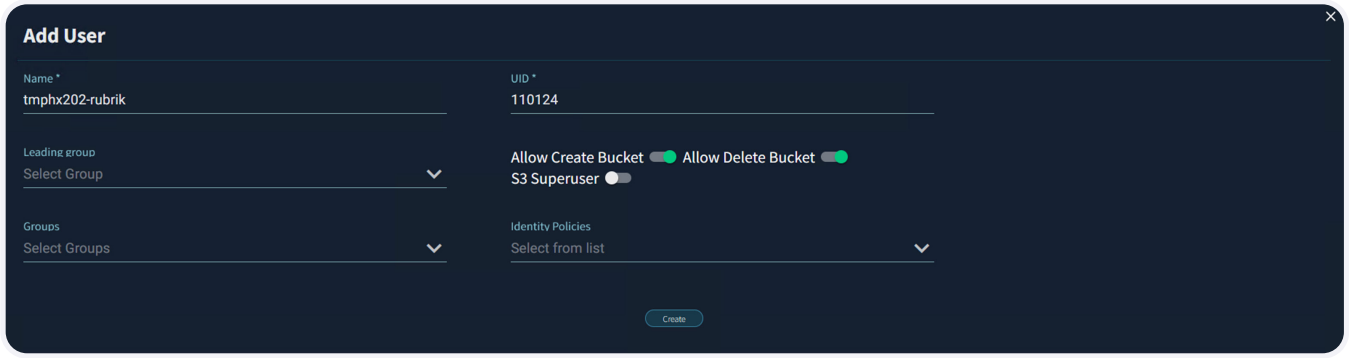


Figure 35 – Creating a New User for S3

After clicking Create the user will show up in the User Management window (Figure 36).

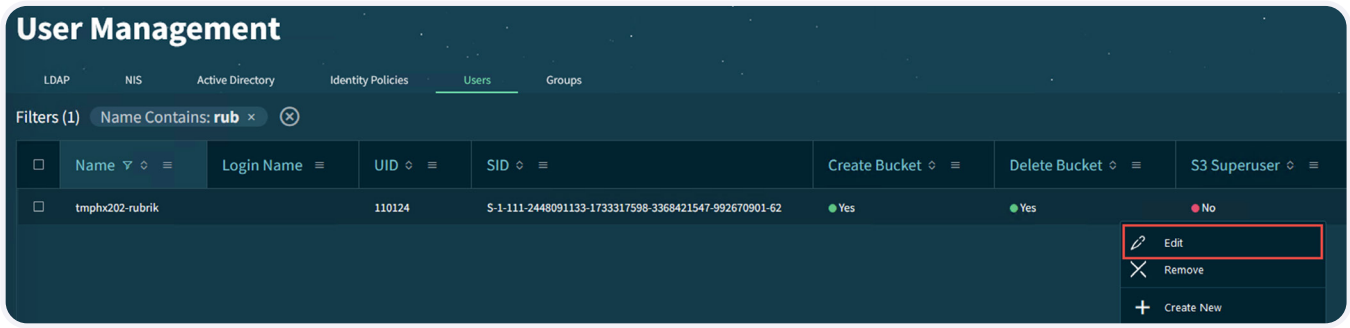


Figure 36 – User Created

Now that the user is created it needs to be edited to create and capture the active and secret keys. Right clicking anywhere on the user will bring up a drop-down menu (Figure 36). Simply select **Edit**.

In the **Update User** window (Figure 37) click on the Create new key button to create a new Access Key. Now copy the active key and secret key some place safe for later. This is the only time when the secret key will be obtainable so be sure to copy it down in some safe location. Click **Update** to close the window.

