WHITE PAPER

Building Confident Cyber Resilience

Regulatory-ready Reference Architecture with Pure Storage and Commvault

Contents

Overview

In today's rapidly evolving digital landscape, the resilience of financial sector infrastructure against cyber threats has evolved from an IT concern to a matter of global financial stability. The European Union's Digital Operational Resilience Act (DORA) introduces a comprehensive framework focused on ensuring financial entities, particularly global banks, are not only prepared for unplanned events but capable of recovering swiftly and effectively. As the potential for disruption grows and cyber threats become more sophisticated, the importance of DORA mandates cannot be overstated—a single, long-term disruption has the ability to destabilize not only a firm's operations but those of all dependent entities, resulting in irreparable damage to reputation.

Under DORA, financial institutions must prioritize rapid service recovery in any unplanned event, which requires robust and well-defined strategies. Institutions that can reliably demonstrate cyber resilience will not only approach regulatory compliance with greater assurance, but will also build lasting trust and confidence among stakeholders that customer assets are being safeguarded. Failure to demonstrate resilience will lead to challenges with required risk management testing and incident reporting, resulting in significant consequences.

Audience

This reference architecture is intended for IT engineers and architects working with end customers, systems integrators, and resellers. It assumes basic familiarity with the involved technologies from Pure Storage®, Commvault®, and VMware, as well as a functional understanding of networking concepts.

Cyber Resilience with Pure Storage and Commvault

Pure Storage and Commvault have engineered a joint solution that enhances companies' cyber resilience postures. Designed around strict DORA regulation requirements and the Pure Storage layered resilience architecture, the solution offers a comprehensive approach to data protection and cyber recovery, addressing the challenges of securing, managing, recovering, and deriving value from data. By combining Commvault's industry-leading cyber resilience software with the innovative Pure Storage data platform, businesses can unlock new levels of efficiency, agility, and resilience in managing their data.

DORA Adherence

No single technology solution can guarantee compliance with all the requirements of DORA, as some articles must be addressed with business policies and processes. This joint solution is designed to address the articles that require technical responses in a way that rapidly restores critical data, applications, and systems needed to restore vital operations during and after a cyberattack. For more information on how the capabilities of this solution address specific DORA articles, see DORA Compliance with Confidence.

Solution Overview

The Commvault and Pure Storage solution follows several key design principles to ensure that data is not vulnerable to attackers or internal bad actors and that recovery is as fast as possible in a cyber event.

- **Data isolation**: Protected data must be separated from production storage and not directly accessible from production networks.
- **Data immutability and indelibility**: Data must be protected from accidental and malicious modification and deletion.
- **Management isolation**: Control planes for the isolated environment must be unreachable from production networks.
- **Microsegmentation**: Systems within the recovery environment should be separated.
- **Role separation**: The same person should not have access to both production and recovery environments.
- **Threat detection and prevention**: Content entering the vault must be checked for threats.
- **Automation**: Wherever possible, automate or allow for automation of repetitive and complex tasks.
- **Cyber recovery testing**: Automate continuous testing within an isolated recovery environment using Commvault Cloud software and Pure Storage FlashArray™ and FlashBlade® systems.
- **Cost optimization**: The cost of meeting regulations should not outweigh the cost of noncompliance.
- **Infrastructure flexibility**: As long as the required capabilities are met, companies can customize infrastructure components such as servers, network devices, and hypervisors.
- **Modular adoption**: Companies should be able to deploy the core solution components and add other components over time.

Solution Components

The joint solution consists of four main building blocks (Figure 1), each with mandatory and optional components. Each block you deploy enhances your cyber resilience posture and addresses some of the DORA articles, but addressing the complete set requires deploying all of the blocks.

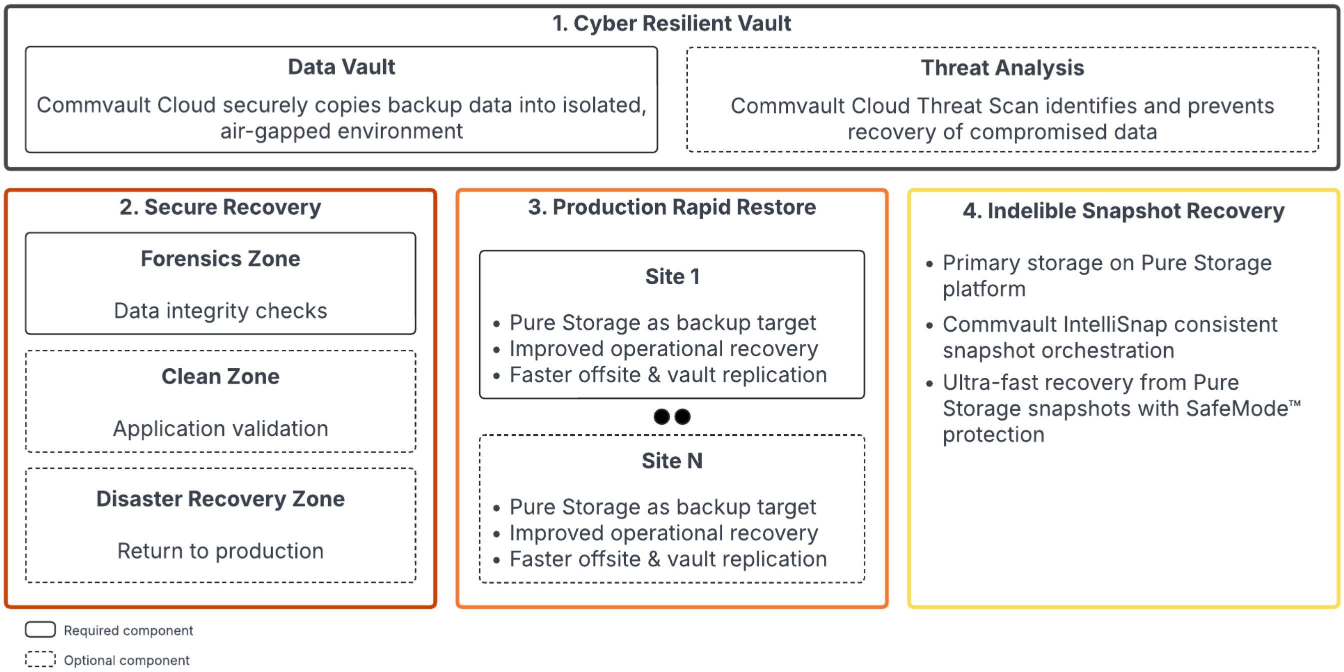
1. Cyber Resilient Vault
Data Vault
Commvault Cloud securely copies backup data into isolated, air-gapped environment
Threat Analysis
Commvault Cloud Threat Scan identifies and prevents recovery of compromised data
2. Secure Recovery
Forensics Zone
Data integrity checks
Clean Zone
Application validation
Disaster Recovery Zone
Return to production
3. Production Rapid Restore
Site 1
• Pure Storage as backup target
• Improved operational recovery
• Faster offsite & vault replication
Site N
• Pure Storage as backup target
• Improved operational recovery
• Faster offsite & vault replication
4. Indelible Snapshot Recovery
• Primary storage on Pure Storage platform
• Commvault IntelliSnap consistent snapshot orchestration
• Ultra-fast recovery from Pure Storage snapshots with SafeMode™ protection
Required component
Optional component

FIGURE 1 Solution building blocks

Solution Architecture

The fully deployed solution (Figure 2) includes isolated vault and recovery environments in a Secure Cyber Resilience Environment (SCRE) as well as components that reside in production. The Pure Storage platform forms the foundation of all the components, providing a secure, performant storage layer. Commvault Cloud adds protection, recovery, and cyber threat handling capabilities to the various blocks.

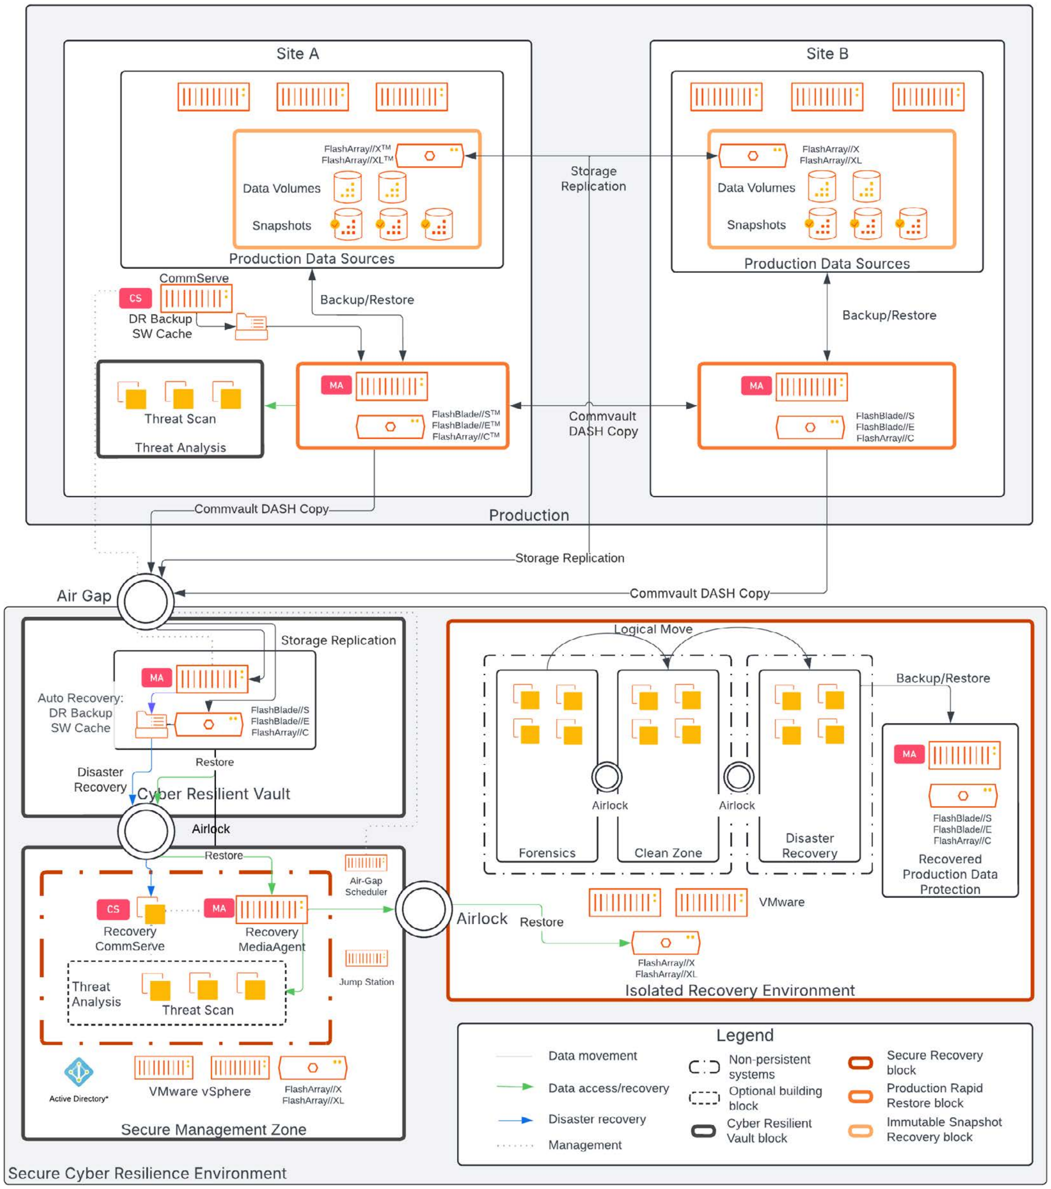

FIGURE 2 Solution architecture

Cyber Resilient Vault

The Cyber Resilient Vault (CRV) is the first block to be deployed (Figure 3). It forms the foundation of the SCRE and—with the Threat Analysis option—supports adherence to DORA articles 8, 9, 10, 12.1, and 12.5. This network zone houses the infrastructure for receiving and storing data from production systems. One or more Commvault MediaAgent systems replicate data from production backup storage into the vault and store it on the Pure Storage platform. Backups are made immutable and indelible through Pure Storage SafeMode™ Snapshots.

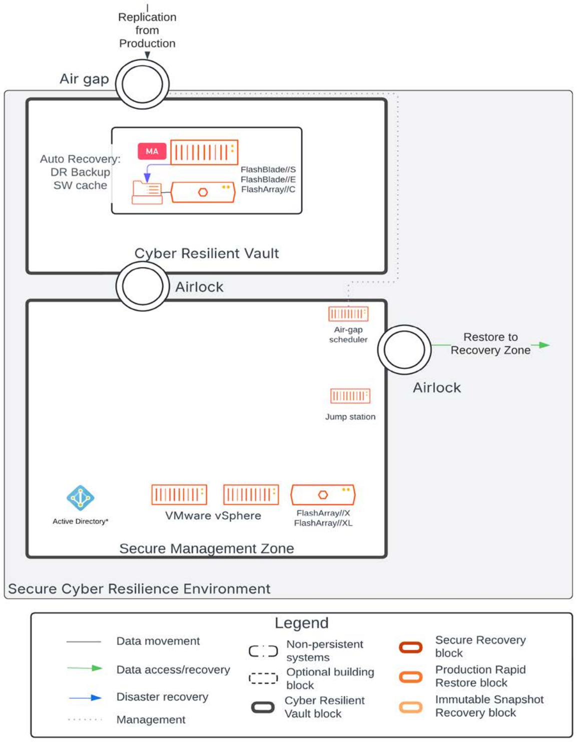

FIGURE 3 CRV and Secure Management Zone

A sometimes-available network connection, typically referred to as an "air gap" or "drawbridge," sits between the CRV and production systems. The connection is periodically enabled and disabled to allow or block network traffic to the CRV. The air gap allows only connections from MediaAgents in the CRV to MediaAgents in the production network over Commvault communication channels to prevent any compromise of production systems from reaching the CRV and affecting data stored there.

Deployed along with the CRV is the Secure Management Zone (SMZ). This network zone houses the infrastructure dedicated to managing the isolated infrastructure and recovery operations.

Network Isolation

Network isolation is one of the core principles of the CRV and SCRE design. To ensure data integrity during recovery, and to prevent disruption of recovery efforts, the SCRE must be properly separated from production. There are several key elements that must be met:

- **Air gap**: Communication between the SCRE and production must be disallowed under normal conditions and only allowed periodically to facilitate data transfer into the CRV.
- **Outbound only**: When connection across the air gap is allowed, all network sessions must initiate from within the CRV. Any sessions to a security operations center (SOC) or between the SMZ and CRV must initiate from the SMZ.
- **Traffic restrictions**: Only Commvault MediaAgents can communicate to production. They can only communicate with other Commvault Cloud systems and only using a Commvault Cloud tunnel port.
- **Isolated management interfaces**: The management interfaces for all hardware in the SCRE must be physically isolated from production and only accessible from within the SMZ.

Note: When you deploy the Indelible Snapshot Recovery block, the MediaAgents in the CRV gain access to the FlashArray to send snapshot management application programming interface (API) commands, but no access is granted with this building block.

Other Network Considerations

Additional connections not covered in this architecture may be required. For example, it might be necessary to feed data to a SOC for cyber recovery or periodic testing. Hardware-monitoring software may need to collect health statistics from isolated systems. These connections need to be reviewed individually to ensure they do not compromise the integrity of the environment or expose SCRE systems to direct access. It may also be necessary to deploy additional services within the SCRE to support these connections securely.

Required Elements

Table 1 lists the required elements for the CRV building block. Your systems integrator can assist with defining specific components to suit your environment.

Zone	Component	Details
N/A	Air-gap switches	Management interfaces not exposed to production network Severable uplinks to production network Bandwidth directly affects data transfer speeds
N/A	Management switch	Hypervisor interconnectivity
N/A	Edge firewall	Sufficient bandwidth to support required data transfer rates
N/A	Virtual private network (VPN)	Control access to jump stations
CRV	Commvault MediaAgent	Physical or virtual system meeting Commvault specifications
CRV	Pure Storage platform	FlashBlade (recommended) or FlashArray//C™
SMZ	Hypervisor	VMware vSphere recommended
SMZ	Active Directory/DNS	Required for: • Name resolution • SMB authentication • Domain-based system access management Prevent from running on same hypervisor host
SMZ	Jump station	Two recommended for redundancy Linux (recommended) or Windows Minimal CPU, RAM, storage
SMZ	Air-gap scheduler	Linux host, should be separate from jump station
SMZ	Internal firewall	Physical or virtual, microsegmentation

TABLE 1 CRV required components

Two logical "airlocks" control management communication between the CRV and SMZ and manage traffic to any optional recovery zone. These airlocks are exclusively controlled from the SMZ to ensure secure management access between critical zones.

Threat Analysis

Commvault Cloud software includes out-of-the-box protections against ransomware and other malware. MediaAgents deploy active defenses against bad actors. Anomaly detection identifies unusual changes in backed-up data, and Commvault Cloud actively monitors for telltale changes in the file system. When anomalies are detected, Commvault Cloud raises notifications on the security dashboard so you can investigate them further.

One of the key capabilities in a cyber resilience strategy is preventing reinfection during recovery. Commvault Cloud Threat Scan provides deep analysis of backup files to identify threats and corruption. Any detected threats are prevented from being recovered from any data copy. Threat Scan operates in the production backup environment (Figure 4) and scans data in the primary backup storage pool.

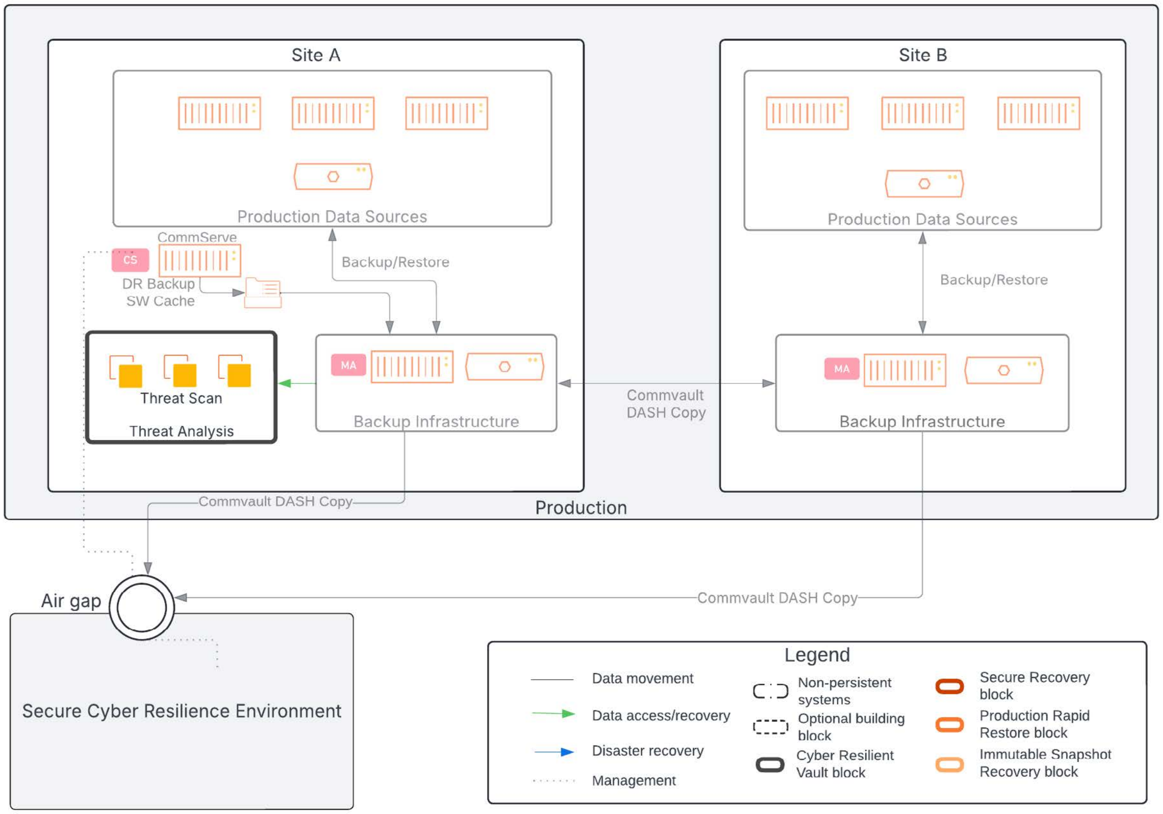

FIGURE 4 Threat analysis in the production environment

Secure Recovery

The second building block to deploy is Secure Recovery. This adds an Isolated Recovery Environment (IRE) to the SCRE, with a recovery clean room for forensic investigation and optional application testing and return-to-production spaces. It also adds on-demand Commvault systems to the SMZ to drive recovery operations. Deploying the Secure Recovery block expands adherence to DORA article 11 and supports adherence to articles 12.2, 12.3, 12.4, and 12.7. Figure 5 highlights the additional components.

