

Datrium Cloud DVX

DATRIUM CLOUD DVX DATASHEET

Cloud DVX Recovery Service Benefits

As-A-Service Simplicity

One-click set-up. Self-healing availability. Automated upgrades. Proactive support.

Lowest Total Cost of Ownership

10x lower AWS costs with global deduplication and forever incremental backups to cloud.

Fastest Recovery Times

Dedupe-aware retrieves. Granular VM, vDisk or guest file recovery. Direct to host retrieves. Always synthetic full.

End-to-End Security

On-wire & at-rest encryption. No VPN required. FIPS-140-2 validated.

Public Cloud: A new backup & archival target

For a long time, the concept of disk to disk to tape (D2D2T) or disk to disk to disk (D2D2D) has defined data distribution across production, backup and archival. Recently, however, public cloud object storage services have become an attractive alternative to disks or tapes for off-site long-term retention use cases.

Public clouds offer three key advantages that are hard to achieve with disks or tapes:

Pay as you go pricing –

Unlike disk or tape systems that need long-term sizing forecasts and upfront CAPEX, public clouds offer pay-per-use pricing and enable massive scale data consolidation.

No media management –

Unlike disk or tape systems, where the media has to be refreshed every few years due to support expiry or media degradation, public cloud infrastructure free up end-customers from risky, costly and time-consuming media management.

Instant Data Access –

Unlike alternative media options, especially tapes, which can take days to make data accessible, data in public cloud is accessible within a few minutes to hours. For example, where this data was not easily accessible previously, the public cloud enabled companies to perform analytics on large amounts of data for business profit.

However, despite the clear benefits of using public cloud as a backup or archival store, existing implementations have severe shortcomings, which have inhibited adoption. Let's review these next.

Challenges with existing backup to cloud offerings

There are numerous solutions that enable leverage of public cloud object storage for retaining backups from on premise systems. Most use direct S3 APIs and a few have done a part of their software to consume cloud compute and storage. However, all of them suffer from one or more of the following challenges:

Challenge 1: High Complexity –

most all existing implementations require end-customers to learn and master AWS to some degree. AWS might offer all its capabilities as a service but operating these vast variety of services in an enterprise-grade way requires deep expertise. Modern hyperconverged primary (HCI) vendors leave monitoring, availability, upgrades and support to the end-customer making it another solution that needs to be learned and managed. In the public cloud, customers want service-based consumption model, and this has clearly not been met. While modern hyperconverged backup (HCI Backup) vendors eliminate this somewhat by using direct S3 APIs, the fact that you need a separate backup-only appliance in addition to production storage to get to cloud adds a lot of management complexity in itself. Using direct S3 APIs also create some inefficiencies.

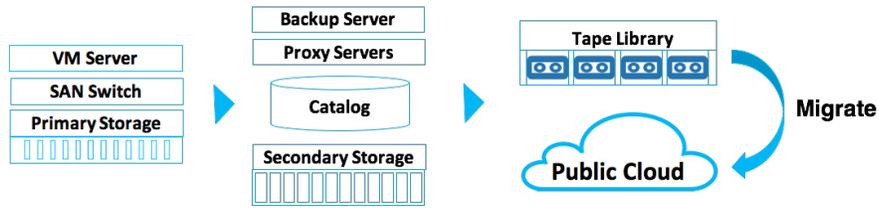


Figure 1: Legacy backup workflow migrating to cloud

Challenge 2: High Total Cost of Ownership –

None of the modern VM-centric HCI or HCI backup vendors offer end-to-end global data reduction in transit from on premise to public cloud and at-rest in the public cloud. This is a critical miss since backup economics depend heavily on reducing the large amount of backup data in the cloud to the absolute minimum. As data sizes approach typical backup and archival capacities, the cost of public cloud storage heavily outweighs other associated costs. Cost suffers even more heavily with HCI Backup vendors since they add another appliance to the cost footprint of on premises infrastructure. In addition, since they write to S3 directly without compute in cloud massaging the data, they require regular full backups of data for every few incremental backups. This increases the total capacity required and significantly increasing cloud costs.

Challenge 3: High Recovery Times –

One of the primary reasons for backup and archival is to ensure timely data recoverability in an unexpected event. However, the lack of global data reduction and forever-incremental backups to public cloud leads to long data copy times back from the public cloud and the RTO suffers. Most primary storage and HCI systems have datastore or LUN level recoverability, which is a lot of data to transfer when a VM or virtual disk recovery would suffice. HCI backup vendors do a nice job of enabling in-guest file recoverability. But, recovery requires a full copy of data from cloud to HCI backup system, and then to the primary system. Combined with the need for regular fulls, synthesization of fulls from incrementals, and hydrated copy to primary, the RTO can be very high.

Datrium has evaluated and addressed these challenges in a transformational way. Let’s review that next.

