

REFERENCE ARCHITECTURE

Rapid Restore with Commvault and Pure Storage FlashBlade//S

Reference architecture for VMware environments.

Contents

- Introduction3**
- How to Use This Guide3**
- Pure Storage FlashBlade//S3**
- Commvault Backup & Recovery and Metallic6**
- Rapid Restore in Partnership with Commvault6**
 - Solution Architecture 7
 - Scaling..... 7
 - Additional Sites 8
- Lab Setup9**
 - Server Details 9
 - Storage Details..... 9
- Source Data Details 10**
 - Adding FlashBlade//S as a Commvault Cloud Storage Pool 10
 - Excluding Commvault Process from AV Scanning21
- Tuning Performance22**
- Testing Details22**
- Test Results 24**
 - Backup 24
 - Full VM Recovery 25
- Ransomware Mitigation 28**
 - Configuring SafeMode Protection.....29
 - Recovering with Object SafeMode 31
- Recommended Practices for Commvault with FlashBlade//S 31**
- Conclusion32**
- Additional Resources 33**
- About the Author 34**

Introduction

The persistent threat of ransomware has forever changed how businesses view data protection. While we used to worry most about meeting a backup window and finding the cheapest possible storage, those concerns are now secondary. We need to be able to recover entire environments as fast as possible, and without leaving backup data vulnerable.

According to Sophos' [State of Ransomware 2022](#), the average total cost of remediating a ransomware attack was \$1.4 million; while down from \$1.85 million in 2021, it is still a significant increase from 2020. In the same period, the average ransom demand increased almost five times, to over \$810,000. Only 4% of the companies that paid managed to decrypt all their data while on average victims lost over a third of their data. The bottom line: If you are attacked, expect to restore a lot of your environment from backups. And the faster you can do it, the less recovery will cost you.

Pure Storage® FlashBlade//S® is here to help. FlashBlade//S supports restore speeds of over 200TB per hour in a single chassis. Immutability options protect your data from attackers, so you can start recovery faster. When you integrate FlashBlade//S with Commvault Backup & Recovery, you get a robust, centralized data management platform for all your data that is simple to deploy with FlashBlade. Commvault adds its own ransomware mitigations to help you reduce, recognize, and respond to attacks and minimize the impact to you.

How to Use This Guide

This guide will show you how to deploy a Pure FlashBlade//S with Commvault for your data backup and rapid recovery needs. You will learn about the FlashBlade//S all-flash storage system and Commvault, a recognized leader of enterprise backup and recovery software. When deployed with Commvault software, FlashBlade//S simplifies and accelerates VMware backup and recovery compared to traditional storage solutions.

This reference architecture is for IT, as well as storage and VMware administrators who are using or considering Commvault software and looking to build a rapid restore data protection architecture with FlashBlade//S. It provides guidelines and techniques to improve VMware backup and restore performance with Commvault and FlashBlade//S.

NOTE: Please note that performance data illustrated in this guide represents performance on our existing lab equipment for illustration purposes and should not be viewed as optimal peak performance figures.

Pure Storage FlashBlade//S

Pure Storage FlashBlade//S is the next generation of enterprise scale-out unified fast file and object (UFFO) storage, delivering rich data services with high density, capacity, performance, and scalability to meet the needs of modern applications. Using a distributed metadata architecture, FlashBlade//S offers multi-dimensional performance on a consolidated platform with NFS, SMB, and S3 protocol access. FlashBlade//S provides a single solution for file and object workloads that's easy to set up,



manage, scale, and upgrade. It's a platform designed to deliver cutting-edge performance capabilities in multiple dimensions to uncomplicate your unstructured data storage—forever.

FlashBlade//S is designed to easily scale out to grow with your unstructured data needs for analytics, artificial intelligence (AI) and machine learning (ML), data protection and rapid restore, among other workloads.

Built using DirectFlash® Modules and all-QLC architecture, FlashBlade//S is the ideal foundation for modern workloads. It's a unified fast file and object (UFFO) storage platform that provides rich data services with higher density and capacity than ever before. FlashBlade//S is designed to easily support the most demanding unstructured data workloads, without compromising on system performance or efficiency.

Architectures that use off-the-shelf solid-state drives (SSDs) have an internal controller to manage the flash media on each specific drive without any knowledge of what's happening at the systems level. In contrast, FlashBlade//S uses Pure's innovative DirectFlash Modules, enabling the storage operating system to manage that media on a global level. The DirectFlash Modules include a small amount of NVRAM that scales as the platform grows. Purity//FB, the operating system for FlashBlade//S, manages all system resources including blades and DirectFlash Modules at a global level. Global media management enables FlashBlade//S DirectFlash Modules to unlock as much as 20% more capacity from NAND when compared to competitors using off-the-shelf SSDs. This delivers more consistent performance, better reliability, and higher media endurance without requiring massive storage class memory (SCM) cache.

Scale compute and storage independently: Designed with a unique modular architecture that allows you to easily increase capacity or performance, FlashBlade//S is a customizable platform that gives you the ability to tailor your configuration for specific workload requirements. It provides the flexibility to easily adapt to your data growth projections and evolving storage needs.

Modular architecture: FlashBlade//S has a unique modular architecture that enables organizations to unlock new levels of power, space, and performance efficiency using an all-QLC design. The architecture disaggregates compute resources from storage, so that capacity and compute can scale independently for extreme flexibility in configuration. FlashBlade//S (Figure 1) is a customizable platform that enables you to tailor your configuration for current workload requirements and non-disruptively upgrade to meet future needs. You can upgrade components on a schedule that is consistent with changing technologies to future-proof the system.

Chassis: The FlashBlade//S chassis is 5RU high and has bays for mounting up to ten blades. Fully populated with high-density blades, a chassis holds 1.92PB of physical flash with headroom for future density increases.

The chassis midplane distributes power and has dual Ethernet links capable of operating at 100Gbps to each blade. Blades connect to two Fabric I/O Modules (FIOMs) on the midplane. The FIOMs have Ethernet switches that connect blades to the client network, or in multi-chassis systems, to external Fabric Modules (XFM)s).

The chassis has four power supply units (PSUs), each rated at 2,400 watts. The PSUs are "2+2 redundant"—any two of them can supply the chassis' maximum rated power of 4,800 watts, to accommodate the greater demands of future components.



FlashBlade//S Chassis




Figure 1. FlashBlade//S

Blade: The blades have CPUs, NICs, DRAM, and four DirectFlash Module (DFM) mounting slots. Blades use NVMe over on-board PCIe to communicate with their DFMs. You can configure blades with either performance-optimized (24TB) or capacity-optimized (48TB) DFMs.

Blades operate with one, two, three, or four DFMs installed. All blades in the chassis must have the same number of DFMs in each blade. Systems can be configured for applications from the very read-intensive (e.g., artificial intelligence and machine learning) to those requiring extremely high capacity (e.g., backup and archiving).

The blade design isolates failure domains. You can replace a failed DFM from the front of the chassis without affecting its blade's ability to function. DFMs can move to different bays in a chassis to minimize the operational impact of recovering from blade failures or replacements.

DFMs with QLC flash: Architectures that use off-the-shelf, solid-state drives (SSDs) have an internal controller to manage the flash media on each specific drive. These systems do not have any visibility into what is happening at the system level. FlashBlade//S takes a different and innovative approach with DirectFlash modules (DFMs) that enable the storage operating system to manage the media on a global level. Global media management unlocks as much as 20% more capacity from NAND, compared to systems that use off-the-shelf SSDs, and delivers more consistent performance, better reliability, and higher media endurance without the need for a massive and expensive storage class memory (SCM) cache.

Network fabric: The integrated networking in FlashBlade//S simplifies large-scale deployments by collapsing three networks (front-end, control, and back-end) into one high-performance software-defined networking (SDN) fabric. This SDN is shared across the two fabric modules in the platform, and it hides the complexity of networking from the administrator.

FlashBlade//S virtualizes the network, so that no matter the size of the platform, it appears as one entity. This virtualization simplifies load balancing and cabling. Each blade can service and restart any client connection and run any protocol, and the platform is stateless because the logic can run anywhere.

Dual Fabric I/O Modules (FIOMs) interconnect blades, connect chassis (in multi-chassis systems), and connect blades to clients. The FIOMs have ethernet switches with eight (8) external ports each capable of 10, 25, 40, or 100Gbps transmission rates. The switches have a total of 2Tbps cross-sectional bandwidth. Each FIOM uses 50Gbps for inter-blade communication in the chassis. Both FIOM switches and blade NICs are capable of 100Gb/s for future expansion.



Purity//FB: The Purity//FB operating system is the heart of FlashBlade//S, enabling scalability in capacity and performance. Purity//FB is all-inclusive software with enterprise-grade data services. The design of Purity//FB optimizes the power of the hardware with its variable block metadata engine and scale-out metadata architecture. It can manage billions of files and objects and deliver unmatched performance for any workload, whether it's sequential or random access with large or small files and objects. Purity//FB delivers a rich set of enterprise capabilities including compression, global erasure coding, always-on encryption, SafeMode™, file replication, object replication, and multiple other features.

Environmental efficiency: Environmental, social, and corporate governance (ESG) initiatives are more important than ever. As a result, space and power constraints are becoming crucial considerations in storage strategy. The architectural design of FlashBlade//S uncomplicates the relationship between data storage and a lower carbon footprint while saving data center space, with streamlined energy consumption and more efficient power and cooling. When combined, this creates a storage solution that has a significant and immediate impact on the environment with lower overall TCO.

FlashBlade//S and Evergreen//Forever: The modular design of FlashBlade//S simplifies adding storage capacity and compute resources with non-disruptive hardware upgrades. This has made it possible for Pure to offer Evergreen//Forever™ service for the new systems. Evergreen//Forever has several advantages, including the ForeverFlash lifetime media guarantee, Ever Agile upgrades with trade-in credits, “flat and fair” service pricing guarantees, and periodic hardware refresh at no incremental cost.

New generations of storage hardware typically appear every three to five years. Historically, this meant that users would effectively repurchase capacity they already owned and had to migrate hundreds of terabytes of data from old to new hardware. FlashBlade//S gets better over time, though. With Evergreen//Forever plus the non-disruptive upgrades on FlashBlade//S, the system can keep pace with hardware evolution.

