

Driving Cloud FinOps at scale with Google Cloud Tagging



Authored by:


Eric Lam - Head of Cloud FinOps, Google Cloud

Sarah McMullin - CostX Product Lead, Google Cloud

Rupa Patel - Product Manager, Google Cloud

Rich Hoyer - Director of FinOps, SADA

Table of contents



04

What is Google Cloud Tag? Why is it game changing?

06

How can I use Tags in Billing?

10

Best practices for implementing Google Cloud Tagging

11

Bringing this all together

13

What's Next

Google Cloud

Technical teams and FinOps practitioners take advantage of a variety of billing organizers to help segment their spend and allocate costs. Billing accounts, Projects, and Folders are all useful, but Tags and Labels are often the best choices when you need more flexibility. Our experience has shown that the greater flexibility Labels in particular offer makes them more prone to problems than formal billing organizers like accounts.

The most common challenges we encounter with labels are poor coverage and a lack of integrity in the labeling data. The latter tends to be a much greater problem than the former, since poor label coverage is clearly visible by tracking a numeric KPI such as “percent of dollars spent without a value populated.” Nonsensical or inaccurate label values, on the other hand, are more difficult to diagnose since verification of accuracy is conducted by executing a time-consuming, case-by-case spot check.

To solve this, Google has just rolled out Tags, which offer much better governance features than Labels — features that should dramatically improve data integrity when applied properly.

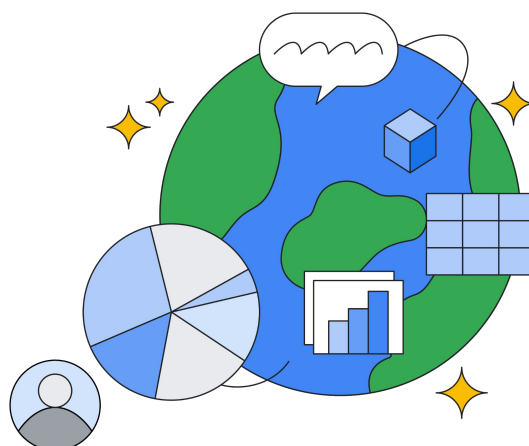
We will examine the differences between Tags and Labels below, but first let’s explore some recommendations around tagging taxonomies in general. To avoid problems with Labels or Tags, we suggest starting with a “demand side” approach to generating tagging policies, rather than a “supply side” approach.

In the “supply side” approach, technical leaders and FinOps practitioners conjure their own tagging taxonomies based on their personal ideas of what Tags or Labels “may” be useful. This can lead to unnecessary key-value pairs or a lack of useful key-value pairs, which is why we suggest avoiding this approach.

In a “demand side” approach, on the other hand, all interested parties in finance, accounting, management, and technical teams share their current and anticipated reporting needs in working sessions to determine Tag and Label taxonomies. From there, practitioners should determine which tag or label key-value pairs will support the desired reporting, and ensure that those key-value pairs and only those key value pairs are enforced as part of the policy. Ideally, teams should clearly document tagging and labeling policies, including details on permitted values, governance measures in place, and rationale for the key-value pairs.

To maintain strong governance, it’s critical that you perform a regular cadence of checks for tagging or labeling coverage to ensure that values are accurate and conform to the policies. Importantly, data integrity checks must accompany any mandatory tagging or labeling policy to avoid “junk” values being used by technical owners to meet the requirement. This is not an uncommon occurrence, and it is a significantly bigger problem to have junk labeling data than to have poor coverage.

Labels are prone to such challenges because Label key value pairs can be created by technical owners with no restrictions on Label values apart from the inability to include certain characters. This can lead to a variety of problems, among them is a lack of consistency in the values populated — such as using the values “Production,” “prod,” “prd,” and various misspellings in the Production State label. These challenges can also lead to inadvertently deleting both keys and values.






Google's new Tags offer significantly improved governance features — relative to Labels — that can greatly reduce error rates. Tag Admins, Tag Users, and Tag Viewers are each defined using the IAM Policy Engine before a single Tag can be created. Admins have the ability to define not only the Tag fields but a full, restricted set of possible values for each. Admins can also define what level of enforcement to employ — such as constraints on untagged resources. Tag Users have the ability to populate the tag fields, and all who are assigned Tag Viewer privileges can report against them. Tag values are inherited, which greatly facilitates reporting at the level of parent billing organizers. Finally, since Tags can be centrally managed through either the console or the API, it's much easier to make large-scale changes.

