

TECHNICAL WHITE PAPER

Asynchronous Replication with Commvault IntelliSnap

Deployment and management techniques for Pure Storage® FlashArray™

Contents

Introduction3

Solution Components3

Solution Architecture.....4

Basic Configuration4

 Configure Asynchronous Array Connection 4

IntelliSnap Setup6

 Adding FlashArray to Commvault..... 6

 Adding FlashBlade to Commvault..... 9

 Creating a Server Backup Plan 10

Checking Backup Workflow 14

Restore Options 15

Conclusion 16

Additional Resources 16

About the Author 17



Introduction

Pure Storage and Commvault joint solutions deliver the data protection and recovery performance needed to protect your modern data center. Organizations need to recover data, applications, or systems with minimal or zero downtime. Both Gartner Magic Quadrant leaders in their respective fields, together Pure Storage and Commvault empower organizations to store and protect their most valuable asset—their data—with speed, simplicity, and scale. This document joins the [other published papers](#) on how those capabilities work and how to deploy them.

Purity//FA, the operating system powering the FlashArray product family, provides powerful yet simple-to-manage block-level protection services. These services range from in-system replication with snapshots and clones suitable for recovery to cross-systems replica with synchronous and asynchronous replication services, ideal for disaster recovery and offloaded backup workloads.

Commvault IntelliSnap is an orchestration engine that combines application-aware data protection and recovery integration with the data protection services of Purity//FA. Commvault Backup & Recovery Software provides a single pane of glass and policy-based administration to manage infrastructure efficiently and alleviate the system administrator’s burden.

This guide will show how to deploy and manage asynchronous replication of Pure Storage FlashArray using only Commvault IntelliSnap as your single pane of glass to address your backup and recovery needs. This technical white paper is for IT, storage, and system administrators who want to leverage FlashArray snapshots and asynchronous replication to simplify their data protection and reduce backup and recovery times.

Solution Components

To demonstrate this solution. Our lab environment consisted of the following:

Component	Details
VMWare ESXi	Version 6.7
Purity FlashArray (//X, and //C)	Purity//FA 6.0.3
Purity FlashBlade®	Purity//FB 3.1.3
Connection Protocol (FlashArray, to FlashArray, and FlashArray to ESXi Hosts)	iSCSI

Table 1. Solution components



Solution Architecture

Figure 1 demonstrates a logical diagram of our VMWare lab environment. A couple of FlashArray//X Arrays serve datastores on the primary site. Pure FlashBlade is the storage target for snapshot metadata and backup copies. While you can use any storage target supported by Commvault, FlashBlade provides a simple, scalable platform for rapid recovery. Please see our [Rapid Restore White Paper](#) for more details on how to achieve rapid restore with FlashBlade.

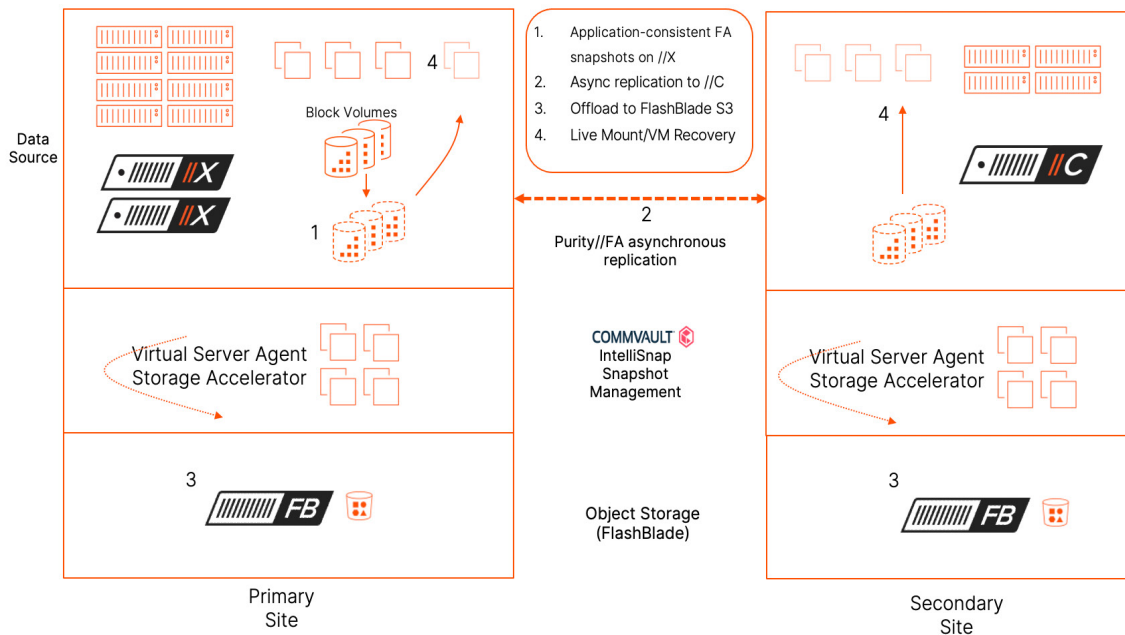


Figure 1. Solution diagram

The FlashArray//X arrays are configured for fan-in asynchronous replication to a single FlashArray//C Array in the secondary site. FlashBlade is the storage library for metadata and backups on the secondary site.

