VAST + COMMVAULT

SOLUTION GUIDE

Best Practices for Modernizing Data Protection With All-Flash Universal Storage

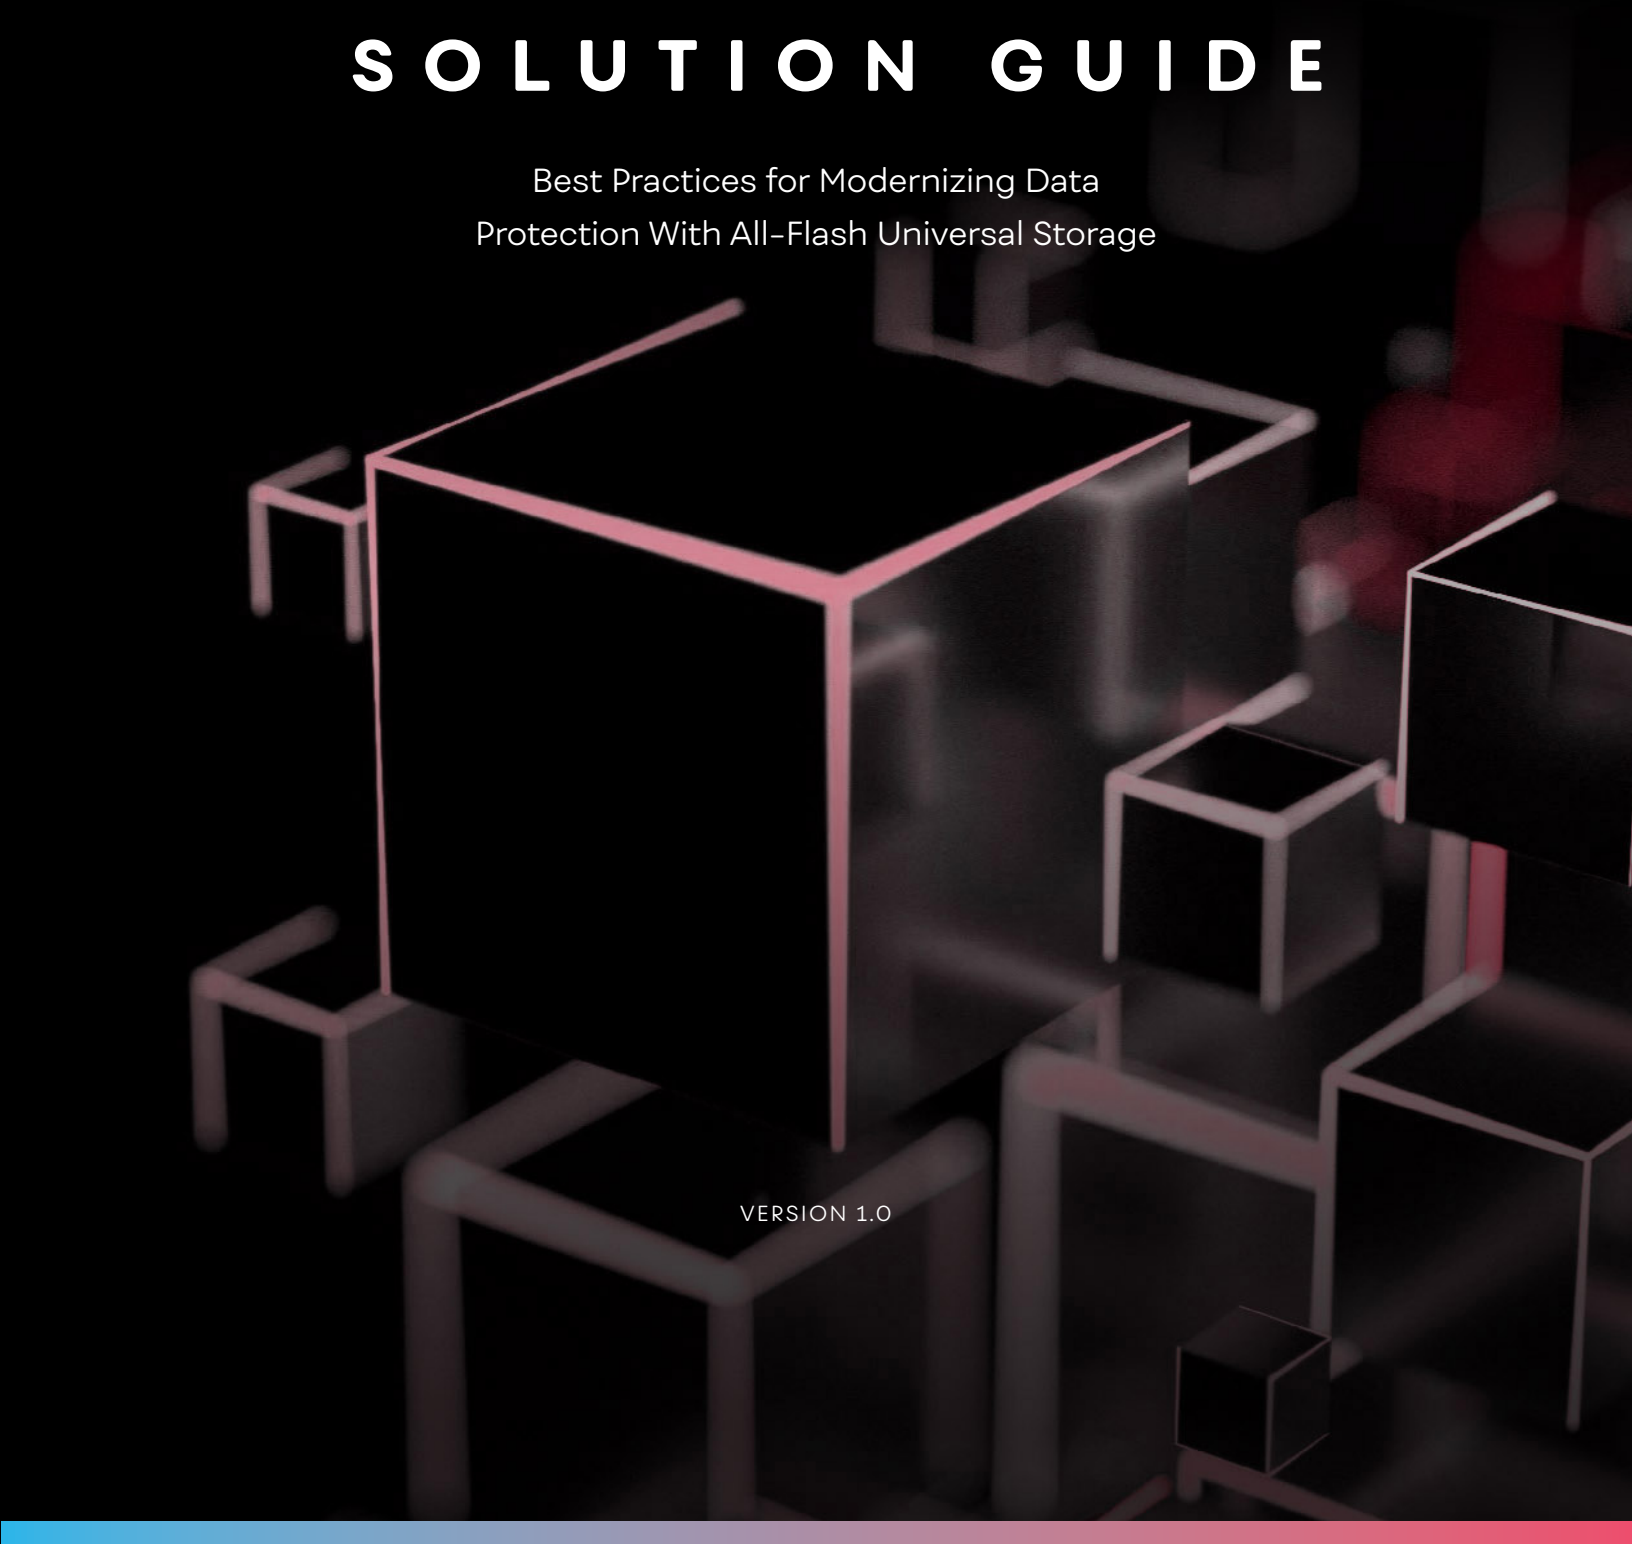

VERSION 1.0

EXECUTIVE SUMMARY

This solution guide outlines technical details, architectural considerations, and recommended best practices when combining VAST Data's Universal Storage with the Commvault data protection suite.

The combination of Commvault and VAST Data provides a powerful solution for users needing to manage secondary copies of their data. While Commvault brings extensive application support and data management functionality, VAST Data's Universal Storage system, based on the revolutionary DASE (Disaggregated Shared Everything) architecture, delivers the all-flash performance users require at a cost below that of many disk-based backup solutions.

In addition to unrivaled performance for fast restores, the testing of this combination revealed that the data reduction capabilities of VAST Data and Commvault are complementary, with VAST systems further reducing Commvault's post-deduplicated and post-compressed data by up to an additional 2X.

TABLE OF CONTENTS

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 disk-based 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 have to perform their first large restore, and discover that the purpose-built backup appliance (PBBA) they spent a million dollars on restores only 1/5th as fast as they back up to it.

VAST Data's Universal Storage provides an all-flash alternative to disk based PBBAs (purpose-built backup appliances) that's not only faster than the disk based alternative but also much more scalable. A VAST cluster provides a single storage pool, namespace and data reduction realm over 100's of petabytes. Where disk based PBBAs are slower restoring than backing up, VAST systems are actually faster reading than writing.

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.

This paper will introduce VAST Data's Universal Storage system, and it's DASE architecture, and describe how Commvault users should configure their Commvault software to best operate with VAST.

Part of this exercise requires a deeper look at the interoperation of Commvault's built-in data compression and deduplication configuration options with VAST's unique data reduction techniques to minimize the storage required while optimizing for network efficiency.

DASE ARCHITECTURE OVERVIEW

VAST Data's Universal Storage is based on a new scale-out architecture concept consisting of two building blocks that are scaled across a common NVMe Fabric.

First, the state (and storage capacity) of the system is built from resilient, high-density NVMe-oF storage enclosures.

Second, the logic of the system is implemented by stateless docker containers that each has the ability to connect to and manage all of the media in the enclosures.