When applied correctly, these new governance features can result in far better tagging coverage and data integrity. For those who prefer labels for their flexibility, the option to use them is still available, giving technical owners and others greater freedom of use.

What is Google Cloud Tag? Why is it game changing?

The Google Cloud resource hierarchy can organize your resources into a tree structure to help you manage resources at scale, but it models only a few business dimensions, including organization structure, regions, workload types, and cost centers. The hierarchy lacks the flexibility to layer multiple business dimensions together.

Tags provide a way to conditionally allow or deny policies based on whether a resource has a specific tag. You can use tags and conditional policy enforcement for fine-grained control across your resource hierarchy. Tags are structured as a key-value pair. A Tag key resource can be created under your organization resource, and tag values are resources that are attached to a key — as a Tag key environment with values production and development, for example.



Google Cloud

Using Tags allows you to create more reliable reporting. There are a few benefits to using Tags versus Labels including:

1. Tags offer more reliable reporting with strong **governance** policies set in IAM Policy Engine. This means for tag values you can assign a Tag Viewer, Tag Admin, and Tag User. With Labels' lack of a strong permissions model, resource owners are able to remove cost-tracking labels set by cloud FinOps operators
2. Tags support **binding and inheritance**. You can centrally manage tags, and through the use of tag inheritance, multiple resources can be tagged through a single tag binding operation. With Labels, someone must individually update all resources if a change is required. Tags can bind to the [following resources](#). For reference, labels support these [resources](#).
3. Tags have **structure**. Tag keys, tag values, and tag bindings are all discrete resources. Labels are not resources; they are attributes of resources which can make it difficult to establish a centralized governance. Tags can be centrally managed through Tags API and Console.
4. Tags have an **enhanced character set**. A tag value can include upper ASCII symbols such as "@" or "." which is useful to assign a tag value as an email of the Cost Center or Team owner.



How can I use Tags in Billing?

Google Cloud Billing now provides access to Tags data in the Billing Export to BigQuery via the Standard and Detailed usage cost data export. Support for Tags in Billing Console is coming in late 2023. Tag data can help perform Cost Management tasks like chargebacks, audits, and other cost allocations. For more information on tags, please see the [Tags Overview](#) section of this technical paper. For a general primer about our Billing data export to BigQuery, see our Cloud Billing export to BigQuery [documentation](#).

What's new with Tags in Billing BigQuery Export?

The Standard and Detailed Cost exports will now include four fields for Tags data: Tags Key, Tags Value, Tags Inherited, and Tags.Namespace. Tags data will appear in BigQuery Export for Resources, Projects, Folders, and Organizations.

New Fields

Tags Key

This field represents the key associated with this particular tag. A tag key resource can be created under your organization resource.

Tags Value

This field represents the values — the attached resources — associated with this particular key. At any given time, exactly one value can be attached to a resource for a given key.

Tags Inherited

This field represents whether this tag binding is inherited (Tags Inherited = True) or direct/non-inherited (Tags Inherited = False). Tags have a robust [inheritance model](#) which allows users to create a tag binding to a parent resource in the [resource hierarchy](#), and that tag binding will inherit down to all of its descendants.

For example, imagine a user creates a tag binding — such as cost-center=abc — against a project that only has one GCE instance running within it. That tag binding will inherit down to the GCE instance and within the billing export; there will be a line for the respective GCE instance SKU with the following:

- Tags Key: cost-center
- Tags Value: marketing-emea
- Tags Inherited: true

Alternatively, if that tag binding was created against the instance directly, the respective line within the billing export would look like:

- Tags Key: cost-center
- Tags Value: marketing-apac
- Tags Inherited: false

Google Cloud

Example 1

In Standard and Detailed Exports, you might have a line item that looks like this:

| Billing Account Id | Service Description | SKU ID | SKU Description | Usage | Cost |
|---------------------|---------------------|----------------|---------------------|-------|-------|
| 00617-12345C-AABBCC | Compute Engine | 2E27-4F75-95CD | N1 Instance Core | 1 hr | \$.20 |
| 00617-12345C-AABBCC | Compute Engine | 5T13-8D95-DJ94 | Storage PD Capacity | .50 | \$.02 |

