NetApp

FIVE WARNING SIGNS
THAT YOU MAY HAVE A PUBLIC CLOUD SPENDING PROBLEM

and 10 things you can do now to control storage costs

CONTENTS

NetApp

INTRODUCTION

Cloud storage is one of the fastest-growing business resource demands on the planet.

In fact, approximately 402.74 million terabytes of data are generated every day[1], and roughly 180 zettabytes will be created by 2025[2]. By our count, that's enough to fill around 116 quadrillion 3.5" floppy disks. (That's a particularly helpful reference if you happen to be reading this in your DeLorean time machine circa 1985.)

Simultaneously, the countless data users across all those data-generating organizations are interacting more frequently with their data. It's estimated that users in the datasphere interact with data approximately once a minute, and by 2025 that frequency is expected to increase to one interaction every 18 seconds.[3]

Don't hit the panic button yet; your organization will be just fine. In fact, data is your key to thriving in whatever next normal gets thrown your way. And cloud is absolutely a good thing for your business—you just need the right strategy to make it work harder for you. In this cloud spending guide, we walk you through strategies designed to help you keep your head above water in the inevitable data tsunami that's headed your way.

[1] Fabio Duarte, Amount of Data Created Daily (2024), Exploding Topics, June 13, 2024.

[2] Petroc Taylor, Volume of data/information created, captured, copied, and consumed worldwide from 2010 to 2020, with forecasts from 2021 to 2025, Statista, November 16, 2023.

[3] IDC, The Digitization of the World: From Edge to Core (2018).

n NetApp

STATE OF THE CLOUD

From data center doldrums to cloud calm

Before the cloud, your organization enjoyed a world in which capital expenditures (capex) bought ample storage that existed within a fixed data center ecosystem. Multiple data types coexisted on similar hardware, making the allocation of storage—regardless of whether you **really** needed it or not— a reasonably simple task. The IOPS were glorious and abundant; the security, airtight like a Vegas casino. You probably had a team of storage administrators on the payroll to attend to your every storage whim.

But then something broke down, and the corridors of bureaucracy and its cascade of red tape hindered quick recovery. You dreamed of the day when your storage would work harder for you, like an overeager intern poring over a stack of spreadsheets.

Then your dream became some version of reality—you made it to the cloud. And today, you've exchanged the complexity of data center management for an easy monthly operating expense (opex) that offers virtually infinite scalability and flexibility.

Sounds great, right? But as any astute accounting major can attest, that opex can sneak up on you at the end of the month, when your profit margin resembles a slowly disappearing oasis. Paying for storage is now based on multiple factors that involve separate storage types, tiers, protocols, and costs.

Inevitably, unexpected costs will bite you in your bottom line. The culprit? Any combination of a lack of tiering, legacy virtual machines (VMs) that aren't deprovisioned, and unnecessary snapshots, to name a few.

Alas, simply living the cloud dream is not enough.

Whether you're seeking to optimize your organization's cloud storage, or just searching for ways to cost-effectively expand your cloud storage use, you're going to be constrained by your organization's insatiable need to lower the cost of cloud storage.

Sound familiar?

Let's talk about the five warning signs that you may have a cloud spending problem, and then we'll give you what you really came for–10 ways you can immediately address out-of-control cloud costs, right now. Like, today.

NetApp

FIVE WARNING SIGNS
THAT YOU MAY HAVE A PUBLIC CLOUD SPENDING PROBLEM

You're still reading, so it's safe to assume that you have a sense that your organization is overspending on cloud.