Since the compute elements are disaggregated from the media across a data center scale Fabric, each can scale independently – thereby decoupling capacity and performance.

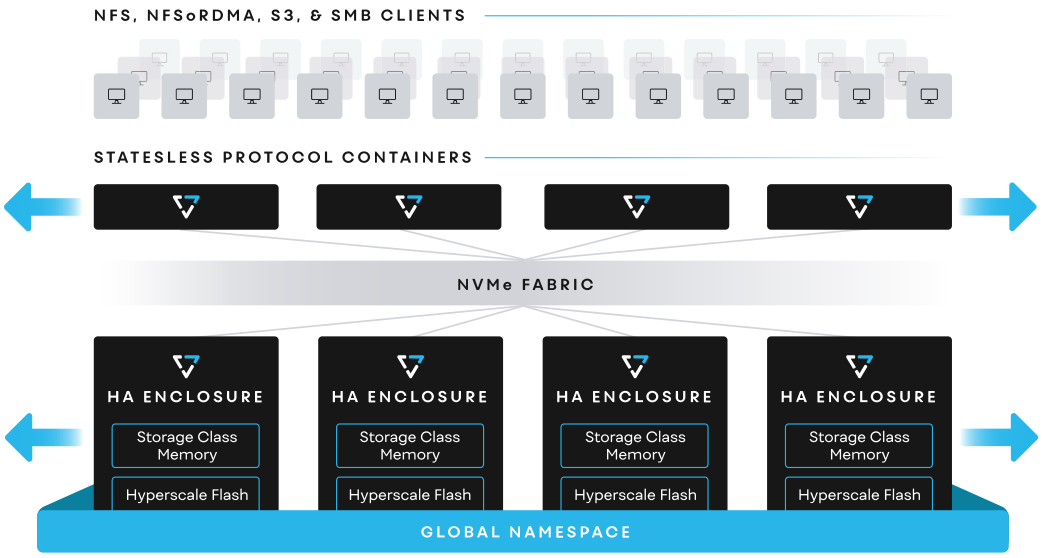

In this Disaggregated Shared Everything (DASE) architecture, every VAST Protocol Server in the cluster has direct access to all the cluster's storage media with PCI-e levels of low latency.

> 'Disaggregated' describes how a VAST cluster separates the VAST Protocol Servers that manage data from the resilient VAST Enclosures that hold the storage media, both flash, and storage class memory (SCM). This enables the cluster compute resources to be scaled independently from storage capacity across a commodity, data center scale network.

> VAST's 'Shared-Everything' concept combines NVMe-oF with a set of consistent data structures that makes it possible for all of the data and metadata, across all the enclosures, to be globally accessible by all VAST Servers in the cluster. This global view enables the system to implement global algorithms that define how the cluster builds an atomically consistent namespace, performs global data reduction and data protection.

SOLUTION DESIGN

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.

Owing to the distributed nature of employing a metadata server (Commserv) and loosely coupled data mover installations referred to as MediaAgents (MA), the Commvault product is able to control operations against storage devices, scaling up to petabytes organically as business needs grow.

The visualization of key elements of the control path versus data path are critical for keeping consistent performance as scale increases. Without these considerations, performance will be limited by individual compute hosts.

Likewise, VAST Data's share-everything architecture and stateless compute (CNodes) allow for the same capability as the data path requirements scale up.

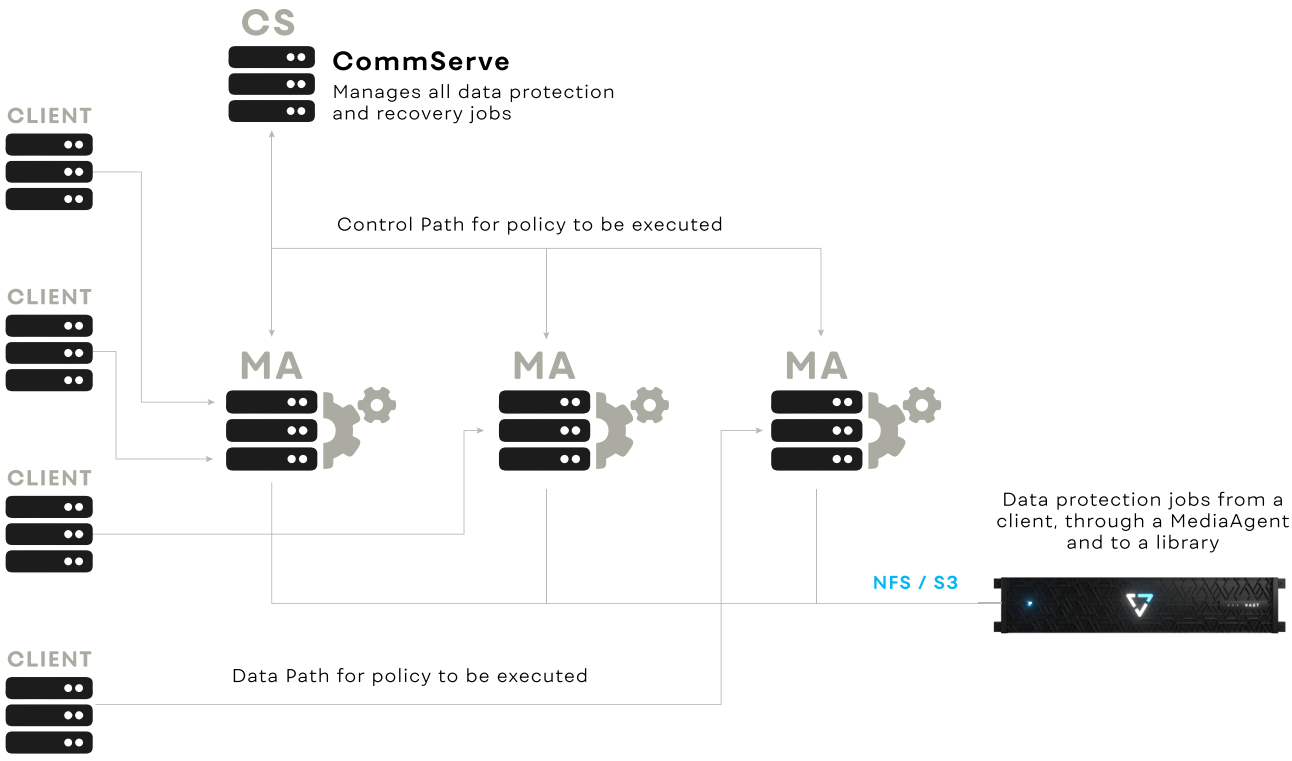

Control path point advantages

The solution uses Commvault to schedule the backup, restoration and replication of data from edge to core as a data mover, SLA enforcer, and single pane of glass for data protection. Commvault interoperates with VAST Data's platform with few changes, taking a product created in the era of the spinning disk into a future-proof, validated all-flash solution.

Data path point advantages

Owing to Commvault's lineage of operating system support for MediaServers, next generation transport capabilities such as the IO-uring buffer within NFS (Network File System) (NConnect) in the 5.3 linux kernel, VAST Data's multi-path driver package and NVIDIA RDMA (ConnectX-5/6), the solution is well positioned to drive exceptional backup and restore performance with minimal effort.

VAST Data's S3 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. The multi-protocol S3 and NFS workflow is well understood and once data has landed on VAST, the contents of a backup are readily available via a single and secure best practice operation that scales up to the 10's and 100's of petabytes that today's business requires.

TEST ENVIRONMENT INFRASTRUCTURE

ESX Host Configuration:

Dual Socket Intel(R) Xeon(R) Silver 4216 CPU @ 2.10GHz

256GB RAM

960GB disk

100GbE

Media Server Configuration:

Dual Socket Intel(R) Xeon(R) Silver 4216 CPU @ 2.10GHz

256GB RAM

960GB disk

100GbE network

Commserv Configuration:

Virtual machine under ESX 6.7

16 VCPUs

32GB RAM

20TB Disk space

Data Network Switch Fabric:

Mellanox SN2700 4x 100GbE ports

VAST Data Storage Configuration (Backup Target):

1x Quad Server Chassis incl 4 x VAST Servers (Part No: QUAD-CX5CX6)

1x HA VAST Enclosure (up to 40GB/s): incl. 676TB NVMe Flash, 18TB SCM (Part No: DF-5615-CX5X2)

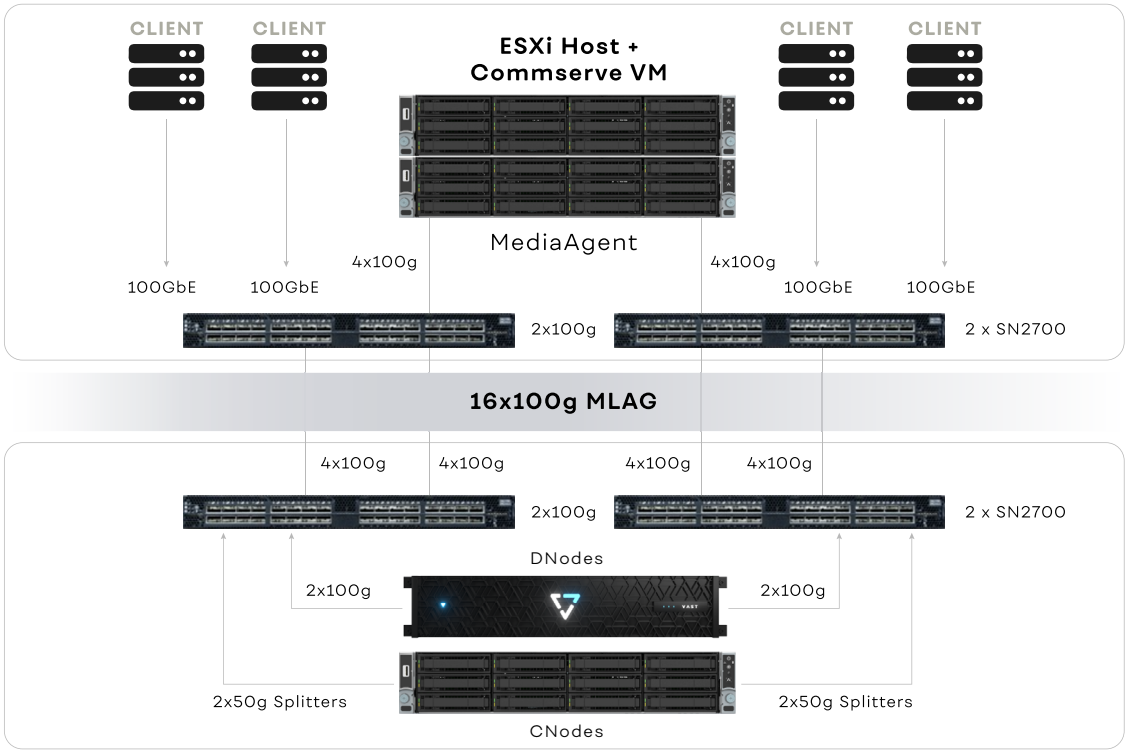

Test Environment: VAST Data "1x1" Storage Configuration Details

VAST DATA'S SIMILARITY-BASED DATA REDUCTION

Similarity-Based Data Reduction in Context

VAST's DASE scales data reduction much further. VAST's DASE architecture breaks the trade-off between data deduplication and scale by storing all the hash tables and other reduction metadata in the large pool of Storage Class Memory (SCM) across the VAST enclosures instead of controller DRAM. Since each enclosure adds more SCM, and therefore more hash space, along with its QLC flash, a VAST cluster can scale to exabytes as a single data reduction realm.

