

WHITE PAPER

Configuring Commvault® RO1105 Remote Office Appliance With Pure FlashBlade®

Deployment and Best Practices Guide.

Contents

Introduction	4
Solution Overview	4
Solution Architecture.....	4
Data Control	6
RO1105 Remote Office Appliance	6
Data Storage	6
Pure FlashBlade	6
How Commvault Uses Object Storage	6
Cloud Accelerator	6
Ransomware Protection	7
Basic Configuration	7
Set up Pure FlashBlade.....	7
Configure Network for Data Access	7
Add FlashBlade Replication Links	10
Create Protection Policy for Replication	11
Provision File System for Commvault DR Backups	11
Provision Object Bucket for Backup Data	16
Optional: Add CA Certificate	18
Set up Commvault RO1105 Appliance	19
Initial Configuration	19
Guided Setup	19
Enable Virtualization Solution	23
Deploy Commvault File System Clients	24
Create Subclients	24
CommServe Availability	25
General Best Practices	25
Configure a Single FlashBlade Bucket and Mount Path	25

Share Mount Paths with Data Servers.....	25
Deploy Cloud Accelerator on All Clients.....	27
Use Commvault Client-Side Compression by Default.....	28
Use Commvault Client-Side Deduplication	28
Match VMware Disk Format and Transport Mode	28
Disable TLS, if Allowed.....	29
Data Servers.....	29
Operating System	29
Virtual Server Agent Mode.....	29
Data Server Hardware Specifications	29
Data Server Count	30
New Deployment.....	30
Deduplication Database Storage.....	30
Index Storage	31
Advanced Configuration.....	31
Run Multiple Client Data Readers	32
Increase Cloud Thread Pool Size	32
Configure IntelliSnap Snapshot Integration.....	34
Additional Resources	34
Common Failure Scenarios	34
Cloud Target Creation Fails.....	35
Backups or Restores Fail to Send Data	35
Throughput.....	36
Terms and Concepts	36
About the Authors	38



Introduction

As IT resources grow more and more constrained, deploying and managing complex solutions consumes a larger portion of budgets. In few places is this more visible than backup and recovery. Protecting data often requires complex solutions that do not meet SLAs for your most critical data, or multiple point products that drive up complexity and cost.

The combination of the Commvault® RO1105 appliance and Pure FlashBlade® gives you a simple yet powerful, jointly engineered data management solution to handle all your recovery needs, one that easily grows along with your business to help you avoid the cost and complexity of data silos.

This guide covers basic deployment and configuration of the RO1105 and FlashBlade systems, as well as best practices to optimize backup and recovery performance.

Solution Overview

This joint reference design ties together Pure FlashBlade Object Store and Commvault Backup and Recovery, with its Cloud Accelerator feature, to provide a fast, simple, scalable data protection platform. Cloud Accelerator lets Commvault shift from a data mover role to data control, with endpoints communicating directly with FlashBlade to protect and recover data.

Commvault client-side deduplication and compression ensure efficient use of network and storage resources, while the throughput of FlashBlade cuts backup and recovery times to ensure you can meet your SLAs. Commvault provides the operational foundation for managing and leveraging your protected data. In addition to traditional data placement, cataloging, and reduction, Commvault enables analytics-driven capabilities, such as actionable reporting, as well as copy data management and other orchestration.

FlashBlade adds a storage layer that brings superior performance to the functionality of Commvault. FlashBlade Object Store takes away the pains of growing backup storage, with seamless growth to multiple petabytes.

Solution Architecture

Figure 1 shows the logical architecture of the solution. Commvault CommServe® and Data Server operate in the data control layer, managing metadata catalogs, data placement and retention, scheduling, and agent health, as well as reporting and data reduction. Pure FlashBlade occupies the data storage layer, serving an object storage target for protected data from the workloads. Commvault agents sit in the workload layer, backing up and recovering to the data storage layer, and sending metadata to the data control layer.

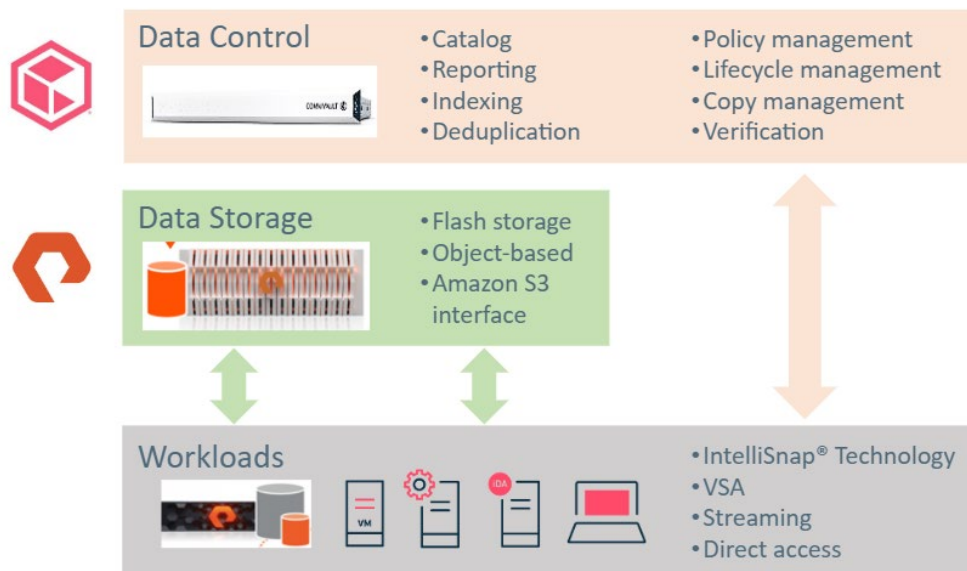


Figure 1. Logical architecture

Figure 2 shows the physical architecture of the solution. A Commvault RO1105 Remote Office Appliance and Pure FlashBlade each sit in the primary and DR sites. In addition to the hosting logical roles as described, the Commvault appliances manage data replication between sites and, optionally, to public cloud storage using Commvault's Deduplication Accelerated Streaming Hash (DASH) Copy process. Backup clients send metadata to the CommServe and Data Server on the RO1105 appliance, and they communicate directly with the FlashBlade to read and write data.

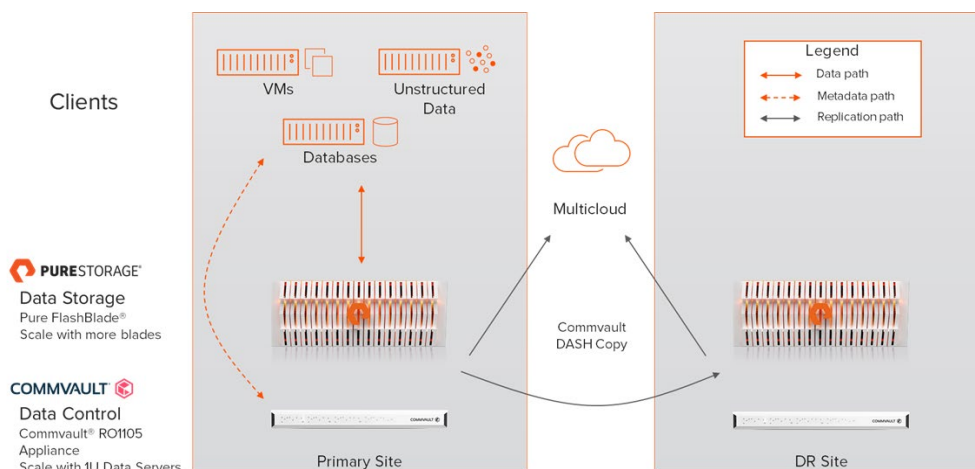


Figure 2. Physical architecture

Both the data control and data storage components can be easily scaled. You can grow the data storage by adding individual blades to the FlashBlade, up to the limits of a single chassis. This does not require any configuration changes or downtime. Scaling to multiple chassis requires additional hardware configuration, but the object storage components on the FlashBlade and in Commvault do not need any changes. As data storage grows, the data control plane needs to expand. You can add 1U [Data Servers](#), with minimal configuration, to manage the additional capacity.

Data Control

RO1105 Remote Office Appliance

The [Commvault RO1105 Remote Office Appliance](#) is the command and control center, or the Data Control Plane, for all data management functionality. The RO1105 is a fully integrated appliance that includes:

- Pre-installed Commvault software
- 4 x 10GbE ports plus 2 x 1GbE ports
- 960GB of metadata storage capacity
- Optional fibre channel HBA cards to enable IntelliSnap® backup copy and VADP SAN transport mode

Once deployed and configured, the RO1105 will manage all data protection, data life cycle management, cataloging, and reporting operations from the Commvault Command Center.

The RO1105 is ideally sized to protect up to 50TB of source data, up to 500 virtual machines, and manage 100TB of FlashBlade storage.

Data Storage

Pure FlashBlade

Pure FlashBlade acts as the main backup storage, providing simple configuration, fast backup, Rapid Restore ([link](#)), and easy capacity and performance scaling. Commvault leverages the native fast object storage on FlashBlade through the Amazon S3 protocol to minimize the configuration effort and take full advantage of the parallelism of FlashBlade.

The base solution configuration includes a FlashBlade chassis with seven 52TB blades for each site, capable of up to 7GB/s read throughput. The FlashBlade capacity and performance can be expanded a blade at a time, up to 150 blades, capable of dozens of GB/s, with the simplicity of a single object bucket.

How Commvault Uses Object Storage

Commvault uses a very different process to read and write to object storage than it uses with file storage. With file storage, each backup or restore stream is broken into chunks, which are written sequentially into large data files.

With object storage, the data chunks are broken into smaller BLOBs (binary large objects) before they are written. The system creates a thread pool that is shared across all streams. As each thread is activated, it opens a TCP connection to storage, and the threads write BLOBs in a highly parallel manner. Commvault automatically expands the thread pool as needed, up to a tunable maximum, to improve throughput. Because the threads each have their own TCP connections, Commvault's model results in excellent load distribution across FlashBlade blades, and distribution improves with more clients.

Cloud Accelerator

The Cloud Accelerator feature in Commvault is at the heart of the solution. It lets clients write directly to FlashBlade without a need for full-fledged MediaAgents. The Data Server controlling the backup and restore jobs can then manage more activity with fewer resources. To use Cloud Accelerator, clients must be able to access the FlashBlade directly over an IP network. Figure 3 illustrates the architectural differences with and without Cloud Accelerator.

