

# Better Target Storage for Backups

## Distributed, Globally Deduped and Cloud-Integrated Platform

### Looking for a better alternative to traditional dedupe appliances?

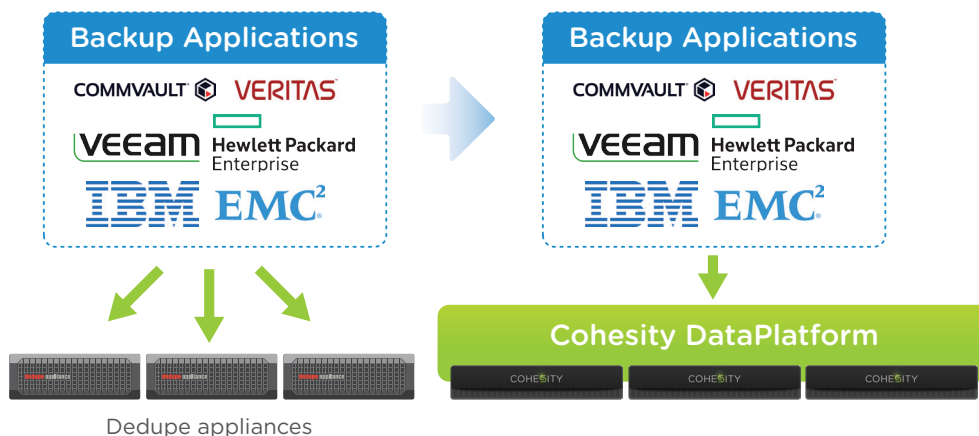
Over the past 10 years, dedupe appliances have seen tremendous success. Compared to tape libraries, these appliances offer faster backup and recovery, much simpler management, and the space efficiency of dedupe.

However, dedupe appliances are mostly based on legacy technology from more than 10 years ago, and have seen only incremental innovations since then. They are typically stand-alone appliances that don't scale out with a distributed architecture. This in turn results in expensive forklift upgrades and data migrations, inefficient management of capacity silos, and lack of global dedupe. Legacy appliances also often lack critical features required for modern private cloud environments, such as native integration with the public cloud, Quality of Service, and multitenancy to isolate data between tenants.

Cohesity addresses these challenges with a distributed platform that provides an ideal target for existing backup software.

### Cohesity as target storage for backups

Cohesity provides the only hyperconverged platform that eliminates the complexity of traditional data protection solutions by unifying your end-to-end data protection infrastructure – including target storage, backup, replication, disaster recovery, and cloud tiering. The Cohesity DataPlatform provides scale-out, globally deduped, highly available storage to consolidate all your secondary data, including backups, files, and test / dev copies. Cohesity DataProtect provides a complete backup and recovery solution, fully converged with the Cohesity DataPlatform. It simplifies backup infrastructure and eliminates the need to run separate backup software, proxies, media servers, and archival.



<p><b>✗ Stand-alone appliances</b></p> <ul style="list-style-type: none"> <li>• Forklift upgrades with data migration</li> <li>• High upfront CAPEX anticipating future growth</li> <li>• Complex capacity planning across silos</li> <li>• Appliance failures lead to downtime</li> </ul>	<p><b>✓ Scale-out platform</b></p> <ul style="list-style-type: none"> <li>• Non-disruptive upgrades</li> <li>• Pay-as-you-go linear scalability</li> <li>• Global resource pooling and load balancing</li> <li>• No single point of failure</li> </ul>
<p><b>✗ Local dedupe</b></p>	<p><b>✓ Global dedupe</b></p>
<p><b>✗ Bolt-on cloud gateway</b></p>	<p><b>✓ Native cloud integration for tiering, archival and replication</b></p>
<p><b>✗ Performance capped by controller</b></p>	<p><b>✓ High performance with parallelized ingest</b></p>
<p><b>✗ Designed for legacy datacenters</b></p> <ul style="list-style-type: none"> <li>• Every backup job gets same priority</li> <li>• No isolation within appliance</li> <li>• Manual management and configuration</li> </ul>	<p><b>✓ Support for hybrid cloud and multitenant environments</b></p> <ul style="list-style-type: none"> <li>• QoS priorities</li> <li>• Isolation between tenants</li> <li>• REST API for orchestration and DevOps</li> </ul>
<p><b>✓ Enterprise ready features such as replication and software-based encryption</b></p>	

## Advantages of a scale-out platform

The scale-out architecture of Cohesity DataPlatform provides inherent advantages compared to conventional stand-alone appliances:

**Non-disruptive upgrades:** Conventional dedupe appliances have to be upgraded or retired every few years. This requires purchasing expensive new hardware, and performing complicated data migrations from the old appliance to the new. Cohesity eliminates this problem by enabling customers to expand their clusters dynamically by simply adding newer appliances. A single cluster can include multiple generations of appliances, and older appliances can be de-commissioned automatically with no data migrations.

