

VAST Data Platform with Commvault® Backup

Configuration and Best Practice Guide



Table of Contents

Introduction	4
Overview	4
Key Takeaways	5
Best Practices	7
Encryption	7
Encryption – Data In-Flight	8
Encryption – Data at Rest	8
Compression and Deduplication	9
Compression	9
Deduplication	9
Client-side Deduplication and Compression	10
DDB Block Size	10
Solution Design	10
Solution Sizing	11
MediaAgent sizing	11
Space Reclamation	12
Conclusion	12
Appendix – General Configuration	13
VAST Configuration	13
Creating an S3 User	13
S3 Identity Policy	15
Creating an S3 Endpoint on VAST	18
Creating an S3 Bucket	20
Enabling S3 Object Lock	22
SSL Certificates	24
Installing SSL Certificate into VAST	25
Installing SSL Certificate into Commvault	26
Configure Commvault with VAST S3	27
Commvault WORM	30
Enabling WORM mode	32
Configure Commvault with VAST NFS	33
Create Commvault NFS Storage Pool	33

Revision History

Revision Number	Description	Date
1.0	Initial Release	May 2021
2.0	Updated DR and Configuration	October 2022
3.0	S3 Object Lock, WORM, Encryption	May 2024

Introduction

The ever-accelerating pace of IT operations, and the shift to all-flash for primary storage, have left many organizations dissatisfied with the performance of their backup targets, especially on restores. While many users still think of backups and restores as sequential, as they were in the days of tape, modern backup techniques, from incremental forever to data deduplication have made backups and restores much more random than they used to be.

Users too often discover this when they must perform their first large restore, they discover restores are only 1/5th as fast as they back up to it. A VAST cluster provides a single storage pool, namespace, and data reduction realm (See [VAST Data Reduction Whitepaper](#)) capable of spanning 100's of petabytes and are actually faster reading than writing.

VAST Data's multiprotocol S3 and NFS protocol support also allows for pluggable integration as a primary copy of data or a replication target without compromising on storage scale capabilities, redundant appliances, or isolated silos of backup targets.

All-Flash performance also unlocks the secondary copies of data for uses beyond just backup. Commvault's data management suite can index and search secondary copies for eDiscovery, ransomware detection, data exfiltration analysis, etc., but only if the repository is fast enough to support these functions.

Commvault has an architecture that employs a balance of scale out and locality options for being able to backup and recover data quickly and efficiently. With its distributed nature of employing a metadata server (CommServe) and loosely coupled data mover installations referred to as MediaAgents (MA), the Commvault product can control operations against storage devices, scaling up to petabytes organically as business needs grow.

With Commvault's lineage of operating system support for MediaAgents and with new advancements in the S3 protocol and immutable backups (WORM), plus innovation from VAST the solution is well positioned to drive resilience from cyber-attack, coupled with exceptional backup and restore performance with minimal effort, protecting a customer's most important asset, their data.

Overview

With cyber security at the forefront of today's data world, immutable backups with S3 Object Lock and WORM are at the center of minimizing breaches and minimizing down time.

Commvault brings extensive application support and data management functionality, and the VAST Data Platform delivers all-flash performance users require and with unique data reduction technology that brings TCO below that of many spinning disk-based backup solutions. Together it is a best in class cyber resilient backup solution.

This is the third revision of this guide, and it covers some new content around VAST's S3 Object Lock and Commvault's WORM, along with many updates. The first half of the document will discuss technical concepts to ensure that the combined solution can achieve the highest performance and maximum capacity usage needed for a user's specifications. The second half of the document highlights specific configuration steps to ensure the solution is setup properly.

For a quick high-level synopsis there is a key takeaway section next followed by the more In-depth examination of the forementioned topics.

Key Takeaways

This paper will describe how Commvault users should configure their Commvault software to best operate with VAST. It outlines technical details, architectural considerations, and recommended best practices when combining VAST Data's Universal Storage with the Commvault data protection suite. For a quick summary of the important takeaways here is the fast list:

Encryption

- To maximize data reduction rates, allow VAST to handle encryption with encryption over the wire and at rest.

Compression and deduplication

- Enable Commvault client-side deduplication to maximize data reduction which minimizes network traffic and accelerates both backup and restores.
- VAST recommends the Commvault default LZO compression scheme where possible.
- VAST Deduplication and Similarity
 - Provides additional capacity savings and is complimentary to Commvault's compression and deduplication.
 - Significantly reduces storage usage across similar datasets – when a DDB is sealed, or WORM mode is enabled.
- VAST can provide significant data reduction between native formats and the backup copy format stored on the same VAST cluster, enabling additional use cases such as disaster recovery and high availability.
- Deduplication Database – When configuring VAST as a cloud library, Commvault defaults to using a 512KB block size for deduplication. This can be used if VAST is the only copy or if there are secondary copies to public cloud (such as Commvault Air Gap Protect). To maximize Commvault's deduplication efficiency, a 128KB block size is recommended, especially if there are secondary copies on other storage devices.

