PURESTORAGE®

TECHNICAL WHITE PAPER

Simplify Data Availability with ActiveCluster and Commvault IntelliSnap

Site-aware backup, recovery, and reuse from FlashArray™ snapshots.

Contents

Introduction

Your critical applications require high availability (HA) to ensure uptime. At the same time, IT budgets demand simple solutions with a low total cost of ownership (TCO). Purity ActiveCluster™ from Pure Storage® meets both these needs, with active-active synchronous replication and the proven simplicity of Pure FlashArray™. Just as important, you need to be able to protect your critical applications without service disruption and recover them as quickly as possible. Commvault IntelliSnap technology leverages Purity//FA snapshots to extend the benefits of ActiveCluster and lower the impact of backup on your critical systems while restoring data fast and enabling simple self-service data copies for non-production use. Together, IntelliSnap and ActiveCluster deliver application-consistent data protection, automated point-in-time recovery, and rapid self-service data copies in a single pane of glass that's simple to deploy and manage.

One of the great features of ActiveCluster is its ability to support active-active applications running in two sites at Metro distance (11ms round trip or less). In a server or site loss, the application will still run and be available without disruption to customers. However, it's still important to protect applications against corruption and data loss. When you create a Purity//FA snapshot of data in a stretched pod, ActiveCluster records the same recovery point on both arrays. Commvault IntelliSnap takes advantage of the automatic replication to enable site-local backup, recovery, and access operations. You can:

- Set policies to back up from snapshots to storage local to one of the FlashArrays.
- Use Commvault for database log backups to get one-step point-in-time recovery from snapshots.
- Restore individual VM files from snapshots, with Commvault automatically choosing the best site for the operation.
- Use Live VM Recovery to instantly power on a virtual machine from a snapshot copy while restoring the VM in whichever site you need.
- Use the Live Mount and Instant Clone features to enable IT admins or even end-users to increase productivity, automatically provisioning and de-provisioning temporary snapshots in one or both sites for non-production use such as development and analytics.

How Commvault IntelliSnap Interacts with ActiveCluster

Commvault IntelliSnap seamlessly integrates Purity//FA snapshots into the data protection lifecycle to minimize the impact of backup on production applications. At the beginning of each backup, the Commvault agent interacts with the source application, operating system, and/or hypervisor to dynamically identify the ActiveCluster storage arrays and stretched pod volumes behind the protected data. The Commvault agent temporarily freezes writes to the pod volumes using the appropriate application or OS APIs—leaving the application online—then creates a snapshot of the pod volume or volumes. ActiveCluster ensures that any snapshots are replicated between arrays. Once the snapshots are created, the Commvault

agent releases the frozen writes. This whole process typically happens in a few seconds or less and isn't noticeable to the application or users. Commvault adds information about the snapshots and the data they contain into its catalog.

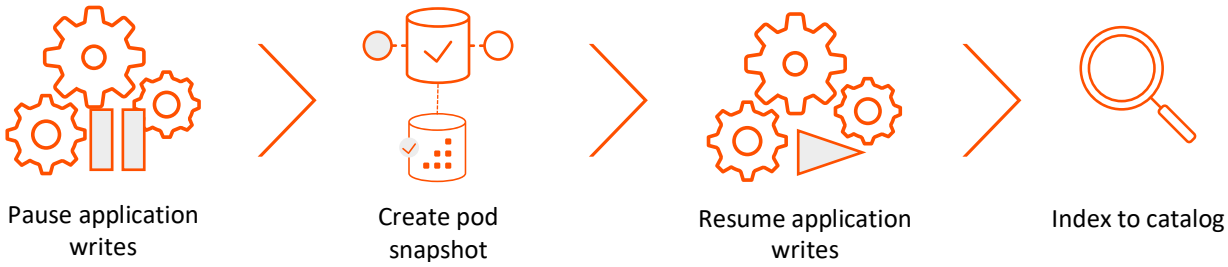

Figure 1. Commvault backup process.

Snapshot management is part of the same policy that manages backups. Commvault will periodically run a backup of the captured snapshot to Pure FlashBlade® or other storage, on a frequency you define, as though it came from the live production system. You can back up every snapshot, or just a subset, letting you use frequent snapshots for short recovery point objectives (RPO) and less frequent backups for longer-term recovery.

When Commvault needs to retrieve data from a snapshot, recovery, or data reuse, it identifies which pod member array is local to the system that will access the snapshot based on the configuration. It tells that array to copy the snapshot to a temporary volume—consuming no extra storage initially—outside the stretched pod (Figure 2). This minimizes the data and network footprint of the operation by keeping all data and traffic local within the site. When the operation is complete, Commvault removes the temporary volume and frees up any storage it consumed.

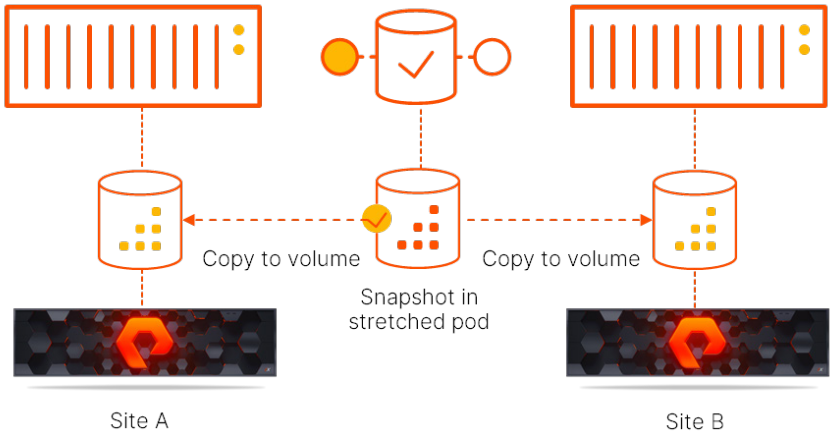

Figure 2. Copying a pod snapshot to a local volume

Solution Architecture

To demonstrate the solution's joint capabilities, we built the architecture illustrated in Figure 3, which follows the same principles as our Pure Validated Design solution. A pair of FlashArrays, logically separated into two sites, were connected in ActiveCluster synchronous replication mode. A pod was created and stretched between the arrays, and a volume was created within the pod to host a VMware vSphere datastore. A vSphere cluster with ESXi hosts in both sites was attached to the stretched datastore volume. VMs were provisioned in the stretched datastore, with half the VMs attached to hosts in each site.

A VM running the Commvault Virtual Server Agent was provisioned in each site. A Pure FlashBlade object bucket in each site was configured as a cloud storage pool, each with a deduplication database (DDB) and server plan. Connection details for the FlashArrays were added to Commvault to enable IntelliSnap. The VMs in each site were added to separate VM groups in Commvault and associated

With this configuration, we were able to create application-consistent snapshots on both FlashArrays for either VM group. We then backed up the VM group to the local FlashBlade bucket, using the snapshot as the source. We were able to run Live Mount and Live VM Recovery operations using the snapshot in either site, regardless of which site hosted the source VM. While we did not configure Commvault auxiliary copy for multi-site availability of backups on FlashBlade, that is easily achieved.

