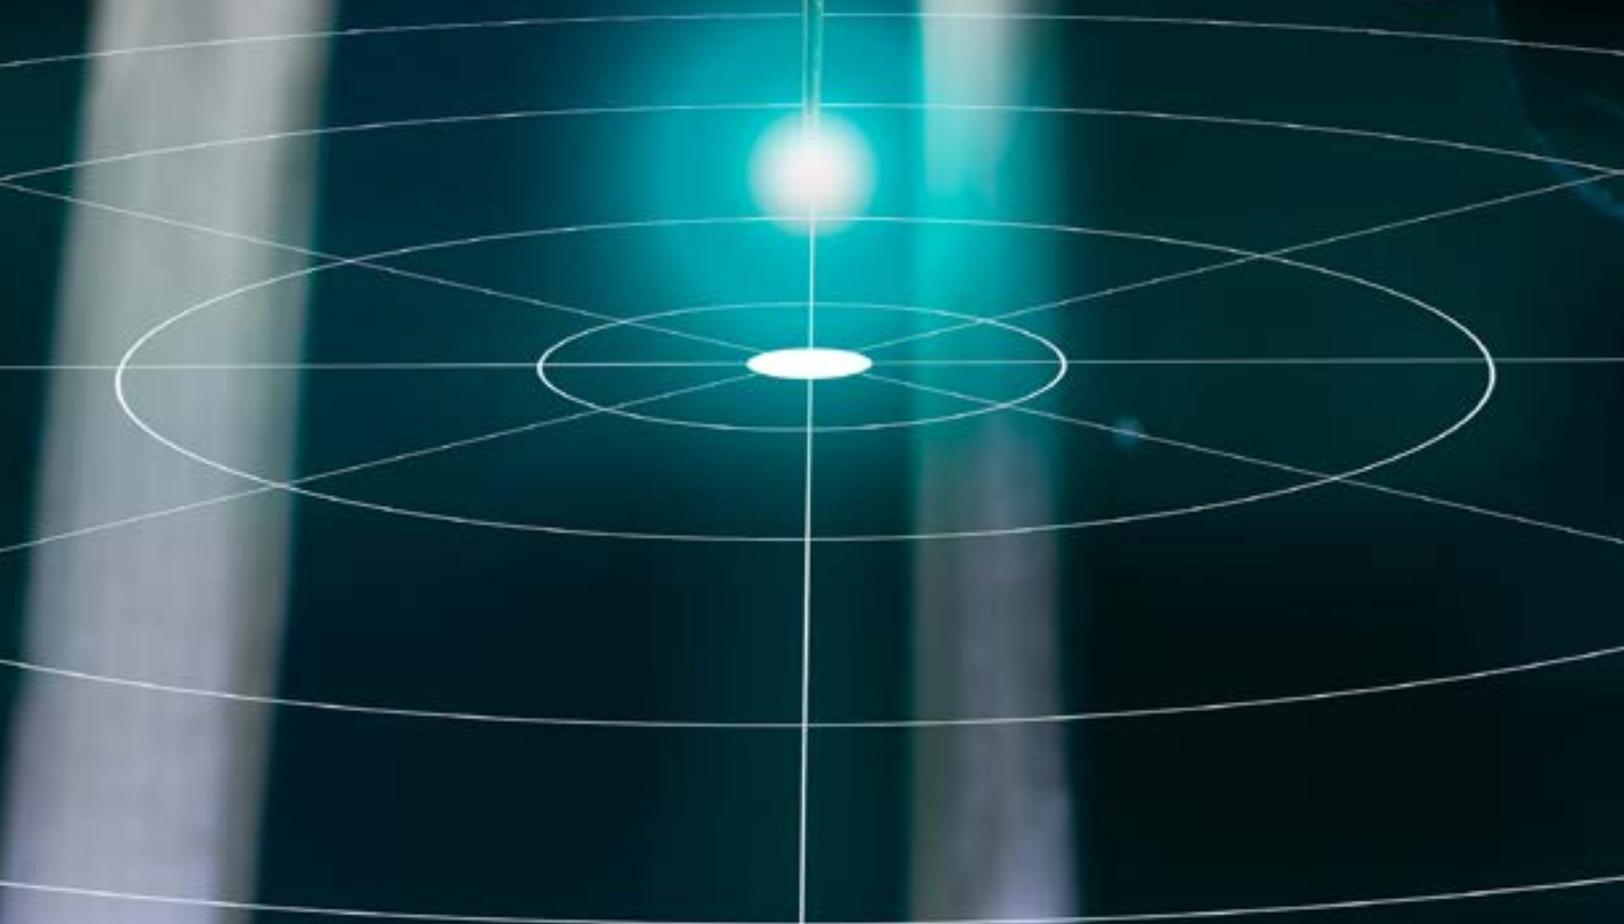