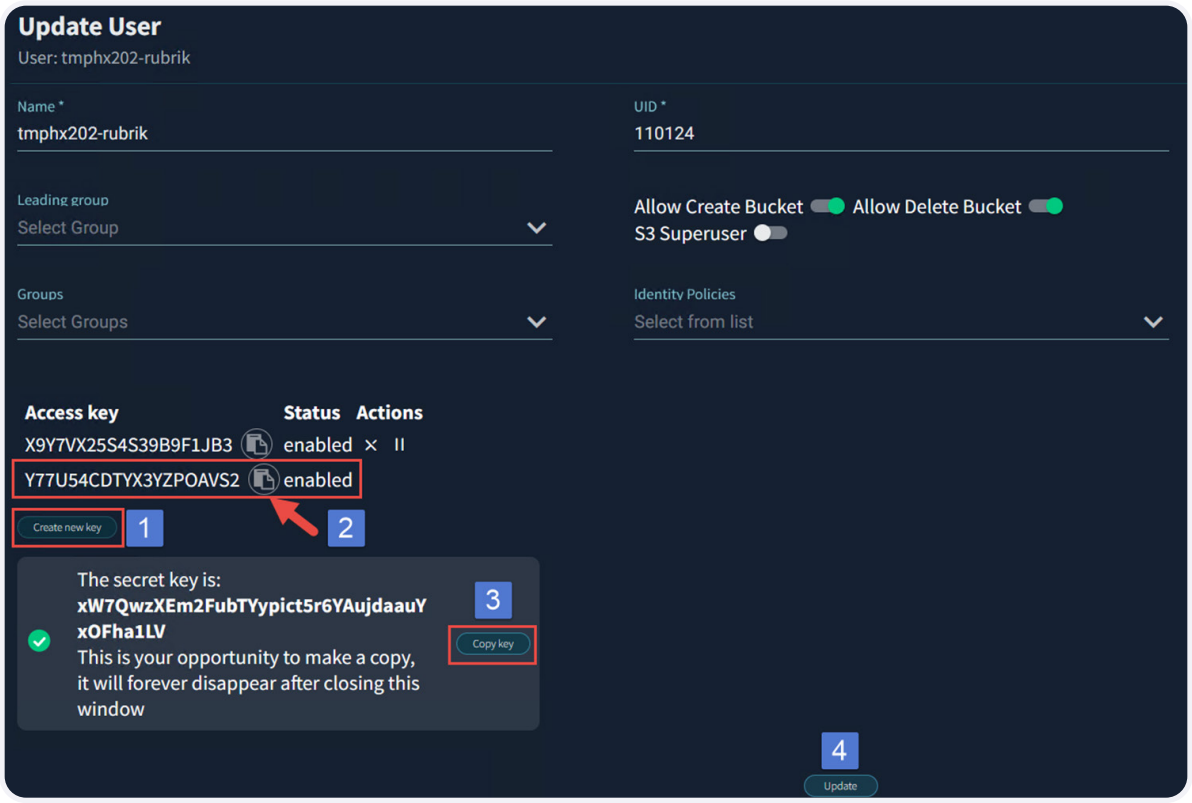


Figure 37 – Capturing Active and Secret Keys

The user can now be used in the creation of a bucket or view within the VAST UI.

Note

This example showed a local user. As mentioned previously, a domain user can also be used if the VAST cluster has been joined to a domain. Any domain user can be searched for and modified in a similar manner for secret and access keys.

Creating an S3 Endpoint on VAST

A good best practice for organizational purposes and ease of use with creating buckets within the Rubrik NCD workflow is to create an endpoint. All new buckets created with the same user assigned to the endpoint will be placed in an orderly manner underneath that endpoint. If one were to just create buckets without an endpoint, they would all land in the root directory of the VAST cluster, and it could become onerous to manage. Buckets can be created within the VAST UI, the Rubrik NCD workflow (version specific), or through any S3 utility.

To create an endpoint, go to the **Element Store** within the VAST UI, select the **Views** tab and then click ***Create View*** as highlighted in the upper right in Figure 38.

