REFERENCE ARCHITECTURE

Commvault Backup with FlashArray//C File Services

Reference architecture for FlashArray File Services with Commvault

Contents

Introduction

Ransomware attacks continue to be a rapidly growing problem. According to FinCEN, more data ransoms were paid in just the first half of 2021 than all of 2020: nearly $600 million[1]. That's after Bitdefender measured a 485% increase in number of attacks in 2020[2]. In a recent KPMG survey of 500 CEOs, 18 percent said that cybersecurity risk would be the biggest threat to their organization's growth through 2024—up from 10 percent in 2020[3]. In other words, it's more important than ever to have a backup and recovery solution that you can rely on to protect your data and get you back online quicklyin the event of a compromise.

Pure Storage® FlashArray//C and Commvault work together to reduce the risk and impact of a ransomware attack and improve overall data recovery. Commvault's proven architecture, combined with Pure's capacity-optimized flash, provide a powerful, enterprise-ready backup platform that's easy to scale. FlashArray™ SafeMode™ mitigates against attacks on your backup data by ransomware, rogue admins, and other bad actors. And your backup data is safe from forklift upgrades thanks to Evergreen Storage™.

This reference architecture is a how-to and best practices guide to assist with the design and implementation of Pure Storage FlashArray//C into Commvault backup and recovery environments. The target audience for this document includes, but is not limited to, system architects, systems engineers, IT managers, and storage administrators.

Solving Backup and Recovery Challenges

Backup and recovery have never been more important to businesses. Data continues to grow exponentially and ransomware attacks grow ever more common and sophisticated. Insurance companies may not be willing to cover payments. The ability to recover data and applications quickly can mean the difference between a temporary disruption, a major hit, or even bankruptcy. As the statistics for cyber attacks continue to rise in 2021 and beyond, it's clear you need to be able to not only restore your services, but do it quickly.

Traditional backup and recovery also benefit from FlashArray//C. All-flash performance and fast file services can help you shrink your backup window and reduce recovery time objective (RTO) for streaming restores. If you have your production data

[1] "Financial Trend Analysis," Financial Crimes Enforcement Network, October 15, 2021, https://www.fincen.gov/sites/default/files/2021-10/Financial%20Trend%20Analysis_Ransomware%20508%20FINAL.pdf

[2] "2020 Consumer Threat Landscape Report," Bitdefender, April 6, 2021. https://www.bitdefender.com/files/News/CaseStudies/study/395/Bitdefender-2020-Consumer-Threat-Landscape-Report.pdf

[3] "KPMG 2021 CEO Outlook Pulse Survey," KPMG, https://home.kpmg/xx/en/home/insights/2021/03/ceo-outlook-pulse.html

on any FlashArray model, you can use Commvault IntelliSnap to incorporate hardware snapshots and asynchronous replication into your backup routines, lessening the VM "stun" effect and backup load on your production systems. And you can make database refreshes more automated and efficient.

As IT budgets continue to be stretched, reducing total cost of ownership (TCO) is important for enabling all your technology goals. You need to get the most value possible out of your solutions. Having a single platform that can manage all your data in a small footprint can significantly lower TCO compared to point solutions for specific applications and data sets.

FlashArray//C

FlashArray//C builds upon the industry-leading success of the FlashArray platform, extending flash performance to workloads that need high capacity more than lowest latency. FlashArray//C optimizes data density through features, such as DirectFlash® managed QLC, inline data reduction, and global deduplication, making it a perfect fit as high-performance backup storage. With integrated Server Message Block (SMB) and Network File System (NFS) file services, FlashArray//C can serve multiple use cases in a very small footprint. Pure Evergreen Storage ensures your environment stays simple and modern. Non-disruptive upgrades of controllers and storage mean the end of forklift refreshes and migrations. SafeMode protection gives you the confidence your data is safe from attackers.

Figure 1: FlashArray//C

FlashArray File Services

Pure Storage FlashArray File Services brings the reliability, data reduction, and simplicity of Purity//FA to network-attached storage (NAS). Where traditional NAS is often complicated to deploy, difficult to manage, and painful to refresh, FlashArray unified all-flash storage delivers the same experience that's been delighting customers for years.

Commvault Backup & Recovery and Metallic

Commvault Backup & Recovery (Commvault) 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.

Pure and Commvault: Better Together

When you deploy Commvault with FlashArray//C File Services as its storage, you get a powerful, scalable platform. FlashArray//C adds a shared storage layer that is easy to deploy with Commvault, drives backup and recovery performance, and enhances resilience. Commvault ransomware protections and FlashArray//C SafeMode protection provide assurance against attackers destroying your backups.

Pairing the solution with FlashArray//X for primary storage makes the story even better. VMware datastores become extremely fast without complexity. Commvault IntelliSnap technology lets you leverage FlashArray snapshots to speed backup and recovery and drive secondary use cases.

Solution Architecture

Figure 2 illustrates the solution architecture. The source VMware vSphere virtual machines (VMs) reside on VMFS datastores on FlashArray//X. For VMware virtual machine backups, two or more Commvault Virtual Server Agents (VSAs) and MediaAgents (MAs) reside together on physical servers; virtual machines are also supported, but VMware SAN mode will not be available. The MAs share one or more FlashArray file systems using SMB protocol. The shared storage allows the MAs to efficiently balance load, and each MA can directly access data written by another MA, so backups are still online in the event an MA becomes unavailable.

During backups, the VSAs read the source VM data, using one of the VMware transport modes or IntelliSnap, deduplicate it, and transfer the unique blocks and metadata to the FlashArray//C volumes. You can use the same CommCell environment to protect and recover any other data type Commvault supports. You can also follow the same architecture template with Metallic, substituting on-prem gateways in place of MAs.

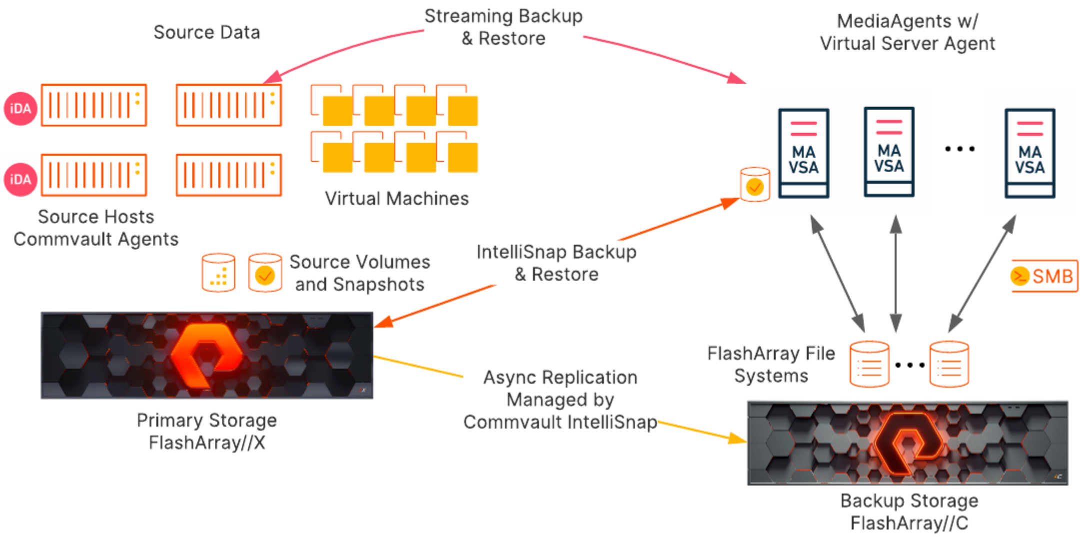
Source Data
Streaming Backup & Restore
MediaAgents w/ Virtual Server Agent
iDA
iDA
Virtual Machines
MA VSA
MA VSA
MA VSA
Source Hosts Commvault Agents
Source Volumes and Snapshots
IntelliSnap Backup & Restore
SMB
FlashArray File Systems
Primary Storage FlashArray//X
Async Replication Managed by Commvault IntelliSnap
Backup Storage FlashArray//C

Figure 2: Solution architecture

Scaling

The architecture allows you to start with one VSA/MA and easily add more as you need more backup and restore streams. Commvault automatically distributes backup load across the infrastructure. We recommend deploying at least two MAs for resilience.

Additional Sites

If you have multiple data centers, you can simply duplicate the architecture in each site and leverage Commvault's DASH Copy feature to replicate backup data efficiently between sites.

Lab Setup

Server Details

For the lab testing, we used a 4-node ESXi cluster and four physical Windows servers for combined Commvault MAs and VSAs. 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	3 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 2019 Datacenter Version 1809 Commvault Release 11.24

Table 1: Server configuration details

Server Tuning

We made the following settings changes in Windows to optimize network performance for our environment. Some or all of these may improve your performance.

TCP Congestion Provider

Windows Server 2019 introduced a new default TCP congestion provider, CUBIC. In our lab this provider is less performant than CTCP in Windows Server 2016. We changed the default providers using the following PowerShell commands.

```
Set-NetTCPSetting -SettingName "InternetCustom" -CongestionProvider CTCP
Set-NetTCPSetting -SettingName "DatacenterCustom" -CongestionProvider DCTCP
Set-NetTCPSetting -SettingName "Datacenter" -CongestionProvider DCTCP
Set-NetTCPSetting -SettingName "Internet" -CongestionProvider CTCP
```

Receive-Side Scaling (RSS) Profile

The RSS profile for a network adapter affects how its incoming traffic is distributed to different CPU cores for processing. The network adapter driver we used defaults the RSS profile to Closest, which matches the behavior of Windows Server 2008 R2. On NUMA-based systems this can be significantly slower than the NUMA and NUMAStatic profiles. You can learn more in the RSS profiles section of Choosing a Network Adapter.

You can use the Get-NetAdapterRSS PowerShell cmdlet to identify the current RSS profile.

```
Get-NetAdapterRSS | Select-Object Name,InterfaceDescription,Profile
```

We changed the RSS profile on each of our Mellanox network adapters to NUMAStatic. If any of your physical adapters use the Closest profile, you can change them with the Set-NetAdapterRSS cmdlet.

```
Set-NetAdapterRSS -Name "<adapter name from Get-NetAdapterRSS output>" -Profile NUMAStatic
```

Storage Details

A FlashArray//C60R3 hosted the volumes for the Commvault disk target. Source VMs were hosted on FlashArray//M70 arrays. The source arrays were connected using 4x10GbE iSCSI. Table 2 shows the array details.

> Important: FlashArray//C must be running Purity 6.1.14 or later and have opportunistic locking (oplocks) disabled. Contact Pure Support for assistance with upgrades and disabling oplocks.

Server Role	Array Model	Purity Releases	Physical Storage	Connectivity	Array Model
Commvault Disk Target	2x Intel Xeon E5-2697 v2 @ 2.70GHz, 24 cores total HyperThreading enabled	FlashArray//C60R3	6.1.14 (opportunistic locking disabled)	191.5TB (usable)	SMB: Single file system
Data Source (x3)	2x Intel Xeon E5-2697 v2 @ 2.70GHz, 24 cores total HyperThreading enabled	FlashArray//M70	6.1.4	21TB (usable)	iSCSI: 8 sessions per target per ESXi host

Table 2: Storage configuration 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 datastore. Each test used some part of the VM set, distributed as evenly as possible across the hosts and datastores. 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	Data Source (x96)

Table 3: Test VM details

The test data set on each VM was a blend of unique data generated on the VM and shared data copied across the VMs, with variable compressibility between 0 and 40% compressible. Between backups we added randomly selected, incompressible data from a shared pool to generate data changes that partially overlapped across the VMs.

Adding FlashArray//C File System as Commvault Disk Storage Pool

To set up the disk storage pool in Commvault, we followed this procedure for each MA.

> Note: You can work with Pure Support to enable File Services on your array.

> Note: FlashArray must be configured with at least one virtual interface and an Active Directory account. Instructions on these steps are available in the quick start guide.

Create Service Account in Active Directory

Create a service account in Active Directory to control share access. Do not make the account a member of any groups or grant it login access on any systems.

> Note: This account will be the sole owner of the share data. To protect the account from compromise, we recommend using a strong password, and a password vault if available.

Grant Service Account Write Access on MediaAgents

The service account needs to be able to write log entries on the MAs that access the network share. For each MA that will use FlashArray//C, you need to grant the service account the Modify permission on the Commvault Log Files directory (Figure 3).