Note: For this paper, there is only one physical FlashBlade. Logically, we have created two S3 buckets, designated as local-storage and remote-storage. In a production deployment, each site would have a separate FlashBlade to keep backups local within the site.

During backup, Commvault will quiesce the VMs being protected. Commvault IntelliSnap will create application-consistent snapshots of the datastores on the FlashArray//X arrays and replicate them to the FlashArray//C in the secondary site. It will also manage the independent retention of each snapshot and the backup copies.

Basic Configuration

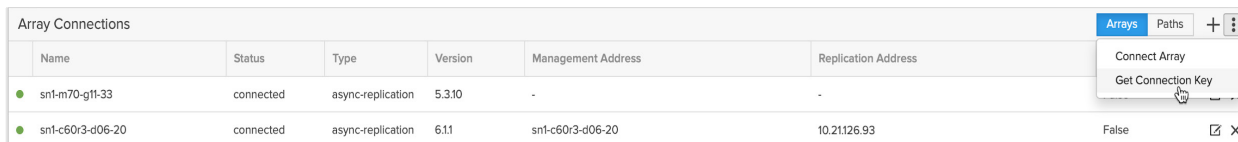
This section will cover the basic configuration to get the environment running, following documented best practices.

Configure Asynchronous Array Connection



Connecting two arrays involves logging in to one of the FlashArray Purity//FA Web interfaces, obtaining a connection key, then using that key to authenticate from the other array. Array connections do not define the direction of replication so that you can create them from either array. For demonstration purposes, we will create the connections from the FlashArray//X arrays.

1. From a web browser, log in to the Purity//FA Web Interface of the FlashArray//C.
2. From the Navigation Pane, select **Storage > Array**.
3. In the Array Connections panel, click **Get Connection Key** as shown in Figures 2 and 3, then copy the connection key to the clipboard.



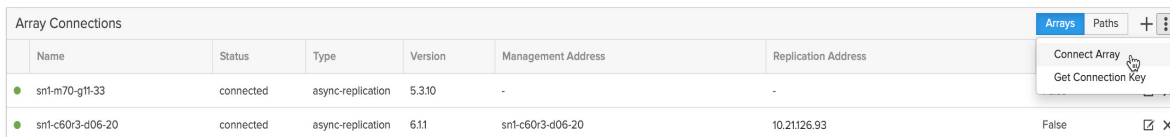
Name	Status	Type	Version	Management Address	Replication Address	
sn1-m70-g11-33	connected	async-replication	5.3.10	-	-	
sn1-c60r3-d06-20	connected	async-replication	6.1.1	sn1-c60r3-d06-20	10.21.126.93	False

Figure 2. Getting the connection key



Figure 3. Connection key pop-up window

4. On the first FlashArray//X, login to Purity//FA Web Interface.
5. From the Navigation Pane, select **Storage > Array**
6. Click the menu item on the Array Connections Panel and select **Connect Array**, as shown in Figure 4.



Name	Status	Type	Version	Management Address	Replication Address	
sn1-m70-g11-33	connected	async-replication	5.3.10	-	-	
sn1-c60r3-d06-20	connected	async-replication	6.1.1	sn1-c60r3-d06-20	10.21.126.93	False

Figure 4. Connect array menu

7. The Connect Array pop-up window will appear; set the following connection details, also shown in Figure 5:
 - a. In the Management Address field, enter the management virtual IP address or FQDN of the FlashArray//C.
 - b. In the Type field, leave it at the default Async Replication.
 - c. In the Connection Key field, paste the connection key string you copied from FlashArray//C in step 3.
 - d. In the Replication Address field, leave blank for auto-discovery unless NAT is in use. (If you use network address translation (NAT) between FlashArrays, you must manually enter the replication address presented by NAT.)