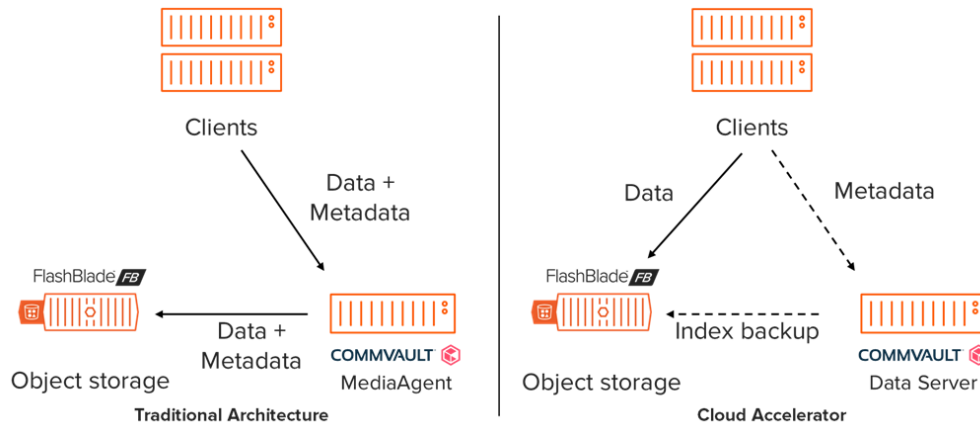


Figure 3. Traditional and Cloud Accelerator architectures

Ransomware Protection

Commvault includes a number of [ransomware detection and mitigation capabilities](#), to not only help ensure you can recover quickly from an attack but also to identify and react to attacks early and minimize potential damage.

You can add an extra layer of defense for your backup data on FlashBlade by leveraging [SafeMode™ Snapshots](#). See [Achieve greater protection against ransomware with Commvault and Pure Storage](#) and [Ransomware Protection with Pure Storage and Commvault](#) for more details on the solution and how to implement it.

Basic Configuration

The basic configuration is intended to help you implement the solution as quickly and simply as possible. Refer to the [General Best Practices](#) and [Advanced Configuration](#) sections for advanced optimization guidance.

Note: This document assumes you have an Active Directory forest deployed.

Set up Pure FlashBlade

Work with your installation team to perform the initial setup of the FlashBlade Purity//FB OS and networking. You will then need to set up data access, create replication links between FlashBlade arrays, and provision storage.

Configure Network for Data Access

1. You must create a subnet for backup data traffic (Figure 4).
 - In the Purity//FB GUI, select Settings > Network.
 - In the Subnets list, click the Add (+) button in the Subnets title bar. The Create Subnet pop-up window appears.
 - In the Name field, type the name of the subnet.
 - In the Prefix field, type the IP address of the network prefix and prefix length in the form ddd.ddd.ddd.ddd/dd for IPv4, or xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx/xxx for IPv6.
 - In the VLAN field, specify the VLAN ID to which the subnet is configured. Valid VLAN ID numbers are between 1 and 4094.

- In the Gateway field, type the IP address of the gateway through which the data vip communicates with the network in the form ddd.ddd.ddd.ddd for IPv4, or xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx for IPv6
- In the MTU field, specify the maximum transmission unit (MTU) of the data vip. If the MTU is not specified during subnet creation, the value defaults to 1500.
- Click Create.

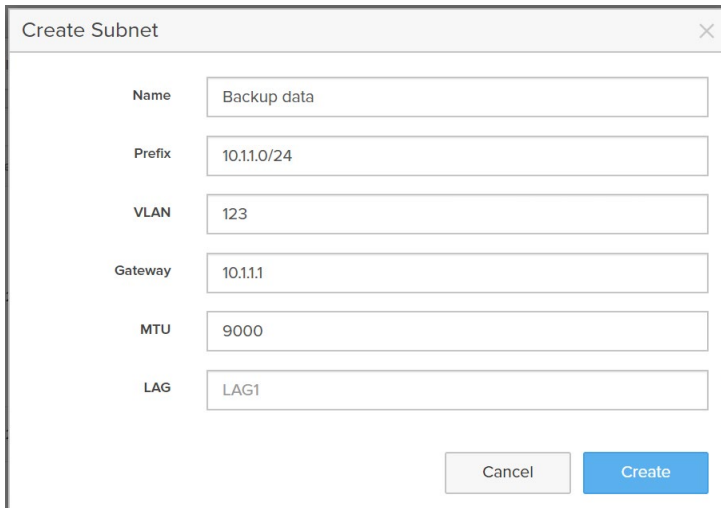
A screenshot of a 'Create Subnet' dialog box. The dialog has a title bar with 'Create Subnet' and a close button. It contains several input fields: 'Name' with the value 'Backup data', 'Prefix' with '10.11.0/24', 'VLAN' with '123', 'Gateway' with '10.11.1', 'MTU' with '9000', and 'LAG' with 'LAG1'. At the bottom right, there are two buttons: 'Cancel' and 'Create'.

Figure 4. Create Subnet form

2. You must create a virtual interface on the backup subnet. (Figure 5).
 - Select Settings > Network.
 - In the Subnets list, find the subnet with the correct network prefix, VLAN ID, and gateway. The data vip will be attached to this subnet for file export purposes. Create a subnet if none of the existing ones meet your requirements.
 - Click the Add interface button (+) belonging to the subnet to which the data vip will be attached. The Create Network Interface pop-up window appears.
 - In the Name field, type the name of the data vip.
 - In the Address field, type the IP address to be associated with the data vip.
 - In the Services field, leave the service type as data.
 - In the Subnet field, leave the subnet name as the default.
 - Click Create.

Create Network Interface

Name

Backup1

Address

10.1.1.2

Services

data

Subnet

Backup

Cancel

Create

Figure 5. Create Network Interface form

System

Network

Users

Security

Subnets

	Name	Enabled	Prefix	VLAN	Gateway	MTU	LAG	Interfaces	Addresses	Services	
<input checked="" type="checkbox"/>	Backup	True	10.1.0/24	123	10.1.1	9000	uplink	Backup1	10.1.2	data	<input checked="" type="checkbox"/>

+ Add Interface

Figure 6. Subnets list

3. If you plan to use FlashBlade to replicate Commvault DR backups or data, you must also configure a replication interface. This can be on the same subnet you just created or a separate subnet. Follow steps 2 and/or 3 as appropriate to create the interface. In step 3, when setting the Services option, select **replication** instead of **data**. (Figure 7)

Create Network Interface

Name

Replication1

Address

10.2.2.2

Services

replication

Subnet

mgmt

Cancel

Create

Figure 7. Creating a replication interface

4. Create forward and reverse lookup records in DNS for the data IP address you assigned.

Add FlashBlade Replication Links

To allow the primary site FlashBlade to replicate to the DR site, you must connect the arrays. To connect two arrays, perform the following:

1. From the **Storage > Array** page, click the add (+) button in the **FlashBlade Array Connections** panel. The **Connect FlashBlade Array** pop-up window appears.
2. Enter the target array's hostname or management address (unless using NAT) in the **Management Address** field. The address can be located from the **Subnets** table by navigating to the **Settings > Network** page.
5. On the target array, create a connection key. Connection keys are created from the FlashBlade Array Connections panel of the Array page. To create a connection key, perform the following:
 - On the target array, navigate to the Storage > Array page.
 - In the FlashBlade Array Connections pane, click More Options > Create Connection Key. The Connection Key pop-up window appears displaying the new connection key.
 - Copy the new connection key.
 - Once created, the key is active for two hours. If the source and target arrays have not been connected within that two hour period, a new connection key must be created.
3. Enter the connection key for the target array in the **Connection Key** field.
4. The replication address is auto-discovered unless using NAT. The source and target arrays each need a replication network interface. If not already created, create the replication network interfaces on the source and target arrays. Enter the target array's replication network address in the **Replication Address** field.
5. (Optional) Enable encryption by setting the **Encrypted** toggle to on. If set, the `_default_replication_certs` CA certificate group is applied. Note that encryption is set on the source array only. If enabled, the encryption setting displays as enabled on the target array.
6. Click **Connect**.
7. If using NAT, enter the source array's replication address on the target array.
 - On the target array, in the **FlashBlade Array Connections** panel on the **Storage > Array** page, click the edit button at the end of the row displaying the connected source array. The **Edit Connected Array** pop-up window appears.
 - Enter the replication address for the source array in the **Replication Address** field and click **Save**.

Repeat this procedure, using the DR site FlashBlade as the source array and the primary site FlashBlade as the target. You will now have replication capability in both directions.

Create Protection Policy for Replication

To use replication in both directions, you must create a protection policy on each FlashBlade. The policy manages scheduling and retention of snapshots for replication. To create a protection policy:

1. From the **Protection > Policies** page, click the **Add (+)** button in the heading of the **Policies** list. The **Create Policy** pop-up window appears.
2. In the **Name** field enter the name of the snapshot policy.
3. By default, the policy is set to enabled (blue). If you wish to disable the policy, set the **Enabled** button to disabled (gray).
4. Click the expand button next to **Create rule for policy** to add rules.
5. In the **Create or replicate 1 snapshot every** field, enter the frequency at which snapshots are created or replicated. The value must be entered in the format `n{mlhldlw}`. For example, 15m, 3h, 2d, 1w, etc. For DR backup replication, the frequency should be 12h or less.
6. In the **At** field, enter the time of day the snapshot is created. The value must be entered in the format `n[am|pm]`, where `n` is a value of 1-12. If entered without `am` or `pm`, `n` must be a value of 0-23. This field is only editable if the value specified in the **Create or replicate 1 snapshot every** field is in days. For example, 24h, 48h, 72h, 1d, 2d, 3d, 1w, etc.
7. In the **And keep for** field, enter the retention period for the snapshot. The value must be entered in the format `n{mlhldlw}`. For example, 30m, 1h, 2d, 1w, etc. The retention period cannot be less than the snapshot interval. The retention for DR backups should be 5d or greater.
8. Click Create.

Provision File System for Commvault DR Backups

FlashBlade provides a simple way to ensure availability of the DR backups Commvault DR backups.

