



Cloud Cost Visibility through Resource Tagging Strategies and Taxonomies

Venkata Sasidhar (Sasi) Kanumuri

Email id: sasinrrt@gmail.com

Abstract Implementing a well-designed tagging strategy is a cornerstone of effective cloud cost management and successful FinOps practices. Tags act as key-value pairs, enabling granular cost tracking, budgeting, and chargeback/showback models. This article provides a comprehensive framework for developing and executing a tagging strategy that aligns with your organization's unique needs. We cover the fundamentals of tagging, the importance of establishing a clear and consistent taxonomy, and best practices for tag creation. The recommended tag categories span business-level attributes (project names, owners), infrastructure details (environment, service names), Kubernetes-specific labels, and security classifications. The article emphasizes the power of automation using Infrastructure-as-Code (IaC) tools like Terraform. IaC ensures tagging consistency, streamlines management, and helps prevent untagged resources from slipping through the cracks. Guidance on tag enforcement mechanisms is provided to maintain data integrity and reliable cost reporting. Finally, the article offers practical methods for discovering untagged resources, employing a mix of native cloud provider tools, APIs, and third-party solutions for optimal visibility.

Keywords: Tagging, tag hygiene, tag health, cloud resource tagging, taxonomy, tagging strategy, cloud cost visibility through tagging, resource allocation, cloud spend attribution, enforce tagging.

Introduction

In essence, tags are unique labels you apply to your cloud services. They function as key-value pairs, such as "Project: Website Launch" and "Department: Marketing". This labeling system offers an effective way to arrange and analyze your cloud environment.

Why is this important for cost containment, then? Thanks to the comprehensive insights that tags unlock, You can see how much you spend on different departments, initiatives, or individual applications. This provides actionable data and removes the cloud-charging "black box" impression. You may easily identify locations where resources may be over-provisioned, locate resources that have been forgotten but are still costing you money, or find out whether any development environments are wasting money on downtime.

Tags also encourage responsible spending and frugal budgeting. Teams are more inclined to choose resources wisely when they see the expenses directly related to cloud usage. Within each team, this visibility may also result in more precise budgeting and forecasting.

Tags may and should be applied to all cloud resources you provide to support your enterprise, not only processing and storage, as seen in Figure 1 below. Identifying organizational cost drivers enables drill-down costs and granular breakdowns in any combination of variables.



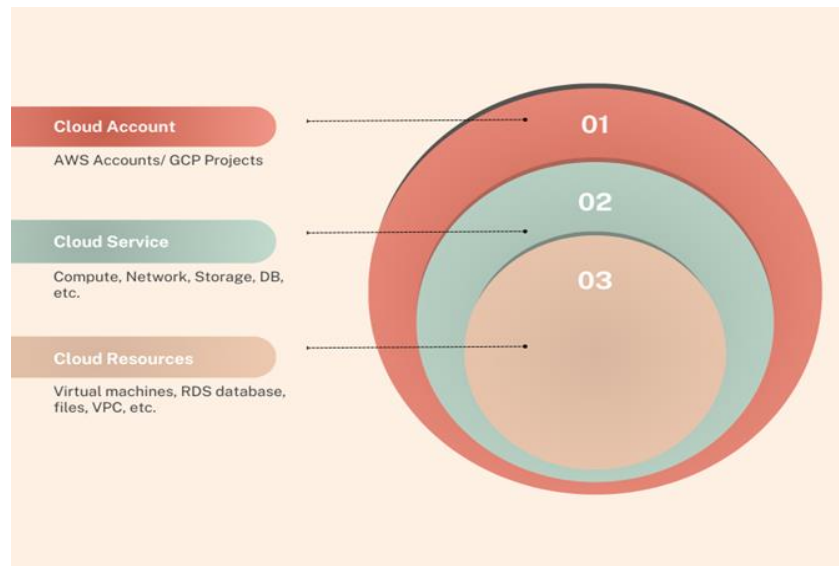


Figure 1: Cost Allocation via Tagging

A. Taxonomy

Consider cloud cost taxonomy as a highly specialized cloud expense filing system. It outlines the labels and categories to arrange expenses consistent with your company objectives. This methodical approach is essential for gaining insightful information and improving cost management; it's not simply about tidying.

Regarding cloud spending, a well-designed taxonomy guarantees that everyone in your company speaks the same language. This consistency allows you to: Organize according to cost center, environment, application, or any other criterion.

