


WHITE PAPER

You Need True Multi-Cloud Cost Management. Here's Why.



The multi-cloud computing environment offers a variety of benefits for businesses and developers alike. Empowering teams to access diverse cloud capabilities and avoid vendor lock-in while taking advantage of competing services, a multi-cloud environment is an effective way to create more flexible and cost-effective software.

However, the multi-cloud environment isn't without its complicating factors. Formulating effective cost management solutions can be challenging, even in single-cloud settings. The additional variables multi-cloud architecture introduces presents novel challenges for organizations concerned with accurate cost allocation.

An effective strategy for multi-cloud cost management extends the same concepts for managing costs in a single-cloud architecture. Incorporating various cloud providers' environment variables, technical details, and terminology into a holistic cost management strategy enables businesses to extend single-cloud cost management's tried-and-true tactics to multi-cloud environments. Cost management has challenged many businesses as they adopt cloud architecture. As multi-cloud architectures become increasingly common, cloud resource costs become substantially more difficult to pin down. It quickly becomes nearly impossible for finance teams to understand the return on investment from their cloud budget without accurate data on actual resource costs.

Yotascale provides an effective means of automating cost management at scale across both single- and multi-cloud deployments. With a robust, granular cost allocation and resource tracking system, Yotascale helps eliminate the constant second-guessing that accompanies cloud accounting.



Defining Effective Cloud Cost Management

The actual cost of a cloud-based product can be challenging to determine. Though all the data to answer expense questions may be available, tracking the price of specific features and services becomes complicated without a clear organizational plan in place.

Tracking the cost of an enterprise's cloud products back to the teams working on them requires a cloud architecture that mirrors its organizational structure. Services operating in the cloud should be structured and labeled, so it's immediately apparent which team is responsible for individual features, containers, and cloud resources.

This organizational structure mirroring allows organizations to "chargeback" individual teams for using cloud resources. Making the teams responsible for using the underlying resources aware of their expenses encourages developers and engineers to make cost-effective decisions.

Understanding resource use across distributed environments enables engineers to make informed decisions about sharing, pooling, and more efficiently using those resources. Analyzing this information allows teams to identify cost-reducing vectors and novel methods of minimizing their cloud resource use.

Creating a labeling and tagging system directly mirroring your internal organization is a crucial bit of groundwork that allows cost management strategies to scale freely. By defining a descriptive approach to tags and labels, you can rapidly correlate resource use back to the source, no matter your organization's size.

Creating Granular Cost Visibility

Cost management is a question of granularity. How specifically can an enterprise track individual cloud resource costs? With automated tools like Yotascale, analysis becomes highly granular using systematic descriptive labeling.

Yotascale's advanced tagging features form the basis of the cost attribution strategy necessary to manage expenses in single- and multi-cloud environments. With the ability to create rules for automatically tagging services dynamically, and machine learning identifying untagged or inaccurately labeled resources, teams can zoom in on microscopic views of their resource spend.



When teams can account for every bit of resource use, it becomes possible to budget with this high level of granularity in mind. Budgeting can happen on the level of individual features rather than entire services or departments.

Similarly, cost predictability increases as resource tracking becomes more granular. Companies can forecast budgets based on intricate details about individual feature costs and refine them to incorporate expenses that would otherwise hide beneath complicated accounting.

Defining your cost allocation strategy's granularity is a critical factor in cost management. With Yotascale, you can leverage dynamic automatic tagging that quickly answers vital questions, such as:



What feature or service is consuming resources?



Which team is responsible for those resources at that particular time?



How much does the resource use cost for this specific feature, service, cluster, or container?

Cost attribution strategies like resource tagging make important resource use metrics available for both human and machine analysis.

Expanding to Multi-Cloud Infrastructure

Competing cloud technologies have created a rich ecosystem of developer tools. As applications become increasingly compartmentalized into microservices, deployments become more widely distributed across various clouds. Cloud architectures that leverage multiple competing platforms are common and offer a wealth of benefits to businesses and developers alike.

The fear of vendor lock-in has become real as modern software continues to become heavily dependent on cloud providers. One strategy for avoiding dependence on any specific cloud platform has been to adopt multi-cloud deployments that keep businesses actively using multiple competing public cloud platforms alongside private clouds.