1. Create three accounts in Active Directory.
 - User account for FlashBlade SMB to bind to Active Directory
 - User account for writing Commvault DR backups
 - Domain group to contain CommServe computer accounts
 - Set a value in the `uidNumber` and `gidNumber` attributes on both users and `gidNumber` attribute on the group
2. Configure directory services for SMB on the FlashBlade.
 - Select Settings > Users.
 - In the Directory Service panel, select SMB.
 - Click the Edit icon to the right of Configuration. The Edit Directory Service Configuration pop-up window appears. (Figure 8)
 - In the URIs field, type the comma-separated list of up to 30 URIs of the directory servers.
For file sharing over SMB, the base DN of the directory service is used in place of the URI to represent the LDAP URL.

Each URI must include the scheme `ldap://` or `ldaps://` (for LDAP over SSL), a hostname, and a domain name or IP address. For example, `ldap://ad.company.com` configures the directory service with the hostname "ad" in the domain "company.com" while specifying the unencrypted LDAP protocol.

We highly recommend either configuring StartTLS by enabling a certificate or certificate group, or configuring URIs by using the `ldaps://` scheme to use LDAP over SSL, unless there is no need for secure communication in your environment.

For SMB, only one domain controller (DC) is supported and its preferences cannot be set when configuring the URI. If specifying a domain name, it should be resolvable by the configured DNS servers.

If specifying an IP address, for IPv4, specify the IP address in the form `ddd.ddd.ddd.ddd`, where `ddd` is a number ranging from 0 to 255 representing a group of 8 bits.

For IPv6, specify the IP address in the form `[xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx]`, where `xxxx` is a hexadecimal number representing a group of 16 bits. Enclose the entire address in square brackets (`[]`). Consecutive fields of zeros can be shortened by replacing the zeros with a double colon (`::`).

If the base DN is not configured and a URI is provided, the base DN will automatically default to the domain components of the URIs.

Optionally specify a port. Append the port number after the end of the entire address. Default ports are 389 for `ldap`, and 636 for `ldaps`. Non-standard ports can be specified in the URI if they are in use.

- In the Base DN field, type the base distinguished name (DN) of the directory service. The Base DN is built from the domain and should consist of only domain components (DCs). For example, for `ldap://ad.storage.company.com`, the Base DN would be: `DC=storage,DC=company,DC=com`.
- In the Bind User field, type the username used to bind to and query the directory.
For OpenLDAP and Active Directory servers, you can use the full DN of the user account that is used to perform lookups. For example, `CN=John,OU=Users,DC=example,DC=com`.
For Active Directory servers, you may instead choose to enter the username - often referred to as the `sAMAccountName` or user login name - of the account that is used to perform directory lookups. The username cannot contain the characters "`[] : ; | = + * ? < > / \`", and cannot exceed 104 characters in length.
- In the Bind Password field, type the password for the bind user account.
- In the Join OU field, enter the relative DN of the organizational unit (OU) within your domain where the system machine account should be created when joining the domain for SMB. For example, `OU=Arrays,OU=Storage,OU=ServiceMachines`.
- Click Save.

Edit Directory Service Configuration

Service Name

smb

Enabled

☐

URIs

ldap://lab.storereduce.com

Base DN

DC=lab,DC=storereduce,DC=com

Bind User

fblookup@lab.storereduce.com

Bind Password

....

Join OU

OU=FB

Test

Cancel

Save

Figure 8. Configure SMB directory service

3. After you configure the directory service settings, test the directory service configuration to verify that the URI can be resolved and that the directory service can successfully bind and query the tree using the bind user credentials.

To test the directory service configuration:

- Select Settings > Users.

– In the Directory Service panel, select the directory service you wish to test.

– Click Test. The Test SMB Configuration pop-up window appears, displaying the output of the test. During the directory service test, Purity//FB tests the directory service configuration to verify that the URI can be resolved and that the directory service can successfully bind and query the tree using the bind user credentials.

Test SMB Configuration

fm1: ldap://lab.storereduce.com

Testing connection ldap://lab.storereduce.com

Binding to ldap://lab.storereduce.com. Type: Windows AD 2016

Searching ldap://lab.storereduce.com

Searching for base_dn

Searching for join_ou

Checking AD Domain Joining. Mode: ad-rc2307

SMB is currently disabled at the system level.

fm2: ldap://lab.storereduce.com

Testing connection ldap://lab.storereduce.com

Binding to ldap://lab.storereduce.com. Type: Windows AD 2016

Searching ldap://lab.storereduce.com

Searching for base_dn

Searching for join_ou

Checking AD Domain Joining. Mode: ad-rc2307

SMB is currently disabled at the system level.

Figure 9. Test SMB Configuration

4. Once the test passes, enable the directory service. (Figure 10)
 - Select Settings > Users.
 - In the Directory Service panel, select the directory service you wish to enable and click the Edit icon.
 - Set the Enabled toggle button to enable (blue) the directory service.
 - Click Save.

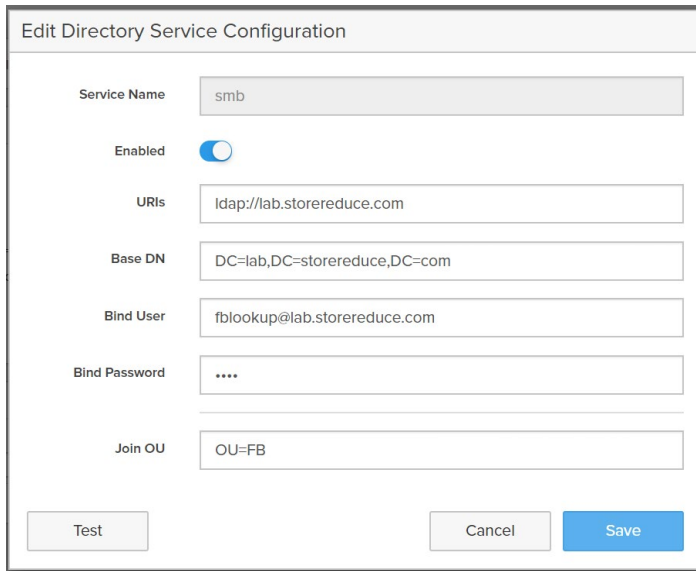
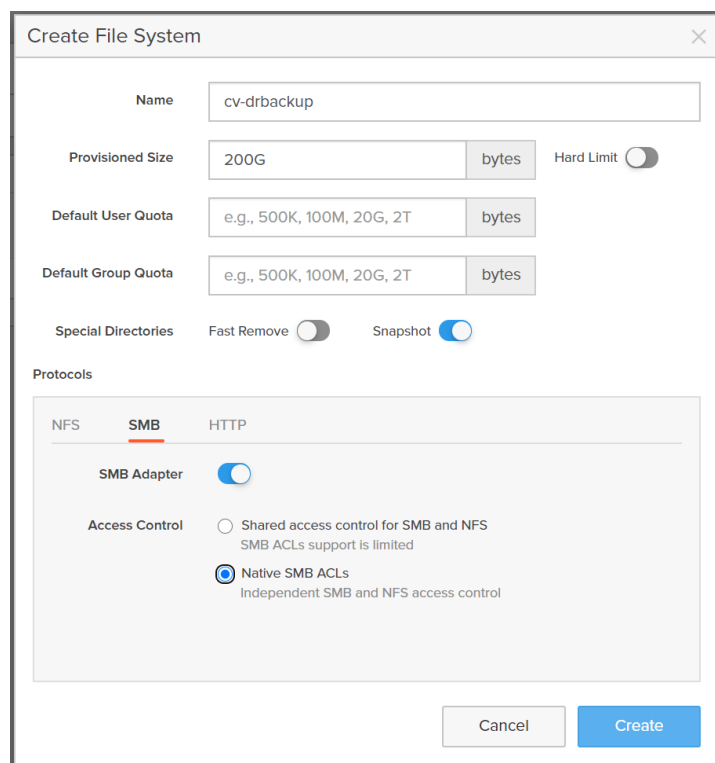


Figure 10 shows the 'Edit Directory Service Configuration' dialog box. The 'Service Name' is set to 'smb'. The 'Enabled' toggle is turned on (blue). The 'URIs' field contains 'ldap://lab.storereduce.com'. The 'Base DN' field contains 'DC=lab,DC=storereduce,DC=com'. The 'Bind User' field contains 'fblookup@lab.storereduce.com'. The 'Bind Password' field is masked with '....'. The 'Join OU' field contains 'OU=FB'. At the bottom, there are 'Test', 'Cancel', and 'Save' buttons.

Figure 10. Enable SMB directory service

5. Create a file system to store the Commvault DR backups. (Figure 11)
 - Create the file system.
 - From the Storage > File Systems page, click the add (+) button in the heading of the File Systems list. The Create File System pop-up window appears.
 - In the Name field, type the name of the directory to be exported.
 - In the Provisioned Size field, specify the provisioned size allocated to the file system. The size is a quota of space that helps gauge the fullness of the file system. If left blank, the provisioned size will default to an unlimited size. You should set the provisioned size to 100G or greater, although DR backups should consume significantly less.
 - Click SMB in the Protocols section.
Click the SMB Adapter toggle button to enable (blue) the SMB protocol adapter.
From the Access Control section, select Native SMB ACLs.
 - Click Create.



Create File System

Name: cv-drbackup

Provisioned Size: 200G bytes Hard Limit: ☐

Default User Quota: e.g., 500K, 100M, 20G, 2T bytes

Default Group Quota: e.g., 500K, 100M, 20G, 2T bytes

Special Directories: Fast Remove ☐ Snapshot ☒

Protocols: NFS SMB HTTP

SMB Adapter: ☒

Access Control:

☐ Shared access control for SMB and NFS

SMB ACLs support is limited

☒ Native SMB ACLs

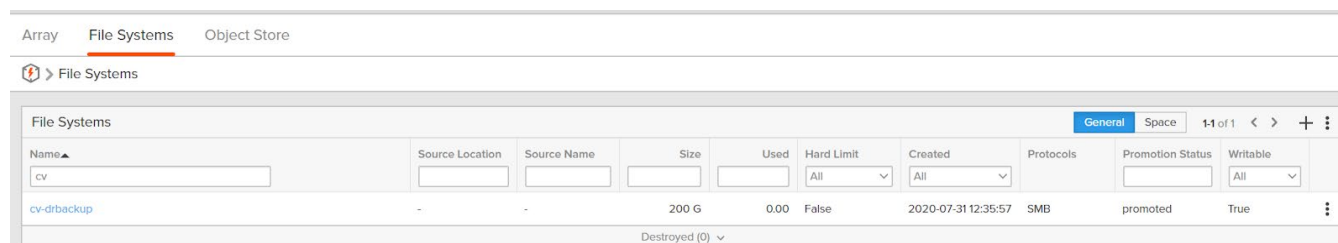
Independent SMB and NFS access control

Cancel Create

Figure 11. Create File System form

Note: The SMB share will have the same name as the file system.

