

Getting Cloud Smart: 7 Steps to Start Managing Multi-Cloud Expenses

Christina Kidd
Technology Evangelist

Keyvan Shirnia
Head of Managed Services & Product Management

5 July 2019



At Fusion Global Business Solutions, we help companies gain measurable business outcomes from their investments in IT Service Management and Operations.

In this white paper, we discuss how companies can implement seven easy, affordable steps to begin managing costs in the multi-cloud. Our conversations and experiences with our partners indicate that establishing affordable, logical governance for the cloud not only promotes immediate cost savings, but begins your process towards cloud governance and cloud maturity.

Table of Contents

Introduction	2
The Multi-Cloud: Benefits and Growth	4
The Problem with the Multi-Cloud	5
Reducing Cloud Expense, Optimising Cost, and Moving Towards Cloud Maturity	6
7 Steps for Multi-Cloud Cost Reduction	7
Step 1. Design and Adhere to a Tagging Plan	8
Step 2. Rightsize your cloud resources	9
Step 3. Choose the appropriate pricing model	10
Step 4. Reclaim orphaned resources	11
Step 5. Throttle poorly utilised resources	12
Step 6. Scrutinise and limit data egress	13
Step 7. Use the free tools	14
Next Steps	15
Resources	16



The Multi-Cloud: Benefits and Growth

In 2019, companies rely on the public cloud for a variety of reasons. These reasons include storing data, testing and developing new apps, utilizing infrastructure-as-a-service (IaaS) and software-as-a-service (SAAS), and experimenting with affordable computing services like machine learning and data analytics. The cloud is rightfully touted for its affordability, flexibility, and scalability – allowing 21st-century companies to keep up with the incessant need to bring well-developed products to market faster than ever.

Companies are embracing the cloud. According to RightScale's 2019 State of the Cloud, which surveys close to 800 global firms, enterprises plan to spend 24 percent more on public cloud use in 2019 compared to just last year. These companies are spending significantly, too: half will spend more than \$1.2 million USD on the public cloud, with 13 percent spending more than \$12 million.

“Public cloud” includes any number of cloud computing options, from industry heavyweights like Amazon Web Services, Microsoft Azure, and Google Cloud Platform (which, combined, comprise 60 percent of the public cloud) to more specific, smaller cloud vendors.

Companies who utilize public cloud resources are often doing so across several vendors – a strategy known as the multi-cloud. Enterprises are utilizing on average nearly five clouds, comprised of 3.5 public and private clouds in regular use and another 1.5 clouds experimentally.

While companies are utilizing the cloud, few are achieving “cloud smart”. This is a level of maturity and experience with public cloud infrastructure wherein companies harness cloud benefits to create a best, custom package that maximizes value at the minimum cost.

“Cloud smart” assumes companies are already in the cloud. Instead, it asks whether they are truly harnessing the benefits of the cloud by using it as astutely as possible.

The Problem with the Multi-Cloud

Enterprises often possess on-premise legacy systems that have been customized over decades to suit their company and enhance their competitive edge. For most enterprises, however, on-prem systems lack the massive scale and computing density of public cloud offerings, such as SaaS, automation, and machine learning.

These cloud services are considered the best way to stay efficient and deliver products to consumers as quickly as the market now requires. No longer must you budget capital expenses over five years, or require a minimum timeframe for each system in order to achieve ROI.

Unfortunately, companies continue to spend heavily on the cloud, particularly in early migration stages, and they struggle to achieve their anticipated total cost of ownership (TCO) calculations.

Although cloud maturity – harnessing the cloud efficiently and affordably – is inevitable, with peak spending estimated to taper off in the next several years, it is heavy expenditure in the cloud that's currently preventing most companies from maturing as quickly as they'd like. In fact, this year, nearly two-thirds of companies report cloud cost savings as their top initiative.

Though the benefits of the cloud are real – affordability, flexibility, and speed – they can be tricky to master. Companies used to legacy systems but migrating to cloud

don't understand inherent fundamentals, resulting in unnecessary spending. Old ways of thinking, like shoring up for an anticipated production peak, no longer apply.

Of course, best practices and governance are essential, but they aren't enough. As the cloud functions in new ways, companies must seek and embrace innovative working methods that take full advantage of the cloud. Companies are routinely missing opportunities to save money, due to unidentified workloads and underutilized resources. Cloud users estimate their wasted resources at 27 percent of cloud cost, but real data indicates that cloud waste accounts for 35 percent of cloud spending on average.

The public cloud is partly to blame, too: pricing plans, incentives, and discounts vary wildly across vendors, and transparency isn't common. Industry experts warn against vendor lock-in. Sticking with a single vendor for all of your cloud needs can lead to increased costs over time or difficult transitions across clouds and vendors as needs change.