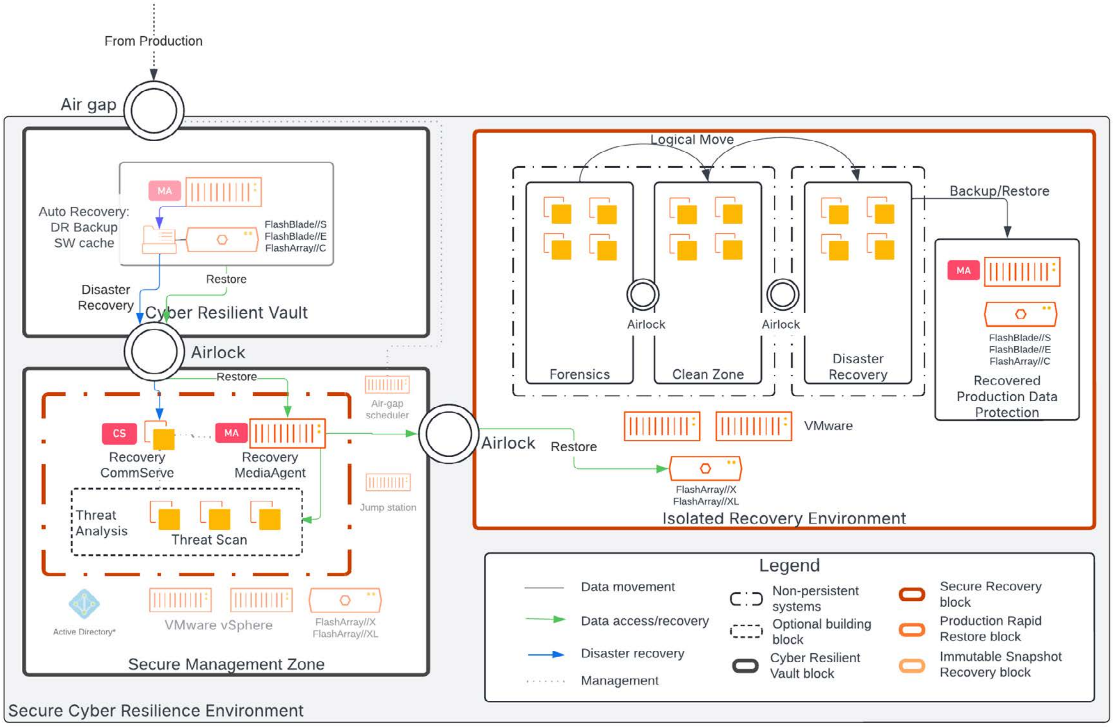

FIGURE 5 Secure Recovery components

There are no firm requirements for additional infrastructure in the SMZ. Factors to consider when deciding whether to add infrastructure include:

- **Physical or virtual MediaAgents**: The recovery MediaAgents can be physical or virtual servers.
- **Data recovery requirements**: The number of servers required depends heavily on the amount of data being recovered and the number of recovered systems that need to operate in parallel.
- **Hypervisor sharing**: The IRE infrastructure can share the same hypervisor servers and storage deployed with the initial building block. For this model to work, virtual networking must isolate the individual recovery zones from each other and from the SMZ, and access for validation users must be controlled with VPN or multifactor authentication (MFA) and separate credentials.

Threat Analysis

You can opt to deploy Threat Scan within the SMZ as part of the Commvault recovery environment. This ensures you can scan data for threats during recovery, if required.

Production Rapid Restore

The SCRE is not dependent on any particular storage for any upstream data copies. However, you can improve your cyber resilience posture by replacing the backup storage targets in your production environment with the Pure Storage platform. In addition to supporting DORA article 11 adherence, using Pure Storage brings several advantages.

- **Improved threat analysis**: Since Threat Scan must read data from backup storage to scan it, the flash-based Pure Storage platform—with its high throughput and low latency—allows Commvault to detect threats faster, helping you respond more quickly.
- **Faster vault transfer**: Your cyber resilience posture is tied to how quickly data can transfer into the CRV. Longer transfer times mean keeping the air-gap connection enabled longer, and having the most recent data available reduces data loss during cyber recovery. The performance of the Pure Storage platform minimizes the data transfer time, making you better prepared for a cyber event.

- **Enhanced data replication**: The high read performance of the Pure Storage platform lets you more quickly replicate your backup data to other recovery tiers such as offsite storage, better positioning you for recovery in a site loss or similar event.
- **Faster operational recovery**: Your operational recovery also benefits from the Pure Storage platform. A high-performance primary recovery tier accelerates your day-to-day restore operations, data refreshes, and any other operations that rely on data recovery.
- **Immutability and indelibility on more recovery tiers**: Pure Storage SafeMode technologies make it easy to add further protection layers to your higher recovery tiers, giving you resilience to more data-impact scenarios.
- **Simpler environment**: The simple configuration and management of the Pure Storage platform make it easy to add to a Commvault Cloud environment, and there's little to no reconfiguration required as your data and storage grow.
- **End storage migrations**: The Pure Storage Evergreen® model lets you upgrade storage hardware and software without disruption or downtime. With no periodic forklift migrations, you can allocate more resources to delivering new value to the organization.

You can easily add the Pure Storage platform to your production environment (Figure 6) with minimal disruption to other operations. You can deploy in any number of sites, in whatever order works best for you.

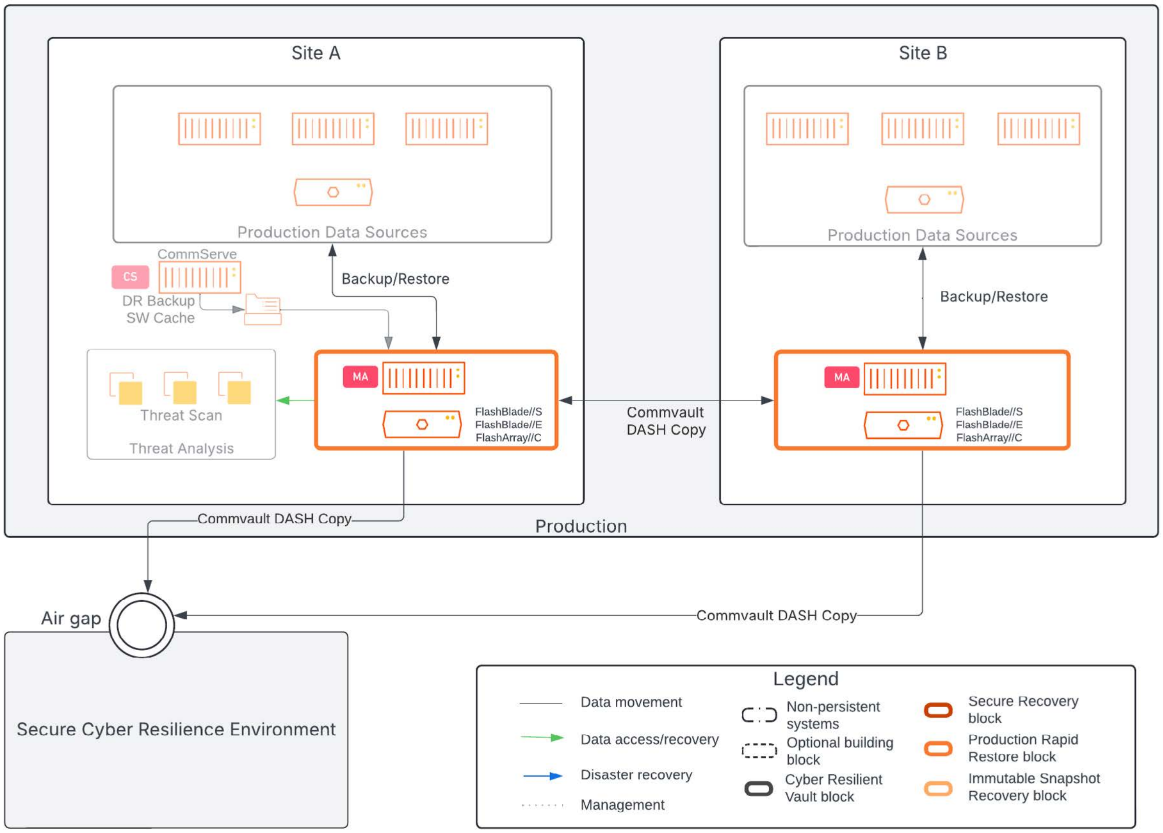

FIGURE 6 Production Rapid Restore building block

Indelible Snapshot Recovery

Some workloads and data sets are too large or time-critical to recover within the recovery time objective using traditional methods. Recovering these systems within the target window requires snapshots on the Pure Storage FlashArray. Snapshots afford the fastest operational recovery, as FlashArray doesn't need to move any data to copy or revert them. However, unprotected snapshots are vulnerable to an attacker who gains administrative control over a storage array. SafeMode protection on FlashArray prevents anyone from permanently deleting a snapshot, even with full system access, so protected snapshots of critical data are always available for recovery.

However, the National Institute of Standards and Technology (NIST) special publication Security Guidelines for Storage Infrastructure recommends that you have data copies available for recovery that are independent of systems that could be compromised, even if the snapshots are intact. To ensure recoverability of your critical systems, you need to replicate data into the CRV, where copies can be made available to the IRE.

To deploy the Indelible Snapshot Recovery block (Figure 7), you need to move your primary workloads to FlashArray//X™ or FlashArray//XL™. Optionally, you can deploy a second array in another site and leverage Purity ActiveCluster™ for high availability. Finally, you need to deploy a FlashArray (FlashArray//C™ is recommended) inside the SCRE and configure replication to it from the primary array(s) using either asynchronous or ActiveDR™ capabilities. Once the workloads reside on the Pure Storage platform, you can enable Commvault IntelliSnap® technology to orchestrate application-consistent snapshots on the primary array(s); FlashArray can then replicate those snapshots to the CRV array. The Indelible Snapshot Recovery block supports adherence to DORA articles 11 and 12.6.

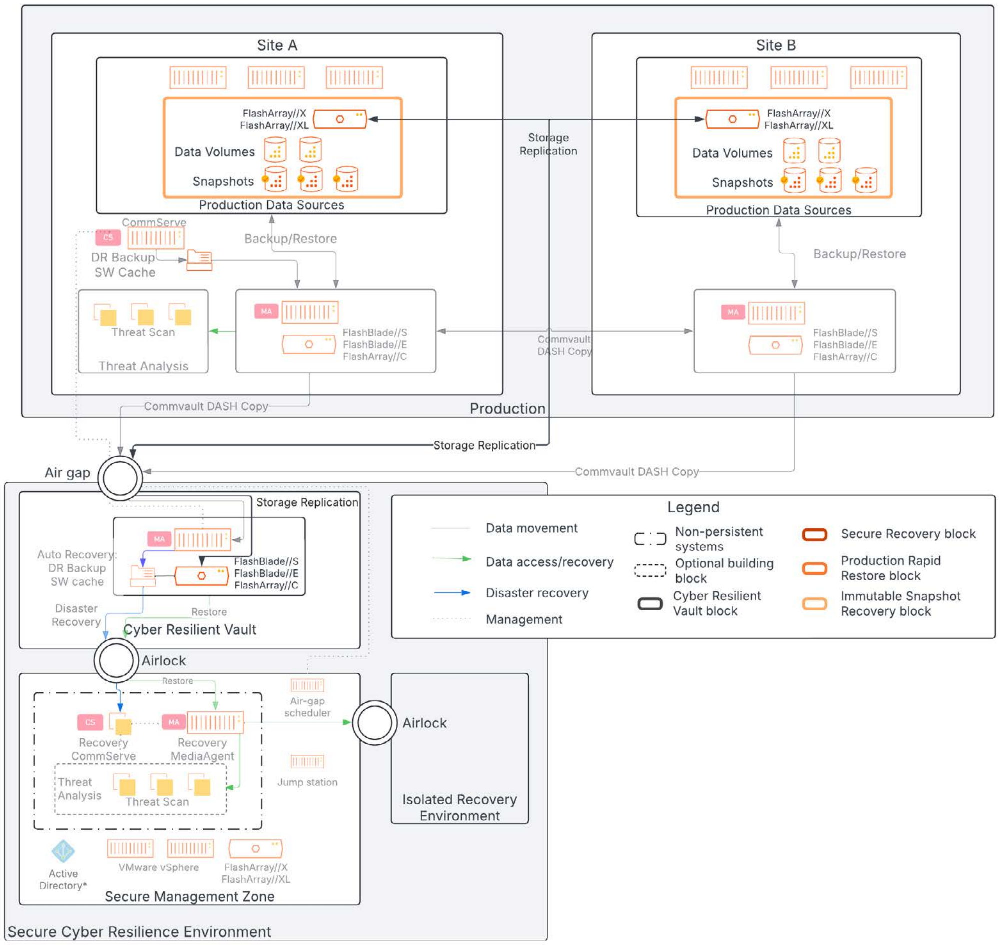

FIGURE 7 Adding Indelible Snapshot Recovery

Flexible Deployment Options

There are several different ways you can deploy Indelible Snapshot Recovery into the solution. You should discuss with your systems integrator what the best options are for your organization.

- You can use ActiveCluster synchronous replication between production arrays. This is not required, but it improves production resilience against site loss.
- You can use Fibre Channel or IP-based replication to the CRV. Fibre Channel does not require as many modifications to the air gap when establishing replication, but it requires Fibre Channel switching infrastructure.
- If you choose to deploy FlashArray//C in the CRV as a Commvault target, you can configure replication to that array or deploy a separate one.
- You can replicate directly to the FlashArray//X or FlashArray//XL in the SMZ that hosts the virtual infrastructure. This is especially useful if you use the same array to run the IRE hypervisor.
- If you deploy separate arrays in the SMZ and IRE, you can replicate data directly to the IRE array. This option requires more complex configurations to ensure the integrity of the zone boundaries.

Validation Environment

We built this solution in a Pure Storage lab (Figure 8) to validate the key architecture components. This section describes the configuration details of the deployment.

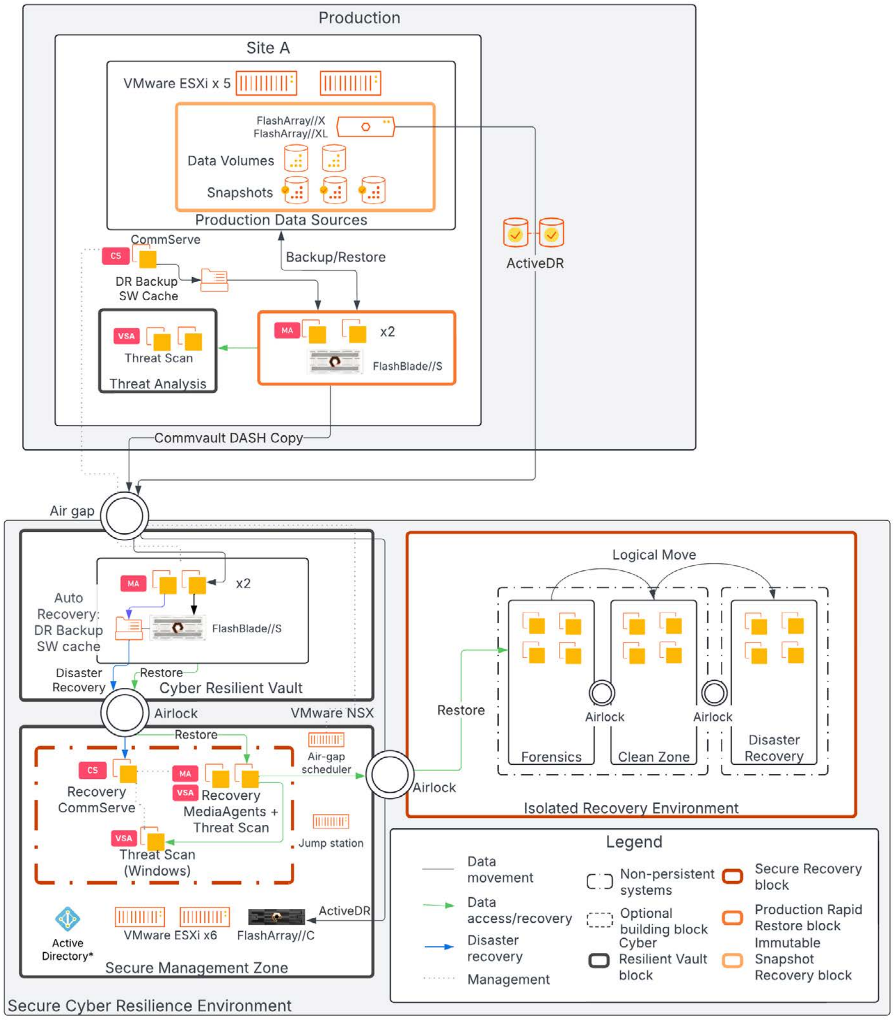

FIGURE 8 Pure Storage lab architecture

Overview

The environment consists of a single-site mock production environment and an SCRE. The source data resides in 50 VMware virtual machines (VMs) hosted on datastores on FlashArray//X and distributed across five ESXi hosts. Two virtual MediaAgents control backups to an object storage bucket on FlashBlade//S™, with object lock enabled on FlashBlade and Commvault configured with the WORM Storage Lock and Compliance Lock features. Two Threat Scan nodes, one Linux and one Windows, provide threat detection.

The SCRE is primarily built on VMware vSphere, with VMware NSX networking providing isolation from production as well as network routing and segmentation within the recovery environment. The recovery environment datastores are hosted on FlashArray//C.

In the CRV, two virtual MediaAgents communicate across the air gap to the production MediaAgents, which also function as proxies for communication with the CommServe®. The Commvault storage target is a FlashBlade//S object storage bucket. As with production, the target is configured with object lock, WORM storage lock, and compliance lock. Commvault Auto Recovery replicates the Commvault disaster recovery backups and software download cache to a file system on the FlashBlade.

The FlashArray//X in production has a replication connection to the FlashArray//C in the recovery environment. To represent critical data, one of the source volumes resides in a pod configured to replicate using ActiveDR to the FlashArray//C, where either the replicated volume or a snapshot can be used for recovery. The FlashArray//X is added to Commvault Cloud so IntelliSnap can create application-consistent snapshots.

All IP addresses within the SCRE are non-routable and unreachable from the production network. Network Address Translation (NAT) rules enable required communication to the production environment while concealing actual IP addresses from a dwelling attacker.