You will see the new file system in the File Systems pane. (Figure 12)



Array File Systems Object Store

File Systems

Name	Source Location	Source Name	Size	Used	Hard Limit	Created	Protocols	Promotion Status	Writable
cv					All	All			All
cv-drbackup	-	-	200 G	0.00	False	2020-07-31 12:35:57	SMB	promoted	True

Destroyed (0)

Figure 12. File Systems pane

6. Create a file replication link for the file system from the primary site FlashBlade to the DR site FlashBlade.
 - From the **Protection > File Replica Links** page, click the add (+) button in the heading of the **File Replica Links** list. The **Create File Replica Link** pop-up window appears.
 - From the **Local File System** list, select the local (source) file system to be replicated.
 - From the **Remote Connection** list, select the remote (target) array.
 - Do not enter a name for the remote file system to which data from the local file system will be replicated in the **Remote File System** field. The target array will create a file system with the same name as the source.
 - From the **Policy** list, select the replication policy you created.
 - Click **Create**.

- 7. Prepare the file system using Windows File Explorer. You may use other tools if you prefer.
 - Set permissions on the base of the SMB share. (Figure 13)
 - Grant full control permissions to the DR backup user account.
 - Grant read & execute access to the CommServe computers group.
 - Remove all other principals from the ACL.

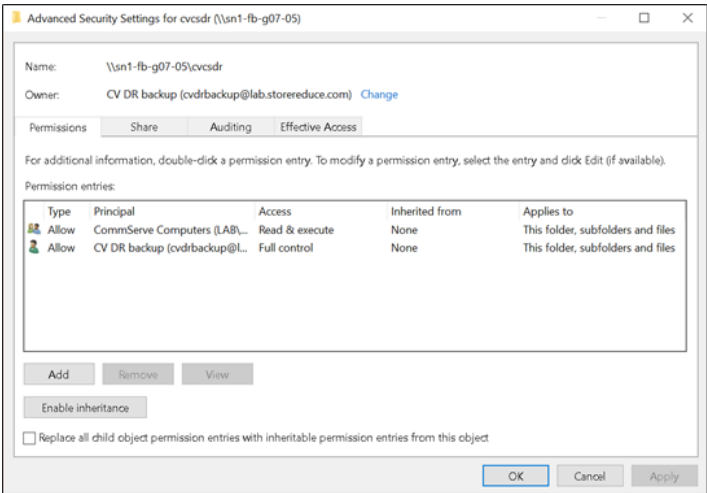


Figure 13. File system ACL for DR backup share

IMPORTANT: Once you commit the ACL, only the DR backup user account will be able to manage data in this file system.

Provision Object Bucket for Backup Data

You must create a bucket for Commvault to write backup data. Buckets are organized into accounts. To create a new account:

1. In the **Accounts** section of the **Storage > Object Store** page, click the add (+) button. The **Create Account** pop-up window appears. (Figure 14)
2. Enter the new account's name in the **Name** field.
3. Click **Create**.




Figure 14. Create Account window

4. You will now see the account listed in the Accounts pane. Click the account name to open its details. (Figure 15)




Figure 15. Accounts pane

Within the account, you must create a user, which will have an access key associated with it. To create a new user:

1. From the **Storage > Object Store** page, click the account to which you wish to add a new user.
2. In the **Users** section, click the add (+) button. The **Create User** pop-up window appears. (Figure 16)
3. Enter the new user's name in the **User Name** field.
4. Do not enable the **Create Access Key** option. You will create a key in a later step.
5. Click **Create**.

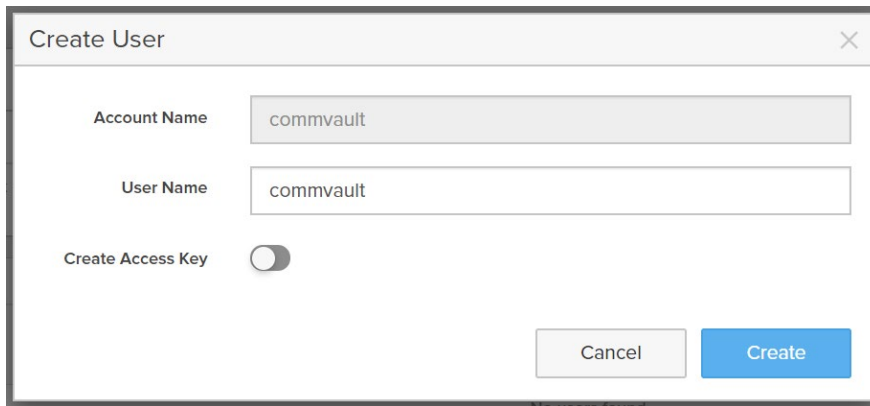


Figure 16. Create User window

Within the same account, you must create a bucket to store the objects Commvault writes. To create a bucket:

1. From the **Storage > Object Store** page, click the account to which you wish to add a bucket.
2. From the **Buckets** pane, click the add (+) button. The **Create Bucket** pop-up window appears. (Figure 17)
3. Enter a name for the bucket.
4. Click **Create**.

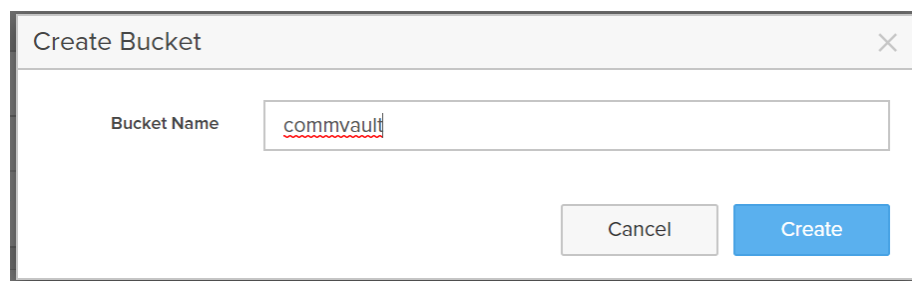


Figure 17. Create Bucket window

You should now see a user and an empty bucket within the object store account you created. (Figure 18)

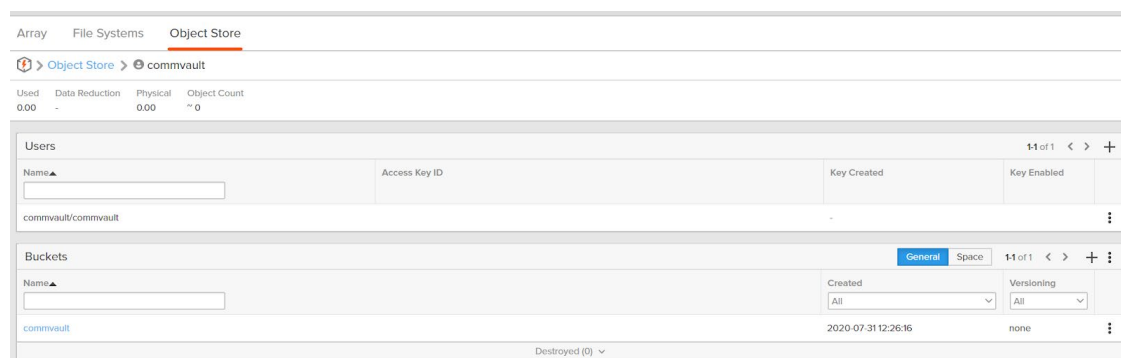


Figure 18. New user and bucket

Repeat this process on the FlashBlade in the DR site.

Optional: Add CA Certificate

If you wish to use TLS to secure the Amazon S3 connection between Commvault clients and the FlashBlade object store, you must update the FlashBlade with a certificate generated by a trusted certification authority (CA). To import a CA certificate:

1. Select Settings > System.
2. Click the **More Options** button from the **SSL Certificate** panel.
3. Click Import Certificate.
4. Complete or modify the following fields:
 - **Certificate** - Click **Choose File** and select the signed certificate. Verify the certificate is PEM formatted (Base64 encoded), and includes the " -----BEGIN CERTIFICATE----- " and " -----END CERTIFICATE----- " lines.
 - **Private Key** - Click **Choose File** and select the private key.
 - **Intermediate Certificate** - (Optional) Click **Choose File** and select the intermediate certificate.
 - **Key Passphrase** - (Optional) If the private key is encrypted with a passphrase, enter the passphrase.
5. Click **Import**. The page will refresh in several seconds.

Set up Commvault RO1105 Appliance

Initial Configuration

Commvault documentation describes in detail the [process to deploy the RO1105](#) in your environment. As you go through the setup procedures, follow the instructions for [Installing the CommServe and MediaAgent Software on Commvault Remote Office Appliance RO1100](#). When complete, you will have the CommServe and MediaAgent packages installed on the RO1105 to support its role as data control.