Immutability

- Leveraging VAST S3 object lock and Commvault WORM mode is supported for backup copies which can be used as part of a cyber recovery solution. Compliance mode immutability is fully supported between Commvault & VAST.
- VAST can deduplicate and similarity reduce across Commvault backup cycles when using WORM mode (DDB seal) for storage immutability. This unique capability removes the WORM lock storage penalty and allows larger data sets to be locked for longer periods of time, saving on requiring 2.5x to 3x extra back end storage.

MediaAgents

- VAST recommends using nConnect with Linux MediaAgents to maximize throughput to/from the VAST cluster if the NFS protocol is leveraged.

Overall Solution

- **Cost-effective:** The combination of data reduction and reduced infrastructure leads to lower solution costs and increased overall TCO.
- **Scalability:** The Storage Accelerator model distributes the data movement and processing across the clients to provide horizontal scaling minimizing the quantity of MediaAgent servers, while the VAST architecture can scale up capacity to the largest enterprise needs.
- **Complementary data reduction technologies:** Commvault's client-side deduplication and compression accelerates data operations and minimizes data movement across the network, while VAST can further reduce the backup footprint and copies of data on the storage synergistically minimizing space consumption.

Best Practices

This section expands on some of the key takeaways by investigating and examining several areas of performance characteristics and how best to configure them. The first area is compression and deduplication which is of the utmost importance for both network bandwidth and storage capacity considerations.

Before discussing those, it is important to understand encryption on the VAST cluster, Commvault encryption and how it can impact deduplication.

Encryption

When considering a solution with Commvault and VAST one needs to look at balancing security, performance, and capacity savings. There are many settings within Commvault and VAST that can help optimize for these.

VAST cluster provides encryption-at-rest and encryption on the wire. VAST encryption-at-rest uses FIPS 140-2 Level 1 ciphers and encrypts both the data and the metadata. Encryption at rest currently needs to be enabled at the time of cluster creation, however, a later release provides the ability to enable this feature on existing clusters that did not initially enable it.

Commvault also has an option to encrypt the data it is protecting. The decision to enable Commvault encryption can have performance and capacity downsides. Enabling the application side encryption on Commvault will result in zero data reduction on the VAST side, which is reflected in a Data Reduction Rate (DRR) of one-to-one (1:1). Also, client-side encryption may affect the MediaAgent's performance as it requires additional CPU to encrypt each block. The performance penalty can be overcome by beefier hardware for the MediaAgent servers, but the capacity savings that would be lost cannot be overcome.

With Commvault's quality data reduction technologies, by way of deduplication and compression, traffic across the wire is minimized and for DASH copy replication which is why VAST encourages our customers to implement these. However, when the encrypted data is sent to VAST, the overall space savings will be much lower than when VAST is chosen to manage the encryption. By choosing to encrypt at the VAST layer, you will benefit from both Commvault's data reduction to create a smaller data set, reduce overall network traffic and benefit from VAST's advanced data reduction techniques without compromising the overall security of your data.

There are scenarios in which Commvault may need to duplicate data even when deduplication is desired. Some of these scenarios include recovering from a missing or corrupted DDB, operations such as sealing the DDB, or adjusting the block size of the Storage Pool or Storage Policy. In these cases, a full copy of data needs to be completely rewritten upon the next backup. The result is that 100% of the base backup set size is sent to VAST, but with Commvault's encryption disabled, the net impact to useable capacity is essentially zero with VAST's advanced data reduction capabilities.

However, if Commvault encrypts the data, VAST cannot reduce it and the result is a doubling of used capacity until previous restore points age out and are deleted. More on DDB best practices are discussed in the [Compression and Deduplication](#) section.

You may still choose to perform the encryption within Commvault and VAST is still an excellent storage platform to store Commvault-encrypted data. Even if VAST's deduplication, compression, and similarity savings cannot further

reduce the data, the overall solution is still very compelling and is denser, faster for backup, recovery & object maintenance tasks and requires less power and cooling than any alternatives available today. Below are several ways to approach encrypting data based on the desired level of security.

Encryption – Data In-Flight

To secure data on the network between Commvault MediaAgents and a VAST cluster, it's recommended to configure the VAST cluster as a cloud library which will use the SSL protocol for communication. SSL will use TLS encryption on all data in-flight to the VAST cluster. This also applies to clients that utilize the Storage Accelerator which will move data directly to and from the VAST cloud library via an SSL (TLS) encrypted tunnel. (**Note:** You may need to add the VAST SSL security certificate to the Commvault CA bundle to enable this configuration – See [SSL Certificates](#) for more information)

For the Commvault clients that will back up through a MediaAgent you can enable [encryption on network traffic in the Commvault software](#) to secure the data in transit between these devices. The MediaAgent will then transmit the data via the encrypted SSL tunnel to the VAST cloud library.

Encryption – Data at Rest

If you need to have data encrypted at rest, both Commvault and VAST offer this capability and you can choose the option which best meets your requirements.

Using VAST encryption: To get the advanced benefits of VAST's data reduction (Adaptive Chunking, Similarity etc.), turn off Commvault encryption on media, and ensure encryption is enabled on VAST (All new systems have encryption at rest turned on by default). Combined with the in-flight TLS encryption from the previous section, this will achieve end-to-end security and still allow VAST to perform its advanced data reduction techniques on the non-encrypted data and write it in a secure format.