This mix of enterprise lack of knowledge, poor framework and practices, and pricing opacity means that cloud pricing is understandably confusing. Add in the many vendors you may use in a multi-cloud situation, and cloud costs are an escalating and significant concern.

Reducing Cloud Expense, Optimising Cost, and Moving Towards Cloud Maturity

Cost is the problem with the multi-cloud, but reducing cost can feel like a futile exercise. Fortunately, many cost-related pain points can be mitigated with straightforward solutions.

The common denominator across organisations we partner with at Fusion GBS is that they don't know what they have migrated to the cloud. If they don't know this information, then no one can track it. For all the discussion around hidden costs versus visible costs, employee skills gaps, and the differences between legacy systems and the cloud, the issue that is hardest to tackle is systemic: your processes and your people.

Companies believe that reducing cost is the goal. And it is, at least for a time. But once companies begin implementing cost-saving steps, the issue becomes how to optimize cost, not simply reduce it. It is this sophistication in the cloud that marks a company's cloud maturity.

Instead of being approached as a one-time objective, cost optimization must be a continual process that focuses on both refining and improving a service over its entire life cycle. The goals of cost-optimization include:

- Meeting functional requirements
- Fully utilising all resources
- Achieving all desired and necessary outcomes at the lowest possible price

As you move closer to cloud maturity, you should recognise that reducing expenses is only one of four pillars of cost-related cloud maturity. You'll also want to focus on aligning supply and demand, maintaining expenditure awareness over time, and optimising cost as a regular occurrence.





7 Steps for Multi-Cloud Cost Reduction

In 2016, Gartner developed seven steps that companies can implement in order to reduce public Infrastructure-as-a-Service (IaaS) multi-cloud expenses. These steps focus on obtaining visibility and affordability more easily. They also serve as a baseline to develop and deploy best practices and governance, turning problems with the cloud into assets.

At Fusion Global Business Solutions, we have embraced these seven steps as we help companies reduce their multi-cloud costs. But, with emerging technologies like Platform-as-a-Service, containers and serverless options such as Google BigQuery and Amazon Relational Database Service (RDS), we have updated these seven steps to reflect the multi-cloud world of 2019. In each section, we've included a Fusion Pro-Tip that moves the theory of the step towards good cloud governance.

While a variety of paid tools and services purport to solve this problem, remember that tools and services are only part of any technology puzzle – you'll rely much more heavily on the right people and the right processes in order to successfully adopt any enterprise-wide initiative. Companies whose goal is to reduce costs can begin by using these seven steps.

FUSION TIPS:

Before implementing any step, first compare your company's current performance against that recommended step. Consider a broad, rudimentary scoring system, such as: not started, immature, mature, and comprehensive. Then, start to envision ways your company may enact the step.

Step 1.

Design and Adhere to a Tagging Plan

Companies are often unaware of what they have in the cloud, stopping any cost-optimization efforts before they can begin. With a comprehensive, easy-to-implement strategy of tagging, you can know exactly what's in the cloud at any point, calling up cross-referenced cloud resources as needed.

The practice of attaching user-defined keywords to cloud resources, known as tagging, must be applied consistently throughout the company. Depending on the complexity and breadth of your cloud resources, tagging activities may be more specific or holistic.

Best practices for tagging include:

- **Employ a variety of attributes** across domains of interest. Think cross-functionally, so you can include tags related to department/team owner, purpose (development, testing, producing, etc.), the related app or service, compliance requirements, cost centre, deadline to shut down the workloads, security, automation, etc.
- **Link tagging to capacity management functionality.** This ensures that the tagging, and the following six steps, are reported accurately for utilisation, services, and other conditions.
- **Require tagging across the entire company.** Encourage this practice by enacting policies such as prohibiting or shutting down untagged resources in the cloud.
- **Run frequent reports.** Tagging makes it easy to see who is using cloud resources and how. Running frequent reports – on a daily or even hourly basis – can quickly identify overspending.

FUSION TIPS:

Draft a tagging plan before enacting one. Consider all the parties that may have resources in the cloud and how they may tag items differently from your team. You will see what areas are over- and under-defined by tagging.

Step 2.

Rightsize your cloud resources

Rightsizing your cloud resources is a tricky endeavour, and companies tend to err on one of two sides when it comes to spending. Cost-conscious companies may under resource their cloud, risking resource starvation, insufficient security, and general performance degradation - all of which result in poor customer experience. On the other side, companies who over-purchase or allocate additional resources to prepare for unlikely heavy loads are wasting resources. This traditional way of managing assets no longer applies in the age of public cloud agility and scalability.

Embracing a new cloud mindset can be difficult. In many situations, sizing a workload for the cloud is almost identical to sizing a virtual machine in your traditional data centres. Build the right resource profile, comprised of CPUs, amount of RAM, and network bandwidth, and align with service-level expectations for performance, reliability, and availability.