SCRE Network Details

Physical Network

The recovery network consists of two 100Gb vault top-of-rack (TOR) Ethernet switches, with uplinks to each other and to two switches on the production network. The ESXi hosts and Pure Storage arrays are connected to both switches to provide redundancy and performance.

The management interfaces for all isolated devices, including the TOR switches, are connected to a stand-alone management switch. This prevents an attacker from gaining administrative access to the TOR switches through the production network. A stand-alone jump station (a real-world implementation would have two for redundancy) is connected to the management switch, enabling network configuration and access to Intelligent Platform Management Interface devices on the servers. The jump station itself also needs to be connected to an accessible network to allow remote access. In practice, this should be a secure connection using a VPN and/or MFA. Due to lab constraints, we connected the jump station to one of the TOR switches and used firewall rules to restrict the access sources and ports; this would not be suitable for a production deployment.

Figure 9 shows the physical connectivity for the recovery network in the Pure Storage lab.

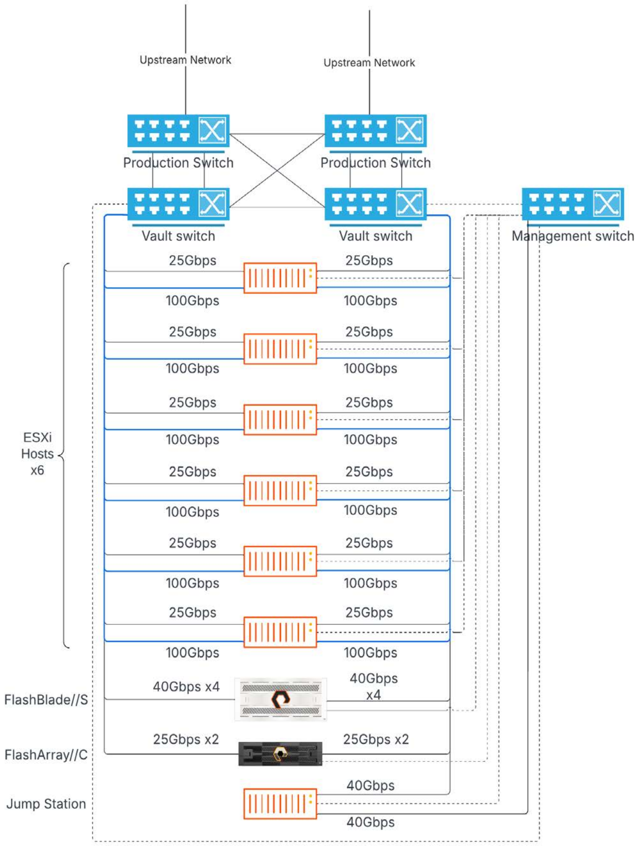

FIGURE 9 Physical network connectivity

Virtual Networking

For the lab environment, we selected VMware NSX to provide networking and security services in the SCRE. NSX, available as part of VMware Cloud Foundation, provides a range of software-defined, scale-out services from layer 2 to layer 7. Its ability to bridge between virtual and physical devices, the ubiquity of VMware deployments, and its declarative policy APIs made it a good fit for this architecture.

We created several virtual network configurations (Figure 10) inside and outside NSX. A management vSphere distributed switch (vDS) provides the foundation for the NSX Manager nodes, NSX Edge nodes (for communication outside the vSphere cluster), vCenter server, and an Active Directory and Domain Name System (DNS) server. The vDS uses two 25Gb ports from each ESXi host as uplinks.

The NSX virtual switch defined in vSphere uses two 100Gb ports from each ESXi host as uplinks. We defined separate topologies for the CRV and SMZ/recovery zones. With this approach, we were able to easily segregate the Commvault auxiliary copy and FlashArray replication traffic between separate NSX Edge nodes, and therefore physical network ports, maximizing throughput for both data transfers. Each of the topologies has a tier 0 (T0) and tier 1 (T1) gateway defined. Distributed firewall services enforce microsegmentation policies to minimize the attack surface. NAT services on the T0 gateway translate all communication with production to present different IP addresses than configured on the actual systems.

The CRV topology defines a single overlay segment, which houses the vault MediaAgents and a virtual jump station for managing the isolated systems. Microsegmentation restricts access between segments and within each segment.

The SMZ/recovery topology has three segments defined. The Commvault segment contains all the systems that get built during recovery, such as the CommServe, MediaAgents, and Virtual Server Agents. The other segments contain the recovered VMs for the forensic and clean zones. Jump stations for forensic and application testers would also reside in these segments.

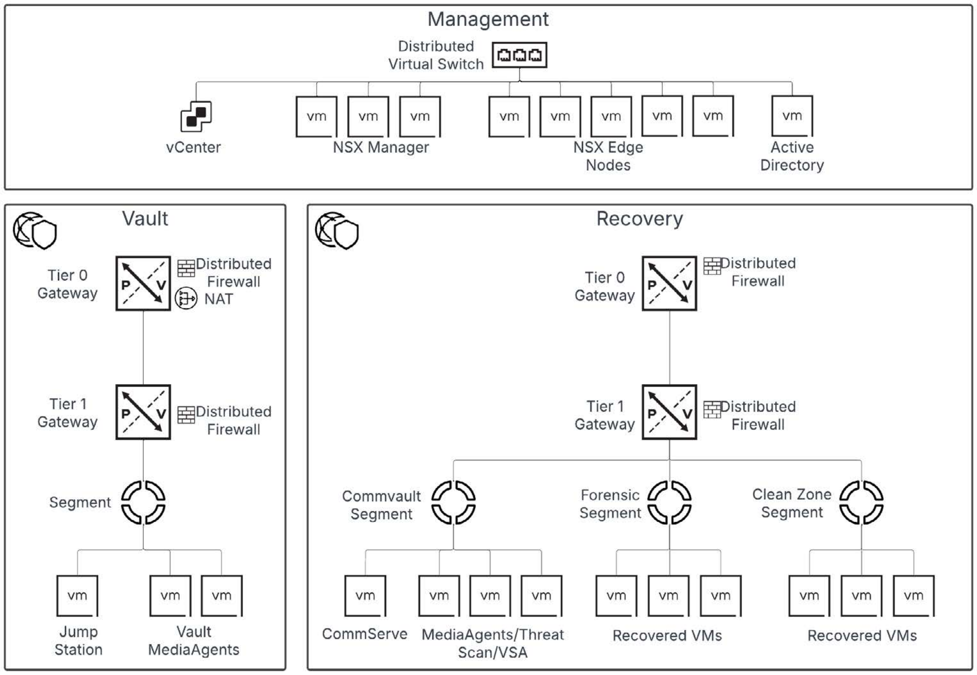

FIGURE 10 Virtual network topology

Network Address Translation

We set up two NAT rules in NSX to prevent identification of private IP addresses in the SCRE. The Commvault MediaAgents in the CRV communicate to the production MediaAgents, with all CRV MediaAgents translated to a single outward-facing IP address. Similarly, the SCRE FlashArray can communicate with the production FlashArray so they can replicate data changes into the SCRE.

Table 2 lists the NAT rules we configured.

Gateway	Action	Source	Destination
CRV T0	SNAT	CRV MediaAgents	Production MediaAgents
SMZ/Recovery T0	SNAT	SMZ FlashArray//C Management IP address Replication IP addresses	Production FlashArray//X Management IP address Replication IP addresses

TABLE 2 NSX NAT rules

NSX Distributed Firewall

To allow and restrict access between the CRV and production and between the systems in the SCRE zones, we configured a number of policies and rules in the NSX distributed firewall. To simplify the firewall rules, we defined services in NSX for FlashArray replication and Commvault tunnels. We created NSX groups for rule assignment, using dynamic groups wherever possible. We set the default layer 3 rule to drop all packets so that only explicitly allowed communication is possible. Table 3 lists the communication we allowed in our policies.

Source Systems	Destination Systems	Services/Ports
Commvault recovery systems	CRV FlashBlade data VIP	HTTPS (TCP/443) SMB (TCP/445)
	SMZ vCenter SMZ ESXi hosts	HTTPS TCP/902
	SMZ shared services DNS Active Directory	DNS (UDP/53) LDAP (TCP/389, TCP/636) HTTPS Kerberos (TCP/88, UDP/88) Active Directory (TCP/135, TCP/445)
	Same segment	Any
Commvault CRV MediaAgents	CRV FlashBlade data VIP	HTTPS, NFS (TCP/2049)
	Production MediaAgents	Commvault (TCP/8403)
	SMZ shared services	DNS LDAP HTTPS Kerberos
SMZ FlashArray management	Production FlashArray management	HTTPS
	SMZ shared services	DNS LDAP HTTPS Kerberos
SMZ FlashArray	Production FlashArray replication ports	FlashArray replication (TCP/8117)
CRV FlashBlade management	SMZ shared services	DNS LDAP HTTPS Kerberos
Physical jump station	Virtual jump station	Remote Desktop Protocol (TCP/3389, UDP/3389)
	vCenter NSX management IPs	HTTPS
Virtual jump station	SMZ FlashArray management	HTTPS
	CRV Commvault MediaAgents	SSH (TCP/22)

TABLE 3 Required communication

FlashArray Configuration

We deployed two FlashArray systems in the lab. A FlashArray//X on the production network hosts the primary data, while a FlashArray//C in the SCRE network houses the virtual infrastructure for the CRV, SMZ, and recovery zones.

Storage Volumes

On the FlashArray//X, we created three volumes for VMware datastores (Figure 11). One volume is used as a workspace for Threat Scan, and the other two volumes host the VMs we're using as a data source.

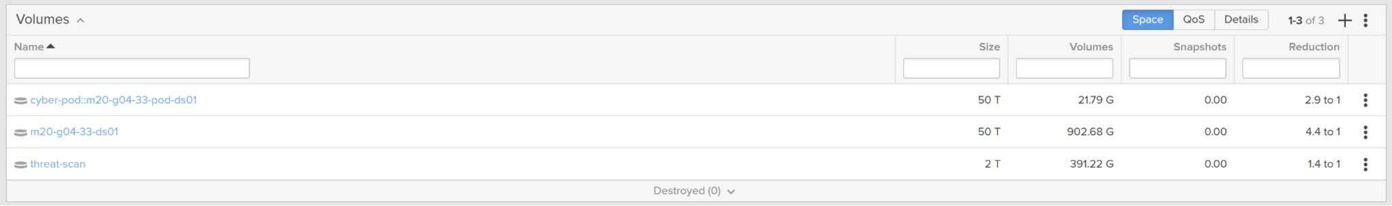

Volumes ^ | Space | QoS | Details | 1-3 of 3

Name ▲	Size	Volumes	Snapshots	Reduction
cyber-pod::m20-g04-33-pod-ds01	50 T	21.79 G	0.00	2.9 to 1
m20-g04-33-ds01	50 T	902.68 G	0.00	4.4 to 1
threat-scan	2 T	391.22 G	0.00	1.4 to 1

Destroyed (0)

FIGURE 11 FlashArray//X primary volumes

One of the volumes represents mission-critical data that must be recovered quickly (using snapshots) in a cyber incident. To facilitate this, we created a replication connection between the FlashArray//X and the FlashArray//C. We then added the volume to a pod and configured ActiveDR to maintain synchronization from the FlashArray//X to the FlashArray//C (Figure 12).

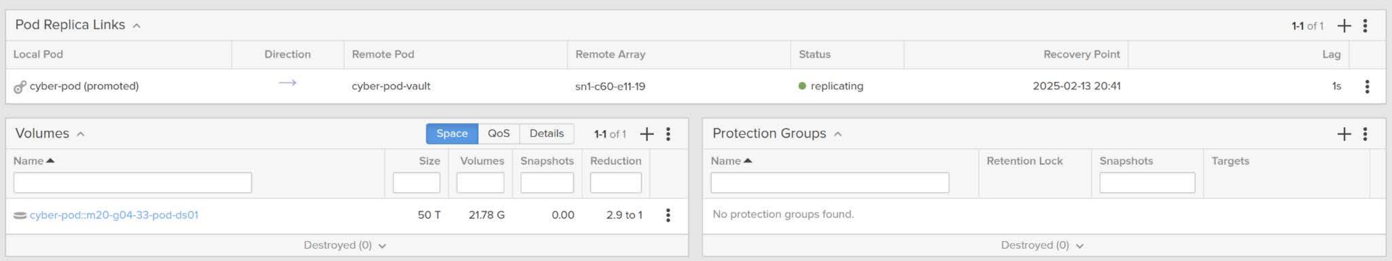

Pod Replica Links ^ 1-1 of 1

Local Pod	Direction	Remote Pod	Remote Array	Status	Recovery Point	Lag
cyber-pod (promoted)	→	cyber-pod-vault	sn1-c60-e11-19	● replicating	2025-02-13 20:41	1s

Volumes ^ | Space | QoS | Details | 1-1 of 1

Name ▲	Size	Volumes	Snapshots	Reduction
cyber-pod::m20-g04-33-pod-ds01	50 T	21.78 G	0.00	2.9 to 1

Destroyed (0)

Protection Groups ^

Name ▲	Retention Lock	Snapshots	Targets
No protection groups found.			

Destroyed (0)

FIGURE 12 ActiveDR pod settings

Storage Admin Account

On the FlashArray//X, we created a local account (Figure 13) to use with Commvault IntelliSnap. We granted the account the storage_admin role and created an API token for it so Commvault Cloud could send commands to the array.

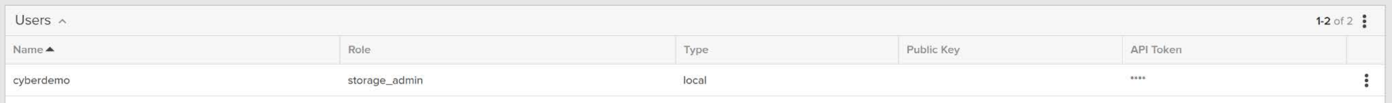

Users ^ 1-2 of 2

Name ▲	Role	Type	Public Key	API Token
cyberdemo	storage_admin	local		****

FIGURE 13 FlashArray storage admin account

FlashBlade Configuration

On the FlashBlade arrays, we configured object storage for Commvault according to Pure Storage best practices. We also created a file system to accept the replicated Commvault disaster recovery backups and software download cache.

Network

To enable data access from the CRV and SMZ, we created a 24-bit subnet on the FlashBlade and an interface on that subnet with data service enabled (Figure 14).

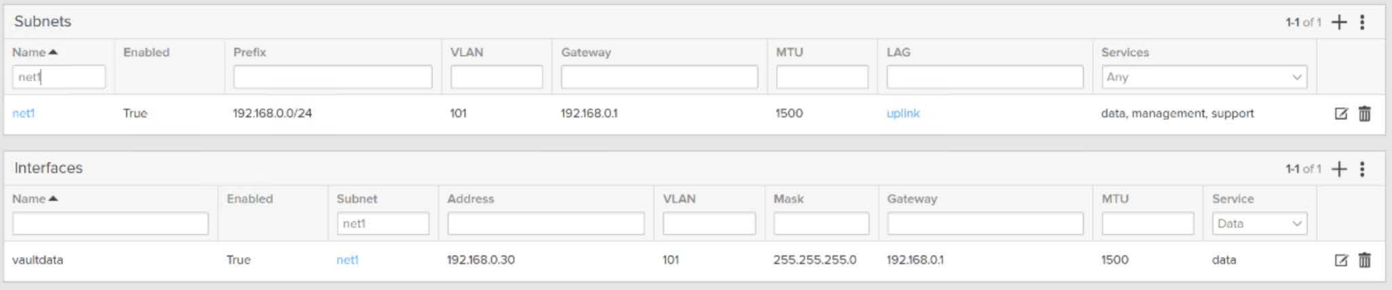

Subnets

1-1 of 1

Name ▲	Enabled	Prefix	VLAN	Gateway	MTU	LAG	Services
net1							Any
net1	True	192.168.0.0/24	101	192.168.0.1	1500	uplink	data, management, support

Interfaces

1-1 of 1

Name ▲	Enabled	Subnet	Address	VLAN	Mask	Gateway	MTU	Service
		net1						Data
vaultdata	True	net1	192.168.0.30	101	255.255.255.0	192.168.0.1	1500	data

FIGURE 14 FlashBlade network interface configuration

Object Storage

For the production FlashBlade, we used the Provision Pure Storage FlashBlade Object Storage workflow available in the Commvault Store to manage the provisioning; the SCRE network restrictions require manual provisioning for the FlashBlade buckets. The final configurations were the same for both approaches.

On each FlashBlade, we created an object storage account. Within each account, we created (Figure 15):

- Two object access policies.
  - One policy grants access to read, write, and delete objects in the bucket from the CRV MediaAgents.
  - The other policy grants only read and write access. This policy is used with Storage Accelerator clients.
  - Both policies allow setting object lock attributes.
  - A rule in each policy grants read-only access from the Commvault Cloud recovery systems to prevent data being changed during recovery.
- Two object users, one for privileged and one for unprivileged access: Each account is associated with one of the object access policies.
- An object access key for each object user.
- Two object buckets, one with and one without object lock enabled.