If additional copies are needed to different non-encrypting storage, enable encryption in Commvault for the auxiliary copy. Commvault will write the data in an encrypted format to the secondary device so that it's also protected at rest on that copy.

Using Commvault encryption on media: If there are specific security needs, Commvault can encrypt the data and write it to VAST in a secure format. This will require additional compute resources on the Commvault MediaAgents, and VAST will not be able to provide the additional storage reduction since encrypted data will look like random data to VAST.

Allowing VAST to do the encryption at rest removes any performance hit on the Media Agents and allows for storage side data reduction. The SSL connections to VAST leverage TLS encryption to secure the data in flight and the Commvault client to Media Agent using a secure connection completes the end-to-end security.

Allowing VAST to do the encryption removes any performance hit on the MediaAgents and maximizes capacity savings.

Compression and Deduplication

Commvault and VAST data reduction techniques are complementary of each other and using 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. This section examines the various components of Commvault's built-in data compression and deduplication configuration options with VAST's unique data reduction techniques.

Compression

By default, Commvault uses LZO compression and is very effective at getting most out of compressing data. Tests performed previously looked at this and an older compression type and LZO was by far the best at maximizing space savings and minimizing MediaAgent compute usage which equates to faster backups.

Since LZO is the default setting this is an 'easy button' and nothing needs to be configured to achieve the best results.

Deduplication

The VAST Data Platform's unique deduplication techniques, including adaptive chunking and similarity, reduce misaligned data well and can give additional reduction savings on top of Commvault's already deduplicated data. There might be some desire to turn off Commvault deduplication and compression to minimize compute resources but the trade-off with that is increased bandwidth usage on the local network as well as the increased data needing to be sent to VAST. All of this will lead to longer backup times. The perfect balance between compute and network usage and backup intervals is to allow both Commvault and VAST to perform deduplication and compression.

Another advantage to the VAST and Commvault solution is readily apparent when using WORM or when the deduplication database (See - [DDB Block Size](#)) needs to be sealed independently of WORM. When a DDB is sealed, a new baseline full backup is run with a brand new DDB, so all the data is now considered new and therefore written again to storage. This typically costs 2.5x or more of total storage capacity. However, due to VAST's superior deduplication it has essentially no capacity cost and in tests shows only 1-3% additional storage consumed.

VAST's byte level adaptive chunking and similarity on data written to VAST maximizes storage efficiency in addition to Commvault deduplication alone.

Additionally, if WORM mode is enabled, or for some reason the DDB needs to be sealed and a new set of baseline full backups are run, the VAST Data Platform can deduplicate the new set of data against the older set leading to reduced storage consumption.

Client-side Deduplication and Compression

Enabling deduplication and compression on the client-side minimizes the data sent over the local network and to the VAST cluster like MediaAgent-based processing but with the potential of reducing backup and restore times as the process is now distributed across each individual client.

The same deduplication and compression technology that is on MediaAgents now runs on the client and the VAST cluster will still further deduplicate, compress and reduce with similarity the data written to storage maximizing storage efficiency.

DDB Block Size

Another aspect of Commvault's deduplication is the Deduplication Database (DDB) that it uses to do hash lookups of file and object data to determine if they have already been written to storage. If the hash of an object or file is in the DDB then the data is not sent again to VAST. This not only minimizes the data sent from the MediaAgents to VAST but also frees up local network bandwidth from clients to the MediaAgents.

When configuring VAST as a cloud library, Commvault defaults to using a 512KB block size for deduplication. This can be used if VAST is the only copy or if there are secondary copies to public cloud (such as Commvault Air Gap Protect).

To maximize Commvault's deduplication efficiency, a 128KB block size is recommended for high-speed local disk storage. 128KB will give better deduplication ratios for Commvault (less data moved through the network) and more granular data aging to minimize storage bloat. 128KB should be used if there are copies on other storage devices so that the storage efficiency is maximized across all copies.

Solution Design

With VAST's unique all-flash DASE (Disaggregated Shared Everything) platform, architecture of the backup environment plays a crucial role and must be considered in the design. Identifying the lowest common denominator in the data flow is essential as it determines the limiting factor.

When considering restore scenarios, the VAST platform will usually not be the bottleneck as its current restore speeds are almost 8x faster than writes, so the target system's write speeds and network links should be considered first.

Considerations:

- VAST currently has 8x read speeds at almost 40GB/s, compared to write speeds (write speeds will increase 2x or more in subsequent releases).
- The network architecture may become a bottleneck if it's not designed for the high performance restore speeds of the VAST platform. VAST provides multiple 100 Gbit/s NIC's for connectivity to the customer network.
- The disks on the clients will always be a major consideration for restore performance when considering Recovery Time Objectives (RTO).

Solution Sizing

MediaAgent sizing

The network architecture is important to consider in the overall solution design to make sure that data can be moved fast enough to VAST. MediaAgents can be sized using the [Commvault specifications for deduplication mode](#) with the back-end sized for cloud storage. Since VAST is a high-speed solution, the MediaAgent should include fast interconnects so that it can get the full benefits of the cluster as this connection will typically be the bottleneck of the solution.

