

# Pure Storage and Cohesity Integration:

## Better Data Protection For Apps Running On FlashArray//M

### Executive overview

Pure Storage has been driving the transition to all-flash storage arrays for mission critical applications. FlashArray//M is typically used to run IO intensive workloads such as production databases and large virtualized environments. Pure Storage customers have the option to use array-based snapshots to provide fast, non-disruptive data protection for their mission critical applications. Compared to traditional backup solutions, array-based snapshots can provide much faster RPOs and RTOs, with much lower performance impact on the production applications.

Cohesity provides a hyperconverged platform designed to converge secondary storage silos – including backup infrastructure and target storage. The Cohesity solution includes distributed scale-out storage with best-in-class space efficiency, and converged data protection. Cohesity integrates natively with Pure Storage FlashArray//M to quickly and efficiently tier longer-term snapshots to the underlying Cohesity DataPlatform. Cohesity manages the snapshot retention policy across both primary and secondary storage, and can automate pre / post-scripts to make the snapshots application consistent.

With the joint solution, customers benefit from a secondary storage tier that reduces overall data protection costs and increases data availability with a second layer of protection. Cohesity provides native integration with all the leading public clouds, enabling customers to leverage the public cloud for long-term data retention. Customers can also use Cohesity to accelerate test/dev processes by instantly provisioning zero-cost clones of their backup snapshots for developers.

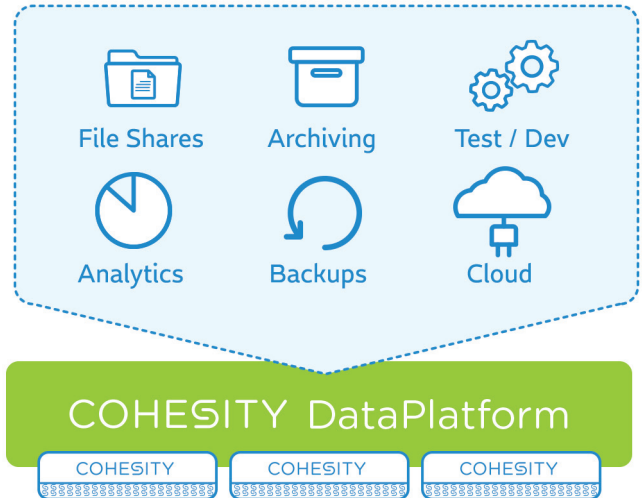
### Pure Storage and Cohesity – Solution Highlights

- Leverage Pure Storage FlashArray//M snapshots for data protection. Provide fast RPO / RTO and minimize performance impact on production apps
- Efficiently tier longer-term snapshots from Pure Storage FlashArray//M to Cohesity DataPlatform
- Increase data availability with two tiers of storage protection
- Reduce overall data protection costs with a space and cost-efficient secondary storage tier
- Automate snapshot creation and tiering process with Cohesity retention policies
- Provide app-aware snapshots with Cohesity automated pre/post-scripts
- Integrate natively with public cloud for long-term retention
- Accelerate test/dev processes by using Cohesity to instantly provision zero-cost clones of backup snapshots

## Cohesity Overview

Cohesity provides a hyperconverged platform designed to consolidate and manage all your secondary data at web scale – including backups, test/dev copies, files, analytics data, and objects. Cohesity provides distributed storage with best-in-class space efficiency, efficient copy data management, and native cloud integration. The platform incorporates data management functions for converged data protection, in-place analytics, and copy data management to automate DevOps workflows. Key capabilities of the Cohesity platform include:

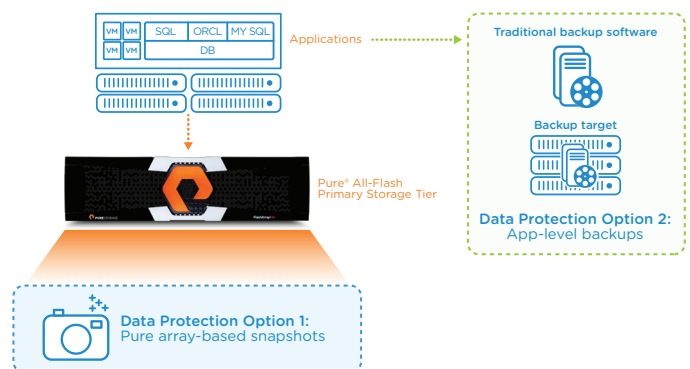
- Distributed platform: Scale performance and capacity linearly by simply adding nodes, eliminating the need for data migrations and forklift upgrades.
- Converged data protection: Replace multiple data protection silos (target storage, media servers, master servers, cloud gateways) with a single converged solution for backup and replication. Use Cohesity data protection for application-level and VM-level backups, or to manage array-based protection with Pure Storage.
- Global space efficiency: Global variable-length dedupe and compression.
- Cloud integration: Native cloud integration with AWS, Google Cloud Platform, and Microsoft Azure.
- Copy data management: Zero-cost snapshots and clones with instantaneous data access to quickly provision test/dev copies in support of DevOps workflows.
- Cost-effective secondary storage: Reduce storage TCO by 50% or more.



## Pure Storage FlashArray//M snapshots for data protection

Customers can utilize two different approaches when it comes to protecting their critical production data. They can use traditional backup applications that protect at the application level. Or they can use storage array-based snapshot solutions such as Pure FlashArray//M snapshots. Both these methods have their pro's and con's and ultimately provide IT professionals with several options which they should select based on business requirements.