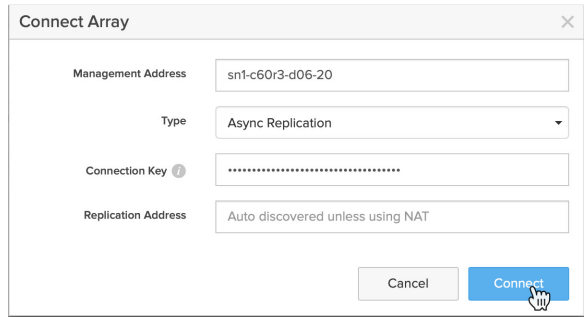
A screenshot of a 'Connect Array' dialog box. It contains four input fields: 'Management Address' with the value 'sn1-c60r3-d06-20', 'Type' with a dropdown menu set to 'Async Replication', 'Connection Key' with a masked value '*****', and 'Replication Address' with the text 'Auto discovered unless using NAT'. At the bottom are 'Cancel' and 'Connect' buttons. A mouse cursor is pointing at the 'Connect' button.

Figure 5. Connect array pop-up window

- 8. Click **Connect**. The Array appears in the list of connected arrays, and a green checkmark appears in the row, indicating that the two arrays are successfully connected.
- 9. Repeat steps 4-8 to establish a connection between the second FlashArray//X and the FlashArray//C.

When complete, the Array Connection Panel on the FlashArray//C will have two connections as shown in Figure 6, one connection per FlashArray//X and FlashArray//C relationship.







Array Connections									+ ⋮
Name	Status	Type	Version	Management Address	Replication Transport	Replication Address	Throttled		
sn1-m70-d06-26	connected	async-replication	6.0.3	-	Ethernet (IP)	-	False		 
sn1-m70-d06-29	connected	async-replication	6.0.3	-	Ethernet (IP)	-	False		 

Figure 6. Array connections panel with successfully connected arrays

IntelliSnap Setup

To use IntelliSnap to manage FlashArray snapshots, Commvault must be made aware of the underlying arrays. Note that you do not need to specify any host or volume details in Commvault since IntelliSnap will dynamically detect the storage tied to protected data.

Adding FlashArray to Commvault

To add FlashArray under Commvault management, for each Array, we need to have access to:

- 1. The virtual IP address or FQDN of FlashArray management interface
- 2. User ID, with at least a storage admin role
- 3. API Token created for the User ID, which is acquired from Purity//FA Web Interface, by going to **Settings > Access**, on Users panel, click **Show API Token** (Figures 7 and 8). Copy the token to the clipboard.



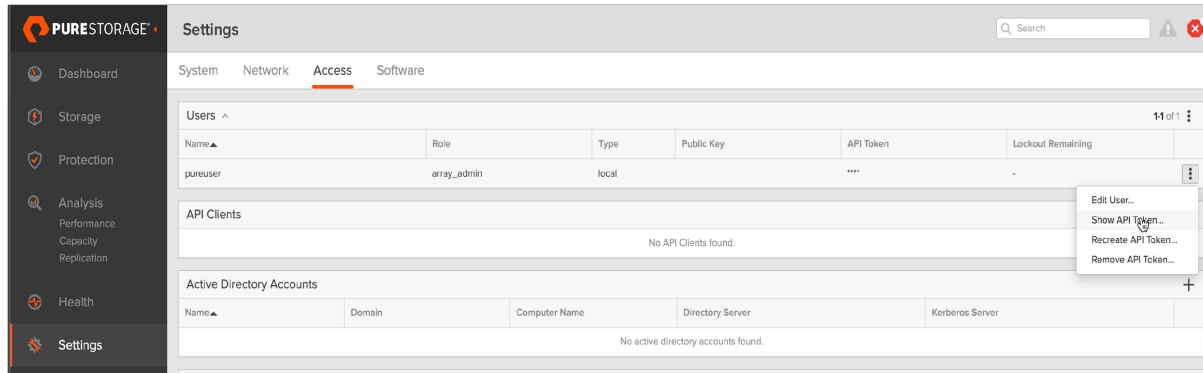


Figure 7. Acquiring API Token for the user

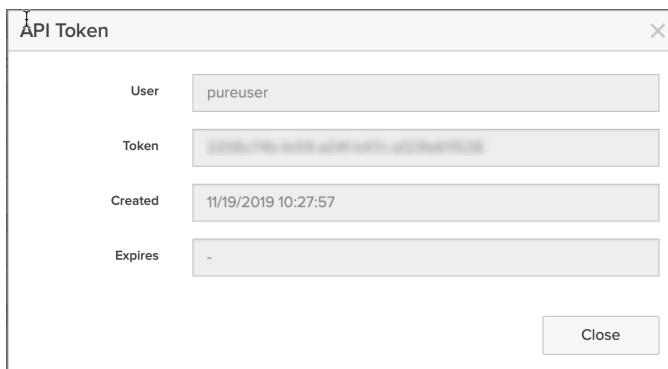


Figure 8. API Token pop-up window

Using Commvault Command Center, add each of the FlashArrays under Commvault management. From the navigation pane:

1. Click on **Manage > Infrastructure**. The Infrastructure page appears.
2. Click the **Arrays** tile.
3. In the upper-right area of the page, click **Add**; the Add arrays form page appears.
4. On the General tab, add the FlashArray information:
 - a. Snap Vendor: Pure Storage FlashArray
 - b. Array Name: Virtual IP address or FQDN of FlashArray//X management interface
 - c. Username: User ID with Admin rights
 - d. API Token: Paste the API Token copied from step c.
 - e. Click **Next**.