Sizing for WORM storage

On typical object storage that lacks duplication but supports the S3 Object Lock API on a per-object level (or equivalent), enabling WORM will result in approximately a 2.2x to 2.5x increase in storage utilization. Two baselines of data must be retained on the device, along with some overhead and potential cycle overlap. Based on testing results, VAST can achieve 96% or more of deduplication across two baselines, reducing the cost of using WORM with VAST to just a few percentage points, lowering storage utilization to around 1x.

The result is that WORM mode becomes cost-effective. This not only reduces the cost of utilizing this feature but also enables customers to retain larger datasets for longer periods, which may not be economically feasible with other storage systems.

Commvault Storage Accelerator

The Commvault Storage Accelerator can be used on clients to directly send the backup data to the VAST cloud storage library (sometimes referred to as Client-Direct mode). Without the Storage Accelerator, backups send the client data to the MediaAgent first, and then the MediaAgent stores the data in the VAST cloud storage library. When used on many clients, this can achieve high aggregate throughput for both backup and restores while lowering MediaAgent requirements. Storage Accelerator does consume more resources on the client, so resource allocation should be considered on systems that are heavily loaded. The Storage Accelerator configuration has the following advantages:

- Speeds up backup and restore process by directly backing up and restoring from the cloud storage library and avoiding additional network transmission.
- Reduces the load on the cloud MediaAgent by eliminating the need for processing backup data. (Pruning, Data Verification (DDB Lookups), and defrag operations would continue to be processed by the destination MediaAgent.)

Any client that has the Storage Accelerator installed and has a backup policy configured to a cloud storage library will automatically try and use the direct mode, and if the connection to the storage fails, it will fall back to sending data through the MediaAgent instead. The same SSL certificate requirements and additional settings that are required by the MediaAgent would also be needed on each client (See [Appendix – General Configuration](#) for more information)

The Storage Accelerator model distributes the data movement and processing across the clients without sending data through the MediaAgent to provide horizontal scaling.

Space Reclamation

With VAST configured as cloud storage (in your DC or in a private cloud DC), Commvault space reclamation can be run regularly to compact the backup files and increase storage efficiency. And due to VAST's performance and IO capabilities, these tasks run swiftly without impacting regular backup schedules.

Conclusion

The combination of Commvault 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 Commvault are complementary. Commvault will compress and deduplicate data minimizing local network traffic and speed up backups and then VAST will further reduce Commvault's post deduplicated and post-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, high availability, and test/dev operations.

Extraordinary capacity savings is also seen between backup cycles when sealing a DDB in conjunction with using Commvault's WORM functionality where VAST's duplication nullifies a typical 2.2x - 2.5x capacity requirement down to 4% or less and allows for larger data sets to be securely stored for extended periods of time.

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 Commvault'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.

Appendix – General Configuration

This section covers the integration between VAST and Commvault. The goal is to cover as many of the critical configuration steps as possible to ensure that the implementation can be up and running quickly.

VAST Configuration

During the process of adding an S3 backup repository the credentials of the S3 will be needed so the creation of an S3 user is first discussed in the next section.

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 1](#)).




Figure 1 - Selecting Users Category

The Add User window appears ([Figure 2](#)) and a username needs to be entered along with any desired UID. This user was not granted full S3 credentials as it is not a requirement during the Commvault process of adding an S3 repository.

Figure 2 - Creating a New User for S3

After clicking **Create** the user will show up in the User Management window (Figure 3).

Figure 3 - User Created

Now that the user is created it needs to be edited to create and capture the active and secret keys. On the far right of the user under the action column click on the three dots and select Edit (Figure 3).

In the Update User window (Figure 4) click on the Create new key button to create a new Access Key. Now copy the active key and secret key some place for a bit 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.

Update User
User: cvlt-TMPHX10

Name *
cvlt-TMPHX10

UID *
80823

Leading group
Select Group

Groups
Select Groups

Access key

Access key	Status	Actions
4VBUNQ6SUWC89HDE7ZT	enabled	✕
AY3VQV8ST25X66IT3TBA	enabled	

Create new key

The secret key is: **Xi0rHQluVSzyxr3+gXFsl/iOwcrOdUtUDlrcbcNE**
This is your opportunity to make a copy, it will forever disappear after closing this window

Copy key

Update

Figure 4 – Capturing Active and Secret Keys

The user can now be used in the creation of a bucket or view within the VAST UI. For S3 implementations with Commvault and Identity policy is also needed and discussed next.

S3 Identity Policy

To ensure that VAST and Commvault S3 interoperates correctly it's necessary to apply all appropriate S3 permissions to the bucket owner.

Commvault's required S3 permissions are:

```
"s3:CreateBucket",
"s3:GetBucketLocation",
"s3:GetObject",
"s3:PutObject",
"s3:PutObjectRetention",
"s3:PutObjectTagging",
"s3:ListBucket",
"s3:ListAllMyBuckets",
"s3:DeleteObject",
"s3:DeleteObjectVersion",
"s3:ListBucketVersions",
"s3:PutBucketObjectLockConfiguration"
```

To apply these to the bucket owner first create an identity policy by going to Identity Policy under User Management (Figure 5).