Commvault Backup & Recovery and Metallic

Commvault Backup & Recovery is an industry-leading data protection software for medium to large enterprises. Commvault is known for a flexible, scalable architecture, broad application integration, and cloud capabilities. Solid performance and ransomware hardening and detection round out the platform capabilities.

Metallic Cloud Storage Service is a fully operational cloud storage backup target for Commvault Backup & Recovery and is completely integrated with Commvault data management software. With this cloud service, customers can simplify their cloud data management with pre-configured networking and storage, reduce costs via efficient deduplication and no egress fees, and mitigate ransomware with secure, air-gapped cloud data protection.

Rapid Restore in Partnership with Commvault

Rapid restore with Commvault and FlashBlade//S replaces or augments both purpose-built backup appliances (PBBAs) and BYO disk-based backup targets. Data reduction in Commvault and FlashBlade//S provide storage efficiency. The high bandwidth of FlashBlade//S gives you the throughput you need to restore your critical data fast, especially when you need it most in a large-scale recovery event, such as a ransomware attack. FlashBlade//S shortens your recovery time objective (RTO), reducing critical downtime costs. You can also take advantage of the performance to run more frequent backups and reduce your recovery point objective (RPO).

For maximum restore performance, it's critical that your primary data reside on high-performance storage such as Pure FlashArray//X™, otherwise you're going to severely limit how fast you can recover it.



Solution Architecture

A pool of MediaAgents (MAs) acts as the backup and restore data movers, writing to FlashBlade//S using Amazon S3 protocol. Virtual Server Agents (VSAs), either standalone or installed on the MA servers, protect the VMs. Application agents (iDAs) protect databases and other data types. The VSAs and iDAs can send backups through the MAs, or they can have [Storage Accelerator](#) installed and communicate directly with the FlashBlade. For long-term retention, you can use Commvault's auxiliary copy (aux copy) feature to transfer data to a lower tier such as FlashArray//C, cloud, or other supported storage types.

With compatible hardware, such as FlashArray//X, you can use [Commvault IntelliSnap](#) to orchestrate application-consistent hardware snapshots for backup and recovery, and further reduce RPO and RTO.

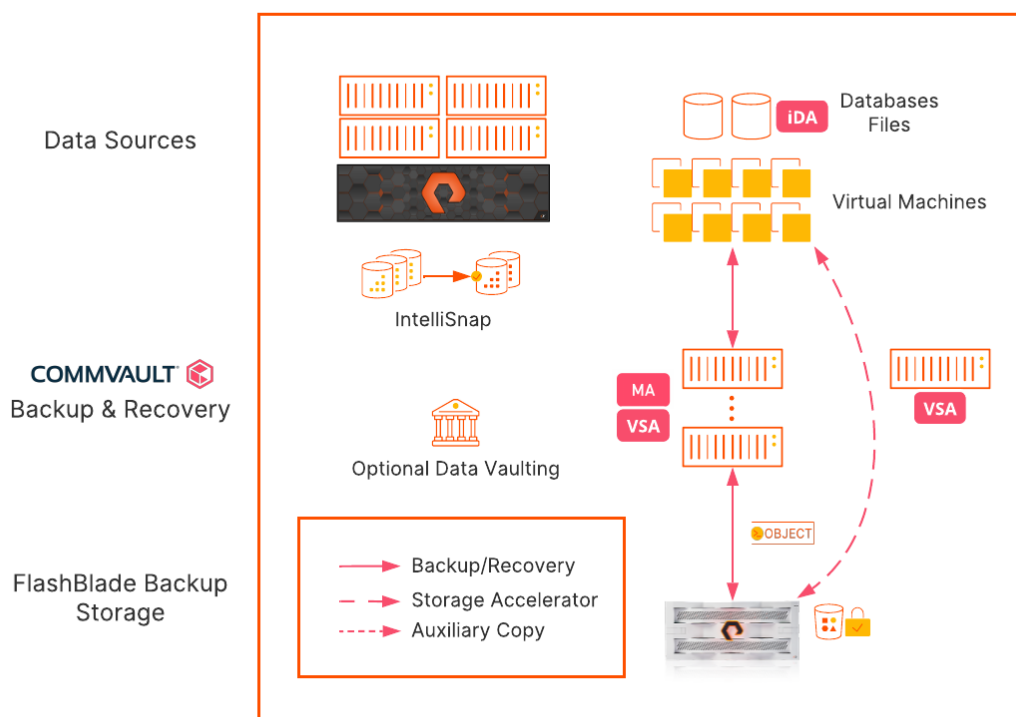


Figure 2. Solution logical architecture

Scaling

As your environment grows, you may need to increase your backup and restore speeds. FlashBlade//S can support speeds more than double the first-generation FlashBlade, to handle even your most demanding requirements. And if you aren't fully utilizing the FlashBlade bandwidth you have, you can expand your Commvault infrastructure to drive more performance. Following are the ways you can scale this architecture:

Expanding FlashBlade//S: You can transparently add blades or storage to the FlashBlade//S to increase performance and capacity. When you expand the FlashBlade, you do not need to make any changes to the FlashBlade or Commvault



configuration to take advantage of the new hardware. Work with your Pure account team to figure out the right upgrades for you.

Adding MediaAgents: You can add MediaAgents to increase the amount of bandwidth available for your backup infrastructure. It is very simple to add MAs into an existing storage pool, and you do not need to reconfigure the FlashBlade. If you reach the performance limit of a storage pool, you can add a new pool on separate MAs, with their own deduplication databases.

Storage Accelerator: Another way to scale your environment without adding MAs is to use Commvault's [Storage Accelerator](#) feature. With Storage Accelerator, clients send and receive data directly to FlashBlade//S, bypassing the MA. This effectively lets the MA, also referred to as the Data Server in this context, manage more data; these clients can use most or all their available bandwidth without significantly increasing the MA network utilization. The clients will still send the index data for the backup jobs to the MA, and the MA will still process synthetic full backups. The number of data streams the MA and deduplication databases (DDBs) can handle and the DDB performance will eventually be the limiters. Figure 3 illustrates the data path differences with Storage Accelerator.

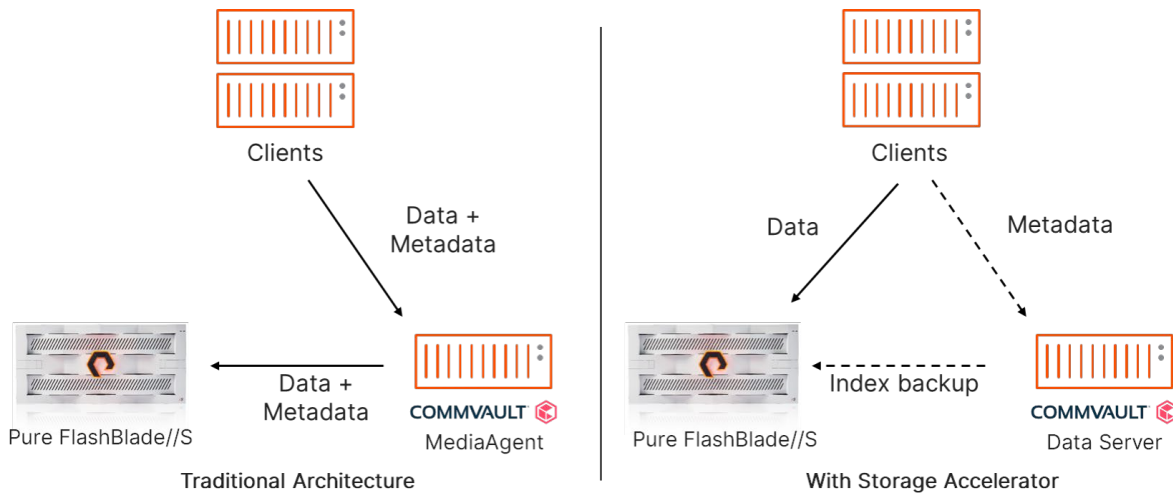


Figure 3. Storage Accelerator overview

NOTE: For CommCells built on release 11.26 or later, by default Commvault will automatically install the Storage Accelerator package on clients sending data to object storage. You can change this behavior in the [media management service configuration](#).

Additional Sites

To protect data to FlashBlade in two sites, for purposes such as disaster recovery or site resilience, you can deploy a second FlashBlade//S and Commvault MediaAgents in the remote site. You can leverage aux copy to efficiently transfer new unique data to the second site, with throttling available if you need to manage intersite bandwidth. Aux copy also lets you define different data retention settings for each copy of the data.

If you need to keep only offsite backups of data housed on FlashArray, you can use IntelliSnap to simplify the environment. IntelliSnap can manage FlashArray asynchronous replication or ActiveCluster synchronous replication to let the arrays move data between sites, then back up the data from the replica to FlashBlade//S. Figure 4 shows an example multi-site deployment using both replication options.



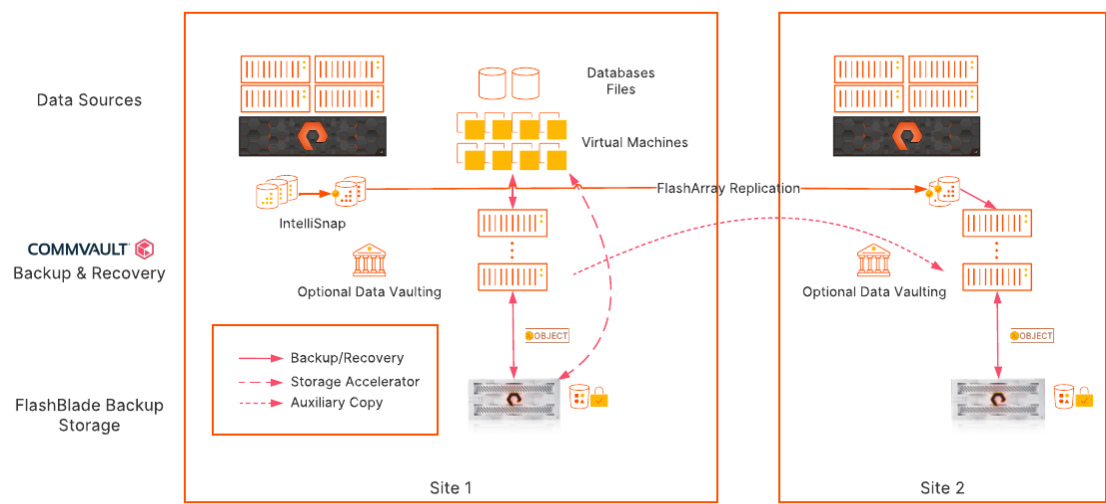


Figure 4. Multi-site deployment

Lab Setup

This section covers the lab environment used for testing the solution, including system details and configuration procedures.