Much like traditional deduplication/backup appliances, VAST's Similarity-Based Data Reduction starts by breaking data into coarse blocks and hashing these blocks after new writes have been persisted into the SCM write buffer. This, however, is where the parallels to deduplication appliances end – as the hash function executed by a VAST cluster is very different from the types of hashing implemented by deduplication systems.

> Deduplication systems use cryptographic hashes, such as SHA-1 or SHA-256, that are designed to be collision-resistant, generating very different hash values from blocks that have small differences in their respective data

> VAST's Similarity Hash, on the other hand, is designed to enable a hashing function to measure the approximate distance between two blocks and generates the same hash for similar blocks of data.

It's a bit of an oversimplification, but Similarity Hashes can be thought of as language clues or semantic markers that suggest a high degree of correlation between different blocks. If two blocks generate similar hashes, they're likely to reduce well when compressed with a common compression dictionary.

Similarity-Based Data Reduction in Practice

The efficiency that can be gained from Similarity-Based Data Reduction is naturally data-dependent. While encrypted data will see almost no benefit from this approach to data reduction, other applications will often see significant gains. Reduction gains are, of course, relative – and where a VAST cluster may be 2X more efficient than a legacy deduplication appliance for backup data (at, say, 15:1 reduction), a reduction of 2:1 may be as valuable in an unstructured data environment where legacy file storage systems have not ever been able to demonstrate any reduction.

To create a cluster of Similarity-hashed and compressed blocks, VAST clusters compress the first block that generates a specific hash value and declare that block a reference block for that hash. When another block hashes to the same

hash ID, the system then compresses the new block against the reference block to create a difference (or delta) block. The data that remains from this global reduction method is a dictionary, which is stored in SCM as an attribute of the compressed reference block, and the delta objects that are stored as reduced symbols that can be as little as a single-byte and who do not also require their own dictionary. If two blocks are exactly the same, there is simply no delta to store.

When applications read from a VAST cluster, Servers traverse to the different compression dictionaries that are distributed across VAST V-Trees, enabling the system to minimize the scope of decompression to just a single reference block, in essence enabling the cluster to read 4KB objects within 1ms because the cluster doesn't have to decompress the whole namespace per every read as would be the case if all of the data was managed in a single dictionary.

Imagine being able to zip your entire namespace and read from this namespace with millisecond latency, this is the design point for VAST's Similarity-Based Data Reduction.

COMMVAULT & VAST'S SIMILARITY-BASED DATA REDUCTION

In order to measure the effectiveness of VAST's Similarity-Based Data Reduction, VAST conducted a series of tests using various Commvault settings across multiple dataset types representative of various real-world backup scenarios. The following section outlines the findings from the tests.

Compression Configuration Options & Best Practices

Commvault leverages GZIP as the default compression codec for unstructured data for filesystem backups and the LZO codec for VMware backups using the virtual agent.

By changing the default compression codec, VAST's data reduction algorithm is able to take better advantage of the duplicate blocks of data post compression that are committed to the backup target, both over NFS and S3 backup targets.

How To Change Commvault's Compression Codec

To customize your compression codec per storage device, use the commserve client, select libraries, your desired library, then properties, additional settings tab.

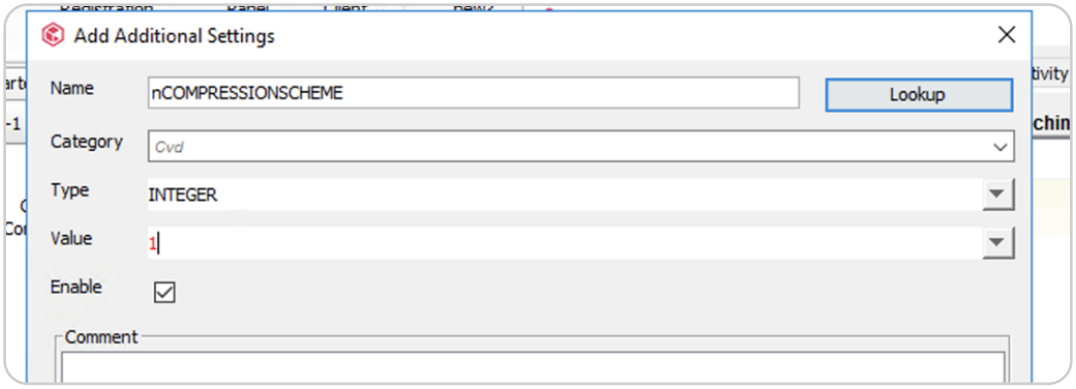

Enter nCOMPRESSIONSCHEME in the name field, and the corresponding integer value for the desired compression codec in the "Value" field:

> 0 = GZIP compression (default)
>
> 1 = LZO compression

Sample Backup Datasets

To demonstrate the impact of Commvault and VAST Data's combined data reduction effectiveness, we created and tested three datasets representative of common backup scenarios:

> **VMware VMs**
>
> 500GB across 15 virtual machines: 2x Windows Server 2016, 1x Windows 10, 5x Centos 7, 1x Centos 8, 6x Ubuntu 20.04.

> **Unstructured Data**
>
> 1.5 TB of files including 4k mp4, CAD drawings, GIS Geospatial Json, Genomics FASTa, Deep learning training sets (Imagenet png +jpg), word, excel, csv, rtf, pdf, mp3, txt and html files.

> **Database Servers**
>
> Two Microsoft Sql servers, each housing 1.8 TB and 1.2 TB of structured data inside Microsoft SQL. The databases are comprised of real datasets used for online analytics tutorials, including GIS, exported web server logs, and public records of topological features.

Test Results: Data Reduction Effectiveness

Total Combined Data Reduction Ratio
(Using various Commvault compression settings with Commvault De-Dupe = ENABLED and VAST Similarity Data Reduction = ENABLED)

Backup Dataset	Commvault Compression: GZIP	Commvault Compression: LZO	Commvault Compression: Disabled
VMware VMs	3.4:1	6:1	8:1
Unstructured Data	3.5:1	4.5:1	4.6:1
Database Servers	3.4:1	5.3:1	9:1

The table above summarizes the total data reduction ratio expected when combining Commvault's various compression codecs with VAST Data's Similarity-based Data Reduction.

Disabling compression within Commvault results in the best total data reduction possible, regardless of the source dataset. By turning off compression entirely, this grants VAST's data reduction techniques full visibility of the dataset, enabling it's algorithms to exploit the totality of the original source data. To disable compression on Commvault, please update these settings.

However, it is common for deployments that are trying to mitigate backup traffic congestion across local area networks (LAN) and wide area networks (WAN) by using compression on the host. In this case, the combination of effective compression and lighter CPU utilization makes LZO the best candidate for use with VAST.

Backup Performance Best Practices

As backup jobs occur, Commvault comes prepared with a multitude of different options to help avoid Wide Area Network / Local Area Network (WAN/LAN) traffic congestion by compressing data on the client or MediaServer.

The two options are GZIP and LZO. While GZIP is known for having good general compressibility for all media types, LZO provides a better balance of CPU effort and compression.

In other words, LZO requires lighter lift on the client (typically the least CPU capable devices in most environments, i.e. Laptops, Desktops, virtual machines, low performance pools), allowing backup speed to increase, lowering the time to completing backups (RPO) and potentially lowering the restore time (RTO).

These entries below are representative of typical backup performance rates for different compression algorithms configurable within Commvault for libraries and backup target storage protocols.

Protocol + Compression	GBytes / Hour
S3 + GZIP	1388.97
NFS + LZO	1676.85
S3 + LZO	1525.21

Our testing against both protocols and different compression schemes indicates the best backup rates occur when using the LZO compression algorithm. Combined with the fact that the LZO compression algorithm is the most compatible with VAST Data's advanced DDR capabilities, **VAST recommends using LZO as the preferred Commvault compression codec wherever possible.**

CONFIGURING MULTI-SITE DR TO VAST REMOTE

Extending the single site data protection policies with Commvault and VAST Data is as simple as using Commvault's Auxiliary Copy (auxcopy) feature to copy jobs that have data landing inside a primary site for the purposes of rapid restore to an external MediaAgent inside the same Commserve. Using auxcopy, all SLA's for data protection are visible inside Commserve's reporting, assuring a single pane view for the secure state of the business' critical assets.