Tag Name	Description	Tag Key and Values
Cost Center	Identifies the department or team responsible for a resource.	cost_center <ul style="list-style-type: none"> Dept A/ C
Environment	Indicates whether the resource is for development, staging, or production.	env <ul style="list-style-type: none"> Development/ Staging
Project	Associate the resource with a specific project or initiative.	project <ul style="list-style-type: none"> Project A/ Project B
Service	Associate the resource with a particular application or service.	service_name <ul style="list-style-type: none"> App A/ App C
Owner	Identifies the individual or team responsible for the resource.	service_owner <ul style="list-style-type: none"> Username 1/ Team Name

- [1]. **Track spending patterns:** Recognize patterns across time, identify irregularities, and comprehend how prices vary with business activities.
- [2]. **Pinpoint waste:** Focus on places where resources might be left running unintentionally, overprovisioned, or underutilized.
- [3]. **Improve forecasting:** Use past cost data grouped by your taxonomy for more accurate budget estimates.
- [4]. **Drive accountability:** When teams or projects see their specific cloud costs, they will likely make responsible usage decisions.

While a simple taxonomy can provide a strong foundation for managing cloud costs, the true potential lies in customizing it to reflect your business's unique needs. Think of your taxonomy as an adaptable tool rather than a rigid set of rules. As your organization grows and changes, so should how you categorize your cloud spending.

B. Tag Health

Tagging is the bedrock for understanding your cloud expenses, but it's not a set-and-forget solution. Think of tag health as maintaining a well-organized garden. Left unchecked, tags can become inconsistent, outdated, or even misleading. This messy tagging can lead to inaccurate cost reporting, making it challenging to spot cost optimization opportunities or hold teams accountable for their cloud usage.

C. What is Tag Health?

Think of tag health as a measure of how well your cloud tagging system is performing. A healthy tagging environment means your tags are complete, accurate, consistent, and follow a clear set of standards. Why does this matter? Because reliable and organized tagging directly translates into trustworthy cost data and smarter cloud management.

D. Key Elements of Tag Health:

- [1]. **Completeness:** Every cloud resource should have the appropriate tags attached. Missing tags create blind spots in your cost reporting, making understanding where your money is going challenging.
- [2]. **Accuracy:** Tag values must reflect the resource's purpose, ownership, or other attributes. Inaccurate tags muddy the waters of your cost analysis and mislead decision-making.
- [3]. **Consistency:** It is crucial to use the same tags in the same way across your entire cloud environment. This ensures everyone speaks the same language regarding costs, enabling comparisons and reliable data aggregation.
- [4]. **Standardization:** A well-defined taxonomy serves as the rulebook for your tagging strategy. It outlines the specific tags you'll use, their valid values, and how they should be applied. This structure prevents confusion and maintains the integrity of your data over time.

E. Why is Tag Health Important?

Poor tag health can lead to several issues:

- [1]. **Erroneous Cost Allocation:** Without clean and accurate tags, it's like deciphering a mixed-up bill at a restaurant. Teams might be charged for resources they don't use, or costs could disappear into an untraceable black hole. This undermines accountability and makes it impossible to pinpoint areas for cost savings.
- [2]. **Wasted Resources:** Imagine finding lost items in a messy attic – that's the challenge of identifying forgotten assets without good tagging. Unused development environments, idle storage, or oversized instances might all be quietly draining your budget, hidden amongst the cloud clutter.
- [3]. **Ineffective Reporting:** Tagging chaos is the enemy of meaningful analysis. Trying to glean insights from inconsistent data is like assembling a puzzle without all the pieces. This hinders your ability to track trends, compare costs across teams, or justify cloud spending to stakeholders.
- [4]. **Compliance Challenges:** Many industries have regulations regarding data storage and access. Inaccurate, missing, or inconsistent tagging can make it impossible to prove that you're meeting security standards or protecting sensitive information, leading to compliance fines and reputational damage.

F. Strategies for Maintaining Tag Health:

- [1]. **Define a Simple and Clear Tagging Taxonomy:** Consider your taxonomy a neatly arranged set of drawers. Start with the essentials – cost center, environment, project – and keep the labels clear. Avoid complex jargon or overly granular tags that only a few people understand. A streamlined and intuitive taxonomy encourages team-wide adoption and sets the stage for accurate reporting.
- [2]. **Case Sensitivity:** Remember, cloud systems see "TeamA" and "teama" as entirely different tags. Enforce consistent capitalization to prevent data from being split incorrectly, which creates confusion and inaccurate cost breakdowns.