But not everything is identical. In fact, the rapid penetration of containers and serverless options, services that don't scale to the cloud in the same way, is completely changing how developers work. Rightsizing containers and serverless options means understanding how these resources relate to each other, which containers communicate, and which cloud is the best fit for each resource.

To rightsize your resources:

- **Understand all components and dependences** involved in each specific service or application and gather baseline resource consumption usages per service. Consider both on-prem systems as well as containers and serverless options.
- **Prioritise application migration.** Not all applications should be moved at the same time. App and service priority can consider criteria such as:
 - Consumption
 - Operating services
 - Criticality to business service
 - Deployed environs
- **Migrate application to the cloud based on actual (not estimated) usage.** Actual usage prevents both under- and overestimations.
- **Continue monitoring resource consumption regularly.** Resource sizing is not a one-time activity. Regular, frequent reviews of your cloud workloads, which are easier to track due to comprehensive tagging, are essential to maximising value. Review workloads on a weekly, daily, or hourly basis, depending on the size and volatility of those workloads.

FUSION TIPS:

Take the long view: rightsize your resources now and frequently, particularly for large and critical services. This practice improves accuracy over time, ensuring maximum benefit when you eventually adapt additional best practices, automation, or cost management services.

Step 3.

Choose the appropriate pricing model

At first glance, the public cloud's pay-as-you-go model seems beneficial: more flexibility, on-demand scaling, and little to no planning or prep work required. But it's rarely that simple.

Cloud vendors typically charge per computation, pennies on pennies. These add up quickly, particularly when you are unaware of all your cloud resources. To complicate matters further, cloud vendors have seemingly endless pricing options and tiers. Depending on your enterprise's cloud presence, you may negotiate your own specific pricing. All this is to say that choosing the appropriate pricing model is very complex.

Fortunately, both large and small public clouds offer incentive pricing models or discounts with built-in capacity planning tools and price calculators. These tools analyse spending patterns to predict future workloads. However, do be aware that the tools are typically specific to each cloud silo, making it difficult to compare costs across multiple providers.

So, with myriad pricing and incentives, how do you choose a cloud provider? Research and test these price calculators and incentive pricing models on small deployments first. The proper use of these tools offers built-in cost savings, but the tools aren't easy to master. In fact, the casual, scattered use of these tools can result in unnecessary spending.

FUSION TIPS:

Two things change often: cloud pricing schedules and your workloads. Frequent, regular rebalancing is necessary, but the more systematic this rebalancing becomes, the easier it is to execute.



Step 4.

Reclaim orphaned resources

Orphaned resources comprise a significant area of waste. The public cloud makes it easy to create many cloud resources. For example, you may stand up a new server and storage. Months later, your developer is finished with that work and relinquishes (shuts down and deletes) the server.

Unfortunately, as many companies learn too late, your storage stays in the cloud until you specifically release it. The result? Orphaned resources that continue to compute and, therefore, incur costs.

Evaluating these underutilized resources helps companies determine whether the orphaned resources may still serve value if reallocated, or whether they have become redundant, in which case best practice will be to shut them down.

How to reclaim orphaned resources: Scan your resources often (on a weekly or daily basis) to identify unallocated computing and storage resources. Specify utilisation targets to define orphaned resources (e.g., those running under 10 percent). Once identified, shut down or reallocate these instances to reduce computing costs.

Long-term best practices can include manual and automated tools that help scan and remove resources, but a well-implemented tagging process is more than sufficient.

FUSION TIPS:

Map the server and all allocated resources. When a server is relinquished, all associated resources, such as storage, are also released.

Step 5. Throttle poorly utilised resources

Next, shift focus to resources that are poorly utilized. Many resources are used only at specific times, such as virtual machines for remote desktops or testing and developing activities. When not computing, these resources are idle, but they may not do well when shut down entirely. Employees using these workloads still expect these works-in-progress to persist and be available until their next computation.

As most cloud computing services are virtual, it's easy to "pause" a resource without entirely shutting it down. A paused resource saves its state, but its virtual machine is turned off, saving computational costs. This practice, known as parking, results in significant savings when consistently applied: pausing a virtual machine outside of a typical eight-hour workday can result in cost savings of up to 75 percent.

To throttle poorly utilized resources, start by running reports to identify them. Again, define what "poorly utilized" means to your company. Perhaps it includes any services or applications that aren't needed 24 hours a day. Next, develop an action plan on how to treat these resources – it is likely you'll need several action plans to be implemented for the different types of resources at play. Implement these action plans, but measure their efficacy by tracking the ensuing costs.

FUSION TIPS:

Require new resources to the cloud to have a use time identified in the change records and their tags.



Step 6.