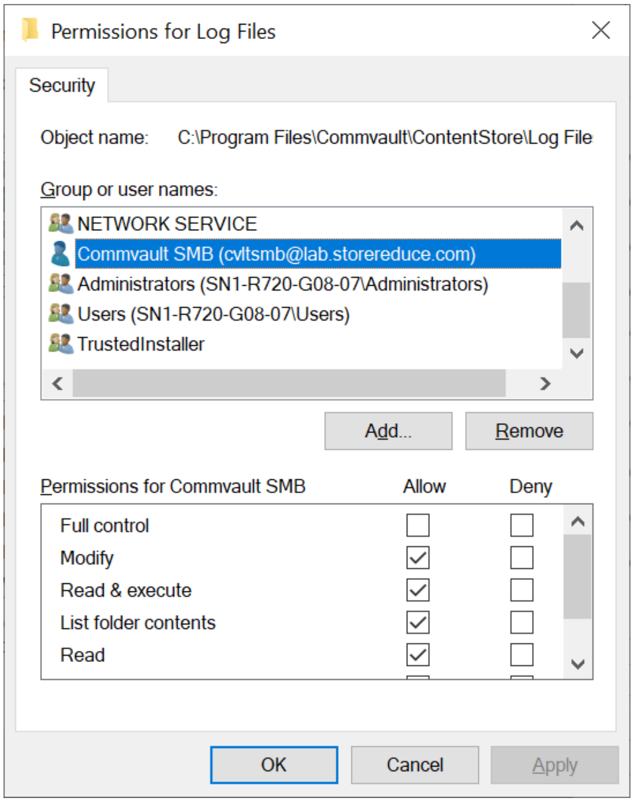

Permissions for Log Files

Object name: C:\Program Files\Commvault\ContentStore\Log File

Group or user names:

- NETWORK SERVICE
- Commvault SMB (cvltsmb@lab.storereduce.com)
- Administrators (SN1-R720-G08-07\Administrators)
- Users (SN1-R720-G08-07\Users)
- TrustedInstaller

Permissions for Commvault SMB	Allow	Deny
Full control	□	□
Modify	☑	□
Read & execute	☑	□
List folder contents	☑	□
Read	☑	□

Figure 3: Service account permissions on Log Files directory

Create and Share Storage Pool File System on FlashArray//C

1. Create an export policy on FlashArray//C that will be used for the Commvault file system (Figure 4). You should not enable access-based enumeration; this is unnecessary overhead with the access restrictions you will apply.

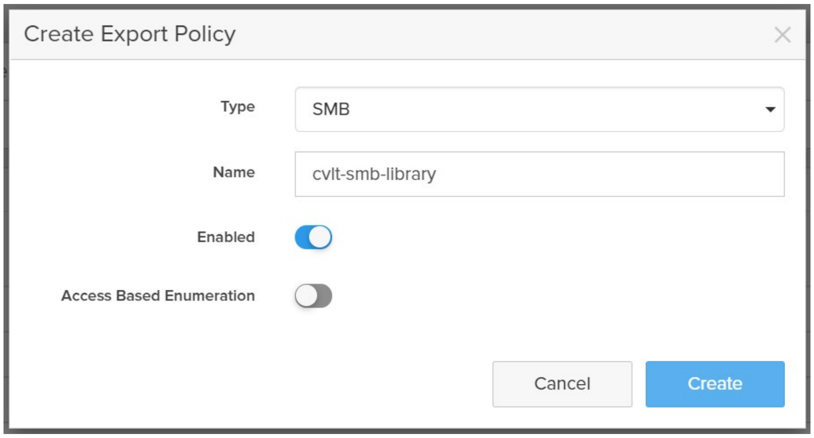

Create Export Policy

Type: SMB

Name: cvlt-smb-library

Enabled: ●

Access Based Enumeration: ○

Figure 4: Creating an export policy on FlashArray//C

2. Within the policy, create a new rule for each MediaAgent (Figure 5). You should create a separate rule for each MediaAgent's IP address so no other systems can access the shares. Do not enable anonymous access. You may enable required SMB encryption, but note that this will reduce performance by up to 20%.

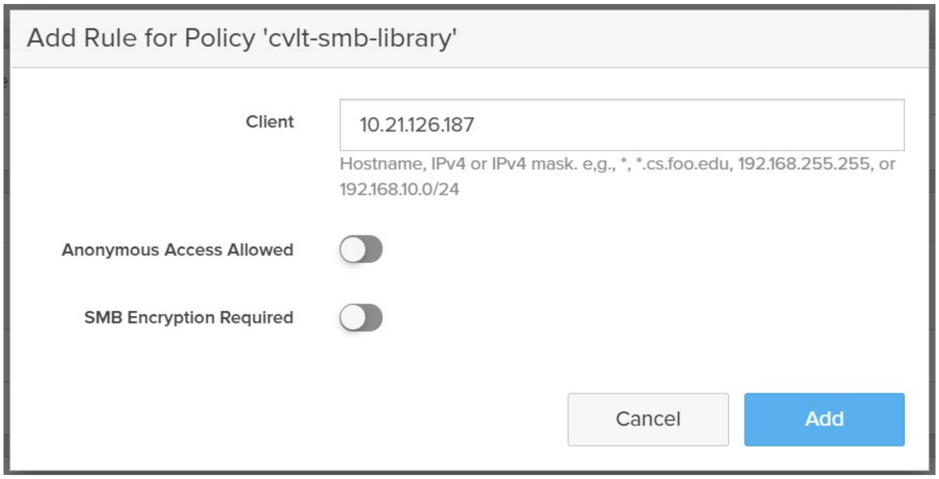

Add Rule for Policy 'cvlt-smb-library'

Client: 10.21.126.187

Hostname, IPv4 or IPv4 mask. e,g., *, *.cs.foo.edu, 192.168.255.255, or 192.168.10.0/24

Anonymous Access Allowed: ○

SMB Encryption Required: ○

Figure 5: Creating an export rule for a MediaAgent

3. Create a new file system on FlashArray//C (Figure 6).

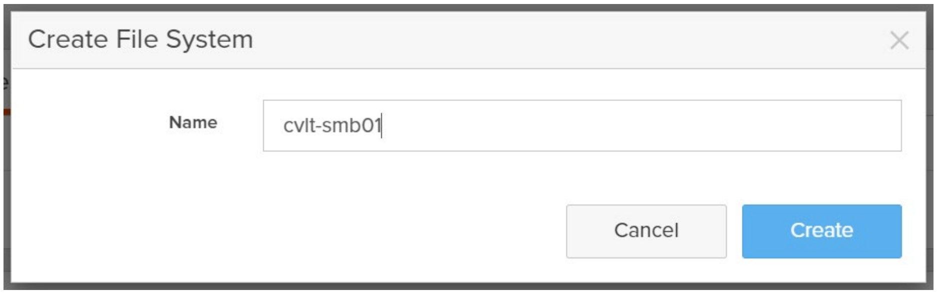

Create File System

Name: cvlt-smb01

Figure 6: Creating a file system

4. In the File Systems tile, click the file system name to access its properties. Click the Create Exports (+) button in the Directory Exports tile to share the file system. Select the file system root and SMB policy you created, then enter an export name and click the Create button (Figure 7).

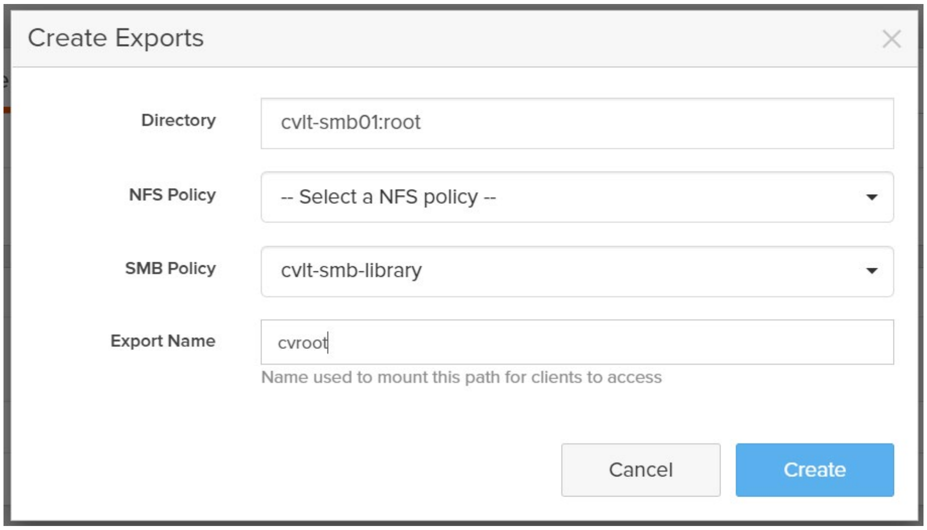

Create Exports

Directory: cvlt-smb01:root

NFS Policy: -- Select a NFS policy --

SMB Policy: cvlt-smb-library

Export Name: cvroot

Name used to mount this path for clients to access

Figure 7: Exporting the file system

5. From one of the MediaAgents, connect to the shared file system using the service account. In the security advanced properties for the share, change the owner to the service account, and grant it exclusive full control on this folder, subfolders, and files (Figure 8). All new managed directories will now inherit these permissions.

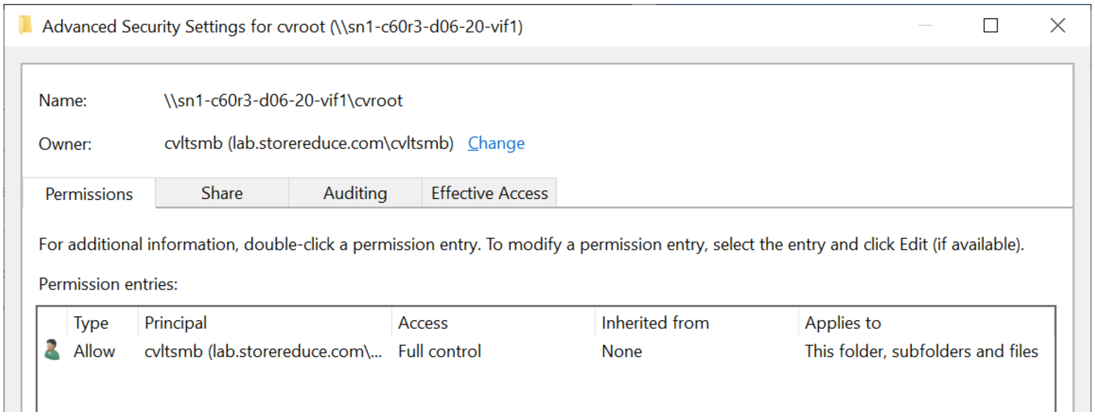

Advanced Security Settings for cvroot (\\sn1-c60r3-d06-20-vif1)

Name: \\sn1-c60r3-d06-20-vif1\cvroot

Owner: cvltsmb (lab.storereduce.com\cvltsmb) Change

For additional information, double-click a permission entry. To modify a permission entry, select the entry and click Edit (if available).

Permission entries:

Type	Principal	Access	Inherited from	Applies to
Allow	cvltsmb (lab.storereduce.com\...	Full control	None	This folder, subfolders and files

Figure 8: Setting file system permissions

6. On the FlashArray//C, create a new managed directory within the file system (Figure 9). This directory will act as the storage pool mount path in Commvault.

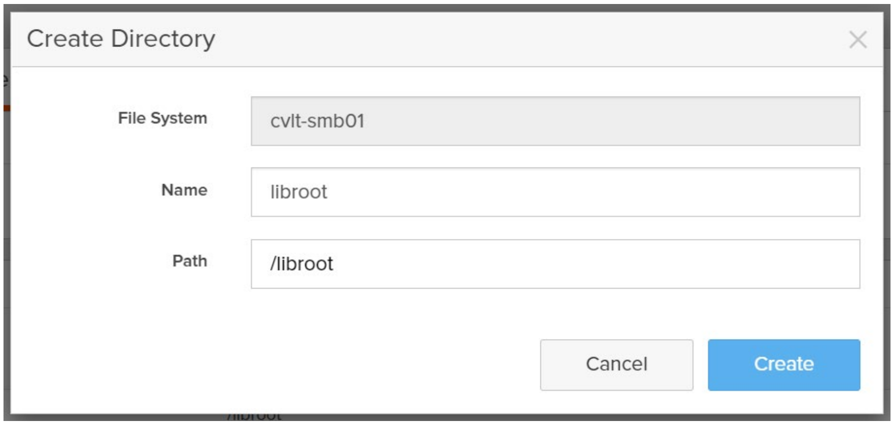

Create Directory

File System: cvlt-smb01

Name: libroot

Path: /libroot

Figure 9: Creating a managed directory

7. Export the managed directory, using the same SMB export policy (Figure 10).

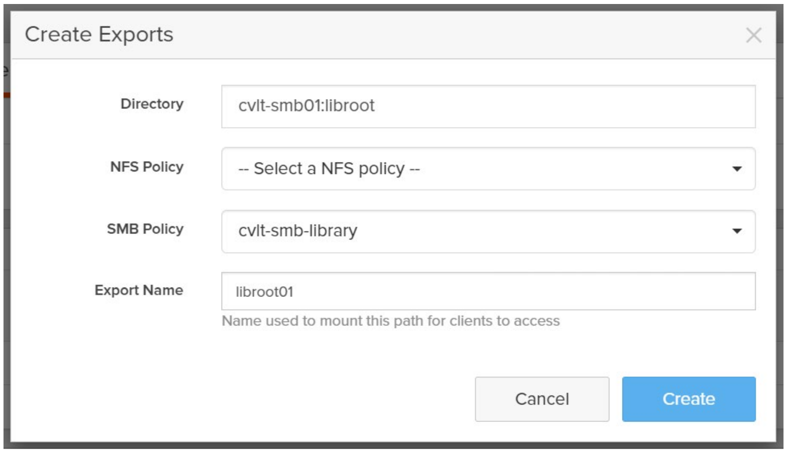

Create Exports

Directory: cvlt-smb01:libroot

NFS Policy: -- Select a NFS policy --

SMB Policy: cvlt-smb-library

Export Name: libroot01

Name used to mount this path for clients to access

Figure 10: Exporting the managed directory

8. Remove the export for the file system root by clicking its trash can icon in the Directory Exports tile (Figure 11). Click the Delete button in the confirmation dialog box that appears.