Figure 9. Add Array, General Tab

5. Optionally, for site awareness, it's recommended to select a MediaAgent to perform operations on the Storage Array, like pruning expired snapshots. Click on the Array Access Nodes tab, and choose which MediaAgent will control these operations for the selected Storage Array.
 - a. Under the Available MediaAgents box, select an array access node. Typically, it should be a MediaAgent in the same data center as the Storage Array.
 - b. Click **Next**.
6. On the Snap Configuration tab, fill in the form with the following:
 - a. IP address or array name of the target array for replication: In our example, this is the FlashArray//C virtual IP address or FQDN.
 - b. Remote Snap MA: (Optional) Choose a MediaAgent in the same Datacenter as the Target FlashArray for replication.

Note: If you do not choose a remote snap MA, each source client will communicate directly with the FlashArray.

7. Click **Save**.

Figure 10. Add array, Snap configuration tab

8. Repeat steps 3-6 for the second FlashArray//X.

For FlashArray//C, repeat steps 3 through 5 to enable Commvault to manage operations like the pruning of expired snapshots.



Adding FlashBlade to Commvault

FlashBlade supports NAS and Amazon S3 Protocols. However, utilizing an object bucket such as Storage Target dramatically simplifies management of the environment and provides non-disruptive scaling. In our environment, we created two object buckets. The first one, "local bucket," stores metadata and backup copies of the snapshots for both FlashArray//X primary site arrays. The second object bucket will store metadata and backup copies of snapshots from the FlashArray//C Array in the secondary site. Each bucket is associated with its own storage pool for deduplication.

To add the storage target for the first bucket to the Commvault environment:

1. Login to Commvault Command Center. Navigate to Storage > Cloud; in the upper-right area of the page, click **Add**.
2. On the Cloud Storage Type Page, click **Cloud Storage Tile**.
3. On Add Cloud Storage Form, fill out the following fields:
 - a. Configure the Cloud Name field, enter a display name for the cloud storage target.
 - b. Type field: Select **S3 Compatible Storage**.
 - c. MediaAgent field: Select the MediaAgent that should manage the pruning and be the preferred data path.
 - d. Server host field: Enter the FlashBlade FDQN or VIP for Commvault to access.



Note: By default, Commvault will use TLS to secure S3 connections. If the FlashBlade does not have a trusted certificate that matches the provided name, TLS will fail. Please refer to related [Pure Storage FlashBlade Best Practice Guide](#), or Commvault's related documentation.

- e. Credentials field: Select or create a saved credential for the FlashBlade bucket user. Using saved credentials simplifies the configuration of additional buckets and rotation of access keys.
- f. Bucket field: enter the name of the target bucket under the configured FlashBlade object user account.

Note: For details on adding users and buckets to FlashBlade, please see [FlashBlade User's Guide](#).

- g. Use deduplication toggle, default enabled. Clicking the Add link will bring a pop-up window to choose the MediaAgent, and the ability to either browse to a path to create the deduplication database for the storage pool or type the full directory path if you already know it. We recommend always using Commvault deduplication with FlashBlade object storage.

Repeat the above procedure to add the second FlashBlade S3 Bucket.

Creating a Server Backup Plan

The following steps describe how to create a backup server plan, where Commvault orchestrates snapshots and replication on the FlashArray.

1. Login to Commvault Command Center.
2. From the navigation pane, go to **Manage > Plans**.
3. In the upper-right area of the page, click **Create Plan > Server Plan**.
4. On the Create Server Plan Form, provide a descriptive name for the plan name.
5. Under Backup Destinations, click on **Add**
 - a. The Name field, which defaults to "Primary," can be changed to something more meaningful. (i.e., data center, location, target, etc.)
 - b. In the **Storage** field, the pull-down menu will show both previously-configured FlashBlade S3 storage libraries. Select the local storage library.
 - c. Under Retention Rules, the **Retention period** is how long to keep backups on the specified storage library. This retention only applies to the Backups stored on this Storage library and does not control how many snapshots will be kept on the



FlashArray. Specify an appropriate period and click **Save**.