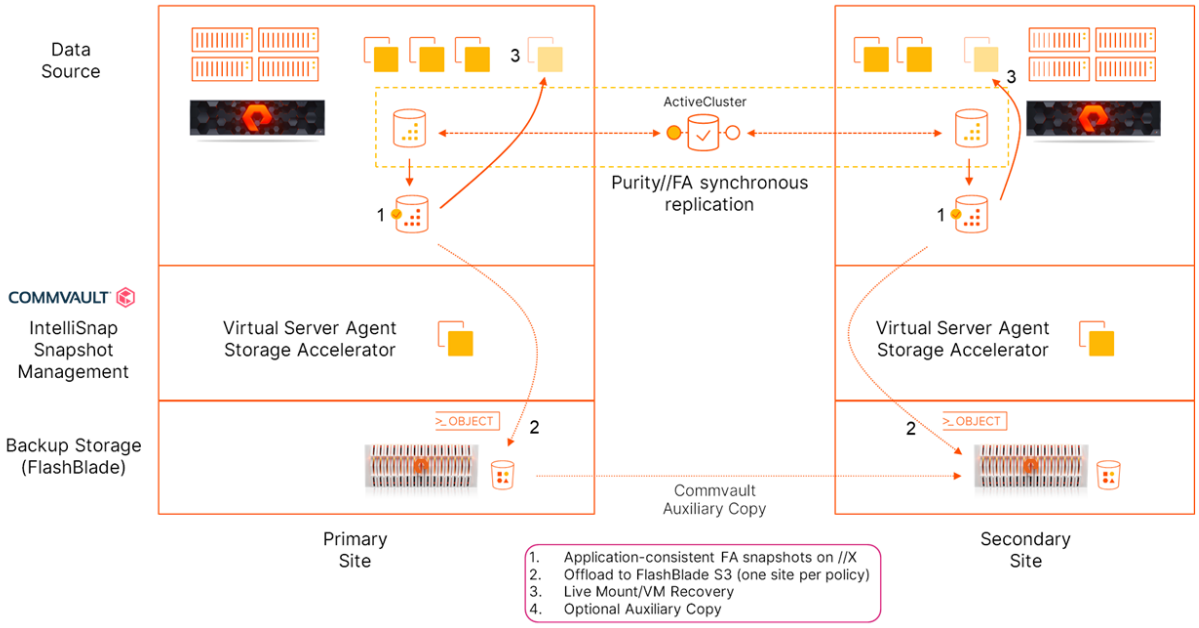

Figure 3. ActiveCluster IntelliSnap solution architecture

Lab Components

Server Details

Four ESXi hosts formed a simulated stretched cluster hosting the source VMs and Commvault components. Commvault Virtual Server Agent (VSA) and MediaAgent (MA) were deployed together on a VM in each simulated site. Each VM also had the MediaAgent installed to provide deduplication and storage management for the FlashBlade object storage target. One of the Commvault VMs also housed the CommServe services.

NOTE: You should always size Commvault components based on your environment.

Server Role	CPU	RAM	Networking	Storage	Operating System
ESXi Host (x4)	2x Intel Xeon E5-2697 v2 @ 2.70GHz, 24 cores total HyperThreading enabled	512GB	2x Mellanox MT27500 family network adapter @ 40Gbps	1x datastore from 2 FlashArray//M70 arrays in stretched pod 2x datastores from FlashArray//M70 not in pod	VMware ESXi 6.7.0
Commvault Virtual Server Agent & MediaAgent (x2)	12 vCPU	24GB	1x vmxnet3	1x 100GB VMDK (boot) 1x 400GB VMDK (index & DDB)	Windows Server 2019 Datacenter build 1809

Table 1. System details

Storage Details

A FlashBlade with 15 17TB blades provided the object storage for Commvault backup data and snapshot indexes. We created two separate object buckets to simulate FlashBlades sitting in separate sites, and we configured separate storage pools within Commvault for each bucket. Source VMs resided on FlashArray//M70 arrays in a stretched pod configuration.

Storage Role	Array Model	Purity Releases	Physical Storage	Networking
Backup Storage	FlashBlade 15x17	Purity//FB 3.1.4	162.46TB (usable)	4x40Gb Ethernet (data) 2x1Gb Ethernet (management)
Data Source (x2)	FlashArray//M70	Purity//FA 6.1.4	21TB (usable)	4x1Gb Ethernet (management) 4x10Gb Ethernet (iSCSI)

Table 2. Storage details

Source Data Details

We built 10 source VMs with the following configuration. All VMs shared the same stretched datastore on the FlashArrays. The VMs were distributed between the ESXi hosts, with five VMs per simulated site.

VM Role	CPU	RAM	Networking	Storage	Operating System
Data Source (x10)	2 vCPU	4GB	vmxnet3 virtual adapter	1x100GB VMDK thin provisioned	Windows 10

Table 3. Source data details

Setup

Configure Replication Requirements

ActiveCluster requires a specific configuration on replication networks before connecting arrays. Most important are replication addresses and Mediator access. To use IP replication, you must have four replication IP addresses configured on each array, not a single IP address applied to bonded replication ports. You must also have access from both arrays to a Mediator, which can be either accessed via Pure1® through the Internet or a local Mediator virtual appliance on your network. The full specifics are detailed in ActiveCluster Requirements and Best Practices on the Pure Support site.

Enable ActiveCluster Synchronous Replication

Once the ActiveCluster requirements are met, contact Pure Support to enable and configure the ActiveCluster feature on your array. Support will confirm that the array is configured correctly for synchronous replication and will assist with creating the connection between arrays. You cannot connect arrays for synchronous replication without Pure Support assistance.

Create and Stretch a Pod

ActiveCluster functions around the concept of pods, virtual storage containers that can be local to an array or stretched across arrays. A pod can contain multiple volumes and protection groups. To create a pod:

1. Open the management UI on either array, navigate to the **Storage** view, then select the **Pods** tab.
   a. To create a new pod, click the **Create Pods (+)** button.
2. Enter a name for the pod in the **Create Pod** dialog box (Figure 4).

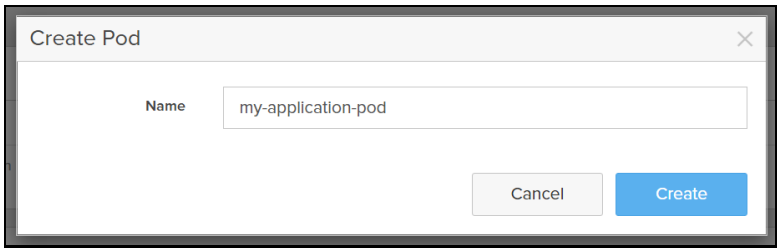

Figure 4. Create Pod dialog box

3. Once the new pod appears in the **Pods** tile, click its name to access its properties page.

You may create and/or move volumes into the pod.

> NOTE: You can only move volumes into a pod that has not yet been stretched. Once the pod has been stretched, you can only create volumes in it.

Creating a volume within a pod is similar to creating any FlashArray volume:

1. From the **Volumes** tile, click the **Create Volumes (+)** button.
2. In the **Create Volume** dialog box, enter the volume name and the size you want to provision, then click the **Create** button (Figure 5). You can also click the **Create Multiple** button to create more than one volume with the same size and naming convention.

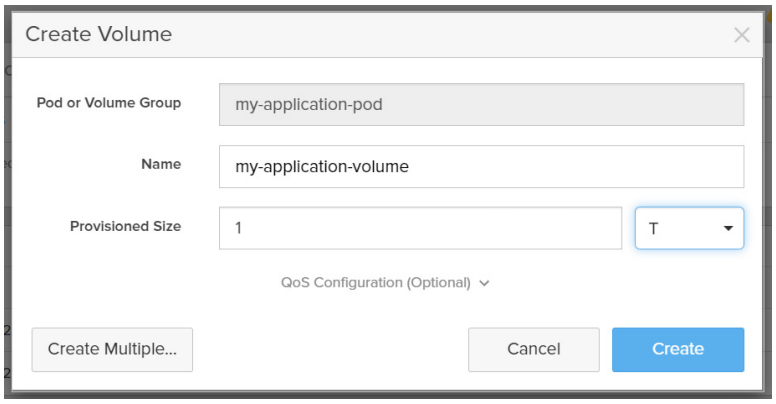

Figure 5. Create Volume dialog box

To move an existing volume into the pod, from the **Volumes** tile click the **menu** button, then select the **Move In** option. In the **Move Volumes In** dialog box, select the volume or volumes you want to belong to the pod, then click the **Move** button (Figure 6).

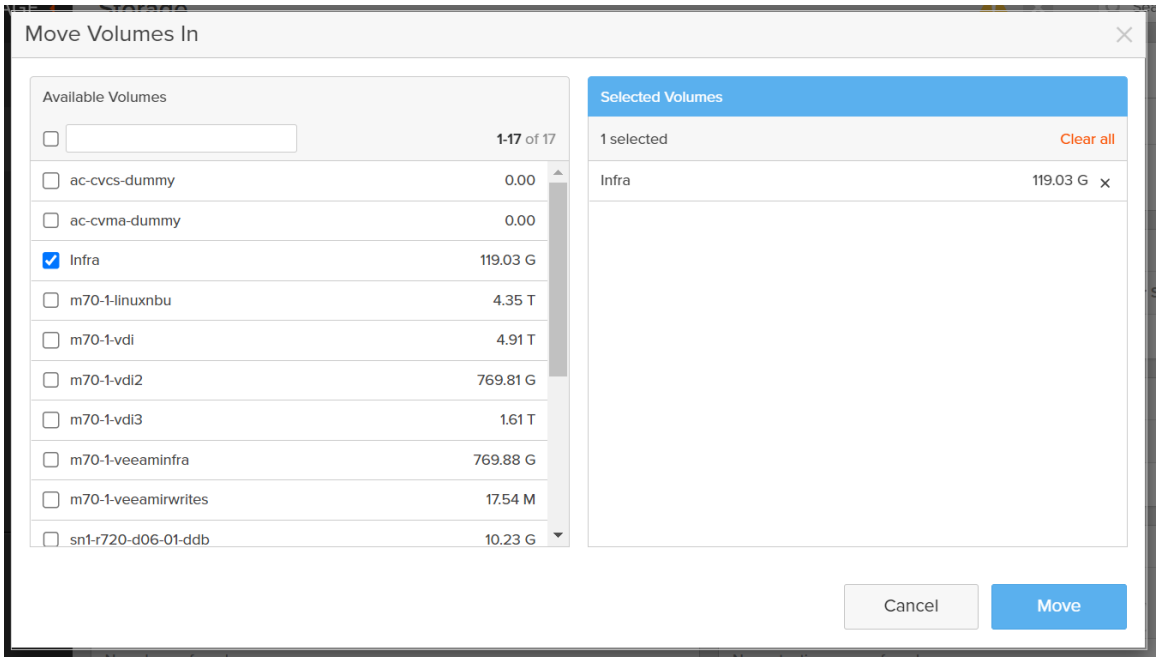

Figure 6. Selecting volumes to move into the pod

Note that the volume name will now include the pod name (Figure 7). This does not affect the host's view of the volume or existing IntelliSnap management, but it could affect scripts that expect only the volume name.