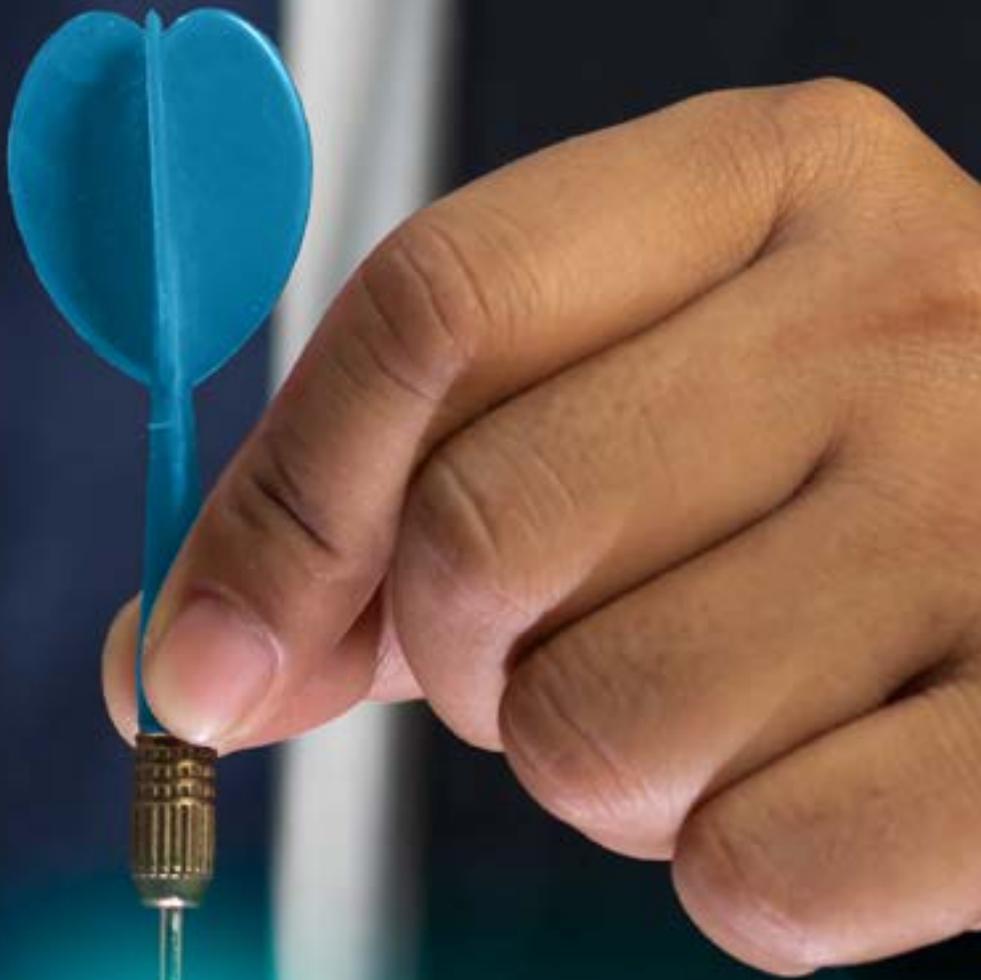