Figure 5 - User Management IdentityPolicies

Then click **Create Policy** and ensure the following JSON code is inserted:

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

Note: The above code has some additional S3 permissions beyond Commvault's requirements to ensure complete coverage.

Once done the Policy should look like the one in [Figure 6](#).

Update Policy

Name *
CVLT-ObjectLock

Policy

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "s3:CreateBucket",
        "s3:GetBucketLocation",
        "s3:GetBucketVersioning",
        "s3:GetObject",
        "s3:PutObject",
        "s3:PutObjectRetention",
        "s3:GetObjectRetention"
      ]
    }
  ]
}
```

Upload JSON file

Users list (3)
CommVault Partner Account
cvlt-TMPHX10

Groups list (0)

Update

Figure 6 – Identity Policy

Note that the User List shows the users that have this policy applied. To add a user to the policy, that user needs to be modified. To do that, click on the **Users** tab which is right next to **Identity Policies** ([Figure 5](#)). Find the user that will be assigned owner of the Endpoint (Each new bucket in the Endpoint will inherit the Endpoint properties – See [Creating an S3 Endpoint on VAST](#)) and edit it by clicking on the dots to the right. In [Figure 7](#) user **cvlt-TMPHX10** is being edited. In the pull-down menu for **Identity Policies** the **CVLT-ObjectLock** policy has been selected which contains all the proper S3 permissions. Simply click on **Update** to complete the edit.

Update User
User: cvlt-TMPHX10

Name *
cvlt-TMPHX10

UID *
80823

Leading group
Select Group

Groups
Select Groups

Identity Policies
CVLTObjectLock

Allow Create Bucket ☒ Allow Delete Bucket ☒ S3 Superuser ☐

Access key	Status	Actions
4VBUNQ6SUWC89HDE7ZT	enabled	× II
AY3VQV8ST25X66IT3TBA	enabled	× II

Create new key

Update

Figure 7 - Update User with Identity Policy

After completing these tasks, an appropriate Cloud Storage pool and policy can be created within Commvault (See - [Configure Commvault with VAST S3](#)).

Creating an S3 Endpoint on VAST

A good best practice for organizational purposes and ease of use with creating buckets within the Commvault workflow is to create an endpoint at the root of where buckets would be placed. Once an endpoint is established the buckets that are created using the same user and secret key will be placed in an orderly manner underneath that endpoint. Buckets can be created within the VAST UI, the Commvault workflow, or through an S3 utility like S3 Browser®.

To create an endpoint, go to the Element Store within the VAST UI, select the Views tab and then click Create View as highlighted with the upper right box in [Figure 8](#).

Element Store

Views View Policies Lifecycle Rules Quotas Locks Tenants QoS Policies

Path	Alias	Share	Bucket	Bucket Owner	Policy Name	Protocols	Physical Capacity	Logical Capacity	Live Monitor	Actions
/commvault	Search...	Search...	Search...	Search...	spl	S3 ENDP ...	1,656 TiB	5,395 TiB	No	Select Live Mon

+ Create View

Figure 8 - Creating an S3 Endpoint on VAST

Add View

General S3 Share-level ACL

Tenant *
default

Path *
/commvault

Protocols *
S3 Endpoint

NFS alias
Type NFS alias...

Policy name *
spock-s3

Create Directory ☒

SMB share name
Type SMB share name...

S3 Bucket Name *
Type S3 Bucket Name...

QoS Policy
Select QoS Policy

Create

Figure 9 - S3 Endpoint General Details

This will open an Add View wizard as shown in [Figure 9](#). Here the process is similar to bucket creation where a path is entered, a policy is given but only the Endpoint protocol is selected. (Don't forget to slide the Create Directory toggle over). On the S3 tab ([Figure 10](#)) of this wizard ensure the appropriate user is given access to create buckets. Click Create when finished. This will show the endpoint as in [Figure 8](#).

Add View

General S3 Share-level ACL

S3 Owner *
Type...

S3 Object Ownership Rule
None

S3 Features

S3 Versioning ☐

S3 Object Lock ☐

Retention Mode
None

S3 Retention Period
1d

S3 Access Control

Bucket Creators (Users)
cvlt-TMPHX10 Type...

Bucket Creators (Groups)
Type...

Anonymous access ☐

Create

Figure 10 - S3 Endpoint S3 Details

Again, a handy feature of having an endpoint is when using any S3 utility along with the same user it will propagate all parameters according to the endpoint and place buckets cleanly under the endpoint path.

For example, the endpoint that was just created here has a path of **/commvault**. A new bucket created with the name – **bucket1** will show up like this in the VAST UI: **/commvault/bucket1**. This becomes especially useful when creating the bucket within the Commvault Cloud Library workflow or other S3 utilities that allow bucket creation.

Creating an S3 Bucket

During the Commvault library creation workflow, a VAST bucket can be configured without having to use the VAST UI. If a user wants to create one ahead of time this section will describe those steps. To create a bucket, go to the dashboard and go to Element Store as shown in [Figure 11](#) and then click Views.




Figure 11 – Selecting Views to Create a Bucket

This brings up the Element Store window with the Views tab selected. In the upper right corner select Create View ([Figure 12](#)).





Figure 12 – Select Create View