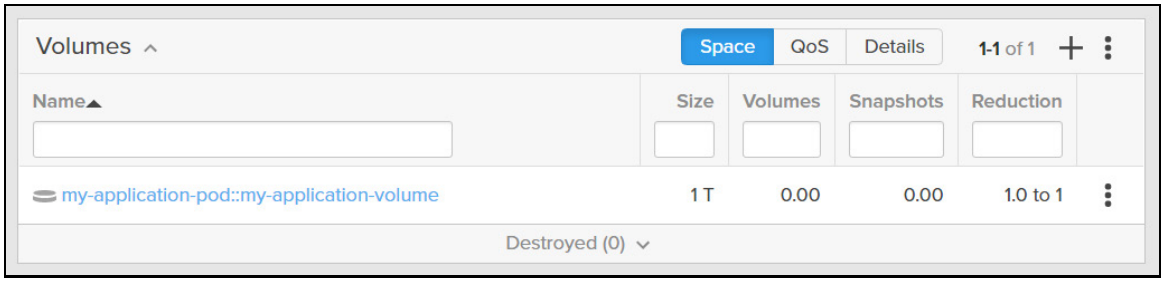

Figure 7. Pod volume naming

Once you have moved any volumes you need into the pod, stretch it to the second array. From the **Arrays** tile, click the **Add Array (+)** button. In the **Add Array** dialog box, select the partner array from the **Remote Array** dropdown, then click the **Add** button (Figure 8).

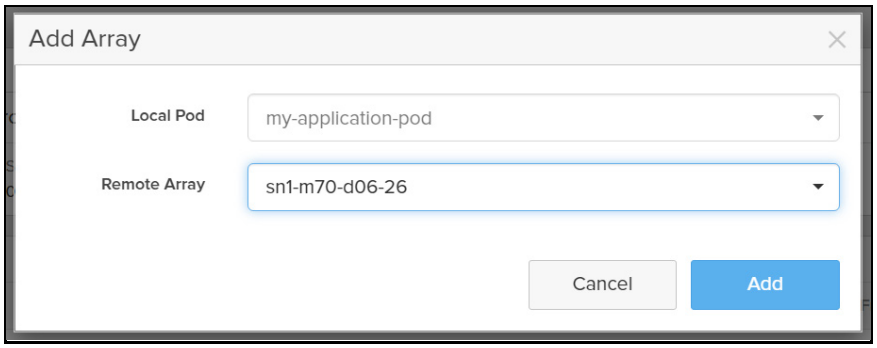

Figure 8. Add Array dialog box

The arrays will automatically connect and synchronize the configuration and data. You can monitor the status in the **Arrays** tile (Figure 9).

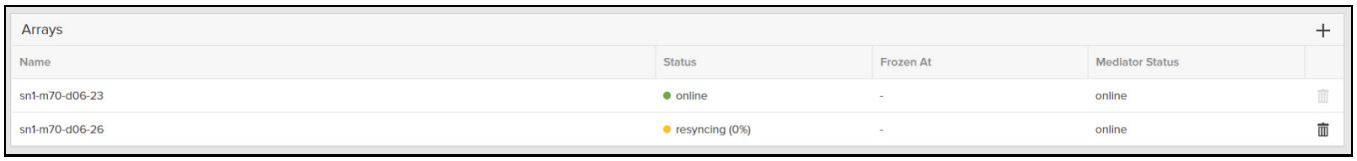

Figure 9. Array synchronization status

Once the sync is complete, both arrays will show "online" status (Figure 10).

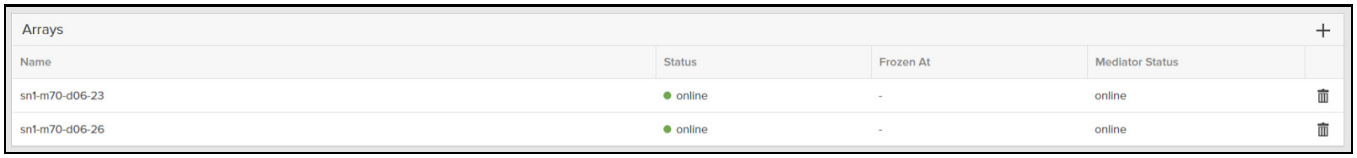

Figure 10. Synchronized arrays

Connect Volumes to Hosts

You may use a uniform or non-uniform configuration when connecting volumes to hosts. In a uniform configuration, each host has access to the volume from both arrays. In a non-uniform configuration, each host has access to only its local array. See ActiveCluster Requirements and Best Practices for more information about uniform and non-uniform configurations to determine what is right for your application.

Once you have determined the best configuration for your application, you must add any hosts and their host ports, as well as any host groups, to the appropriate arrays before you can connect the volumes to them. See the online help in the FlashArray management interface for instructions on setting up hosts and host groups.

You may connect hosts or host groups to the volume. Navigate to the properties page for the volume to update its connections. You can connect hosts from the **Connected Hosts** tile, and host groups from the **Connected Host Groups** tile. Click the menu button in the appropriate tile, then click the **Connect** option. In the **Connect Hosts** dialog box, select the hosts or host groups that need access to the volume (Figure 11). You may allow the arrays to assign a LUN ID automatically or specify one; some applications require matching LUN IDs for automatic failover, others do not. Click the **Connect** button to complete the configuration.

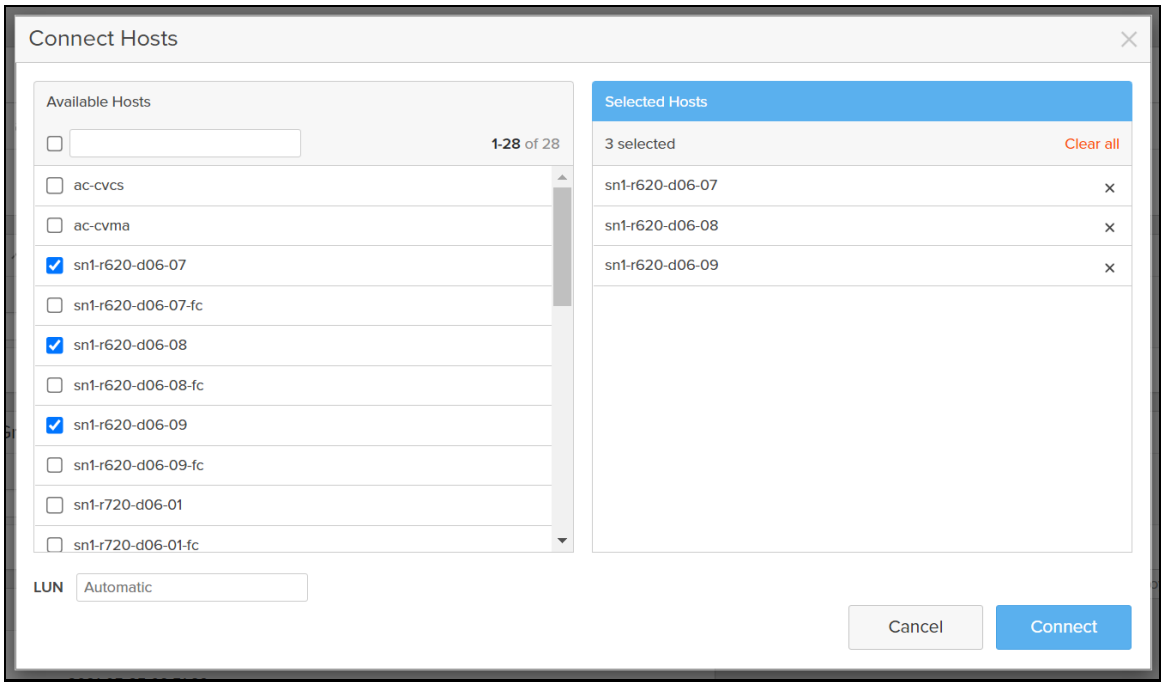

Figure 11. Selecting hosts to connect to a volume

The **Details** pane will show yellow status if the stretched volume is connected on only one array (Figure 12).