Introducing Datrium Cloud DVX

Datrium DVX, a fundamentally new way to converge virtualization, compute, storage and backup for on premise workloads, now runs natively in Amazon Web Services (AWS). Datrium Cloud DVX delivers the same unmatched, always-on data reduction (global dedupe and compression) efficiencies and scaling flexibility of its on-prem sibling.

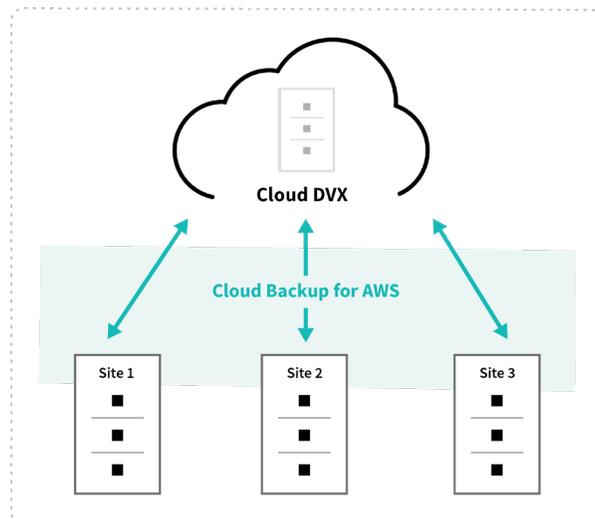


Figure 2: Datrium Cloud DVX

DVX split-provisioning architecture enables IO processing to scale with compute resources independent of capacity. This provides the ultimate flexibility in right sizing speed to capacity and is a natural fit for hyper-scaler cloud architectures featuring a similar separation of compute (EC2) and capacity (S3).

Datrium Cloud DVX recovery service is the first modern converged platform that offers backup to public cloud as-a-service with global deduplication & forever-incremental backups. End customers can simply consume the cloud capacity for backup at the lowest cost of ownership.

Let's dive into the key breakthroughs and differentiation of Cloud DVX:

Cloud DVX Advantage 1: As-A-Service Simplicity –

Cloud DVX automates all the key day one and day two tasks so customers don't have to specialize in AWS infrastructure mechanics. Compared to a leading HCI vendor, customers perform 75% fewer administrative tasks with Cloud DVX¹.

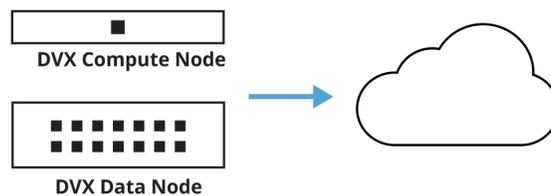


Figure 3: Backup from on-premise tier 1 DVX straight to Cloud DVX minus separate backup hardware or gateways

Key simplicity features of Cloud DVX include:

- 1] One-click set-up:** Setting up Cloud DVX in AWS is as simple as pairing another replication target to an on premise DVX. Just select the AWS region, enter credentials and click finish. In just a few minutes, Cloud DVX is setup in AWS automatically and customers can start replicating snapshots immediately.
- 2] Self-healing availability:** Cloud DVX leverages server-less management of AWS resources to ensure continuity of data replication tasks. In case an EC2 instance suffers an outage, another EC2 instance is automatically spun so that ongoing replication tasks can finish successfully.
- 3] Automated upgrades:** Software upgrades for Cloud DVX are handled automatically without any customer intervention.
- 4] Proactive Support:** Just like on premise DVX, Cloud DVX sends a range of telemetry inputs to the Datrium support teams for proactive issue identification and resolution.

Cloud DVX Advantage 2: Lowest Total Cost of Ownership –

On premise and Cloud DVXs combine to deliver unmatched management and data efficiency. Compared to Cloud DVX, a leading HCI vendor requires 3X the AWS costs while a leading HCI backup vendor requires 10x AWS costs¹!

Key cost optimization feature of Cloud DVX include:

- 1] Flash/disk to cloud:** DVX collapses three-tier D2D2C or F2D2C into a two-tier model by converging primary and secondary in one system and connecting to the cloud directly. By delivering better performance, greater data efficiency and granular data management in one system, DVX enables customers to save over 70% in purchase costs alone, compared to separate primary & HCI backup vendor solutions.

2] Global deduplication and compression on-wire/at-rest: Cloud DVX leverages compute in the cloud for in-cloud deduplication of data across multiple sites or systems. Only unique data is replicated from on premise to the cloud and back. All data stays compressed right from initial ingestion. With an average 3X local data reduction and 1.5X across systems or sites, Cloud DVX enables an average of 4.5X data reduction of all data stored in the cloud leading to huge capacity and cost savings.