- [3]. **Automate Tagging:** Remove the human factor wherever possible. Tools like AWS CloudFormation or Terraform allow you to define tagging rules automatically applied as resources are created. This minimizes the risk of typos, missing tags, and inconsistencies that can creep in with manual input.
- [4]. **Remember Tags:** Make tagging an essential part of your cloud resource lifecycle. Whether spinning up a new virtual machine or deploying a storage bucket, build tagging into your provisioning processes. Untagged resources are invisible to cost analysis, making it difficult to pinpoint where your money is truly going.
- [5]. **Tagging Enforcement:** Don't just recommend good tagging practices; make them mandatory! Use policies that prevent resources from being created without the required tags. This drives consistency and can act as a safety net, catching potential miss configurations and preventing untraceable costs from surfacing down the line.
- [6]. **Tag Governance:** Consider this as appointing 'librarians' for your cloud cost data. A designated team or individual should be responsible for defining your tagging rules, updating the taxonomy, and educating others on best practices. This ensures your tags don't become a free-for-all, preventing inconsistencies that can erode the value of your cost insights.
- [7]. **Review and Update:** Your cloud environment isn't static, nor should your tagging be! As new projects emerge, teams reorganize, or technology changes, you should introduce new tag categories or redefine existing ones. Schedule regular reviews to ensure your taxonomy meets your evolving business needs. This proactive approach prevents your tagging from becoming outdated and hindering accurate analysis.

Tagging Strategy for Finops at Scale

Businesses that wish to establish and grow their FinOps operations require a good tagging strategy. To effectively manage your company spending, you must have a tagging system that accurately categorizes and allocates all cloud resources you deploy. However, even the most brilliant taxonomy will only succeed if tagging is managed and enforced consistently across your teams.

A tagging technique can considerably increase financial visibility while optimizing resource allocation and cost analysis. Defining, sharing, reviewing, automating, and enforcing the plan is critical to developing a successful tag strategy. Each component contributes significantly to the strategy's effectiveness and efficiency.

- [1]. **Define:** A well-structured tagging taxonomy is the cornerstone of an effective strategy. Think of this step as creating a custom labeling system for your cloud resources. Work with key stakeholders across the organization to develop a set of tags that capture the most important cost-related information. While every business differs, we'll explore some common tag examples shortly.
- [2]. **Communicate:** Once you have a solid plan, it's time to spread the word! Document your tagging strategy and explain the meaning behind each tag. Make this documentation easily available across the company, particularly to teams involved in cloud resource creation (developers, IT admins, finance). Training sessions or open Q&A forums can be great ways to reinforce the importance of tagging and clear up any confusion.
- [3]. **Review:** Your tagging strategy should be flexible. Schedule regular reviews to assess how well your current tags are working. Collect feedback from finance, IT, and other teams that rely on cost data. Analyze trends in your cost reports – can you easily pinpoint where spending is increasing or decreasing? These insights help you identify tags that need to be added, modified, or retired.
- [4]. **Automate:** Embrace the power of automation to minimize errors and make tagging a seamless part of your cloud workflows. Your cloud provider likely has native tools for automating tag applications, and Infrastructure-as-Code (IaC) tools like Terraform allow you to embed tagging directly into your resource definitions. Look for third-party solutions to help monitor and correct tagging mistakes across your cloud environment.
- [5]. **Enforce:** Establish clear governance around your tagging standards. Use policies that prevent untagged resources from being created, or at least trigger alerts so they can be addressed quickly. Consider real-time notifications to help teams correct tagging oversights on the spot. This promotes accountability and keeps your cost data clean and reliable.





Figure 2: Tagging Strategy

Tagging can be very tedious, especially without supporting automation, and it requires coordination with cross-functional partners to ensure accuracy, which may be why your teams haven't tackled this yet. However, the gains to your business in terms of cost visibility, which will drive hard dollar COST SAVINGS for your business, justify the effort involved in creating your tagging strategy. The tagging strategy you create has to be well-defined and communicated for tagging to be fruitful.