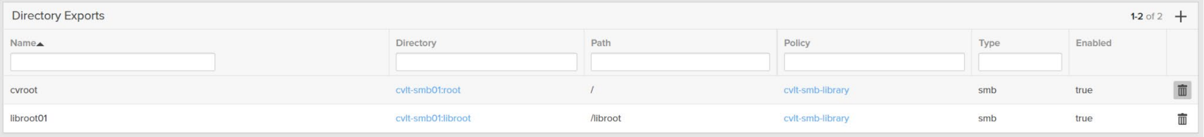

Directory Exports 1-2 of 2

Name▲	Directory	Path	Policy	Type	Enabled
cvroot	cvlt-smb01:root	/	cvlt-smb-library	smb	true
libroot01	cvlt-smb01:libroot	/libroot	cvlt-smb-library	smb	true

Figure 11: Deleting the root directory export

Create a Disk Storage Pool in Commvault

We created a disk storage pool on the managed directory using Commvault Command Center.

> Note: The appearance and steps may vary between Commvault releases.

1. In the left-hand navigation pane, expand Storage, then click **Disk**. Click the **Add** link in the upper right corner of the Disk page to open the Add storage form (Figure 12).

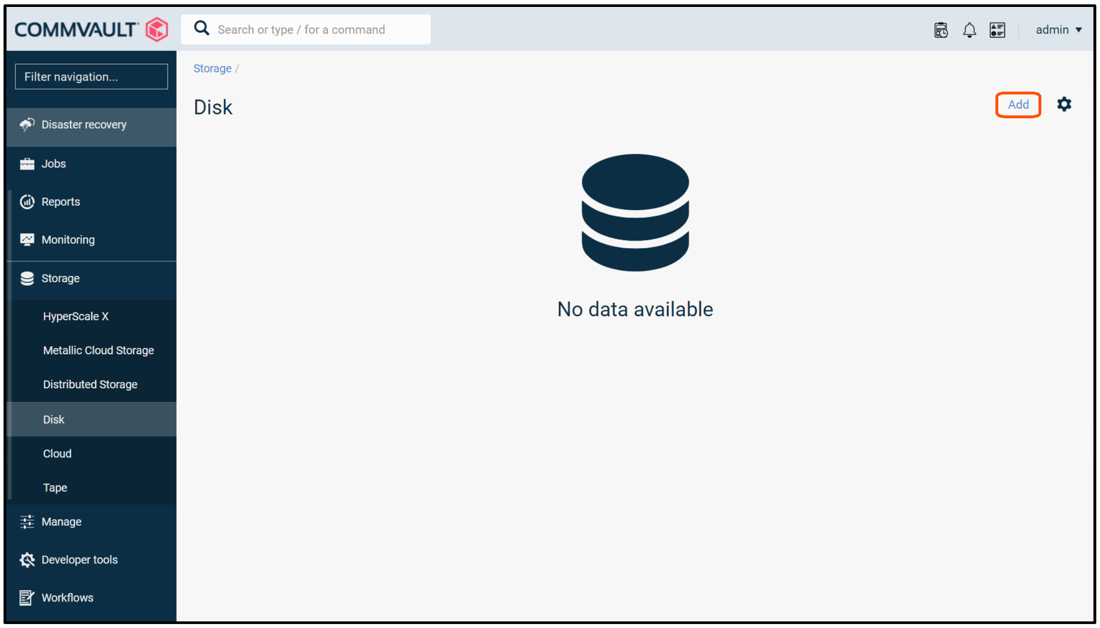

No data available

Figure 12: Adding a disk storage pool in Commvault Command Center

2. In the Add disk form, enter a display name for the storage pool. Click the **Add** link to add a MediaAgent and storage path.
3. In the Add storage form:
   a. In the MediaAgent dropdown, select the first MA you want to add to the pool.
   b. Set the Type option to Network.
   c. In the Credentials section, add a saved credential for the Active Directory service account (Figure 13).

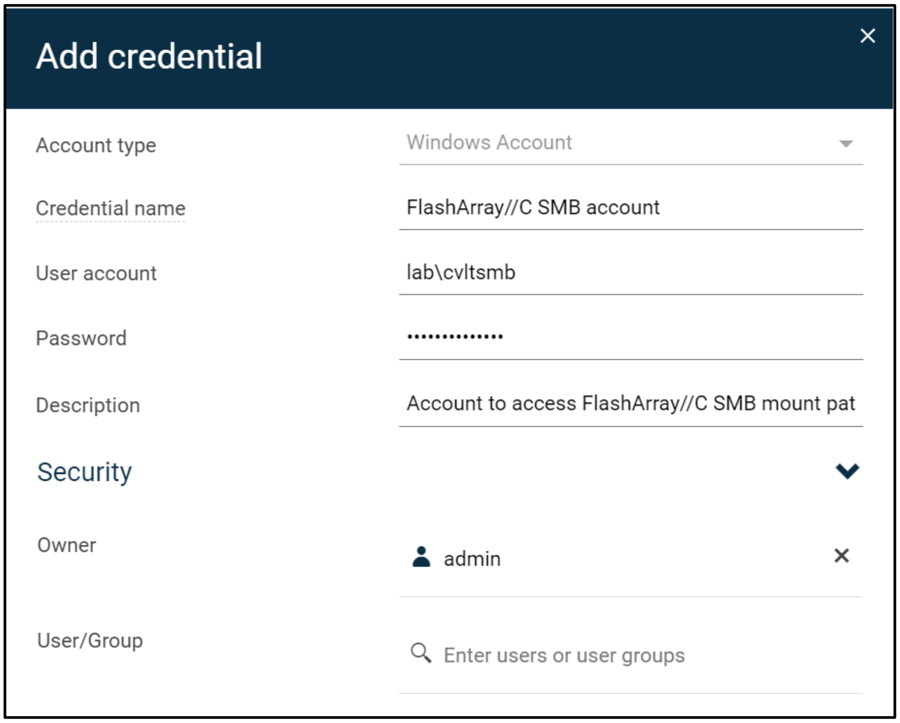

Add credential

Account type	Windows Account
Credential name	FlashArray//C SMB account
User account	lab\cvltsmb
Password	••••••••••••••
Description	Account to access FlashArray//C SMB mount pat

Security

Owner	admin
User/Group	Enter users or user groups

Figure 13: Adding a saved credential

d. In the Backup location field (Figure 14), enter the UNC path to the managed directory export on the FlashArray//C.

e. Enable the Use deduplication option.

f. In the Deduplication DB location field, enter or browse to the MediaAgent folder where you want to store the DDB. This should be on high-performance, low-latency storage We did record the data reduction, ideally using NVMe flash.

g. Click the **Save** button.

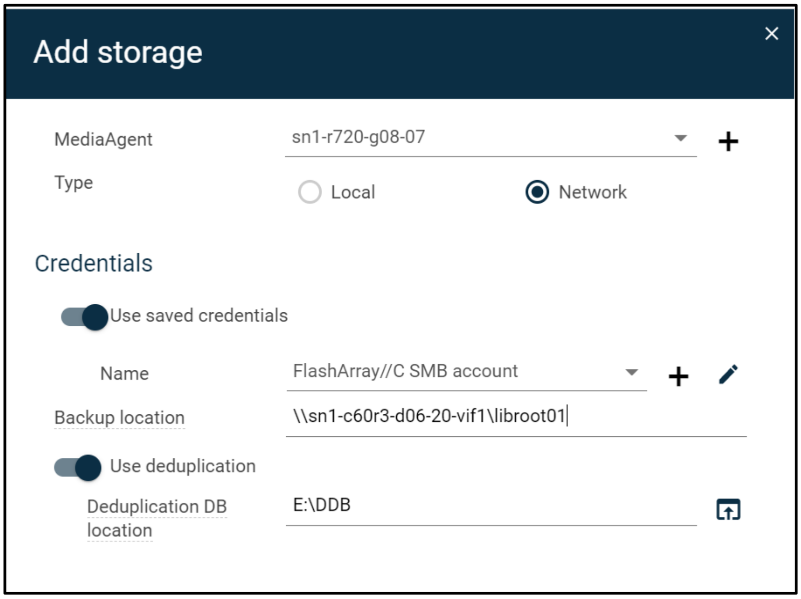

Add storage

MediaAgent	sn1-r720-g08-07
Type	○ Local ● Network

Credentials

Use saved credentials

Name	FlashArray//C SMB account
Backup location	\\sn1-c60r3-d06-20-vif1\libroot01

Use deduplication

Deduplication DB location	E:\DDB

Figure 14: Add storage form

4. For each MA that will belong to the storage pool, click the Add button and repeat step 3. Use the same UNC path and saved credential for each MA.

5. When all the MAs are listed (Figure 15), click the Save button to create the storage pool.

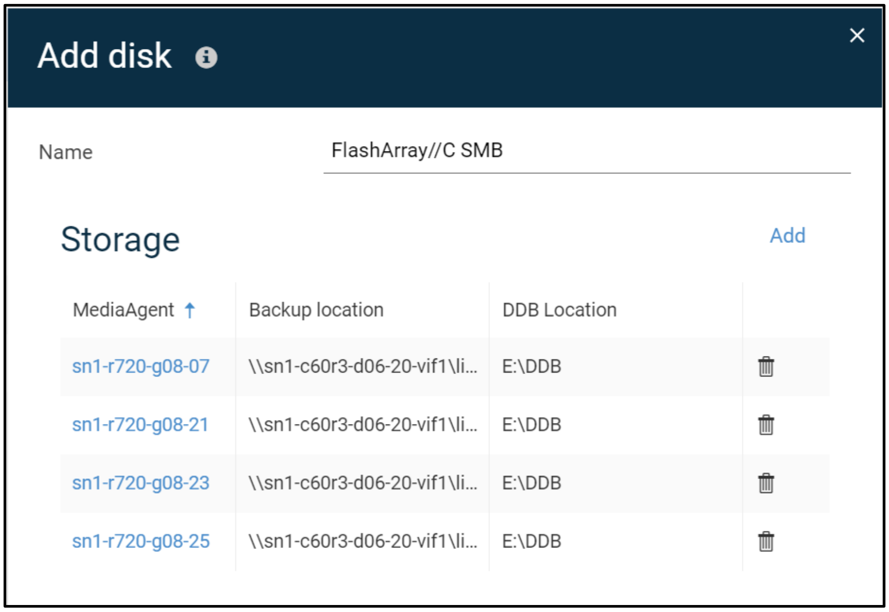

Add disk

Name	FlashArray//C SMB

Storage

MediaAgent ↑	Backup location	DDB Location
sn1-r720-g08-07	\\sn1-c60r3-d06-20-vif1\li...	E:\DDB
sn1-r720-g08-21	\\sn1-c60r3-d06-20-vif1\li...	E:\DDB
sn1-r720-g08-23	\\sn1-c60r3-d06-20-vif1\li...	E:\DDB
sn1-r720-g08-25	\\sn1-c60r3-d06-20-vif1\li...	E:\DDB

Figure 15: Add disk form with four MediaAgents

The storage pool properties and mount path will be displayed (Figure 16). From this page you can verify that the mount path shows a “Ready” status.

