

CLOUD COST OPTIMIZATION: FOUR STEPS TO SUCCESS



Adopting cloud computing provides a number of challenges for organizations, but managing cloud spend may be one of the most demanding. [A survey of over 750 enterprises revealed that over 30% considered their cloud spend wasteful](#), and 80% stated that they found managing their cloud spend challenging.

Most organizations move their system to the cloud to reduce costs. A common mistake that many enterprises make when moving to the cloud is trying to follow the fastest path. They simply lift and shift the applications they were using from their in-house data center to a [cloud infrastructure](#).

Most of the time, they find that this has not made these applications cheaper, scalable, faster, or more flexible. In fact, many enterprises are surprised to find that their costs go up instead of down. In many cases, the users within the organization may not even realize that the company invested a large amount of time and resources to move operations to the cloud because everything works predominantly the same.

Cloud cost optimization is a new discipline that organizations need to deploy cloud computing effectively. Here is some guidance to help you create a comprehensive and standardized optimization process to improve cloud cost.

(This article is part of our [IT Cost Management Guide](#). Use the right-hand menu to navigate.)

Challenges to cloud optimization

When it comes to lowering the cost of cloud spend, many organizations run into significant issues such as:

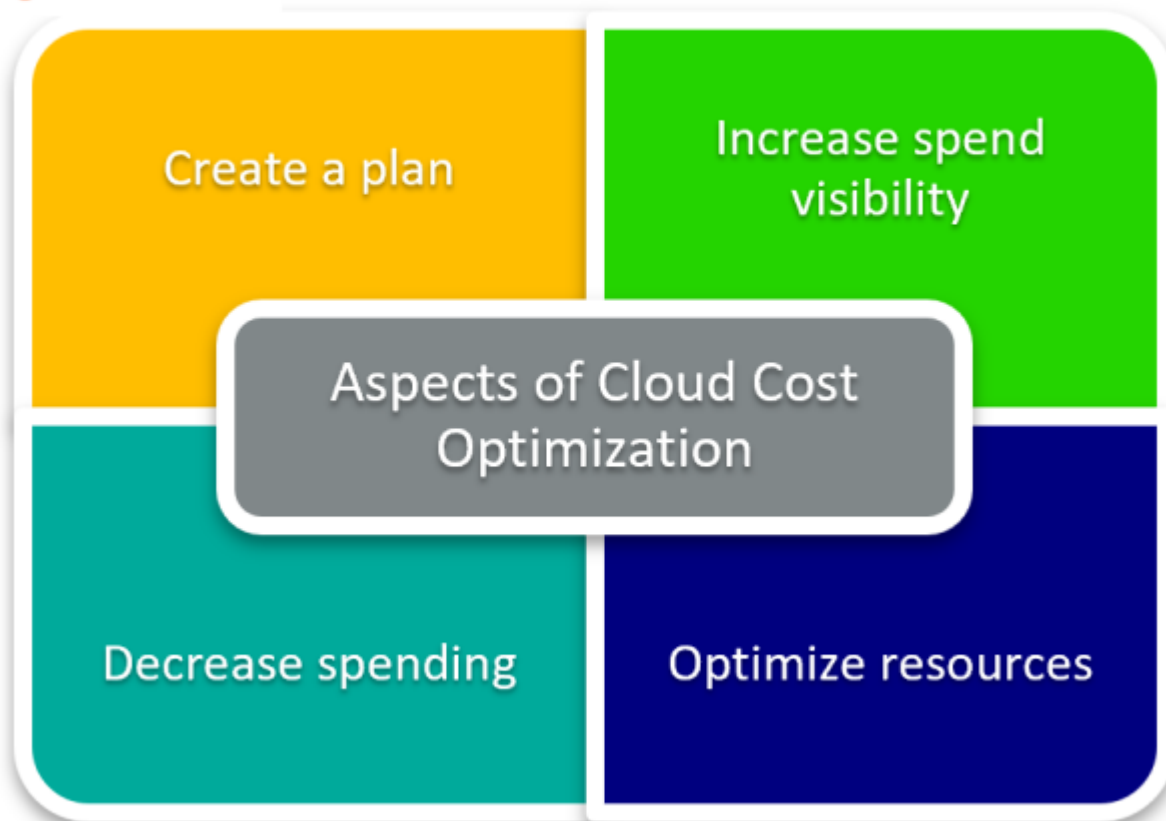
- Complicated and multifaceted pricing structures
- Extreme cloud bill detail
- Ease of cloud service provisioning
- Continuous modifications to cloud offerings
- Excessive alternative architectures
- Lack of coherence across cloud platforms

Many times, advice on cost management provides a list of tasks like turning off unused instances and deleting unused storage. However, these practices do not give a comprehensive view of cost management. In fact, some of these tasks (like turning off unused instances) can even lead to disruption, frustration, and [shadow IT](#).

