



Cost Management Service 1-latest

Managing cost data using tagging

Organize resources and allocate costs with tags

Cost Management Service 1-latest Managing cost data using tagging

Organize resources and allocate costs with tags

Legal Notice

Copyright © 2024 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

XFS[®] is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

Learn how tagging works in cost management, and outlines strategies for managing your cost data with tags and labels. Cost management is part of the Red Hat Insights portfolio of services. The Red Hat Insights suite of advanced analytical tools helps you to identify and prioritize impacts on your operations, security, and business.

Table of Contents

CHAPTER 1. MANAGING COST DATA USING TAGGING	3
CHAPTER 2. PLANNING YOUR TAGGING STRATEGY	4
2.1. WHY USE TAGS?	4
2.2. CONSIDERATIONS FOR YOUR TAGGING STRATEGY	4
CHAPTER 3. CONFIGURING TAGS AND LABELS IN COST MANAGEMENT	7
3.1. HOW COST MANAGEMENT ASSOCIATES TAGS	7
3.1.1. Tag matching hierarchy in cost management	7
3.1.2. OpenShift label inheritance in cost management	8
3.1.3. Direct resource matching (instance ID)	8
3.1.4. Special OpenShift tags	8
3.1.5. Custom tags	9
3.2. CONFIGURING TAGS ON YOUR INTEGRATIONS	9
3.2.1. Adding tags to an AWS resource	9
3.2.2. Adding tags to a Microsoft Azure resource	10
3.2.3. Adding tags to a Google Cloud resource	10
3.2.4. Viewing labels in an OpenShift namespace	10
3.2.5. Enabling and Disabling tags in cost management	10
3.2.6. Configuring Amazon Web Services cost categories in cost management	11
CHAPTER 4. VIEWING AND EXPORTING YOUR COST DATA	12
4.1. FILTERING YOUR COST DATA VIEW	12
4.2. GROUPING COST DATA BY TAG CATEGORY	12
4.3. EXPORTING COST DATA TO A REPORTING TOOL	13
CHAPTER 5. ADDITIONAL RESOURCES	15
5.1. TAG SPECIFICATIONS BY INTEGRATION TYPE	15
5.2. FURTHER READING	16
PROVIDING FEEDBACK ON RED HAT DOCUMENTATION	17

CHAPTER 1. MANAGING COST DATA USING TAGGING

This guide explains in depth how tagging works in cost management, and how you can use tagging to best organize and view your resources to manage your costs.

CHAPTER 2. PLANNING YOUR TAGGING STRATEGY

2.1. WHY USE TAGS?

Tags (also sometimes called labels) are a string of custom metadata assigned to resources that your organization can use in multiple ways.

In the context of cost management, tags help you to differentiate and allocate costs between various parts of your environment for a more accurate view of your cost data, filling the gap between business logic and the resource. This allows you to assign organizational-specific details to help digest information later.

As many projects and technical services (for example, environment, region, cost center) support a business service, tags can map business concepts to reports providing clarity. In that way, a properly allocated tag can help you to show costs grouped by those business concepts. Tags should be used to complement hierarchies in clouds, or specific hierarchies already present in your clusters (like projects).

Tags are sometimes utilized in more applications than cost management. For example, it is possible and recommended to use them to define certain operations, like business automation, operational profiles, or access and security controls. You can then apply policies based on those tags. However, this document does not consider those use cases. Cloud providers place limits on the number of tags or labels associated with resources, so you should consider all of those uses before defining your tagging strategy for cost management.

Additionally, tags are sometimes used to split resources into smaller units when it is not possible to directly organize resources into projects or subprojects. For example, a shared cluster running many services to provide different business capabilities can use tags to differentiate applications rather than splitting them into projects. Additionally, an AWS account can run multiple different services for projects, and sometimes those services are shared between several applications (like the load balancer or any other shared resource).

Furthermore, tags also help you identify relationships between integrations, allowing you to group applications across multiple clusters tagged with the same environment, cost center, or team. Therefore, making it possible to identify the costs of an application when running on the development, quality, or production clusters. Tags can also help to identify dependencies when there is not a direct link between a resource, such as the link between an RDS database and the OpenShift project using it.

2.2. CONSIDERATIONS FOR YOUR TAGGING STRATEGY

When planning your tagging strategy, these considerations can help you decide how to organize and report costs for your integrations.

Methodology

It is best practice to implement a Crawl, Walk, Run strategy when initially assigning tags for your resources. Start with the minimum number of tags you feel are required to accomplish your organizational objectives. Then slowly build on that foundation over time. As your understanding of how your organization leverages cost management grows, so will your skill in implementing a tagging strategy.

Map business to reporting

Define the business perspectives you want to report on. For example, your taxonomy for cost management could consider these different perspectives:

Ownership and usage:

Defining the owner and the user of the resource: for example, the unique identifier of the user who requested the resource, and the one that is actually consuming the resource.

Tenancy:

If your environment is shared, it can be beneficial to understand which group or business unit has requested the resource. When the user is part of different groups, one group needs to be selected. For cost reporting, you can achieve this in many cases using cost center; but department, project, or partner are also good candidates.

Location:

For organizations with software deployed globally, cloud providers already identify the region where your resources are running, but your private cloud may be different.

Environment or stage:

You may want to differentiate between development and production, so that different cost decisions can be made depending on the environment where you are creating or running the resources. If your development pipeline already includes stages, such as development, testing, staging, pre-production and production, this is a good candidate.