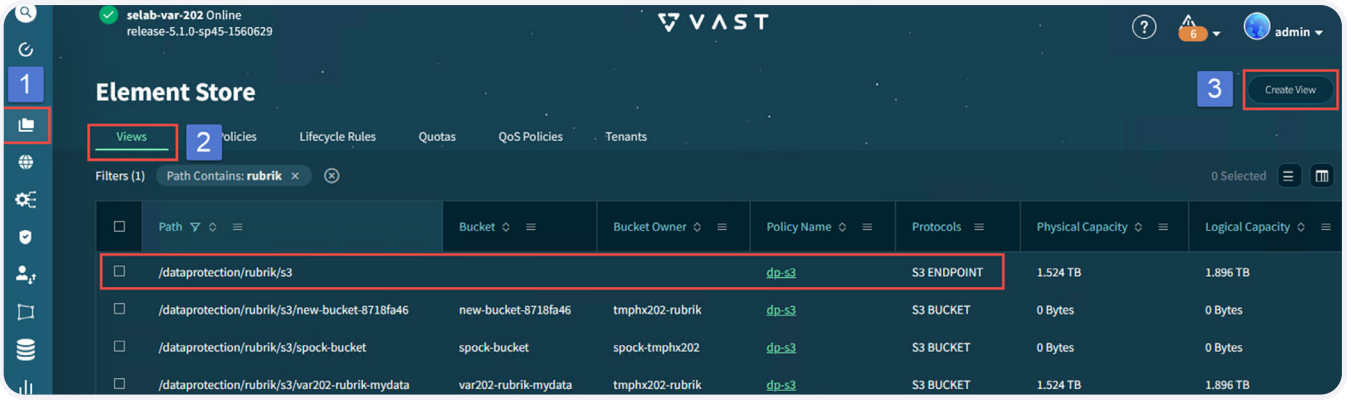


Figure 38 – Creating an S3 Endpoint on VAST

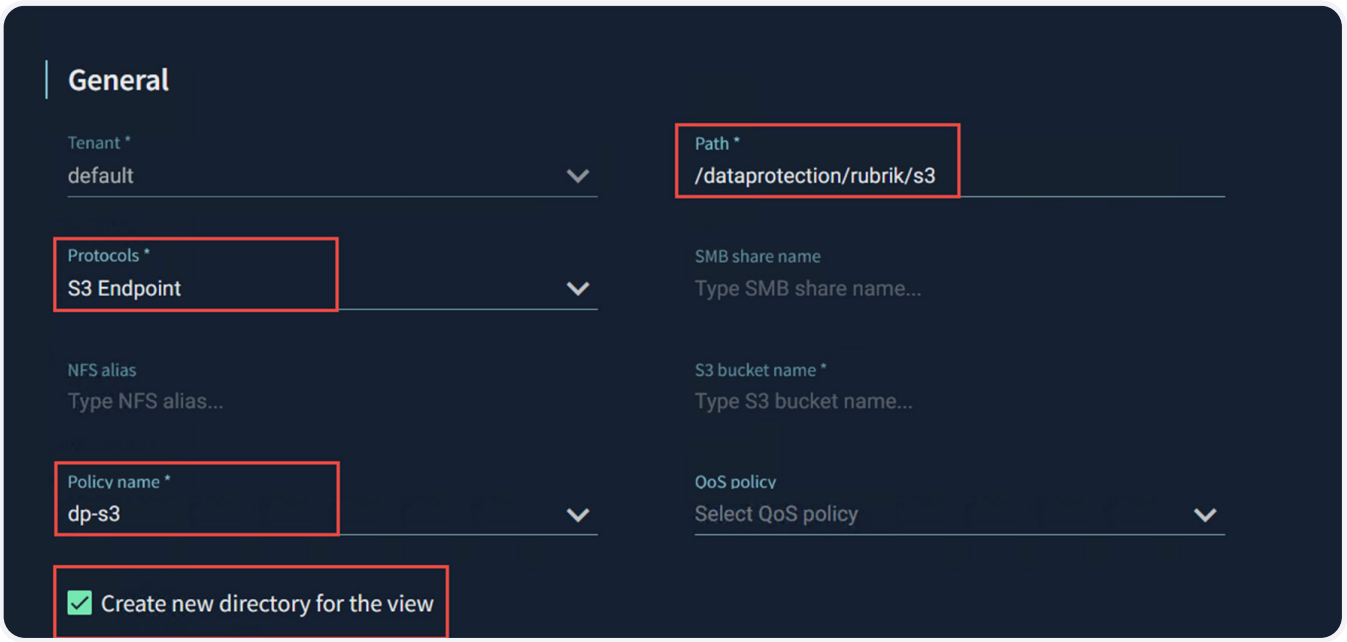


Figure 39 – S3 Endpoint General Details

This will open an Add View wizard as shown in Figure 39. Here the process is similar to creating a bucket, but the Endpoint protocol is selected instead. (Don't forget to slide the Create Directory toggle over). In the S3 section of this wizard (Figure 40) ensure the appropriate user is given access for **Bucket creators**. Click Create when finished. This will show the endpoint as in Figure 38.

> **Note**
>
> An endpoint can have multiple users added as **Bucket Creators**, but a user can only be added to one endpoint.

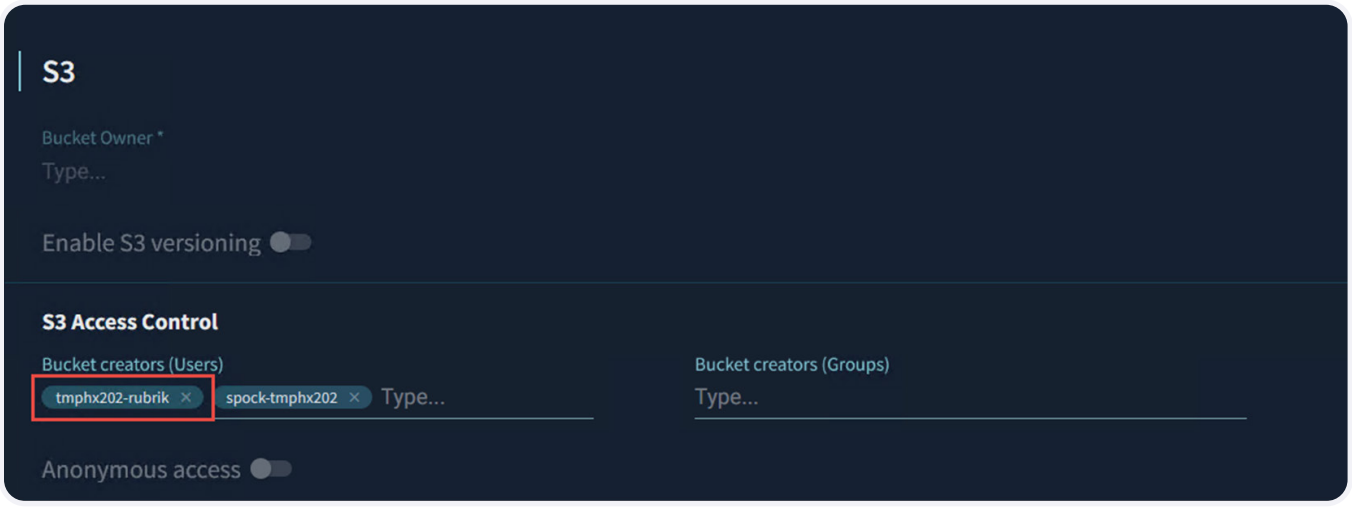


Figure 40 – S3 Endpoint Bucket Creators

Again, the benefit of having an endpoint is that during bucket creation all endpoint properties will propagate to the buckets and cleanly place them under the endpoint path.

For example, the endpoint that was just created here has a path of

```
/dataprotection/rubrik/s3
```

A new bucket created with the name ***– bucket1*** will show up in the VAST UI like this:

```
/dataprotection/rubrik/s3/bucket1.
```

This helps to organize the VAST system and becomes especially useful when creating the bucket within the Rubrik policy workflow (version specific) or other S3 utility that allows for bucket creation.

Creating an S3 Bucket

During the NCD backup policy creation workflow (version specific), a VAST bucket can be configured without having to use the VAST UI. However, if a user wants to create one ahead of time from the VAST UI, this section will describe those steps.