Figure 12. Volume status when connected on only one array

Repeat the connection process for any additional volumes in the same pod.

Change to the second array and repeat the procedure to connect the appropriate hosts or host groups, based on your choice of uniform or non-uniform configuration.

Add Arrays to Commvault

To use IntelliSnap, you must add information to Commvault on how to connect to the FlashArrays. You will need to obtain an API token from each array.

In the FlashArray console, click **Settings** in the left-hand navigation pane. Click **Access** on the top navigation bar. Click the menu icon for the user you want Commvault to use, then select **Show API Token** (Figure 13). You may first create a user if you wish to use a dedicated account for Commvault; the user must have the Storage Admin or Array Admin role.

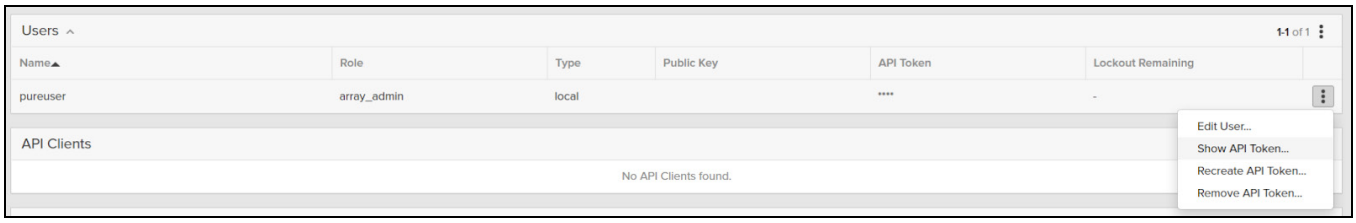

Figure 13. User menu

The **API Token** dialog box appears (Figure 14). Select the token text and copy it to the clipboard.

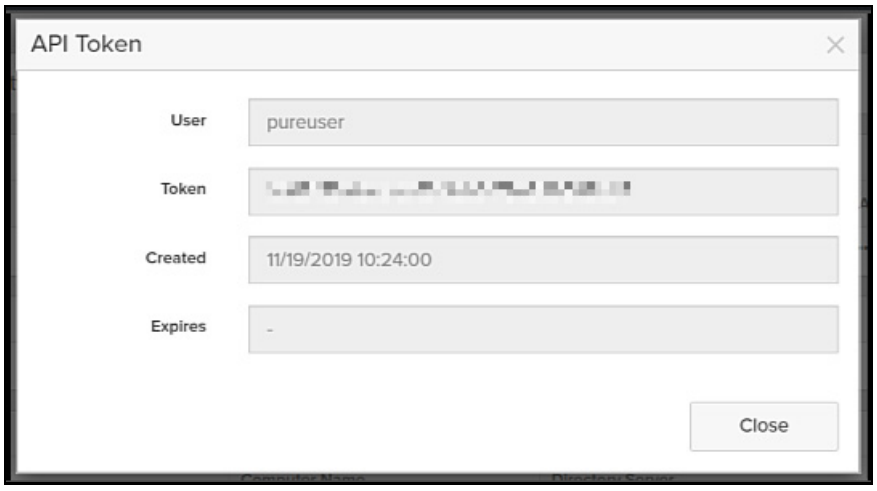

Figure 14. API Token dialog box

In Commvault Command Center, click **Manage** in the left-hand navigation pane, then click **Infrastructure**. Click the **Arrays** tile to open the list of managed arrays. Click the **Add** link to add a new array.

On the **General** form (Figure 15), select "Pure Storage FlashArray" from the **Snap vendor** dropdown. You can either enter the array credentials directly in the **User name** and **Password** fields, or you can select the **Use saved credentials** option and create a credential. Using saved credentials lets you centrally manage your Commvault passwords if you update them regularly.

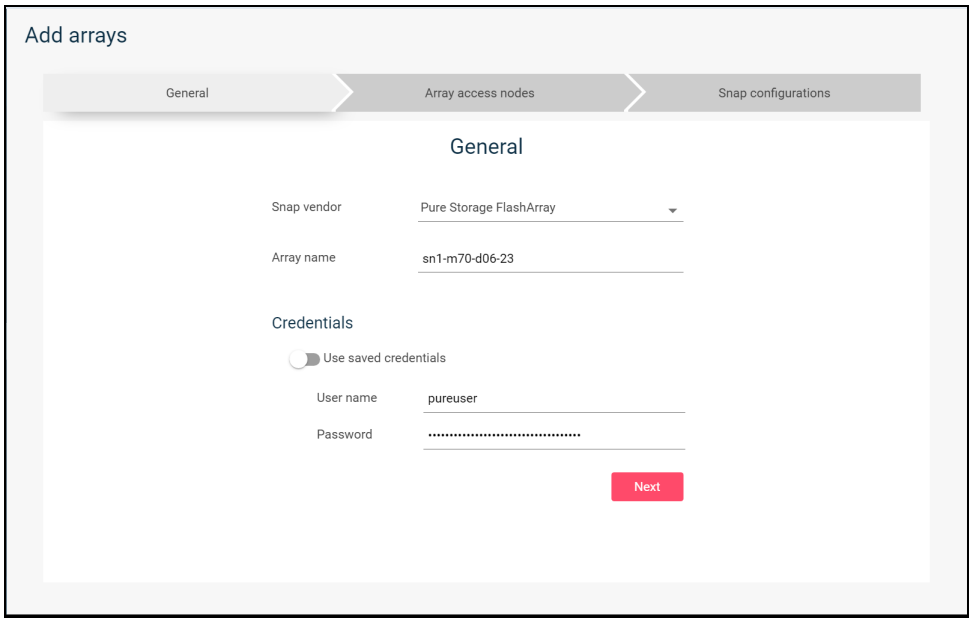

Figure 15. Adding FlashArray to Command Center - General form

On the **Array access nodes** form (Figure 16), from the **Available MediaAgents** dropdown, select any systems that will access data copies from this array. Selected systems will appear in the **Array access nodes** list. Avoid selecting systems that are in a different site from the array.

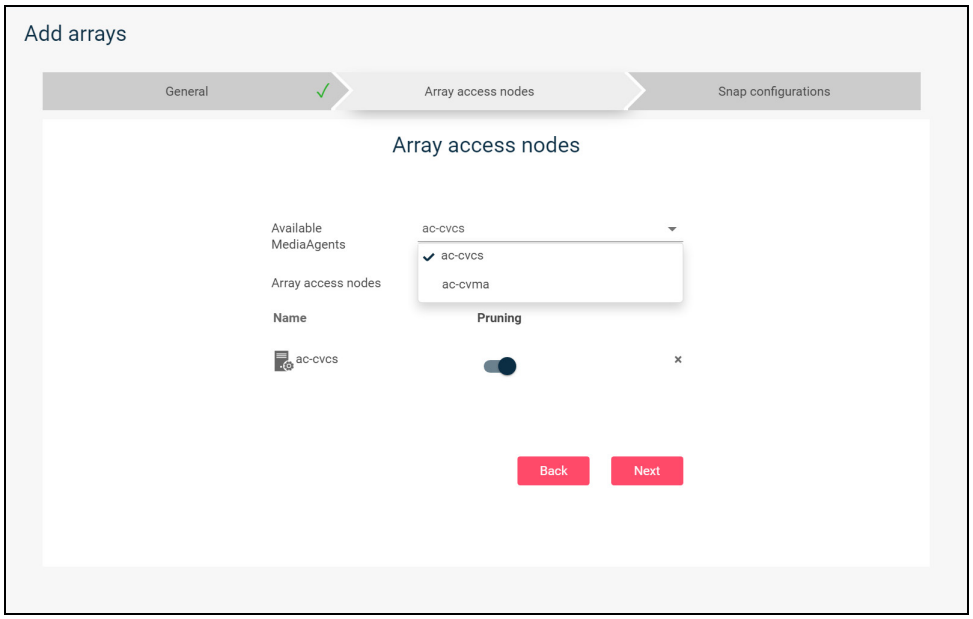

Figure 16. Adding FlashArray to Command Center - Array access nodes form

Settings on the **Snap configurations** step (Figure 17) are optional. However, we recommend setting the **Connect to a Host Group** and **Use Host if Host Group is not available** options when running VMware vSphere and clustered applications. We also recommend leaving the **Do not track Pod volume secondary snapshots** option disabled if you want to be able to perform operations such as Live Mount or database cloning in specific sites.