With the introduction of Tags in Standard and Detailed Export, you'll see the new fields as shown below:

| Billing Account Id | Service Description | SKU ID | SKU Description | Tags.Key | Tags.Value | Tags.Inherited | Tags.namespace | Usage | Cost |
|---------------------|---------------------|----------------|---------------------|------------|-------------|----------------|----------------|-------|-------|
| 00617-12345C-AABBCC | Compute Engine | 2E27-4F75-95CD | N1 Instance Core | costcenter | advertising | false | 324908234 | 1 hr | \$.20 |
| | | | | env | staging | true | 324908234 | | |
| 00617-12345C-AABBCC | Compute Engine | 5T13-8D95-DJ94 | Storage PD Capacity | env | prod | true | 324908234 | .50 | \$.02 |

Example 2

Say you have two services, one tagged as associated with “costcentera” and one tagged as “costcenterb” on February 1st, 2022. Due to an internal reorganization on February 16th, 2022, the tags for “costcentera” and “costcenterb” have been removed from the services and “costcenterc” and “costcenterd” have been assigned. The history for these resources will show the tags associated with the resources during that usage time, as shown in the following example:

Your BigQuery Export might look like this:

| Service Description | SKU ID | Usage Start Time | Usage End Time | SKU Description | Tags.Key | Tags.Value | Tags.Inherited |
|---------------------|--------------|------------------|----------------|-----------------|-------------|-------------|----------------|
| Compute Engine | 2E27-4F75-95 | 2022-02-01T2 | 2022-02-01T2 | N1 Instance | costcentera | advertising | false |
| | CD | 3:00:00 | 4:00:00 | Core | | | |
| Compute Engine | 2E27-4F75-95 | 2022-02-16T2 | 2022-02-16T2 | N1 Instance | costcenterc | marketing | false |
| | CD | 3:00:00 | 4:00:00 | Core | | | |
| Compute Engine | 5T13-8D95-D | 2022-02-01T2 | 2022-02-01T2 | Storage PD | costcenterb | advertising | false |
| | J94 | 3:00:00 | 4:00:00 | Capacity | | | |
| Compute Engine | 5T13-8D95-D | 2022-02-16T2 | 2022-02-16T2 | Storage PD | costcenterd | marketing | false |
| | J94 | 3:00:00 | 4:00:00 | Capacity | | | |

Best practices for implementing Google Cloud Tagging

By following these best practices and using Google Cloud Tagging to its full potential, organizations can better manage and organize their resources within the Google Cloud Platform, resulting in improved efficiency and cost savings.

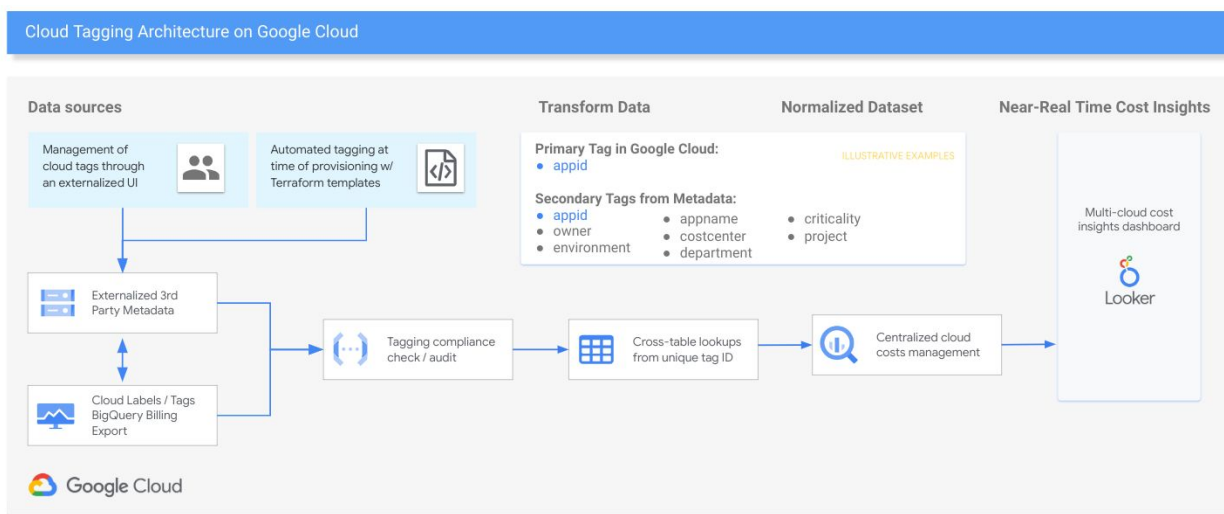
- **Establish a tagging strategy:** Before implementing tags, it's important to establish a clear strategy for how they will be used. This can include deciding on the types of Tags that will be used, the format of the Tags, and the naming conventions for Tags.
- **Use consistent naming conventions:** Consistently naming tags will make it easier to find and manage resources. For example, using a standardized prefix for all Tags, such as "env:", can help identify the environment a resource is in.
- **Use Tags to organize resources by business groups:** Tags can be used to organize resources by business groups like marketing, customer-facing applications, and shared services. As such, group cloud resources first by workload and environment, and then the workloads and environments by business group.
- **Use Tags to identify ownership:** Tags can be used to identify the owner or team responsible for a resource. This can help with identifying who to contact for support or troubleshooting.
- **Automate tagging:** Google Cloud Tagging can be automated using tools like Google Cloud API or Terraform templates. This allows for resources to be tagged as they are created, reducing the need for manual tagging.
- **Continuously adjust and optimize:** Monitor for nonstandard cases and changes in organization structure that may require redefining tags. Set budgets by business group and set up alerts to notify when workloads go over budget.