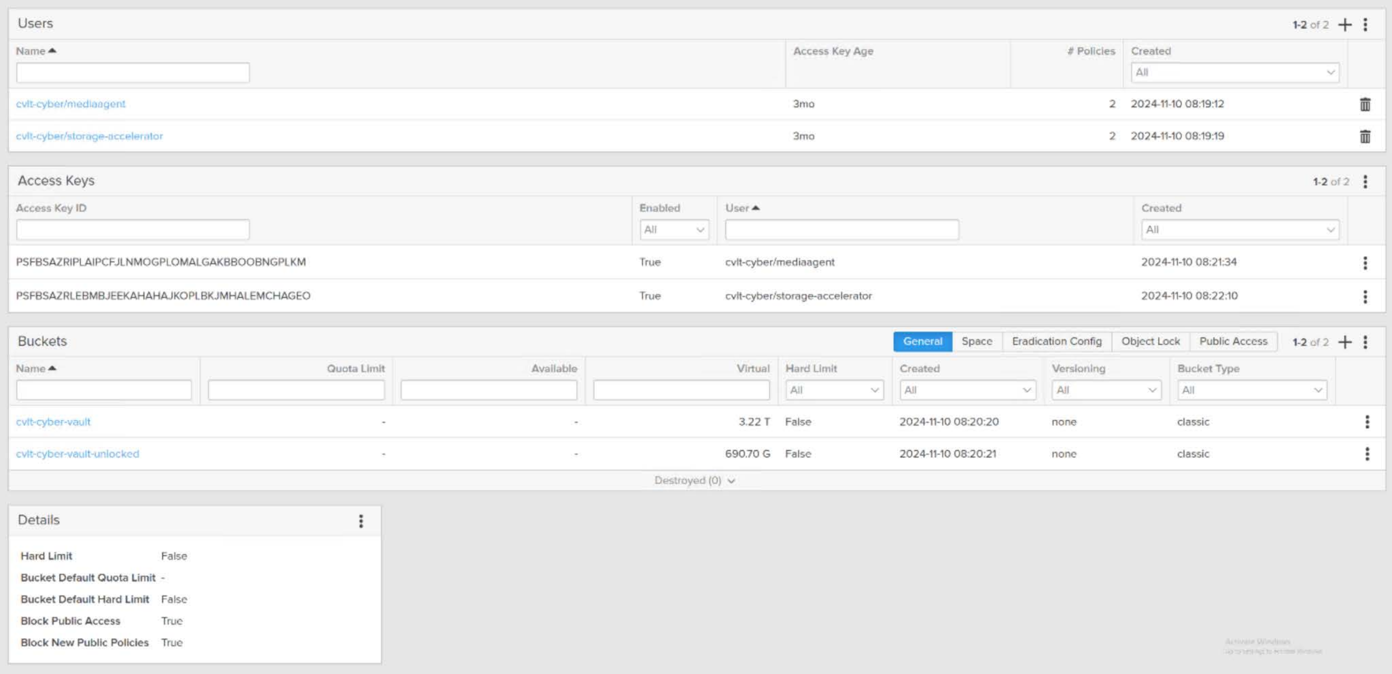

Users — 1-2 of 2

Name	Access Key Age	# Policies	Created
cvlt-cyber/mediaagent	3mo	2	2024-11-10 08:19:12
cvlt-cyber/storage-accelerator	3mo	2	2024-11-10 08:19:19

Access Keys — 1-2 of 2

Access Key ID	Enabled	User	Created
PSFBSAZRIPLAIPCFJLNMOGPLOMALGAKBBOOBNGPLKM	True	cvlt-cyber/mediaagent	2024-11-10 08:21:34
PSFBSAZRLEBMBJEEKAHAHAJKOPLBKJMHALEMCHAGEO	True	cvlt-cyber/storage-accelerator	2024-11-10 08:22:10

Buckets — General | Space | Eradication Config | Object Lock | Public Access — 1-2 of 2

Name	Quota Limit	Available	Virtual	Hard Limit	Created	Versioning	Bucket Type
cvlt-cyber-vault	-	-	3.22 T	False	2024-11-10 08:20:20	none	classic
cvlt-cyber-vault-unlocked	-	-	690.70 G	False	2024-11-10 08:20:21	none	classic

Destroyed (0)

Details

Hard Limit	False
Bucket Default Quota Limit	-
Bucket Default Hard Limit	False
Block Public Access	True
Block New Public Policies	True

FIGURE 15 FlashBlade object account

Following our best practices, on the buckets with object lock enabled, we configured the following (Figure 16):

- Versioning: disabled
- Object Lock: enabled
- Default Retention: none
- Default Retention Mode: none
- Freeze Locked Objects: enabled
- Retention Lock: ratcheted
- Eradication Mode: permission-based

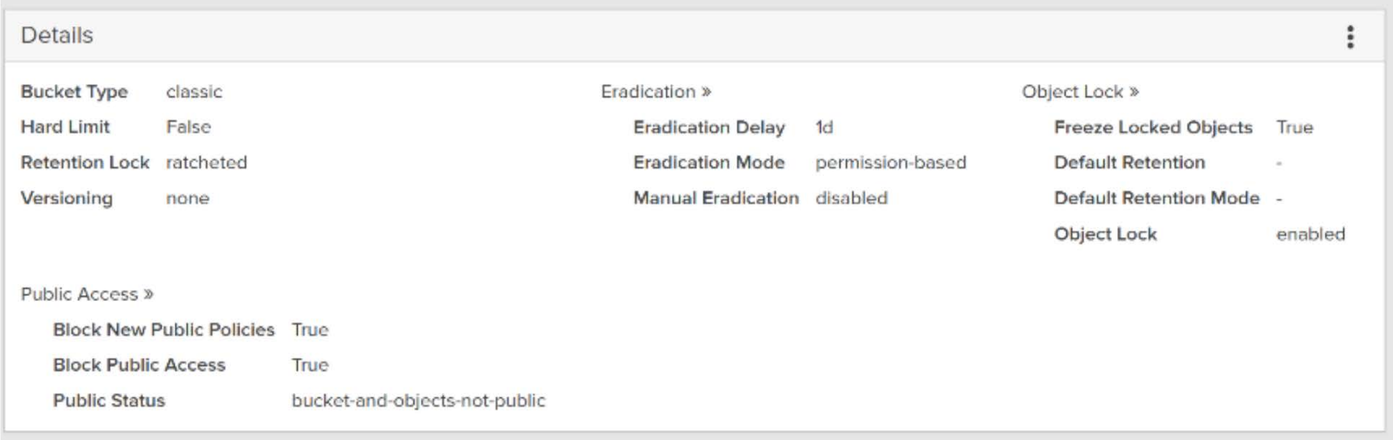

Details

Bucket Type	classic
Hard Limit	False
Retention Lock	ratcheted
Versioning	none

Eradication »

Eradication Delay	1d
Eradication Mode	permission-based
Manual Eradication	disabled

Object Lock »

Freeze Locked Objects	True
Default Retention	-
Default Retention Mode	-
Object Lock	enabled

Public Access »

Block New Public Policies	True
Block Public Access	True
Public Status	bucket-and-objects-not-public

FIGURE 16 Bucket object lock settings

File System for Auto Recovery (with Mount)

We created a file system on the SCRE FlashBlade to provide Commvault Cloud a location for synchronizing the disaster recovery backups and software download cache.

We first joined the array to the Active Directory domain in the SMZ to enable Server Message Block (SMB) authentication (Figure 17).

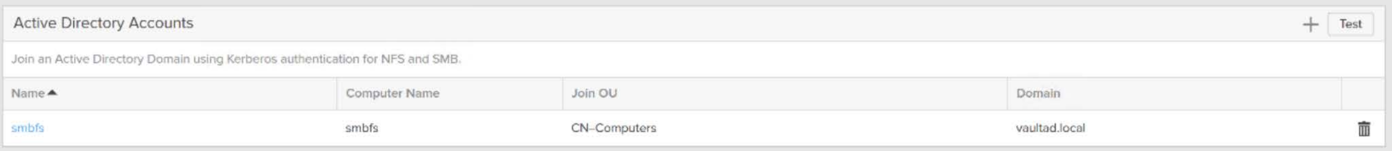

Active Directory Accounts

Join an Active Directory Domain using Kerberos authentication for NFS and SMB.

Name ▲	Computer Name	Join OU	Domain
smbfs	smbfs	CN-Computers	vaultad.local

FIGURE 17 SCRE FlashBlade joined to Active Directory

We created a client policy for SMB and added a rule (Figure 18) to allow only the recovery CommServe as a client, with read-only access.

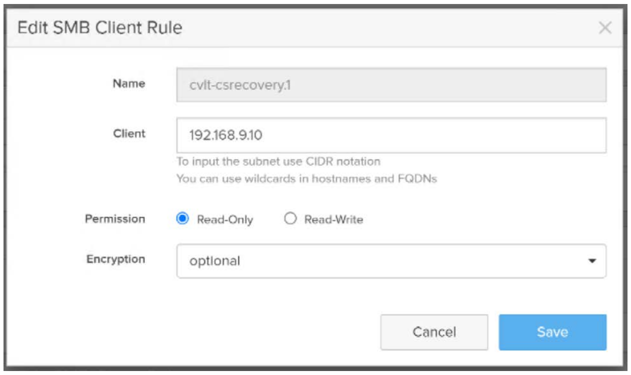

Edit SMB Client Rule

Name: cvlt-csrecovery.1

Client: 192.168.9.10

To input the subnet use CIDR notation
You can use wildcards in hostnames and FQDNs

Permission: ● Read-Only ○ Read-Write

Encryption: optional

FIGURE 18 SMB client rule

We created the file system and enabled NFSv3 exports, with export rules configured to allow only the CRV MediaAgents (Figure 19).

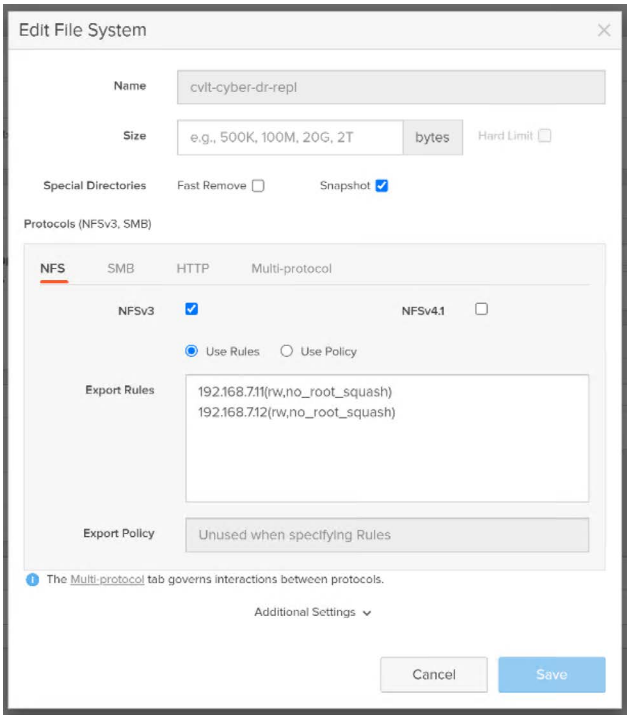

Edit File System

Name: cvlt-cyber-dr-repl

Size: e.g., 500K, 100M, 20G, 2T bytes Hard Limit □

Special Directories: Fast Remove □ Snapshot ☑

Protocols (NFSv3, SMB)

NFSv3 ☑ NFSv4.1 □

● Use Rules ○ Use Policy

Export Rules:

```
192.168.7.11(rw,no_root_squash)
192.168.7.12(rw,no_root_squash)
```

Export Policy: Unused when specifying Rules

The Multi-protocol tab governs interactions between protocols.

FIGURE 19 FlashBlade file system with Network File System (NFS) export rules

We also enabled SMB on the FlashBlade file system, applying the client policy we created (Figure 20). We applied the default SMB share policy.

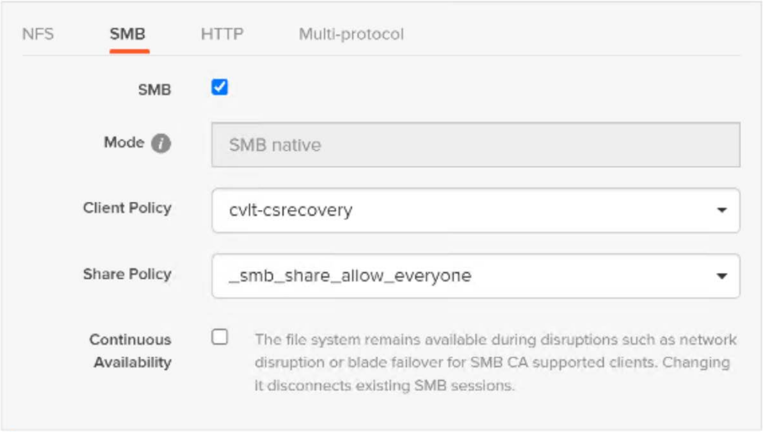

SMB ☑

Mode: SMB native

Client Policy: cvlt-csrecovery

Share Policy: _smb_share_allow_everyone

Continuous Availability ☐ The file system remains available during disruptions such as network disruption or blade failover for SMB CA supported clients. Changing it disconnects existing SMB sessions.

FIGURE 20 FlashBlade file system SMB settings

We set the Multi-protocol behavior to Shared and left the Safeguard ACLs setting enabled (Figure 21).

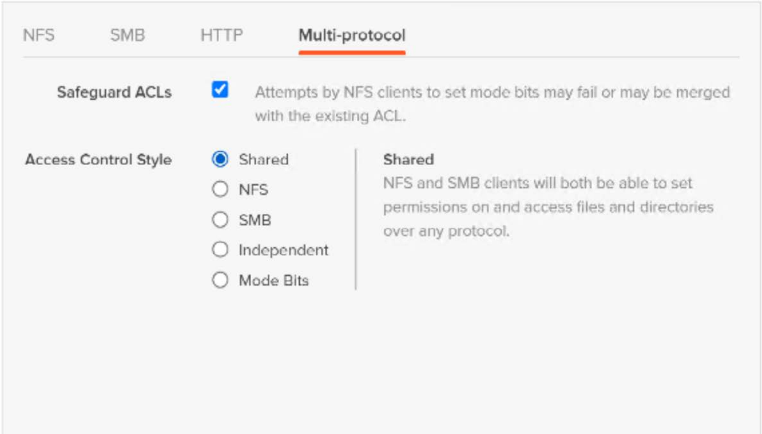

Safeguard ACLs ☑ Attempts by NFS clients to set mode bits may fail or may be merged with the existing ACL.

Access Control Style

● Shared
○ NFS
○ SMB
○ Independent
○ Mode Bits

Shared
NFS and SMB clients will both be able to set permissions on and access files and directories over any protocol.

FIGURE 21 FlashBlade file system multi-protocol settings

VMware vSphere

In both the production and SCRE environments, we used iSCSI to connect the respective FlashArray to the ESXi hosts. From the FlashArray, we connected the volumes we created earlier to those hosts. We performed a storage rescan on the vSphere cluster, and we created VMware VMFS datastores on the newly detected volumes.

Source Data

We created a test set of 50 VMs with a mix of Windows and Linux operating systems and placed them on the two production datastores. Five of the Windows VMs went on the datastore configured with ActiveDR, and we placed the rest on the independent datastore. To simplify the backup configuration in Commvault Cloud, we also created VM folders for ActiveDR and non-ActiveDR VMs and placed the VMs in those folders.

Commvault Cloud Configuration

This section details the production Commvault Cloud elements that we configured while implementing the solution.

Network Topologies

To allow the CRV MediaAgents to communicate outbound to the production MediaAgents and CommServe, we defined two network topologies in Commvault Cloud. Prior to creating the topologies, we created server groups for production and CRV that contain the MediaAgents for that zone.

The first topology (Figure 22) is of the one-way type and controls the direct communication between the MediaAgents across the air gap. The Servers option is set to the CRV MediaAgent server group, and the production MediaAgent group is set as the Infrastructure machines option.

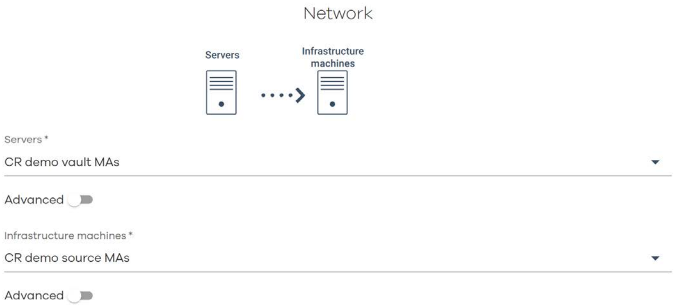

Servers *
CR demo vault MAs

Advanced

Infrastructure machines *
CR demo source MAs

Advanced

FIGURE 22 Commvault Command Center™: one-way network topology for MediaAgent-MediaAgent traffic

The second topology (Figure 23) is of the one-way forwarding type. It allows communication from the CRV MediaAgents to the CommServe, with the production MediaAgents acting as gateways. The Servers option is set to the CRV MediaAgent server group, the production MediaAgent group is set as the Network gateways option, and the Infrastructure machines option is set to the special group My CommServe Computer.

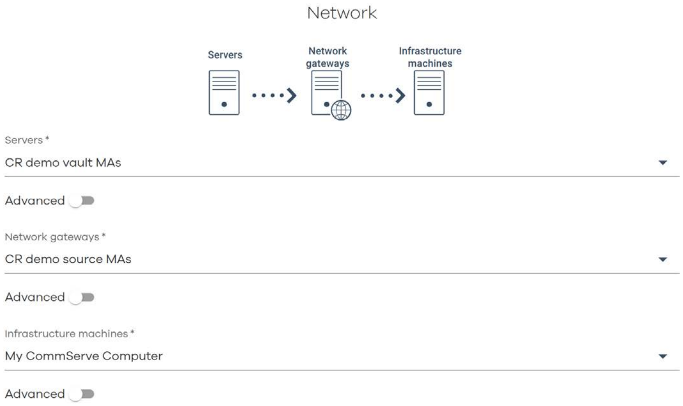

Servers *
CR demo vault MAs

Advanced

Network gateways *
CR demo source MAs

Advanced

Infrastructure machines *
My CommServe Computer

Advanced

FIGURE 23 Commvault Command Center: one-way forwarding network topology for MediaAgent-CommServe traffic

Stored Credentials

To simplify password management, we created stored credentials in the Commvault Cloud credential vault for each of the accounts on the FlashBlade and FlashArray systems.

For the FlashArray//X, we created a credential of the Storage Array Account type, entering the API token we created earlier into the Password field (Figure 24).

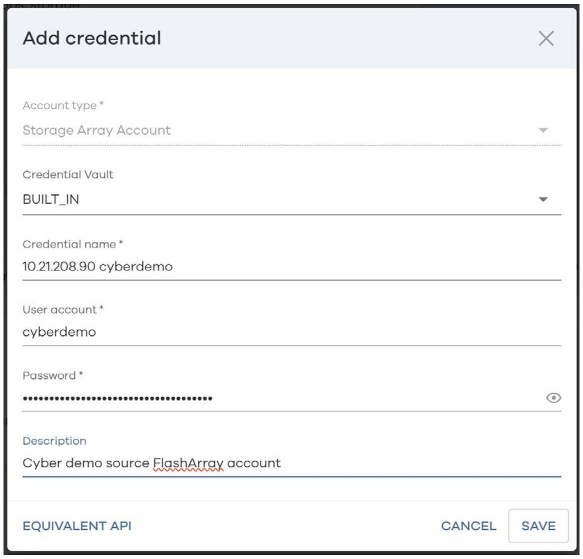

Add credential

Account type *
Storage Array Account

Credential Vault
BUILT_IN

Credential name *
10.21.208.90 cyberdemo

User account *
cyberdemo

Password *
••••••••••••••••••••••••••••••••••••

Description
Cyber demo source FlashArray account

FIGURE 24 Stored credential for FlashArray//X storage admin account

For the FlashBlade arrays, we created Cloud Account credentials using the Pure Storage FlashBlade vendor type. We added a credential for each privileged and unprivileged user on each array. Figure 25 shows an example credential.

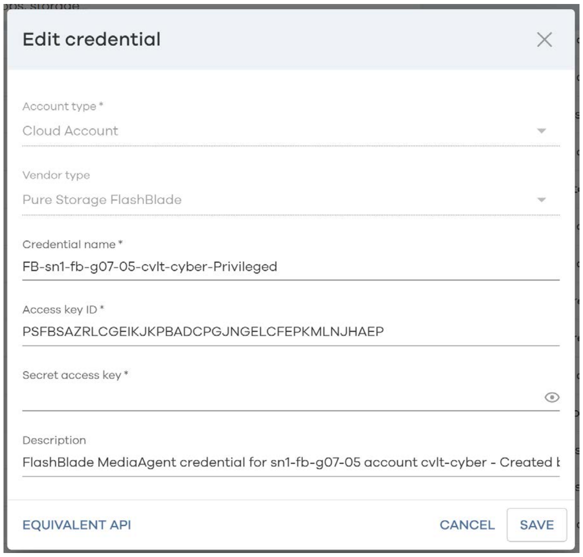

Edit credential

Account type *
Cloud Account

Vendor type
Pure Storage FlashBlade

Credential name *
FB-sn1-fb-g07-05-cvlt-cyber-Privileged

Access key ID *
PSFBSAZRLCGEIKJKPBADCPGJNGELCFEPKMLNJHAEP

Secret access key *

Description
FlashBlade MediaAgent credential for sn1-fb-g07-05 account cvlt-cyber - Created b

FIGURE 25 Stored credential for FlashBlade object user

Storage Pools

We created two storage pools with WORM Storage Lock enabled to manage our test backup data. For the storage targets, we used the FlashBlade object buckets that had object lock enabled. We also created two storage pools without WORM storage lock to manage IntelliSnap index data, index backups, and deduplication database (DDB) backups.

Figure 26 shows the bucket configuration in Commvault Cloud for the production storage pool. Each production MediaAgent has read/write access, and the Storage Accelerator credentials drop-down is set to the unprivileged user credential.