If you have one or more of the following symptoms, a cloud spending problem is a safe diagnosis.

$$$

1 No one knows all the public cloud accounts the company has open

2 No one can fully explain the company's monthly cloud bills

3 There's no way to trace approved purchases to initiatives or cost centers

4 All capacity is purchased at on-demand prices

5 No one is regularly reviewing consumption or egress efficiency

NetApp

1

No one knows all the public cloud accounts the company has open

A typical enterprise maintains multiple accounts with multiple cloud providers. Because the public cloud is designed to empower anyone in an organization to set up an account and allocate resources, those with spending discretion often do so without any centralized notification of account creation. Shadow IT accounts pop up everywhere. Each one can be traced back to the lack of central control, and there never seems to be a single person who knows what accounts exist and what's in use. The typical result is a finger-pointing match rather than a productive discussion to help you place your finger on the real issue.

2

No one can fully explain the company's monthly cloud bills

Your organization's monthly cloud bill probably spans a dizzying number of pages that contain thousands of lines and reference many service names, instance types, and regions, none of which are necessarily self-explanatory. Sure, storage volume pricing may seem straightforward at first glance, but look at your latest bill and try to comprehend—let alone explain—all the additional charges and how they relate to specific business initiatives. It's often next to impossible.

NetApp

3

There's no way to trace approved purchases to initiatives or cost centers

Warning sign number 3 often goes back to the anonymous overspender. You have a blind spot when it comes to who is swiping the credit card—and what's behind these spending sprees. Building on the second warning sign, organizations with multiple teams, projects, and budgets often don't take the time to map their cloud storage and service consumption to projects and/or cost centers. The result is regular recurring spending that no one can explain.

4

All capacity is purchased at on-demand prices

The beauty of the cloud is that you buy whatever storage (and other cloud resources) you need... on demand. Although this pay-as-you-go method gives your organization the easiest means to curate additional storage, the reality is that more often than not, you don't need what you just purchased. That confounding mess of a bill we just mentioned? It shows that there's no managed procurement process and reveals a complete lack of blended purchasing.

Take the example of purchasing AWS On-Demand Instances. Even though On-Demand lets you spin up a VM at will, using AWS Reserved Instances instead gives you a discount of up to 75%. Taking it one step further, Spot by NetApp™ instances (where you pay for a spare Elastic Compute Cloud instance that meets your needs) can save you as much as 90% over On-Demand. See? If you're indiscriminately paying On-Demand prices, you can safely assume that there's no spending strategy or usage plan—just a bunch of random purchases made by individuals acting much like your friendly neighborhood shadow IT. They just expect the organization to pay the overages. And they're right, you will pay, but at what cost to the business?

NetApp

5

No one is regularly reviewing consumption or egress efficiency

Breaking news: Using data can be as costly as allocating it in the first place. There are two issues here. First up: How efficiently is your storage being used? You don't want to pay for completely unused storage, and you definitely don't want to use the most expensive tier of storage for data that could be stored on a less costly one. Costs can drop from $.20-$.30/GB all the way down to $.02/GB per month in "colder" storage tiers.

Egressing data from the cloud (or among clouds) is one of those pricing components that you probably fully comprehend in concept, but you can't translate into actual monthly cost. This ties back to aligning costs with initiatives—the forensics work. Someone needs to trace data that's egressed and determine whether the constant egress makes sense. For example, unsuspecting developers might be simply copying test data to their on-premises endpoint to speed up the testing of some code. Doing this regularly adds up to huge fees, and it's often completely avoidable. Someone needs to be on top of reviewing, questioning, and adjusting best practices around these charges.

NetApp

It's not difficult to find yourself with a cloud storage spending problem.

All organizations, regardless of size, can easily find themselves in these cloud storage quandaries. The overarching challenge is twofold.

First, everyone is using the cloud as if it exists exclusively to meet their needs, and no one is concerned about the resulting cost to the business.

Second, there's a visibility challenge; it's difficult to tie spending back to business needs and the user who's responsible. And without visibility comes a lack of accountability. The unfortunate result is that organizations like yours end up spending far more than they should.

But it doesn't have to be like that.

NetApp

10 THINGS
YOU CAN DO NOW TO CONTROL STORAGE COSTS IMMEDIATELY

The bad news is that you have a cloud storage spending problem. The good news is that the first (and often hardest) step is admitting that you have a problem.