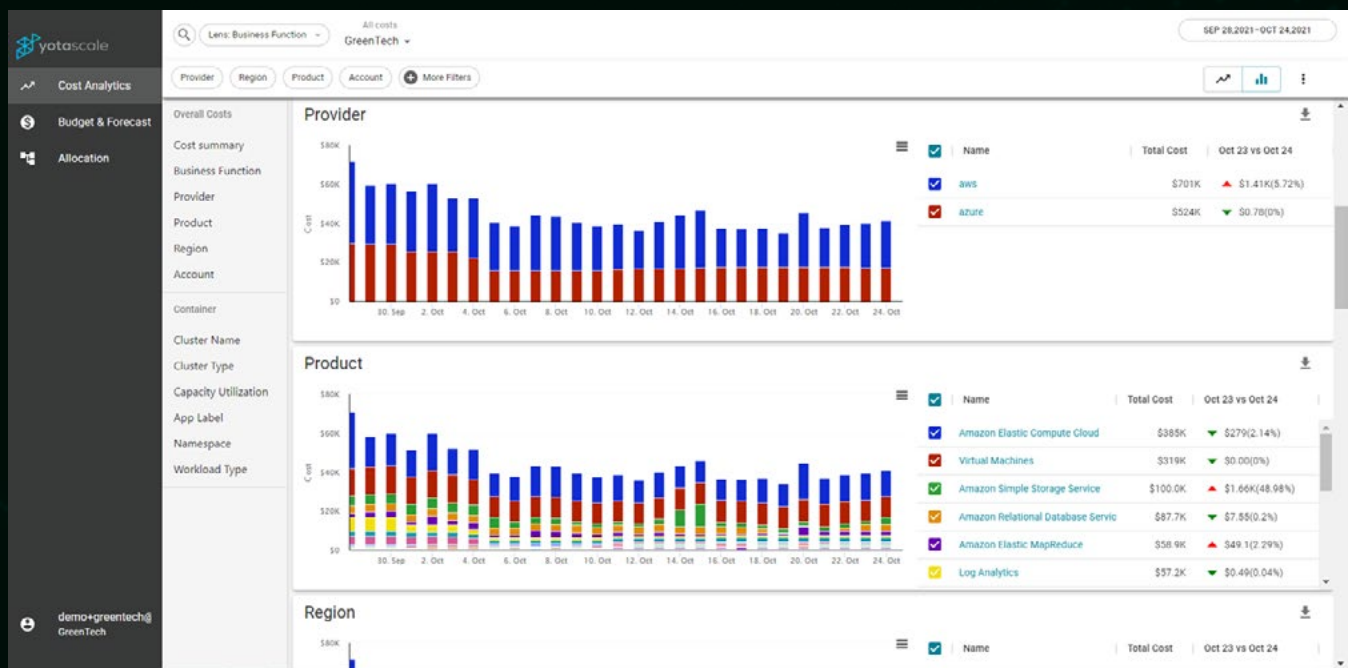
Even for businesses unconcerned with vendor lock-in, competitive pricing across different cloud providers can reduce the overall resource cost of application deployments when appropriately leveraged.

Companies often transition to a multi-cloud environment to access a compelling feature one platform offers that's unavailable elsewhere. Multi-cloud architectures enable developers to incorporate new technologies without necessitating radical infrastructure shifts.

Companies can extend their single-cloud cost management strategy to the multi-cloud environment without too much trouble. Successful cost allocation still depends on descriptive service tagging and defining a granularity level that mirrors your organizational structure.

While tagging and cost allocation principles remain the same, multi-cloud architecture introduces a few challenges. In particular, cloud platforms each use unique terminology and have their methods of presenting data. Combining data from disparate sources into a single analysis also creates an added layer of complexity when aggregating multi-cloud costs into a single report.

You can handle all these challenges with a bit of planning. Yotyscale helps businesses tackle their multi-cloud costs every bit as effectively as their single-cloud costs.



Yotyscale presents a unified view of all your multi-cloud costs, in the context of how you do business.

Accounting for Dynamic Infrastructure

Modern applications based on containers scale dynamically in response to demand. Tools such as Kubernetes orchestrate provisioning resources in real-time, responding to changes in resource use by allowing containerized applications to adjust resource availability as needed. Dynamic scaling capabilities are great for keeping services functioning. However, they come with fundamental questions about the cost of shared resources — and which teams are responsible for those resources at a particular moment.

Kubernetes provides a suite of tools for tracking resource use in real-time. But without a strong cost allocation structure correlating shared resources use back to individual teams and applications, translating the Kubernetes' raw resource monitoring into an effective cost-analysis is a real struggle. With the proper cost allocation framework in place, teams can quickly trace the source of shared resource use — even in dynamically scaling environments.


Determining the source of cost anomalies and collecting information about the specific shared clusters, containers, and developers responsible for these cost anomalies enables businesses to respond to cost spikes and rapidly address their root causes.

Yotascale provides right-sizing suggestions that minimize resource over-allocation across Kubernetes clusters through diligent tagging and automated recommendations. By carefully analyzing the precise use requirements of individual features and services, engineers have the opportunity to dial the efficiency of every cluster to its maximum potential, and finance teams have the chance to correlate resource use to cost rapidly.

Properly tagging and tracking the dynamically scaling Kubernetes clusters requires a wealth of context. Which team is responsible for a given container can change depending on the container's development life cycle position. Companies can provision or decommission clusters automatically in response to load or on the fly by developers spinning up new features.