You can configure the **Remote Snap MA** setting if you wish to limit which systems communicate with the array. By default, each Commvault client creating or accessing a snapshot will send commands directly to the array. Setting a Remote Snap MA forces all array commands to go through the configured Commvault client instead.

> **IMPORTANT:** Make sure to enter the client name as it appears in Command Center, not the fully-qualified domain name (FQDN) or IP address. Entering the FQDN, IP address, or mismatched client name will cause failures.

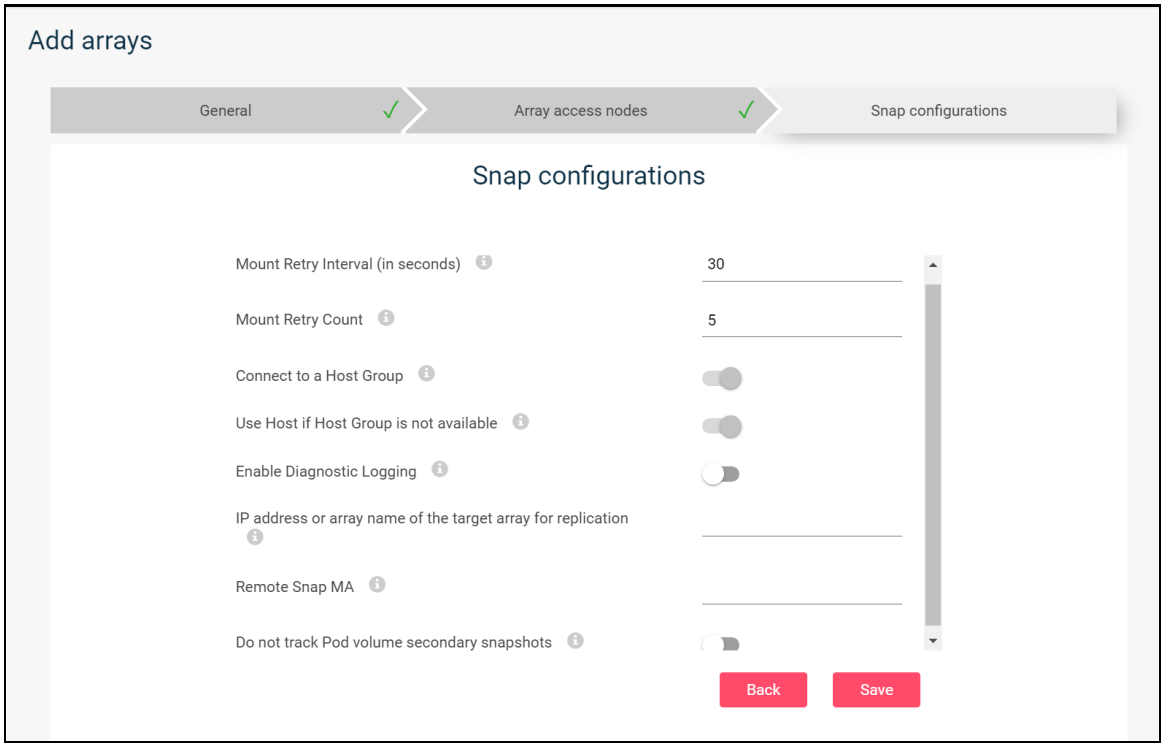

Figure 17. Adding FlashArray to Command Center - Snap configurations form

Repeat the setup for the second array.

Configure Snapshot Retention

Once the arrays are added, you need to make sure snapshot retention is configured in the server plan or storage policy. If you use server plans, you should already have snapshot retention enabled. You can look in Command Center at the plan's **Overview** page and confirm that a snapshot primary copy is listed in the **Backup destinations** pane (Figure 18). If one does not exist, you can add one by clicking the **Add** link and selecting **Snap copy**, then entering the necessary information.

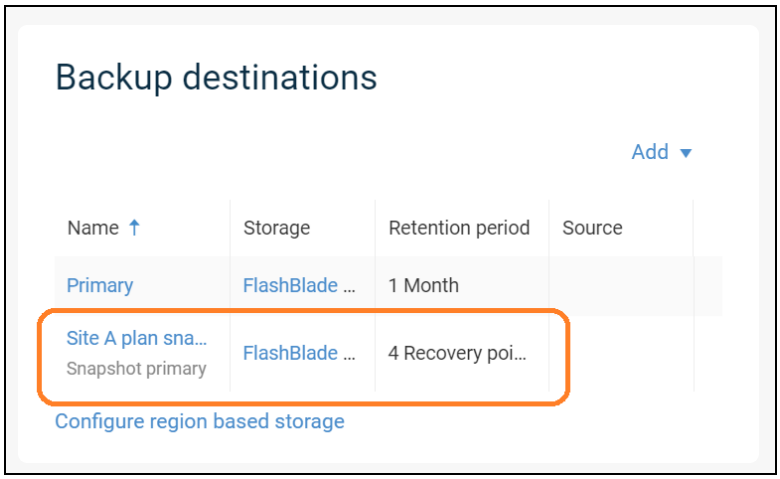

Figure 18. Server plan backup destinations pane

If you use storage policies instead of server plans, see Commvault documentation for instructions on how to add a snapshot copy if the storage policy does not already have one.

Enable IntelliSnap for Protected Data

You need to enable IntelliSnap for the data set to allow Commvault to orchestrate snapshots. This is done by setting the **Enable hardware snapshot** option on the data set (Figure 19). The placement of this option varies between data types. For VMware, it is on the **Configuration** page of the VM group; you may have to scroll down the page to find it. For other data types, it is usually in the subclient properties. Consult Commvault documentation for instructions on enabling IntelliSnap for specific data types. Click the switch icon to enable IntelliSnap and open the snapshot management options form.

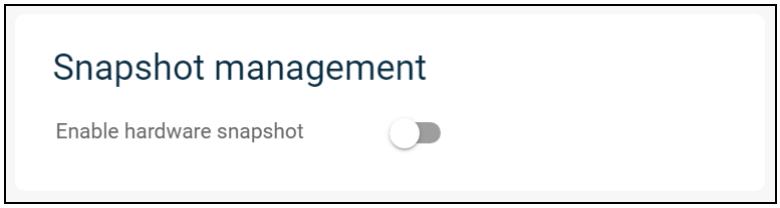

Figure 19. Enable hardware snapshot option

Figure 20 shows the available snapshot options for VMware VM groups. Select “Pure Storage FlashArray Snap” from the **Engines** dropdown. Under the **Snap mount esx host** section, expand the tree and select an ESXi host to perform snapshot mounts.

> NOTE: The available snapshot management options and appearance vary significantly between data types. See Commvault documentation for information on available options for specific data types.

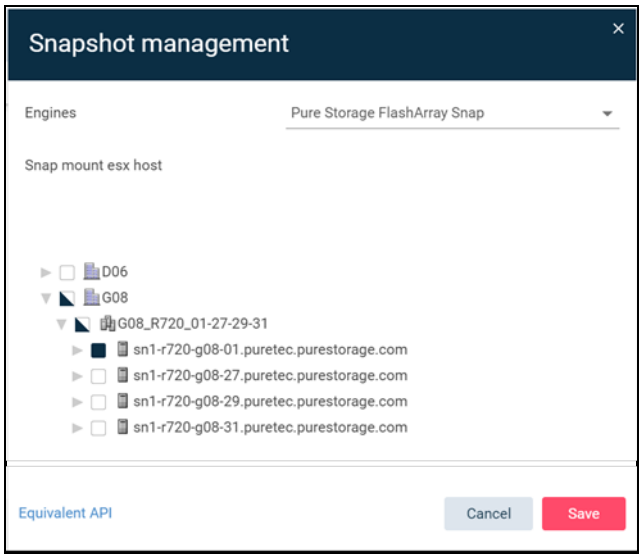

Figure 20. IntelliSnap options for VMware VM groups

Once IntelliSnap is enabled on the data set, future backups will start with a Purity//FA hardware snapshot. Backups to FlashBlade or other disk and cloud storage will use the snapshot as the source. Clients and access nodes that are associated to one array as an array controller will access snapshots through that array.

> NOTE: Support for using snapshots with different backup types varies by agent. See Commvault documentation for supported combinations of agent and snap backup types.

Live Recovery

During Live Recovery, Commvault will use the source snapshot to boot the VM being restored, then use Storage vMotion to migrate it to a production datastore while it's running. IntelliSnap will decide which ActiveCluster array to use based on the recovery settings.