Scrutinise and limit data egress

After remedying major cloud waste, move into areas with less significant waste. Applications that are particularly lively, with data flowing in and out 24 hours a day, may be doing so without your knowledge. For example, companies are often surprised at the countless sources of data egress, often unnecessary or redundant, such as a web page improperly using HTTP caching directives, or workloads overcommunicating with other external clouds.

Limiting data egress may seem fundamentally easy, but most companies have issues tracking data flow because it requires visibility into several essential areas: how applications connect and communicate, the direction and amount of data flowing to and from each application, and whether data packets traverse cloud zones or exit the cloud altogether.

FUSION TIPS:

Map all data dependencies, both ingress and egress – automation is essential for this process. Measure the data flow (ingress and egress) from each application. Next, overlap the data flow volumes with the mapped dependencies, removing unnecessary data transfer points.



Step 7. Use the free tools

Now that your multi-cloud has a method for cost optimization, your company is perfectly situated to take advantage of free built-in tools that are native to most cloud vendors.

These tools can take on tedious tasks like cost planning, consumption, and monitoring. They may seem primitive – emailed alerts based on historical data or reasonable projected values – but they're the first line in defence of your cost-optimization processes. Hourly or daily dashboard insights provide real-time insights into your cloud resources and the related costs.

However, free tools have one major drawback: their visibility is limited to the specific vendor. Companies in the multi-cloud will have to rely on a variety of tools across all vendors. With free tools in place for each vendor, you may want to explore whether additional, fee-based services can support your cost optimization, particularly in providing holistic, multi-cloud visibility.

FUSION TIPS:

Not all tools are the same. Experiment with combining native and add-on tools for the exact dashboard you need, at a reasonable price.



Next Steps

These seven steps are the starting point. Once they have been put into place, both holistically and consistently, your company will begin reaping the benefits of significant cost reduction. But cloud maturity isn't inevitable, and as your company's knowledge of the cloud matures, you can take on longer-term opportunities to maximize your cloud value. In no particular order, here are additional best practices to consider adopting:

- **Automate.**

In small workloads typical of early cloud adaptation, manual processes are plenty. As you scale, automation becomes necessary to ensure consistent, successful application of best practices.

- Automation also drives data analytics – take advantage of machine learning capabilities in your cloud, which can optimize cost recommendations over time.
- Automation frees up employees to innovate. The less time employees have to spend on manual, repetitive processes, the more they can focus on innovative ways to deploy the cloud.
- Weigh the benefits of automating part or all of each of the seven steps outlined above.

- **Consider cost management and/or hybrid cloud solutions.**

Once you have your baseline cost management processes in place, additional solutions may solve unique problems.

- Standalone cloud cost monitoring and optimization (CCMO) tools are more cost effective than hybrid tools, with many reaching ROI within six months.
- Hybrid cloud management (HCM) tools go beyond cost optimization by adding solutions to complex issues like compliance, orchestration, and added levels of automation. These are pricier, so be sure the cost is worth the value.

- **Refactor apps for cloud-native operations.**

Consider switching your standard software and on-prem operations for IaaS, SaaS, and PaaS offerings that are cloud native. These investments mean you won't be trying to make a legacy system work with the cloud.

- **Explore alternates to the public multi-cloud.**

The public cloud is only one solution. Just as you had legacy, on-prem systems, a private cloud may be worth the investment. These are complex, time-consuming tasks, but the level of customisation a private cloud can provide may be just what your company needs.

As companies attain cloud intelligence, one thing is clear: success in the cloud, whether public or private, depends more on the maturity of your cloud management and practical governance strategies than on the nature of your cloud workloads. Technology is one tool among many, and it is never a solution on its own. The right technology can support the right people and processes, resulting in a continuous cycle of planning, tracking, reducing, optimising, and maturing. Cloud maturity, then, is linked directly to how wisely you use your cloud in support of your people, your product, and your processes.



Contact:

www.fusiongbs.com

London: +44 208 814 4888
New York: +1 844 456 1342
Madrid: +34 91 790 1106

Enquiries@fusiongbs.com

Resources:

- [2019 State of the Cloud Report](#)
- [10 Tips for Managing Cloud Costs](#)
- [Amazon whitepaper: Cost-Optimization Pillar: AWS Well-Architected Framework](#)
- [BMC Cloud Cost Optimization Strategies and Tools](#)
- [Fusion GBS: "7 steps for managing multi-cloud expenses"](#)
- [Gartner: Fund Innovation through Smarter IT Spending](#)
- [Gartner: Advance Cloud Computing Capabilities](#)
- [Gartner whitepaper: 7 Steps to Reducing Public Cloud IaaS Expense](#)
- [RightScale 2019 State of the Cloud Report from Flexera Identifies Cloud Adoption Trends](#)