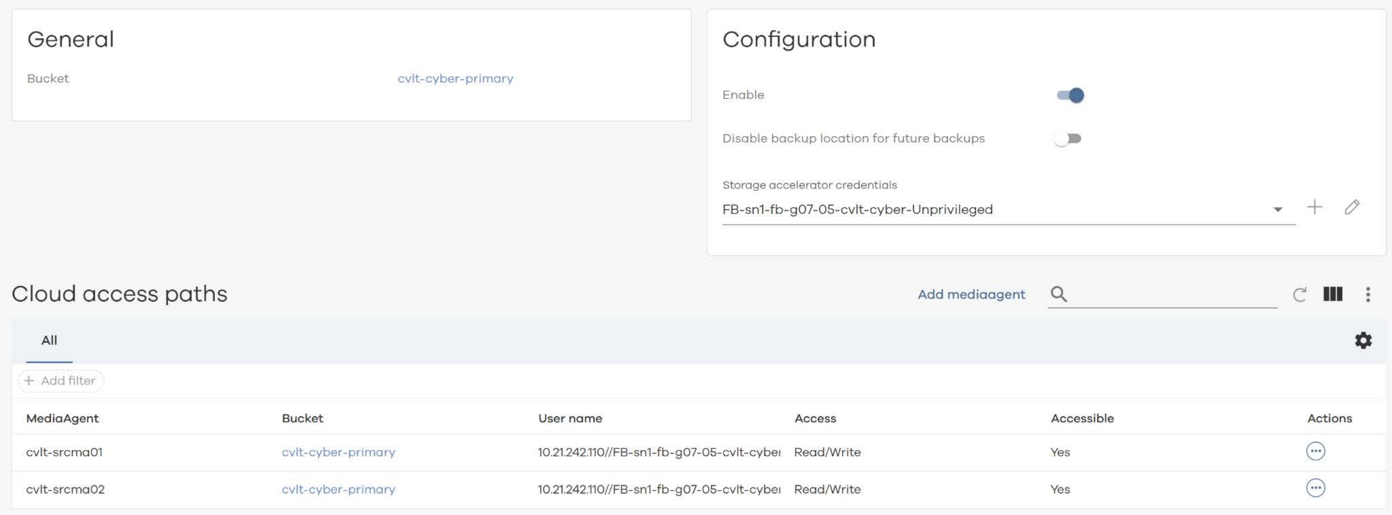

General

Bucket: cvlt-cyber-primary

Configuration

Enable: ●

Disable backup location for future backups: ○

Storage accelerator credentials: FB-sn1-fb-g07-05-cvlt-cyber-Unprivileged

Cloud access paths

MediaAgent	Bucket	User name	Access	Accessible	Actions
cvlt-srcma01	cvlt-cyber-primary	10.21.242.110//FB-sn1-fb-g07-05-cvlt-cyber	Read/Write	Yes	⋯
cvlt-srcma02	cvlt-cyber-primary	10.21.242.110//FB-sn1-fb-g07-05-cvlt-cyber	Read/Write	Yes	⋯

FIGURE 26 Source storage pool bucket configuration

Figure 27 shows the WORM settings for the production storage pool. Retention for the storage pool is set to one month, and every server backup plan that uses the pool will inherit that setting.

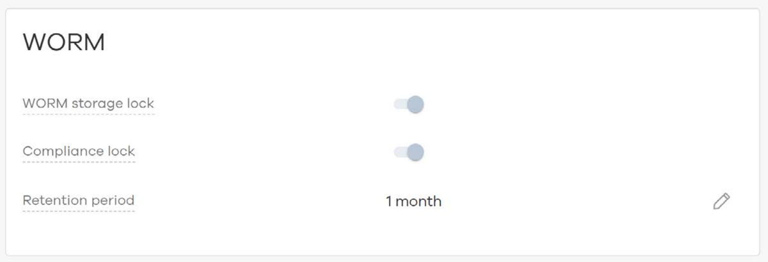

WORM

WORM storage lock: ●

Compliance lock: ●

Retention period: 1 month

FIGURE 27 Production storage pool WORM settings

Figure 28 shows the CRV storage pool bucket configuration. As with production, both CRV MediaAgents have read/write access. We don't use the Storage Accelerator feature in the SCRE, but we still set the Storage Accelerator credential for extra protection against privileged key compromise.

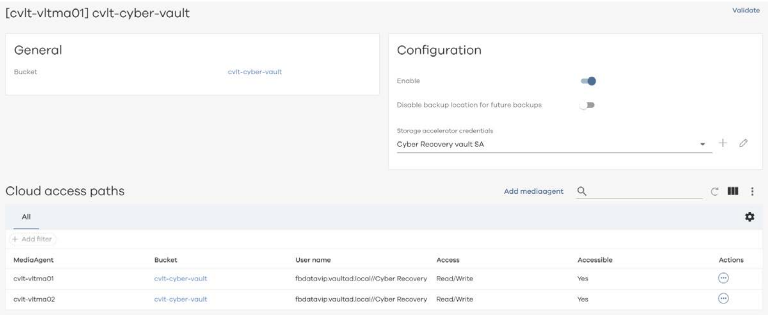

[cvlt-vltma01] cvlt-cyber-vault

General

Bucket: cvlt-cyber-vault

Configuration

Enable: ●

Disable backup location for future backups: ○

Storage accelerator credentials: Cyber Recovery vault SA

Cloud access paths

MediaAgent	Bucket	User name	Access	Accessible	Actions
cvlt-vltma01	cvlt-cyber-vault	fbdatavip.vaultad.local//Cyber Recovery	Read/Write	Yes	⋯
cvlt-vltma02	cvlt-cyber-vault	fbdatavip.vaultad.local//Cyber Recovery	Read/Write	Yes	⋯

FIGURE 28 CRV storage pool bucket settings

Figure 29 shows the WORM settings for the CRV storage pool. We set a lower retention than production, but you can customize the value based on your own requirements.

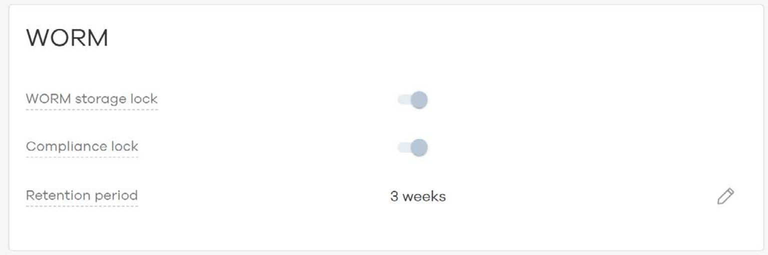

WORM

WORM storage lock	●
Compliance lock	●
Retention period	3 weeks

FIGURE 29 CRV storage pool WORM settings

Auto Recovery for File System Replication

The Auto Recovery feature in Commvault Cloud keeps data in sync between two locations. It supports a number of data types, with some variation in the transfer mechanisms for different types. We used file system synchronization to transfer the CommServe disaster recovery backups and software download cache to a file system on the CRV FlashBlade//S. This method lets us share the same data flows that transfer the production backups into the CRV, so we do not need to allow any additional traffic across the air gap. It is also platform agnostic, so it doesn't matter whether the CommServe server runs Windows or Linux, whether the disaster recovery backups and software cache are on local or network storage, or whether the CRV storage target is FlashBlade or FlashArray.

To synchronize the data, Auto Recovery for file systems performs an out-of-place restore from the predefined source copy. It's triggered when each backup job completes. To set up Auto Recovery, we first created a file system backup subclient on the CommServe and set its content to the disaster recovery backup and software cache locations (Figure 30). These could be defined in separate subclients or even separate servers, but having a single subclient simplifies the operations.

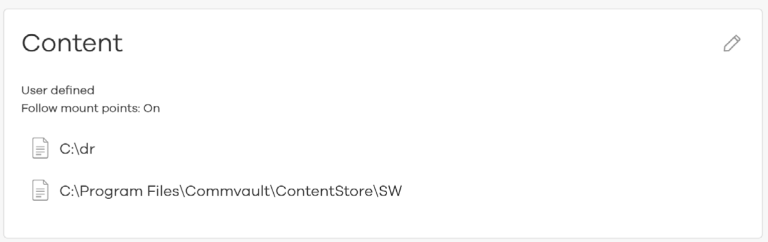

Content

User defined
Follow mount points: On

C:\dr

C:\Program Files\Commvault\ContentStore\SW

FIGURE 30 CommServe file system subclient for Auto Recovery

We associated the subclient to the same server backup plan as our other data. Backups land on the production FlashBlade and are copied to the CRV FlashBlade, with object lock applied on both targets. After creating the subclient, we ran a backup so the data would be available for Auto Recovery.

We then created two Auto Recovery replication groups. For each group, we selected one of the source paths as the content, set one of the CRV MediaAgents as the target (with the destination path on the file system we mounted earlier), and selected the CRV copy as the synchronization source (Figure 31).

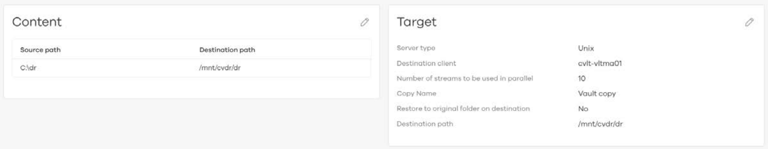

Content

Source path	Destination path
C:\dr	/mnt/cvdr/dr

Target

Server type	Unix
Destination client	cvlt-vltma01
Number of streams to be used in parallel	10
Copy Name	Vault copy
Restore to original folder on destination	No
Destination path	/mnt/cvdr/dr

FIGURE 31 Auto Recovery replication group configuration

Storage Arrays

Configuring Commvault IntelliSnap capabilities allowed us to enable extremely fast recovery for critical data sets. Commvault Cloud created application-consistent snapshots on the production FlashArray//X. The ActiveDR pairing between the production array and the SCRE FlashArray//C automatically replicated those snapshots whenever the air gap was allowed; when the air gap was blocked, the production array queued the changes until the FlashArray//C was reachable again.

To enable IntelliSnap, we first added the FlashArray//X to the managed arrays in Commvault Cloud. We configured the Remote Snap MA setting to make Commvault send all array commands through one of the production MediaAgents. We also enabled the setting to not track pod volumes separately between arrays. We left all the other settings as their defaults. Figure 32 shows all the configuration settings.

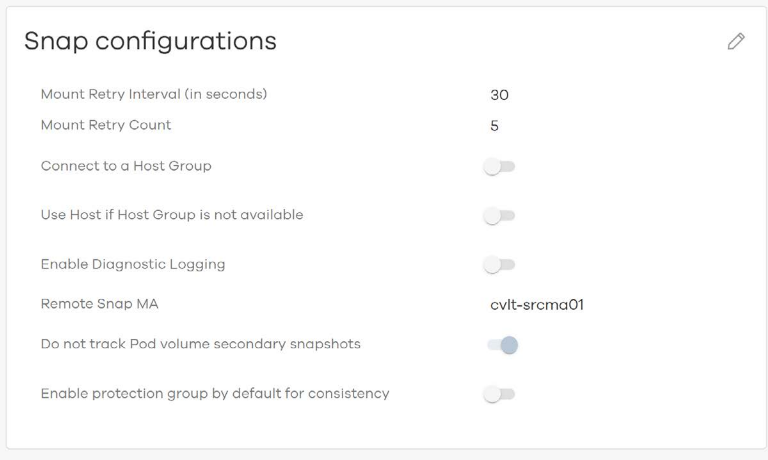

Snap configurations

Mount Retry Interval (in seconds)	30
Mount Retry Count	5
Connect to a Host Group	○
Use Host if Host Group is not available	○
Enable Diagnostic Logging	○
Remote Snap MA	cvlt-srcma01
Do not track Pod volume secondary snapshots	●
Enable protection group by default for consistency	○

FIGURE 32 Snap configuration settings for production FlashArray

Server Backup Plans

To manage data protection and transfer into the CRV, we created a server backup plan. We used the default recovery point objective (RPO) settings (Figure 33), so backups ran daily, synthetic full backups ran as needed, and secondary copies—which manage the auxiliary copies across the air gap—ran on an automatic schedule. For a production deployment, we would consider setting a more rigid secondary copy schedule to align with the air-gap schedule.

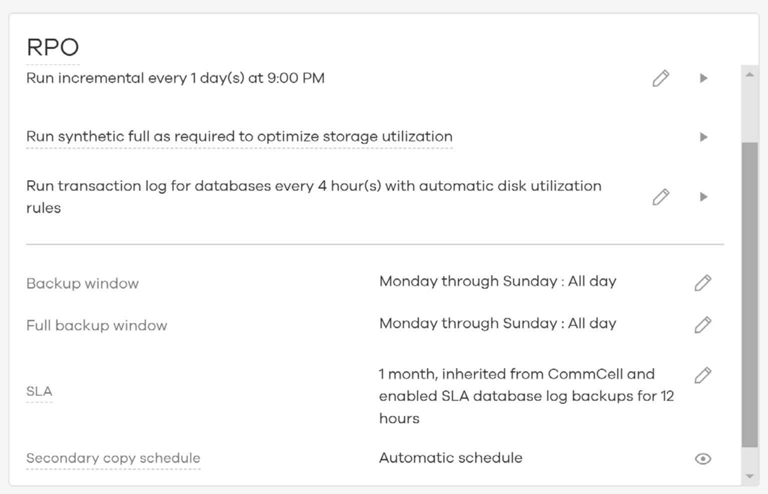

RPO

Run incremental every 1 day(s) at 9:00 PM

Run synthetic full as required to optimize storage utilization

Run transaction log for databases every 4 hour(s) with automatic disk utilization rules

Backup window	Monday through Sunday : All day
Full backup window	Monday through Sunday : All day
SLA	1 month, inherited from CommCell and enabled SLA database log backups for 12 hours
Secondary copy schedule	Automatic schedule

FIGURE 33 Server backup plan RPO settings

We also left the snapshot options (Figure 34) set to the plan defaults. For any snapshot-enabled data sets, Commvault Cloud would create FlashArray snapshots, back up those snapshots on a four-hour interval, and send the backup data to the CRV.

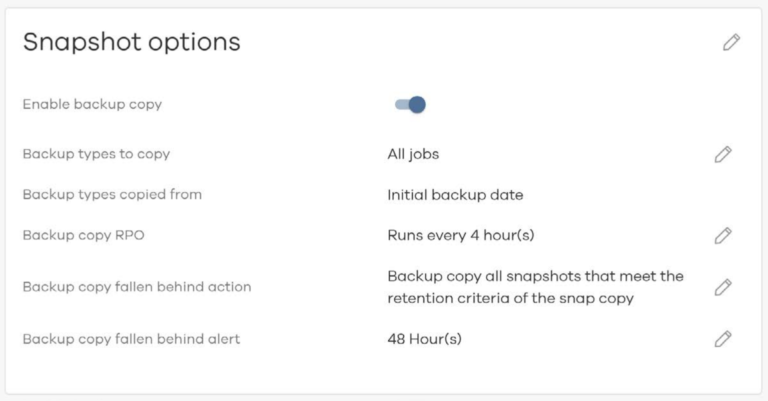

Snapshot options

Enable backup copy	●
Backup types to copy	All jobs
Backup types copied from	Initial backup date
Backup copy RPO	Runs every 4 hour(s)
Backup copy fallen behind action	Backup copy all snapshots that meet the retention criteria of the snap copy
Backup copy fallen behind alert	48 Hour(s)

FIGURE 34 Server backup plan snapshot options

We configured three backup destinations, or copies (Figure 35). The primary snap copy managed retention of snapshots on FlashArray//X, keeping them for four days. The production and vault copies used the object lock–enabled FlashBlade storage pools. Both FlashBlade copies inherited their retention settings from their respective storage pools. The snapshot copy uses the unlocked production pool; as of this writing, Commvault Cloud does not support WORM pools for snapshot index data.

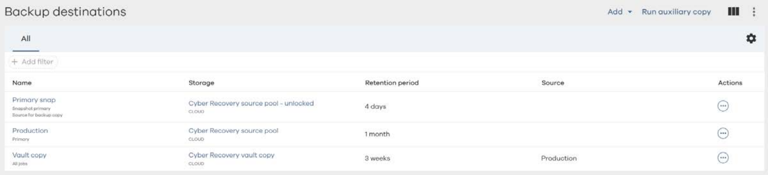

Backup destinations

Name	Storage	Retention period	Source	Actions
Primary snap Snapshot primary Source for backup copy	Cyber Recovery source pool - unlocked CLOUD	4 days		⋯
Production Primary	Cyber Recovery source pool CLOUD	1 month		⋯
Vault copy All jobs	Cyber Recovery vault copy CLOUD	3 weeks	Production	⋯

FIGURE 35 Server backup plan backup destinations

Hypervisor

We added the production vCenter as a hypervisor in Commvault Cloud and configured the production Virtual Server Agent VMs as its access nodes.

Data Sources

We created two VM groups under the vCenter hypervisor to manage the test backup jobs. The first group protected the VMs on the stand-alone datastore with streaming backups to FlashBlade.

The other group protected the replicated datastore. In this configuration, IntelliSnap created application-consistent FlashArray snapshots, which the production FlashArray//X automatically replicated to the SCRE FlashArray//C. Commvault Cloud then used the snapshot on FlashArray//X as the source for a streaming backup to FlashBlade, which it later copied to the CRV.

We set the same general options on both VM groups (Figure 36). We increased the number of readers to 10 to drive more parallelism in the backup jobs, and we kept the defaults for the other settings.

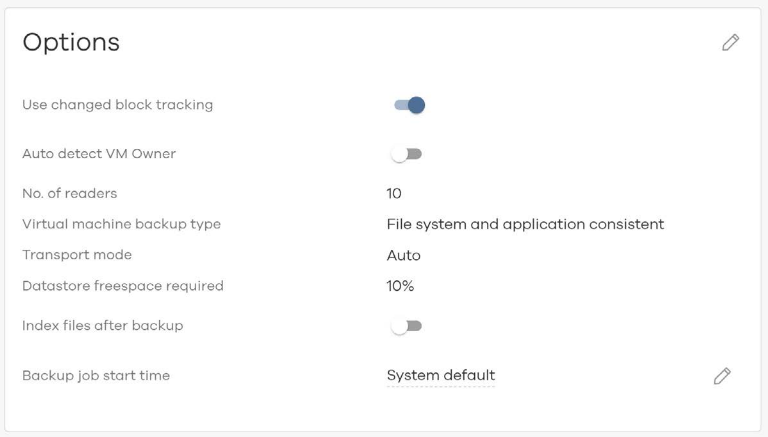

Options

Use changed block tracking	●
Auto detect VM Owner	○
No. of readers	10
Virtual machine backup type	File system and application consistent
Transport mode	Auto
Datastore freespace required	10%
Index files after backup	○
Backup job start time	System default

FIGURE 36 Commvault Cloud VM group general options

For the ActiveDR VM group, we also configured the snapshot management settings (Figure 37). We enabled snap backup, set the snapshot engine to Pure Storage FlashArray Snap, and defined a snap mount host. We did not need to define any specific arrays; Commvault Cloud automatically detects array relationships during snap backup jobs.