To perform Live VM Recovery:

1. In Command Center, click **Protect** in the left-hand navigation pane, then click **Virtualization**. Locate the VM you want to recover. Click the **Actions (...)** button for the VM, then select the **Restore** option.
2. Select the Live recovery restore type (Figure 21).

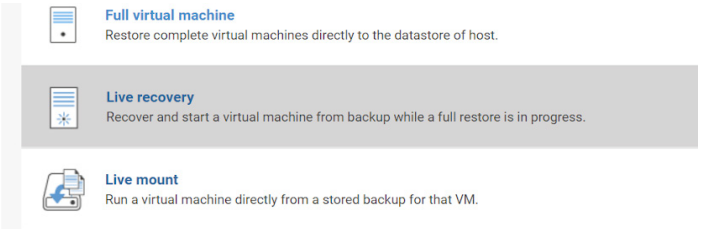

Figure 21. Live recovery restore type

3. The **Restore options** form appears. You can choose to restore in place, recovering to the original ESXi host and datastore, or out of place. For either option, you should always select an access node. Commvault chooses which array to use as a

recovery source based on the associations between the access nodes and arrays. If you do not select an access node, Commvault will choose one automatically, which may result in recovery failure or storage access between sites.

For in-place restore (Figure 22), select an access node in the same site as the original ESXi host. If the original VM still exists, you must enable the **Unconditionally overwrite if it already exists** option. You can select a datastore to redirect writes and delay the start of the Storage vMotion migration for a set number of hours. Click the **Submit** button to begin the recovery.

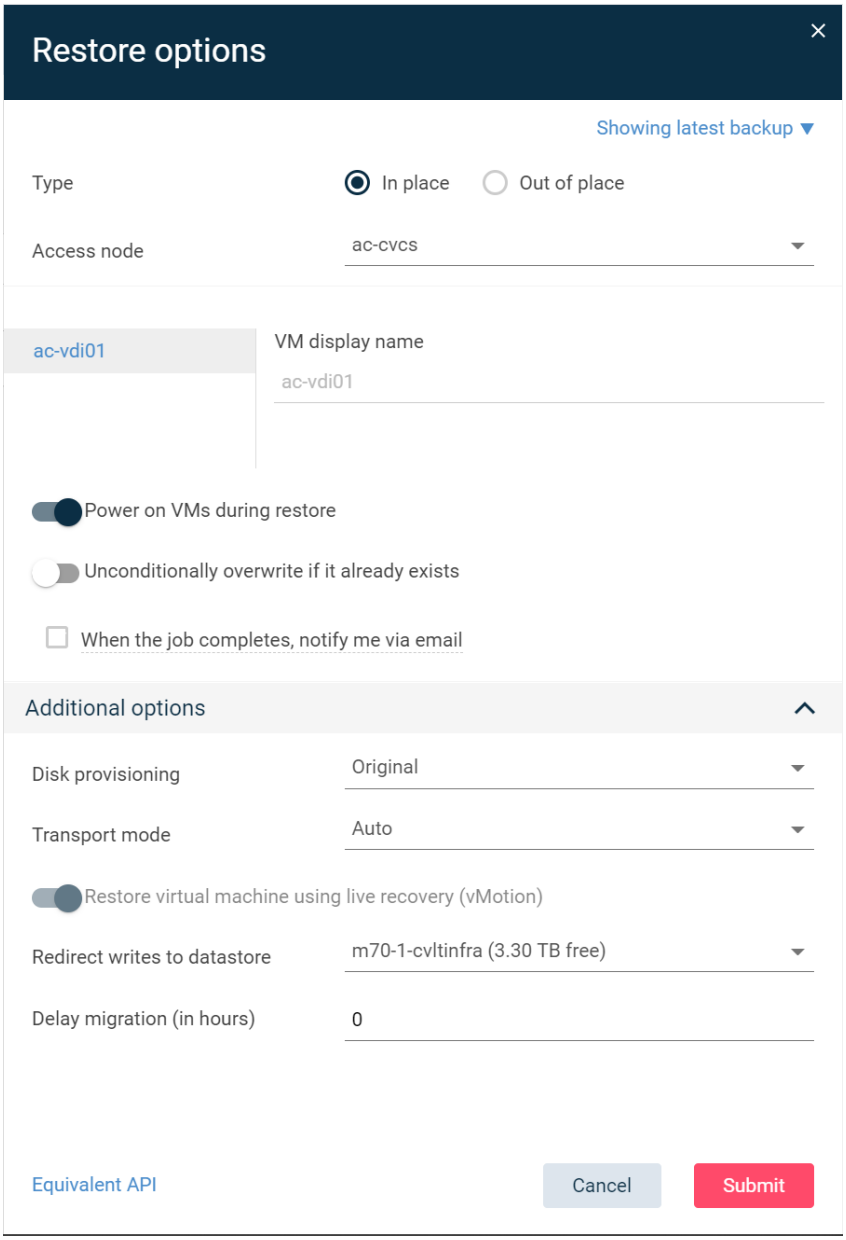

Figure 22. In-place recovery options

Out of place restore (Figure 23) has more available options, such as renaming the VM, setting a destination ESXi host and datastore, and remapping networking settings. You should select an access node in the same site as the ESXi host where you will recover the VM. You can change the VM display name if needed, and you can also change the destination ESXi host, datastore, resource pool, and VM folder.

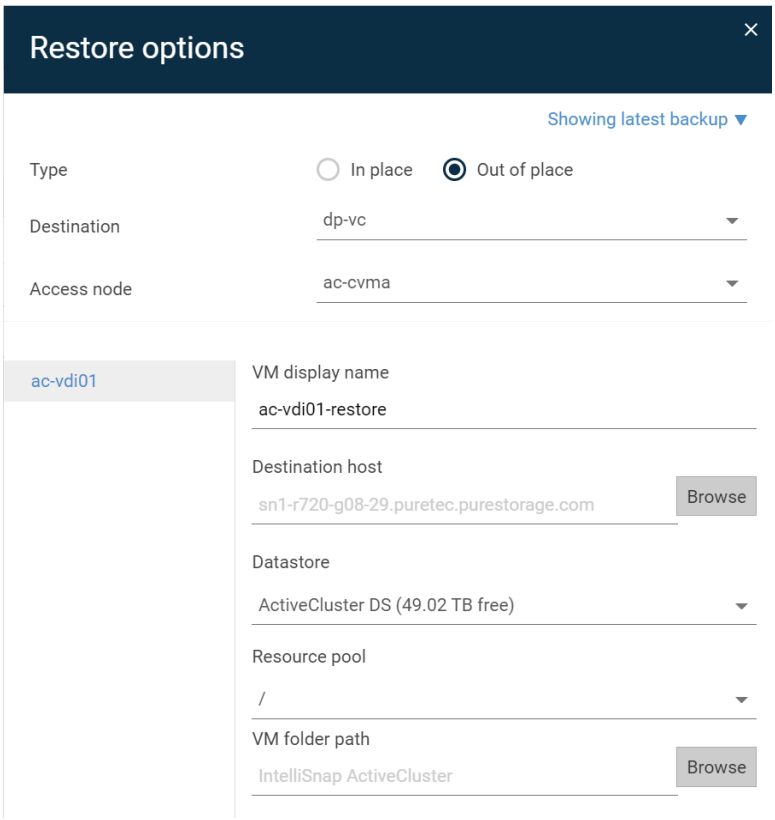

Figure 23. Out-of-place recovery basic options

If you need to change virtual networks or IP addresses, you can add those mappings by expanding the appropriate section and clicking the **Add** link (Figures 24 and 25).

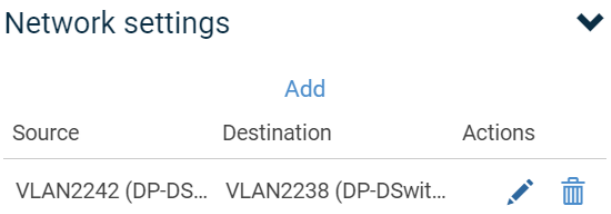

Figure 24. Remapping virtual networks

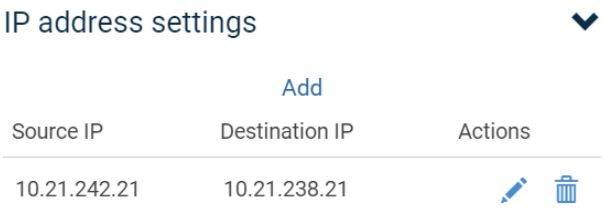

Figure 25. Remapping IP addresses

Figure 26 shows the remaining recovery options. To prevent VM naming collisions, you can enable the **Unconditionally overwrite if it already exists** option. Select the datastore that will receive redirected writes, but do not choose the same datastore where the VM will run after migration. You can choose to delay the migration for a set number of hours, which is

useful if you want to spread out the I/O from a large number of VMs or prevent migration during production hours. Click the **Submit** button to begin the recovery.