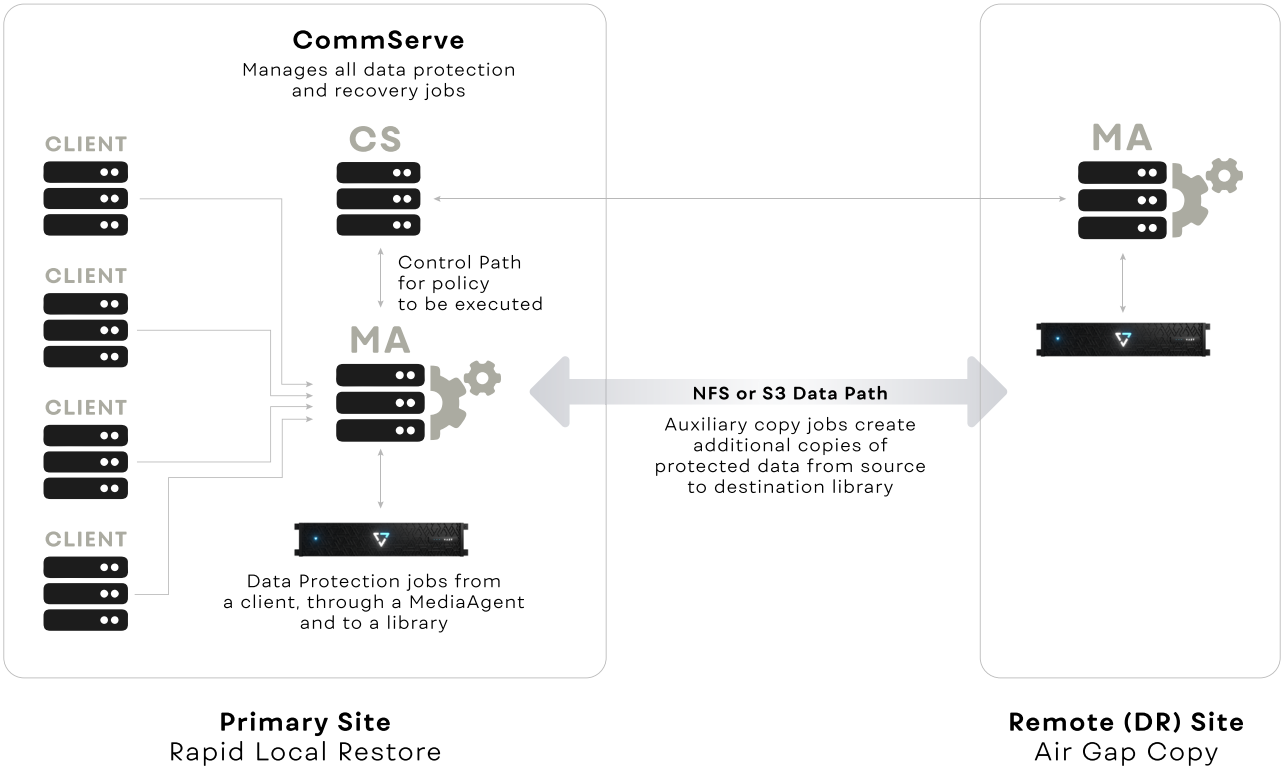

The following instructions provide a configuration that will replicate local NFS backup data to a remote site via S3, this is typical as NFS allows for portions of a backup to be accessed for rapid recovery, while using S3 for remote auxcopy operation over WAN allows for enhanced security hardening.

1 Establish backup to primary VAST Data library

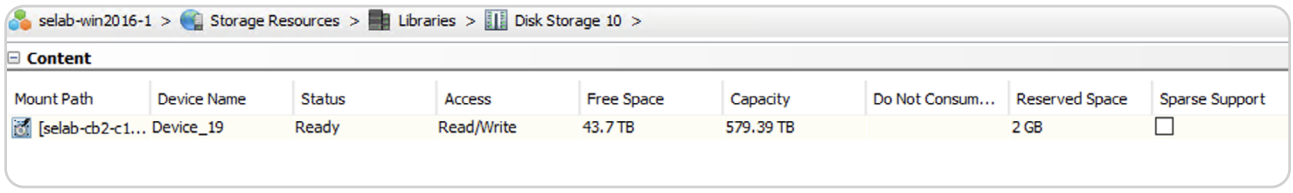

2 Configure disk library with instructions from Install / Configure NFS Backup Target On VAST Data section

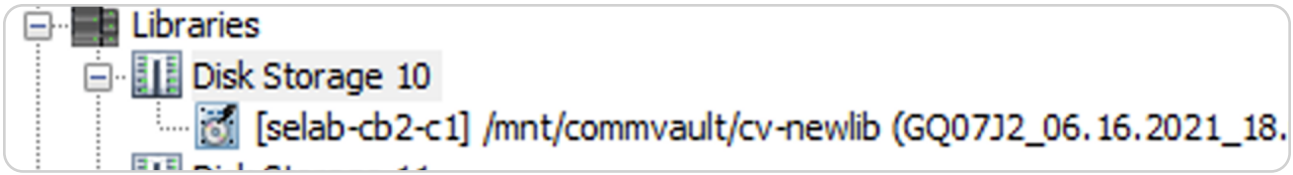

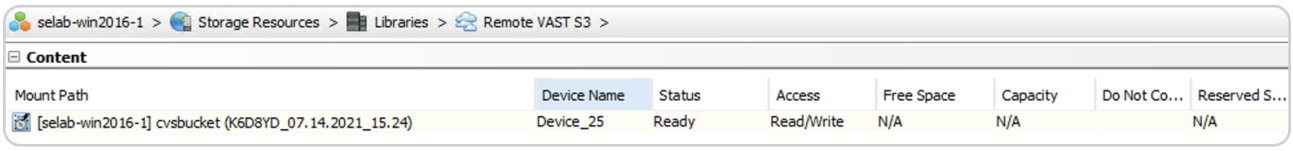

3 Configure cloud library with instructions from Install / Configure S3 Backup Target On VAST Data section

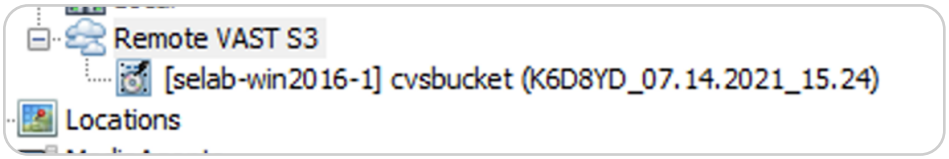

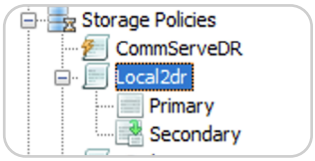

4 Create a storage policy with two copies, primary with the Disk library target, secondary with the S3 Cloud target

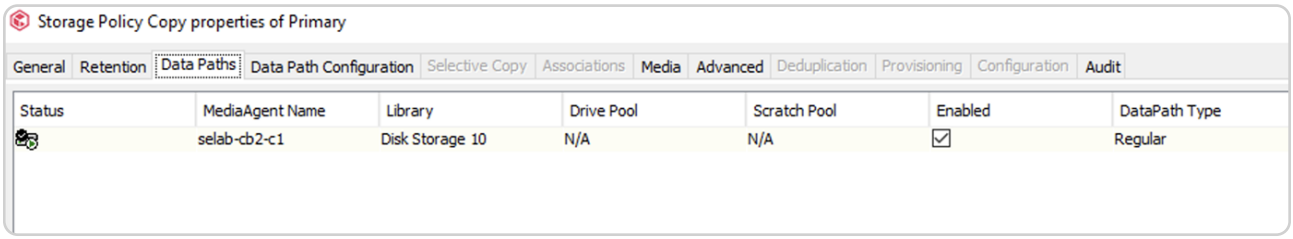

View of NFS target policy (Primary) above

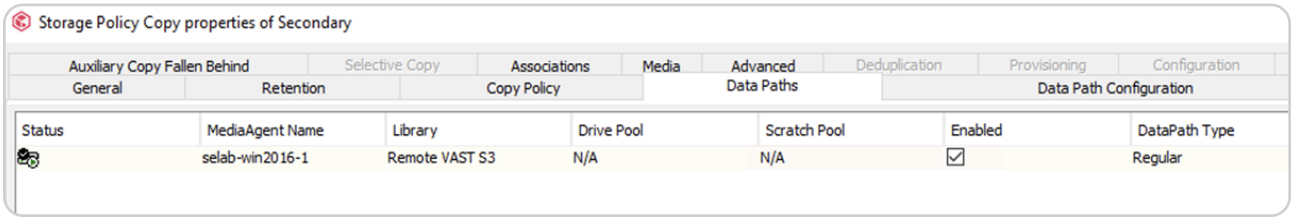

View of S3 target policy (secondary) above

5 Associate backup client with the storage policy containing both targets for your subclient

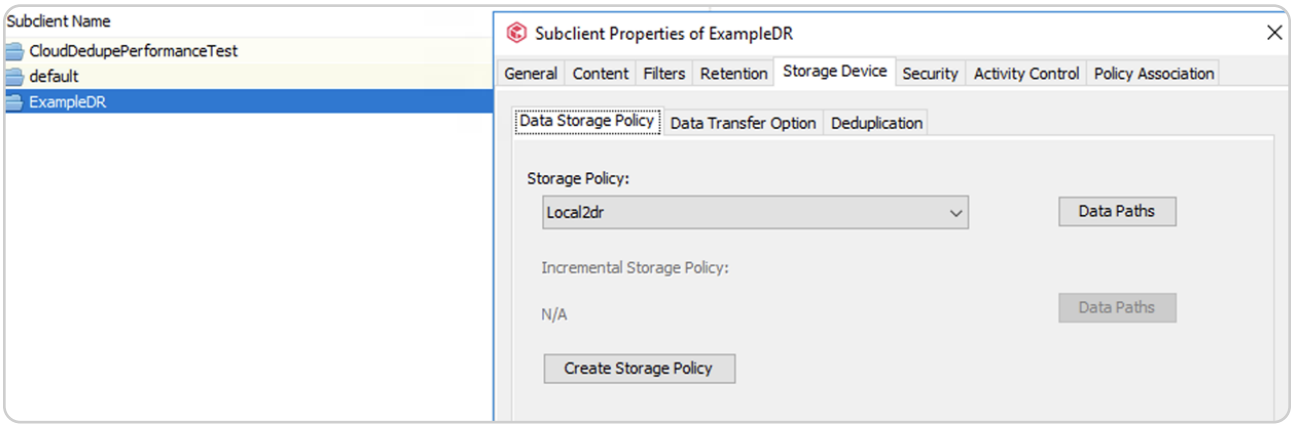

6 Run a backup on the subclient, afterward, run an AUX copy job on the storage policy, validate it has started in the job tab under active jobs