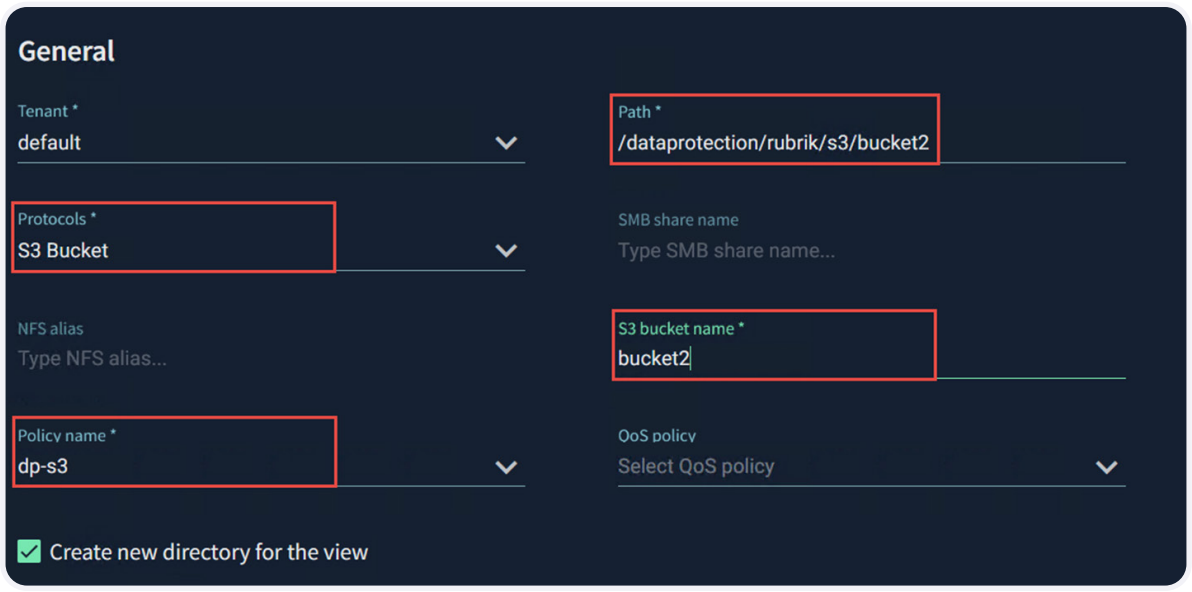


Figure 41 – Creating an S3 Bucket

To create a bucket (Figure 41), mimic the same steps for creating an endpoint shown in Figure 38 and Figure 39 but ensure that only **S3 Bucket** is selected as the protocol. To ensure proper alignment with an existing endpoint make sure the path, policy, and bucket creator follows the same format. The example in Figure 41 shows '**bucket2**' being created with the proper sub-path underneath the endpoint example given earlier. The view (bucket) now shows up in the **Element Store** as shown in Figure 42.

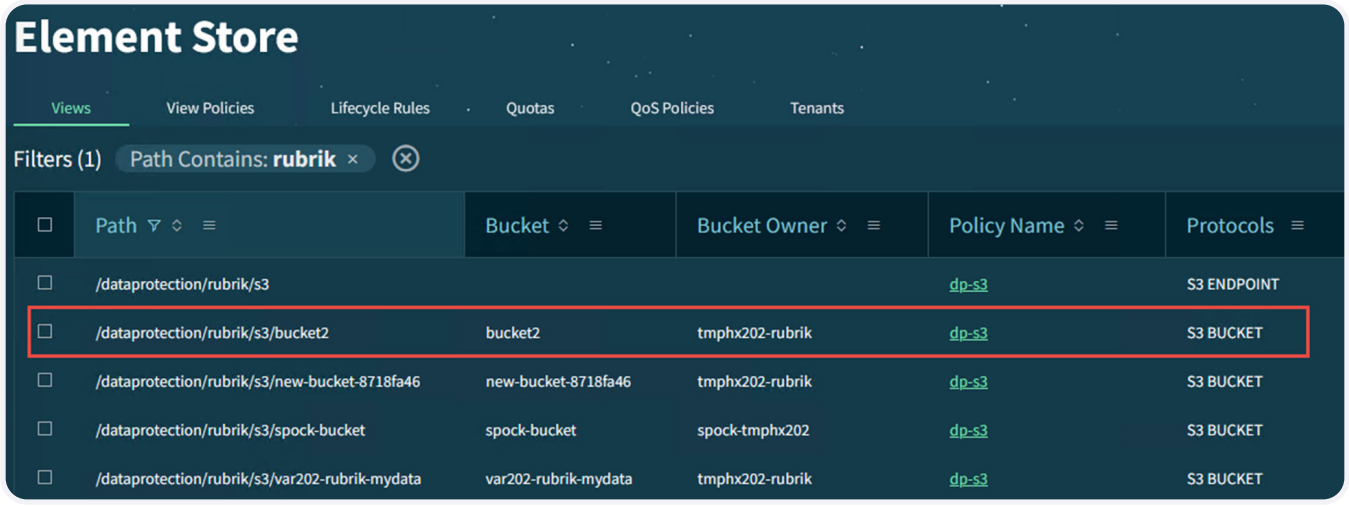


Figure 42 – S3 Bucket Created

With the bucket now created it can be used during the policy creation within the Rubrik NCD UI.

S3 Identity Policy

While there were no issues executing any backup and recovery to and from an S3 bucket without an Identity Policy – it is prudent to cover the concept. Every environment is different so there may be an environment that will need explicit S3 API permissions applied to a user. For users on a VAST cluster, this is achieved through an identity policy. Rubrik NCD refers to them as S3 IAM permissions. Creating an identity policy is simple and once completed should be assigned to a user. To create an identity policy, go to Identity Policy under User Management (Figure 43).