3] Forever Incremental Backups: After initial seeding of snapshot data to the cloud (which is locally de-duplicated and compressed before transfer), all subsequent snapshots result in only the differences being transferred. This eliminates regular full backups that are required by leading HCI backup vendors reducing capacity and cost in the cloud.

4] Built-in end-to-end encryption (no VPN required): DVX natively encrypts all data in-transit to the public cloud using FIPS-140-2 validated SSL eliminating the need for costly public cloud VPN services, which are charged by transferred data capacity and connection-hours basis. This ensures data security while keeping ingest network cost zero.

Cloud DVX Advantage 3: Fastest Recovery Times – Cloud DVX reduces the amount of data and the number of hops from cloud back to on premise. Cloud DVX provides 2X lower recovery times compared to leading HCI and HCI backup vendors².

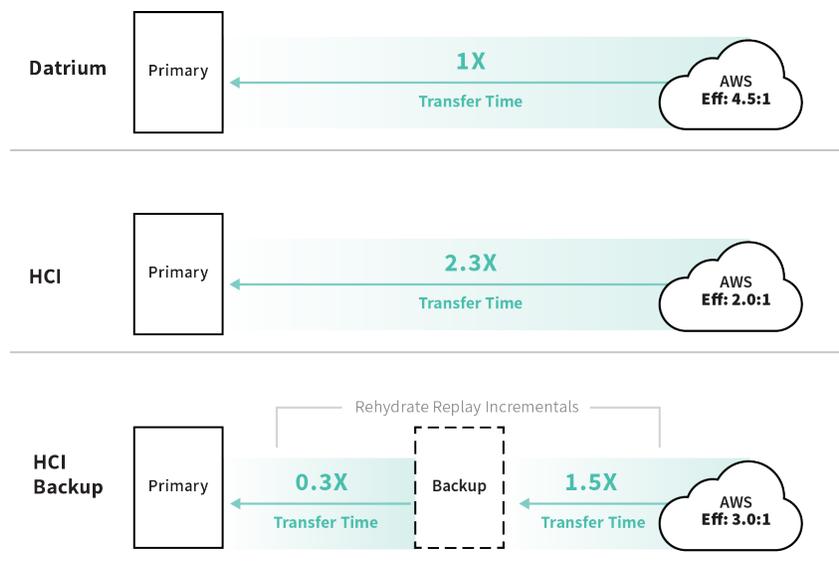


Figure 4: Recovery Time Compare: DVX vs HCI primary vs HCI backup

Key recovery time optimization features of Cloud DVX include:

1] Dedupe-aware retrieves: Cloud DVX minimizes the amount of data transferred from public cloud to on premise by sending only unique data not already available on premise.

2] Granular recovery: Cloud DVX further reduces the amount of data transferred by allowing end-customers to recover not just virtual machines (versus LUNs or Datastores) but also individual virtual disks, datastore-files (OVAs, ISOs), persistent container volumes and even guest files.

3] Direct to host: With a two-tier on prem to cloud model enabled by on premise DVX and Cloud DVX, data is recovered from the cloud straight to the host where it can be instantly utilized on the primary infrastructure.

4] Always synthetic full: Cloud DVX auto synthesizes incrementals while retaining the data efficiency benefits. This helps in the recoveries where restores are instant and do not need to wait for incrementals to be first applied like in case of HCI backup vendors. This enables instant restores as soon as data has finished transferring to the on premise DVX.

The Bottom Line

With Cloud DVX, consuming public cloud resources for long-term off-site retention of backup data is easier than ever before. With global data efficiency, granular object management, end-to-end security and massive consolidation of on premise infrastructure into a single converged system, customers can realize dramatic total cost of ownership savings while meeting their RTO SLAs.

Datrium Cloud DVX Quick Specifications

Public Cloud	Amazon Web Services
Resources	Resources = EC2 - 1 i3.xlarge (30 TB usable) or 1 i3.4xlarge (115 TB usable), S3 standard, DynamoDB
Cloud DVX software	Resources = EC2 - 1 i3.xlarge (30 TB usable) or 1 i3.4xlarge (115 TB usable), S3 standard, DynamoDB
Simplicity	Serverless upgrades, availability & proactive support
Data Efficiency	Global dedupe and compression across sites, systems & objects. Forever incremental snapshots. Synthetic fulls.
Retrieve Granularity	Guest Files, Virtual Disks, OVAs, ISOs, Virtual Machines
Security	Always-on SSL (prem to cloud); S3 envelope encryption (AWS). No VPN required

Learn more about Datrium at www.datrium.com.

¹ <https://www.datrium.com/resources/esg-cloud-dvx/>

² Based on amount of data transfer needed from cloud to on-premise (in case of HCI backup, full hydrated from secondary to primary system at 10X speed). Data reduction assumptions: Datrium - 3x local, 1.5x cloud; HCI - 2x local, 1x cloud; HCI Backup - 3x local, 1x cloud