**A Holistic
Approach to
Cloud Cost
Optimization**





Introduction

Cloud cost optimization offers significant savings to customers engaged with the cloud, whether they are at the beginning of their cloud journey or experienced multi-cloud consumers.

As the current business environment places ever-increasing financial demands on organizations, cloud operating costs are under the microscope. Taking a holistic look at cloud cost optimization is a critical part of an organization's cloud operations.

The growing remote workforce and prevalence of shadow IT has led to greater levels of unmonitored cloud usage, leading to more wasted cloud spend. For organizations looking to ensure this doesn't happen, taking a holistic look at cloud cost optimization will help eliminate cloud waste, while enabling firms to streamline expenditure and enjoy significant savings.

It's important to note that cloud cost optimization is a journey—an ongoing journey that should be reviewed regularly. Cloud costs are, by nature, fluid, so regular monthly or quarterly reviews will enable cost optimization to remain relevant and accurate.

What exactly is cloud cost optimization? Cloud cost optimization offers businesses the opportunity to conduct a comprehensive audit of their cloud spend and find ways to maximize their investment through identifying and addressing inefficiencies.

This is achieved by looking at a number of optimization pillars to ensure the environment you've built in the cloud is running efficiently from both a performance and cost perspective. Key pillars to examine include auto-parking and on-demand services, cost models, rightsizing, instance family/VM type refresh, addressing waste and shadow IT. To accomplish this, it is critical to look at a cloud environment holistically and with the key objective of eliminating waste by identifying where resources may be mis-managed, underused, or redundant.

CLOUD COST OPTIMIZATION OFFERS BUSINESSES THE OPPORTUNITY TO CONDUCT A COMPREHENSIVE AUDIT OF THEIR CLOUD SPEND AND FIND WAYS TO MAXIMIZE THEIR INVESTMENT THROUGH IDENTIFYING AND ADDRESSING INEFFICIENCIES.



Why Do You Need Cloud Cost Optimization?

CLOUD COST OPTIMIZATION IS NOT A NEW PHENOMENON. BUT WHAT IS IT ABOUT PUBLIC CLOUD THAT MAKES COST OPTIMIZATION ESSENTIAL?

One key reason for its necessity is that public cloud vendors continuously release new services, features and products, and knowing how to navigate the sometimes-overwhelming sea of options is a skill in and of itself. Even with the mainstream cloud products that many customers use, such as AWS EC2 and Azure Virtual Machines, staying on top of cloud environments can prove complicated when new instance families or VM types are released.

Add to this the fact that customers themselves often make changes to their cloud environments—especially those with larger, more evolved environments—and the complexity is compounded. But with this complexity comes opportunity—opportunity for multiple areas of optimization, which is why organizations need to take a holistic look at their cloud cost optimization strategy.

This need for holistic optimization is only magnified with so many organizations now operating remote teams, leading to additional resources being built on cloud environments, such as Amazon WorkSpaces or Azure Windows Virtual Desktop—resources that drive additional costs and maintenance.



Cloud Optimization Areas of Examination

To approach cloud cost optimization holistically, you must first understand where the problems lie. What areas of your cloud environment are leaving you exposed to unused cloud services and wasted expenditure?

Identifying these areas will enable you to implement the key cost optimization pillars that can meet these challenges head on and offer significant savings on your cloud infrastructure, while enabling you to maintain a streamlined cloud environment.

WHAT AREAS OF YOUR CLOUD ENVIRONMENT ARE LEAVING YOU EXPOSED TO UNUSED CLOUD SERVICES AND WASTED EXPENDITURE?

Over-provisioned Resources

Over-provisioning of resources is one of the most prevalent issues found in cloud cost optimization, especially for organizations that are new to cloud. For those companies coming from a traditional hosting world with the experience and mindset of using physical servers, it can be a challenge to understand the scope of resources required when moving to the cloud. These companies tend to overestimate the resources they'll need in their cloud environment, basing their need off resource data from their physical datacenter.