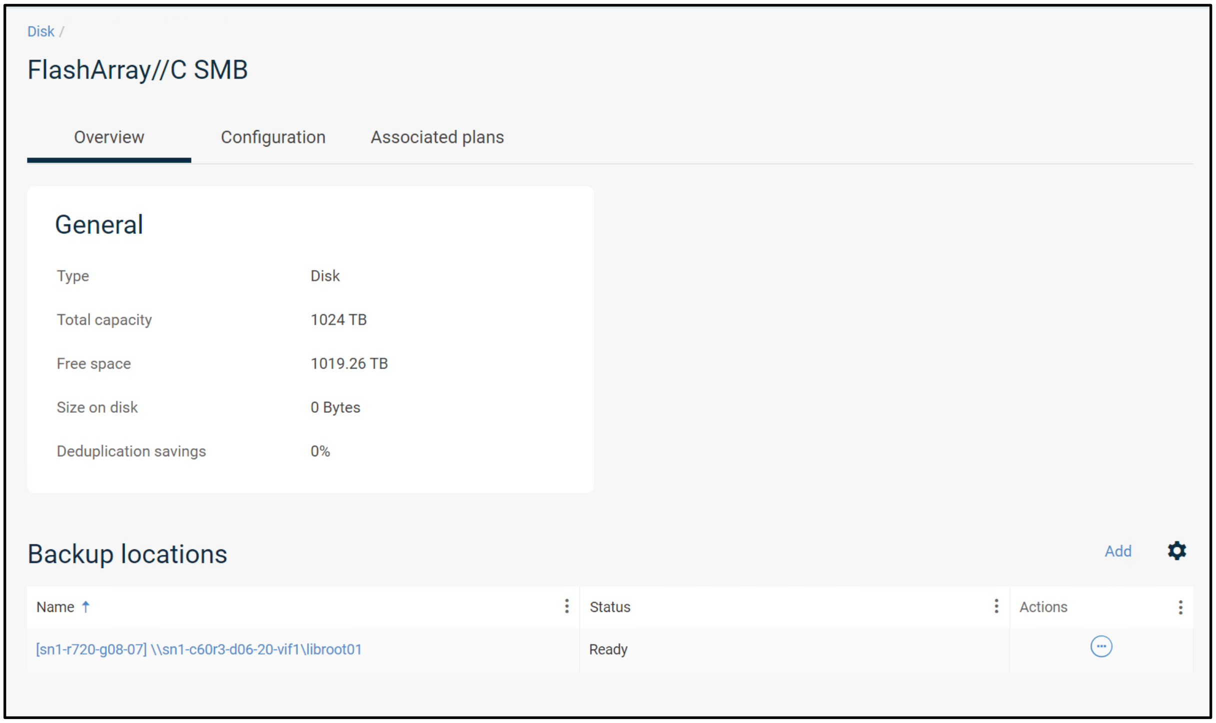

Disk /

FlashArray//C SMB

General

Type	Disk
Total capacity	1024 TB
Free space	1019.26 TB
Size on disk	0 Bytes
Deduplication savings	0%

Backup locations

Name ↑	Status	Actions
[sn1-r720-g08-07] \\sn1-c60r3-d06-20-vif1\libroot01	Ready	⋯

Figure 16: Storage pool details

You can click the backup location path name and see the status of the disk access paths (Figure 17).

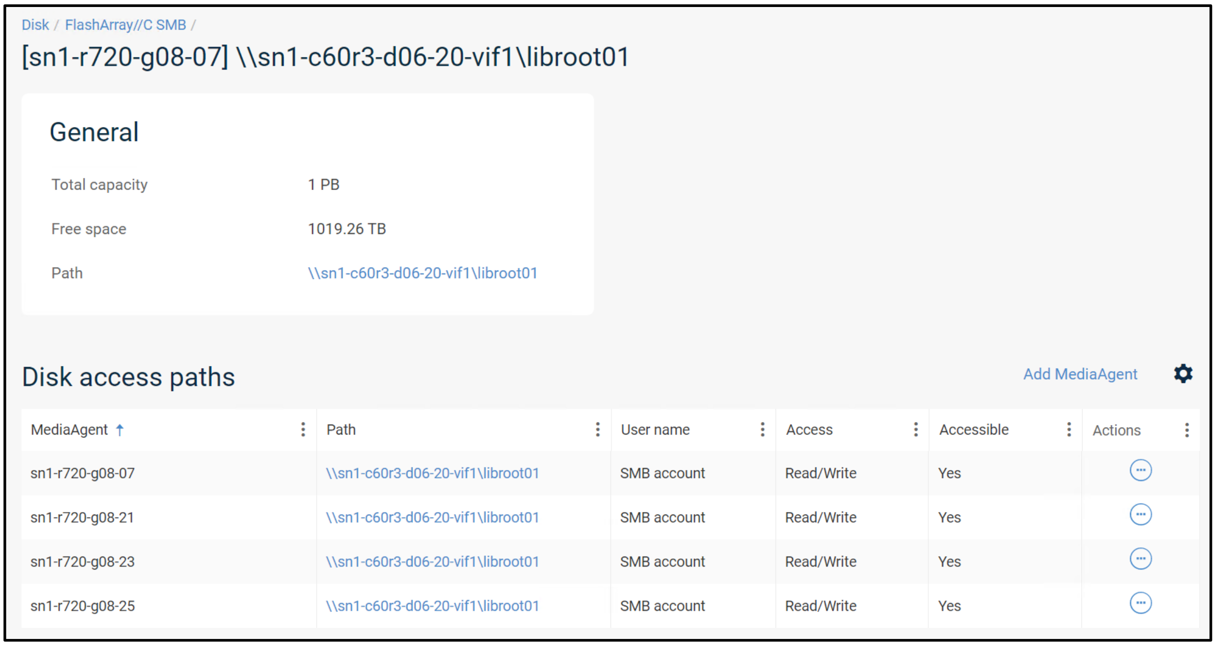

Disk / FlashArray//C SMB /

[sn1-r720-g08-07] \\sn1-c60r3-d06-20-vif1\libroot01

General

Total capacity	1 PB
Free space	1019.26 TB
Path	\\sn1-c60r3-d06-20-vif1\libroot01

Disk access paths

MediaAgent ↑	Path	User name	Access	Accessible	Actions
sn1-r720-g08-07	\\sn1-c60r3-d06-20-vif1\libroot01	SMB account	Read/Write	Yes	⋯
sn1-r720-g08-21	\\sn1-c60r3-d06-20-vif1\libroot01	SMB account	Read/Write	Yes	⋯
sn1-r720-g08-23	\\sn1-c60r3-d06-20-vif1\libroot01	SMB account	Read/Write	Yes	⋯
sn1-r720-g08-25	\\sn1-c60r3-d06-20-vif1\libroot01	SMB account	Read/Write	Yes	⋯

Figure 17: Disk access path details

Create Server Backup Plan

We added the storage pool to a server backup plan (Figure 18) to define SLAs and retention for protected data. We set the disk pool as the backup target, with retention set to 1 month. We set the recovery point objective (RPO) at 1 day, with a backup start time of 9:00 PM. We set snapshot retention to keep four snapshots and run backup copy automatically every four hours.

> Note: Because we executed all test jobs using Commvault workflow, we disabled the plan's schedule policies to avoid any interference with test jobs.

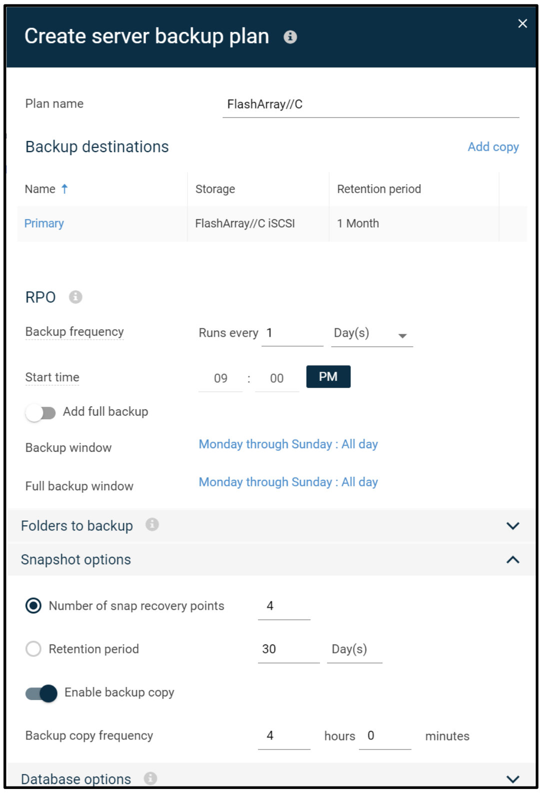

Create server backup plan

Plan name: FlashArray//C

Backup destinations — Add copy

Name ↑	Storage	Retention period
Primary	FlashArray//C iSCSI	1 Month

RPO

Backup frequency: Runs every 1 Day(s)

Start time: 09 : 00 PM

○ Add full backup

Backup window: Monday through Sunday : All day

Full backup window: Monday through Sunday : All day

Folders to backup

Snapshot options

● Number of snap recovery points: 4

○ Retention period: 30 Day(s)

● Enable backup copy

Backup copy frequency: 4 hours 0 minutes

Database options

Figure 18: Create server backup plan form

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 4 from AV scanning gave the biggest improvement in backup and recovery performance.

C:\Program Files\Commvault\ContentStore\Base\3dnfsd.exe

C:\Program Files\Commvault\ContentStore\Base\CLBackup.exe

C:\Program Files\Commvault\ContentStore\Base\ClMgrS.exe

C:\Program Files\Commvault\ContentStore\Base\cvd.exe

C:\Program Files\Commvault\ContentStore\Base\CVDistributor.exe

Table 4: 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.

Testing Details

We set out to measure several factors with Commvault deployed on FlashArray//C. We wanted to determine how backup and restore performance changed with parallel VMs, as well as what the process to recover the storage pool from SafeMode snapshots would look like.

We created a set of VM groups (Figure 19) 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.

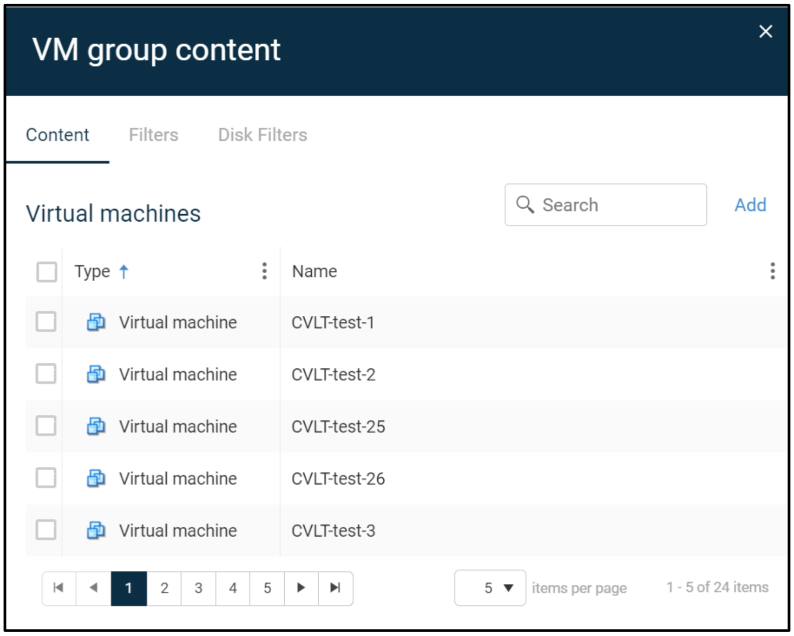

VM group content

Virtual machines

□	Type ↑	Name
□	Virtual machine	CVLT-test-1
□	Virtual machine	CVLT-test-2
□	Virtual machine	CVLT-test-25
□	Virtual machine	CVLT-test-26
□	Virtual machine	CVLT-test-3

1 - 5 of 24 items

Figure 19: Example VM group content

We enabled hardware snapshots on the VM group and set the number of readers to be higher than the number of VMs in the group (Figure 20).

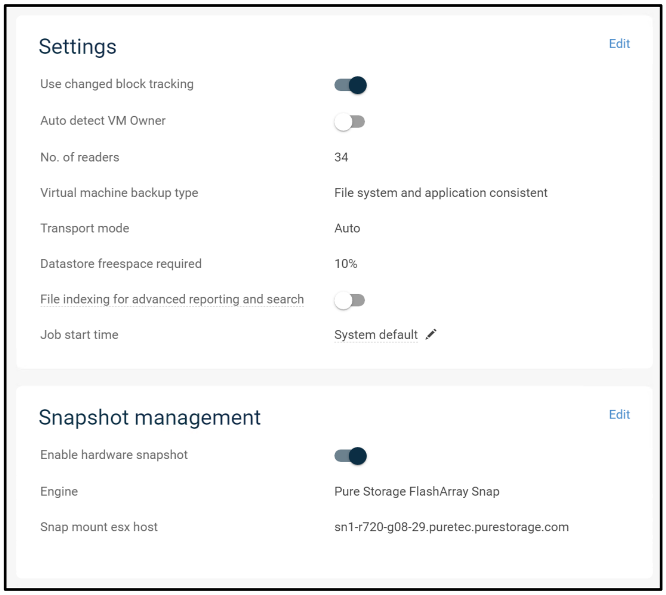

Settings

Use changed block tracking	●
Auto detect VM Owner	○
No. of readers	34
Virtual machine backup type	File system and application consistent
Transport mode	Auto
Datastore freespace required	10%
File indexing for advanced reporting and search	○
Job start time	System default ✎

Snapshot management

Enable hardware snapshot	●
Engine	Pure Storage FlashArray Snap
Snap mount esx host	sn1-r720-g08-29.puretec.purestorage.com

Figure 20: 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, we did not measure the data reduction rates on the FlashArray for each test. We did record the data reduction rate and size of data written Commvault reported 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 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 11.24.