And there are practical steps you can take today to begin fixing it.
The following 10 tips, although they're largely IT-centric, can help the owners of each moving part of cloud storage to review their usage and reduce their storage footprint, enabling the organization to rein in spending.

NetApp

Remove unattached cloud storage

When your VMs are terminated, normally only the root volume associated with the VM is automatically deleted. The additional storage volumes remain intact and incur storage costs—in some cases by design, to avoid accidental deletion. An easy way to reduce cloud costs is to find and delete unattached volumes. This one's a no-brainer—if you're not using the storage, get rid of it. But, as with any untethered resources, that involves some detective work to identify who owns the storage and have them attest to its need to exist or not.

Purchase the right storage tier

Every public cloud provider offers a number of storage tiers. And yet everyone seems to choose the fastest (and most expensive) level, with little or no regard for cost. If it costs more, it must be better, right? Not necessarily. The GB/month tier price is generally based on how often and how quickly you need to access your data. "Hot" storage tiers (which usually hold frequently accessed data that requires low latency, high performance and throughput, and high availability) can be priced as much as 5 times higher than their "cold" tier counterparts (where infrequently accessed data, such as backups and archives, should reside). Consider both your performance and cost requirements when assessing storage tiers, balancing your needs with attention to budgetary targets. You can always move the data to another tier later.

*"Hot" storage tiers (which usually hold frequently accessed data) can be priced as much as **5 times higher** than their "cold" tier counterparts.*

NetApp

*Mental math: By choosing the right storage tier based on capacity and throughput, and jumping to a higher storage tier only at peak usage times, you can **save up to 70%** on your monthly cloud bill.*

Rightsize underused storage volumes

The easiest way to waste money on cloud storage is to create a storage volume that isn't used appropriately. Cloud providers don't let you shrink storage volumes. NetApp recommends that you first go through the process of identifying oversized volumes, then create a new volume with the space you actually need, migrate the existing data, and delete the oversized volume. In the future, simply implement better storage need assessments for the volume creation process.

Downgrade storage based on throughput required

Cloud providers also offer performance tiers to meet your throughput needs. You can reduce storage costs by monitoring the actual read/write access of a given volume; if throughput is low, then downgrade it to a lower performance tier. This reduces the storage IOPS to be better aligned with the workload using it, while reducing cost.

NetApp

Determine the level of redundancy needed for storage

When people have the option to replicate data anywhere, they tend to get panicky and choose a far-away location. But do you really need your data in the United Kingdom to protect against a loss of data due to, say, a hurricane in the United States? The answer, of course, is no. These sorts of decisions have a material impact on cost; for example, redundancy across different geographies can be twice as expensive as local redundancy. It's important to plan your redundancy requirements wisely, taking a page from business impact analysis and risk assessments to determine what's actually needed.

Delete old snapshots

The snapshot is a staple of any VM recovery strategy. Multiple snapshots give IT organizations the ability to restore to a specific point in time, based on a specific disaster recovery scenario. The last thing you want is to delete something that the owner of the workload needs. But when you have hundreds of VMs, each creating snapshots daily without deleting the previous day's snapshot, your cloud storage costs grow exponentially. You must establish a strategy around snapshot expiration. Fortunately, most cloud providers have some form of snapshot lifecycle policy to automate deletions, eliminating the need to rely on any one individual.