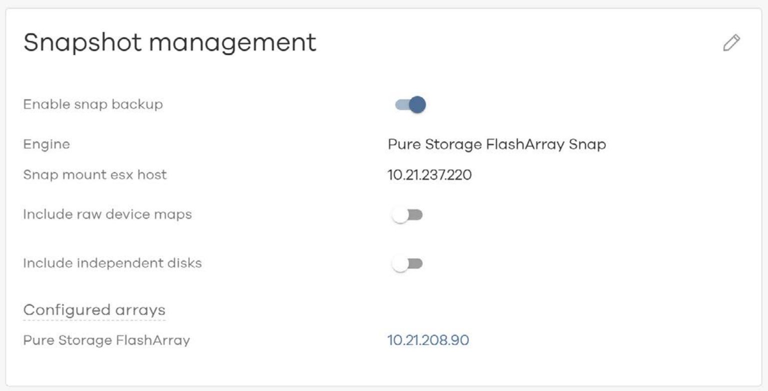

Snapshot management

Enable snap backup	●
Engine	Pure Storage FlashArray Snap
Snap mount esx host	10.21.237.220
Include raw device maps	○
Include independent disks	○

Configured arrays

Pure Storage FlashArray	10.21.208.90

FIGURE 37 VM group snapshot management settings

Threat Scan Nodes

We deployed two Threat Scan nodes in the production environment. Commvault threat analysis must run on the same operating system family as the source of the data being scanned, so we installed Threat Scan on one Windows VM and one Linux VM. We created a server group (Figure 38) containing both nodes that simplifies the threat analysis process.

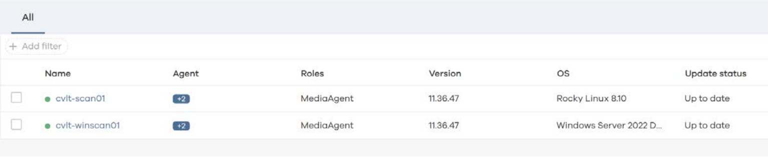

	Name	Agent	Roles	Version	OS	Update status
□	● cvlt-scan01	+2	MediaAgent	11.36.47	Rocky Linux 8.10	Up to date
□	● cvlt-winscan01	+2	MediaAgent	11.36.47	Windows Server 2022 D...	Up to date

FIGURE 38 Production Threat Scan server group

Validation Tests

To validate the recovery environment, we ran a series of tests. We confirmed the network isolation was working as intended to ensure the components were resilient to periodic disconnection. We ran threat detection with Threat Scan to ensure its functionality in our environment. Once the basic validation was complete, we simulated a cyber incident recovery, first by using FlashArray snapshots and ActiveDR for critical data recovery, and then by using Commvault Cloud to recover data from its backups to FlashBlade.

Network Isolation Tests

We ran several tests to confirm that the network isolation was working as expected. We ran port scans using nmap to confirm whether the NAT addresses exposed any ports. The following scan results showed that no ports were listening on the NAT address.

```
NSE: Script scanning 10.21.207.40.
Initiating NSE at 19:57
Completed NSE at 19:57, 5.01s elapsed
Initiating NSE at 19:57
Completed NSE at 19:57, 0.00s elapsed
Initiating NSE at 19:57
Completed NSE at 19:57, 0.00s elapsed
Nmap scan report for 10.21.207.40
Host is up (0.012s latency).
All 65535 scanned ports on 10.21.207.40 are in ignored states.
```

Next we ran readiness checks of the vault MediaAgents within Commvault Cloud. The checks passed (Figure 39), indicating that the one-way communication was working as expected.

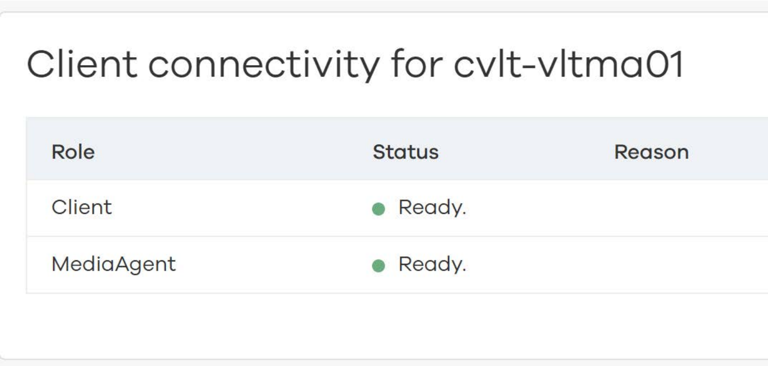

Client connectivity for cvlt-vltma01

Role	Status	Reason
Client	● Ready.	
MediaAgent	● Ready.	

FIGURE 39 MediaAgent readiness check results

To confirm that Commvault communication sessions are only allowed in one direction, we ran the Commvault cvping utility. First, we ran the test from the production MediaAgents that proxy communication between the CRV MediaAgents and the CommServe, attempting to connect to the services in the CRV. The connections failed as expected (see the following page).

```
[root@cvlt-srcma01 packer]# /opt/commvault/Base/cvping 10.21.207.40 8403
Input:
Interface : 10.21.207.40
Port      : 8403
Family    : IPv4
Thu Feb 20 06:01:52 2025
Trying to connect to 10.21.207.40(10.21.207.40) port 8403
Failed to connect to 10.21.207.40(10.21.207.40) port 8403.
Error: [Connection refused]
```

We then confirmed that the same test succeeded in the other direction.

```
[root@cvlt-vltma01 packer]# /opt/commvault/Base/cvping 10.21.208.11 8403
Input:
Interface : 10.21.208.11
Port      : 8403
Family    : IPv4
Wed Feb 19 22:10:56 2025
Trying to connect to 10.21.208.11(10.21.208.11) port 8403
Successfully connected to 10.21.208.11(10.21.208.11) port 8403.
```

We confirmed from the production MediaAgents that the tunnel was establishing Transfer Control Protocol (TCP) sessions and that they were initiating from the NAT address (shown on the following page).

```
[root@cvlt-srcma01 packer]# ss | grep 10.21.207.40
tcp   ESTAB 0      0                 10.21.208.11:admind2   10.21.207.40:45280
tcp   ESTAB 0      0                 10.21.208.11:admind2   10.21.207.40:42341
tcp   ESTAB 0      0                 10.21.208.11:admind2   10.21.207.40:13888
tcp   ESTAB 0      0                 10.21.208.11:admind2   10.21.207.40:44659
tcp   ESTAB 0      0                 10.21.208.11:admind2   10.21.207.40:24039
tcp   ESTAB 0      0                 10.21.208.11:admind2   10.21.207.40:65092
tcp   ESTAB 0      0                 10.21.208.11:admind2   10.21.207.40:9581
tcp   ESTAB 0      0                 10.21.208.11:admind2   10.21.207.40:29877
tcp   ESTAB 0      0                 10.21.208.11:admind2   10.21.207.40:49858
tcp   ESTAB 0      0                 10.21.208.11:admind2   10.21.207.40:39250
tcp   ESTAB 0      0                 10.21.208.11:admind2   10.21.207.40:22687
tcp   ESTAB 0      0                 10.21.208.11:admind2   10.21.207.40:21409
tcp   ESTAB 0      0                 10.21.208.11:admind2   10.21.207.40:34442
tcp   ESTAB 0      0                 10.21.208.11:admind2   10.21.207.40:40963
tcp   ESTAB 0      0                 10.21.208.11:admind2   10.21.207.40:20910
tcp   ESTAB 0      0                 10.21.208.11:admind2   10.21.207.40:59150
```

We then severed the air-gap connection. Commvault Cloud was no longer able to communicate to or from the CRV systems. Figure 40 shows the readiness test failing for lack of connectivity.

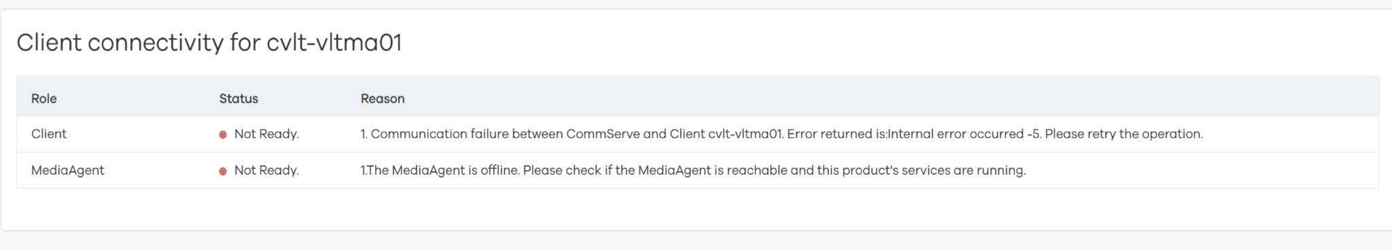

Client connectivity for cvlt-vltma01

Role	Status	Reason
Client	● Not Ready.	1. Communication failure between CommServe and Client cvlt-vltma01. Error returned is:Internal error occurred -5. Please retry the operation.
MediaAgent	● Not Ready.	1.The MediaAgent is offline. Please check if the MediaAgent is reachable and this product's services are running.

FIGURE 40 Failed MediaAgent readiness check

The cvping test failed from the CRV to a production MediaAgent, as shown below.

```
[root@cvlt-vltma01 packer]# /opt/commvault/Base/cvping 10.21.208.11 8403
Input:
Interface : 10.21.208.11
Port      : 8403
Family    : IPv4
Thu Feb 20 09:06:32 2025
Trying to connect to 10.21.208.11(10.21.208.11) port 8403
Failed to connect to 10.21.208.11(10.21.208.11) port 8403.
Error: [Connection timed out]
```

Severing the air-gap connection made the CRV storage pools appear offline (Figure 41), as expected.

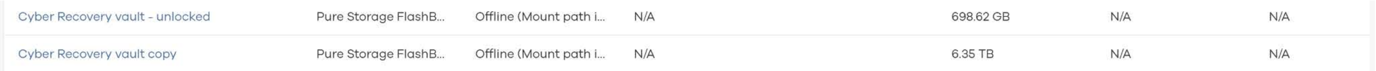

Cyber Recovery vault - unlocked	Pure Storage FlashB...	Offline (Mount path i...	N/A	698.62 GB	N/A	N/A
Cyber Recovery vault copy	Pure Storage FlashB...	Offline (Mount path i...	N/A	6.35 TB	N/A	N/A

FIGURE 41 Offline CRV storage pools

The scheduled auxiliary copy job still ran, but without the MediaAgents available, it stayed in the waiting state (Figure 42).

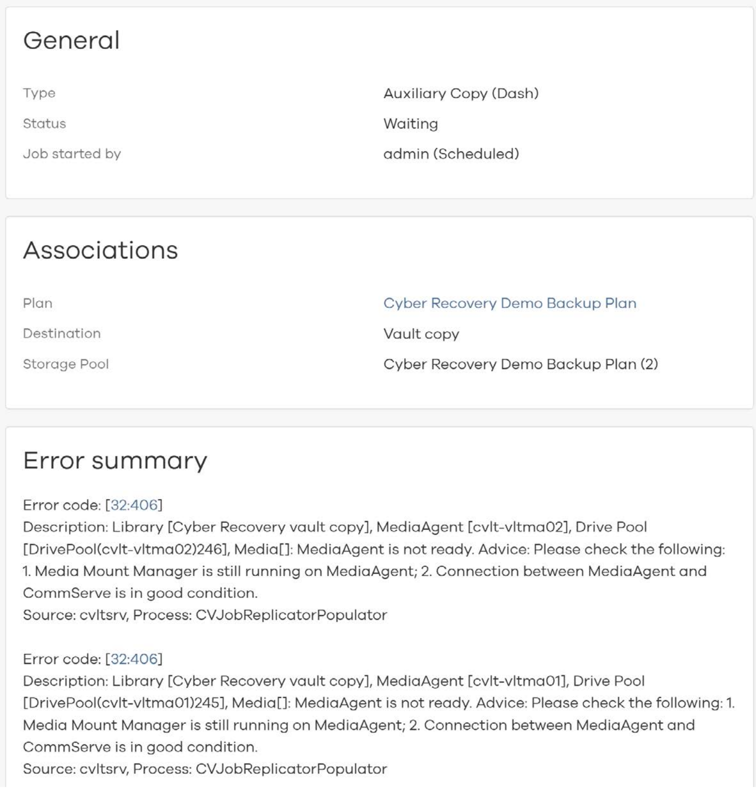

General

Type	Auxiliary Copy (Dash)
Status	Waiting
Job started by	admin (Scheduled)

Associations

Plan	Cyber Recovery Demo Backup Plan
Destination	Vault copy
Storage Pool	Cyber Recovery Demo Backup Plan (2)

Error summary

Error code: [32:406]
Description: Library [Cyber Recovery vault copy], MediaAgent [cvlt-vltma02], Drive Pool [DrivePool(cvlt-vltma02)246], Media[]: MediaAgent is not ready. Advice: Please check the following: 1. Media Mount Manager is still running on MediaAgent; 2. Connection between MediaAgent and CommServe is in good condition.
Source: cvltsrv, Process: CVJobReplicatorPopulator

Error code: [32:406]
Description: Library [Cyber Recovery vault copy], MediaAgent [cvlt-vltma01], Drive Pool [DrivePool(cvlt-vltma01)245], Media[]: MediaAgent is not ready. Advice: Please check the following: 1. Media Mount Manager is still running on MediaAgent; 2. Connection between MediaAgent and CommServe is in good condition.
Source: cvltsrv, Process: CVJobReplicatorPopulator

FIGURE 42 Auxiliary copy job in waiting state

Next, we reestablished the air-gap connection. To get the storage pool online without waiting for the Commvault Cloud storage refresh interval, we restarted the Commvault services on the CRV MediaAgents. This forced a storage status update, and the storage pools came back online right away (Figure 43).

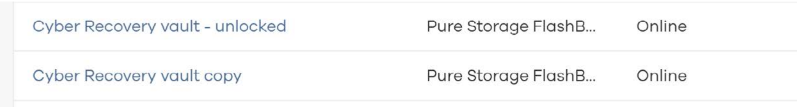

Cyber Recovery vault - unlocked	Pure Storage FlashB...	Online
Cyber Recovery vault copy	Pure Storage FlashB...	Online

FIGURE 43 CRV storage pools returned to online status

The waiting auxiliary copy job resumed automatically, and it was able to complete successfully (Figure 44).

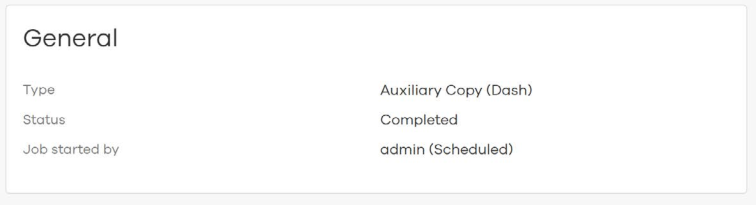

General

Type	Auxiliary Copy (Dash)
Status	Completed
Job started by	admin (Scheduled)

FIGURE 44 Completed auxiliary copy job

Array Replication

We also tested ActiveDR replication with the air gap open and closed. When we severed the air-gap connection, both FlashArrays showed the ActiveDR relationship as unhealthy (Figures 45 and 46).

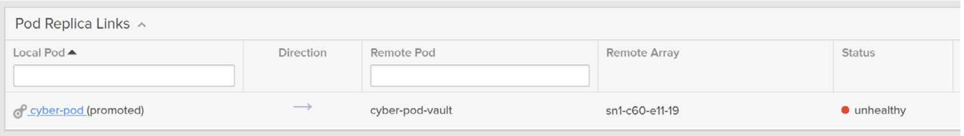

Local Pod	Direction	Remote Pod	Remote Array	Status
cyber-pod (promoted)	→	cyber-pod-vault	sn1-c60-e11-19	● unhealthy

FIGURE 45 Production FlashArray unhealthy ActiveDR replication

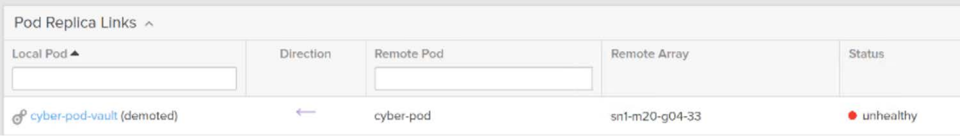

Local Pod	Direction	Remote Pod	Remote Array	Status
cyber-pod-vault (demoted)	←	cyber-pod	sn1-m20-g04-33	● unhealthy

FIGURE 46 SCRE FlashArray unhealthy ActiveDR replication

While the air gap was severed, we were able to use IntelliSnap to create snapshots on the production array (Figure 47), but they were unable to replicate to the SCRE.

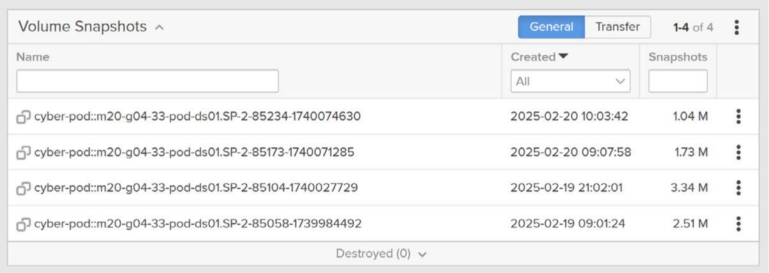

Name	Created	Snapshots
cyber-pod::m20-g04-33-pod-ds01.SP-2-85234-1740074630	2025-02-20 10:03:42	1.04 M
cyber-pod::m20-g04-33-pod-ds01.SP-2-85173-1740071285	2025-02-20 09:07:58	1.73 M
cyber-pod::m20-g04-33-pod-ds01.SP-2-85104-1740027729	2025-02-19 21:02:01	3.34 M
cyber-pod::m20-g04-33-pod-ds01.SP-2-85058-1739984492	2025-02-19 09:01:24	2.51 M

FIGURE 47 Production snapshots created with air gap severed

Once the air gap was reconnected, ActiveDR automatically resumed replication (Figure 48).

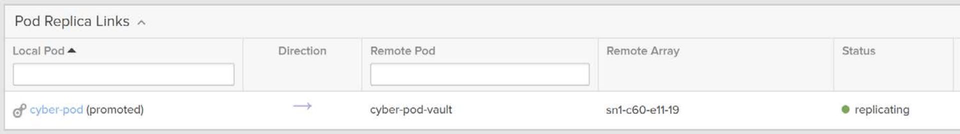

Local Pod	Direction	Remote Pod	Remote Array	Status
cyber-pod (promoted)	→	cyber-pod-vault	sn1-c60-e11-19	● replicating

FIGURE 48 ActiveDR replication resumed

Snapshots synchronized within a few minutes and became available in the SCRE for recovery (Figure 49).

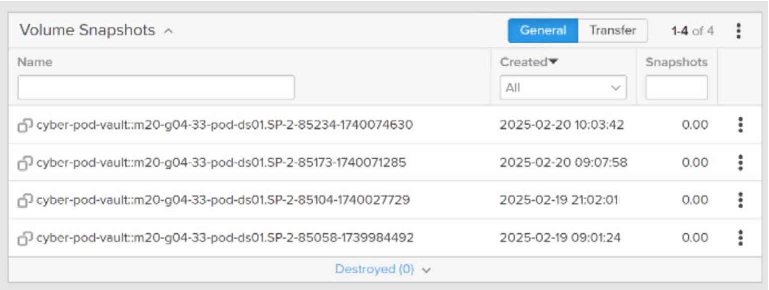

Name	Created	Snapshots
cyber-pod-vault::m20-g04-33-pod-ds01.SP-2-85234-1740074630	2025-02-20 10:03:42	0.00
cyber-pod-vault::m20-g04-33-pod-ds01.SP-2-85173-1740071285	2025-02-20 09:07:58	0.00
cyber-pod-vault::m20-g04-33-pod-ds01.SP-2-85104-1740027729	2025-02-19 21:02:01	0.00
cyber-pod-vault::m20-g04-33-pod-ds01.SP-2-85058-1739984492	2025-02-19 09:01:24	0.00

FIGURE 49 Replicated FlashArray volume snapshots after synchronization

Threat Detection

We tested Commvault Cloud Threat Scan in production and the SCRE by performing out-of-place recoveries of VMs. First, we tested in the production environment. For the recovery destination options (Figure 50), we chose the production vCenter. From the Access node drop-down, we chose the Threat Scan server group we created earlier, not an individual node. This approach allowed us to recover Windows and Linux VMs in a single job, with Commvault Cloud automatically assigning VMs to appropriate Threat Scan nodes based on the operating system.