Figure 11. Add backup destination

6. Under RPO (Recovery Point Objective), Backup frequency, specify how often the incremental backup jobs will run.
7. Optional: If you would like to create periodic full backups rather than just incremental and synthetic full backups, you can enable full backup frequency with Add full backup, toggle button, then specify the frequency.
8. Optional: Under Folders to backup, you can more granularly specify which folders to backup. The default is all folders.
9. Under Snapshot options, select Number of snap recovery points field, specify how many snaps to keep on production FlashArray//X Arrays.
10. Optional: Under Database Options, you can specify shorter RPO or leave the default 4 hours.
11. Optional: Under Override restrictions, allow override of the plan. The default is to disable derived Plans that use this plan as the base plan from overriding this plan's settings. For our purposes, we're leaving this option at default.
12. Click on **Save**.

Congratulations! You just enabled the server backup plan for IntelliSnap to manage and orchestrate hardware snapshots of your production environment on the FlashArray//X Arrays. The plan will also make a copy of the Snapshots to the local storage library sitting on FlashBlade.

Next, we will add a snapshot destination on the secondary site where IntelliSnap will orchestrate and manage Asynchronous snapshot replicas to the FlashArray//C.

1. In the Commvault Command Center navigation panel, click on **Manage > Plans**.
2. The Plans page will appear, with a list of all backup plans created. Under the Plan name column, click on the just-completed plan.
3. The plan's properties will appear. Scroll down on the right-hand side to the Backup destinations panel. (see figure 12 below)
4. Click **Add** and select **Snap Copy**.
5. In Add Backup Destination form:
 - a. In the Name field, specify the name of the new destination appropriate to the site.



- b. In the Copy Type field, leave default Vault/Replica selected.
 - c. In the Source field, select the primary site Snapshots to be the source for the remote Snapshots.
 - d. In the Storage field, select the remote site storage library; in this case, it's the FlashBlade Amazon S3.
 - e. In Retention rules, select the number of snap recovery points field and specify how many snapshots recovery points on the remote FlashArray//C you want to keep. Typically, this number will be higher than the number of snapshot recovery points set for the primary site. For example, to minimize storage costs, you might only keep a small number of snapshots in the primary site while keeping more snapshots on the lower-tier storage in the secondary site.
 - f. Click **Save**.
6. The Plan Backup Destinations panel will update with the third backup destination. Figure 12 shows what the backup destination panel looks like in a multi-site Intellisnap Plan.

Name ↑	Storage	Retention period	Source
Local Primary	Local Storage	1 Day	
Primary Site snap Snapshot primary	Local Storage	4 Recovery points	
Remote Snap Replica Default vault/replica	Remote-Storage	24 Recovery points	Primary Site sn...

[Configure region based storage](#)

Figure 12. Backup destination setup with multi-site Intellisnap

Next, we will associate a VM group with the newly created plan. If this is a net new Commvault environment, a VM group with the name Default will automatically be created as is the case in my Lab environment. You should avoid using the default group since it's a catch-all group, to avoid backups of unnecessary data, you should create your own VM group.

Note: If you don't have vCenter configured with your Commvault environment, please follow the instructions in [Commvault Documentation](#).

1. From Commvault Command Center, Navigation Panel, click on Protect > Virtualization.
2. Click on the VM groups tab, a list of all VM groups will appear; select the proper VM group.



3. VM group properties will appear, with Overview Tab Active. On the Policies Tile, click on edit.

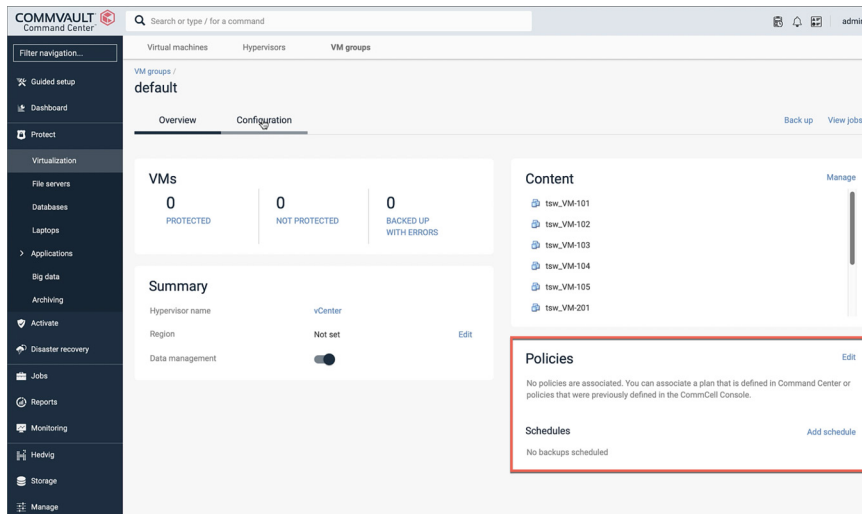


Figure 13. Associating Server Plan to VM Group

4. In the Edit Plan form, click on the pull-down menu next to **Plan**, select the Plan we have just created, and click **Save**.
5. Click on Configuration Tab, next to the Overview Tab.