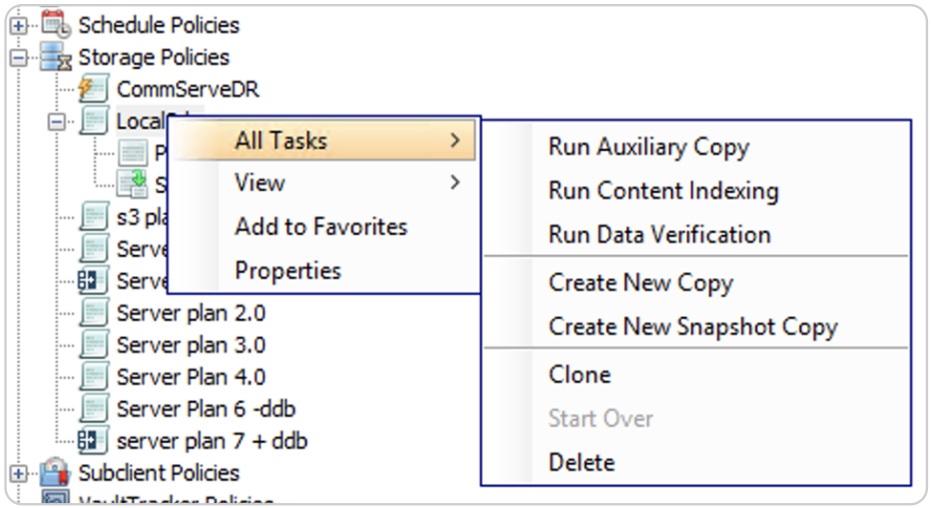

SETTING UP S3 ON VAST DATA FOR COMMVAULT

A typical Commvault integration into VAST Data is broken down into the typical steps Commvault would use with local attached flash for the deduplication database, and VAST Data for NAS, or S3.

Populate the S3 account with a bucket name

Use your s3 credentials with your s3 client of choice to create an S3 bucket.

Create an S3 account on VAST Data

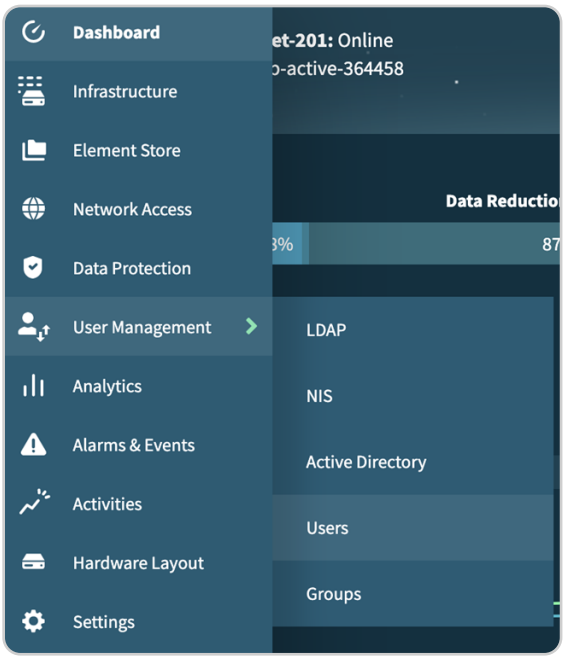

1 Select the User Management section from the right hand side of the VAST UI

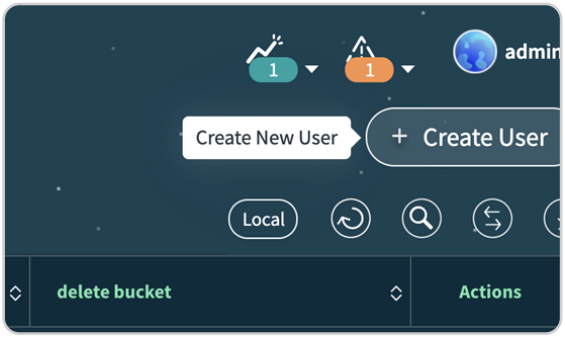

2 Pick the top left option to create a user

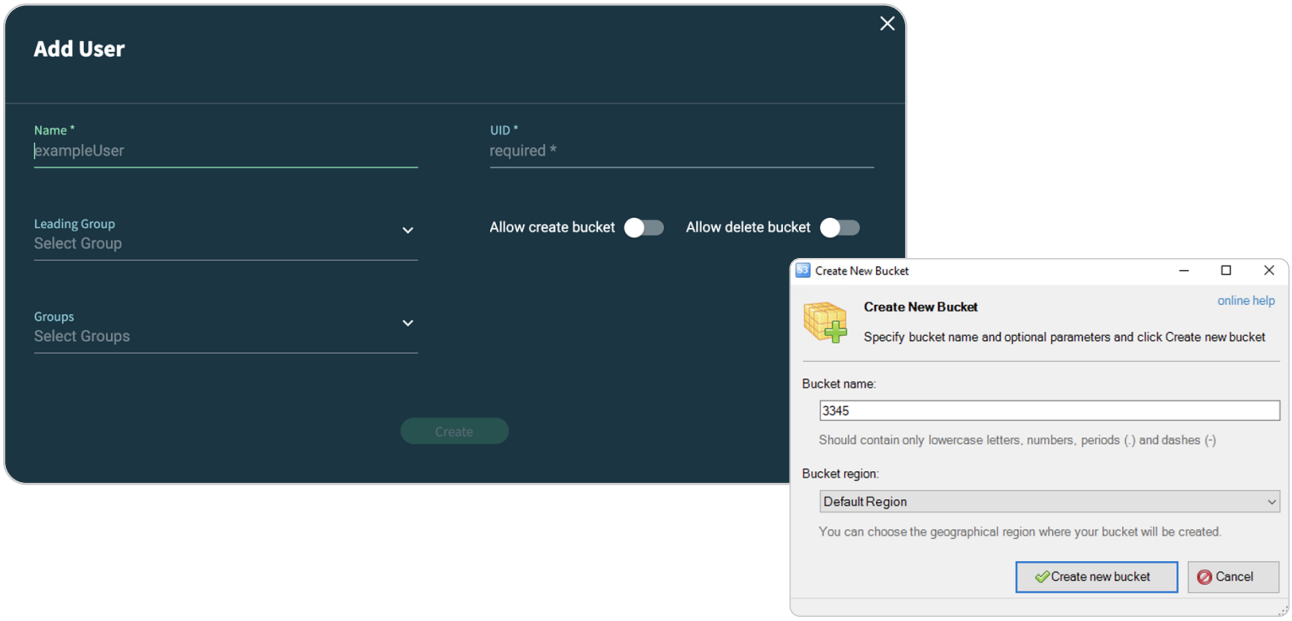

3 Populate an account name and UID, slide the allow bucket delete and create buttons for the account

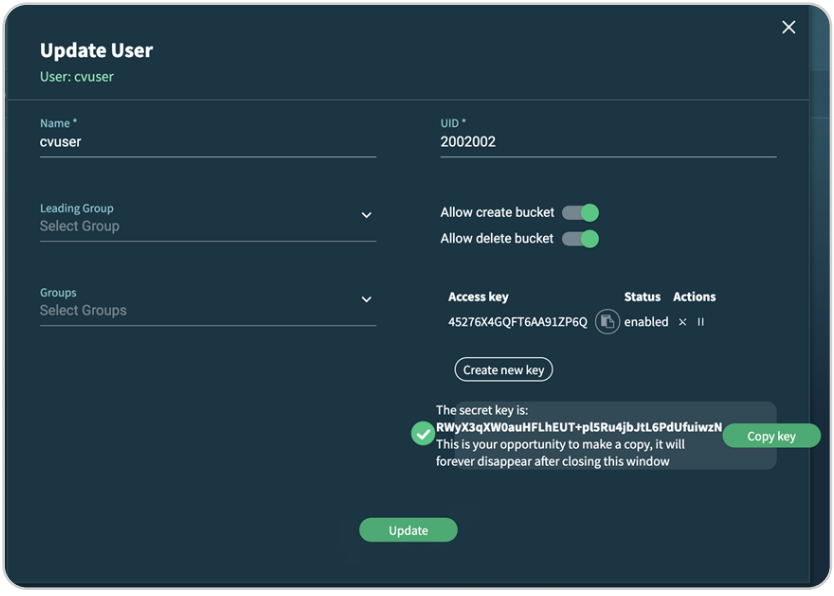

4 Click generate key for both the access key and secret key to be generated. Once this dialog is closed, the secret will be gone, so retain it if you require it for configuration**

**If you lose the secret, you can create a new key. Both keys can coexist for your user and will both have access to the buckets the account contains.

INSTALL / CONFIGURE S3 BACKUP TARGET ON VAST DATA

1 Take the bucket name and credentials populated in the previous section (Setting Up S3 on VAST Data for Commvault) , and populate the client with the access key, secret key and a VAST Data vip (found under network access, vips in the VAST UI)