This view is being configured as a bucket only using just the S3 Bucket protocol. Multiple protocols (NFS, S3) can be used on a view. [Figure 13](#) shows all the settings needed to create the new bucket. Since this is a new view the create directory toggle is selected.



Add View

General S3

Path *
/commvault/S3Bucket

Protocols *
S3 Bucket

SMB share name
Type SMB share name...

NFS alias
Type NFS alias...

S3 Bucket Name *
s3bucket


Policy name *
s3_default_policy

Create Directory ☒

Create

Figure 13 – S3 Bucket Configuration

The user created previously needs to be given access or ownership of this view. This is done on the S3 tab as shown in [Figure 14](#). When finished click Create from either tab.



Add View

General S3 Share-level ACL

S3 Owner *
cvlt-TMPHX10 Type...

S3 Object Ownership Rule
None

S3 Features

S3 Versioning ☒


S3 Object Lock ☒

Retention Mode
None

Create

Figure 14 – Adding Bucket Owner (User)

The view (bucket) now shows up in the Element Store as shown in [Figure 15](#).



Path	Alias	Share	Bucket	Bucket Owner	Policy Name	Protocols	Physical Capacity	Logical Capacity	Live Monitor	Actions
bucket-new	Search...	Search...	Search...	Search...	Search...	Select Protoc...	0 Bytes	0 Bytes	Select Live Mon	
/commvault/bucket-new			bucket-new	cvlt-TMPHX10	spock-s3	S3 BUCKET	0 Bytes	0 Bytes	No	

Figure 15 – S3 Bucket Created


With the bucket now created a cloud storage pool using this bucket can be created within the Commvault UI.

Enabling S3 Object Lock

When the intention is to use Commvault WORM, object lock will need to be enabled on the bucket within the VAST UI. During the creation of a Commvault cloud storage library (See – [Configure Commvault with VAST S3](#)) a VAST bucket is either created within that workflow or manually as shown in the previous section. Once the bucket has been chosen or created then that is the time to enable object lock.

To do that start by selecting the Element Store and find the appropriate bucket and click on the three dots on the right to edit ([Figure 15](#)).

Within the bucket on the General tab, ensure that this is a bucket and not an endpoint ([Figure 16](#)).



Update View

General

S3

Share-level ACL

Tenant *

default

Path *

/commvault/bucket-new

Protocols *

S3 Bucket

SMB share name

Type SMB share name...

NFS alias

Type NFS alias...

S3 Bucket Name *

bucket-new

Policy name *

spock-s3

QoS Policy

Select QoS Policy

Create Directory

☐

Update

Figure 16 – Editing Bucket General Tab

Click on the neighboring S3 tab and toggle S3 object lock on (this will also turn on S3 Versioning). The key part here is the Retention Mode setting.

Update View

General

S3

Share-level ACL

S3 Owner *

cvlt-TMPHX10 Type...

S3 Object Ownership Rule

None

S3 Features

S3 Versioning

S3 Object Lock

Retention Mode

None

S3 Retention Period

0d

S3 Access Control

Bucket Creators (Users)

Type...

Bucket Creators (Groups)

Type...

Anonymous access

Update

Figure 17 - Editing Bucket S3 Tab

Since Commvault will be managing all aspects of the locked objects and their corresponding retention, ensure that the Retention Mode (Figure 17) is set is set to **None**. Enabling object lock at the time of cloud creation ensures that it is not overlooked as backups will fail without it.

SSL Certificates

One of the advantages of S3 is that it can use TLS encryption for the data in flight from either the MediaAgent or a client with Storage Accelerator. Combine that with VAST's encryption at rest and Commvault's network encryption and you have an end-to-end solution without encrypting the data (See - [Encryption](#)). This allows for VAST to give additional deduplication and compression savings on top of Commvault.

To be able to properly use an encrypted connection to a VAST cloud library over S3 an SSL Certificate will need to be generated and added to any client or MediaAgent 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 ([Figure 19](#)) but this one will be needed for S3. The same certificate could be used for all the 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 Commvault.

Installing SSL Certificate into VAST

From the VAST UI and from the left menu select Settings and Certificates as shown in [Figure 18](#).




Figure 18 - Select VMS From Menu

Within the Certificates window ensure that S3 has been selected as shown in [Figure 19](#).




Figure 19 - Selecting S3 Certificate Definitions

Now upload or paste the certificate contents into the left section – appropriately labeled *Certificate* (Figure 20).

Figure 20 – Paste Certificate and Private Key

Next, upload or paste the .pem file (private key) into the **Key** section within VMS (Figure 20). To save the settings click **Update**.

Installing SSL Certificate into Commvault

The easiest approach is if the Commvault server (Windows) is a member of the domain from which the VAST S3 certificate was issued because it will already have the root CA included as a Trusted Root Certification Authority for User, Local Computer and Services certificate stores. So, if the Commvault server (Computer) is already part of the same domain as the VAST cluster, then by default it's trusted, and configuring secure cloud storage will work without issues.

This functionality is not currently available for Linux with Commvault releases through v11.34. Starting in Commvault v11.36 it will introduce domain level certificates. Previous versions will need to have each MediaAgent and client manually modified.