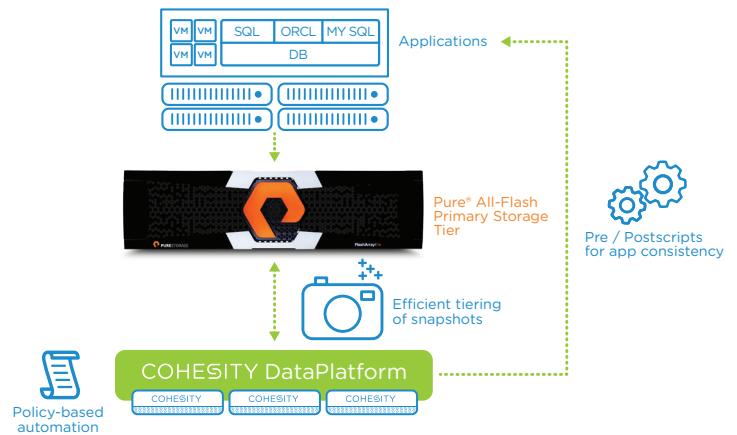
Pure Storage FlashArray//M snapshots offer the advantage of enabling fast RPOs (with frequent snapshots), fast RTOs (quick recovery) without impacting the performance of the primary application. The downside of using array-based snapshots is that the snapshot data resides locally, meaning that a loss of the primary storage array can result in data loss. Furthermore, as data changes over time, older snapshots will consume more space on primary storage arrays.



## Cohesity and Pure Storage Integrations

Cohesity integrates natively with Pure Storage FlashArray//M to improve snapshot-based data protection. The integration enables customers to dynamically tier snapshots between Pure Storage primary storage and Cohesity secondary storage to provide the best possible data protection. Customers can retain their most recent snaps on Pure Storage for faster data protection, while copying or moving older snapshots dynamically to Cohesity secondary storage. The integration consists of the following key capabilities:

- Efficient tiering of snapshots between Pure Storage and Cohesity using array snapshot APIs
- Automation across both Pure Storage and Cohesity using common policies created in Cohesity
- Pre and post-scripts to provide application-consistent data protection
- Ability to restore snapshots from Cohesity to any Pure Storage array (primary storage or alternate Pure Storage array)



## Cohesity + Pure = Better Data Protection

### Faster RPOs and RTOs

Traditional backup software typically provides only weekly full and daily incremental backups. In other words, up to a day of data could be lost between two backups. Backup frequencies are limited by the performance impact on production applications and limited backup windows. Recovery times are typically in the 10s of minutes or hours, because full backups have to be reconstructed and data has to be moved back to primary storage.

By using Pure Storage snapshots instead of traditional backup software, customers can provide much more frequent backups (faster RPOs) and faster recovery times (faster RTOs). Backups can be taken as frequently as every few minutes without impacting production apps. And recovery is almost instantaneous as the snapshot is immediately available on the primary storage array.

### Improve protection with two layers of storage

When data is only protected using snapshots on Pure Storage primary arrays, there is only one layer of protection. In other words, if the Pure array becomes unavailable, both the production data and the snapshots become unavailable.

By tiering snapshots to Cohesity, customers benefit from a second level of protection. Even if the Pure array becomes unavailable, data can still be recovered from Cohesity secondary storage.

### Reduce overall data protection costs

Keeping long-term snapshots on Pure can be an expensive proposition. Business imperatives and compliance / regulatory mandates require data to be retained for longer periods of time making it very expensive and thus impractical to retain them in all-flash primary environments. By tiering snapshots to Cohesity, which provides best-in-class space efficiency, customers can reduce their overall data protection costs.

## Automate snapshot creation and retention

Managing snapshots manually can be an operational burden. The number of applications to be protected, multiplied by the retention period, can lead to a very large number of snapshots to manage.

The Cohesity and Pure solution automates snapshot creation and retention. Customers assign simple policies in Cohesity, which automate the snapshot and creation process across both Pure Storage arrays and Cohesity.

## Provide application-consistent snapshots

Storage array snapshots are typically not application consistent, but crash consistent. Recovering an application from a crash consistent snapshot could take a long time as the data has to be recovered to an app-consistent state first. Cohesity automatically manages pre- and post-snapshot scripts to put applications in an app-consistent state prior to initiating the snapshot.

## Recover to any Pure Storage FlashArray//M

Snapshots can be recovered from Cohesity to any PureStorage FlashArray//M. Applications can be recovered not only to the original primary storage device, but also to an alternate array sitting either in the same datacenter or in a DR site.

## Integrate natively with the public cloud for long-term data retention

Cohesity provides native integration with the public cloud for long-term data retention. Customers can use Cohesity to move their longer-term snapshots to AWS, Microsoft Azure, Or Google Cloud. These snapshots can then be recovered by Cohesity and restored to any Pure Storage FlashArray//M, providing effective long-term data archival and Disaster Recovery protection.

## Accelerate test/dev processes with Copy Data Management

Cohesity provides the ability to instantaneously clone any snapshot at zero cost. These snapshots can then be presented to a test/dev Pure Storage array to provide a copy of the production application to developers. This process can be done very quickly and eliminates any potential impact on production applications. Developers can have faster access to their development environments. And the test/dev copies can be provided to any point-in-time available in the backup snapshots, for more efficient debugging and problem solving.