**Pay as you grow:** The Cohesity scale-out architecture enables a simple, pay-as-you-grow purchasing model. With legacy appliances – customers typically have to purchase the amount of capacity they anticipate needing until the next refresh cycle. This is typically an inaccurate guessing game that requires high upfront CAPEX investment. With Cohesity, you can start with a single block (3 or 4 nodes in a single chassis). And then add small, individual nodes to the cluster as capacity needs increase. No more second guessing your long-term capacity requirements and spending high upfront CAPEX in anticipation of those needs.

**Global resource pooling and load balancing:** With traditional appliances, capacity has to be managed in silos and backup jobs have to be manually distributed between appliance silos. Cohesity provides a global pool of capacity across an entire cluster, thereby eliminating the need to manage capacity at the appliance level. Storage capacity and IO performance is automatically balanced across all the nodes in the cluster. Capacity planning becomes much simpler and capacity utilization increases.

**No single point of failure:** Cohesity DataPlatform keeps two copies of each block of data across multiple nodes in a cluster. Even in the event of a complete node failure, the data continues to be available with no downtime. With traditional appliances, an appliance failure would result in data being inaccessible for backups or restores.

## Increased space efficiency

**Global deduplication and compression:** Cohesity provides variable-length global dedupe and compression. Data is deduplicated across all the nodes in a cluster. With global dedupe, the data reduction ratios are significantly higher than when dedupe is only performed locally within an appliance. Dedupe can be configured to be either in-line or post-process.

## Native cloud integration

Cohesity provides native integration with all the leading public cloud providers including Google Cloud Storage Nearline, Microsoft Azure, Amazon S3 and Glacier. Traditional dedupe appliances don't provide this integration natively, and a separate storage gateway has to be bolted on to leverage the public cloud. Cohesity can leverage the cloud for the following use cases:

**CloudTier:** Leverage compelling cloud economics with native cloud integration for data tiering. Colder data is tiered to the cloud to increase the capacity of the Cohesity cluster.

**CloudArchive:** Replace tape archive and off-site data protection with long-term archival to all the leading cloud providers.

**CloudReplicate:** Replicate your data to the cloud for disaster recovery and seamless cloud migrations.

## Increased performance with parallelized data ingest

With traditional appliances, multiple backup jobs have to run sequentially on a single dedupe appliance. This can lead to performance bottlenecks and long backup windows. Cohesity DataPlatform parallelizes backup jobs across multiple nodes in a cluster, resulting in faster backups and shorter backup windows.

## Support for private clouds and multitenant environments

To support large scale environments with multiple tenants and DevOps requirements, Cohesity provides Quality of Service, data isolation between tenants and fully documented REST API.

**QoS:** Set the relative priority between jobs and control the quantity of flash allocated to any job, thereby ensuring that high priority backup jobs perform within the required SLA.

**Tenant isolation:** Tenant isolation provides data isolation between backup jobs, ensuring no data leakage and keeping potential malware confined within a single job.

**REST API:** Cohesity DataPlatform exposes a full set of REST APIs that can be used for integration with leading DeOps tools like Chef and Puppet.

## All the enterprise features of conventional dedupe appliances

In addition to all the unique differentiated features of Cohesity DataPlatform, DataPlatform also matches all the leading dedupe providers in terms of supporting enterprise-class features. Specifically, Cohesity also provides the features below out of the box:

- Unlimited replication with flexible topologies
- Real-time analytics on data and backup job utilization
- Software-based encryption meeting the AES 256 standard

## Beyond target storage: Cohesity DataProtect

Cohesity DataProtect is an end-to-end backup and recovery solution that is fully integrated with the platform to provide simple data protection with 15 minute RPOs and instantaneous RTOs. Cohesity DataProtect converges all your data protection infrastructure on Cohesity DataPlatform - including target storage, backup, replication, disaster recovery, and cloud tiering. Cohesity DataProtect ensures strict application SLAs by providing fast Recovery Points and near-instantaneous Recovery Times, all while cutting data protection costs up to 50%.

Customers can choose to protect some of their applications with DataProtect, while continuing to use their legacy backup software for the remainder of their applications. A non-disruptive transition strategy is to start using DataPlatform with existing backup software, and over time migrate more applications to DataProtect.