2 Load the cloud library dialog under storage

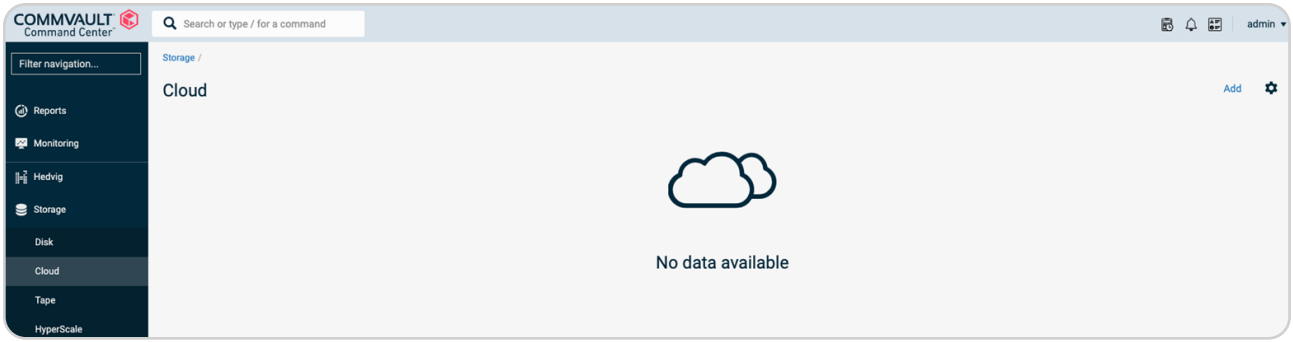

3 Select Cloud Storage

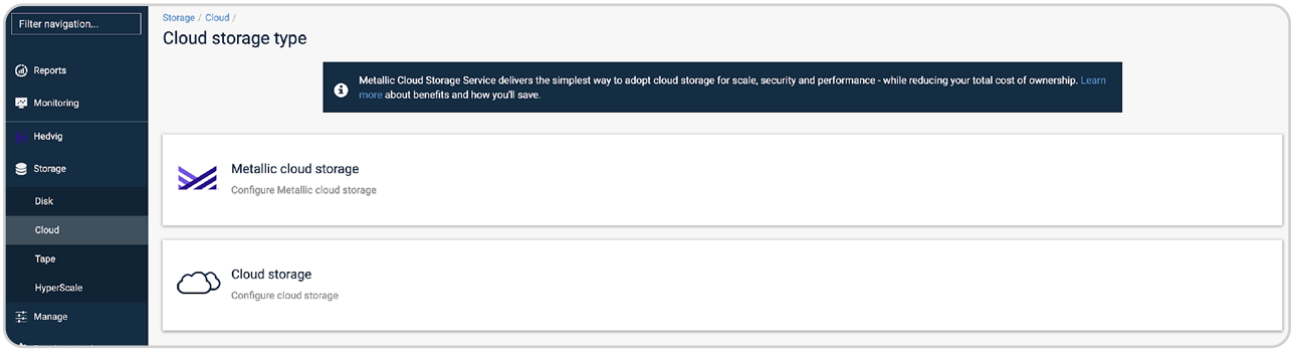

4 Populate the S3 credentials created in the Create an S3 account on VAST Data section

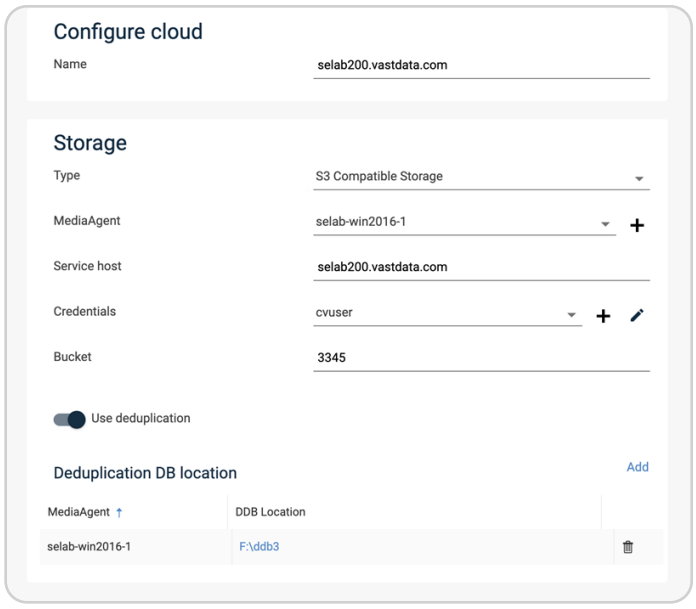

5 Cloud library is now visible

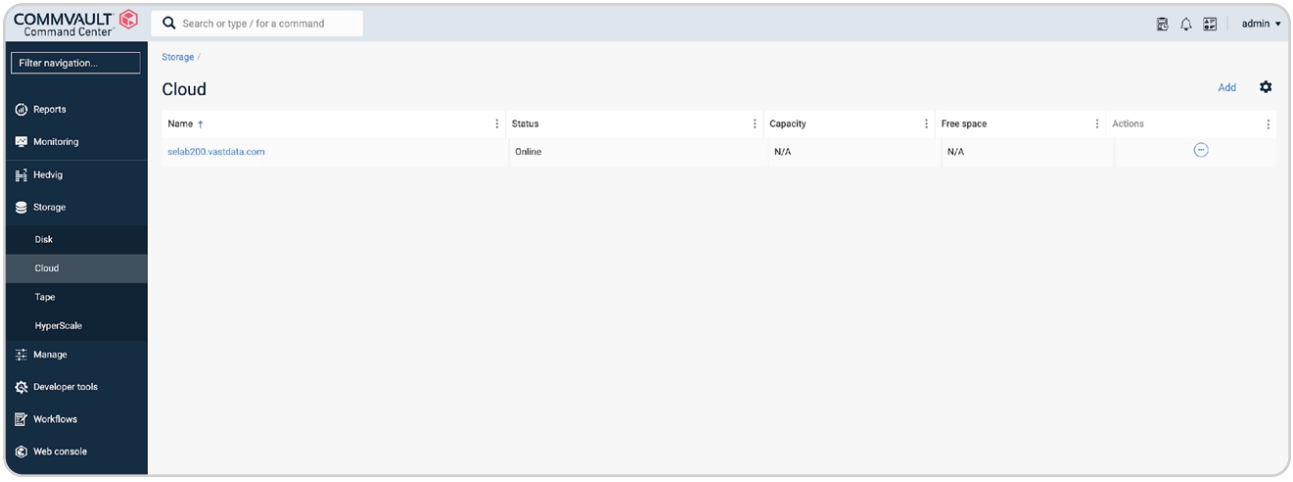

6 Create storage plan

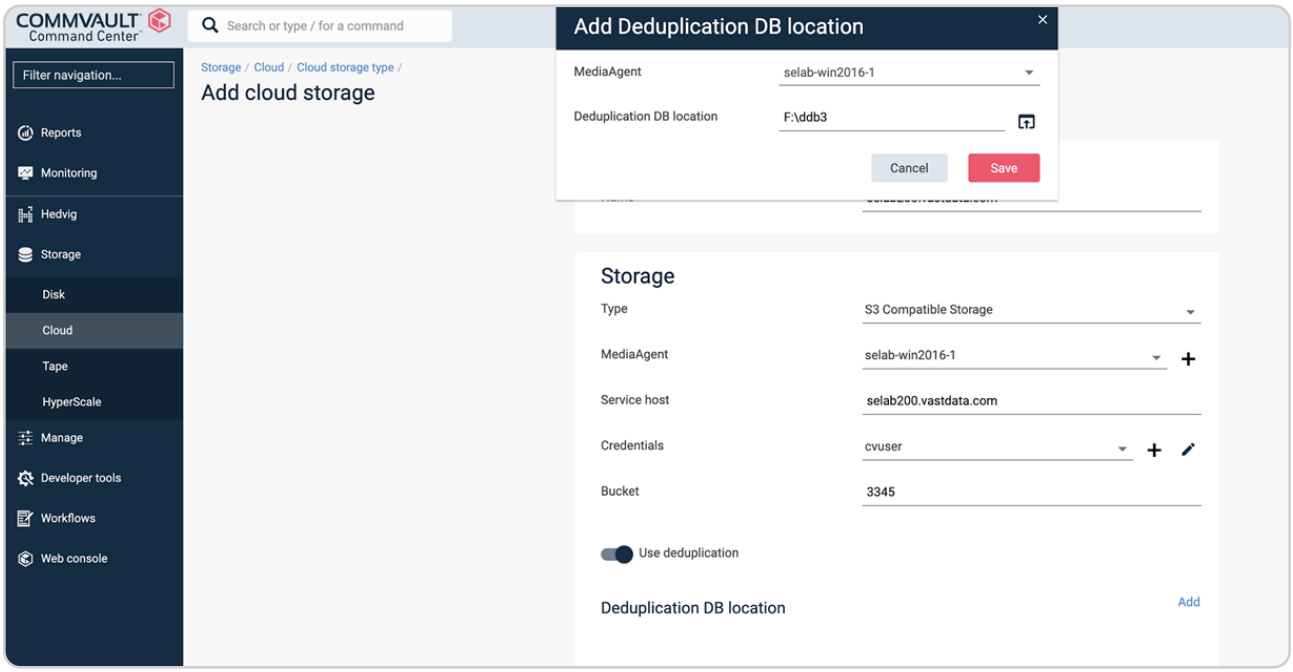

7 Select dedicated DDB location or use global DDB policy.

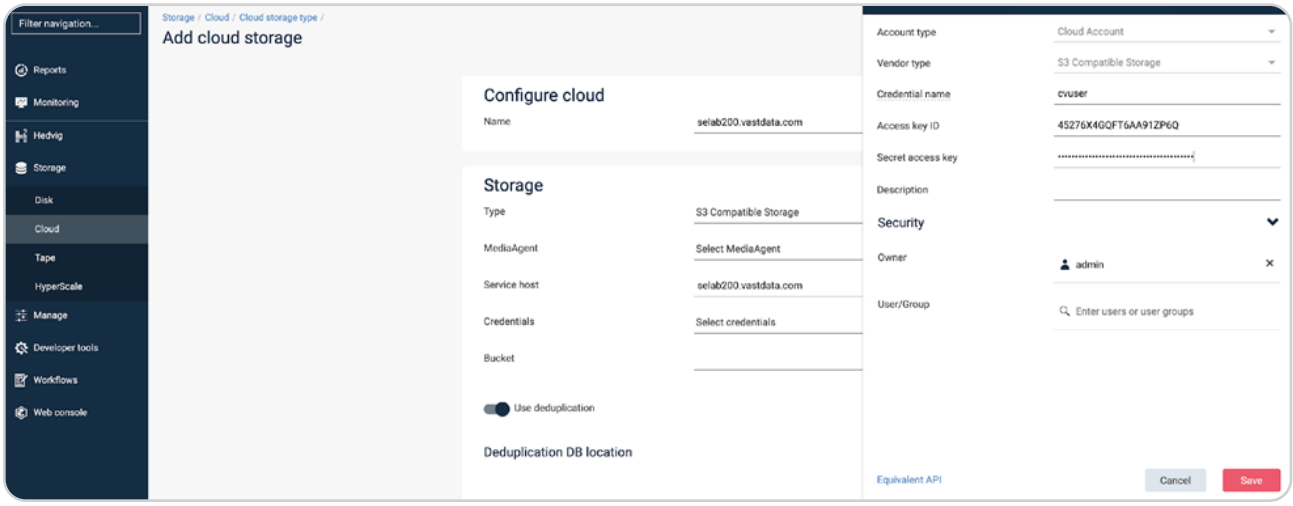

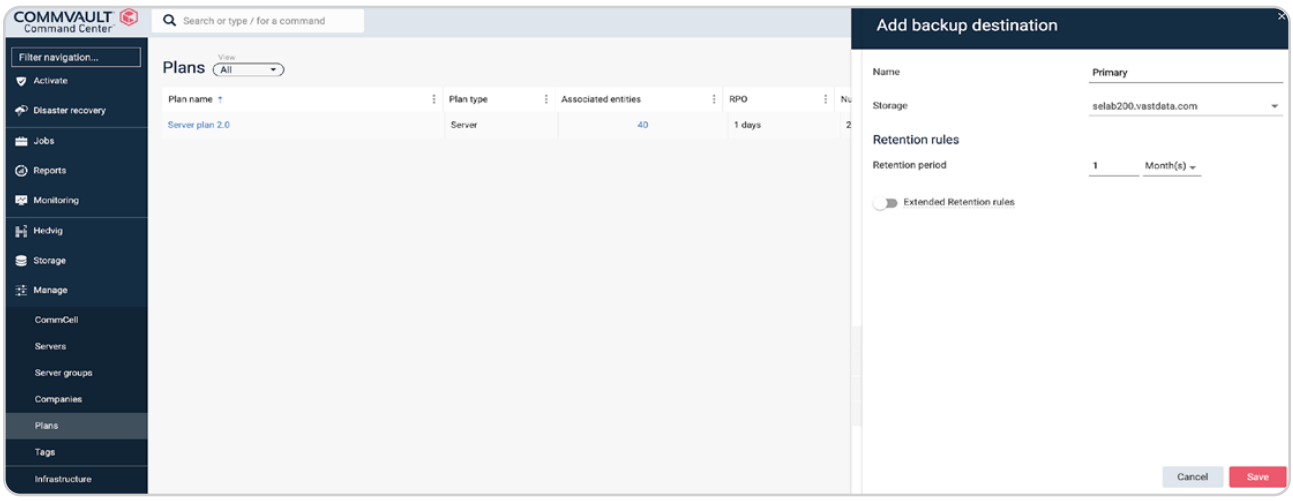

8 Create a sub client, browse to the content you wish to protect and associate it with a plan, then run the backup, check inside of jobs tab under active jobs to verify success