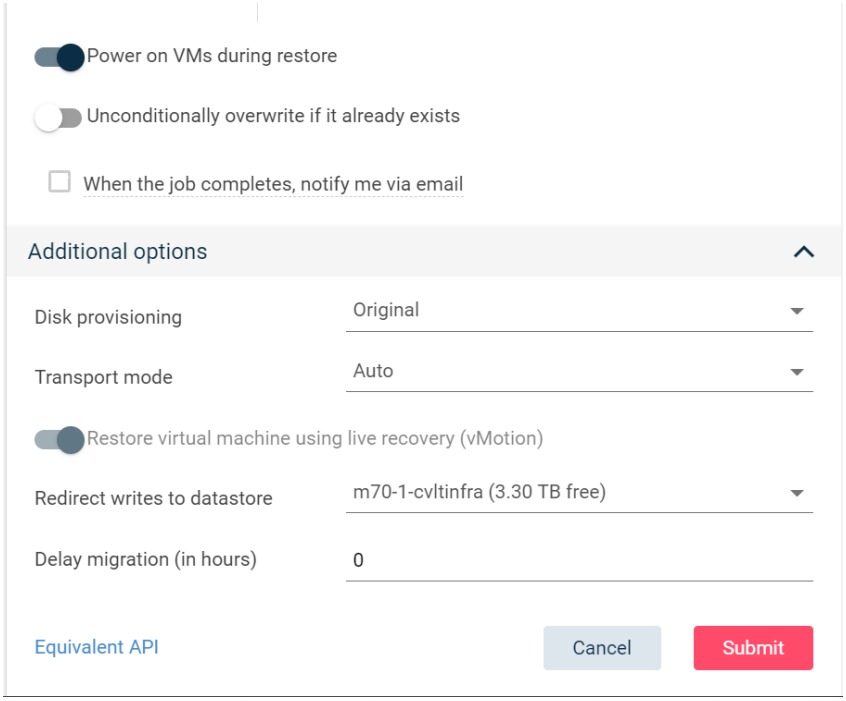

Figure 26. Final VM recovery options

Once the recovery job is submitted, Commvault will create a datastore from a temporary copy of the snapshot, on the FlashArray associated with the access node you selected. It will register the VM, apply any configuration changes you set, then bring the VM online. If you set a migration delay, Commvault will wait that many hours and then trigger a Storage vMotion migration to move the VM to the destination datastore. After the migration, Commvault will remove the temporary datastore.

Live Mount

Live Mount lets you streamline operations around DR testing, development, and other use cases. You can grant access for IT administrators and end-users to boot temporary VM copies that will run from FlashArray for performance and automatically clean up to avoid VM sprawl. You can configure policies that will run VM copies in a specific site, regardless of which site hosts the original VM.

> NOTE: To use Live Mount in Command Center, you must have a license for either Commvault Disaster Recovery or Commvault Complete Data Protection.

Create Recovery Targets

Before you can run Live Mount, you must create at least one recovery target. These define what options a user can choose when they mount a VM copy. Options include which vSphere resources they can use and how long they can reserve the VM copy before it is cleaned up. To avoid storage access across site links, you should create recovery targets that are tied to specific sites. We created one recovery target for each simulated site, but you can create as many as you need to meet your own Live Mount requirements. To create a recovery target:

1. In Command Center, navigate to **Disaster recovery**, then click **Recovery targets**. Click the **Add** link to open the **Add recovery target** form.
2. Complete the **Add recovery target** form as follows (Figure 27):
   a. Select "VMware vCenter" from the **Vendor** dropdown.
   b. Select "Regular" from the **Application type** dropdown.
   c. Enter a display name in the **Recovery target name** field.
   d. Select the appropriate vCenter instance from the **Destination hypervisor** dropdown.
   e. From the **Access node** dropdown, select the VSA that will access vCenter. This selection does not affect which FlashArray will be used for snapshot access.
   f. From the **Users and groups** dropdown, select the users and groups that should have access to the recovery target. You can set a default prefix or suffix for VM display names. The default can be overridden during the Live Mount process.
   g. For the **Destination host** field, click the **Browse** button and select the ESXi host that will host VMs during Live Mount.
   h. Select the datastore where you want writes to be redirected during Live Mount operations.
   i. If you want Live Mount VMs to use a specific resource pool, select it from the **Resource pool** dropdown.
   j. To place Live Mount VMs in a specific folder, click the **Browse** button for the **VM folder** field, then select the appropriate folder.
   k. From the **Destination network** dropdown, select the virtual networks you want users to choose from for Live Mount operations.
   l. Expand the **Live mount options** section of the form.

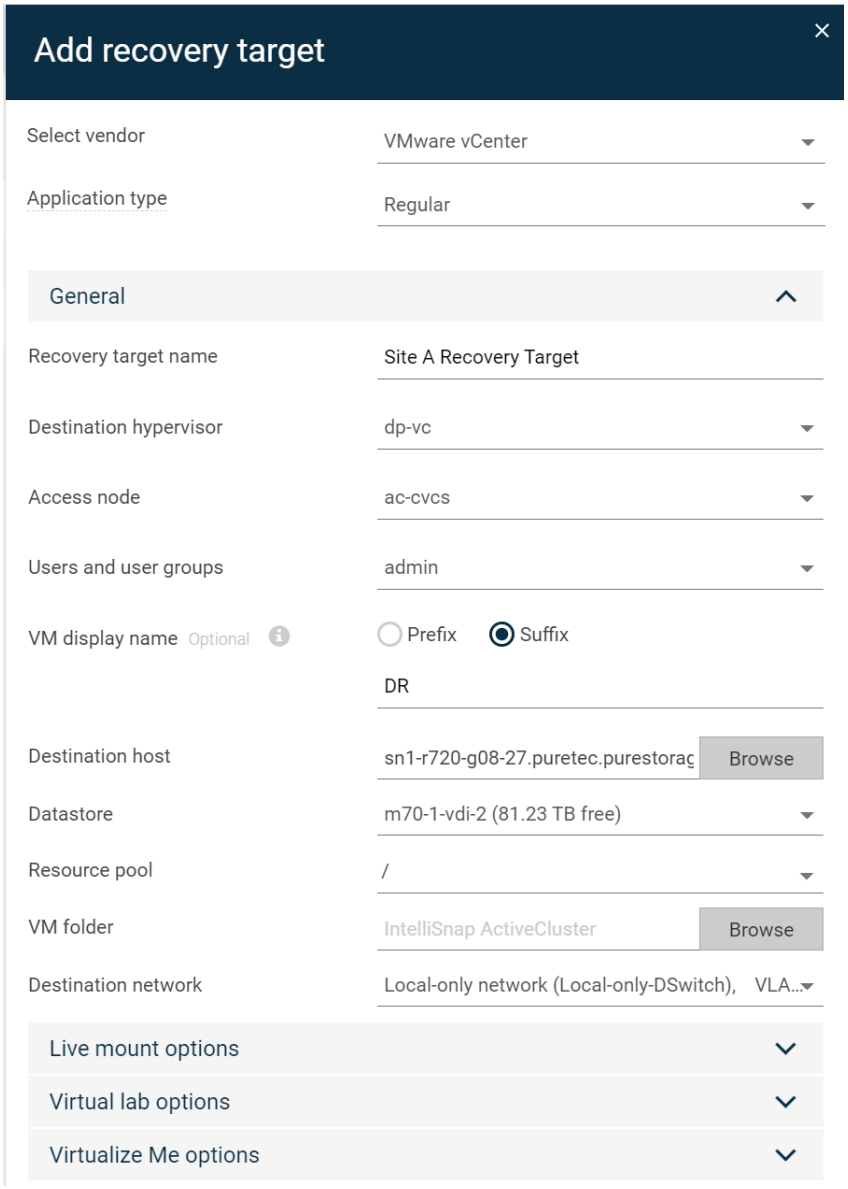

Figure 27. Add recovery target form

3. Configure the Live Mount options (Figure 27):
   a. In the **Expiration time** fields, set the desired number of hours or days for Live Mount reservations. Requesters cannot override this period, but they can extend reservations after Live Mount.
   b. In the **No of VMs** field, set the number of concurrent VMs users can run under this recovery target.
   c. In the **MediaAgent** dropdown, select an array access node associated with the FlashArray where you want mounts to occur. All Live Mount operations under this recovery target will be tied to that site.
   d. Click the **Add** button to create the recovery target.