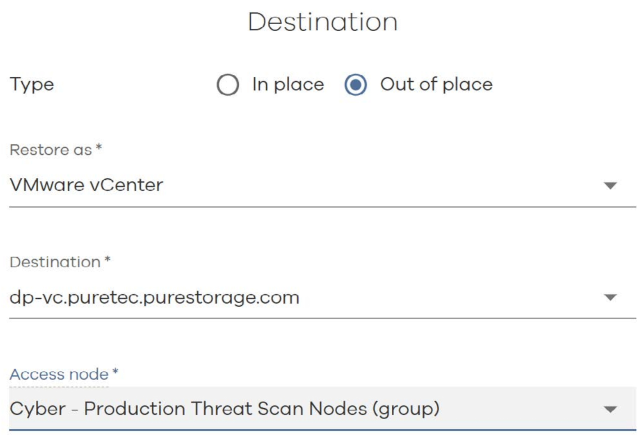

Destination

Type ○ In place ● Out of place

Restore as *
VMware vCenter

Destination *
dp-vc.puretec.purestorage.com

Access node *
Cyber - Production Threat Scan Nodes (group)

FIGURE 50 Recovery destination options

We configured additional restore options for the recovered VMs (Figure 51). Since the original VMs were still running, we added a suffix to avoid any naming collisions. We configured a destination host, set the destination datastore to the one we created specifically for Threat Scan, and configured a VM folder to contain the VMs in vCenter.

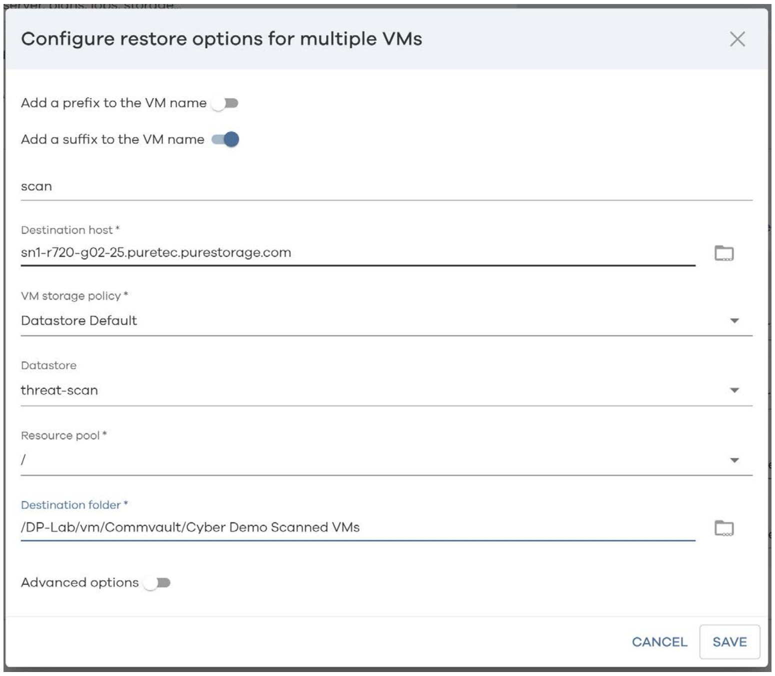

Configure restore options for multiple VMs

Add a prefix to the VM name (off)

Add a suffix to the VM name (on)

scan

Destination host *
sn1-r720-g02-25.puretec.purestorage.com

VM storage policy *
Datastore Default

Datastore
threat-scan

Resource pool *
/

Destination folder *
/DP-Lab/vm/Commvault/Cyber Demo Scanned VMs

Advanced options (off)

FIGURE 51 Additional VM restore options

In the post-restore options (Figure 52), we chose not to power on the VMs since we did not want any risk of network conflicts between the scanned VMs and their sources. We enabled the Run Threat Analysis option to trigger Threat Scan as part of the restore job.

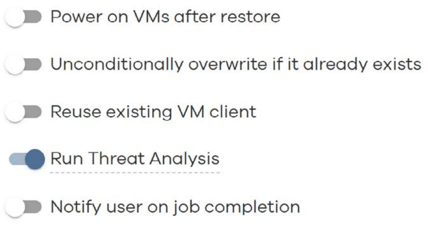

○ Power on VMs after restore

○ Unconditionally overwrite if it already exists

○ Reuse existing VM client

● Run Threat Analysis

○ Notify user on job completion

FIGURE 52 Post-restore options

Enabling threat analysis expanded the VM information in the restore job details (Figure 53) to include scan statistics.

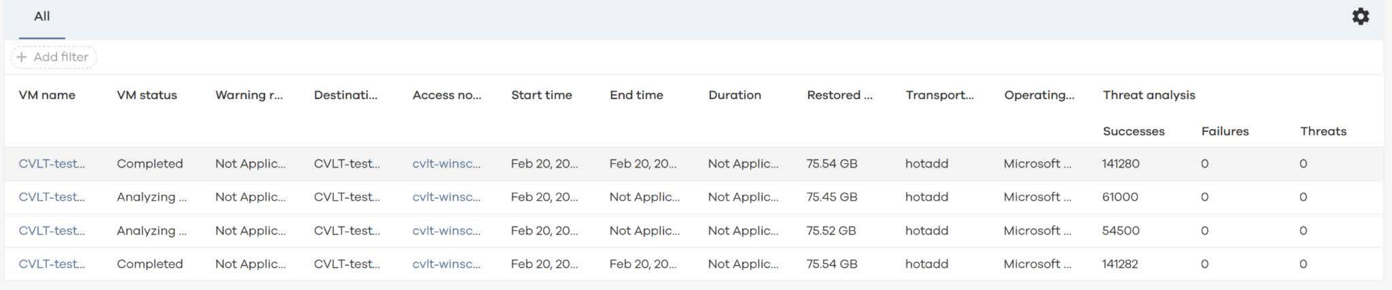

VM name	VM status	Warning r...	Destinati...	Access no...	Start time	End time	Duration	Restored ...	Transport...	Operating...	Threat analysis		
											Successes	Failures	Threats
CVLT-test...	Completed	Not Applic...	CVLT-test...	cvlt-winsc...	Feb 20, 20...	Feb 20, 20...	Not Applic...	75.54 GB	hotadd	Microsoft ...	141280	0	0
CVLT-test...	Analyzing ...	Not Applic...	CVLT-test...	cvlt-winsc...	Feb 20, 20...	Not Applic...	Not Applic...	75.45 GB	hotadd	Microsoft ...	61000	0	0
CVLT-test...	Analyzing ...	Not Applic...	CVLT-test...	cvlt-winsc...	Feb 20, 20...	Not Applic...	Not Applic...	75.52 GB	hotadd	Microsoft ...	54500	0	0
CVLT-test...	Completed	Not Applic...	CVLT-test...	cvlt-winsc...	Feb 20, 20...	Feb 20, 20...	Not Applic...	75.54 GB	hotadd	Microsoft ...	141282	0	0

FIGURE 53 VM details with Threat Scan statistics

One of our test VMs had a positive scan result. When Threat Scan detected malicious files, it noted the number detected and generated a report on the Monitoring/Threat indicators page (Figure 54). Had we chosen the option to power on VMs, Commvault Cloud would have left the affected VM powered off to prevent malware from spreading.

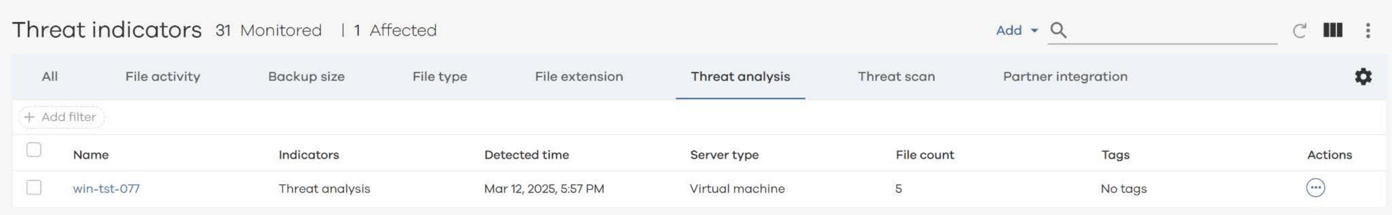

Threat indicators 31 Monitored | 1 Affected

□	Name	Indicators	Detected time	Server type	File count	Tags	Actions
□	win-tst-077	Threat analysis	Mar 12, 2025, 5:57 PM	Virtual machine	5	No tags	⋯

FIGURE 54 Detected threat on Threat indicators page

The report (Figure 55) listed the detected threats and the recovery point where they were detected.

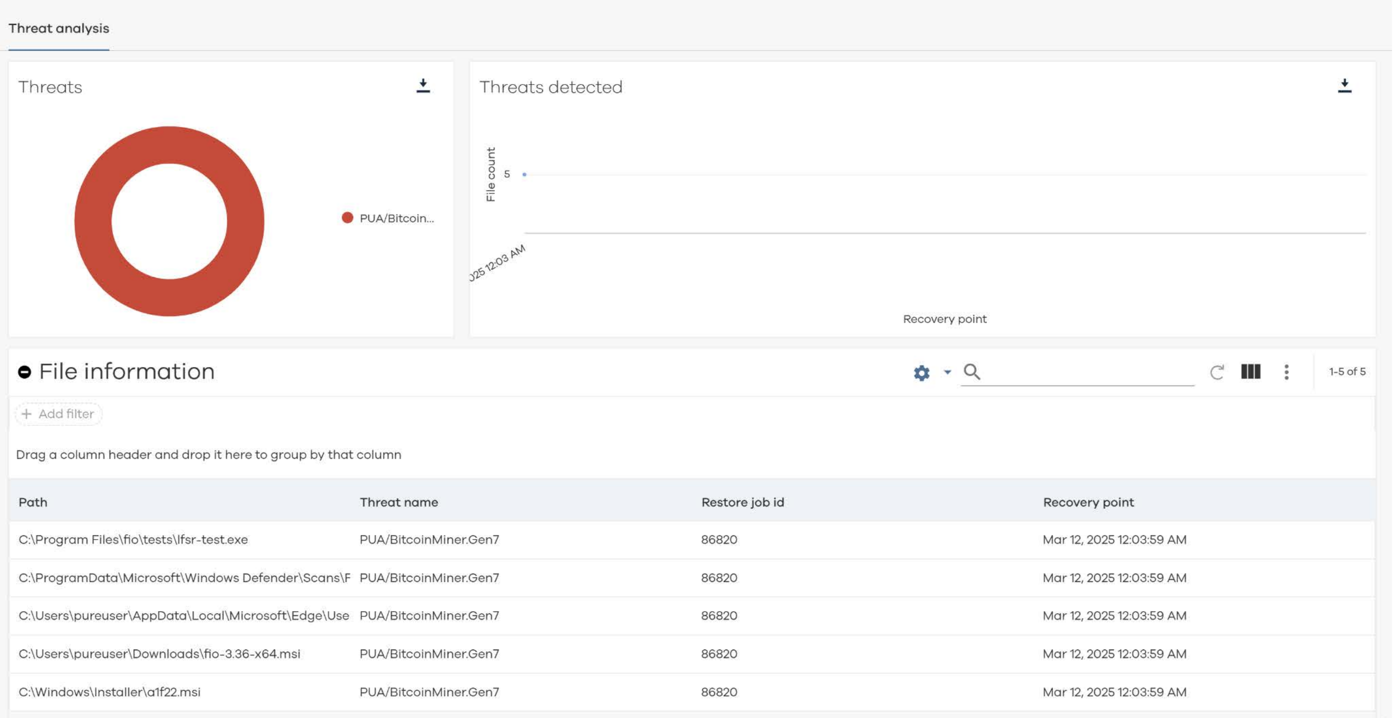

File information

Drag a column header and drop it here to group by that column

Path	Threat name	Restore job id	Recovery point
C:\Program Files\fio\tests\lfsr-test.exe	PUA/BitcoinMiner.Gen7	86820	Mar 12, 2025 12:03:59 AM
C:\ProgramData\Microsoft\Windows Defender\Scans\F	PUA/BitcoinMiner.Gen7	86820	Mar 12, 2025 12:03:59 AM
C:\Users\pureuser\AppData\Local\Microsoft\Edge\Use	PUA/BitcoinMiner.Gen7	86820	Mar 12, 2025 12:03:59 AM
C:\Users\pureuser\Downloads\fio-3.36-x64.msi	PUA/BitcoinMiner.Gen7	86820	Mar 12, 2025 12:03:59 AM
C:\Windows\Installer\a1f22.msi	PUA/BitcoinMiner.Gen7	86820	Mar 12, 2025 12:03:59 AM

FIGURE 55 Detected threat report

Critical Data Recovery with ActiveDR

Before beginning any recovery tests, we severed the air-gap connection to prevent the production environment from introducing any data changes. On the SCRE FlashArray, we copied the latest replicated snapshot to a new volume located outside the pod (Figure 56). In the copy options, we chose not to enable default protection with SafeMode so it would not affect cleanup between tests; enabling it would not have interfered with any of the recovery steps.

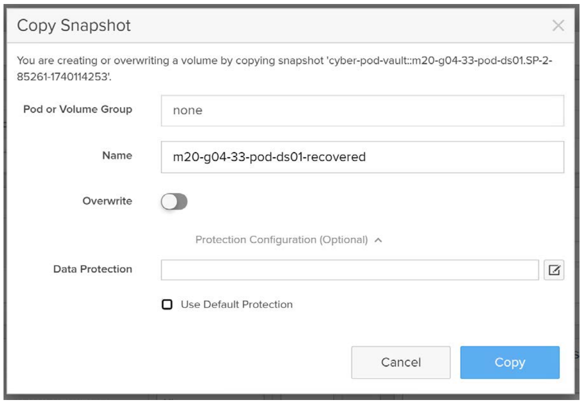

Copy Snapshot

You are creating or overwriting a volume by copying snapshot 'cyber-pod-vault::m20-g04-33-pod-ds01.SP-2-85261-1740114253'.

Pod or Volume Group: none

Name: m20-g04-33-pod-ds01-recovered

Overwrite: ○

Protection Configuration (Optional)

Data Protection:

□ Use Default Protection

FIGURE 56 Copying FlashArray snapshot to new volume

We connected the new volume to our ESXi host group (Figure 57).

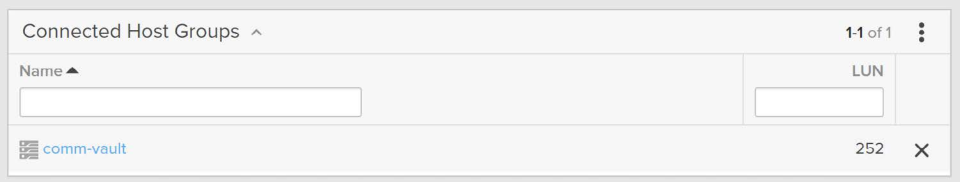

Connected Host Groups

Name	LUN
comm-vault	252

FIGURE 57 Copied FlashArray volume connected to host group

We performed a storage rescan on the vSphere cluster, detecting both new devices and VMFS volumes. vSphere flagged the copied volume as a snapshot, so this process couldn't automatically bring it online as a datastore. To resolve the issue, we ran the vCenter wizard to create a new datastore, selecting the copied volume (Figure 58).

Note: The datastore name in this step is arbitrary. In the next step, vSphere will import the name from the original datastore.

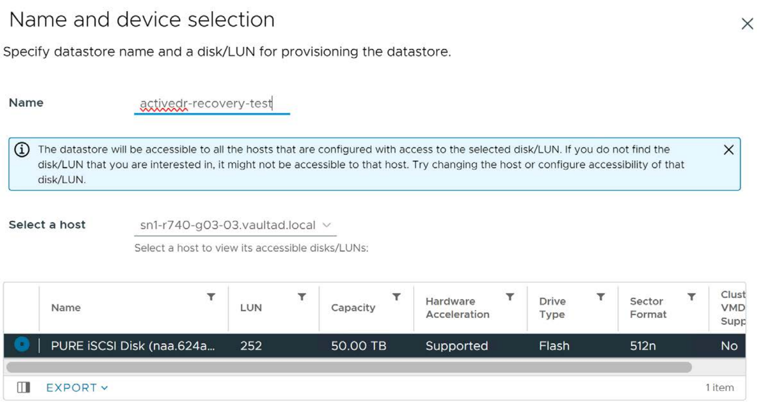

Name and device selection

Specify datastore name and a disk/LUN for provisioning the datastore.

Name: activedr-recovery-test

ⓘ The datastore will be accessible to all the hosts that are configured with access to the selected disk/LUN. If you do not find the disk/LUN that you are interested in, it might not be accessible to that host. Try changing the host or configure accessibility of that disk/LUN.

Select a host: sn1-r740-g03-03.vaultad.local

Select a host to view its accessible disks/LUNs:

	Name	LUN	Capacity	Hardware Acceleration	Drive Type	Sector Format	Clust VMD Supp
●	PURE iSCSI Disk (naa.624a...	252	50.00 TB	Supported	Flash	512n	No

FIGURE 58 New datastore name and device selection

vSphere detected the existing VMFS volume and prompted us for mount options. Since there is no possibility of communication with the production vSphere environment, we chose to keep the existing signature (Figure 59). This also preserves the original datastore name.

Note: In our lab, vSphere would change the datastore name when we chose to assign a new signature.

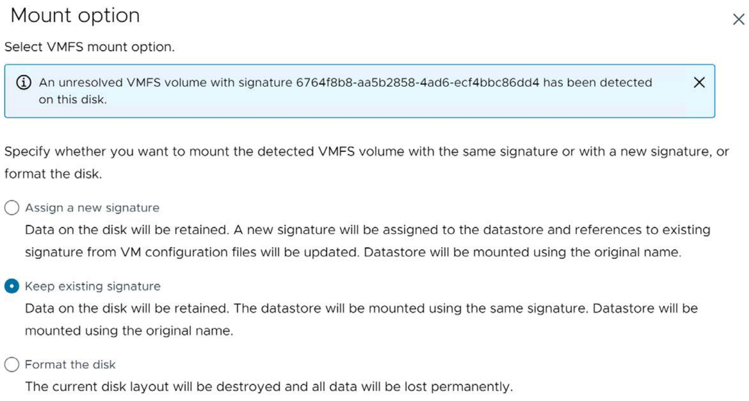

Mount option

Select VMFS mount option.

ⓘ An unresolved VMFS volume with signature 6764f8b8-aa5b2858-4ad6-ecf4bbc86dd4 has been detected on this disk.

Specify whether you want to mount the detected VMFS volume with the same signature or with a new signature, or format the disk.

○ Assign a new signature
Data on the disk will be retained. A new signature will be assigned to the datastore and references to existing signature from VM configuration files will be updated. Datastore will be mounted using the original name.

● Keep existing signature
Data on the disk will be retained. The datastore will be mounted using the same signature. Datastore will be mounted using the original name.

○ Format the disk
The current disk layout will be destroyed and all data will be lost permanently.

FIGURE 59 Datastore mount options

Once the datastore became available, we browsed its files and used the Register VM action to add the VMs. From the folder list in the left pane, we selected the VM folder. From the right pane, we selected the VMX file (Figure 60) and clicked the OK button to start the Register Virtual Machine wizard.