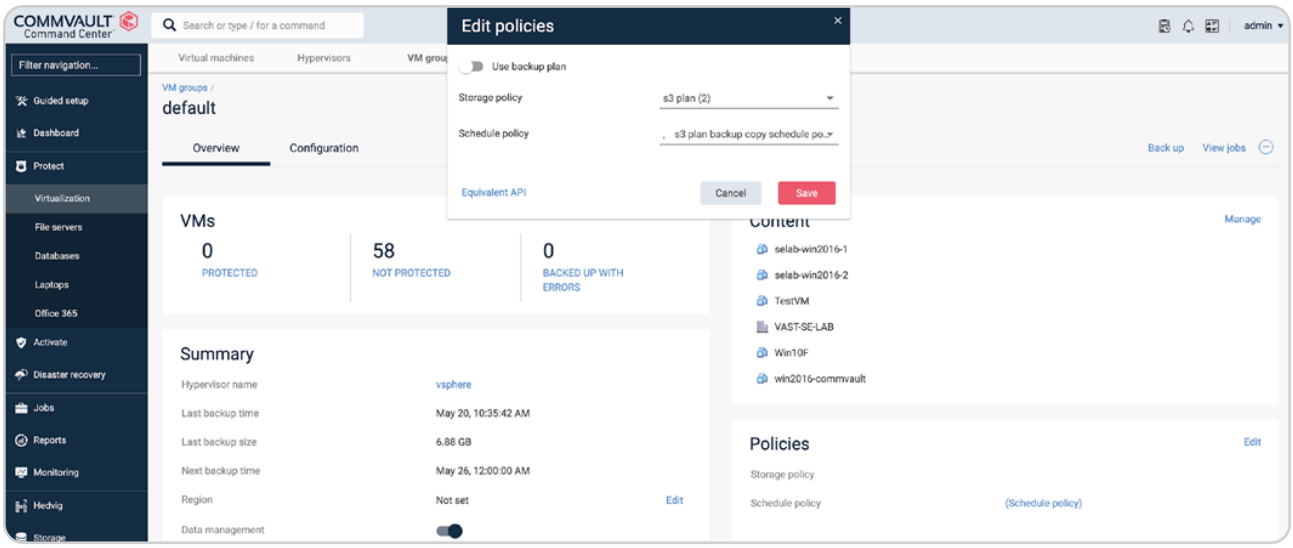

INSTALL / CONFIGURE NFS BACKUP TARGET ON VAST DATA

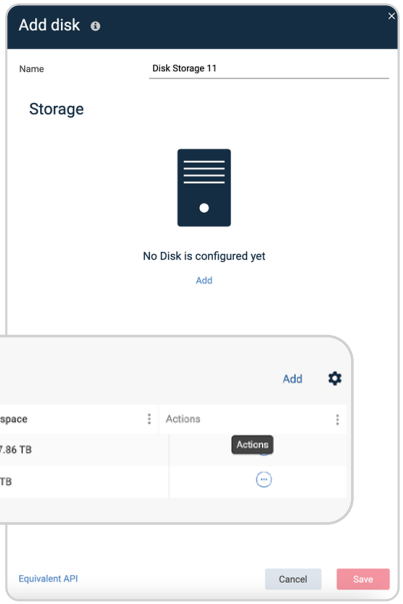

1 Open the storage menu, select disk, click add, add again, this will create a form to put details for the NFS Backup Target

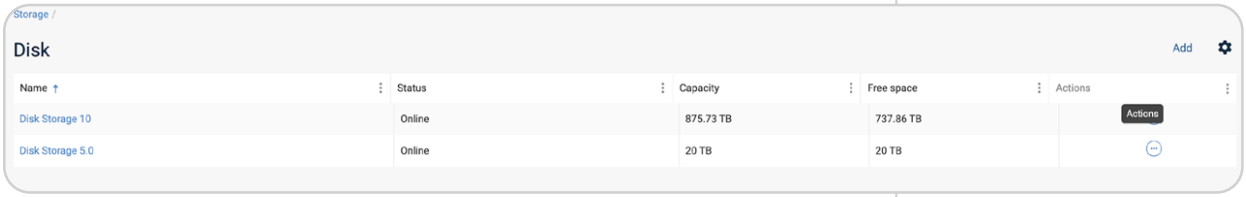

2 Select media agent with NFS mounts (See mount options section)

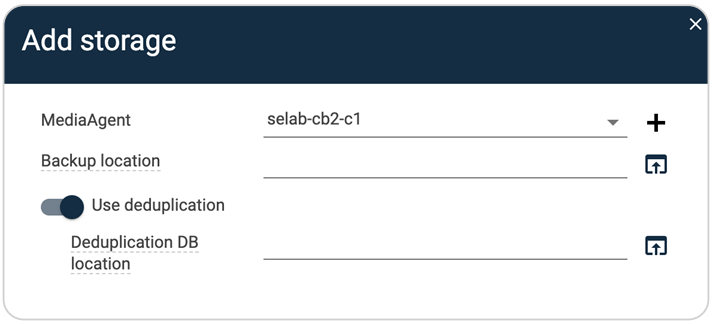

3 Select backup target location folder within or at the mount point, uncheck local deduplication database (should be dedicated flash DAS w/ global deduplication policy)