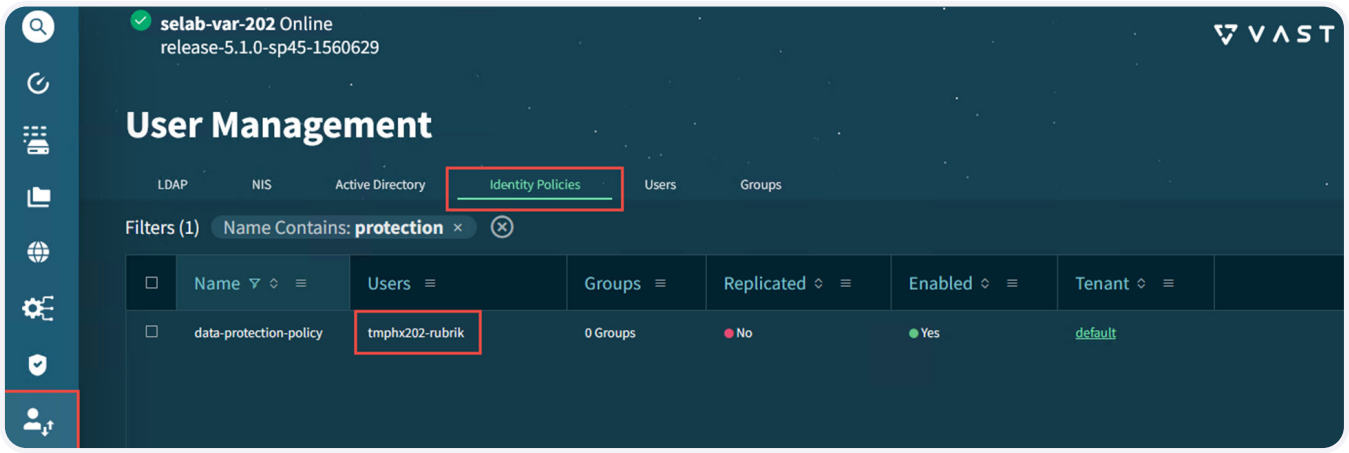


Figure 43 – User Management Identity Policies

Then click **Create Policy** (not shown) and ensure the appropriate S3 API permissions are inserted in JSON code as shown in the example below:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "VisualEditor0",
            "Effect": "Allow",
            "Action": [
                "s3:CreateBucket",
                "s3:GetBucketLocation",
                "s3:GetObject",
                "s3:PutObject",
                "s3:PutObjectTagging",
                "s3:ListBucket",
                "s3:ListAllMyBuckets",
                "s3:PutObjectLegalHold",
                "s3:GetObjectLegalHold",
                "s3:DeleteObject",
                "s3:AbortMultipartUpload",
                "s3:ListBucketMultipartUploads"
            ],
            "Resource": "*"
        }
    ]
}
```

Note

The above code is just used as illustration and each environment will need the appropriate permissions for their needs. Also take special care that the quotes used in the code are straight quotes and not curly quotes.

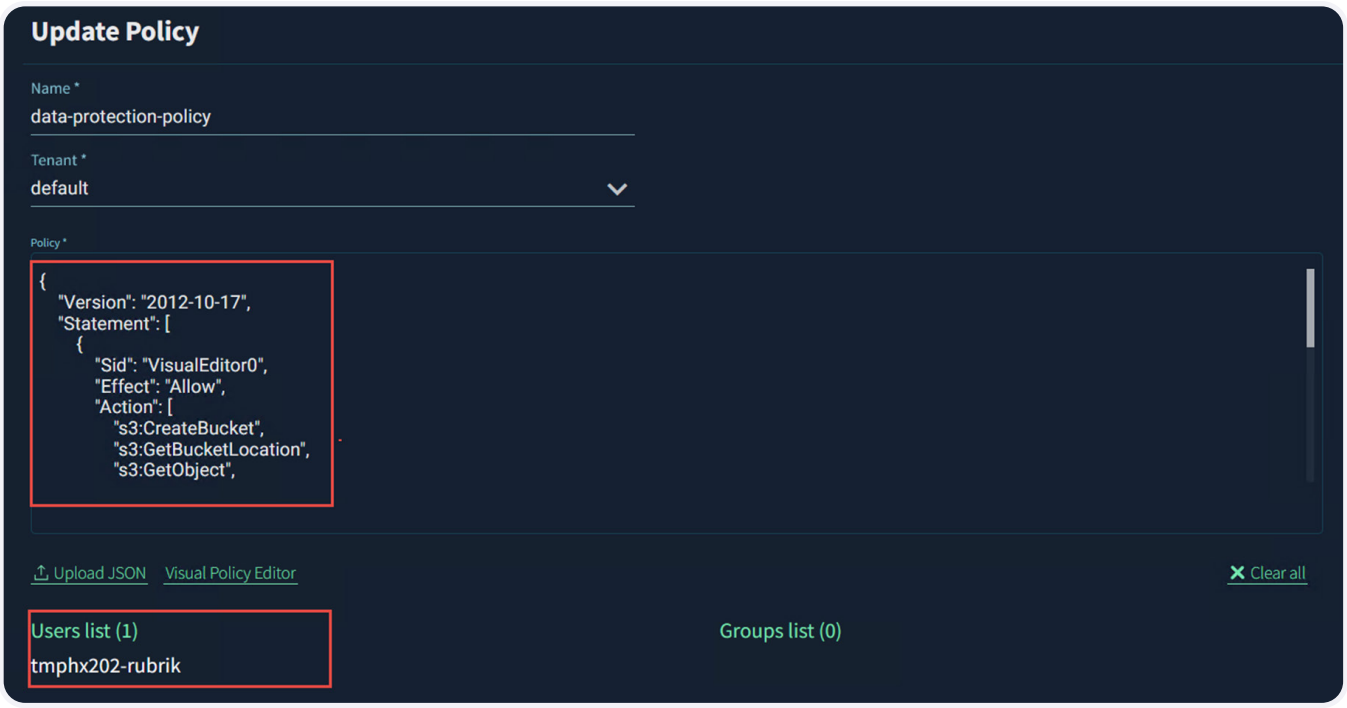


Figure 44 – Assigned Identity Policy

The **User List** shows the users that have this policy applied. (Figure 44) To add a user to the policy, that user needs to be modified. To do that make sure to save the identity policy, then click on the **Users** tab which is right next to **Identity Policies** (Figure 43). Find the user that