For example, in a physical environment, organizations typically buy a set number of servers based on what they need to run not just today, but also for the following one to two years. This means having to pre-provision and buy all the relevant gear upfront. If these organizations followed the same mentality when provisioning for their cloud environments, they would size their requirements significantly higher than necessary, like utilizing EC2 instances that are greater than the applications require, considerably driving up their cloud bill. Instead, these organizations need to break out of their traditional thought process of the old hosting world and transition to building their cloud environment in accordance to only what is necessary to run their applications.

Storage

With a wide array of storage options in the public cloud, from block, file and object storage to hot and cold storage, assessing how often you'll need to access your data is a critical first step in selecting the right storage solution for your needs. There are benefits to each type of storage solution, but a lack of understanding or attention to the right storage service for your requirements can lead to significant waste.

Shadow IT

Larger companies are at particular risk of shadow IT impacting their cloud costs. They have numerous departments using cloud resources, and the people tasked with overseeing cloud usage and costing often don't know who is using what.

The organization may have a development environment or group building an application that those tasked with overseeing cloud expenditure know nothing about. Or Marketing may be investing in cloud options they don't fully disclose, understand, or need, and soon you start to build a shadow cloud-procurement environment that may be going unmonitored, unsupervised, and unsecured. This leads to an ever-increasing cloud spend problem that can be extremely challenging to manage without holistic cloud cost optimization measures in place.



Showback

The complexity of showback can be difficult to manage, especially for companies with high cloud usage and multiple accounts. If you're a large corporation with 200 internal business units that are all responsible for cloud, for example, allocating expenditure to the right department, project or group can seem an insurmountable task. If you're a large enterprise with monthly cloud bills upwards of 1 million, how do you allocate, control and reduce cost? How do you know who's billing what? Without an effective process with which the company can digest the bill and allocate spend accordingly, cloud cost optimization can be impacted heavily.

Lagging Technology Cultures

Organizations that have yet to embrace the latest technology cultures, such as a DevOps, are likely to have challenges when it comes to cloud cost optimization.

The forward-thinking focus that comes with a DevOps culture enables companies to adopt an application-first mentality. This means assessing application flow to see where parts of that process can be refactored to take advantage of enhanced

cloud services, such as containers, serverless, AI or data warehousing—things that are outside the traditional hosting customer that can offer cost savings and additional performance enhancements. Implementing new optimization practices such as Application Performance Monitoring (APM) gives companies visibility into their application layer and the ability to accelerate delivery and optimize their applications.

Not keeping up with this cultural change results in missed opportunities to accelerate operations and take advantage of technologies you might not necessarily have been comfortable with previously.

Maintaining proficiency with technology cultures like DevOps enables you to translate those efficiencies, like those that come from application refactoring, into the cloud cost optimization journey.



Cloud Cost Optimization Pillars

Where do you start in addressing these cost optimization challenges? We recommend examining the Pillars of Cloud Cost Optimization to prevent over-expenditure and wasted investment.

While a number of the pillars can be managed internally, you should also look to your vendors for discounts or rewards, generally offered based on commitments or prepays. These pillars will help you examine all facets of your cloud environment ensuring it's run efficiently and cost-effectively.

EXAMINE THE PILLARS OF CLOUD COST OPTIMIZATION TO PREVENT OVER-EXPENDITURE AND WASTED INVESTMENT.



Auto-parking and On-demand Services If you have a test, staging, or development environment your engineers work on during business hours, you will benefit from auto-parking and on-demand cloud services. Auto-parking, simply stated, is the ability to automate the powering on and off of your instance / VM based on your business requirements. You have to weigh the various pricing models—Savings Plans, RI Purchases, Spot, etc.—when choosing to auto-park resources because they run using on-demand rates, which is the most expensive option but gives the flexibility to stop those charges when the instance / VM is powered off. If you can turn your cloud off at 5pm on Monday through Friday and don't need it running over the weekend because it's not mission-critical, on-demand is likely the better choice for you.

With both auto-parking and on-demand services, there are still storage costs that need to be taken into consideration. Organizations must also be mindful that on-demand rates are the most expensive cloud rates you can sign up for, but the flexibility of being able to turn off services when not in use provides an excellent opportunity for cost optimization in businesses that simply don't need their cloud compute running 24/7.