NOTE: Performance maximums and scaling will vary depending on the infrastructure and resources available. In this lab test, we demonstrate backup and recovery but do not test for or achieve the maximum FlashBlade//S throughput with the equipment available.

Server Details

For the lab testing, we used a 4-node ESXi cluster and four physical Windows servers with both Commvault MAs and VSAs installed. Table 1 shows the hardware and configuration details.

Server Role	CPU	RAM	Networking	Storage	Operating System/Software
ESXi Host (x4)	2x Intel Xeon E5-2697 v2 @ 2.70GHz, 24 cores total Hyperthreading enabled	512GB	2 x Mellanox MT27500 family network adapter @ 40Gbps	9 datastores from 3 FlashArray//M70 arrays	VMware ESXi 6.7.0
Commvault MediaAgent and Virtual Server Agent (x4)	2x Intel Xeon E5-2697 v2 @ 2.70GHz, 24 cores total Hyperthreading enabled	512GB	2 x Mellanox MT27500 family network adapter @ 40Gbps, in LACP team	300GB Intel MLC SSD on 6Gb SATA (boot) 1TB Toshiba MLC SSD on 6Gb SATA (Index/DDB) iSCSI connection to FlashArray//M70	Windows Server 2022, 21H2

Table 1. Server configuration details

Storage Details

A FlashBlade//S S200 hosted the object bucket for the Commvault cloud target. Source VMs were hosted on FlashArray//M70 arrays.

The source arrays were connected using 4x10GbE iSCSI. Table 2 shows the array details.

Storage Role	Array Model	Purity Release	Physical Storage	Connectivity
Commvault Cloud Target	FlashBlade//S S200 7x2 (capacity optimized)	Purity//FB 4.0.0 (beta)	159.6TiB (usable)	8x100Gb Ethernet
Data Source (x3)	FlashArray//M70	Purity 6.1.4	21TiB (usable)	4x10Gb Ethernet (iSCSI)

Table 2. Storage details

Source Data Details

We built 96 source VMs with the following configuration. The VMs were evenly distributed across the ESXi hosts and FlashArray//M70 datastores, with 24 VMs per host and 32 VMs per array. Each test used some part of the VM set, distributed as evenly as possible across the hosts and arrays. Table 3 shows the VM details.

VM Role	CPU	RAM	Networking	Storage	Operating System
Data Source (x96)	2 vCPU	4GB	vmxnet3 adapter	1x100GB VMDK Thin provisioned	Windows 10

Table 3. Source data details

Adding FlashBlade//S as a Commvault Cloud Storage Pool

Configure FlashBlade Object Storage

In the FlashBlade management interface, navigate to the Storage page, then click the **Object Store** tab. In the Accounts tile, click the **Create (+)** button to add an account. In the **Create Account** dialog box (Figure 5), enter a name for the account. Click the **Create** button to commit the account.

NOTE: This is a display name internal to FlashBlade and is not used by Commvault

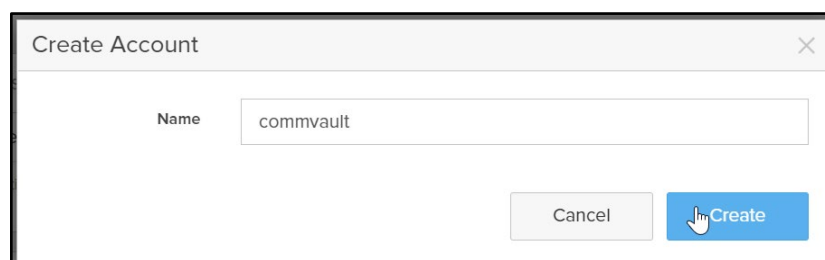


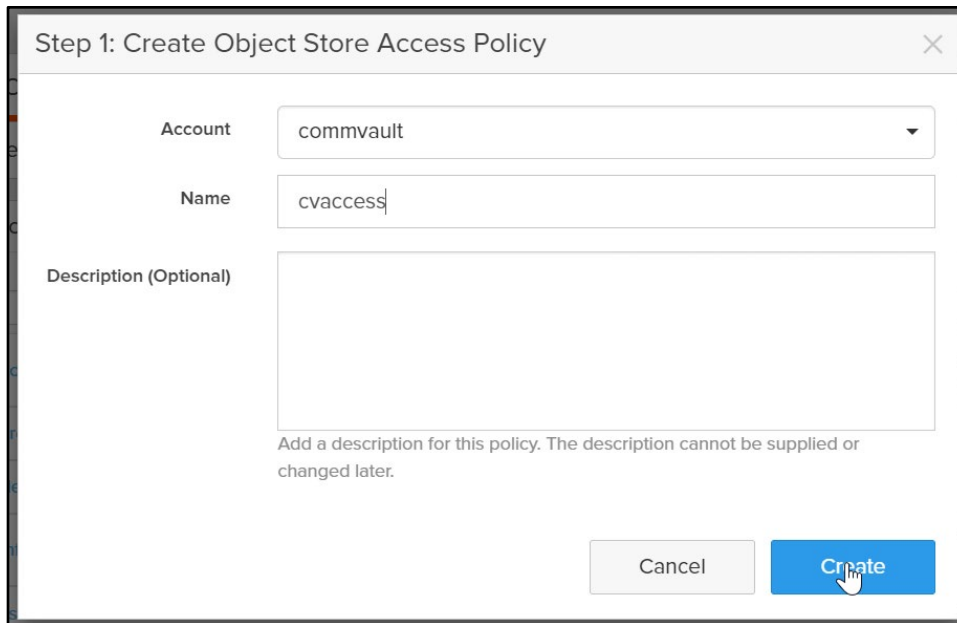
Figure 5. Create Account dialog

To manage Commvault's bucket access, create an object access policy with the required permissions. Navigate to the Policies page, then click the **Object Store Access** tab. In the Object Store Access Policies tile, click the **Create (+)** button to add a policy. Complete the Step 1 dialog (Figure 6) as follows.

1. From the **Account** dropdown, select the object account you created.
2. In the **Name** field, enter a name for the access policy.
3. Optionally, enter a description for the policy in the **Description** field.



4. Click the **Create** button to create the policy and continue to step 2.

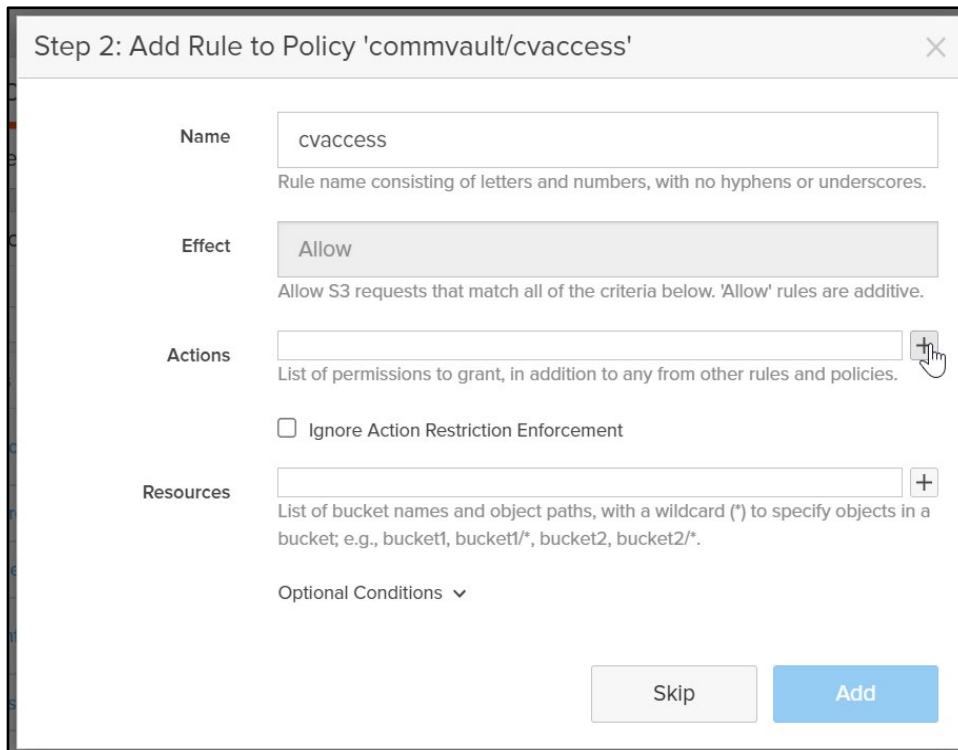


The dialog box is titled "Step 1: Create Object Store Access Policy". It contains three input fields: "Account" with a dropdown menu showing "commvault", "Name" with a text box containing "cvaccess", and "Description (Optional)" with a large empty text area. Below the text area is a note: "Add a description for this policy. The description cannot be supplied or changed later." At the bottom right are two buttons: "Cancel" and "Create". A mouse cursor is clicking the "Create" button.

Figure 6. Creating an access policy: Step 1

Complete the Step 2 dialog (Figure 7) as follows.

1. In the **Name** field, enter a name for the rule. This is separate from the policy name.
2. Click the **+** button next to the **Actions** field to add allowed actions.