Test Results

Backup

Scaling by Number of VMs

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 21 shows the results.

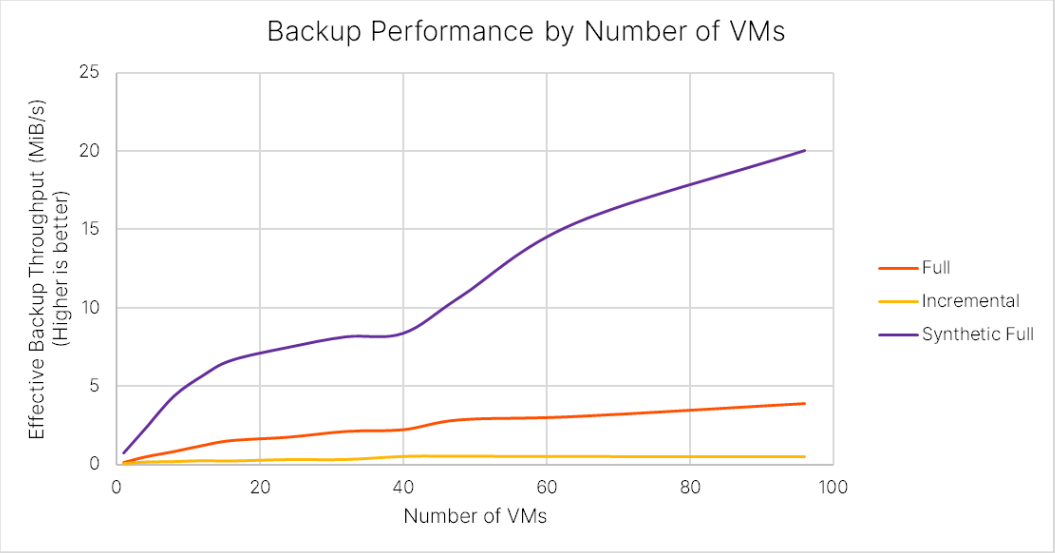

Figure 21: Backup performance scaling by number of VMs

All the backup types saw steady improvement as the number of VMs increased. Baseline full backups reached an effective 3.88GiB/s, and synthetic full backups exceeded 20GiB/s. Effective backup throughput is tied to DDB performance more than storage write speed. Actual write speeds on FlashArray peaked at around 2.8GiB/s. As noted earlier, for VM sets above 32, we divided the VMs into multiple parallel backup jobs.

Full and Synthetic Full Backups

We compared relative performance and space savings between synthetic full backups and full backups with no data change. The source data was identical between the tests for a given VM set. Average data size for this test was 82GiB per VM.

For performance, we compared the effective throughput of full and synthetic full backups for each VM set. Figure 22 shows the results.

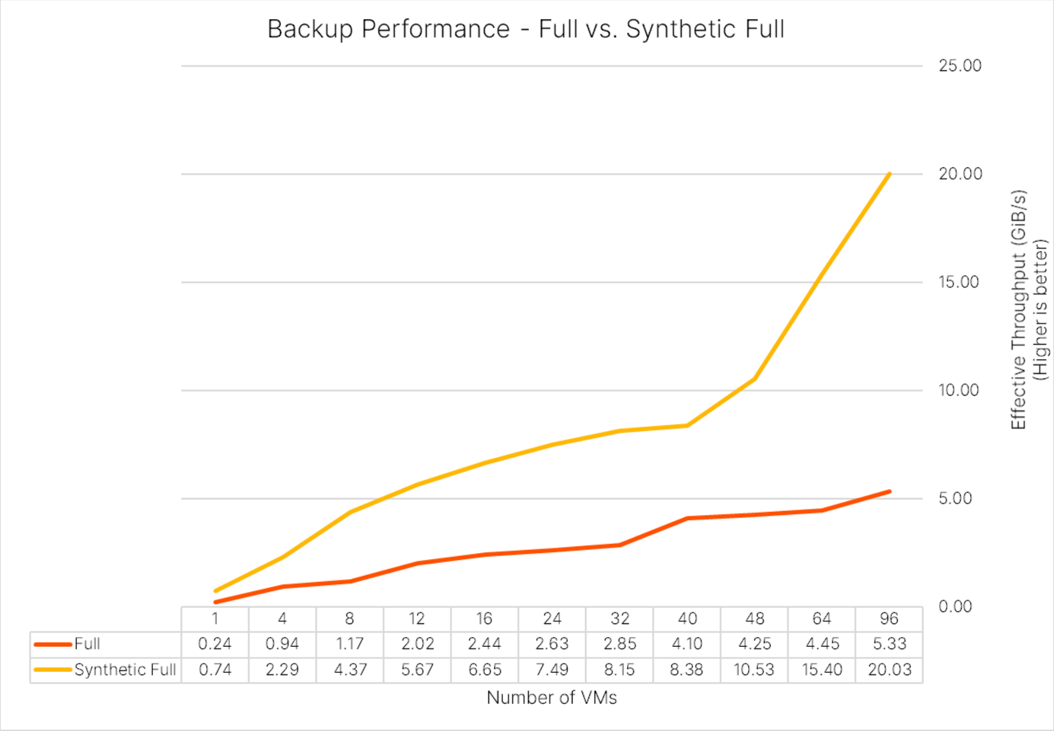

	1	4	8	12	16	24	32	40	48	64	96
Full	0.24	0.94	1.17	2.02	2.44	2.63	2.85	4.10	4.25	4.45	5.33
Synthetic Full	0.74	2.29	4.37	5.67	6.65	7.49	8.15	8.38	10.53	15.40	20.03

Figure 22: Full and synthetic full backup performance comparison

With Commvault deduplication, synthetic full backups have very little interaction with backup storage. The process reads from the index and DDB and writes a new set of pointers into the DDB and backup storage. It does not need to write any data blocks. As a result, the effective throughput of synthetic full backups was consistently triple or more compared to full backups. It's important to note that Commvault limits the number of concurrent synthetic full backups for a storage policy, so the gap will close as the number of jobs increases, but if the MediaAgents can handle the load, you can increase the number of concurrent synthetic full jobs. Performance is also heavily dependent on the DDB and index storage IOPS capabilities. SSDs and NVMe drives will enable much higher effective throughput than spinning disk.

We also measured the relative data reduction Commvault was able to achieve, and the amount of data actually written during each backup job. Figure 23 shows the results.

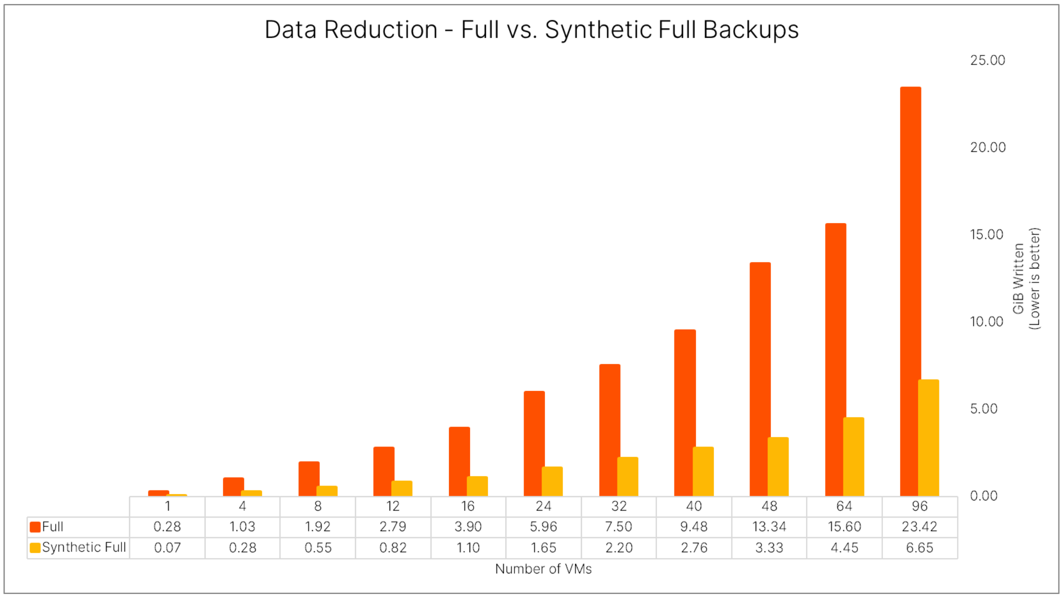

	1	4	8	12	16	24	32	40	48	64	96
🟥 Full	0.28	1.03	1.92	2.79	3.90	5.96	7.50	9.48	13.34	15.60	23.42
🟨 Synthetic Full	0.07	0.28	0.55	0.82	1.10	1.65	2.20	2.76	3.33	4.45	6.65

Figure 23: Full and synthetic full backup data reduction comparison

Both full and synthetic full backups were extremely efficient, each writing less than one percent of the source data to disk. Even so, synthetic full backups were more efficient and wrote only a third or less compared to full backups of the same data.

Between the performance and efficiency, plus zero production impact, synthetic full processing has a significant advantage over traditional full backups. We strongly recommend using an incremental forever approach with synthetic full backups wherever possible.

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 24 shows an example of the restore options.

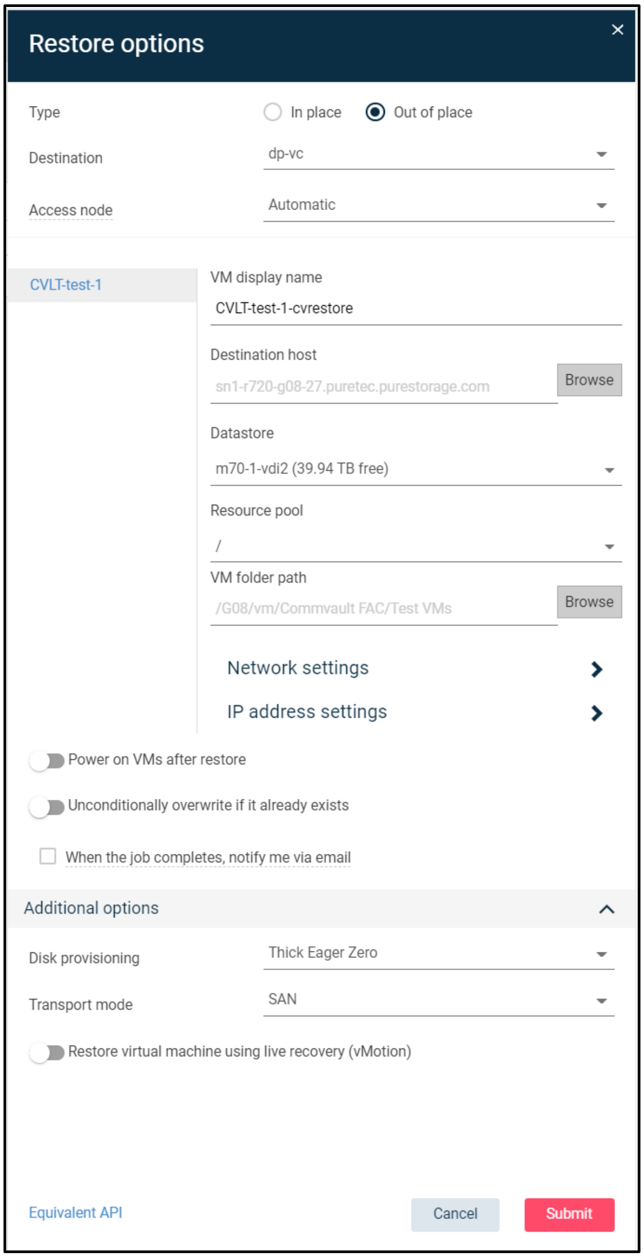

Restore options

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

Datastore: m70-1-vdi2 (39.94 TB free)

Resource pool: /

VM folder path: /G08/vm/Commvault FAC/Test VMs

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)

Figure 24: Full VM restore options

We measured throughput as the average rate at which data was transferred until all the concurrent jobs completed. Figure 25 shows the results.