Instead, it's critical to create a strategic approach to cloud cost management. A set expectation is vital to keep spending within a certain budget. Plus, tracking and organizing costs around applications and centers will provide insights that cause cloud consumers to take a proactive approach to their spending.

Organizations need to develop sound financial management processes to prevent cloud overspending and create more efficient cloud service consumption.

Developing these processes will affect multiple departments and roles within the organization, such as the Cloud Center of Excellence (CCoE), I&O, finance, and users of the cloud services. These processes will translate into concrete management requirements and involve adopting new tools.



Four aspects of cloud cost optimization

Managing and optimizing cloud costs requires a multifaceted and comprehensive approach. Here are four steps to overseeing and improving cloud spending.

Create a plan

Random tactics will not guarantee that spending remains within expectations if no one knows what the expectations are. Organizations need to forecast their consumption and create budget expectations to stay within that forecast.

Organizations need to develop this capability and run processes before deploying applications, workloads, and projects in the cloud.

CCoE needs to define requirements to identify the exact outcomes that impact cloud services design and keep from overengineering applications. They need to gain a comprehensive understanding, question assumptions and clarify what each application is accomplishing.

This step entails collaborating with product owners and stakeholders to gain an understanding of the value each application offers the organization. They also determine the key metrics that require the utilization of specific architectural principles.

Ask questions to help determine workload requirements:

- How [sensitive is the data](#) the workload handles?
- What happens if the data remains unavailable for a certain amount of time or is lost altogether? Does this application require service-level objectives?
- What is the desired performance target?
- Does the data require compliance with any industry-specific regulations, such as [HIPAA](#)?

This step also required choosing pricing models that best suit your organization. Model prices will vary depending on factors such as [data integrity](#), [service availability](#), embedded license-based software, and performance targets. It's critical to find which pricing model meets your budget and your organization's needs.

Increase spend visibility

When you've established your budget, landed on your pricing model, and deployed the application, maintaining visibility is critical. Create an organized view of costs outside of the bill itself. The amount of data makes manually managing each line item most likely unsuccessful. Plus, it won't allow you to get a daily view of spending, which is critical for optimization.

Define which [metrics are critical](#) to track for your organization. These metrics are crucial to building dashboards and reports, as well as maintaining automation workflows to optimize spending. Some metrics to track include:

- Cost of services
- Capacity
- Utilization
- Availability
- Performance

Every major cloud provider utilizes tags (or labels) as a fundamental governance construct. These tags are a critical tool for cost tracking. They allow customizable names, multilateral structures and can be implemented across multiple clouds.

While tags will appear in bills after implementation, they do not appear on previous bills, so it's critical to implement tags as quickly as possible to enable cost tracking.

Building dashboards and reports and updating them daily are critical for staying on top of spend. These reports include:

- Top and least spenders
- Trending daily, monthly, quarterly, and annual patterns
- Actual versus planned spending
- Overall spending
- Estimated spending waste

Decrease spending

Once you've set goals and enabled visibility, you can better identify ways to reduce waste. Because you've completed the first two steps, reducing waste won't require changing application code or architecture. Plus, it's also easier to implement and calculate ROI with tracking already in place.

Some ways to reduce spending include:

Discard unused resources. Although it may seem like an obvious measure, it is an uncommon practice for many traditional data centers. Find any resources that your organization has deployed but does not use. In particular, find allocation-based resources that accrue cost regardless of their usage. Some examples include old snapshots, unused storage volumes, idle compute instances and unassigned IP addresses. To ensure that it is truly unused, mark them first and then notify owners to solicit an action from them.

Schedule services. It is common for organizations to have resources idle at specific hours or days. Reduce waste by scheduling cloud services based on these expected patterns. Do this by describing it with a "duty schedule" tag. Any cron-like scheduler can then read that tag value and schedule services correspondingly.