needs to be assigned to the identity policy. In Figure 45 user **tmphx202-rubrik** is being edited. In the pull-down menu for **Identity Policies** the **data-protection-policy** has been selected which contains all the proper S3 permissions. Simply click on **Update** to complete the edit.

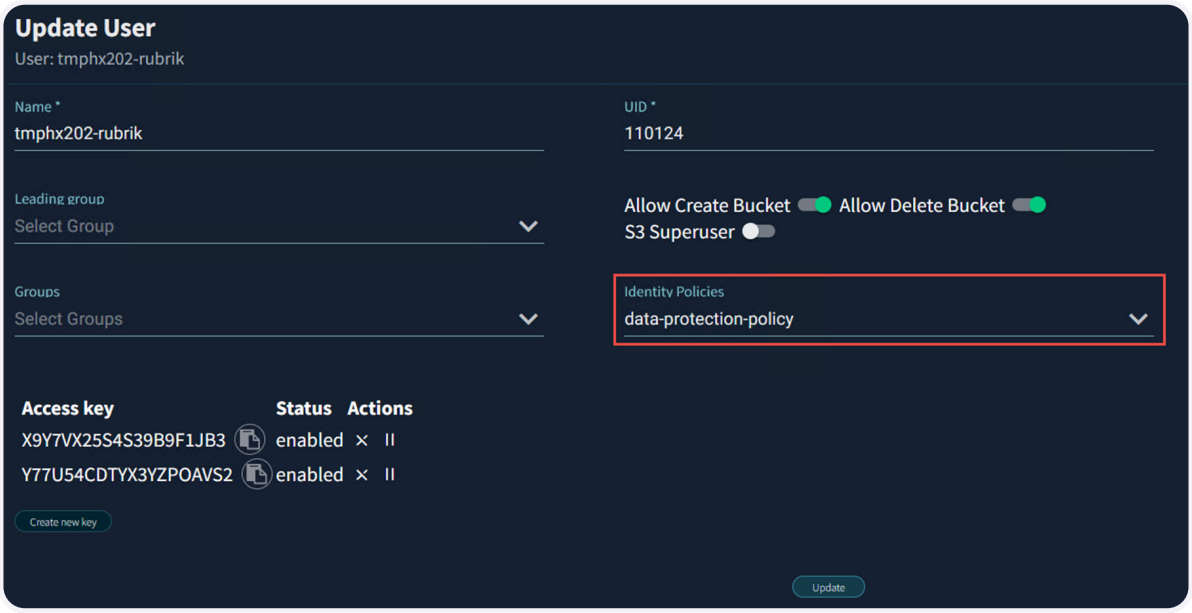


Figure 45 – Update User with Identity Policy

Once the user has the appropriate identity policy assigned, the Identity Policy will look similar to the one in Figure 44.

VAST NFS Configuration

This section addresses configuring an NFS view which is similar to creating other views within VAST. Start by selecting the **Element Store** in the left windowpane, pick **Views** and then select **Create View** as shown back in Figure 38. Creating an NFS view with the wizard only needs a couple of fields populated – the path and policy (Figure 46). Remember to check the box if this is a new directory. An alias could also be used here to further simplify the path within Rubrik NCD. When finished, simply click on **Create**.

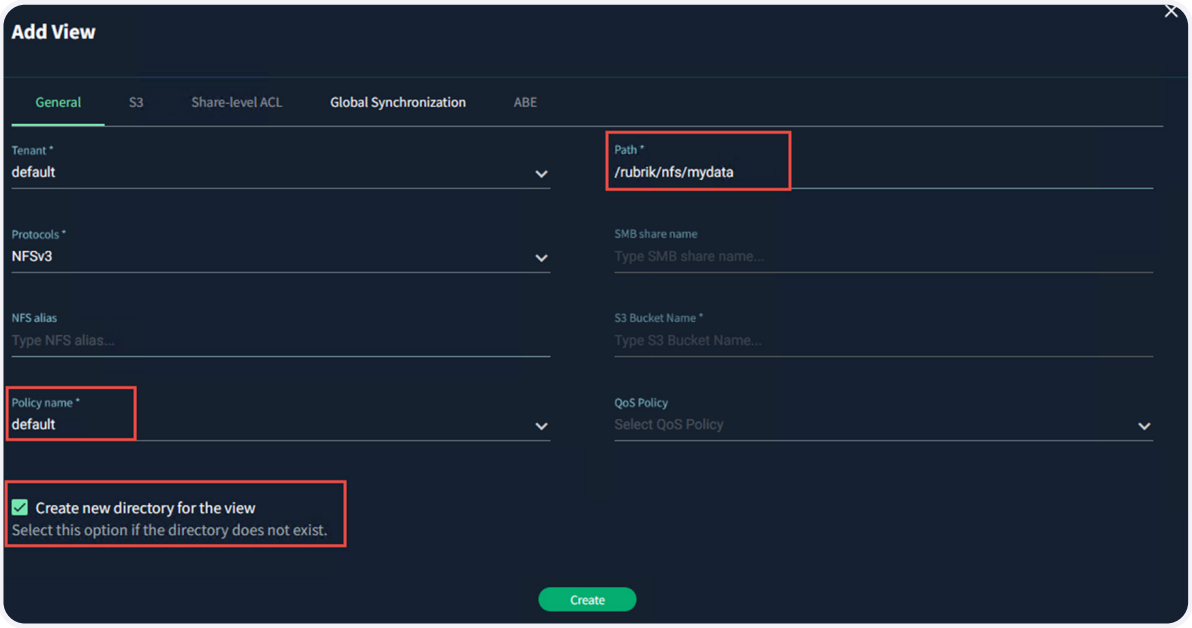
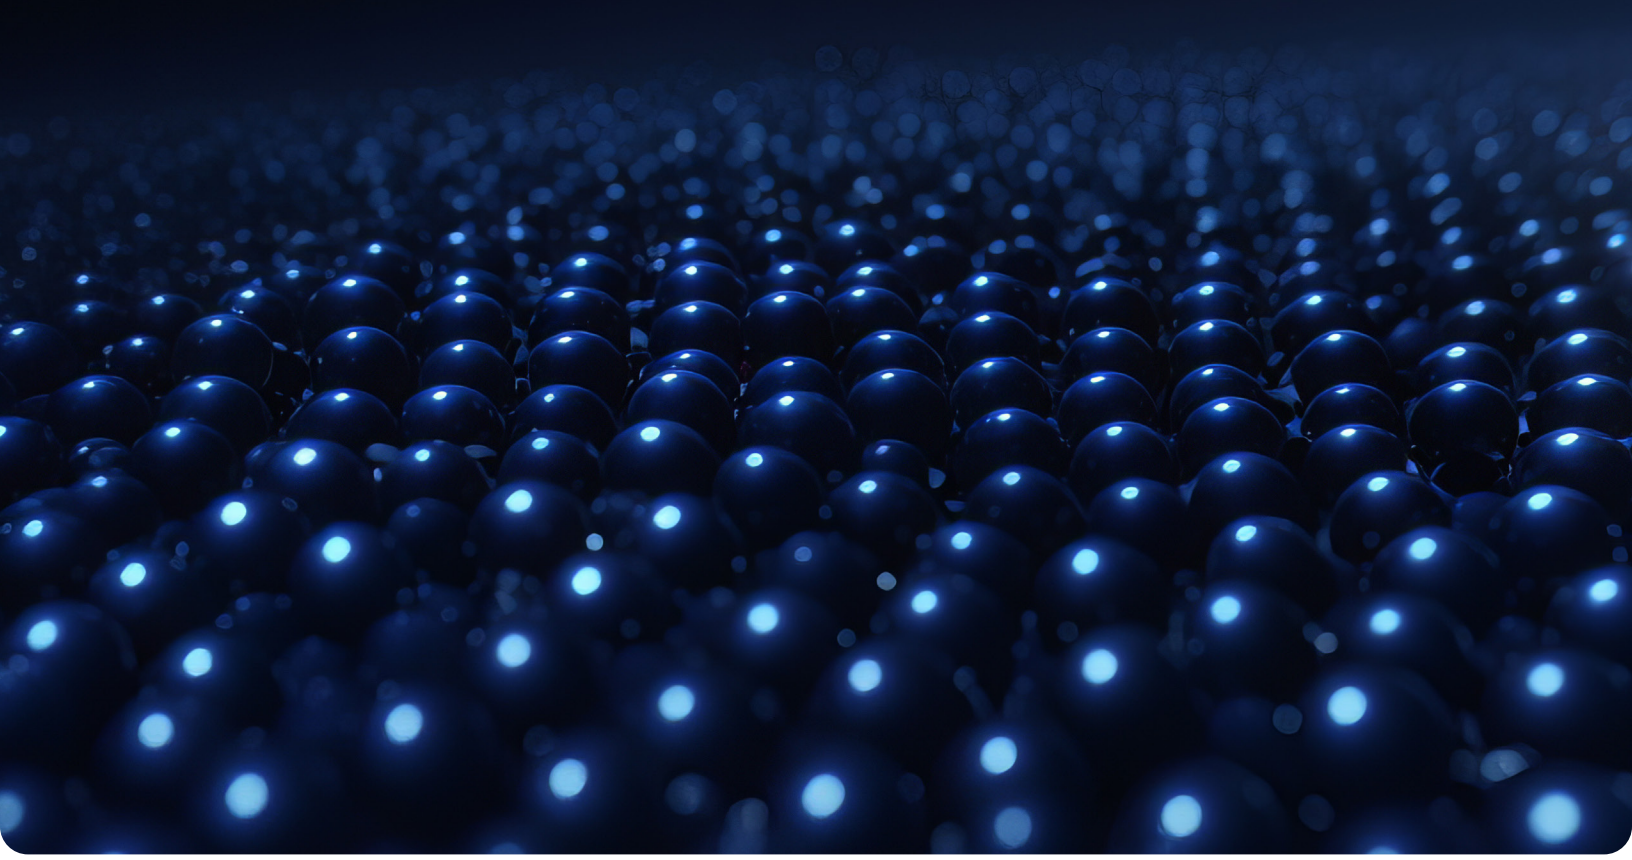


Figure 46 – Creating VAST NFS View

For further information on creating views and policies, please refer VAST Support Documentation.



SSL Certificates

