

Storage Tiering with Cloud Volumes ONTAP Cuts Costs on AWS, Azure, and Google Cloud

While the cloud presents a less expensive way to store data than the traditional data center, enterprise data sets are growing so large that bill shock is a common event during cloud migration.

One way to reduce cloud costs is to tier infrequently used data to less-expensive storage formats. However, public cloud providers only offer data tiering between classes of their object storage offerings. While tiering between object storage classes can lower object storage costs, that won't work for production workload data that needs to be accessible for high-performance systems based on block storage. Fortunately, NetApp has a way to do this.

Cloud Volumes ONTAP data tiering makes it possible to automatically and seamlessly move data between block storage and object storage on AWS, Azure, or Google Cloud. Data tiering with Cloud Volumes ONTAP can also take advantage of the lower-cost storage classes within each object format service. In this eBook, we'll take a closer look at the data tiering capabilities of Cloud Volumes ONTAP.


We'll cover the ins and outs of data tiering with Cloud Volumes ONTAP, including:

- What Is Data Tiering?
- How Data Tiering Works on Cloud Volumes ONTAP
- Tiered Storage on AWS
- Tiered Storage on Azure
- Tiered Storage on Google Cloud

What Is Data Tiering?

Data tiering is the concept of moving data from one type of storage system to another. This technology isn't inherently tied to the cloud—for example, moving data from a performant on-prem system to another, less-performant system—but the cloud does make it more agile and less expensive, as cloud object storage is available on demand and does not require CAPEX investment.

The primary reason for moving data between tiers is to find and keep data storage in the most optimal storage type for its usage pattern. For instance, data that is always in use (which is typically referred to as “hot data”) should stay on lower-latency storage types, such as block storage. Data that isn't frequently accessed but may still be needed again (which is called “cold data”) is more appropriately stored on higher latency, less-expensive storage types, such as object storage. Any data can qualify as cold data. It can be the “usual” infrequently used data such as snapshots, DR environments, long term backups, etc., but it can also include active, live data that isn't accessed all that frequently.



The main challenge can be identifying which data belongs in which tier. Another challenge is finding a way to move that data between two tiers without impacting performance.

Cloud Volumes ONTAP can intelligently detect which data is cold in your block storage on Amazon EBS, Azure Disk, or Google Cloud Persistent Disk and move that data seamlessly and automatically to object storage on Amazon S3, Azure Blob, or Google Cloud Storage, respectively.

Since object storage costs are significantly lower than block storage costs, the savings that can be expected with this capability are significant – depending on the use case and type of data, some workloads can reduce their cloud storage costs by 90%.

This data tiering function can also be leveraged to work around the storage limits of native cloud managed services, such as RDS, by sending the less-frequently used data to object storage tier.

How Data Tiering Works with Cloud Volumes ONTAP

Cloud Volumes ONTAP tiers data between these low-latency performance tiers based on block storage and these object storage-capacity tiers:

- Amazon EBS – Amazon S3 (Standard, Standard-IA, One Zone-IA, and Intelligent-Tiering)
- Azure Disks – Azure Blob (Hot access and Cool access)
- Google Persistent Disks – Google Cloud Storage (Nearline, Coldline, Archive)

There are three different policies to tier data with Cloud Volumes ONTAP: Auto, Snapshot Only, and All, and each volume can have its own unique data tiering policy.



Tiering data to object storage

Volume Tiering Policy

- ☐ All - Immediately tiers all data (not including metadata) to object storage.
- ☒ Auto - Tiers cold Snapshot copies and cold user data from the active file system to object storage.
- ☐ Snapshot Only - Tiers cold Snapshot copies to object storage
- ☐ None - Data tiering is disabled.

Auto

Using the Auto tiering policy will tier any cold data that Cloud Volumes ONTAP detects in your block storage volume to object storage, including both active file system data and snapshot copies. The default cooling period is 31 days, which can be changed by the user.

When Cloud Volumes ONTAP detects random reads, the data will be termed hot and is automatically brought back up to the performance tier. Sequential reads, such as those done by virus scans, will be ignored and the data will remain cold and in object storage.

Snapshot Only

With the Snapshot-only data tiering policy, only NetApp Snapshot™ data is tiered to the capacity tier. The cooling period for the data in this option is 2 days by default, but that period can be changed by the user. If this snapshot data is read, the data will be considered hot and automatically moved back to the performance tier.

All

This is the fastest way to tier data. Using the All policy immediately tiers all data to the capacity tier. The ideal use case for this policy is for secondary data copies such as disaster recovery, backup, or long-term retention data copies. If activated through a read operation, the data is not moved back to the performance tier. Instead, it will be read directly from the capacity tier and cached on the performance tier.

Tiered Storage on AWS

With Cloud Volumes ONTAP data tiering for AWS storage, the performance tier is located on Amazon EBS and the capacity tier on Amazon S3. The capacity tier can be changed between the different AWS storage tiers used by Amazon S3.

The initial storage class for your tiered data will be Amazon S3 Standard by default. However, if you do not plan to

access the cold data frequently, you can further reduce your Amazon S3 storage costs by changing a system's tiering level. Currently, Cloud Volumes ONTAP data tiering supports the Standard-Infrequent Access (Standard-IA), One Zone-Infrequent Access, and S3 Intelligent-Tiering classes.