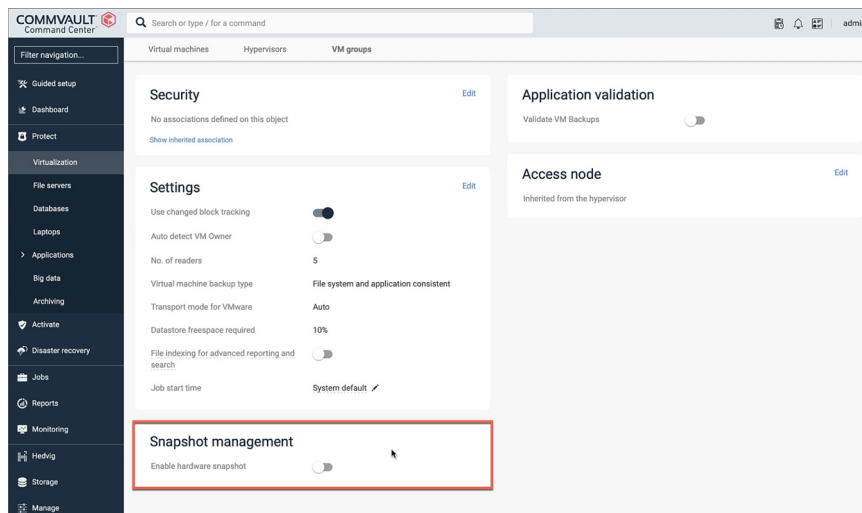


Figure 14. Enabling Hardware Snapshot with IntelliSnap on VM Group

6. In the Configuration tab, enable the Enable Hardware Snapshot option under the Snapshot Management Panel.
 - a. Switch the toggle for Enable Hardware Snapshot to On (Enabled).
 - b. The Snapshot Management Popup will appear.
 - i. Next to Engine Field, click on the pull-down menu, and select "Pure Storage FlashArray Snap."
 - ii. Below Snap mount Esx host, browse through vCenter cluster, and select determine which ESXi host can mount snapshots for recovery operations. Figure 15 is an example of how it should look.
 - iii. Click **Save**.



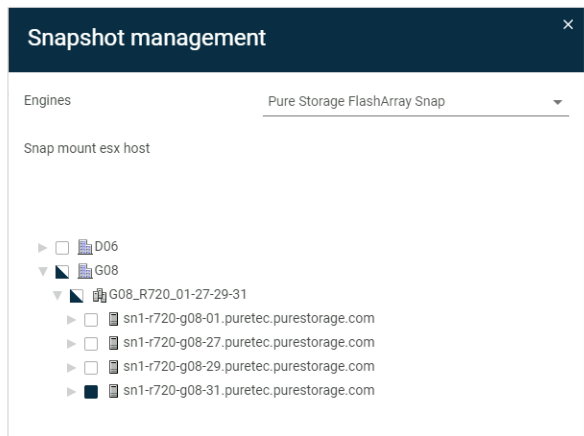


Figure 15. Enable snapshot management for VM group

7. The plan should start a new snapshot job shortly after the settings are saved to meet configured RPO. Snapshot jobs will be visible by clicking on either Jobs or Dashboard from the Navigation panel on the left.

Checking Backup Workflow

There are various ways besides investigating the job logs to check whether the Plan workflow is working.

Commvault Command Center will show all snapshots created from the Storage Arrays view:

- 1. On the Navigation Panel, click **Manage > Infrastructure**, click on the Arrays Panel.
- 2. Next, there will be a menu item on the right-hand side next to each of the managed Arrays, as shown in Figure 16, for one of the managed arrays.