Tracking the context of these dynamic Kubernetes states is much easier with a system of automated tagging. Apply automated tags when provisioning clusters and explore detailed labels to help track their ownership, purpose, and overall resource consumption throughout the cluster's lifetime.

Because it offers the capability to tag Kubernetes resource use automatically, Yotascale plays an instrumental role in creating a cost management strategy that integrates Kubernetes deployments as part of a multi-cloud architecture. Yotascale empowers teams to track their Kubernetes spend using the same tools available for monitoring public cloud costs.



Analyzing Multi-Cloud Costs

One of the most significant challenges in managing multi-cloud costs is the ongoing need to aggregate data from multiple disparate sources. With the correct data available, traditional cost management methods like those applied to single-cloud deployments can be effective in multi-cloud architectures. But this has proven challenging at scale, requiring employees to merge and analyze data sourced from competing services.

Yotascale helps businesses consolidate their multi-cloud costs into a single interface. Extending Yotascale's automated tagging and detailed labeling to every cloud in a deployment creates granular cost allocation that leaves no stone unturned, no matter the multi-cloud environment's complexity. Yotascale automates tagging and labeling through a system of user-defined rules and compiles data into detailed visualizations capable of tracking costs throughout every cloud in a deployment.

Yotascale's holistic view of a multi-cloud environment provides machine learning insights for right-sizing and termination recommendations to keep clouds running lean. It also centralizes anomaly

and cost spike reporting, providing rapid response capabilities through real-time monitoring and automated notifications to relevant team members.

Yotascale dramatically simplifies cost management for multi-cloud environments by exploring resource costs across your entire multi-cloud deployment from a single interface and drilling down into detailed reportings correlated to your organizational structure with just a few clicks.

Platform-Specific Nuances

Competing cloud providers offer similar services. Functionally, they're often essentially identical, accomplishing more or less the same tasks through slightly different procedures. But each platform has its specific nuances, terminology, and communication methods.

A modern multi-cloud deployment may consist of any number of public clouds operating in concert. Each cloud platform or provider likely has its own glossary of terms relating to its functions and services. These terms are often similar enough to seem interchangeable, yet specificity becomes extremely important when working in a multi-cloud environment.

Tagging for cost allocation needs to incorporate and acknowledge terminology variations. In Yotascale, that means configuring automated labeling to correspond to the environment's terminology. Using generic labels quickly confuses subtle differences in language across clouds, making it difficult to pinpoint the precise meaning of a generic term.

Differences also present themselves in data. The exact methods, metrics, and scales cloud solutions record for resource consumption can vary. Cloud providers may use different systems for collecting this information and various means of reporting it.

Rather than attempting to abstract these differences away, incorporate them into your cost allocation strategy's granularity through descriptive labels. Acknowledging the peculiarities of each cloud solution operating in a multi-cloud architecture simplifies the process of tracing costs back to their source.

Embrace Multi-Cloud with Yotascale

Yotascale offers a rich system of visualization and observability. These presentations make it simple for everyone, from engineering, to management, to finance to quickly visualize how resource use correlates to costs and extract detailed insights about individual teams and services.

Complementing these visibility tools is Yotascale's state-of-the-art machine learning. Yotascale's automated tools help teams stay on top of their expenses. They provide actionable insights for improving efficiency by automatically detecting anomalies, such as cost spikes, and leveraging predictive analytics to forecast upcoming budgets.

Combined with a well-defined cost attribution strategy diligently tagging your cloud resources, Yotascale targets cost-reduction and right-sizing recommendations at individual services and teams. It also offers budgeting alerts for teams, products, and applications, creating real-time feedback loops for detecting and eliminating cost spikes.

As it diligently monitors the load across your multi-cloud services, Yotascale provides automated right-sizing recommendations based on actual usage and performance data. Right-sizing enables teams to define the most efficient operating parameters for their cloud services, minimizing wasted resources and unnecessary over-allocations.



Yotascale understands that many different discrete teams and individual employees need access to cost data, so it automates generating reports relevant to specific groups. Employees receive rich visualizations tailored to their particular interests, such as the tags relevant to their team's projects.

Putting all teams on the same page is an ongoing challenge in cost management. Without a consistent dataset to reference, engineering and finance may find themselves in disagreement over resource use or actual feature costs. With a robust interface, Yotascale's visualizations are equally helpful for developers, engineers, managers, and financial analysts, providing a single point of truth for everyone collaborating on the project.

Whether you're looking to extend your existing cost allocation strategy to multiple clouds or the idea of cost management is new for your business, Yotascale has the capabilities to demystify your cloud product costs rapidly.

**Ready to discover the actual
cost of your cloud resources?**
Try a free demo of Yotascale's dynamic
cloud cost management capabilities.