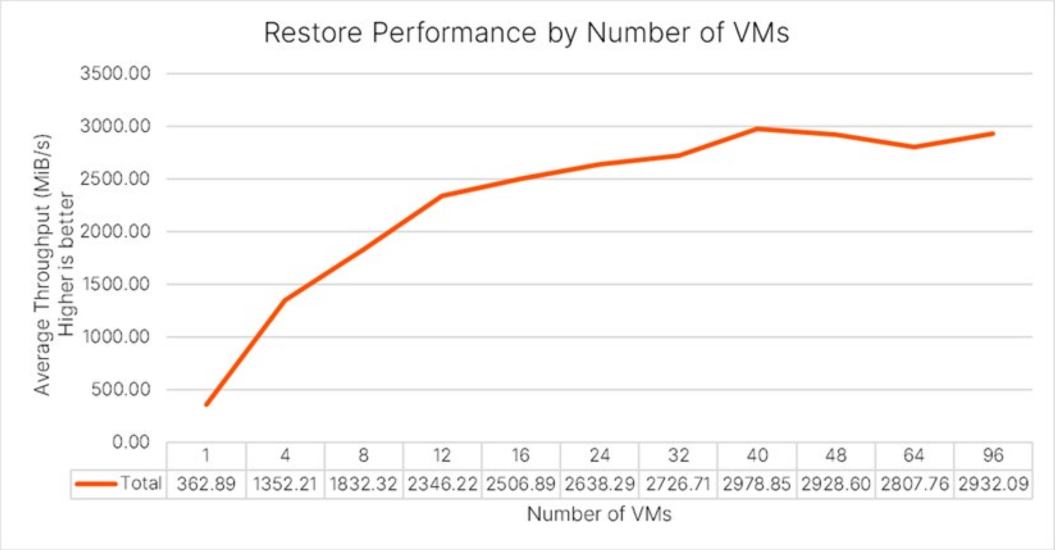

	1	4	8	12	16	24	32	40	48	64	96
Total	362.89	1352.21	1832.32	2346.22	2506.89	2638.29	2726.71	2978.85	2928.60	2807.76	2932.09

Figure 25: Restore performance scaling by number of VMs

Restore speed scaled smoothly up to 16 concurrent VMs. Between 16 and 48 VMs it still increased, but less quickly, peaking at slightly below 3GiB/s for 40 VMs. Above 40 VMs, performance remained relatively steady between 2.74GiB/s and 2.9GiB/s.

Ransomware Mitigation

Configuring SafeMode Protection

Commvault includes a number of features to detect and prevent ransomware and other attacks against your backup data. FlashArray//C includes SafeMode protections that complement Commvault's capabilities and give you an extra layer of protection against malicious or accidental deletion of your backup data. You can use FlashArray protection policies to periodically capture a consistent view of your storage pool, and SafeMode will ensure that it is available in the event an attacker does gain access and attempts to remove your data at the file system level.

FlashArray uses a two-step deletion process for removing file systems. Any file system being deleted must first be destroyed, then eradicated. When you destroy a file system, an eradication timer starts, after which the array will permanently delete the file system. You can recover the file system at any point until eradication. SafeMode prevents anyone from manually eradicating any object from the array, so the most an attacker can do is destroy file systems to disrupt service. An administrator can then recover the destroyed data and restore service.

When combined with protection policies and FlashArray managed directory snapshots, SafeMode provides an effective way to ensure you can access a previous point in time if your backup data is attacked directly. Protection policies let you schedule snapshots at a consistent time for a set of related directories and define retention settings for them. SafeMode will prevent an attacker from deleting the snapshots before the retention period expires. Between snapshot schedules and retention and eradication prevention, SafeMode further hardens Commvault's layered ransomware defense.

Protection Policy

You can use a protection policy to manage regular snapshots for your Commvault storage pool directories. To create a protection policy, in the FlashArray GUI, navigate to the Protection page, then click the **Policies** tab. In the Snapshot Policies tile, click the **Create Policy (+)** button. Enter a name for the policy, then click the **Create** button (Figure 26).

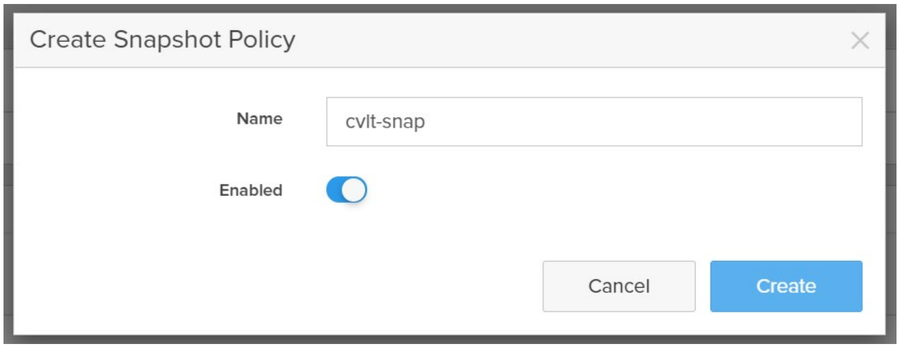

Create Snapshot Policy

Name: cvlt-snap

Enabled: ●

Figure 26: Creating a FlashArray//C protection policy

You can add one or more rules to the policy to control snapshot frequencies and retention. In the Snapshot Policies tile, click the policy name to access its settings. In the Rules tile, click the **Create Rule (+)** button. In the Add Rule for Policy dialog box (Figure 27), enter the frequency in the Create 1 snapshot every field. We recommend creating snapshots at least daily. Creating snapshots more frequently reduces potential data loss in an attack, but it also increases storage consumption and system load.

Enter the retention period you want in the And keep for field. Longer retention gives you more recovery points to choose from after an attack, but the recovery effort increases the farther back you go. We recommend keeping snapshots between five and seven days.

Enter a unique, identifiable client name in the Client Name field. The snapshots the policy creates will be named <client name>.<suffix>, where client name is the value you configure in the policy and suffix is a numeric suffix the array assigns. Click the **Add** button to create the rule. If you want more than one frequency, you can click the **Create Rule** button again and add new rules.

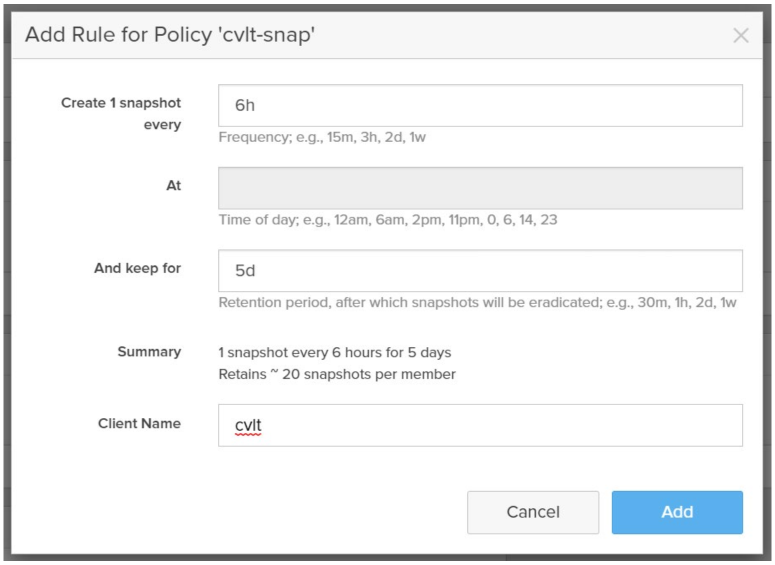

Add Rule for Policy 'cvlt-snap'

Create 1 snapshot every: 6h
Frequency; e.g., 15m, 3h, 2d, 1w

At:
Time of day; e.g., 12am, 6am, 2pm, 11pm, 0, 6, 14, 23

And keep for: 5d
Retention period, after which snapshots will be eradicated; e.g., 30m, 1h, 2d, 1w

Summary: 1 snapshot every 6 hours for 5 days
Retains ~ 20 snapshots per member

Client Name: cvlt

Figure 27: Adding a protection policy rule

You must add managed directories as members to the policy for it to take effect. To add a member, in the Members tile, click the menu icon, then select **Add Member**. In the Add Members dialog box that appears (Figure 28), select all Commvault managed directories, then click the **Add** button. You do not need to select the file system root.

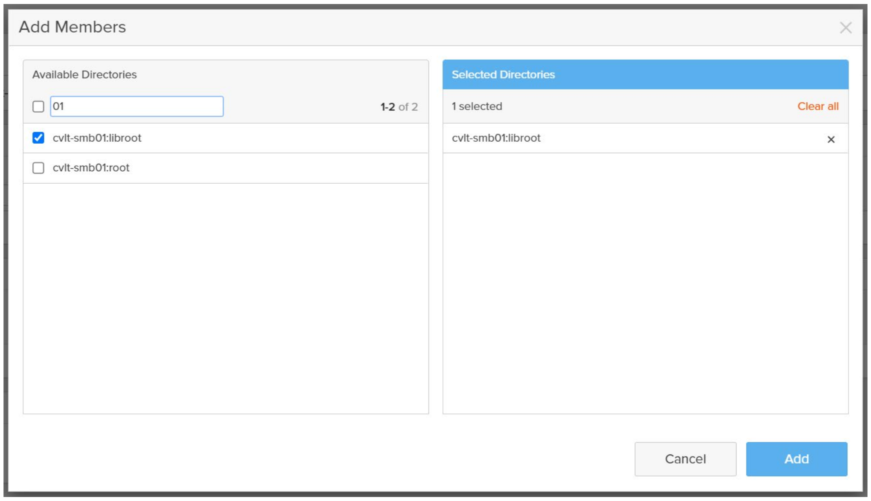

Add Members

Available Directories

□ 01 — 1-2 of 2

☑ cvlt-smb01:libroot

□ cvlt-smb01:root

Selected Directories

1 selected — Clear all

cvlt-smb01:libroot ×

Figure 28: Adding members to a protection policy

Only backups that complete before the snapshots are captured will be recoverable. You should schedule snapshots to occur at the end of your typical backup window to maximize your protection and storage efficiency. You may take more frequent snapshots, but remember that this will require more storage and will make it more complicated to determine what backups are valid for recovery after an attack.

Recovering with SafeMode

This section details the process to recover from an attack using FlashArray snapshots protected by SafeMode. Before beginning, ensure your CommCell is online. If you need to restore from Commvault DR backups or rebuild any MediaAgents, do so before beginning this procedure.

> Important: If you need to recover jobs that exist within a snapshot but have aged from Commvault, you must first perform a DR recovery on the CommCell, using the appropriate DR backup.

Contact Pure Support to Increase Eradication Timer

After a ransomware attack, you should contact Pure Support at your earliest opportunity. You should work with Support to increase the eradication timer to prevent snapshots you might need from being removed from the system. Having Support engaged early is also helpful for faster response should any issues arise during recovery.

Disable Protection Policy

Before making any changes, you should disable the protection policy. This will prevent the system from capturing any new snapshots that capture invalid or compromised data. To disable a policy, navigate to the policy settings, then click the menu icon in the Details tab and select **Edit**. Click the **Enabled** switch to disable the policy, then click the **Save** button (Figure 29).

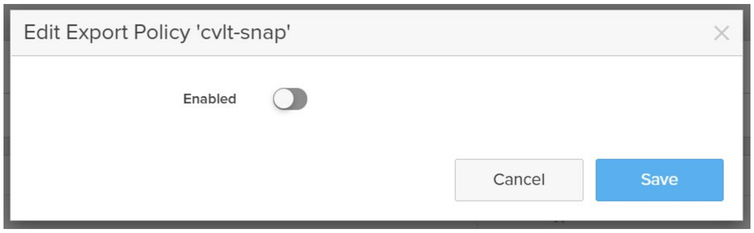

Edit Export Policy 'cvlt-snap'

Enabled ○

Figure 29: Disabling the protection policy

Download Recovery Script

Download the Execute-CvSafeModeRecovery.ps1 PowerShell script from the Pure Code Modern-Data-Protection repository and copy it to one of your MediaAgents. You will use the script later to quickly make the data from the snapshot available to Commvault.

Stop Commvault Services

To prevent storage changes during recovery, you must stop all Commvault services on each MediaAgent that connects to FlashArray//C.

To stop Commvault services, launch Commvault Process Manager. Select the Services tab (Figure 30). Ensure that the All Services item is highlighted, then click the **Stop** button (red square). Wait for all the services to show red squares as their status icons.

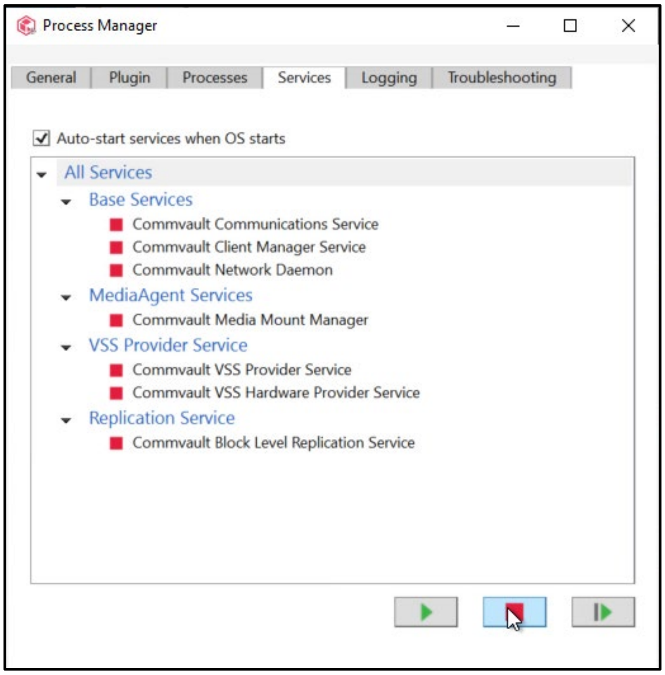

☑ Auto-start services when OS starts

- All Services
  - Base Services
    - 🟥 Commvault Communications Service
    - 🟥 Commvault Client Manager Service
    - 🟥 Commvault Network Daemon
  - MediaAgent Services
    - 🟥 Commvault Media Mount Manager
  - VSS Provider Service
    - 🟥 Commvault VSS Provider Service
    - 🟥 Commvault VSS Hardware Provider Service
  - Replication Service
    - 🟥 Commvault Block Level Replication Service

Figure 30: Stopping Commvault services in Process Manager

Enable Symbolic Link Evaluation on MediaAgents

By default, Windows will not evaluate symbolic links that reside on or reference remote paths. You must use the fsutil tool to enable this behavior on each MediaAgent that accesses the FlashArray//C disk target.

```
fsutil behavior set symlinkEvaluation R2R:1
```

Create Managed Directories

> Important: To prevent failures due to data inconsistency, you must perform this procedure for all managed directories where Commvault stores backup data.