One of the advantages of S3 is that it can use TLS encryption for data inflight. VAST supports S3 over TLS and has encryption at rest. If the source data is also using S3 with TLS (or using another encrypted transport method) and is combined with VAST's S3 over TLS and its encryption at rest, then you have an end-to- end solution without the need to encrypt the data. This allows VAST to give additional deduplication savings on top of Rubrik NCD compression.

To be able to effectively use an encrypted connection to a VAST cloud library over S3 an SSL Certificate will need to be generated and added to any Rubrik NCD node connecting directly to VAST.

The first step is to create an SSL certificate for the VAST cluster. There are several different certificates that are allowed on the system but only the S3 is needed for this solution. (The same certificate could be used for all the VAST categories if desired).

The following sections do not cover the creation of the SSL certificate – Please consult the appropriate team within your organization to acquire an SSL certificate. Ensure that you are provided with two files: a certificate file and a key file and that they are in an X.509 output file format containing ASCII (Base64) encoded data.

To make uploading the document easy, open the 2 files and ensure there is only content from:

[----Begin Certificate-----]

... through ...

[----End Certificate---]

Once an SSL certificate is acquired and the files have been inspected then the following steps can be used to install them into VAST and within Rubrik NCD.

Installing SSL Certificate into VAST

From the VAST UI and from the left menu select **Settings** and **Certificates** as shown in Figure 47.

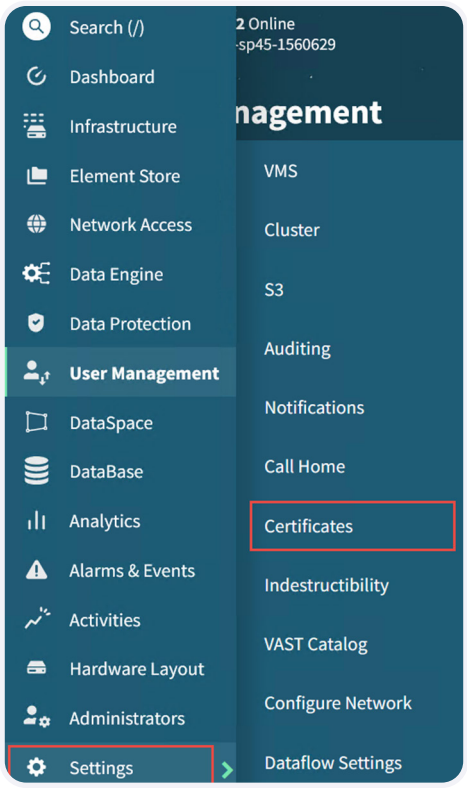


Figure 47 – Select VMS From Menu

Within the Certificates window ensure that S3 has been selected as shown in Figure 48.

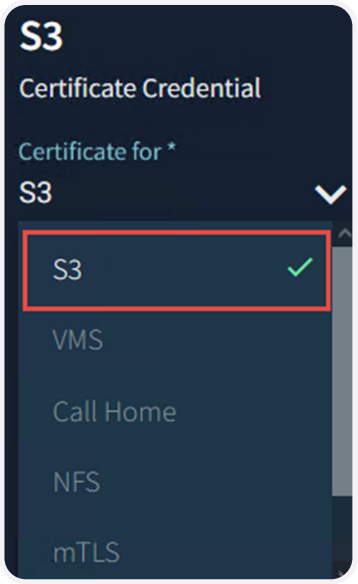


Figure 48 – Selecting S3 Certificate Definitions

Now upload or paste the certificate contents into the left section – appropriately labeled *Certificate* (Figure 49).

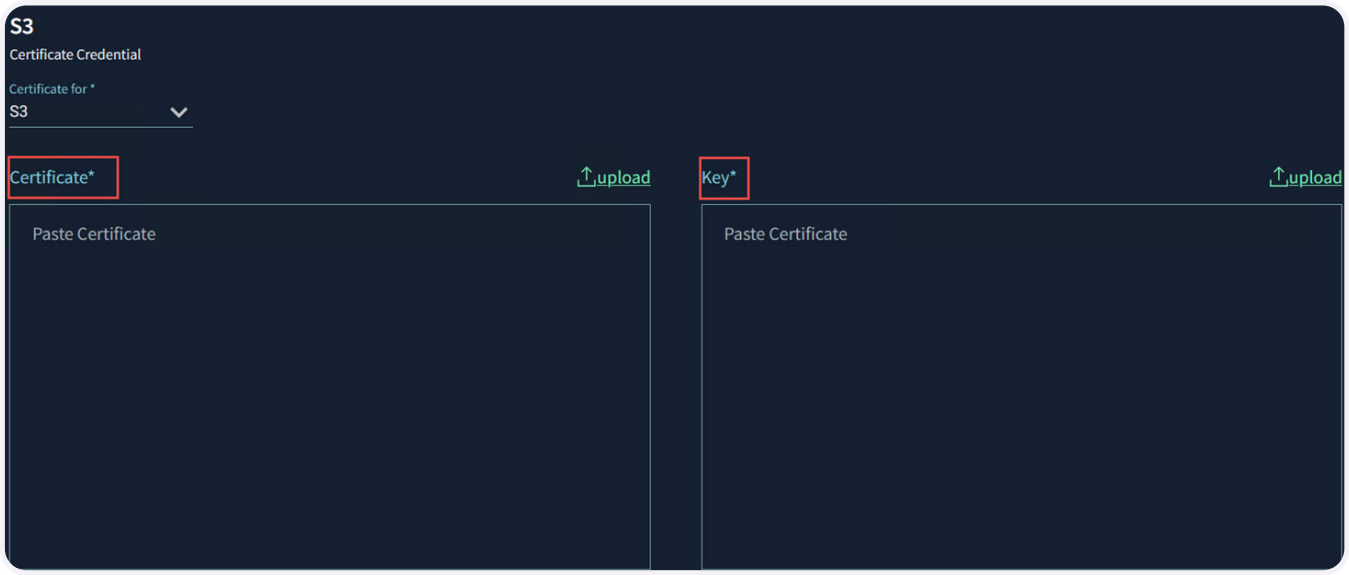


Figure 49 – Paste Certificate and Private Key

Next, upload or paste the .pem file (private key) into the ***Key*** section (Figure 49). Before clicking **Update** to save the settings ensure one last time that this is for the S3 certificate.

Installing SSL Certificate into Rubrik NCD

The CA certificate that is created and added to the VAST cluster for S3 needs to be the same certificate that is uploaded to Rubrik NCD cluster VMs. The **F6** page within the Rubrik NCD console describes how to upload the certificate. For more details, consult the NAS Cloud Direct User Guide.

Once the VAST cluster has a proper certificate installed and the same certificate is installed within each of the Rubrik NCD VMs, SSL communication should function seamlessly.



For more information on the VAST Data Platform and how it can help you solve your application problems, reach out to us at ***hello@vastdata.com***

©2025 VAST Data, Inc. All rights reserved. All trademarks belong to their respective owners.