If not using domain method, then the **curl-ca-bundle.crt** file must be modified manually for both windows or Linux MediaAgents and clients for SSL communication to function properly. This file is in the Commvault **Base** folder. (Before you modify the file it's a good idea to back up the file and verify the host isn't currently running any backups.)

In Windows the file is located under

***\Commvault\ContentStore\Base**

and in Linux under

***/Commvault/Base or**

***/Commvault/Base64**

Simply copy the content of the certificate contents (not the private key) into the file and save it. If there are still issues this knowledge base article can help resolve them – [Fixing SSL Certificate Problems](#)

Once the VAST cluster has a proper certificate installed and the client/MediaAgents have had their certificate filed updated SSL communication should function seamlessly.

Configure Commvault with VAST S3

Configuring a cloud-based storage pool with VAST is an easy task. The methodology differs slightly beginning in Commvault version 11.36. In that release, VAST has been integrated into the workflow and is available as a pull-down menu option. To begin, from Commvault Command Center in the left column click **Storage** and then select **Cloud** as shown in [Figure 21](#).




Figure 21 – Creating a Cloud Storage Pool

Now on the Cloud page click the **Add** button in the upper right ([Figure 22](#)).



Figure 22 – Add Cloud Storage

This will bring up the Add Storage Cloud window (Figure 23). The first entry to select is the object storage type – use the pull-down menu to select VAST Data Platform (v11.36+) or **S3 Compatible Storage** (Figure 24) if the choice for VAST is not listed. Select a MediaAgent that will be used as the data mover. In the **Service Host** field ensure the entry includes **http://** or **https://** for it to effectively communicate with the VAST cluster.

The screenshot shows the 'Add cloud storage' window in a dark-themed interface. The 'Type' dropdown menu is open, showing 'VAST Data Platform' as the selected option. Below the dropdown, there is a search bar with 'Q_vast' entered. The 'Service host' field is empty. The 'Credentials' field contains 'tmphx10-cvlt'. The 'Bucket' field is empty. There is a toggle for 'Use deduplication' which is turned on. Below it, the 'Deduplication DB location' section shows 'No deduplication db location added' and an 'Add' button. At the bottom, there is an 'EQUIVALENT API' link and 'CANCEL' and 'SAVE' buttons.


Figure 23 – Add VAST Data Platform (FR 11.36+)

The **Credentials** setting is the user that was set up and assigned as the bucket owner on the VAST cluster. Enter in the name of the bucket to be used – this will either use a preconfigured bucket or create it on the fly. This cloud storage example is also using a DDB which can be configured by clicking **Add**. This brings up the configuration window as shown in Figure 24. Enter a MediaAgent and location for the DDB database and click **Add**.

The screenshot shows the 'Add cloud storage' window in a dark-themed interface. The 'Type' dropdown menu is open, showing 'S3 Compatible Storage' as the selected option. Below the dropdown, there is a search bar with 'Q_s3' entered. The list of options includes 'Amazon S3', 'Ceph Object Gateway (S3-compatible)', 'Dell EMC ECS (S3-compatible)', 'IBM Cloud Object Storage (S3-compatible)', and 'S3 Compatible Storage', which is highlighted with a red box. The 'Service host' field is empty. The 'Credentials' field is empty. The 'Bucket' field is empty. There is a toggle for 'Use deduplication' which is turned on. Below it, the 'Deduplication DB location' section shows 'No deduplication db location added' and an 'Add' button. At the bottom, there is an 'EQUIVALENT API' link and 'CANCEL' and 'SAVE' buttons.


Figure 24 – Add VAST as S3 Compatible Storage (Pre FR11.36)

The **Credentials** setting is the user that was set up and assigned as the bucket owner on the VAST cluster. Enter in the name of the bucket to be used – this will either use a preconfigured bucket or create it on the fly. This cloud storage example is also using a DDB which can be configured by clicking **Add**. This brings up the configuration window as shown in [Figure 25](#). Enter a MediaAgent and location for the DDB database and click **Add**.



The image shows a configuration window titled "Deduplication DB location". It has a dark blue background. At the top right is a close button (X). Below the title, there are two dropdown menus. The first is labeled "MediaAgent *" and has "cvm" selected. To its right is a plus sign (+). The second dropdown is labeled "Select a Deduplication DB location *" and has "E:\ddb\VAST-Cloud-10-ddb" selected. To its right is also a plus sign (+). Below these dropdowns, a line of text says "Free space available on selected Deduplication DB Disk: 221.28 GB". At the bottom right, there are two buttons: "CANCEL" and "ADD".

Once everything has been properly entered for the Cloud Storage it should be similar to [Figure 26](#). Simply click **Save**.