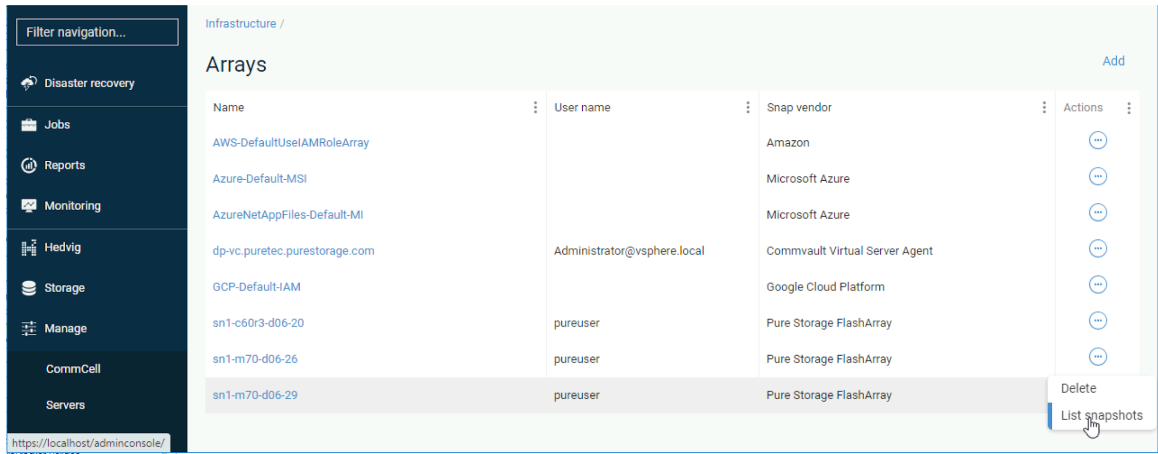


Figure 16. Getting the list of snapshots

3. Click on List Snapshots; you will get a list of the snapshots on that Array which IntelliSnap has created.

We can also check from the FlashArray//X arrays, or FlashArray//C's Purity//FA web interface.

Figure 17 below is a screenshot from the Purity//FA Protection page on FlashArray//C, showing the two protection groups created by Commvault IntelliSnap for the two primary FlashArray//X Arrays. Protection group names are the combination of the source FlashArray name, followed by a colon, "CV-" with a randomly generated number.





Target Protection Groups ^				1-2 of 2	
Name▲		Snapshots	Targets		
 sn1-m70-d06-26:CV-9DD2AF8-5		31.18 M	Allowed on this array		
 sn1-m70-d06-29:CV-9DD2AF8-5		43.70 M	Allowed on this array		

Figure 17. Protection Groups created by IntelliSnap

Restore Options

Using FlashArray with Commvault IntelliSnap to manage multi-site Asynchronous snapshot replication dramatically expands your ability to keep more recovery points. You also don't have to compromise with restore types. Figure 18 shows the list of recovery options available for a VM guest.

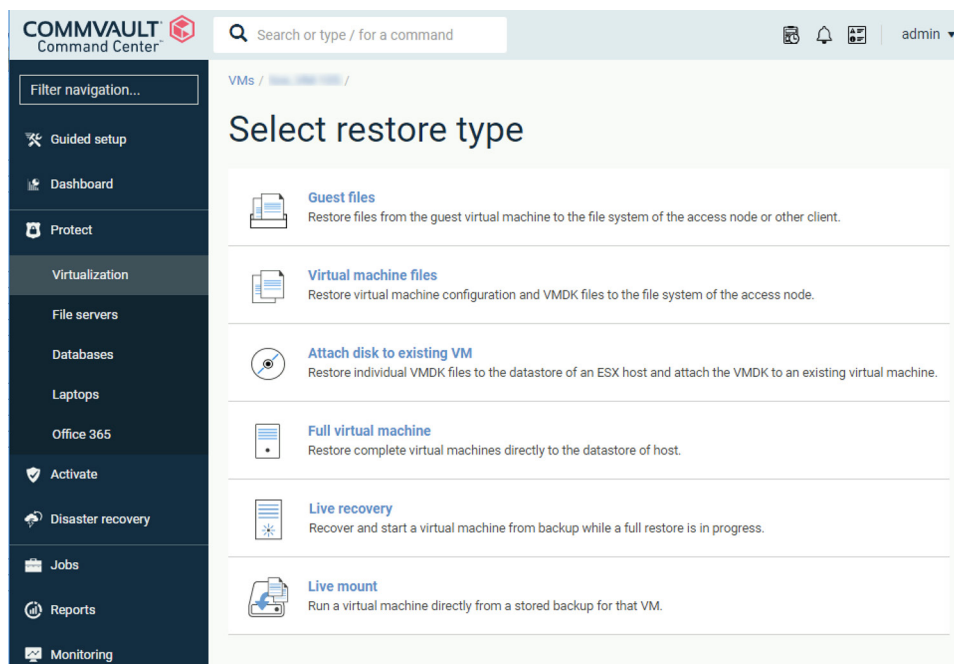


Figure 18. Recovery option for a VM guest

As an example, Figure 19 shows performing a file-level restore operation on a VM guest. If you select "Restore from default copy" (default option), the restore operation searches for the requested data in the primary copy in our case it's the snapshot copy sitting on the FlashArray//X arrays. Suppose the data requested for restore is not found on the primary copy, for example, due to the short-term retention. In that case, it will automatically select the next copy, which again in our case is the snapshot copy on the FlashArray//C array.