Guided Setup

1. Complete Core setup:
 - Go to the Command Center URL: `http://webhost/adminconsole`.
Note: webhost is the host name assigned to the RO1105 during CommServe and MediaAgent installation.
 - Enter Commvault administrator user name and password.
 - Click **Login**.
 - From the navigation pane, go to **Guided setup**.
The initial application setup page appears.
 - Click **Let's get started**.
The **Core Setup** wizard appears.
2. On the **Add storage pool** tab of the wizard, add the disk library created during RO1105 configuration to storage pool settings. Select the 960GB SSD volume for DDB and then click **Save**.
3. Create an access key for the object user you created on FlashBlade.
 - Switch to the Purity//FB GUI.
 - From the Storage > Object Store page, click the account to which you recently added a new user.
 - From the Users section, click the More Options > Create access key button on the same row as the user for which you wish to create an access key. (Figure 19)




Figure 19. More Options menu

- A confirmation popup will display. (Figure 20)

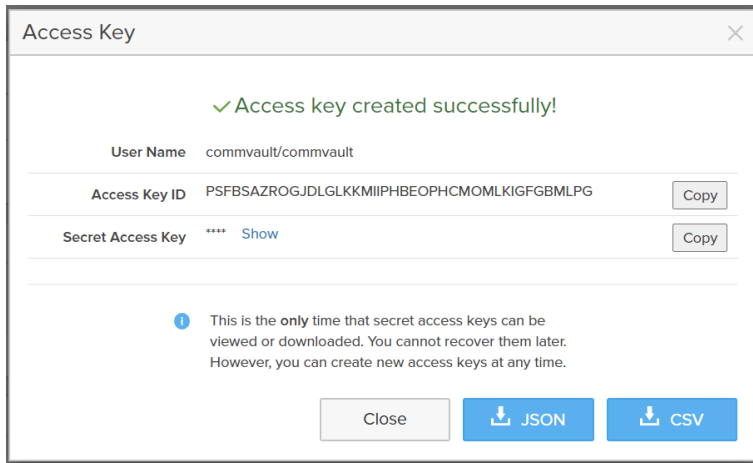


Figure 20. Access Key confirmation popup

- Leave the Access Key confirmation open while adding storage in Commvault Command Center.

4. Create the cloud library in Command Center

- From the Command Center navigation pane, click **Storage > Cloud**. The **Cloud** page appears.
- In the upper right of the page, click **Add** to add a cloud storage. The **Add cloud** dialog box appears. (Figure 21)
- In the **Name** field, enter a display name for the cloud storage target. Click **Add** to create a new cloud storage
- In the **Type** field, select “S3 Compatible Storage.”
- In the **MediaAgent** field, select the RO1105 client name.
- In the **Server host** field, enter the DNS name or vip for Commvault to use to access FlashBlade.
 Note: By default, Commvault will use TLS to secure the Amazon S3 connection. If the FlashBlade does not have a trusted certificate applied that matches the provided name, TLS will fail. To work around this, include “http://” in the **Server host** field. This will prevent Commvault from attempting TLS.
- In the **Bucket** field, enter the name of the target bucket under the configured user account.
- In the **Credentials** field, click the **Create new (+)** button.
 - In the **Credential name** field, enter a descriptive display name for the new stored credential.
 - Copy the access key ID from the Purity//FB GUI, using the **Copy** button, and paste it into the **Access key ID** field.
 - Copy the secret access key from the Purity//FB GUI, using the **Copy** button, and paste it into the **Secret access key** field.
 - Click the **Save** button when all fields are configured.
 Note: To reduce the risk of key compromise, copy the key values directly from the Purity//FB GUI and avoid saving them in a file. Close the **Access Key** popup in the Purity//FB GUI once you have successfully created the cloud library.
- In the **Deduplication DB location** field, select the 960GB SSD volume for DDB.
- Click the **Save** button when all fields are configured. Commvault will connect to FlashBlade and validate the keys and bucket name. Once validated it will create the cloud storage.

Add cloud

Name: FlashBlade S3

Type: S3 Compatible Storage

MediaAgent: cvltma1

Server host: http://10.21.242.110

Bucket: cvlt-s3

Credentials

☒ Use saved credentials

Name: FlashBlade S3 - 10.21.242.110

☒ Use deduplication

Deduplication DB location: /ddb/ddb-fb

Cancel Save

Figure 21. Add cloud form

5. On the Create server backup plan tab, modify the plan according to your requirements, and then click Save.
6. Download latest Feature Release and hotfixes:
 - From **Guided setup** select **Download/copy software**
 - Ensure select **Upgrade to Latest Release**
 - If additional operating systems are required, select those from **Operating system options**
 - Click **Download**
7. As part of the Command Center configuration, a basic Server Plan will be created using the local RO1105 disk library. Subsequent plans should all use the cloud library created above as the Backup Destination. Refer to Commvault documentation for plan options:
 - [File servers](#)
 - [Configure Hypervisor](#)
 - [Databases](#)
 - [Applications](#)

8. Once you finish setup, CommServe DR backups will be stored on local storage on the RO1105 appliance.
You will need to configure Commvault to leverage the replicated file system you set up earlier on the FlashBlade.
 - Add the Active Directory computer account for the RO1105 appliance into the CommServe Computers group you created earlier. You will need to reboot the appliance after changing the group membership.
 - From the Command Center navigation pane, click **Manage > System**. The **System** page appears.
 - Click the **Maintenance** tile to open the **Maintenance** page.
 - Click the **DR backup (Daily)** tile to load the configuration details. Click the **Edit** button (gear icon) to access the **DR backup (Daily)** form.
 - Edit the settings as follows (Figure 22):
 - For the **Backup metadata destination** option, select **Network share**.
 - In the **Path for back up metadata** field, enter the UNC path to the file system you created.
 - Enter the credentials for the DR backup account you created in Active Directory in the **User name**, **Password**, and **Confirm password** fields.
 - If desired, enable the **Upload backup metadata to Commvault cloud** option.
 - Enable the **Upload backup metadata to cloud library** option.
 - In the **Cloud library** dropdown, select the cloud storage target you created on FlashBlade.
 - Click the **Save** button to commit the changes.

Figure 22. DR backup (Daily) form

Enable Virtualization Solution

The guided setup for virtualization creates a hypervisor configuration and enables the virtualization solution to protect VMware virtual machines. A VMware hypervisor can be a vCenter server or a standalone ESXi host. The required Virtual Server Agent (VSA) is installed by default on the RO1105 as well as the Data Servers.

1. If the setup page is not displayed, from the navigation pane, click **Guided setup**.
2. After you complete the core setup, on the **Protect** tab, click the **Virtualization** tile.
3. On the **Create server backup plan** page, type a name for the plan, then provide information about storage, retention, and backup schedules.
If you configured a server backup plan as part of the Core Setup, the wizard skips this page.
4. Click **Save**.
The **Add hypervisor** page appears.
5. Provide the required information for the VMware hypervisor:
 - Select vendor: Select VMware vCenter.
 - vCenter server name: Enter a fully qualified hostname or IP address for the hypervisor.
 - Hypervisor display name: Type a descriptive name for the hypervisor.
To provide access to the hypervisor, enter credentials that provide administrative access to the hypervisor:
 - **User name**: Enter the user name for the vCenter user.
 - **Password**: Enter the password for the vCenter user.
 - **Access nodes**: To identify access nodes (VSA proxies) that can manage backups and restores for the VMware hypervisor, select one or more previously deployed access nodes, and then click **OK**.
6. Click **Save**.
7. On the **Add VM group** page, type a descriptive name to identify the VM group, and then select virtual machines to be protected.
8. Click **Save**.
9. To finish, choose one of the following options:
 - Click **Back up Now** to perform an immediate backup of the virtual machines in the VM group (without requiring confirmation). The **Job details** page appears and displays job status information.
 - Click **Do it later** to go to the hypervisor page without performing a backup.

Additional hypervisor configuration options are available at:

<https://documentation.commvault.com/commvault/v11/article?p=86646.htm>

Deploy Commvault File System Clients

1. From the navigation pane, go to **Manage > Servers**.
The **Servers** page appears.
2. In the upper-right, click **Add server**, and then click **Add file server**.
The **Add server** dialog box appears.
3. Select one of the following installation methods:
 - To install software on the computer, click **Select this method to install software packages on your computer**, and then enter the following information:
 - a) In the **Host name** box, type the host name.
 - b) In the **User name** and **Password** boxes, type the credentials for the server.
 - c) Next to **OS Type**, select the operating system that is installed on the server.
 - d) From the **Select package(s)** list, click each software package to install on the server, and then click **OK**.
Some packages require additional information. For example, if you select DB2, you must enter a DB2 log path.
 - e) **Optional:** In the **Installation location** box, enter the installation location path.
 - f) To reboot the servers in the server group after the installation, move the **Reboot if required** toggle key to the right.
4. Click **Install**.

IMPORTANT: You must also [enable Cloud Accelerator on all clients](#).

Create Subclients

Subclients contain information about what data is backed up. You can create user-defined subclients to manage and back up specific data files and folders. For example, you can create user-defined subclients for folders and files that change frequently or that require a different backup schedule.

If you define overlapping content in multiple subclients within the same backup set, then only one subclient backs up that content. Other subclients do not back up the content.

The subclient that contains the parent folder backs up all the files and folders under the parent folder. However, if any of the child folders are specified as content in another subclient, then the child folder is automatically filtered when the backup runs.

1. From the navigation pane, go to **Protect > File servers**.
The **File servers** page appears.
2. Click the file server.
The file server properties page appears.
3. Under **Subclients**, on the right of the page, click **Add subclient**.
The **Create new subclient** page appears.
4. Enter the following information:

- **Name:** Enter a name for the subclient.
- **Backup set:** Select the backup set that you want to add the subclient to.
- **Backup plan:** Select the backup plan to use for the subclient, and then specify which content you want to back up.

5. Click **OK**

CommServe Availability

Commvault has several options for ensuring availability of CommServe services in the event of a site or server loss. Consult with your Commvault or partner account team to determine which option is best for you.

General Best Practices

The best write performance on FlashBlade storage is achieved by spreading the load across as many blades as possible. Configuring Commvault for maximum writers and distributed streams will give the best outcome with the simplest setup. For simple configuration and scaling, the Commvault agents leverage the Amazon S3 protocol to communicate with FlashBlade. Specific best practices for the solution configuration are:

- Configure a single FlashBlade bucket and mount path.
- Share mount paths between Data Servers.
- Deploy Cloud Accelerator on all clients.
- Use Commvault client-side compression by default.
- Use Commvault client-side deduplication.
- Set maximum writers on the library and Data Servers.
- Configure multiple data paths and round robin in storage policies.
- Match VMware disk format and transport mode.
- Disable TLS, if allowed.

Some best practices include [Commvault Additional Settings](#).

Configure a Single FlashBlade Bucket and Mount Path

Commvault requires only a single bucket. Any cloud library can create mount paths under the same bucket, simplifying the configuration. Unless required for multitenant environments or other business reasons, Commvault should be configured to use a single bucket, in a single cloud library, for each FlashBlade, with a single mount path.