Recommended Tags and Taxonomy

The taxonomy you develop for your FinOps strategy will be specific to your organization's requirements. It will allow you to link costs to teams and provide detailed cost information across several categories or organizations. When creating your tagging taxonomy, think about how you want your expenses to be organized for reporting and drilling down. As shown in Figure 3, construct your cloud cost hierarchy according to the organizational reporting structure, with applications assigned to teams (shared services in blue), teams reporting to business units, and so on.

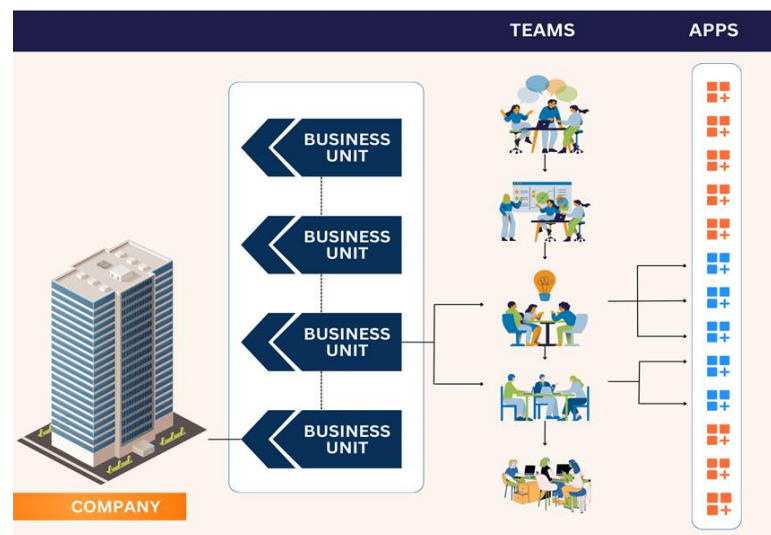


Figure 3: Cloud Costs in a Business Unit (BU) Hierarchy

While a business unit (BU) cost hierarchy is standard, consider alternative ways to structure your cloud cost tracking for maximum impact. Consider categorizing spending based on infrastructure (e.g., Development vs. Production) to gain insights into resource usage patterns. Alternatively, an executive leadership (VP) hierarchy

aligns costs directly with your management structure. This offers simplified cost allocation, greater executive spending transparency, and a heightened sense of team ownership and accountability. The optimal approach for your organization might involve a combination of these hierarchies, providing the flexibility to tailor your cost tracking to match your internal structure and reporting needs.

Whichever hierarchy (or hierarchies—you can use more than one) you choose will help you identify the tags required for your taxonomy. The tagging guidelines below will include the tag name and whether it is required for your taxonomy and the tag's function. We'll sort the suggestions by

A. Business Level Tags

Business-level tags can aid in cost attribution at the account or project level (AWS, Azure, GCP) and serve as the foundation for budgeting, forecasting, and the development of chargeback/show-back models

Tag Name	Req'd?	Purpose
account_name project_name	Yes	Name the cloud account or project, using the naming conventions of the organization
account_owner	Yes	Name of the engineering team that owns the account. A team alias works best in this scenario.
point_of_contact	Yes	Username or email of the primary point of contact (POC) for the account
executive_sponsor	Yes	Represent spend from an executive perspective. Helps to promote top-down budget alignment for cloud spend.
cost_center	Yes	Identifies the cost center that the resource belongs to
cob (cost_of_business) cogs (cost_of_goods)	Yes	Indicates whether it's direct production costs vs. R&D (non-prod) with a TRUE/FALSE value. Use the term that matches your finance reporting standard.

B. Infrastructure Level Tags

A cloud account or project allows one or more teams to collaborate on services and applications. The following tags provide greater granularity for cost attribution across several teams.

Note: When deploying cloud resources, it is recommended to use the environment (Dev, QA, Staging, Prod, etc.) to identify resources in certain environments.