4 Click save and associate a plan and subclient to utilize for backup

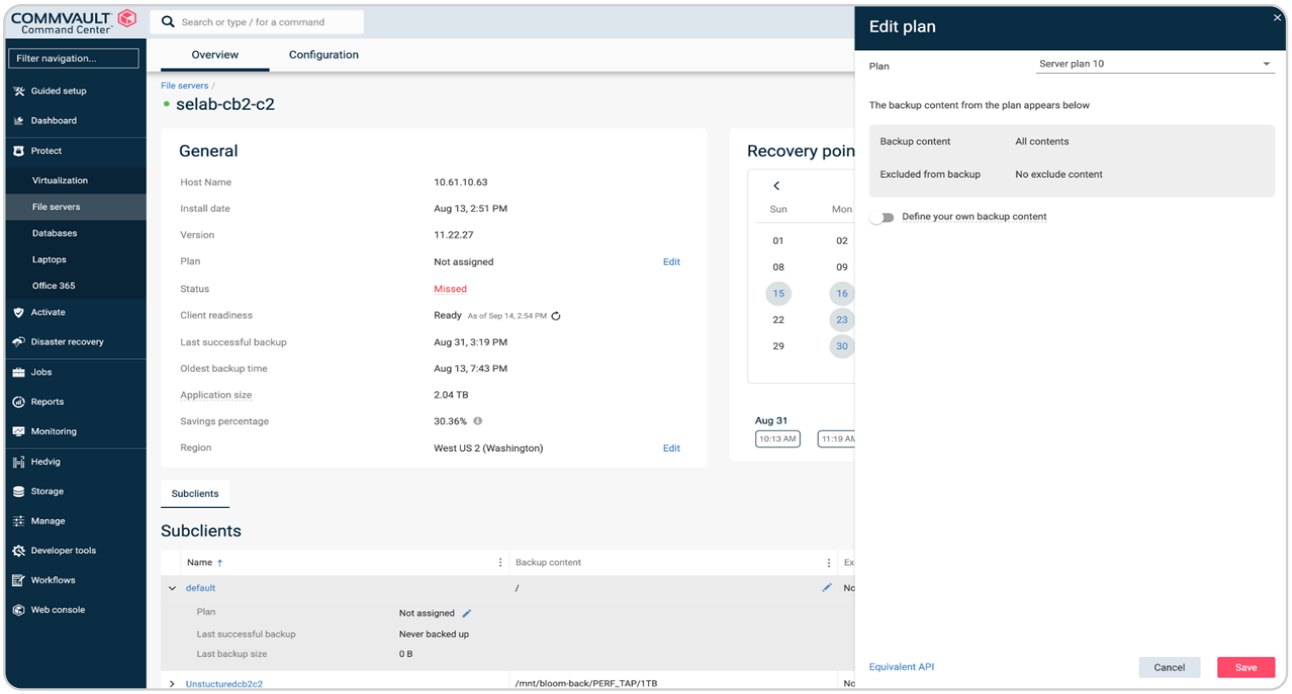

5 Click save to launch backup

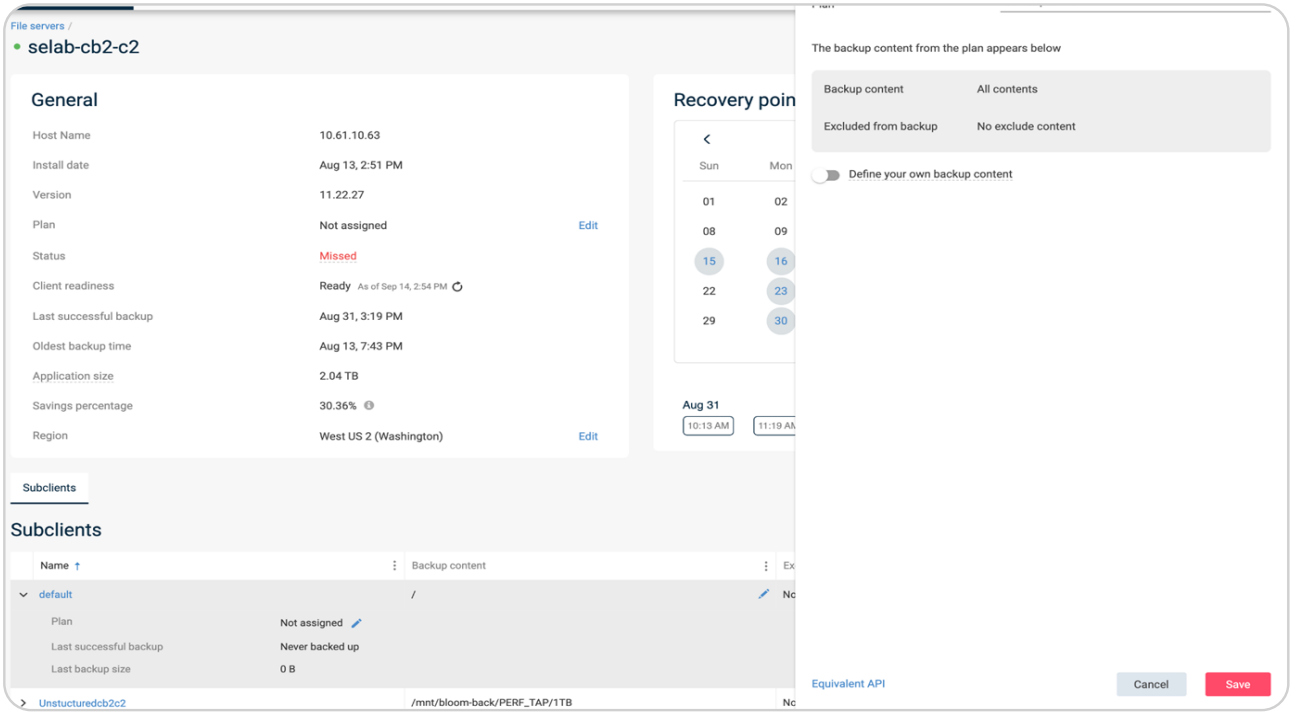

6 Do a test backup

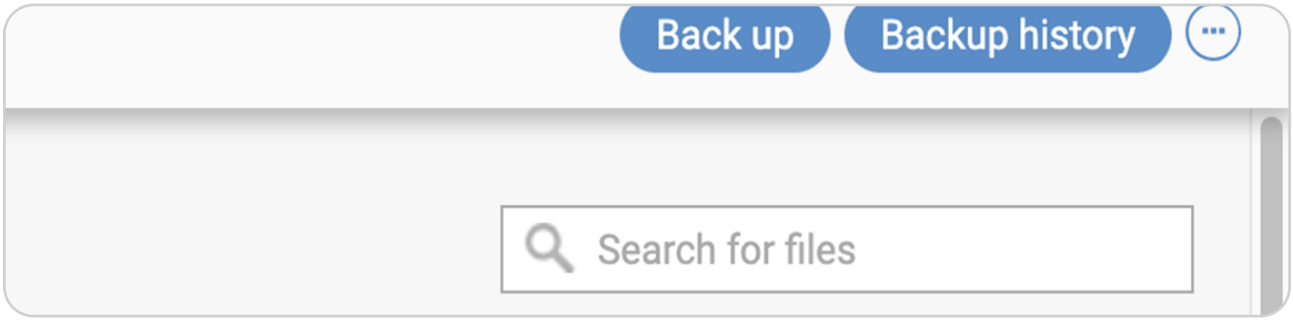

7 Check inside of jobs tab under active jobs to

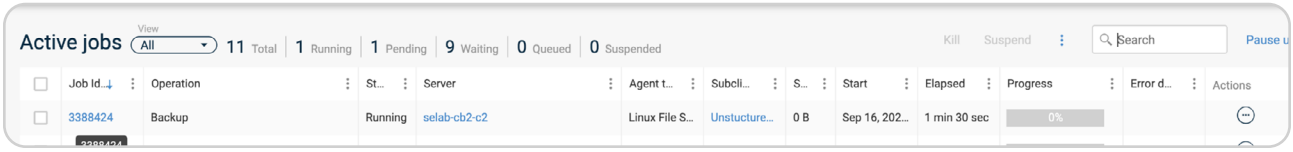

APPENDIX A: VALIDATING MEDIA SERVER PERFORMANCE ON LINUX MEDIA SERVER MOUNT OPTIONS

VAST data can attach to a Linux MediaAgent using NFS mount options. Different options provide different characteristics. To show how a single mount vs multipath can dramatically improve both backup and restore operations, use these steps:

1. Use the mount options below for your test
2. Create an NFS library as outlined above in Install / Configure NFS Backup Target On VAST Data
3. Open the Commcell Browser, expand all libraries until you find the created NFS library
4. Expand the library until you see the host/path information exposed
5. Right click, select Validate Mount Path
6. Populate the criteria for file size, number of writers/threads, file block size and file count to align with your best estimate of your backup target you will be deploying

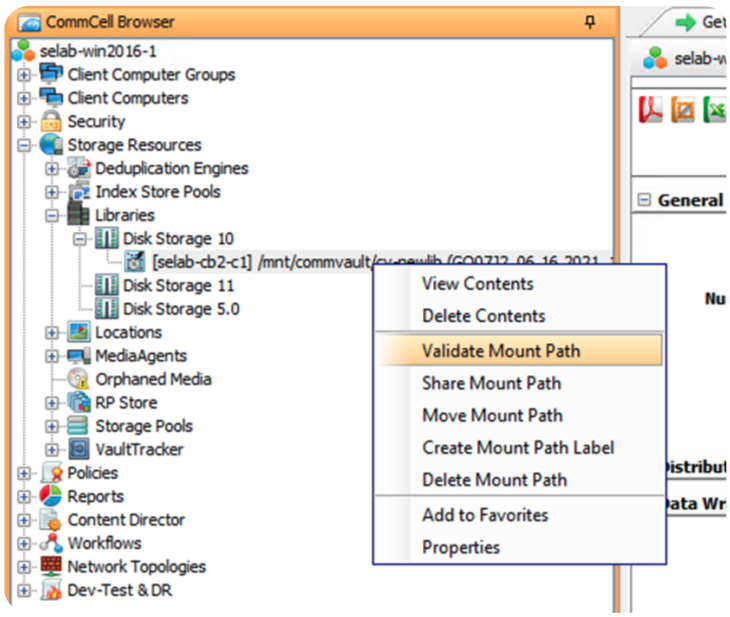

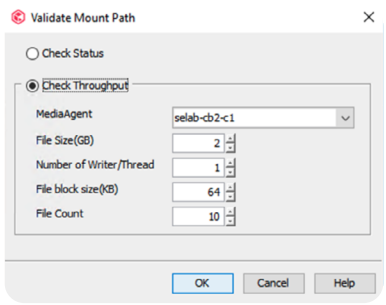

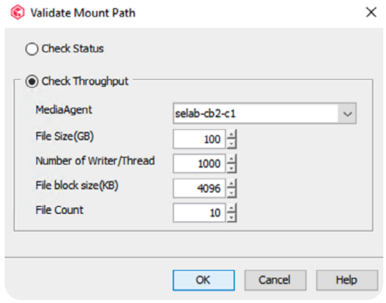

Mount options for nfs/tcp

```
mount -v -o vers=3,proto=tcp 172.200.3.1:/Commvaultlibrary /mnt/Commvaultlibrary/
```

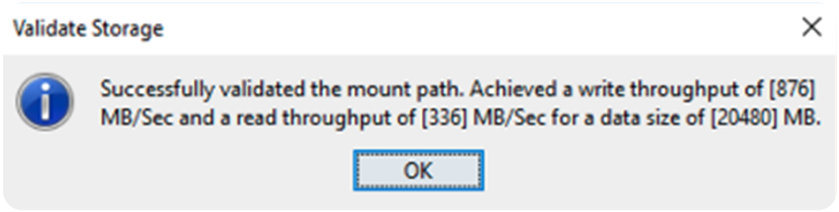

Mount options for nfs/tcp/nconnect

```
mount -v -o vers=3,nconnect=16,proto=tcp 172.200.3.1:/Commvaultlibrary /mnt/Commvaultlibrary/
```

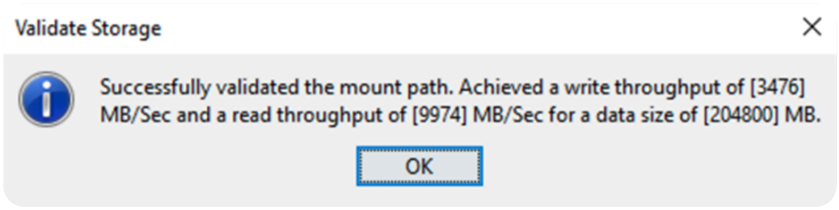

Mount options for nfs/multipath/tcp

Multipath tcp requires a package be installed specific to your MediaAgent's kernel and base OS. Contact support for access to the driver.

```
mount -v -o
vers=3,proto=tcp,port=20049,nconnect=16,localports=172.200.5.61,remoteports=172.200.3.1-172.200.3.8
172.200.3.1:/Commvaultlibrary /mnt/Commvaultlibrary/
```