Share Mount Paths with Data Servers

Sharing mount paths with data servers enables management load distribution without having to configure multiple storage targets. This does not affect client data paths, since clients will communicate directly with FlashBlade. To share a mount path:

1. In Commvault Command Center, navigate to the storage target.
2. In the storage target configuration, click the **Actions** button for the bucket, then click **Add MediaAgent**, as shown in Figure 23.

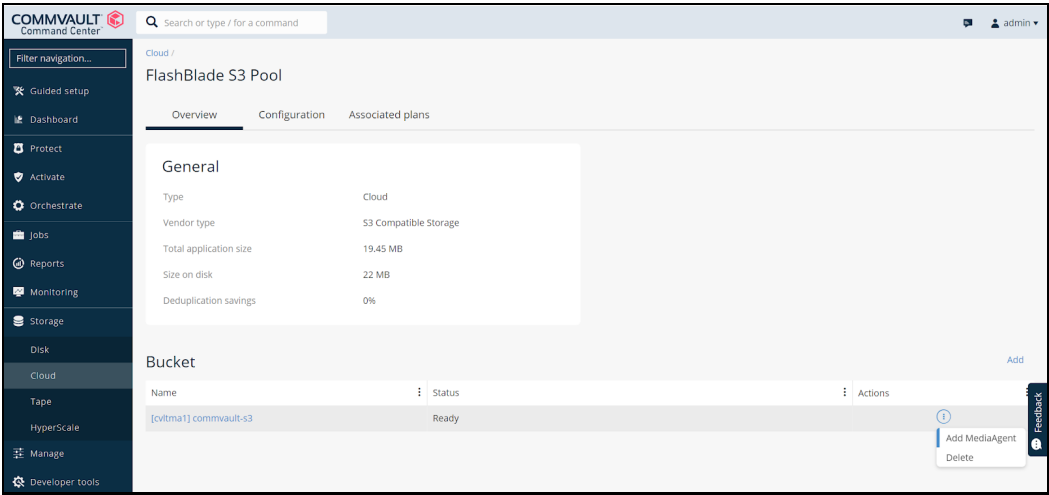


Figure 23. Cloud buckets

3. As shown in Figure 24, select the Data Servers that should manage access to the mount path, then click the **Save** button. Data Servers will be added with read/write access.

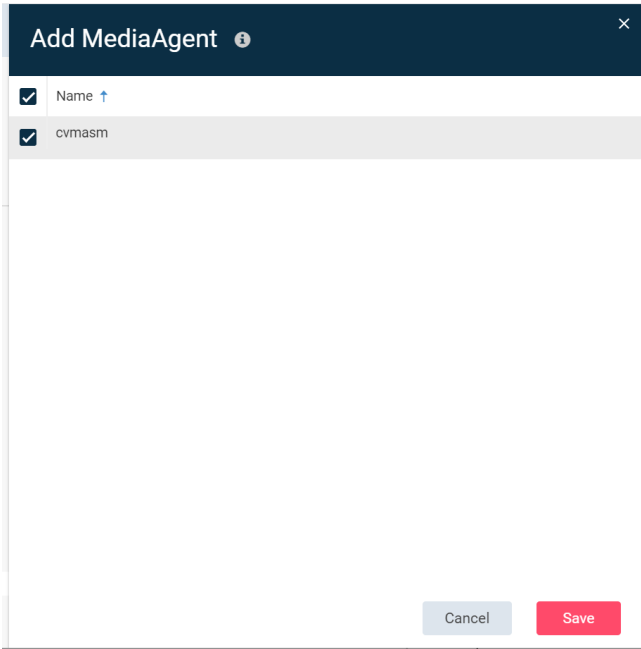


Figure 24. Add MediaAgent form

Deploy Cloud Accelerator on All Clients

The Cloud Accelerator feature lets clients write directly to object storage without having the MediaAgent package installed and without having to share mount paths to them in the storage target configuration. The MediaAgent controlling jobs can become a Data Server and manage more streams with fewer resources. Because clients are accessing the FlashBlade without a MediaAgent, network restrictions are critically important. If a security policy does not allow endpoints to directly access backup storage, you must use a traditional consolidated MediaAgent architecture.

To enable Cloud Accelerator, install the Storage Accelerator package.

1. In the CommCell Console, navigate to **Client Computers**. As shown in Figure 25, right-click the client, then select **All Tasks>Add/Remove Software>Install Software**.

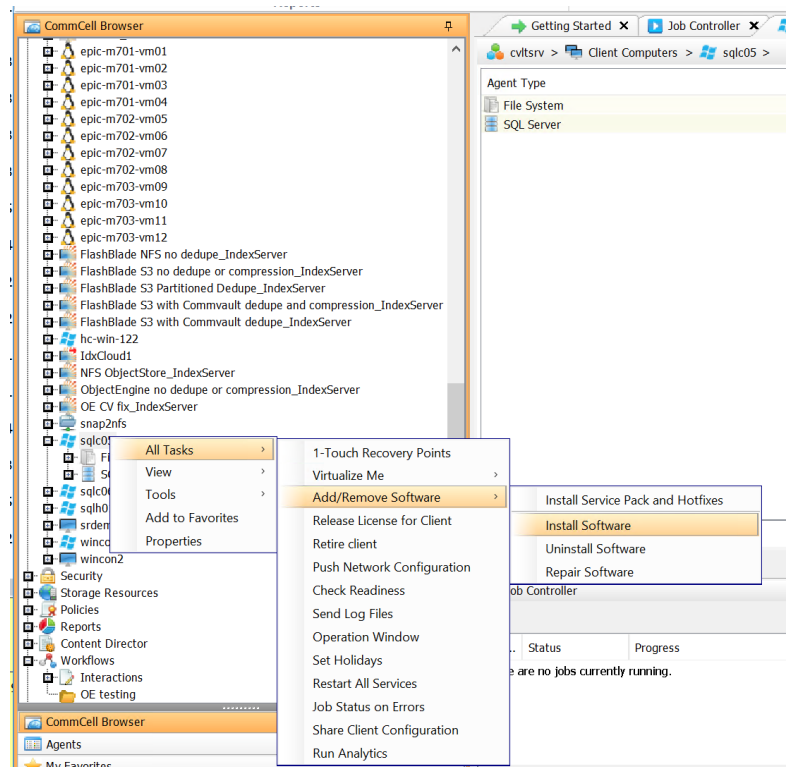


Figure 25. Install Software option

- Follow the wizard, accepting default options. As shown in Figure 26, on the **Select Packages to Install** screen, select the **Storage Accelerator** option under the **Tools** section. Click **Finish** at the end of the wizard to begin the installation. Once the installation completes, Cloud Accelerator will be enabled automatically.

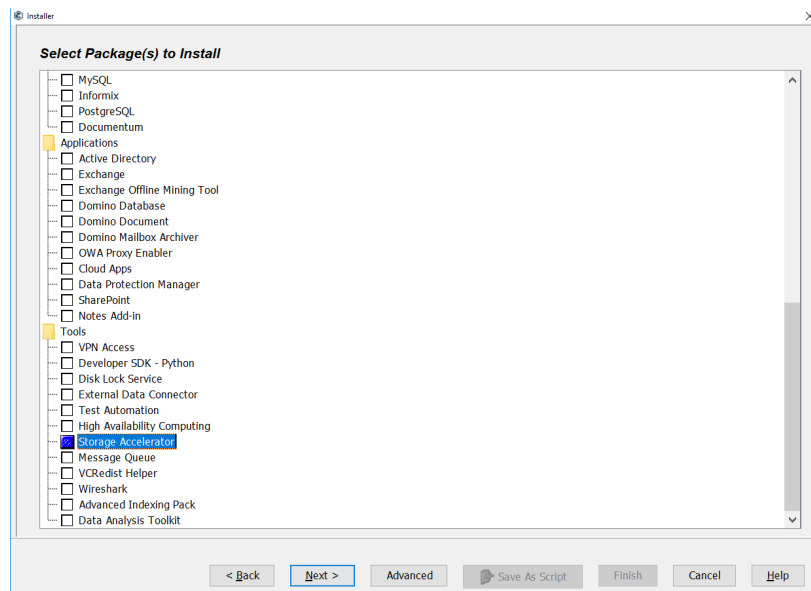


Figure 26. Software package selection

Use Commvault Client-Side Compression by Default

While FlashBlade has effective hardware compression, most clients do not have enough network bandwidth available to offset the data reduction from client-side compression. Commvault's deduplication algorithm can also reduce the effectiveness of FlashBlade compression. With Commvault compression enabled, backups are usually faster and consume a comparable amount of physical storage. If backups are underperforming, compression can be disabled, but generally it should be left in the default enabled state.

Use Commvault Client-Side Deduplication

Commvault deduplication provides data reduction across large data volumes, for improved storage efficiency. Deduplication at the client side will reduce the amount of data sent over the network to FlashBlade. Note that because most data is removed at the client, only initial full backups will send large amounts of data. Subsequent backups will be bound by how fast the client can remove data.

Match VMware Disk Format and Transport Mode

VMware virtual disks allow several formats. When using SAN transport mode, restore performance is best with thick-provisioned eager zero format. Lab testing has shown restore throughput with eager zero format at more than double that with thick provisioned lazy zero format. Thin-provisioned format is the slowest with SAN mode, according to both VMware and Commvault. For thick-provisioned, lazy zero, or thin-provisioned virtual disks, use HotAdd or NBD transport for best throughput.

Disable TLS, if Allowed

Enabling TLS can lower maximum write throughput for a data stream by up to 30%. It should be disabled unless required by your security policy. The [guided setup process](#) explains how to disable TLS when you add FlashBlade to Commvault.

Data Servers

As your backup footprint grows, you may grow beyond the limits of the RO1100. To manage more data, you need to add Data Servers. These 1U servers expand the management layer and allow you to scale to hundreds of TB of storage. Data Servers are an important factor in optimal backup and recovery performance, as they store all metadata, manage deduplication, and retain individual job and logging information. Important factors to consider are:

- Operating system
- Data Server hardware specifications
- Data Server count
- Deduplication database storage
- Index cache storage

Note: Data Servers are managed as MediaAgents in the Commvault consoles.