The dialog box is titled "Step 2: Add Rule to Policy 'commvault/cvaccess'". It contains several fields: "Name" with a text box containing "cvaccess" and a note "Rule name consisting of letters and numbers, with no hyphens or underscores."; "Effect" with a dropdown menu showing "Allow" and a note "Allow S3 requests that match all of the criteria below. 'Allow' rules are additive."; "Actions" with a text box and a "+" button, and a note "List of permissions to grant, in addition to any from other rules and policies."; "Ignore Action Restriction Enforcement" with an unchecked checkbox; "Resources" with a text box and a "+" button, and a note "List of bucket names and object paths, with a wildcard (*) to specify objects in a bucket; e.g., bucket1, bucket1/*, bucket2, bucket2/*."; and "Optional Conditions" with a dropdown arrow. At the bottom right are two buttons: "Skip" and "Add". A mouse cursor is clicking the "+" button next to the "Actions" field.

Figure 7. Creating an access policy: Step 2



3. In the Add Actions dialog box (Figure 8), select the actions listed in Table 4, then click the **Add** button to return to step 2.

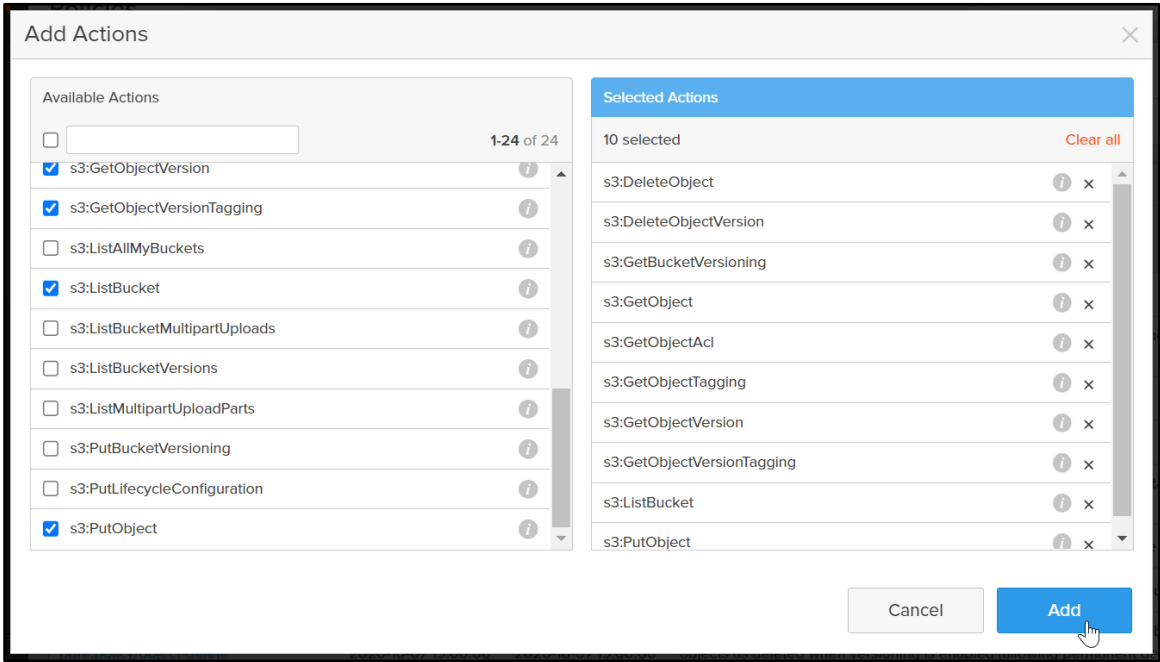


Figure 8. Add Actions dialog

Required Actions

- s3:DeleteObject
- s3:DeleteObjectVersion
- s3:GetBucketVersioning
- s3:GetObject
- s3:GetObjectAcl
- s3:GetObjectTagging
- s3:GetObjectVersion
- s3:GetObjectVersionTagging
- s3:ListBucket
- s3:PutObject

Table 4: Required actions

Complete Step 2 as follows (Figure 9):

1. Click the **+** button next to the **Resources** field. In the dialog that appears, enter **"*"** to grant access to all buckets within the account.
2. Click the **Add** button to finish creating the rule.

Step 2: Add Rule to Policy 'commvault/cvaccess'

Name
Rule name consisting of letters and numbers, with no hyphens or underscores.

Effect
Allow S3 requests that match all of the criteria below. 'Allow' rules are additive.

Actions

s3:DeleteObject X	s3:DeleteObjectVersion X
s3:GetBucketVersioning X	s3:GetObject X
s3:GetObjectAcl X	s3:GetObjectTagging X
s3:GetObjectVersion X	s3:GetObjectVersionTagging X

List of permissions to grant, in addition to any from other rules and policies.

☐ Ignore Action Restriction Enforcement

Resources X

List of bucket names and object paths, with a wildcard (*) to specify objects in a bucket; e.g., bucket1, bucket1/*, bucket2, bucket2/*.

Optional Conditions ▼

Figure 9. Creating an access policy: Step 2 completed

Navigate to the Storage page, then click the Object Store tab. In the Buckets tile, click the **Create** (+) button to add an object bucket. In the Create Bucket dialog box (Figure 10), in the **Bucket Name** field, enter a name for the bucket, then click the **Create** button.

Create Bucket

Bucket Name

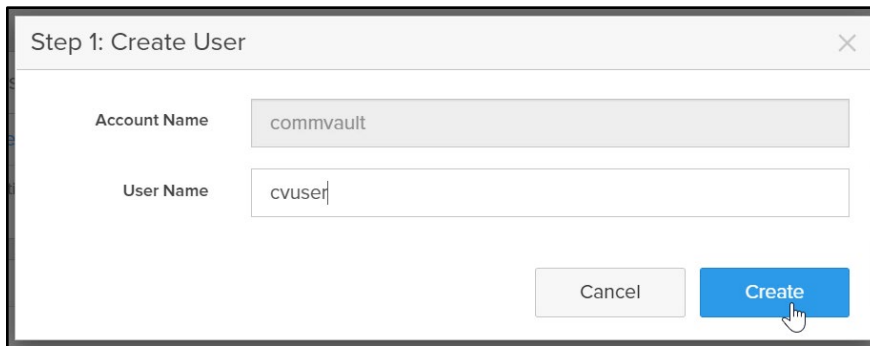
Figure 10. Create Bucket dialog

In the Users tile, click the **Create** (+) button to open the user creation wizard.

On Step 1 (Figure 11), in the **User Name** field, enter a name for the object user, then click the **Create** button.

NOTE: This name will not need to be entered in Commvault.





Step 1: Create User

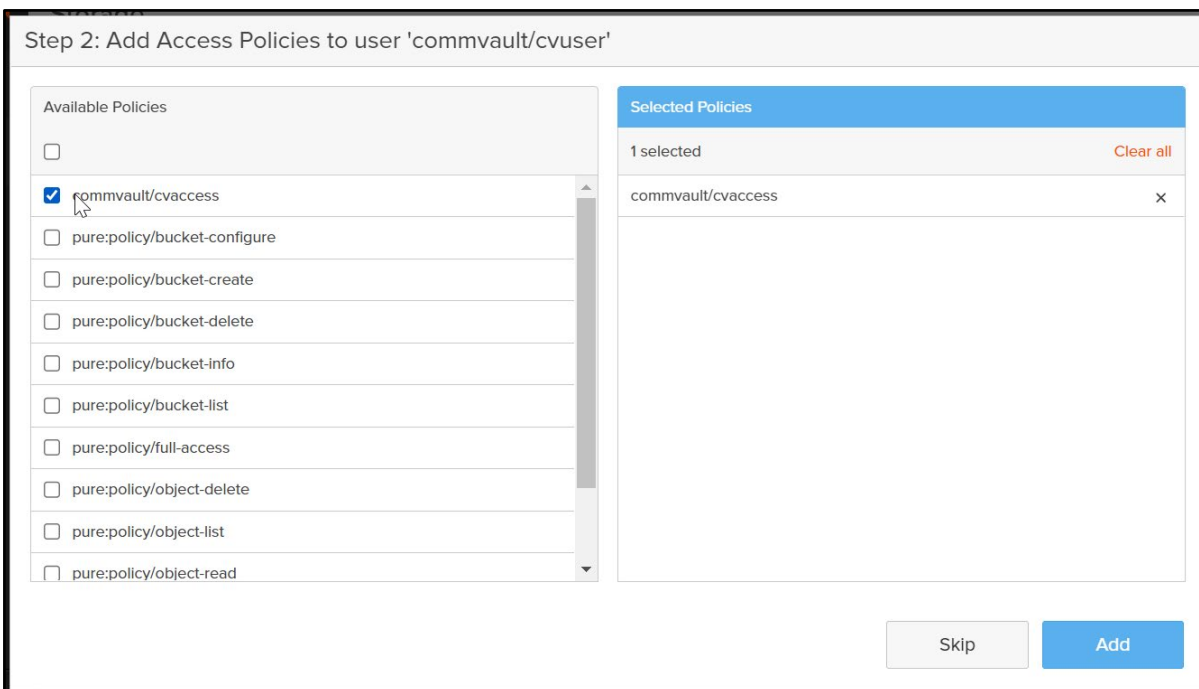
Account Name: commvault

User Name: cvuser

Buttons: Cancel, Create

Figure 11. User creation wizard: Step 1

On Step 2 (Figure 12), select the access policy you created earlier, then click the **Add** button.



Step 2: Add Access Policies to user 'commvault/cvuser'

Available Policies:

- ☐
- ☒ commvault/cvaccess
- ☐ pure:policy/bucket-configure
- ☐ pure:policy/bucket-create
- ☐ pure:policy/bucket-delete
- ☐ pure:policy/bucket-info
- ☐ pure:policy/bucket-list
- ☐ pure:policy/full-access
- ☐ pure:policy/object-delete
- ☐ pure:policy/object-list
- ☐ pure:policy/object-read

Selected Policies:

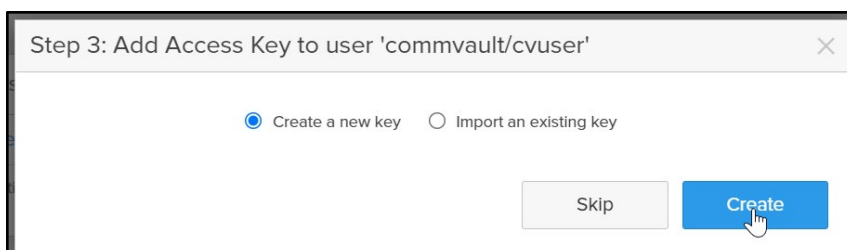
1 selected [Clear all](#)

commvault/cvaccess [x](#)

Buttons: Skip, Add

Figure 12. User creation wizard: Step 2

On Step 3 (Figure 13), select the **Create a new key** option, then click the **Create** button.