Cost Models Cloud vendors offer a variety of cost models that can contribute to optimization efforts. For example, AWS' cost model, Savings Plans, released in November 2019, spans the breadth of instance families AWS offers and provides a discount for pre-paying and committing to resources over a one or three-year period.

A high-production environment with cloud compute that needs to be on consistently will benefit greatly from a cost model like Savings Plans. Alternately, running on-demand rates for a static environment of this kind will result in a significantly higher cloud bill than is necessary.

It's important to fully explore the different cost model options available to ensure you choose the most cost effective for your environment. There is a wide variety of choices, and not taking the time to understand which one is the most financially beneficial, specific to your environment, leaves you susceptible to ineffective cloud cost optimization.



Rightsizing Rightsizing offers ample opportunity for cloud cost optimization. If you're new to cloud, it can be difficult to determine currently used resources and map those to the most cost-effective solution in the cloud.

New cloud users often don't know exactly how much horsepower they'll need to run their applications in their new cloud environment. Rightsizing can address this dilemma and contribute significantly to your optimization strategy. And once you've completed your cloud migration, it's easy to scale resources up or down, change instance sizes, add more memory, add more compute or reduce it, and so on.

Whether you need to increase or decrease usage, rightsizing is an easy and effective way to get precisely the right amount of resources for your needs.



Instance Family/VM Type Refresh Cloud vendors offer a variety of cost models that can contribute to optimization efforts. For example, AWS' cost model, Savings Plans, released in November 2019, spans the breadth of instance families AWS offers and provides a discount for pre-paying and committing to resources over a one or three-year period.

A high-production environment with cloud compute that needs to be on consistently will benefit greatly from a cost model like Savings Plans. Alternately, running on-demand rates for a static environment of this kind will result in a significantly higher cloud bill than is necessary.

It's important to fully explore the different cost model options available to ensure you choose the most cost effective for your environment. There is a wide variety of choices, and not taking the time to understand which one is the most financially beneficial, specific to your environment, leaves you susceptible to ineffective cloud cost optimization.



Addressing Waste One of the simplest ways to address cost optimization is to address wasted resources and storage. Especially with larger customers, there are likely elements of your cloud infrastructure that aren't being used or that you don't even realize you have.

There can be load balancers that are idle or old snapshots that are sitting around, for example, and the vendor will charge you whether you use them to capacity or not. If your snapshot policy is to keep data for 30 days, and you have data being stored for 60-90 days, that's unnecessary and wasted costs.

Wasted storage can impact cloud cost optimization significantly. Customers must make sure they are using the right type of storage for their needs to achieve true cost optimization. Factors such as how frequently data is being accessed should dictate the type of storage used. For example, files that are archived and not pulled frequently can be moved to cold storage such as AWS Glacier, where the costs are significantly cheaper than in typical block storage.



Shadow IT As companies with multiple business units continue to grow their public cloud usage, it becomes challenging to keep track of who has cloud access and who is spending what. A lot of times, CIOs will receive cloud invoices and have no idea who is building what in their cloud account. All they see is the ever-increasing cost affecting their bottom line.

This usage may or may not be malicious in nature, and taking preventative measures such as tagging, applying proper IAM policy/permissions, sub-account invoicing / payback, and performing regular optimization reviews can ensure you catch these problems before they become a burden.

Conclusion

Managing cloud cost optimization holistically is possible with a commitment to ongoing examination covering the key pillars of cost optimization. For organizations that incorporate cloud cost optimization into their cloud infrastructure management, significant savings can be found, especially in larger organizations with considerable cloud spend.

To maximize its effectiveness and simplicity, cloud cost optimization should be implemented at the beginning of a cloud journey. However, it's never too late to introduce and implement an optimization strategy.

Businesses need to look at all their different data points and all their cloud variables to understand where and how they can save money. Taking a holistic view of cloud cost optimization will enable you to see significant cost savings. The time and money invested in doing so, whether in-house or with your service provider, will be returned three or four-fold when the relevant optimizations are established and acted upon.