Optimize allocation-based services. These types of services require users to request particular allocation when provisioning. Instead, adjust the size of your allocation to fit the actual workload demand to reduce costs. Container, compute, storage, database, and application services are all good places to start.

Utilize discounts. Pay-as-you-go (PAYG) models don't work for every organization. If your workload is relatively stable and your utilization in the future is predictable, you might be able to get discounted prices. Some cloud providers offer discounts that you can purchase programmatically. You can also contact your sales rep to negotiate a deal by committing to a minimum spend.

Upgrade instance generation. [Cloud providers](#) have refreshed their computer platforms over the years to provide renewed power to certain use cases. These instances are also often less expensive because they are more efficient than previous generations. As a result, you will want to optimize across instance families with these updates. You may be able to opt for a new instance generation and get the same performance with a smaller size.

Optimize resources

Reducing spending is only part of the strategic techniques to lower cloud costs. Organizations also need to optimize their processes to reduce their need for certain resources.

Find preemptible instances. Particular cloud providers offer a much lower price point for compute instances instead of the PAYG model. However, preemptible are riskier in that providers can terminate them at any time based on demand. Identify components and use cases in your application's architecture that may be suitable in the case of suddenly unavailable infrastructure. Batch workloads, for example, can simply be put on pause if infrastructure becomes unavailable and resume when it's available again.

Adjust data storage. Not all data is equally important, and some require less access as it ages. For example, social media data becomes less important as time goes on. Optimize costs by storing old or less important data in a less expensive service or tier.

(Explore [data storage types](#) & [storage temperatures](#).)

Utilize serverless. [Serverless computing](#) requires organizations to give up ownership of their application infrastructure to a third-party cloud provider. However, they get to experience zero-touch autoscaling, dynamic deployment, and enhanced efficiencies in resource utilization.

While serverless computing services can seem cost-effective, there is a certain point where it does become cost-prohibitive and offers a diminishing return. As a result, organizations should aim to leverage serverless technology for the right use cases.

Utilize horizontal autoscaling. Autoscaling, or enabling applications to increase or decrease in response to events, can significantly optimize costs. However, autoscaling is either "vertical," making one instance bigger, or "horizontal," adding more instances of a similar type and distributing them across.

Both forms of autoscaling can reduce costs and have their place in optimizing. Horizontal autoscaling needs certain design principles built into the application architecture. Applications must be allowed to run multiple instances in parallel. Plus, it needs to be to start and shutdown seamlessly and not rely on local dependencies. It's an effective practice that makes the application more resilient and must be used alongside optimizing allocation-based services because they apply to different sets of applications.

While the initial process requires following each step, this process is iterative. Lowering cloud costs is not a one-step process and does not always require following the steps in order. Instead, it requires going back over each aspect of optimization to find areas to reduce spending.

Best practices for cloud cost optimization

Let's review some tips for how to optimize cloud cost.

Identify stakeholders

Optimization is not limited to your finance team. Successfully optimizing costs requires everyone's buy-in and should be spread across your organization. Some stakeholders to consider include:

- Managers

- Product owners
- DevOps
- Finance department
- CCoE

Implement accountability

Keeping costs down is a part of everyone's job. The IT and financial departments have long had separate goals. However, IT needs to be accountable for its budget. Chargeback models that calculate and charge based on unit costs can help increase accountability. It can also make costs a part of every department.

Implementing a fair chargeback model can increase awareness and help lower costs.

Utilize tools

The right tools can help your organizations gain visibility and critical insights to increase optimizations. Some tools to consider include:

- [Dashboard visualization](#)
- Heat maps
- [Cloud cost intelligence](#)
- Container cost optimization

Summing up cloud costs

As IT and cloud operations teams start to work on optimizing their cloud implementations, many find that the efforts required are more complicated than they anticipated. By applying the principles above, you can start to control spending, establish realistic cloud operations budgets, and ultimately reduce costs and waste.

Related reading

- [BMC Multi-Cloud Blog](#)
- [How to Accurately Estimate the Cost of Cloud Migrations](#)
- [How Multi-Cloud Can Enable Cost Savings](#)
- [Cloud Compliance: Best Practices for Success](#)