Step 3: Add Access Key to user 'commvault/cvuser'

☒ Create a new key ☐ Import an existing key

Buttons: Skip, Create

Figure 13. User creation wizard: Step 3



The Access Key dialog appears. Switch to Commvault Command Center to continue.

IMPORTANT: Do not close the Access Key dialog before inputting the keys into Commvault. If you do, you will need to create new access keys for the user.

Create Cloud Storage Pool in Commvault

NOTE: The steps below apply only to Commvault release 11.26. Specific steps and screen appearance may vary between Commvault releases.

In the left-hand navigation pane, expand **Storage**, then click **Cloud**. From the **Cloud** page, click the **Add** link. Complete the **Add cloud storage** form (Figure 14) as follows:

1. In the **Name** field, enter a display name to help identify the FlashBlade.
2. From the **Type** dropdown, select "S3 Compatible Storage."
3. From the **MediaAgent** dropdown, select the MediaAgent that will act as primary controller for the bucket.
4. In the **Service host** field, enter the DNS name or IP address for the FlashBlade data VIP. If you wish to disable TLS, include "http://" in the field.

Cloud / Cloud storage type

Add cloud storage

Configure cloud

Name *
FlashBlade//S

Storage

Type
S3 Compatible Storage

MediaAgent *
sn1-r720-g08-07

Service host *
http://10.21.237.25

Credentials *
FlashBlade//S S3

Bucket *
cvbucket

Figure 14. Add cloud storage form

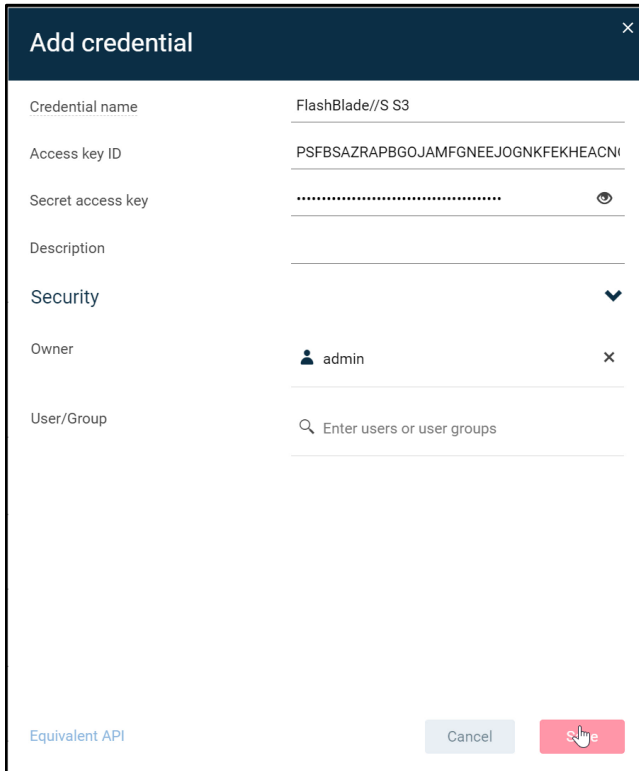
Next to the **Credentials** dropdown, click the **Create new (+)** button to add the FlashBlade keys into a stored credential.

In the **Add credential** form (Figure 15), enter a name for the credential. Switching between the FlashBlade and Commvault interfaces, use the **Copy** button in the FlashBlade interface (Figure 16) to copy the access key ID and secret key and paste



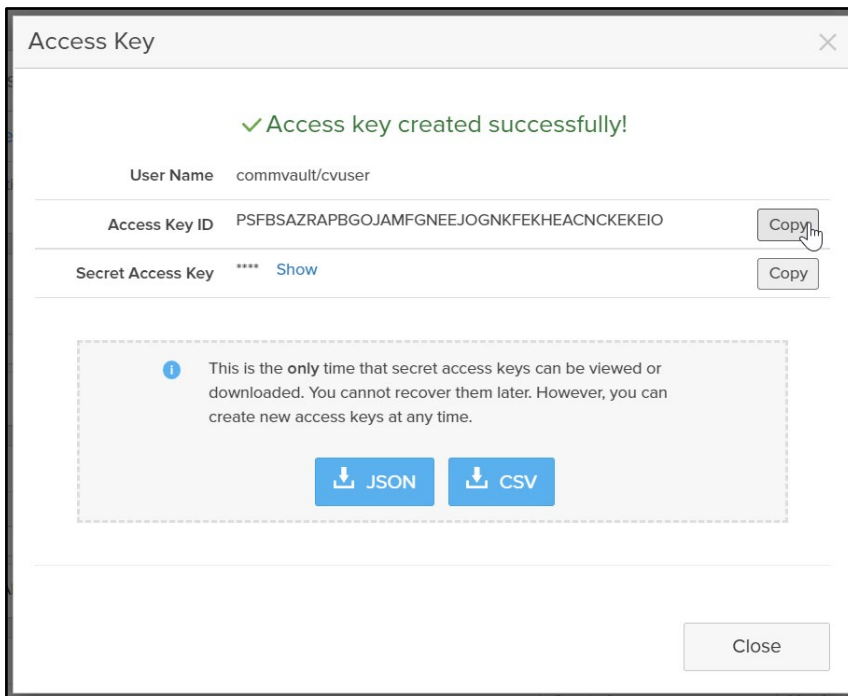
them into the appropriate Command Center fields. You may enter descriptive text in the **Description** field. Click the **Save** button to create the stored credential and return to the **Add cloud storage** form.

NOTE: You should always clear the clipboard after copying sensitive data such as object storage access keys.



The 'Add credential' form is a modal window with a dark blue header. It contains several input fields: 'Credential name' with the value 'FlashBlade//S3', 'Access key ID' with a long alphanumeric string, 'Secret access key' with a masked value and an eye icon, 'Description' (empty), 'Security' (dropdown arrow), 'Owner' with a user icon and 'admin', and 'User/Group' with a search icon and placeholder text 'Enter users or user groups'. At the bottom left is a link 'Equivalent API', and at the bottom right are 'Cancel' and 'Save' buttons.

Figure 15. Add credential form



The 'Access Key' dialog shows a success message '✓ Access key created successfully!'. It displays the 'User Name' as 'commvault/cvuser', the 'Access Key ID' as a long alphanumeric string with a 'Copy' button, and the 'Secret Access Key' as masked text with a 'Show' button and a 'Copy' button. Below this is an information box with a warning icon and text: 'This is the only time that secret access keys can be viewed or downloaded. You cannot recover them later. However, you can create new access keys at any time.' At the bottom of the information box are 'JSON' and 'CSV' download buttons. A 'Close' button is at the bottom right of the dialog.

Figure 16. Copying access keys



Complete the **Add cloud storage** form.

1. In the **Bucket** field, enter the name of the object bucket you created on the FlashBlade (Figure 17).

Cloud / Cloud storage type

Add cloud storage

Configure cloud

Name *

FlashBlade//S

Storage

Type ○ ○

S3 Compatible Storage

MediaAgent *

sn1-r720-g08-07

Service host *

http://10.21.237.25

Credentials *

FlashBlade//S S3

Bucket *

cvbucket

Figure 17. Completing the Add cloud storage form

2. For each DDB partition you want to create, click the **Add** link next to the **Deduplication DB location** section. In the dialog that appears, select the appropriate MediaAgent, and enter or browse to the path in the file system where you want the DDB to reside (Figure 18). We recommend creating partitions on at least two MAs to provide performance and resilience. You can create up to four partitions per storage pool.

Add Deduplication DB location

MediaAgent *

sn1-r720-g08-25

Deduplication DB location *

D:\DDB

CANCEL ADD

Use deduplication ☒

Deduplication DB location

Figure 18. Adding a DDB

Click the **Save** button to create the cloud storage pool.



Configure 128KB Deduplication Block Size

By default, Commvault configures cloud storage pools with a 512KB deduplication block size. Best practice with FlashBlade is to decrease the block size to 128KB for better space efficiency. Deduplication block size is configured using the CommCell Console.

In the CommCell Browser pane, expand **Storage Resources**, then **Storage Pools**. Right-click the storage pool you want to change, then click **Properties>Storage Pool**.

In the storage pool properties dialog (Figure 19), select the **Advanced** tab, then change the **Block level Deduplication factor** dropdown to 128. Click the **OK** button to commit the change.

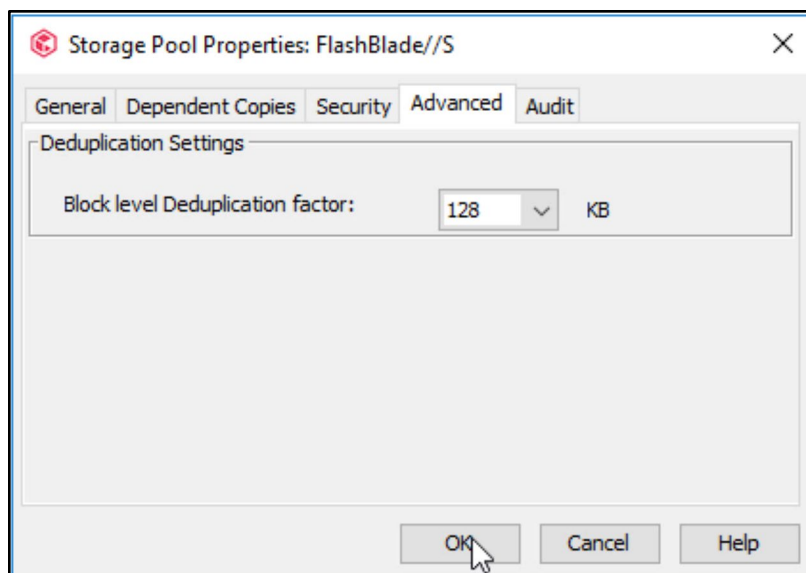


Figure 19. Changing deduplication block size

In the confirmation dialog box that appears, enter the confirmation text, then click the **OK** button.