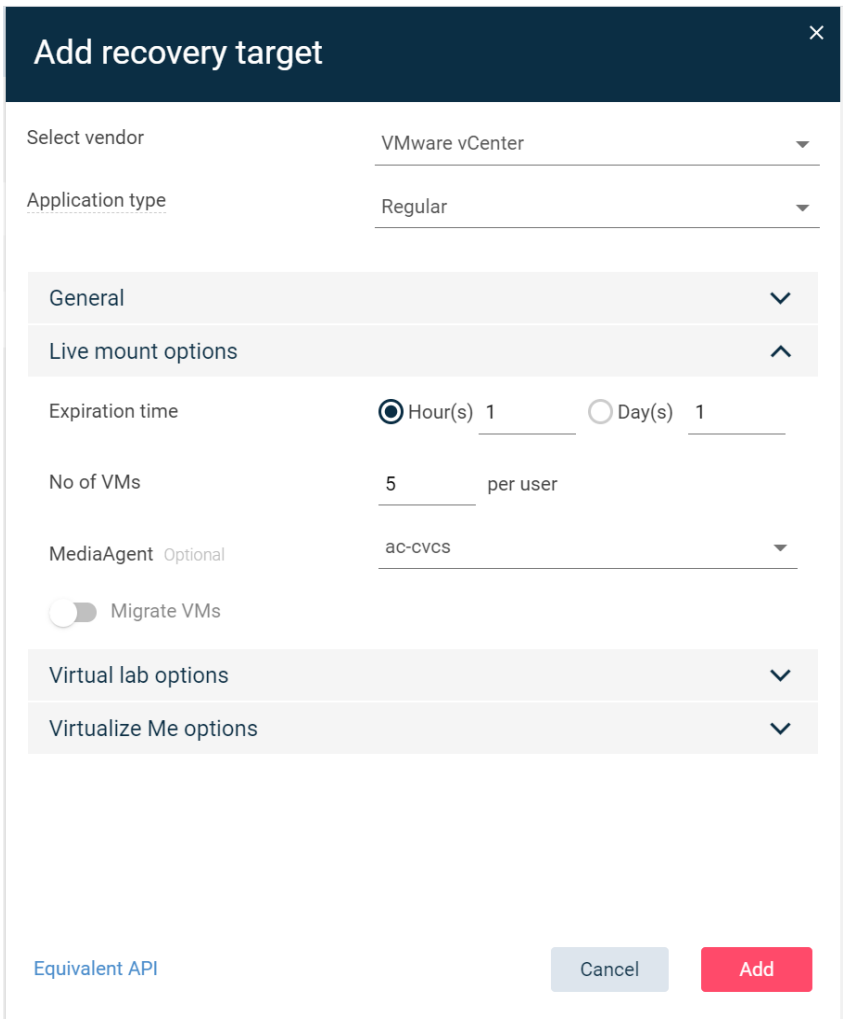

Figure 28. Recovery target options for Live Mount

Perform Live Mount

Live Mount is available as a restore type for single VMs. To perform a Live Mount:

1. In Command Center, click **Protect** in the left-hand navigation pane, then click **Virtualization**. Locate the VM you want to mount, then click its **Actions (...)** button. Select **Restore** from the actions menu.
2. Select **Live mount** as the restore type (Figure 29).

Figure 29. Selecting Live Mount

3. The **Live mount** form appears (Figure 30). Complete the form as follows:
   a. From the **Recovery target** dropdown, select the desired recovery target. The Live Mount will follow all the configured settings for hosts, datastores, access nodes, and by extension FlashArray site, as well as the initial reservation period.

b. In the **Virtual machine name** field, enter the desired name for the temporary VM. The name will default to the source VM name, with any prefix or suffix defined in the recovery target, but you may set it to any value that does not conflict with an existing VM.

c. From the **Network** dropdown, select the virtual network the VM should attach to. The recovery target definition limits the available options.

d. You do not need to change the **Copy precedence** dropdown. By default, it will use the ActiveCluster pod snapshot, with the specific FlashArray determined by the recovery target definition. If you wish to mount the VM from a backup instead of a snapshot, you can select the backup copy.

e. Click the **Submit** button to start the Live Mount process.

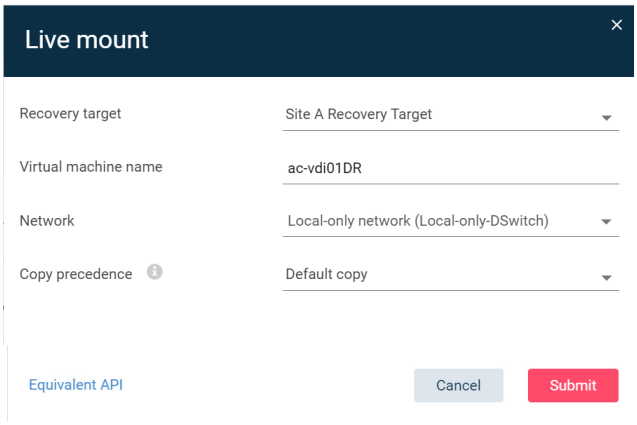

Figure 30. Live Mount options

Once the Live Mount is submitted, Commvault will create a datastore from a temporary copy of the snapshot, on the FlashArray associated with the array access node you defined in the recovery target. It will register the VM copy under the temporary name, apply any configuration changes you set, then bring the VM online. At the end of the reservation period, Commvault will power off the VM, remove the temporary datastore, and delete the temporary snapshot copy.

You can run Live Mount for multiple VMs concurrently, and you can run Live Mount multiple times for the same VM.

Monitor Active Mounts

You can see the active Live Mounts for a given VM or recovery target. From there you can expire or renew the reservations. From either the VM or recovery target, click the **Actions (...)** button, then select the **View active mounts** option. The **Active mounts** page will open, filtered to the VM or recovery target.

From the active mounts list, you can delete or extend the mount by clicking the **Actions (...)** button and select **Delete** or **Renew** from the menu (Figure 31).

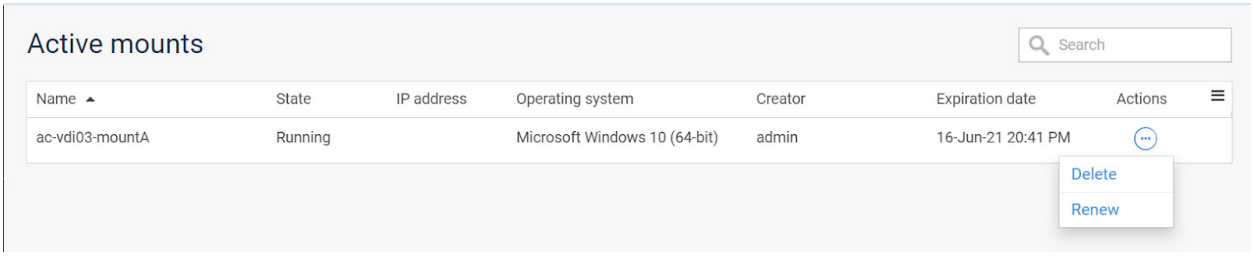

Figure 31. Active mount available actions

Conclusion

FlashArray ActiveCluster is a simple, reliable solution for adding high availability to applications across metro distances. Commvault IntelliSnap complements ActiveCluster with simple, reliable data protection that creates application-consistent FlashArray snapshots. You can use those snapshots automatically for backup that doesn't impact your high-value applications; Live Recovery for near-zero recovery time; and efficient, high-performance data copies. Integration between IntelliSnap and ActiveCluster brings site awareness, so you get these benefits while keeping data local. You can define policies that streamline backups, simplify DR testing, and let your users drive faster database refreshes, with less infrastructure and higher performance than with dedicated second-tier storage.

To learn more about how you can use ActiveCluster to improve availability without complexity, see Active-Active Clustering. Visit Commvault to learn more about the benefits of their data protection solutions. When you're ready to talk, reach out to your Pure account team.

Additional Resources

Next Steps

- Learn more about FlashArray ActiveCluster.
- See how to improve data protection with IntelliSnap snapshot management.
- Visit Pure's Commvault page for information on other joint Pure and Commvault solutions.

Supporting Information

- Pure Validated Design with Commvault
- ActiveCluster Requirements and Best Practices
- Whitepaper: Asynchronous Replication with Commvault IntelliSnap
- Whitepaper: Rapid Restore of VMware with Commvault and Pure Storage FlashBlade
- Commvault IntelliSnap documentation

About the Author

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

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