You must create a new managed directory that will hold the backup data after recovery. You can create the directory in the original file system or a separate one, but we recommend using the original file system. The new directory will inherit the ACL you already set on the file system root.

Remap Managed Directory Exports

To ensure Commvault is writing data to the new managed directory, you must export the original managed directory under a new name, then export the new directory as the original export name.

For example, assume you have a managed directory named libroot, which is exported as libroot01. You could create a new managed directory named libroot-restored. You would delete the libroot01 export, then export the libroot-restored directory as libroot01. You could then export libroot as libroot01-old (Figure 31).

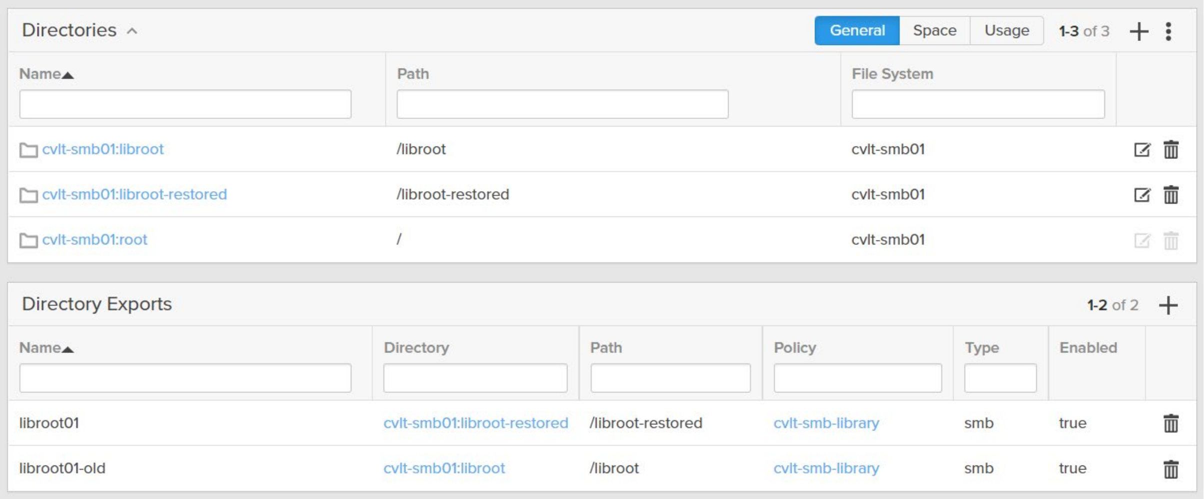

Directories

Name	Path	File System
cvlt-smb01:libroot	/libroot	cvlt-smb01
cvlt-smb01:libroot-restored	/libroot-restored	cvlt-smb01
cvlt-smb01:root	/	cvlt-smb01

Directory Exports

Name	Directory	Path	Policy	Type	Enabled
libroot01	cvlt-smb01:libroot-restored	/libroot-restored	cvlt-smb-library	smb	true
libroot01-old	cvlt-smb01:libroot	/libroot	cvlt-smb-library	smb	true

Figure 31: Remapped directory exports

Identify the Snapshot to Recover

You will need to decide what directory snapshot to use as your recovery source based on the data you need to restore. There are many criteria you can consider, such as backup completion times, snapshot sizes, analytics, and forensic investigation. You can compare the snapshot creation timestamps to your target recovery point to determine which snapshot best fits your needs.

> Important: If you need to recover jobs that exist in the snapshot but have aged from Commvault, you must first perform a DR recovery on the CommCell from the appropriate DR backup.

Extend Snapshot Retention

Before you perform recovery, you should extend retention on the snapshot you will use as your source. Otherwise FlashArray could remove the snapshot before you finish recovery. For snapshots created by a protection policy, you must decouple the snapshot from the policy before you can change the retention.

```
puredir snapshot setattr --policy ‘’ <snapshot name>
puredir snapshot setattr --keep-for <retention length> <snapshot name>
```

For example, to extend retention to 14 days for a snapshot named cvlt-smb01:libroot.cvlt.245, you would run the following commands.

```
puredir snapshot setattr --policy ‘’ cvlt-smb01:libroot.cvlt.245
puredir snapshot setattr --keep-for 14d cvlt-smb01:libroot.cvlt.245
```

Execute Recovery Script

Execute the script to create mappings in the new file system for all existing chunk files to the snapshot of your choice. Once the script completes, you can begin restoring data. You can choose to have the script begin right away to copy files from the snapshot to the new managed directory, or you can rerun the script later to copy the files. Either way, deduplication on FlashArray//C will prevent the data copy from consuming additional physical storage.

```
Execute-CVSafeModeRecovery.ps1 [-SnapshotShare] <String> [-LinkShare] <String> [-CopyOnly] [[-
LogFile] <String>] [[-SummaryFile] <String>] [<CommonParameters>]
```

Continuing the example, assume the DNS name for the FlashArray virtual interface is “fa.” You would run the following command to process the failover.

```
Execute-CVSafeModeRecovery.ps1 -SnapshotShare \\fa\libroot-old -LinkShare \\fa\libroot
```

The script runs in two phases. The first phase will create symbolic links in the new managed directory that point to the snapshot in the old directory. The links are laid out in a way that allows Commvault to read existing data from the snapshot but write new data to the new directory. Figure 32 shows an example of the links that are created.

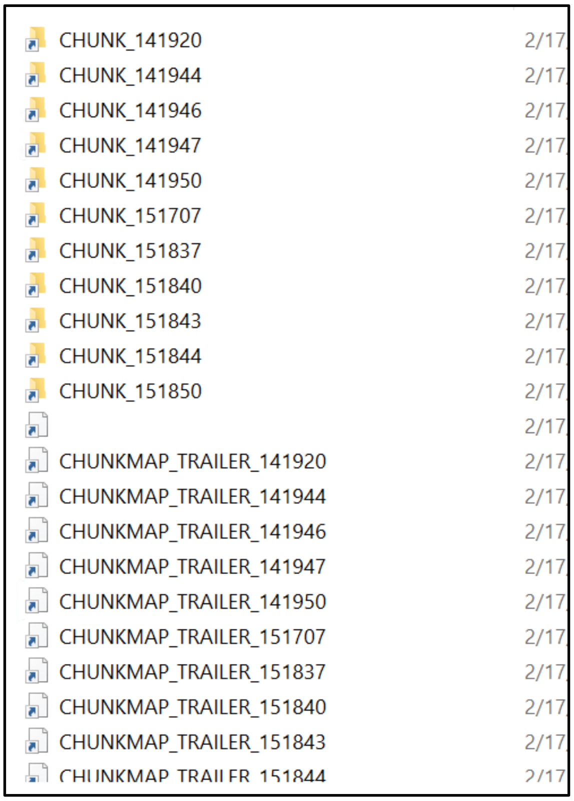

CHUNK_141920	2/17
CHUNK_141944	2/17
CHUNK_141946	2/17
CHUNK_141947	2/17
CHUNK_141950	2/17
CHUNK_151707	2/17
CHUNK_151837	2/17
CHUNK_151840	2/17
CHUNK_151843	2/17
CHUNK_151844	2/17
CHUNK_151850	2/17
	2/17
CHUNKMAP_TRAILER_141920	2/17
CHUNKMAP_TRAILER_141944	2/17
CHUNKMAP_TRAILER_141946	2/17
CHUNKMAP_TRAILER_141947	2/17
CHUNKMAP_TRAILER_141950	2/17
CHUNKMAP_TRAILER_151707	2/17
CHUNKMAP_TRAILER_151837	2/17
CHUNKMAP_TRAILER_151840	2/17
CHUNKMAP_TRAILER_151843	2/17
CHUNKMAP_TRAILER_151844	2/17

Figure 32: Symbolic links

The second phase copies the files and directories from the snapshot to the new directory and removes the links. You have the option to skip this phase so you can start restoring systems right away without impacting performance.

> Note: We recommend against copying data at the same time Commvault is restoring data.

Start Commvault Services

Once the script has completed and the snapshot data is available, start the Commvault services on all MediaAgents. On each MediaAgent, launch Commvault Process Manager (Figure 33). Select the Services tab. Ensure that the All Services item is highlighted, then click the start button (green arrow). Confirm that all services start successfully before moving on.

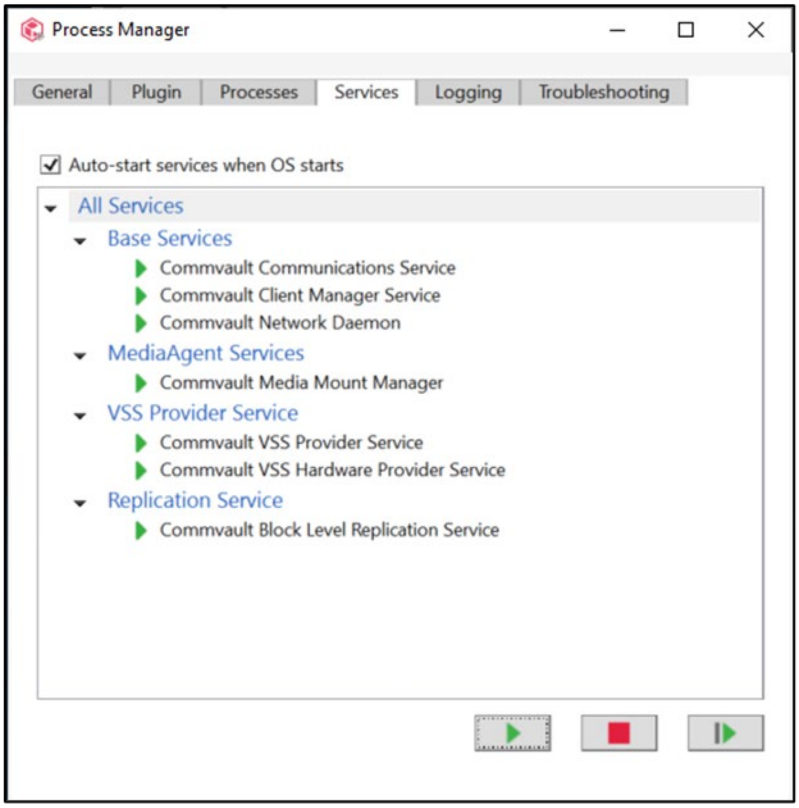

☑ Auto-start services when OS starts

- All Services
  - Base Services
    - Commvault Communications Service
    - Commvault Client Manager Service
    - Commvault Network Daemon
  - MediaAgent Services
    - Commvault Media Mount Manager
  - VSS Provider Service
    - Commvault VSS Provider Service
    - Commvault VSS Hardware Provider Service
  - Replication Service
    - Commvault Block Level Replication Service

Figure 33: Starting Commvault services

Seal Deduplication Databases

You should seal all your deduplication databases after a data attack. This will prevent Commvault from creating any new references to potentially lost unique blocks, ensuring you don't inadvertently put new data at risk. DDB sealing is performed in the CommCell Console. To seal a DDB:

1. Expand Storage Resources, then Deduplication Engines. Expand the appropriate engine.
2. For each database, right-click the database and select All Tasks, then Seal Deduplication Database.
3. Click the Yes button on the Confirm Seal Deduplication Database dialog box.
4. In the Enter Confirmation text dialog box, enter the specified text, then click the OK button.

Disable Backups and Data Aging

Disabling backups will prevent resource contention that can slow down recovery. Preventing data aging prevents Commvault from trying to delete data you may need for recovery or causing failures that might slow you down. You can disable both backups and data aging using the CommCell Console. Open the CommCell Properties dialog (Figure 34), then click the Activity Control tab. Deselect the Enable Data Management and Enable Data Aging checkboxes, then click the OK button to commit the change.