NOTE: Any storage policies you create after the change, including those associated with server backup plans, will inherit the new setting, but existing storage policies will keep the current settings. This only affects copies that use local deduplication, not copies tied to storage pools.

Create Server Backup Plan

Return to Command Center to create a server backup plan. In the left-hand navigation pane, expand **Manage**, then click **Plans**. On the **Plans** page, click the **Create plan** dropdown, then select "Server backup."

On the first page of the wizard (Figure 20), select the **Create a new plan** option. Enter a name for the plan. Click the **Next** button to continue.



1 General

2 Backup destinations

3 RPO

4 Backup content

5 Options & restrictions

General

☒ Create a new plan

New backup plan from scratch

☐ Use existing base plan

Create plan by inheriting setting from base plan

Plan name *
FlashBlade//S backup

CANCEL

NEXT

Figure 20. Create server backup plan wizard: Step 1

On the **Backup destinations** step (Figure 21), click the **Add copy** button to add FlashBlade//S as a storage target. In the **Add copy** dialog box, you may optionally change the **Name** field. In the **Storage** dropdown, select the FlashBlade//S cloud storage pool you created. If you need to change the retention rules or set extended retention, you can do that here as well. Click the **Save** button to add the target.

Add copy

Name *

Primary

Storage *

FlashBlade//S

Retention rules

1

Month(s)

Extended Retention rules

CANCEL

SAVE

Figure 21. Create server backup plan wizard: Step 2

If you need to add destinations for any secondary storage targets, click the **Add copy** button again and complete the dialog again. When you have created all the destinations, click the **Next** button to continue.

On the **RPO** step (Figure 22), you can add or edit backup schedules and windows to meet your needs. Once the options are configured appropriately, Click the **Next** button to continue.

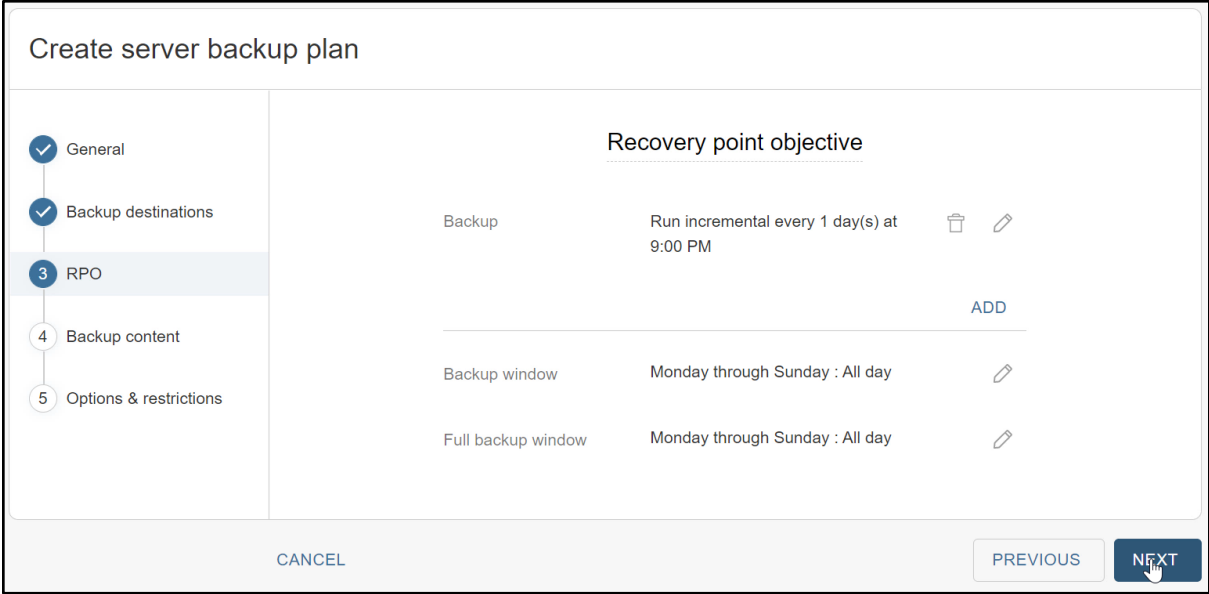


Figure 22. Create server backup plan wizard: Step 3

On the **Backup content** step, you can set any backup options that will apply to default subclients, or data sets, when you link them to the backup plan. The default options are appropriate for most environments. Click the **Next** button to continue.

On the **Options & restrictions** step (Figure 23), you can change retention and RPO for snapshots managed by IntelliSnap, RPO and disk cache options for database log backups, and allow or disallow overrides of plan options. Click the **Submit** button to create the server backup plan.



✓ General

✓ Backup destinations

✓ RPO

✓ Backup content

5 Options & restrictions

Options & restrictions

Snapshot options

☒ Retention period

1

Month(s)

▼

☐ Number of snap recovery points

8

Enable backup copy

☒

Backup copy RPO

4

hours

0

minutes

Database options

Log backup RPO

4

hours

0

minutes

Use disk cache for log backups

☐

Override restrictions

Allow plan to be overridden

☐

CANCEL

PREVIOUS

SUBMIT

Figure 23. Create server backup plan wizard: Step 4

Excluding Commvault Processes from AV Scanning

The high I/O levels involved in backup and deduplication mean antimalware (AV) real-time scanning can have a large impact on performance. We found that excluding the processes listed in Table 5 from AV scanning gave the biggest improvement in backup and recovery performance.

Exclude from AV Scanning

C:\Program Files\Commvault\ContentStore\Base\3dnfsd.exe
C:\Program Files\Commvault\ContentStore\Base\CIMgrS.exe
C:\Program Files\Commvault\ContentStore\Base\cvd.exe
C:\Program Files\Commvault\ContentStore\Base\cvfwd.exe
C:\Program Files\Commvault\ContentStore\Base\CVMountd.exe
C:\Program Files\Commvault\ContentStore\Base\CVODS.exe
C:\Program Files\Commvault\ContentStore\Base\InstallUpdates.exe
C:\Program Files\Commvault\ContentStore\Base\SIDB2.exe
C:\Program Files\Commvault\ContentStore\Base\vsbkp.exe
C:\Program Files\Commvault\ContentStore\Base\vsrst.exe
C:\Program Files\Commvault\ContentStore\Base\vsrst.exe

Table 5. Processes to exclude from AV scanning

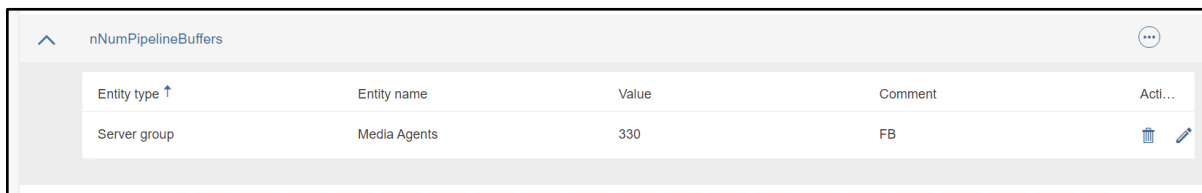
The paths to these processes will change if you install Commvault to a different location. See Commvault best practices for [Windows](#) and [Linux](#), as well as your AV documentation, for details on excluding processes from scanning.

Tuning Performance

We found that some tuning options showed considerable performance benefits with FlashBlade//S in our environment. The level of improvement when using these options will vary widely and may actually degrade performance, so make sure to test before implementing in production.

Pipeline Buffers

We found there was a network bottleneck between the VSA and MA components on our servers. To improve throughput, we [increased the number of pipeline buffers](#) from the default of 90 to 330. Beyond 330 buffers we saw degradation. Figure 24 shows the setting.



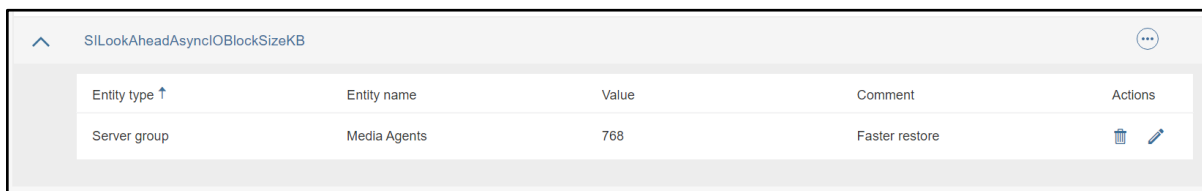
Entity type ↑	Entity name	Value	Comment	Acti...
Server group	Media Agents	330	FB	

Figure 24. Increasing pipeline buffers

This setting increases the RAM usage and the optimal setting is highly dependent on the specific environment. Use caution when changing it.

Look-Ahead Block Size

During restore, Commvault will [read additional deduplicated blocks](#) under the assumption that it will need them, saving time having to fetch them from storage. By default, with object storage it will fetch twice the deduplicated block size. We found that setting the look-ahead size to 768KB improved restore performance by up to 20% with FlashBlade//S. While the impact may vary between deployments and data sets, it is worth testing in your own environment. Figure 25 shows the setting.



Entity type ↑	Entity name	Value	Comment	Actions
Server group	Media Agents	768	Faster restore	

Figure 25. Increasing look-ahead block size

Testing Details

We set out to determine how backup and restore performance changed with parallel VMs. We created a set of VM groups (Figure 26) with increasing numbers of VMs, from one to 32 VMs per group. For test sets larger than 32 VMs, we divided the VMs evenly across two or three VM groups and backed up the groups in parallel.

Each VM group used VMs from no more than one datastore per FlashArray. VM groups for multiple parallel jobs did not overlap datastores. This separation prevented potential issues from mounting multiple copies of the same datastore at the same time.