If you click on the Restore from the default copy list, in the upper-right corner of the page, you can manually restore from a specific copy of backup data residing on the primary snapshot, the replica snapshot, or backup on FlashBlade or other backup storage.



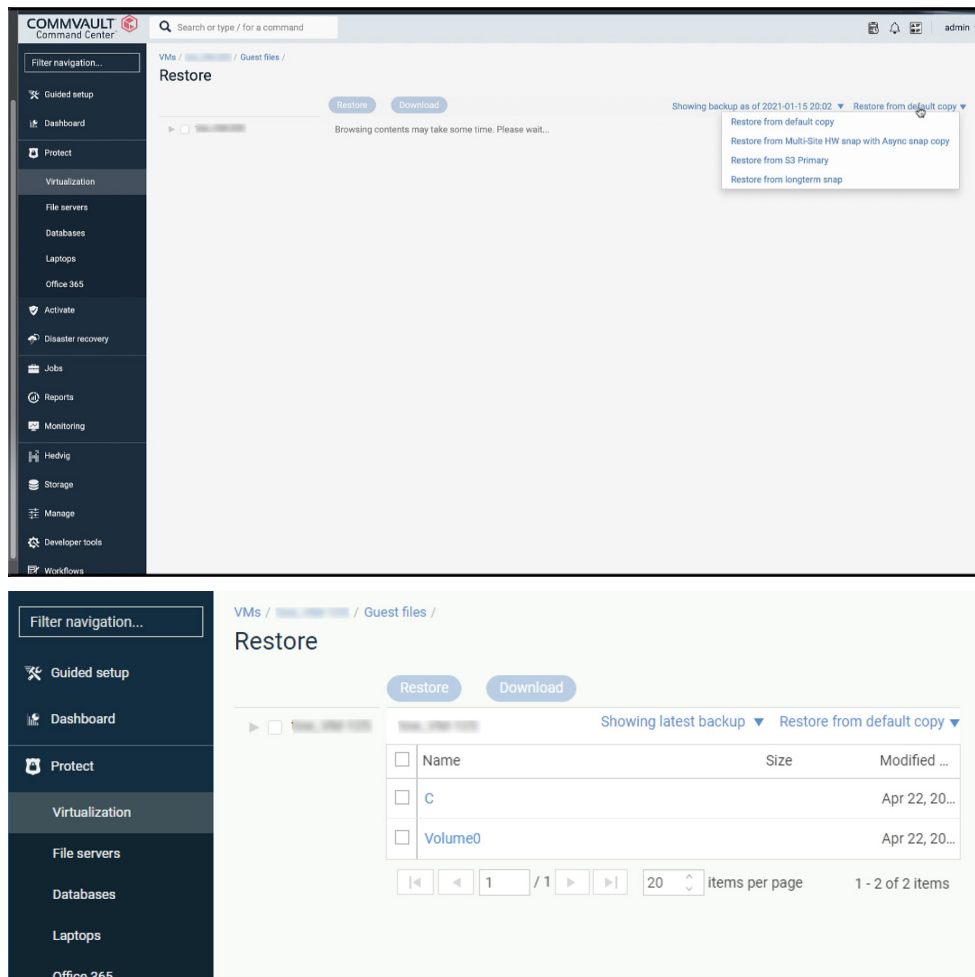


Figure 19. VM Guest files level restores

Conclusion

Creating Snapshots on the FlashArray is fast; they are also efficiently stored, as they are deduplicated and compressed. Even when replicated, snapshots are still deduplicated and compressed, saving a lot of expensive bandwidth.

At Pure, we understand that data protection is one of the main pillars of the modern data center. When it's critical to get your data fast, snapshots are among the fastest options around, and there is nothing better than using a single tool to get it. We worked diligently with Commvault to optimize the integration with IntelliSnap, so our customers can achieve their digital transformation goals.

Additional Resources

- Learn about Pure Storage [FlashArray//C](#), [FlashArray//X](#), and [FlashBlade](#).
- Discover [Commvault Solutions for Pure Storage](#).



About the Author



Tamer Swidan is a senior solution architect with Pure Storage. He is responsible for defining Pure Storage solutions and reference architectures for protecting and recovering primary workloads such as Oracle, SQL, and VMware. Tamer has 19 years of experience working in and with data protection hardware and software solutions, from an end user to subject matter expert consultant, followed by a data protection solution architect.

©2021 Pure Storage, the Pure P Logo, and the marks on the Pure Trademark List at <https://www.purestorage.com/legal/productenduserinfo.html> are trademarks of Pure Storage, Inc. Other names are trademarks of their respective owners. Use of Pure Storage Products and Programs are covered by End User Agreements, IP, and other terms, available at: <https://www.purestorage.com/legal/productenduserinfo.html> and <https://www.purestorage.com/patents>

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

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

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