The image shows a configuration window for the VAST Data Platform. It has a dark blue background. At the top, there is a "Type" dropdown menu with "VAST Data Platform" selected. Below that is a "Name *" field with "VAST-Cloud-10" entered. Then is a "MediaAgent *" dropdown menu with "cvm" selected and a plus sign (+) to its right. Below that is a "Service host *" field with "https://8vips.tmpfx-10.vastdata.lab" entered. Then is a "Credentials *" dropdown menu with "tmpfx10-cvlt" selected, a plus sign (+) to its right, and an edit icon (pencil). Below that is a "Bucket *" field with "bucket-test" entered. Then is a "Use deduplication" toggle switch, which is currently turned on. Below that is a "Deduplication DB location" section. It has a table with two columns: "MediaAgent" and "Deduplication DB location". The table has one row with "cvm" in the first column and "E:\ddb\VAST-Cloud-10-ddb" in the second column. To the right of the table is a close button (X). To the right of the "Deduplication DB location" section is an "Add" button. At the bottom left, there is a link "EQUIVALENT API". At the bottom right, there are two buttons: "CANCEL" and "SAVE".

Figure 26 - VAST

With the cloud storage configured it should now show on Cloud window as shown in [Figure 27](#).




Figure 27 – Cloud Storage Configured

Commvault WORM

Within Commvault you can enable WORM storage mode on the VAST cloud library to provide immutable storage for backups. This feature will leverage S3 object lock to ensure that all backup data is immutable for the duration of the Commvault retention period.

In some combinations of Commvault and VAST an additional setting is required to allow for proper WORM functionality.

Commvault Version	VAST OS	Additional Setting Required
v11.28	Any	No
v11.30 – v11.36	Before 5.0 sp11	Yes
v11.30 – v11.36	5.0 sp11+	No

The additional setting needed is – bDisableNativeFileForMacroPruning and has the following parameters:

Description

In V11.30+ when WORM mode is enabled, we write in “native cloud file” format which defaults to using multi-part uploads.

This may create problems with WORM storage where multi part upload is not supported. To revert to the older methods that do not use multi-part upload set this key to true: bDisableNativeFileForMacroPruning

Name: **bDisableNativeFileForMacroPruning**

Entity: **MediaAgent**


Created On: **MediaAgent**

Type: **Boolean (true or false)**

Value: **true**

This key must be set for all MediaAgents and Storage Accelerator clients that will access the VAST library directly.

This is done under the Additional Settings tab under the client or MediaAgent properties setting.



System / Additional settings

Add entity settings

Name
bDisableNativeFileForMacroPruning

Entity
Media Agents ▼

Value*
True

Comment *
Disable Native File Mode so that WORM mode is supported on VAST Storage

[EQUIVALENT API](#)

Figure 28 – Additional Setting Necessary for proper WORM Functionality

Enabling WORM mode

Under the configuration tab in the cloud library properties, you can click the slider named **WORM Storage Lock** to enable immutability on the storage ([Figure 29](#)).

WORM Storage Lock will leverage the S3 object lock API to set retention on the files on the storage. Compliance lock will be also enabled automatically and will prevent data and job deletion within the Commvault software.




Figure 29 - Enabling Commvault WORM

Configure Commvault with VAST NFS

Create Commvault NFS Storage Pool

Creating an NFS Storage Pool for Commvault is just as simple as creating a cloud storage pool. From Commvault Command Center in the left column click **Storage** and then select **Disk** which is just above **Cloud** as shown in [Figure 21](#).

Now on the **Disk** page click the **Add** button in the upper right similar to what is shown in [Figure 22](#) for **Cloud**. This will bring up the **Add Disk Storage** window ([Figure 30](#)).

The screenshot shows the 'Add disk storage' window in the Commvault Command Center. The window has a title bar 'Disk' and a subtitle 'Add disk storage'. The main content area contains the following fields and controls:

- Name ***: A text input field containing 'VAST-NFS'.
- Backup location**: A section with an 'Add' button in the top right. It contains a table with two columns: 'MediaAgent ↑' and 'Backup location'. The first row shows 'cv-media63' and '/mnt/vast-nfs' with a trash icon in the right margin.
- Use deduplication**: A toggle switch that is currently turned on.
- Deduplication DB location**: A section with an 'Add' button in the top right. It contains a table with two columns: 'MediaAgent ↑' and 'Deduplication DB location'. The first row shows 'cvmb' and 'E:\ddb\VAST-NFS-ddb' with a trash icon in the right margin.
- EQUIVALENT API**: A text label at the bottom left.
- CANCEL** and **SAVE** buttons: Located at the bottom right of the window.

Figure 30 - Adding NFS Disk Storage

Now enter a name for the disk storage and pick a MediaAgent that has access to the NFS share. In this example the share is mounted to the MediaAgent and is selected from a browse window (Figure 31).




Figure 31 – Selecting Backup Location

Entering the DDB information is the same configuration method as is done for cloud storage. When finished it should be like [Figure 30](#). Simply click **Save** to finish and the new NFS Disk storage will as in [Figure 32](#).




Figure 32 - NFS Disk Storage Created

The disk can now be used in a Plan or Storage Policy.



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.