Bringing this all together

As we continue to expand the capabilities and robustness of our tagging capabilities and feature sets, customers have implemented cloud labeling and tagging through various approaches. One of the leading practices that we have seen time and time again is to also leverage an external metadata repository to manage cloud tagging and the associated attributes to the applications, products, and services.

Reference Architecture of an externalized metadata tagging repository:



Advantages of an externalized metadata tagging repository:

- Ease of management
- Flexibility for updates
- Referenceable lookup
- Minimal source tagging
- Multi-cloud platform support

Google Cloud



Ease of Management

By using an externalized metadata repository for tagging, the tags can be easily managed through an external database rather than through the Google Cloud native console or Terraform scripts updates.



Flexibility for updates

Updates to the cloud labeling and tagging can be performed through a user friendly interface that is accessible by the end users rather than access to Google Cloud native console or Terraform scripts which often are limited to improve security posture.



Referenceable lookup

Many customers implement a unique identifier through a single cloud label or tag that is applied directly to the Google Cloud services. By using this unique identifier, customers can do a lookup of the unique identifier against the metadata repository to retrieve the additional attributes.



Minimal source tagging

The leading practice is not to have a plethora of tagging or labels applied to each of the cloud services, but to only apply the tags where they are needed. Externalizing the additional attributes required for the applications effectively enables ease of management and flexibility to make updates.



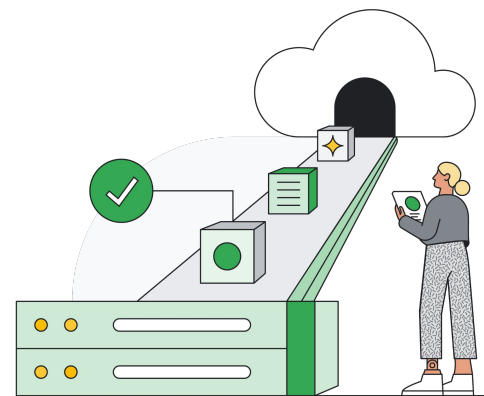
Multi-cloud platform support

Many customers have a multi-cloud strategy, and the implementation of the cloud tagging and labeling needs to be standardized and consistent across all cloud platforms for management and reporting purposes. Inconsistent implementation of tagging and labeling can cause reporting nightmares and loss of confidence by the end users.

What's next?

Google Cloud Tagging is a powerful tool for organizing and managing resources within the Google Cloud Platform. By establishing a proper tagging strategy, using consistent naming conventions, and automating the tagging process, companies can better manage and organize their resources, resulting in improved efficiency and cost savings. It is the fundamental capability of Cloud FinOps to enable cost transparency and drive financial accountability across the organization.

If you are interested in implementing [Google Cloud Tagging](#) for your organization, we recommend reaching out to the Google Cloud Professional Services team or a Google Cloud Partner like [SADA](#) for assistance in developing a strategy and implementation plan that best fits your needs. Additionally, you can check Google Cloud documentation and tutorials to learn more about the features and capabilities of [Google Cloud Tagging](#). With the right approach and guidance to implement a tagging strategy, you can greatly accelerate your [Cloud FinOps](#) practices at scale across your organization.



Driving Cloud FinOps at scale with Google Cloud Tagging

February 2023

Learn more about [Google Cloud Tagging](#).