*In the cloud, inbound traffic is normally free because providers benefit directly when they ingest your data. But when data is transferred outside the cloud provider's network (aka data egress), **it gets expensive fast.***

❚❚ NetApp

13

> *"Companies can **save hundreds of thousands of dollars a quarter, sometimes even $100,000 a month** simply by considering how they architect their cloud-based architecture."*
>
> *Spencer Hamons, NetApp Professional Services*

Manage outbound data transfer requests

Moving data costs money. That's a fact. But those costs are not all created equal. In the cloud, the cost of transferring data depends on the location of the source and the destination cloud servers. Inbound traffic is normally free (or close enough). But when data is transferred outside the cloud provider's network (aka data egress), it gets expensive fast. And remember, as far as the data owner is concerned, the transfer is all about "I need it done," not "it's cost effective to do it this way." To address this issue, encourage users to store data as close to where it's actually used as possible, to eliminate the need to move it somewhere else. Also consider compressing and deduplicating data before egress, using incremental synchronization to save on transfer costs. Finally, delete data or move it to an archive tier to save space and reduce cost.

Minimize cross-region and cross-zone data transfer

When data moves across regions, countries, or availability zones, cloud providers get to charge you more. These data transfers can be part of an application's architecture, used by DevOps to maintain test data, or they can be a part of a redundancy strategy. So data transfer needs to be both purposeful and sanctioned. Make it your goal to host necessary data as geographically close to its user base as possible. Over the long run, consider rearchitecting solutions to minimize the path your data needs to travel.

■ NetApp

Monitor storage pricing tier

Pricing for both storage and data transfers often includes additional cost tiers based on consumption. If you reach specified levels in the cloud provider's pricing table—as indicated by terms such as "greater than <storage amount>"—you're likely to be able to negotiate a better price. Keep in mind that higher discounts apply only to stored data that meets the pricing requirement. Finally, you may be locked into a multiyear contract. You should try to stay within the constraints of that contract, while looking for ways to negotiate costs down based on increased usage overtime.

Clean up incomplete uploads from storage

Some workloads require users to upload files. In this scenario, interrupted uploads sometimes result in partial objects that linger in cloud storage as unusable data that costs you real money. Depending on size, incomplete uploads can add up to tons of waste because admins tend to be afraid to delete or move anything (see step 6). The best course of action is to back up incomplete uploads and then delete them.

NetApp

15

NEXT STEPS TO ELIMINATING YOUR PUBLIC CLOUD SPENDING PROBLEM

Although these 10 steps can materially reduce cloud storage costs, organizations need a way to gain more visibility into what storage has been purchased, allocated, used, and wasted across the entire organization.

What's needed is a means to empower one or more members of the organization who are concerned with what's being spent and how well their cloud storage works.

Embracing reporting, calculators, and other tools offered by cloud providers is a solid first step toward visibility and centralized management of your cloud spending.

Once you've outgrown those basics, you'll need to look at third-party solutions designed to assess, collect, aggregate, analyze, and report on your cloud spending—including storage.

By first acknowledging that cloud spending problems exist and then taking immediate action with the 10 steps just described, you can start to gain control over your organization's cloud storage costs. To ensure that the savings are sustainable, policy, process, and additional technology hold the answers to reining in your organization's out-of-control cloud spending.

NetApp

MORE RESOURCES TO GET YOU STARTED

→ **Learn more at www.netapp.com**

Compare these benchmarks to your cloud storage as it is today.

→ aws

→ ▦ Microsoft Azure

→ Google Cloud

Check out these stories of customers who pioneered change at their companies with NetApp® cloud services:

→ CONA Services

→ Restaurant Magic

→ Optimize Cloud Performance

About NetApp

NetApp is the intelligent data infrastructure company, combining unified data storage, integrated data services, and CloudOps solutions to turn a world of disruption into opportunity for every customer. NetApp creates silo-free infrastructure, harnessing observability and AI to enable the industry's best data management. As the only enterprise-grade storage service natively embedded in the world's biggest clouds, our data storage delivers seamless flexibility. In addition, our data services create a data advantage through superior cyber resilience, governance, and application agility. Our CloudOps solutions provide continuous optimization of performance and efficiency through observability and AI. No matter the data type, workload, or environment, with NetApp you can transform your data infrastructure to realize your business possibilities. www.netapp.com

𝕏 in f ▶ 🎮   Contact Us

■ NetApp