Operating System

Commvault supports Windows and Linux on Data Servers. Functionality, performance, and configuration are similar across operating systems, so you should choose the option that works best for your environment. Data Server and client operating systems do not have to match. For example, you can use a Linux Data Server to back up a Windows client.

Virtual Server Agent Mode

When using Data Servers, you can run VSAs in one of two transport modes. You can use VSAs on the Data Servers themselves, using VADP SAN transport mode or NBD transport mode, or you can deploy separate virtual machines and use VADP HotAdd transport mode. The deployment model affects the amount of data and number of VMs you can manage with the data server.

Data Server Hardware Specifications

You should choose your Data Server specification based on the VSA model and the amount of storage you expect you will need to manage.

Data Server selection and sizing is based on two factors:

1. VSA transport mode - SAN or NBD on the Data Server (referred to below as “Converged”) or HotAdd on the ESXi server and;
2. FlashBlade storage capacity.

Table 1 below provides the Data Server hardware requirements for both Converged and HotAdd VSA transport modes based on the FlashBlade storage capacity referred to as “Cloud Lib Size.”

VSA Mode	Add'l Cloud Lib Size	Data Server size	CPU Cores	RAM	Dedupe SSD Size	Index Volume Size	Add'l FlashBlade Blades
Converged	75TB	M	8	32G	1TB	1TB	2
Converged	150TB	L	12	64G	2TB	2TB	4
Converged	250TB	XL	16	128G	2TB	2TB	7
Converged	500TB	XL x 2	16	128G	4TB	4TB	15* (Add'l chassis required)
Converged	150TB	M	8	32G	2TB	2TB	4
Converged	300TB	L	12	64G	4TB	4TB	8
Converged	500TB	XL	16	128G	4TB	4TB	15* (Add'l chassis required)

Table 1. Data Server specifications

You can add fibre channel HBAs to your Data Servers to enable IntelliSnap backup copy and VSA SAN transport mode.

Data Server Count

You should try to deploy the fewest required Data Servers for maximum density and efficiency. Deploy based on your current and projected needs. For example, if you want to use VSAs in HotAdd mode, and you need 290TB of FlashBlade storage but are growing 5TB per month, you should deploy a single 300TB configuration, rather than two Data Servers.

New Deployment

You will need to deploy a base operating system on the Data Server prior to Commvault installation. You will need to format the [DDB](#) and [index](#) drives according to Commvault best practices.

Follow the instructions for [Adding a MediaAgent](#) to deploy the Data Server software using Command Center.

Deduplication Database Storage

You should add a DDB partition on the new Data Server to expand the deduplication performance and capacity. See [Configuring Additional Partitions for a Deduplication Database](#) in Commvault documentation for instructions on using the CommCell Console to add a partition on the DDB drive.

Index Storage

The index is used by the Data Server for catalog access during recovery, and for temporary storage in Live Mount and Live Recovery cases. A fast index is critical to performance. The Data Server specification includes an SSD specifically for index storage. However, you must explicitly choose to locate the index there. Once you have deployed the MediaAgent software on the Data Server, you can use Command Center to change the index path.

1. From the Command Center navigation pane, click **Manage > Infrastructure**. The **Infrastructure** page appears.
2. Click the **MediaAgents** tile. The **MediaAgents** list appears.
3. Click the name of the Data Server in the list. The MediaAgent details page appears.
4. In the **Index Cache** tile, click the **Edit** link to open the **Edit Index Cache Properties** form. (Figure 27)
 - In the **Index Directory** field, enter or browse to a path on the index SSD.
 - Click the **Save** button to commit the change.

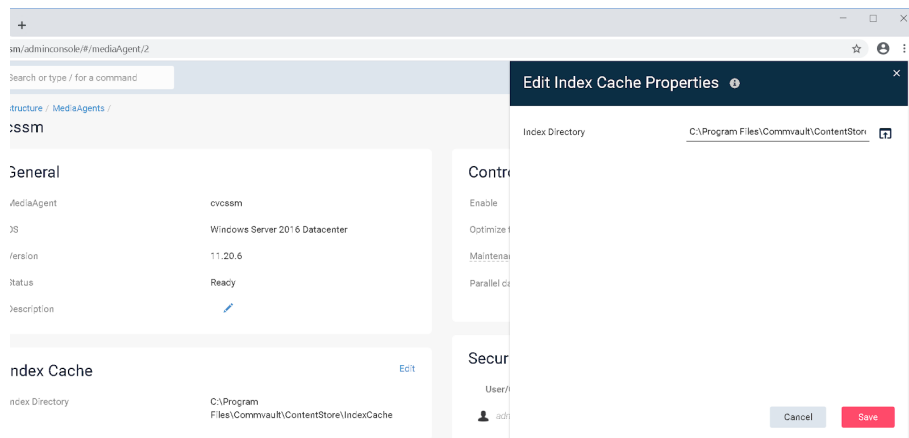


Figure 27. Edit Index Cache Properties form

Advanced Configuration

Amazon S3 is an object-based protocol that Commvault can use with FlashBlade as a cloud storage target. Amazon S3 has the advantages of simplicity and scale compared with NFS. With enough available bandwidth, a single MediaAgent with a single Amazon S3 mount path can reach nearly the maximum FlashBlade write performance. NFS requires multiple mount paths and more complex Commvault configuration to achieve the same result. For best Amazon S3 results:

- Run multiple client readers for large data sets.
- Increase cloud thread pool size.
- Optional: Configure IntelliSnap snapshot integration.

Run Multiple Client Data Readers

As parallel streams are important to getting the best throughput, configuring multiple data readers can improve backup and recovery throughput. With multiple data readers, Commvault can run parallel processes on the client system to pull data from primary storage. Data readers are usually configured on subclients. Specific GUI option names and locations vary by agent type, and optimal values will vary by agent type, client hardware, deduplication database performance, and data profile. Tuning will typically be required. Data readers are configured in the CommCell Console interface. The following are recommended as starting points when the source data is on FlashArray™.

- **SQL Server:** Four backup streams
 - Configured in subclient properties, using the **Number of streams for data backup** field on the **Storage Device** tab.
- **Oracle:** Four backup streams
 - For database backups, configured in subclient properties, using the **Number of Data Backup Streams** field on the **Storage Device** tab.
 - For archive log backups, configured in the instance properties, using the **Number of Archive Log Backup Streams** field on the **Log Backup** tab under the **Storage Device** tab.
- **VMware:** Three data readers x number of VSA proxies
 - Configured in the subclient properties, using the **Number of Data Readers** field on the **Advanced** tab.
- **File system:** Four data readers, enable the **Allow Multiple Readers Within a Drive or Mount Point** option
 - Configured in the subclient advanced properties, using the **Number of Data Readers** field on the **Performance** tab.

Increase Cloud Thread Pool Size

Commvault uses a pool of threads to connect to object storage. By default, a process on the client running Cloud Accelerator can create up to 50 connections to the storage. All jobs on that client will share the connections, and the number of active streams does not directly affect the number of threads in use. In most cases, Commvault does not benefit from increasing the maximum thread count, but in certain resource-limited environments, it can increase backup throughput. If connection counts from a client to the FlashBlade regularly exceed 40, and the network interface is not saturated, increase pool size to 100.

You can measure connection counts using the netstat command on both Windows and Linux clients, by filtering for established TCP connections to the FlashBlade data vip. Refer to the documentation within the operating system for specific command options.

To set the thread pool for a client or MediaAgent, you need to apply the nCloudGlobalUploadThreadPoolMaxCount additional setting. This does not force the system to use a certain number of threads. Instead, it simply allows it to go beyond the default 50. (Get [more information on the setting](#).) Additional settings are managed in the CommCell Console.

The best way to apply the setting is through a client computer group. To create a group to control the thread count:

1. In the CommCell Console, right-click **Client Computer Groups** and select **New Group** from the context menu, as shown in figure 28.

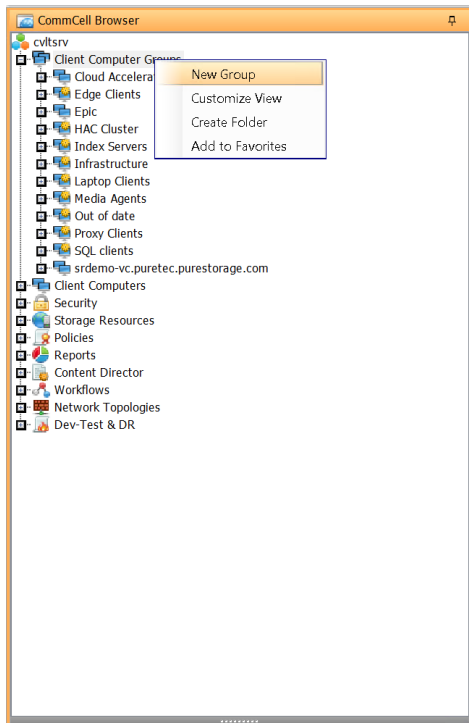


Figure 28. Creating a client computer group

2. As shown in Figure 29, give the group a descriptive name. Select the Manual Association option. Select the desired clients in the left pane, then click the Include button.

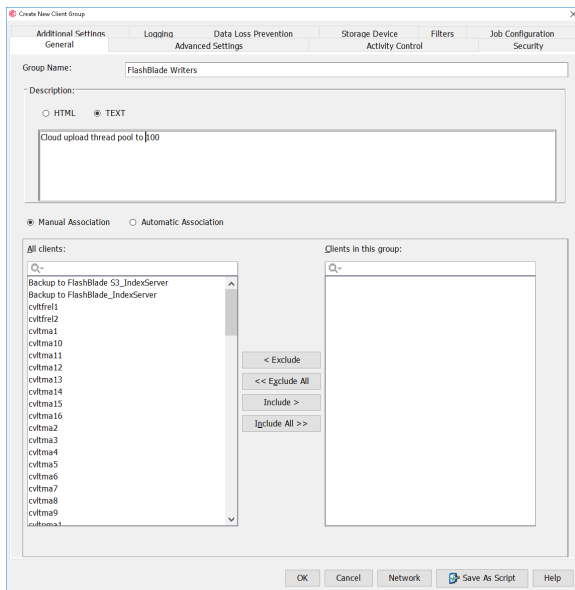


Figure 29. Client computer group general options

3. Select the **Additional Settings** tab. Click the **Add** button. As shown in Figure 30, enter nCloudGlobalUploadThreadPoolMaxCount. When the setting appears in the search list, click it to automatically prepopulate all the fields. Enter the desired maximum number of threads per system, as calculated earlier, in the **Value** field. Enter a text value in the **Description** field. Click the **OK** button to create the group. The setting will automatically apply on the group members and be honored in the next job writing to FlashBlade.