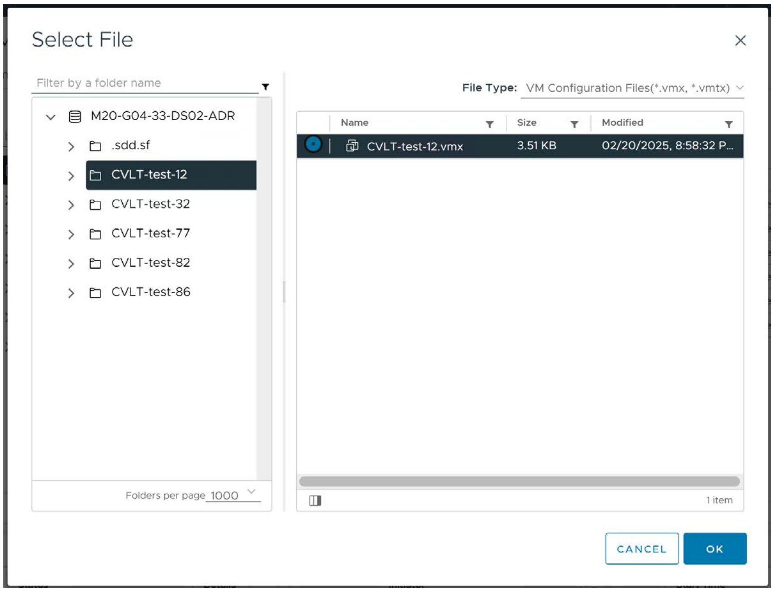

File Type: VM Configuration Files(*.vmx, *.vmtx)

Name	Size	Modified
CVLT-test-12.vmx	3.51 KB	02/20/2025, 8:58:32 P...

FIGURE 60 Selecting VMX file to register VM

We registered the VMs into a VM folder we had created. We selected the entire vSphere cluster as the compute resource to allow the distributed resource scheduler to choose host placement based on ESXi host load.

Once the VMs were registered, we powered them on (Figure 61). After this stage, in an actual recovery, the VMs would be turned over to the team performing forensic analysis or, if you determined analysis was not necessary, to the application owner or database administrator to complete recovery and validation.

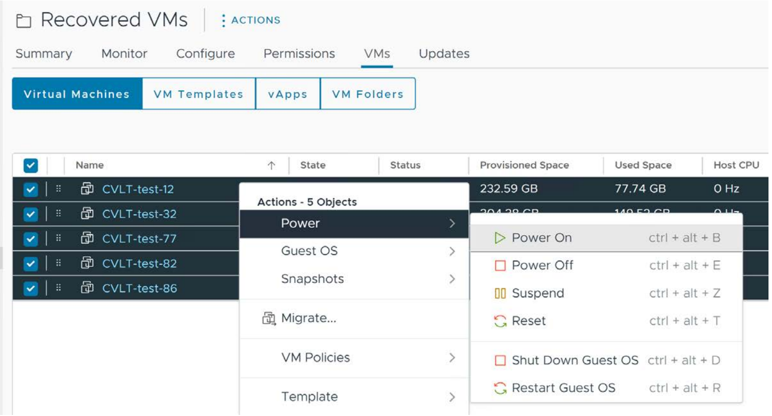

Name	State	Status	Provisioned Space	Used Space	Host CPU
CVLT-test-12			232.59 GB	77.74 GB	0 Hz
CVLT-test-32					
CVLT-test-77					
CVLT-test-82					
CVLT-test-86					

FIGURE 61 Powering on recovered VMs

CommCell Recovery

To validate recovery with Commvault Cloud, we first had to recover the CommCell itself within the SMZ.

To begin, we deployed the Commvault Cloud 11.36 Windows virtual template from the Commvault Store. We could have used the newer 11.38 template, but we decided it was simpler to match our production CommCell version. (Accessing the virtual templates requires a store login.) We logged into the deployed VM connected to the replicated software cache on the FlashBlade file system and used the cache to update the new CommServe to the latest maintenance release. We also copied the latest replicated disaster recovery backup to the new CommServe and used the CSRecoveryAssistant utility to recover the CommCell. After recovery, we disabled the scheduler, data backup, and data aging (Figure 62).

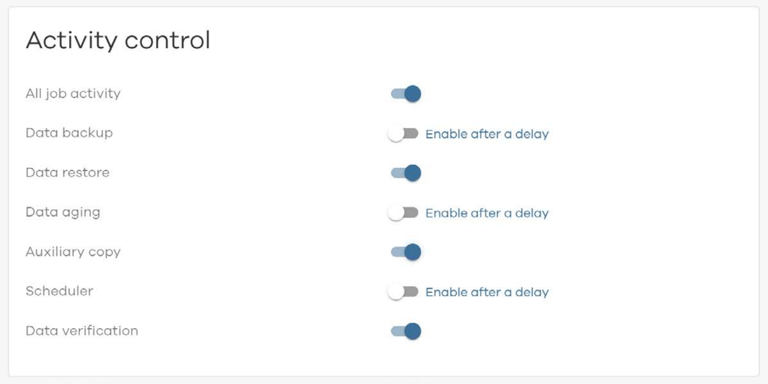

Activity control

All job activity	●
Data backup	○ Enable after a delay
Data restore	●
Data aging	○ Enable after a delay
Auxiliary copy	●
Scheduler	○ Enable after a delay
Data verification	●

FIGURE 62 Disabled CommCell activities

With the CommServe recovered, we prepared to deploy recovery data movers. To begin, we used vSphere to deploy Windows and Linux VMs from prepared templates. We then copied the replicated software cache from the FlashBlade to the local cache on the recovery CommServe. Next, we used Commvault Command Center to push the Commvault software packages to the VMs. All VMs received the Virtual Server package along with the required packages for Threat Scan, and the Linux VMs also received the MediaAgent package.

Once the software installations were complete, we used Command Center to add the new Linux MediaAgents to the vault storage pools' buckets (Figure 63). Even though the FlashBlade access policy only grants read access, we configured the access in Commvault Cloud as read/write; at least one MediaAgent must have read/write access in the storage pool configuration for the pool to come online.

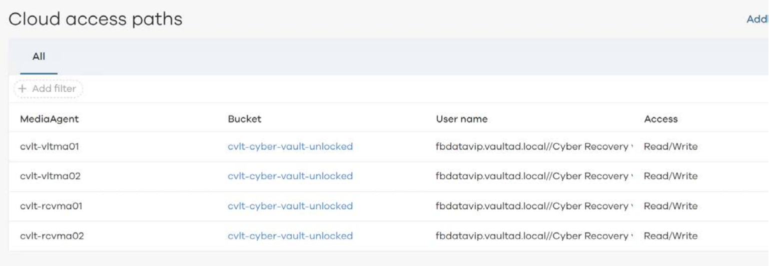

Cloud access paths

MediaAgent	Bucket	User name	Access
cvlt-vltma01	cvlt-cyber-vault-unlocked	fbdatavip.vaultad.local//Cyber Recovery '	Read/Write
cvlt-vltma02	cvlt-cyber-vault-unlocked	fbdatavip.vaultad.local//Cyber Recovery '	Read/Write
cvlt-rcvma01	cvlt-cyber-vault-unlocked	fbdatavip.vaultad.local//Cyber Recovery '	Read/Write
cvlt-rcvma02	cvlt-cyber-vault-unlocked	fbdatavip.vaultad.local//Cyber Recovery '	Read/Write

FIGURE 63 Storage pool with additional MediaAgent access

Similar to our production configuration, we created a server group (Figure 64) that we used to assign hypervisor access nodes and execute Threat Scan during VM recovery. We added all three VMs to the group.

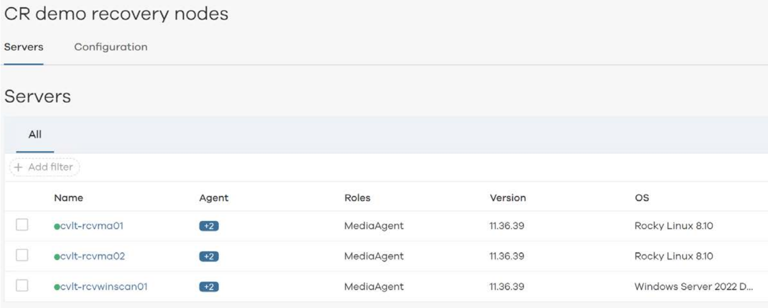

CR demo recovery nodes

Servers

	Name	Agent	Roles	Version	OS
□	cvlt-rcvma01	+2	MediaAgent	11.36.39	Rocky Linux 8.10
□	cvlt-rcvma02	+2	MediaAgent	11.36.39	Rocky Linux 8.10
□	cvlt-rcvwinscan01	+2	MediaAgent	11.36.39	Windows Server 2022 D...

FIGURE 64 Recovery server group in Commvault Cloud

Next, we added our SMZ vCenter instance as a hypervisor (Figure 65), giving Commvault Cloud a place to restore our VMs. From the Access nodes drop-down, we selected the server group we had just created. We skipped the step for creating VM groups, as we did not plan to take new backups in this CommCell.

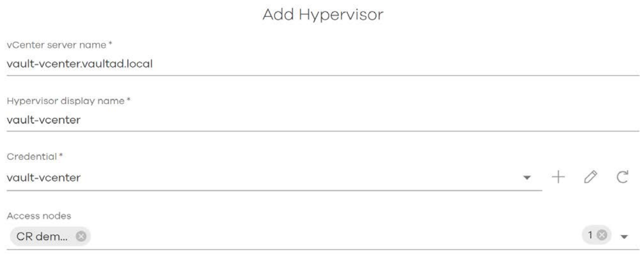

Add Hypervisor

vCenter server name *
vault-vcenter.vaultad.local

Hypervisor display name *
vault-vcenter

Credential *
vault-vcenter

Access nodes
CR dem... 1

FIGURE 65 Adding SMZ vCenter to Commvault Cloud

Data Recovery with Commvault Cloud

To test recovery from backup, we restored the VMs that came from the stand-alone datastore. While Commvault Cloud supports a number of ways to restore VMs, bulk recovery is easiest when starting from the VM group. In Commvault Command Center, we navigated to the Protect/Virtualization page and accessed the VM groups tab. We located the source VM group, clicked its Action button, and selected Restore from the context menu.

In the recovery wizard, we changed the recovery source (Figure 66) to the CRV copy and the MediaAgent to one we had just deployed.

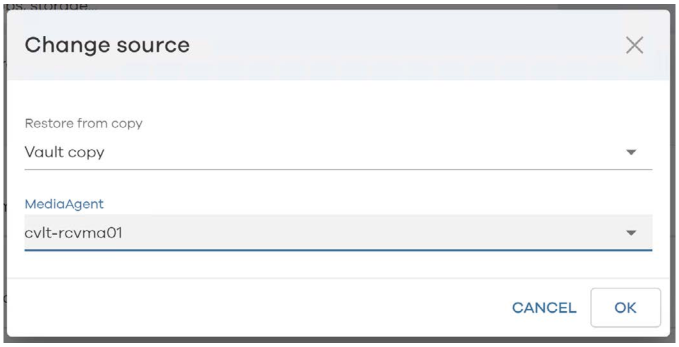

Change source

Restore from copy
Vault copy

MediaAgent
cvlt-rcvma01

FIGURE 66 Changing recovery source

We selected the option to recover the full VM. On the VM selection page, we selected all VMs for recovery and clicked the Restore button. On the Destination page (Figure 67), we selected the SCRE vCenter as the destination hypervisor and the recovery server group as the access node.

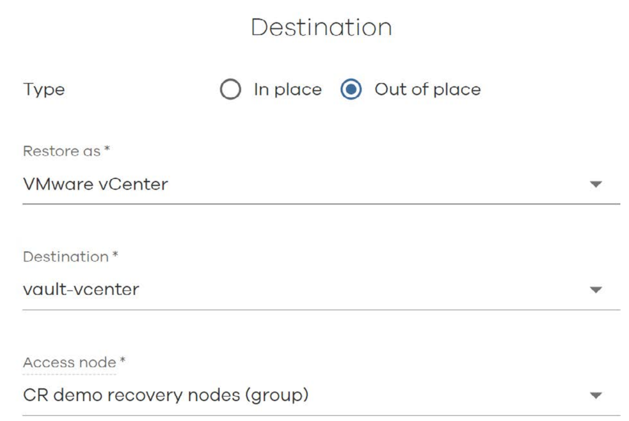

Destination

Type ○ In place ● Out of place

Restore as *
VMware vCenter

Destination *
vault-vcenter

Access node *
CR demo recovery nodes (group)

FIGURE 67 Recovery destination options

We customized the network and storage destination for each VM as part of the recovery. To make the required changes, we configured the VMs individually by selecting each VM and clicking the Configure restore options link. In the restore options form (Figure 68), we changed the destination host, datastore, destination folder, and destination network to match the recovery environment.

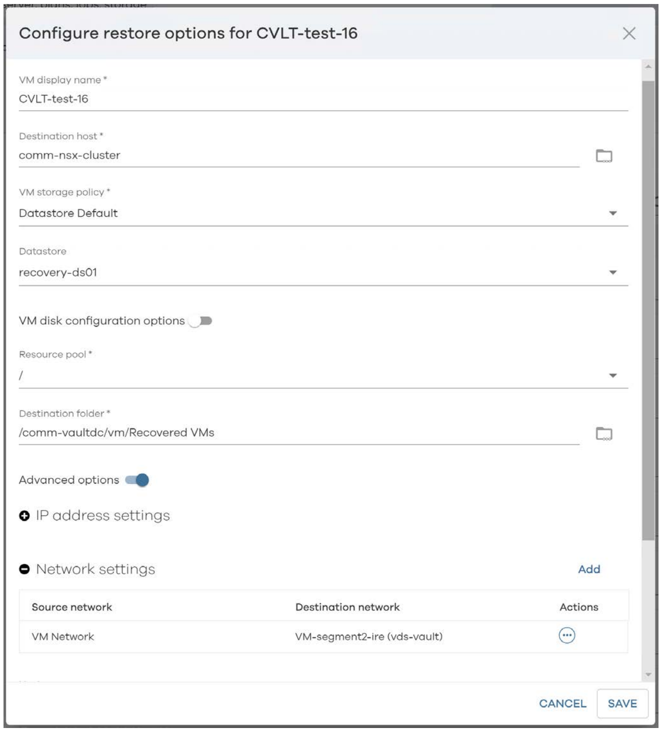

Configure restore options for CVLT-test-16

VM display name *
CVLT-test-16

Destination host *
comm-nsx-cluster

VM storage policy *
Datastore Default

Datastore
recovery-ds01

VM disk configuration options

Resource pool *
/

Destination folder *
/comm-vaultdc/vm/Recovered VMs

Advanced options

IP address settings

Network settings

Source network	Destination network	Actions
VM Network	VM-segment2-ire (vds-vault)	⋯

FIGURE 68 VM recovery options

The VM recovery completed successfully (Figure 69).

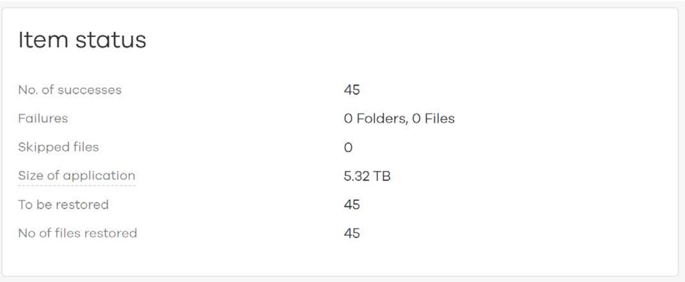

Item status

No. of successes	45
Failures	0 Folders, 0 Files
Skipped files	0
Size of application	5.32 TB
To be restored	45
No of files restored	45

FIGURE 69 Successful VM recovery

SCRE Threat Detection

While Threat Scan is designed to run primarily in the production environment, we tested to make sure it would function in a cyber recovery incident. We repeated the same production Threat Scan tests in the recovery environment following the same procedure as before but selecting the SCRE vCenter, datastore, and server group. We got comparable results to the production test, showing that Threat Scan can be part of the recovery process where needed.

Optimizing for Your Environment

Several parts of the solution require repetitive or labor-intensive tasks to maintain data protection, ensure data security, and enable recovery after a cyber incident. As part of helping develop your cyber resilience policies and procedures, a systems integrator can assist with customizing and automating these tasks to best fit your operations and optimize recovery times. Key areas to consider include:

- Air-gap state management
- Threat response
- VM recovery from ActiveDR replicas
- Building and customizing operating system templates and images
- Deploying and configuring Commvault Cloud recovery systems
- Application recovery plans and automation

Conclusion

Architecture alone does not and cannot make you compliant with DORA or any operational resilience regulation. Architecture is the critical foundation for compliance, but there's much more. DORA and similar regulations contain many articles, and a number of those are non-technical or require you to develop and validate processes.

Attestation is a key component of the DORA articles. Even if you have the architecture and processes to recover from a cyber incident successfully, you can still face fines and other penalties if you can't prove it. You have to demonstrate to regulators, through periodic testing and reporting, that you can actually meet service level agreements. Commvault Cloud and Pure Storage both have reporting capabilities that help meet the requirements, but tying it all together on your own can be daunting. Our partners offer services that can connect the dots and make attestation much simpler so you can be confident in your compliance.

This is not a one-size-fits-all architecture. Every environment is different, and the various vendors, products, and policies you have in place may require a different configuration than we built in our lab. We can't anticipate compatibility challenges you may run into, and a miss could put the integrity of the entire SCRE at risk. That's why we designed the solution around capabilities, not specific components, so that it's customizable for you. As you work with Pure Storage and Commvault to understand your needs, we'll make sure you are connected with a systems integrator that can guide you to the optimal configuration for your environment and budget.

Beyond DORA

Although this solution addresses specific DORA regulations applicable in the European banking sector, operational resilience is applicable to organizations in all industries. Governments recognize this reality and have either published or are drafting regulations that apply across a broad range of industries. The joint Commvault and Pure Storage solution ensures digital operational resilience to align with these current and future regulations. By combining industry-leading data management with innovative data platform technology, we provide organizations with a comprehensive approach to protecting, recovering, and deriving maximum value from their data, even in the face of disruption.

The Pure Storage and Commvault Partnership

Pure Storage and Commvault have partnered since 2010 to solve real-world challenges in securing, managing, and recovering data of all types. The combination of Commvault's industry-leading cyber resilience software and the high-performance, secure Pure Storage platform enables organizations to protect their mission-critical data and applications and deliver uninterrupted services to their customers and employees in the face of increasingly complex and sophisticated cyber threats.

By choosing Commvault and Pure Storage, you can unlock the full potential of your data, drive operational excellence, address compliance obligations, and gain a competitive edge.

Additional Resources

- Learn about the Digital Operational Resilience Act (DORA).
- Read the joint Pure Storage and Commvault solution brief Compliance with Confidence.
- Learn more about Pure Storage and cyber resilience.
- Read more about Commvault solutions for cyber readiness and recovery.
- Check out our other Pure Storage and Commvault joint solutions and integrations.
- Review NIST Special Publication SP800-209 Security Guidelines for Storage Infrastructure.

For future product roadmap purposes, the development, release, and timing of any features or functionality described for Pure Storage products remains at the sole discretion of Pure Storage. The information provided is for informational purposes only and is not a commitment, promise, or legal obligation to deliver any material, code, or functionality and should not be relied upon in making purchasing decisions or incorporated into any contract.

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.

purestorage.com

800.379.PURE