Tag Name	Req'd?	Purpose
service_name	Yes	Identifies the name of the service the resource belongs to. Example: Front-end
service_owner	Yes	Name of the engineering team that owns the service. A team alias works best in this scenario. If the resource is a shared service, add multiple owners separated by “,”.
shared_service	Conditional	Takes Boolean values (yes or no as a value). If the value is yes, we'll have to add all the service owners under the service_owner tag separated by “,”
managed_by	Yes	Team alias/IC email of the service's team.
point_of_contact	Yes	Username or email of the primary point of contact (POC) for the



		service
requestor	Conditional	The name of the team that requests the service; might be required in creating an account
env	Yes	Dev, QA, Stage, Prod, Sandbox, etc. If the infrastructure falls under one of these categories, you must add this tag.
name	Conditional	Any unique name with which the team can identify a resource meaningfully. This is for service owners to decide how to name their services to make them easier to identify. Example: us-west-2a-front-end-01.
tf_managed or cf_managed	Conditional	Indicates if cloud resource is managed via Terraform (tf), CloudFormation (cf), or other Infrastructure as Code (IaC) tool. (set to TRUE if tf / cf managed, otherwise FALSE)

C. Resource Cleanup-related Tags

You can utilize tags to help with resource housekeeping, especially if you use automation to find and act on these tags. Tags like this can help with both cost optimization and compliance.

Tag Name	Req'd?	Purpose
remove_after_date	Conditional	Required for resources created outside of IaC (regular) process and other temp. environments: If any additional infrastructure is created in response to an incident or for testing purposes, this tag helps remove cloud resources after the specified period when no longer needed. Example: remove_after_date = "12/21/2021"
shutdown tag		This tag will be used for non-prod workloads where resources can be turned off during non-business hours and weekends.

Tagging Enforcement

Enforcing tagging is important for data integrity, but before you roll it out throughout your company, make sure you plan and communicate well. This reduces the possibility of disturbances and guarantees that everyone agrees.

Not every tag has the same significance. Considering your business environment, concentrate on choosing a limited number of essential tags required for cost allocation, reporting, and compliance. For recommendations, see the "Recommended Tags" section.

Use tflinter and other similar tools to ensure that your resources have necessary tags if you are using Terraform for infrastructure as code (IaC). By doing this, possible problems with data quality are avoided and uniform tagging procedures are ensured.

You might present this concept to your SRE/Cloud team and request that they apply an obligatory (service_name = "xyz") tag to all of your resources, akin to tflinter. Tflinter can provide both validation and tagging enforcement. A sample code snippet showing how to utilize tflinter for tagging compliance can be found in Figure 4.

```

# Add your tfliinter rule in a location like- docs/rules/aws_resource_missing_tags.md
# aws_resource_missing_tags
Require specific tags for all AWS resource types that support them.
## Configuration
```hcl
rule "aws_resource_missing_tags" {
 enabled = true
 tags = ["service_name", "service_owner", "cost_center"] #note that these are case sensitive
 exclude = ["aws_autoscaling_group"] # (Optional) Exclude some resource types from tag checks, here, aws_autoscaling_group
resource type will be ignored.
}
```

#Assume you have a instance.tf file as follows:
resource "aws_instance" "instance" {
instance_type = "m5.large"
tags = {
service_name           = "xyz"
cost_center    = "abc@org.net"
}
}

```

Once you have tfliinter installed, you will now be able to execute tflint command in the terraform repository to check for any violations.

```

$ tflint
1 issue(s) found:
Notice: aws_instance.instance is missing the following tags: "service_owner" (aws_resource_missing_tags)

```

Figure 4: tflinter to enforce tagging compliance

Reference

- [1]. M. Bayramusta and V. A. Nasir, "A fad or future of IT? A comprehensive literature review on cloud computing research," International Journal of Information Management, vol. 36, pp. 635-644, 2016.
- [2]. ProsperOps, "AWS Cost Allocation Tags," ProsperOps Blog, [Online]. Available:<https://www.prosperops.com/blog/aws-cost-allocation-tags/> [Accessed: Apr. 26, 2023]
- [3]. York, "Tagging to Optimize Show Backs and Chargebacks In the Cloud," CloudSaver Resources, [Online]. Available:<https://www.cloudsaver.com/resources/articles/tagging-to-optimize-show-backs-and-chargebacks/> [Accessed: Apr. 26, 2023]
- [4]. Densify, "AWS Tagging Best Practices," FinOps, [Online]. Available:<https://www.densify.com/finops/aws-tagging-best-practices/> [Accessed: Apr. 26, 2023]
- [5]. J. DeRose, "AWS Cost Allocation Guide: Tagging Best Practices," Duckbill Group Blog, [Online]. Available:<https://www.duckbillgroup.com/blog/aws-cost-allocation-guide-tagging-best-practices/> [Accessed: Apr. 26, 2023]