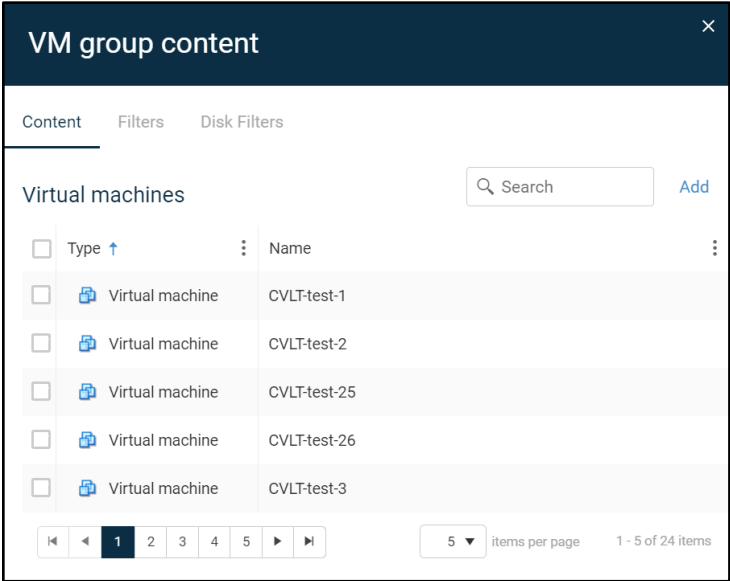


Figure 26. Example VM group content

We enabled hardware snapshots on the VM group and set the number of readers to be equal to the number of VMs in the group (Figure 27).

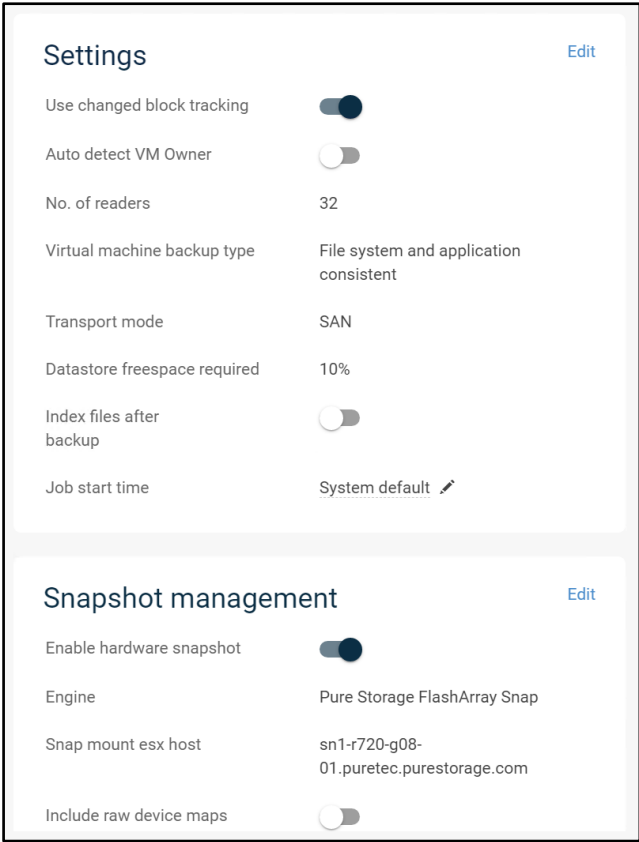


Figure 27. VM group configuration

We ran a series of operations against each VM group, using a custom workflow. We simulated a two-week backup cycle for each VM group, with each week comprising a baseline full backup, six simulated days of incremental backups with 2GiB data

change on each VM, and a synthetic full backup. To compare full and synthetic full backup performance, we ran a final full backup with no data change. We then ran parallel full VM restores for all the VMs in the group. We sealed the deduplication database (DDB) between tests to ensure Commvault captured a complete baseline for each group.

Because Commvault was performing deduplication and compression, we did not measure the data reduction rates on the FlashBlade//S for each test.

To limit the influence of ESXi servers, VM load, and networking on backup and restore throughput, all full and incremental backups used FlashArray snapshots orchestrated through Commvault IntelliSnap. We measured the throughput as reported per VM by [backup copy](#) jobs and [synthetic full backups](#). Backup copies all used SAN transport mode and [multi-node backup copy](#), where copies of snapshots were attached to all of the VSA/MA servers and processed in parallel. All full-VM restores also used SAN mode. For more information on transport modes, see [Commvault documentation](#).

All tests were performed on Commvault release 11.26.

Test Results

Backup

We measured the effective throughput of the different types of VM backups using a single VM. We then repeated the test with increasing numbers of parallel VM backups and measured the performance changes. Figure 28 shows the results.

NOTE: Please note that this test was designed to show real-world backup speeds on available hardware in our lab, and not maximum performance of the FlashBlade//S line. Achieving maximum performance would require additional servers, networking, and primary storage beyond what was available.




Figure 28: Backup performance scaling by number of VMs



All the backup types saw steady improvement as the number of VMs increased, up to a peak. Baseline full backups reached an effective 3.6GiB/s (12.76TiB/hr.), and incremental backups reached 2.97GiB/s (10.4TiB/hr.).

Effective backup throughput is tied to both storage write speed and DDB performance, and data reducibility plays a large part in the balance between the two. Past the peaks we started to see server resource contention that lowered effective throughput and adding more concurrency on the same hardware would further reduce throughput. Actual write speeds on FlashBlade//S peaked at around 3.2GiB/s on full backup and 2.8GiB/s on incremental backup. As noted earlier, for VM sets above 32 we divided the VMs into multiple parallel backup jobs.

Full VM Recovery

For recovery, we measured how throughput scaled as we increased the number of concurrent VMs. We restored full VMs to the same ESXi hosts and datastores from which they were backed up. Commvault imposes a limit of ten concurrent streams in a single VM restore job, so we divided the VMs into sets of up to eight and restored the sets in parallel. We limited the impact of ESXi hosts and networking by using SAN transport mode. Since SAN mode is slower with thin-provisioned virtual disks, we converted them to thick eager zero format during the restore. Figure 29 shows an example of the restore options.

NOTE: Please note that this test was designed to show real-world backup speeds on available hardware in our lab, and not maximum performance of the FlashBlade//S line. Achieving maximum performance would require additional servers, networking, and primary storage beyond what was available.



Restore options

Type

☐ In place

☒ Out of place

Destination

dp-vc

Access node

Automatic

CVLT-test-1

VM display name

CVLT-test-1-cvrestore

Destination host

sn1-r720-g08-27.puretec.purestorage.com

Browse

Datastore

m70-1-vdi2 (39.94 TB free)

Resource pool

/

VM folder path

/G08/vm/Commvault FAC/Test VMs

Browse

Network settings

>

IP address settings

>

☐ Power on VMs after restore

☐ Unconditionally overwrite if it already exists

☐ When the job completes, notify me via email

Additional options

^

Disk provisioning

Thick Eager Zero

Transport mode

SAN

☐ Restore virtual machine using live recovery (vMotion)

Equivalent API

Cancel

Submit

Figure 29: Full VM restore options

We measured throughput as the average overall rate at which data was transferred until all the concurrent restores completed. Figure 30 shows the results.




Figure 30. Restore performance scaling by number of VMs

In our lab, restore speed scaled quickly up to 24 concurrent VMs, where it reached 3.3GiB/s (11.6TiB/hr.). With more than 24 VMs it continued to increase, although less quickly, reaching 3.98GiB/s (14TiB/hr.) with all 96 VMs restoring in parallel, with an upward trend. The primary bottleneck was the CPU in the VSA/MA servers; at 96 VMs they were overloaded. With additional updated hardware and resources, we would expect a dramatic increase in speed and scale.

To that end, we did some quick restore tests using a single newer server with a 64-core AMD EPYC CPU. We found that this single VSA could handle about double the concurrent VMs as the older generation servers. It was able to restore 32 VMs at 2.35GiB/s, more than half the speed of the entire 4-node grid of older servers—which for that test had only restored 8 VMs each. For comparison, we extrapolated from those results, adding 32 VMs and one new-generation VSA at a time and assuming near-linear scaling. With enough bandwidth available on the network and primary storage, six of our new-generation servers would be able to restore 192 concurrent VMs at over 14GiB/s, or 50TiB/hr. This is still well below the maximum read speed we’ve observed for the tested FlashBlade//S configuration, so we would expect to scale even higher with more VSA servers. Figure 31 compares the actual results with the extrapolation, including the number of servers involved.






Figure 31. Restore performance comparison between server CPU generations

Ransomware Mitigation

FlashBlade//S includes the Object SafeMode™ immutability feature for ransomware mitigation. Object SafeMode is a system-wide policy that prevents anyone from deleting an object until it has reached a specific age. Combined with a layered vaulting approach in Commvault, Object SafeMode provides an immutability window where data is protected from an attacker for a minimum period of time. You can tie the immutability window to your SLAs to ensure you can restore systems quickly in the event of a cyberattack.

IMPORTANT: You should always work with your Pure account team before enabling Object SafeMode to ensure you are meeting your objectives and have adequate system capacity.

To support the dependencies over time that global deduplication creates, immutability with Commvault and Object SafeMode uses a layered vaulting approach. Commvault periodically “closes” the vault, sealing the DDB and starting a new, coherent copy of the data. The vaulting frequency matches the immutability window. Object SafeMode retention is configured at twice the immutability window, so that backups taken at the end of the vault period have all the dependencies protected for the full immutability window. Figure 32 illustrates the relationship between the layers.






Figure 32. Layered vaulting immutability solution

When you use Object SafeMode to protect your backups, we also recommend sending your daily CommServer DR backups to a FlashBlade file system and using SafeMode snapshots to protect them.

Configuring SafeMode Protection

There are four key elements used to configure Object SafeMode: automatic DDB sealing, disabling micro pruning, Commvault retention policy, and Object SafeMode retention policy.

You can enable automatic DDB sealing in the CommCell Console. In the CommCell Browser pane, expand **Storage Resources**, then **Deduplication Engines**. Right-click the engine for the DDB you wish to vault, then click **Properties**. The Properties dialog box will open, with the Deduplication tab and Settings sub-tab selected.