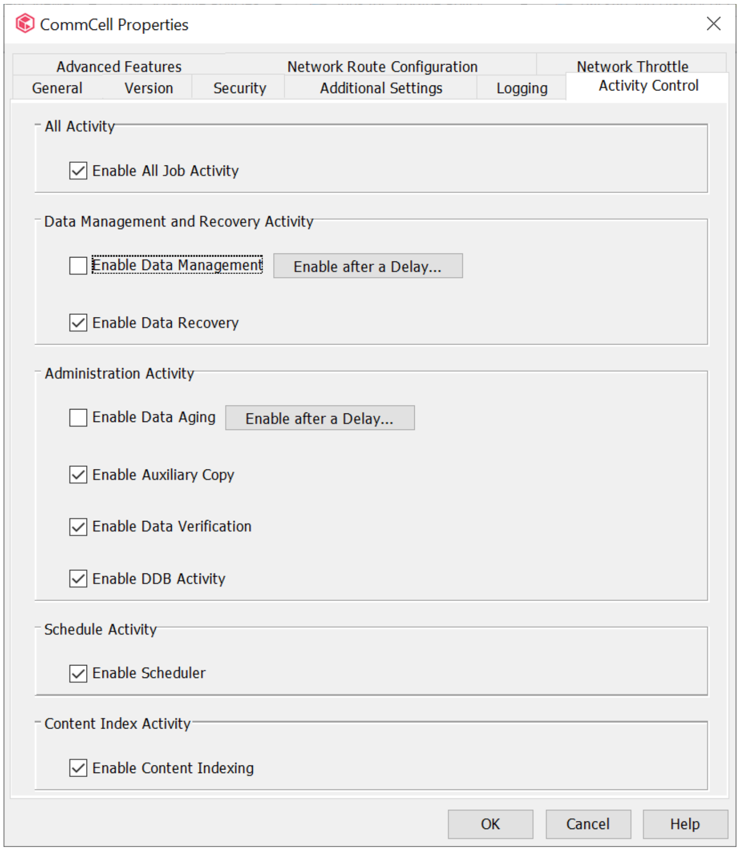

All Activity

☑ Enable All Job Activity

Data Management and Recovery Activity

□ Enable Data Management

☑ Enable Data Recovery

Administration Activity

□ Enable Data Aging

☑ Enable Auxiliary Copy

☑ Enable Data Verification

☑ Enable DDB Activity

Schedule Activity

☑ Enable Scheduler

Content Index Activity

☑ Enable Content Indexing

Figure 34: Disabling backups through Commvault CommCell Console

Restore Data from Backups Prior to Snapshot

Each FlashArray//C snapshot represents a time window in Commvault. You should assume that any backup jobs that completed or aged after the snapshot was captured will fail to recover.

To restore data, simply follow the normal process for the application type to select the data. Make sure to choose a recovery point from before the snapshot was captured (Figure 35).

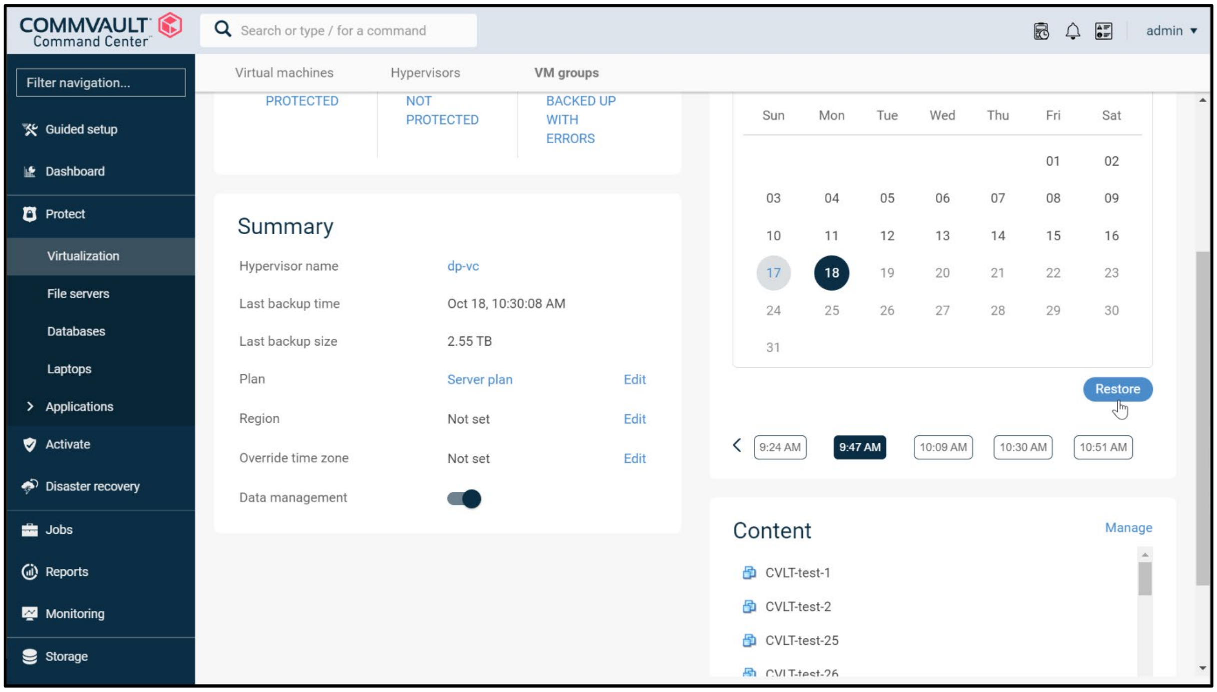

Summary

Hypervisor name	dp-vc	
Last backup time	Oct 18, 10:30:08 AM	
Last backup size	2.55 TB	
Plan	Server plan	Edit
Region	Not set	Edit
Override time zone	Not set	Edit
Data management		

Figure 35: Selecting an earlier recovery point

Run Data Verification for All Deduplication Databases

Once you have completed critical recoveries, you should run full data verification against all the DDBs. This will identify and flag inconsistencies to help prevent failures during restore and auxiliary copy jobs.

> **Important:** Data verification will generate significant CPU load on the MediaAgents housing the DDB partitions and high I/O load on the FlashArray//C and DDB storage. We recommend against running more than two or three concurrent verification jobs.

Data verification is launched through the CommCell Console. To run data verification:

1. Expand **Storage Resources**, then **Deduplication Engines**. Expand the appropriate engine.
2. For each database, right-click the database and select **Run Data Verification** (Figure 36).

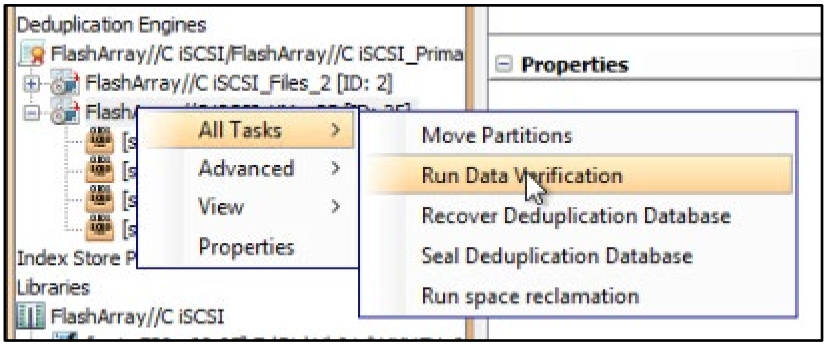

Figure 36: Running data verification

3. In the Data Verification for dialog box (Figure 37), deselect the **Run Incremental Verification** checkbox. Select the **Verification of existing jobs on disk and deduplication database** option. Click the **OK** button to begin the process.

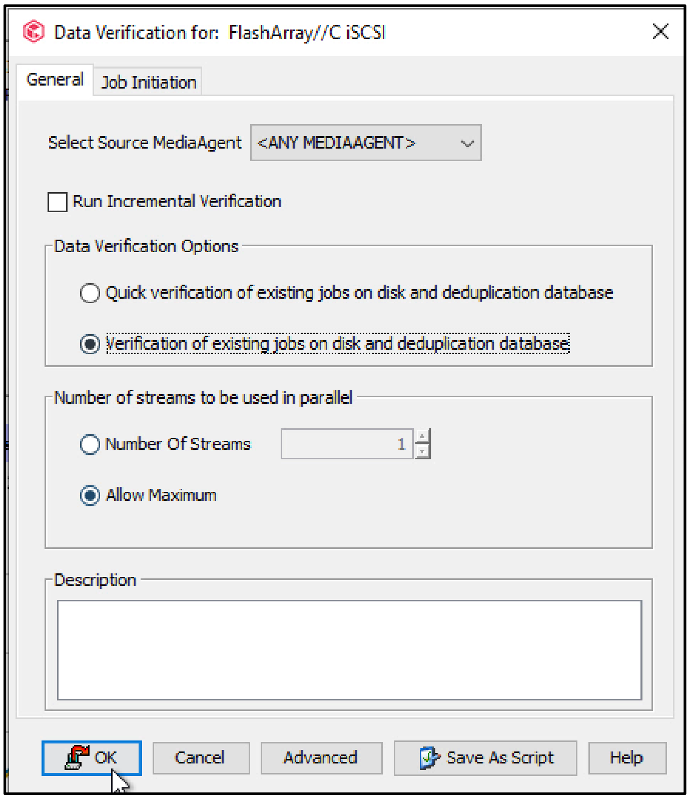

Data Verification for: FlashArray//C iSCSI

Select Source MediaAgent <ANY MEDIAAGENT>

□ Run Incremental Verification

Data Verification Options

○ Quick verification of existing jobs on disk and deduplication database

● Verification of existing jobs on disk and deduplication database

Number of streams to be used in parallel

○ Number Of Streams 1

● Allow Maximum

Description

Figure 37: Data verification options

4. Repeat this process for all the DDBs.

Returning to Normal Service

Copy Snapshot Data

If you chose to skip copying data from the snapshot during the initial recovery phase, you can copy it later using the same recovery script and adding the -CopyOnly argument.

```
Execute-CVSafeModeRecovery.ps1 -SnapshotShare \\fa\libroot-old -LinkShare \\fa\libroot -CopyOnly
```

Remove Old Managed Directory Export

Once the snapshot is no longer required, you should remove the export from the old managed directory.

Disable Symbolic Link Evaluation on MediaAgents

You should disable evaluation of remote symbolic links on all the MediaAgents.

```
fsutil behavior set symlinkEvaluation R2R:0
```

Enable FlashArray Protection Policy

To avoid SafeMode protection gaps, before resuming backups you should enable the protection policy.

Enable Backups and Data Aging

You will need to use the CommCell Console to enable the backup and data aging operations you disabled earlier, following the same steps as before.

Contact Pure Support to Restore Eradication Timer

The final step is to work with Pure Support to reduce the eradication timer back to its original setting.

Recommended Practices for Commvault with FlashArray//C

Following the practices listed in Table 5 will help improve performance for most environments.

Recommendation	Explanation
Required: Run Purity 6.1.14 or later and disable opportunistic locking	Opportunistic locking can cause intermittent failures at high restore load. Contact Pure Support to disable opportunistic locking to prevent these issues. FlashArray//C must be running Purity 6.1.14.
Share mount paths across MediaAgents	Share mount paths MediaAgent. Each MediaAgent can address up to 500TB of back-end storage. Refer to Hardware Specifications for Deduplication Mode for required CPU, RAM, DDB, and index storage for different capacity levels.
Deploy in grids of at least two MediaAgents 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. The same FlashArray//C array can act as both backup storage and a replication target for workloads running on FlashArray//X and FlashArray//C.
Tune networking on Windows Server 2019	Default TCP settings on Windows Server 2019 are not optimal for the solution. See the Server Tuning section for more details on recommended tuning parameters.
Disable default tunneling behavior in Commvault for Windows	In release 11.23, Commvault changed the default behavior for communication between clients and MediaAgents. Clients will tunnel connections to MediaAgents and the CommServe through a single TCP port. On Windows MediaAgents with more than 10Gbps bandwidth available, this can artificially limit backup and restore throughput. Set the nCLNT_FORCE_TUNNEL additional setting to 0 on all MediaAgents with more than one 10GbE connections that will manage mount paths on FlashArray//C. This limitation will be addressed in a future Commvault release.
Limit snapshot protection policies to 7 or fewer days	Snapshots of the storage pool capture a point-in-time view of data in the pool. Over time, old backup data ages out and new data comes in, and the snapshot no longer represents the time window Commvault is managing. The farther out of sync the snapshot and database are, the more complexity and uncertainty you face when recovering from it.

Table 5: Recommended practices

Conclusion

Commvault and FlashArray//C make a great combination for fast, efficient backup and recovery. The solution is simple to configure and easy to scale. Commvault's data reduction lets you protect petabytes of data in as little as 3U, with economics that rival disk-based solutions. SafeMode on FlashArray adds an extra layer of protection against malicious and accidental destruction of your backup data, helping you get back online faster after a ransomware attack.

When you're ready to see how FlashArray//C and Commvault Backup & Recovery can improve data protection for you, visit Pure's Commvault solutions page, and reach out to your Pure account team.

Additional Resources

Next Steps

- Learn more about FlashArray//C.
- See how Commvault Backup & Recovery can modernize data protection in your environment.

Supporting Information

- FlashArray//C Data Sheet
- 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