The screenshot shows a dialog box titled "Add Additional Settings". It contains the following fields and controls:

- Name:** A text box containing "nCloudGlobalUploadThreadPoolMaxCount" and a "Lookup" button.
- Category:** A dropdown menu showing "MediaAgent".
- Type:** A dropdown menu showing "INTEGER".
- Value:** A text box containing "100".
- Enable:** A checked checkbox.
- Comment:** A text area containing "Increase available connections to FlashBlade".
- Details:** A section containing a description: "Description: Use this setting to set the total number of the concurrent upload threads globally for all files to the Cloud Servers in a single process. Default is 50." and other metadata: "Type: INTEGER", "Categories: MediaAgent", "Default Value: 50", "Minimum Value: 0", and "Maximum Value: 1000".

At the bottom of the dialog are "OK", "Cancel", and "Help" buttons.

Figure 30. Additional Settings configuration

Configure IntelliSnap Snapshot Integration

You can leverage the RO1105 and Data Servers as backup proxies for IntelliSnap snapshot management. Since IntelliSnap supports a wide variety of primary storage systems, including Pure FlashArray, refer to [Commvault documentation](#) for configuration requirements and instructions.

Additional Resources

Supporting Information

- [Commvault public documentation](#)
- [Best Practices for Configuring Commvault with FlashBlade](#)
- [Ransomware Protection with Pure Storage and Commvault](#)
- [Commvault Remote Office Appliance 1100 Technical Specifications](#)

Common Failure Scenarios

This section describes potential causes for common implementation failures.

Cloud Target Creation Fails

During setup for the cloud storage target, Commvault connects to the bucket and creates a folder. There are several reasons this may fail, shown in Table 2. If the GUI reports an error and no folder was created in the FlashBlade bucket, it may be due to one of these reasons.

Cause	Error Messages	Explanation	Resolution
Unsuccessful TLS certificate validation	Could not connect Failed to check cloud server status, error = [[Cloud] The server failed to do the verification. Error = 44037]	By default, FlashBlade uses a self-signed certificate, and Commvault will fail the TLS handshake	Use a trusted certificate Disable TLS or certificate validation check for the MediaAgent. See step 4 of Guided Setup on page 20 and Optional: Add CA Certificate on page 18.
Incorrect keys provided	Access denied	The wrong access key ID or secret key was provided. Keys with leading or trailing spaces can be interpreted incorrectly and lead to failures.	Confirm correct keys.
Incorrect bucket name	Bucket not found	The bucket name was entered incorrectly. Bucket names with leading or trailing spaces can be interpreted incorrectly and lead to failures.	Confirm and retype the bucket name

Table 2. Cloud target creation failure reasons

Backups or Restores Fail to Send Data

Backup or restore jobs may fail or change to Pending state without writing any data to FlashBlade. This may be due to one of the issues listed in Table 3.

Cause	Error Messages	Explanation	Resolution
Unsuccessful TLS certificate validation	Could not access mount path	Disabling certificate validation must be performed on each MediaAgent. Backups will fail for a shared mount path if a MediaAgent is selected that has validation enabled and the default self-signed certificate is still in use.	Use a trusted certificate Disable TLS or certificate validation check for the MediaAgent See Disable TLS, If Allowed and Configure a Single Bucket and Mount Path
Network error	Could not access mount path	The writer MediaAgent or Cloud Accelerator client is unable to reach the FlashBlade.	Check for firewall between writer and FlashBlade

Table 3. Failure reasons during data transmission

Throughput

With Commvault client-side deduplication enabled, many environments rarely if ever see network saturation. Commvault reports effective throughput, or the rate at which data is processed at the source, rather than actual throughput, or network utilization. Typically, effective throughput is lower than the available network bandwidth on the initial backup for a data set. It tends to be higher than available bandwidth on subsequent backups because data reduction removes data faster than the network could transmit it. However, this varies across environments. Implementing FlashBlade typically has little impact on backup throughput unless the IP network can transmit data faster than the existing storage can write it. Backup throughput should not be the only factor considered when diagnosing a performance issue. If backups to FlashBlade are slower than an existing storage product, with no other changes, there may be an issue.

If restore throughput is lower than 80% of available bandwidth at the slowest link in the data path, which is not always the IP network, there may be an issue. For example, a restore across a 25Gbps network would not be expected to sustain more than 80MiB/s to a SATA disk that can only write 100MiB/s.

If backup or restore throughput is lower than expected, the cause may be one of the issues listed in Table 4. Increase CPU resources on the client or reduce the number of streams.

Cause	Error Messages	Explanation	Resolution
Clients cannot open sufficient connections to FlashBlade	TCP connection count is equal to thread pool maximum	Commvault thread pool is not large enough to maximize throughput.	Increase the thread pool maximum size. This also suggests a CPU and memory resource limitation on the MediaAgent. Consider upgrading the MediaAgent hardware. See Increase Cloud Thread Pool Size
Insufficient client CPU resources	Client CPU runs at 100% during backup or recovery	Saturated client CPU will limit throughput and may slow backups down.	Increase CPU resources on the client or reduce the number of streams.

Table 4. Reasons for lower throughput than expected

Terms and Concepts

Terms	Explanation
Commvault Command Center	Commvault's HTML console rebranded from Admin Console as of SP14 in December 2018. Simple web interface with some advanced configuration capability.
CommCell Console	Commvault's Java-based console. More capable but less simple than Command Center. Procedures in this guide that use Command Center can also be performed in the CommCell Console.
MediaAgent	Commvault data mover and distributed index store. MediaAgents are the primary communicators with back-end storage such as FlashBlade.

Data Server	Commvault role that performs MediaAgent metadata functions but does not directly transfer data to and from agents. Requires clients to run Cloud Accelerator and leverage object storage.
Agent (also iDataAgent)	Commvault client software that interfaces to a specific data type for backup and recovery. Represented in Commvault GUIs as a configuration object.
Subclient	A Commvault configuration object within the agent configuration that defines a data set and the options and policies used to protect it. An agent can have multiple subclients.
<u>Amazon S3</u>	Amazon Simple Storage Service. Industry standard object-based protocol for writing data over HTTP. Generally considered slower than file-based protocols.
<u>SMB</u>	Server Message Block. Industry standard, file-based protocol for writing unstructured data to network storage. Common on Windows, it is not as widely used on other platforms.
HotAdd transport mode	VMware mechanism for backing up virtual machines. Virtual disks are attached to a VM running backup software and read within the VM as local disks.
SAN transport mode	VMware mechanism for backing up virtual machines. Datastore disks on SAN storage are attached to a physical server running backup software, and virtual disks are read from the datastore.
NBD transport mode	VMware mechanism for backing up virtual machines. Backup agents read data directly from the ESXi host.
Cloud Accelerator	A Commvault feature that lets agents write data directly to object storage, bypassing the MediaAgent.
Live VM Recovery	A Commvault feature for instant VM recovery. A VM is powered on from a virtual datastore Commvault presents, then migrated to a permanent datastore.
VM Live Mount	A Commvault feature for DevOps and similar use cases. Similar to Live VM Recovery, a VM is powered on from a virtual datastore Commvault presents; However, it is not migrated, and its configuration and lifecycle are managed by policy.
Mount Path	A mount path is a unique back-end storage target, accessible by one or more MediaAgents. Multiple mount paths can be grouped within a storage target, or library.
Client data reader	A component of the Commvault agent processes that reads data from the client's storage for backup. By default, agents use one to two readers per subclient (data set) and a single reader per local volume.
Client-side deduplication	A Commvault feature to reduce network traffic and improve backup performance. The backup agent communicates with the MediaAgent to determine which data is already known and discards duplicates.
Target-side deduplication	Similar to client-side deduplication, but performed on the MediaAgent. The agent sends the full data without removing duplicates, resulting in more network utilization but lower CPU and memory load on the client. CPU and memory utilization increases on the MediaAgent.
Software compression	Commvault can compress data being backed up to reduce network and/or storage consumption. Files are analyzed to look for repeated patterns that can be consolidated. Compression can be performed at the backup client or the MediaAgent. Client-side compression reduces network utilization but increases CPU load on the client. Compression at the MediaAgent reduces client load but increases network consumption and CPU load on the MediaAgent.

Table 5. Terms

About the Authors



Roy Child is a Senior Solution Architect with Pure Storage, responsible for defining data recovery solutions around various data protection applications. He is responsible for defining Pure Storage solutions and reference architectures for protecting and recovering primary workloads, such as Oracle, SQL, and VMware. Roy has worked in and with the data protection industry for 20 years, from end user to IT architect with multiple backup and recovery products, followed by product management with Commvault. Roy joined Pure Storage in April 2019.



Shawn Smucker is a Director of Product Management with Commvault primarily responsible for Commvault HyperScale Software. In addition to HyperScale Software, he is also responsible for defining reference designs and reference architectures for both primary and secondary data management and storage. Shawn has been with Commvault for over 12 years but involved in data management for nearly 16 years.

©2020 Pure Storage, the Pure P Logo, and the marks on the Pure Trademark List at <https://www.purestorage.com/legal/productenduserinfo.html> are trademarks of Pure Storage, Inc. Other names are trademarks of their respective owners. Use of Pure Storage Products and Programs are covered by End User Agreements, IP, and other terms, available at: <https://www.purestorage.com/legal/productenduserinfo.html> and <https://www.purestorage.com/patents>

The Pure Storage products and programs described in this documentation are distributed under a license agreement restricting the use, copying, distribution, and decompilation/reverse engineering of the products. No part of this documentation may be reproduced in any form by any means without prior written authorization from Pure Storage, Inc. and its licensors, if any. Pure Storage may make improvements and/or changes in the Pure Storage products and/or the programs described in this documentation at any time without notice.

THIS DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. PURE STORAGE SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

Pure Storage, Inc.
650 Castro Street, #400
Mountain View, CA 94041