Note that while these storage classes have lower storage rates than S3 Standard, they do have higher data retrieval costs. One Zone-Infrequent Access also provides less redundancy. Also note that Cloud Volumes ONTAP storage tiering does not support use of the Amazon S3 free tier.

To see the business case for Cloud Volumes ONTAP storage tiering on AWS, consider the following: Provisioned storage for Amazon EBS General Purpose SSD in the US-East region would cost \$0.10 per GB-month. That means keeping 10 TB of disaster recovery environment data would cost \$12,000 annually. By tiering the 10 TB to Amazon S3 with One Zone-IA, the least expensive of the AWS storage tiers, where storage rates are the lowest, normal storage for the data without access would cost just \$1,200 per year (excluding the 30 days on Standard storage during the cooling period).

Tiered Storage on Azure

Cloud Volumes ONTAP for Azure storage takes advantage of different Azure storage tiers. For the performance tier, single node Cloud Volumes ONTAP uses Azure managed disks, while Cloud Volumes ONTAP HA uses Azure premium page blobs. In both cases the capacity tier is on Azure Blob

****Note that Hot access and Cool access are branded terms for Azure and the data stored in these tiers should not be confused with what NetApp considers hot and cold data; both the Hot access and Cool access tiers on Azure Blob are considered parts of the cold data capacity tier used by Cloud Volumes ONTAP.**

storage. Cloud Volumes ONTAP supports two of the Azure Blob storage tiers: the Hot access tier and the Cool access tier.

The Hot access tier provides low-cost access for data that may need access at some point for read/write operations. The Cool access tier offers lower-cost rates for storage than the Hot access tier, although its transaction and access costs are higher. Because of these extra fees, the data on the Cool access tier should only be data you need to access on rare occasions.

Cloud Volumes ONTAP will use the Hot access tier by default as the capacity tier. Users who want to take advantage of the lower costs of the Cool access tier can choose to further tier their cold data to that storage class. Doing so will first store the data in the Hot access tier, and if the data remains un-accessed there for 30 days (or another period set by the user), it will tier down automatically to the Cool access tier.

To see the business case for data tiering with Azure, consider the following: In the East US region, a Premium SSD of 1TB costs \$135.17 per month. To store 10 TB of data that would require 10 disks, costing \$1,351.70 per month or around \$16,220 annually. Tiering the 10 TB of data to Azure Blob's Cool access tier, where data costs \$0.0152 per GB per month, and leaving it unaccessed would cost just \$1,867 per year.

Tiered Storage on GCP

With Cloud Volumes ONTAP for Google Cloud, the performance tier uses zonal Persistent Disks, and the capacity tier uses Google Cloud Storage, in either Standard, Nearline, Coldline, or Archive class storage. Cloud Volumes ONTAP supports both Standard and SSD zonal Persistent Disks so that the choice will depend on your performance requirements.

Choosing the correct capacity tier depends on your use case. Nearline, Coldline, and Archive storage cost less but have minimal storage duration, data retrieval, and higher operational charges. If you have data that will infrequently, rarely, or never change once written, using these storage classes could further reduce your costs. But besides the lower costs, there are other factors to keep in mind: the Archive Storage class has no availability SLA.

To see the business case for Cloud Volumes ONTAP storage tiering on GCP, consider the following: Provisioned storage for Google SSD Zonal persistent disks in the us-east1 region costs \$0.17 per GB-month; therefore, storing 10 TB of data would cost your company \$20,400 annually. Tiering those 10 TB to Google Cloud Storage Standard class, which costs \$0.02 per GB-month excluding access charges, would cost just \$2,400 per year.



Keeping data in the Coldline storage class costs \$0.004 per GB-month, which for 10 TB comes out to \$480 per year, excluding other charges. The Archive storage class's GB-month cost is \$0.0012, excluding additional charges, which totals a relatively small \$144 per year for 10TB.

Before you choose a storage class, ensure that your use case for that class is valid. Consider an extreme example of 10 TB of compliance data, which you're required to retain for three years, stored in Archive class storage. Retrieval costs for this data is \$0.05 per GB, which means that if you needed to access all of that data at once it would total \$500. While that's a considerable savings over three years of storage on Coldline or Nearline, if you need to access this data several times a year, Coldline or Nearline may better suit your needs.

Conclusion

Businesses thrive on data, so it's crucial for your most important data to be available for rapid, performant use. However, storing that data in high performance disk types comes at a cost. These costs can be mitigated by tiering infrequently used data to object storage, where the price for storage is a fraction of the block storage cost, with a latency that fits data not expected to be used often.

With Cloud Volumes ONTAP, you can create a tiered storage architecture for seamless and automated storage tiering between AWS storage tiers, Azure storage tiers, or GCP storage tiers. This makes sure that you not only save costs on storage but also eliminate timely and possibly costly errors in moving data between tiers.

For more information about Cloud Volumes ONTAP and NetApp's other solutions, visit www.masterclass.eplus.com



Not many partnerships are as tried-and-true as ours.

No need to start from scratch when you choose ePlus and NetApp for your data management and integrity needs. From ransomware resiliency and data protection to maximizing the value of hybrid cloud and integrating new technologies, we'll bring broad experience mixed with strategic vision that will help you build upon and extend your success.