In the Deduplication Database creation section (Figure 33), enable the first **Create new DDB every** checkbox, then set the number of days to match the immutability window. For example, if you require your most recent 14 days of backups to be protected from attack, set the frequency to 14 days. Click the **OK** button to commit the change.

Figure 33. Automatic DDB sealing options

Micro pruning is configured using the CommCell Console, at the bucket or mount path. To disable micro pruning, expand Storage Resources, then Libraries. Expand the library for the FlashBlade. For each mount path, right-click the mount path and select Properties. In the Properties dialog box, select the Allocation Policy tab. In the lowest section (Figure 34), disable the **Enable Micro Pruning** checkbox. Click the OK button to commit the change and click the OK button on the warning dialog that appears.






Figure 34. Disabling micro pruning

You can change Commvault retention for server backup plans using Command Center. In the left-hand navigation, expand Manage, then click Plans. In the list of plans, click the name of the one you want to change to open the plan details. In the Backup destinations tile (Figure 35), click the name field for the FlashBlade//S copy.

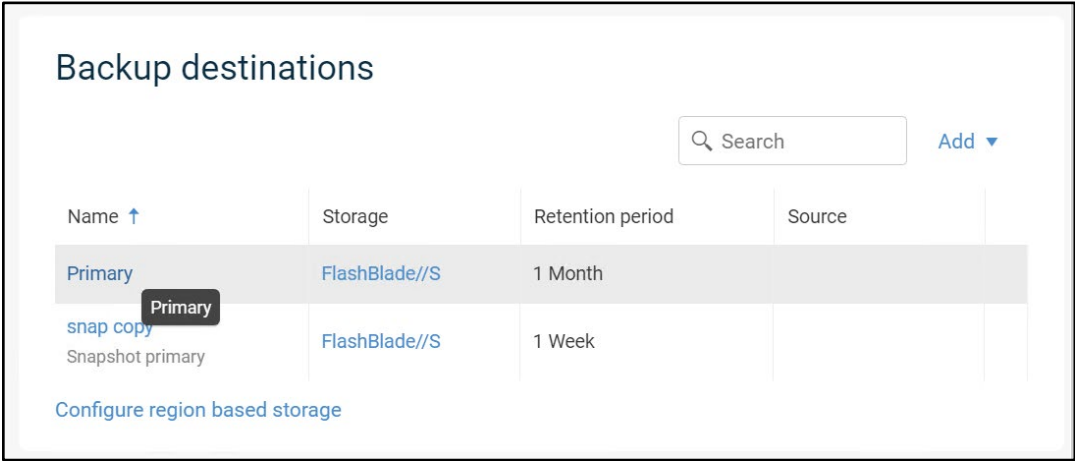


Figure 35. Backup destinations tile

In the Edit backup destination form (Figure 36), set the appropriate retention period. As explained earlier, the optimal period is double the immutability window. Continuing the previous example, if the immutability window is 14 days, Commvault retention should be set to 28 days.

IMPORTANT: The solution will function with retention set to be longer or shorter. However, you should discuss the potential operational and capacity impacts with your Pure account team before implementing.

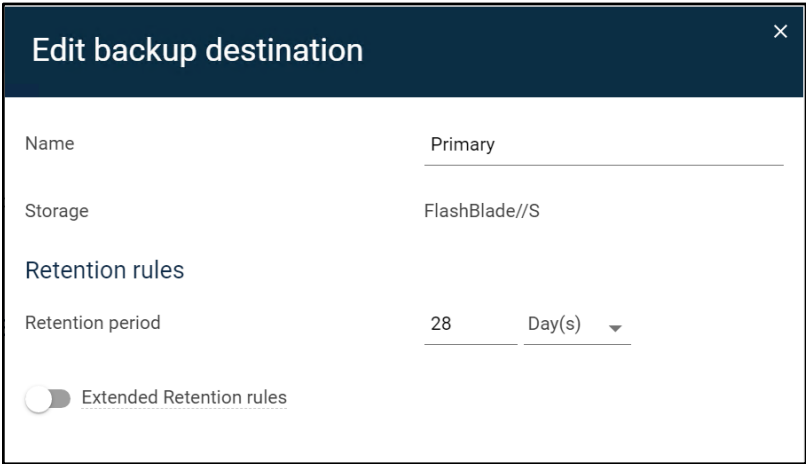


Figure 36. Editing backup retention

NOTE: You can change retention on storage policy copies using the CommCell Console.

To enable and configure Object SafeMode, contact Pure Support. There is a one-time process to designate authorized contacts for your company who are allowed to initiate SafeMode changes. For proper protection, Object SafeMode must be set to double the immutability window. If you plan to use SafeMode snapshots to protect CommServe DR backups, you can request that configuration at the same time.

IMPORTANT: You should not enable Object SafeMode until the Commvault configuration steps are complete.

Recovering with Object SafeMode

Since Object SafeMode protects your most recent backups from alteration or deletion, there is no special procedure required on FlashBlade to make the data available for recovery. Restoring the CommCell from DR backup does not have any effect on data availability. If you need to rebuild a MediaAgent or deploy a new one, you can simply configure the new MA in the existing storage pool, and backups will immediately be available for normal recovery or restore into a “clean room” environment for sanitization.

Recommended Practices for Commvault with FlashBlade//S

Table 6 lists the recommended practices when using FlashBlade//S with Commvault. See [Best Practices for Configuring FlashBlade with Commvault](#) for a full set of best practices and more details.

Recommendation	Explanation
Set 128KB deduplication block size	By default, Commvault sets cloud storage pools to use 512KB deduplication blocks. For some data sets this can significantly increase the physical footprint. Setting the block size to 128KB will ensure Commvault is writing with maximum efficiency.
Share mount path across MediaAgents	Share the object mount path between MediaAgents to simplify horizontal scaling. Each MediaAgent can directly access the FlashBlade to participate in backup and restores.
Deploy MediaAgents in grids of at least two MAs per grid	Distributing workloads across multiple MediaAgents improves parallelism and performance. Adding MediaAgents also increases resilience.
Deploy multiple DDB partitions	Partitioned deduplication allows Commvault to scale performance and capacity for a deduplication store. It also improves resilience in the case of MediaAgent downtime. Commvault defaults to creating two partitions for new DDBs, and you can create up to four.
Use NVMe for DDB and index storage	Fast internal storage is critical to the performance of Commvault deduplication, especially for synthetic full backups.
Use Commvault IntelliSnap	For primary storage with snapshot integration, such as Pure FlashArray, IntelliSnap can reduce production impact, speed backup and recovery, manage replication between arrays (where supported), and drive secondary use cases.
Run space reclamation weekly, with moderate aggressiveness	In release 11.26 , Commvault included a scheduled process to reclaim space by rewriting partially stale objects with only valid data. By default, the schedule runs reclamation daily. Most environments don't age data daily, so little is gained by checking daily for staleness. Weekly reclamation aligns well with weekly full backups and aging of backup cycles.



	<p>The aggressiveness controls how much data must be stale in an object before it is rewritten. More aggressive settings will reclaim space faster, but they will also place a heavier load on both the FlashBlade and the MediaAgent. Setting the value to 2 or 3 will have Commvault ignore objects with less than 60% or 40% stale data, respectively.</p> <p>You can edit the reclamation job settings in the system-created schedule.</p>
Avoid space reclamation with Object SafeMode	<p>Running space reclamation with an immutability solution like Object SafeMode will not yield any benefits and can increase storage consumption. Space reclamation generally addresses a buildup of stale data over time. In most cases, Object SafeMode retention periods are short enough that staleness is minimal. In addition, the stale objects are likely to not have reached the immutability age, so Commvault would not be able to delete them after writing the new objects.</p>
Avoid increasing upload thread pool limits	<p>As explained in Best Practices for Configuring Commvault with FlashBlade, Commvault uses a pool of connection threads when writing to object storage, which it scales dynamically up to a configured maximum.</p> <p>In most cases, the MA should be able to reach its maximum throughput within the default limit of 50 connections. If you are seeing 50 connections but less throughput than you expect, you should confirm that you don't have a bottleneck somewhere else in the pipeline before you increase the limit. More connections may improve performance, or they may simply add latency.</p>
Consider increasing pipeline buffers with high-bandwidth networks	<p>Increasing the number of pipeline buffers can improve data transfer between agents and MediaAgents, particularly when data streams can't fully utilize the network bandwidth. You should test thoroughly, making small increases until performance stops improving.</p>
Consider increasing look-ahead block size	<p>When reading from storage, Commvault proactively fetches additional data to reduce I/O. Increasing this look-ahead size can improve performance by as much as 20%. However, you should test the impact in your environment before implementing, as it may also reduce performance.</p>

Table 6. Recommended practices

Conclusion

For organizations that need to recover large environments extremely quickly, FlashBlade//S delivers. Restore speeds, measured with a small Commvault infrastructure at over 14TiB/hr. with room to grow to many times that rate, will not only accelerate your recovery today, but continue to support your future needs as you grow. Immutability with Object SafeMode ensures you have the data you need, when you need it, and it's ready to go on your fastest backup storage tier.

When deployed with FlashBlade//S, Commvault Backup & Recovery adds leading enterprise-level data management that takes full advantage of FlashBlade//S architecture and fast object capabilities. And for ransomware mitigation, combining Commvault with Object SafeMode adds Commvault's advanced detection and prevention features on top of SafeMode immutability to help spot an attack sooner, limit the impact, and help you recover faster.



Additional Resources

Next Steps

- Learn more about [FlashBlade//S](#).
- See how [Commvault Backup & Recovery](#) can modernize data protection in your environment.

Supporting Information

- [FlashBlade//S Data Sheet](#)
- [Best Practices for Configuring Commvault with FlashBlade](#)
- [Metallic Data Management as a Service](#)
- [Commvault Documentation](#)





About the Author

Roy Child is a Senior Solution Architect with Pure Storage, responsible for defining data recovery solutions and reference architectures for primary workloads such as Oracle, SQL, and VMware. Roy has worked in and with the data protection industry for over 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.

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

[purestorage.com](https://www.purestorage.com)

800.379.PURE