Application / Project / Service / Event:

Perhaps your environment is providing a service, such as a group of transient resources for an event (for example, your yearly customer-focused conference). You could even include the application version.

Standardize your labels

Consistency is the most important aspect of a tagging strategy to deliver accurate and comparable cost reporting results.

Create a clear tagging policy that defines what resources need to be tagged, what tags are mandatory and what tags are optional, making sure that there is no room for interpretation.

If the values need to be chosen between a list, verify that those values are defined, consistent, and easily accessible, or that the list is presented to the user. For example, if you are defining development with the key "Development", do not also use variations such as "Dev", "DEV", or "R&D" to also identify resources as "Development".

Tag all elements on your integrations (manually or through automation)

Since untagged resources cannot be reported, tag as many elements as possible, ideally using automation to prevent human error. Integrations have different automation features to take advantage of for tagging:

- In Azure, you can use Azure Policies to enforce tagging rules and conventions and avoid resources being deployed that do not comply with your expectations. You can create a policy that automatically applies necessary tags during provisioning, that enforces a predefined format for dates, or that makes some tags mandatory for some resource type.
- In AWS, you can use IAM policies for the same. Additionally, you can use an automation tool such as Ansible to add the necessary tags during provisioning and to verify that all the resources have been properly tagged.
- OpenShift Container Platform does not presently have a method of automating labelling.

Review your tags often and refine as needed

Define your tags and use them as early as possible with cost management, even if you need to adjust your tagging scheme afterwards.

Review the resulting reports with your business owner and stakeholders early on to ensure your tags are helping you to generate the desired reports, and review your tagging strategy every few weeks to optimize it.

Select your tag terminology

- Name your resources with names that allow you to identify your resources without accessing metadata, and then continue by adding metadata to it. Many clouds have a guide about how to do this properly; see [Chapter 5, Additional resources](#) for links.
- Map your resources into keys and values. Keys will map to perspectives, while values will define the different options allowed for each key. In some cases, the value will be *Null*.



NOTE

Not all integrations allow the same identifiers, and have different limitations. See [Section 5.1, “Tag specifications by integration type”](#) for limitations by integration.

CHAPTER 3. CONFIGURING TAGS AND LABELS IN COST MANAGEMENT

You must configure tags in each integration before cost management can use the tags to automatically organize your cost data.

After adding your integrations to cost management:

1. Tag or label resources on each of your integrations. See [Section 3.2, “Configuring tags on your integrations”](#).
2. Refine and add to your tags to optimize your view of cost data. See [Section 2.2, “Considerations for your tagging strategy”](#).



NOTE

See the [Getting started with cost management](#) guide for instructions on configuring integrations.

3.1. HOW COST MANAGEMENT ASSOCIATES TAGS

Tags in AWS and Microsoft Azure and labels in OpenShift consist of key:value pairs. When the key:value pairs match, the AWS/Azure and OpenShift costs are automatically associated by cost management. Tag matching in cost management is not case sensitive: for example, an AWS resource tagged *APP* and an OpenShift resource tagged *app* are a match:

Table 3.1. Example: Tag matching

Source and resource type	Key	Value
AWS resource (RDS)	APP	Cost-Management
OpenShift pod	app	cost-management

If an AWS resource tag matches with multiple OpenShift projects, the cost and usage of that resource are split evenly between the matched projects.

This is not the case with AWS compute resources that are matched through the instance ID-node relationship. In that case, cost and usage are broken down using information about a project’s resource consumption within the OpenShift cluster.

By default, cost management tracks AWS compute usage and costs by associating the Amazon EC2 instance ID or Microsoft Azure virtual machine instance ID with the OpenShift Container Platform node running on that instance.

3.1.1. Tag matching hierarchy in cost management

To identify your OpenShift resources running on AWS or Azure instances, cost management matches tags between integrations in the following order:

1. Direct resource matching (AWS EC2 instance ID or Azure virtual machine instance ID)
2. Special OpenShift tags

3. Custom tags

3.1.2. OpenShift label inheritance in cost management

OpenShift labels follow an inheritance pattern from cluster to node, and from project to pod. You can associate costs at the node or project level without labeling every pod in your cluster.

Key-value pairs from node and project labels are inherited at the pod level for pod metrics in cost management. Key-value pairs from the cluster and node labels are inherited at the project level by the persistent volume claims (PVC) at each level. You can group by cluster, node, or project labels to see relevant PVCs in those workloads.

If a key already exists in the pod, then the value for that key in the pod remains. Cost management does not overwrite the pod value with the project or node value. A similar process happens from node to project.

Consider the following examples.

Example 1: Your organization assigns the label **app** and the value **costpod1** to a pod. The project for this pod has the label **app** and the value **cost-project**. These resources are running on a node with the label **us-east-1**. The label **app** and the value **costpod1** remain associated with the pod.

Example 2: Your organization has a project with the label **app** and the value **cost-project**. The project has three pods running and they are not labeled. Cost management associates the label *app* and the value **cost-project** with these pods.

3.1.3. Direct resource matching (instance ID)

The integrations apply these identifiers automatically. This form of tagging provides a direct link between Microsoft Azure or AWS instances and OpenShift nodes.

AWS assigns every EC2 instance a resource identifier (a number such as **i-01f44b3d90ef90055**). OpenShift nodes are matched directly to the AWS EC2 instance the cluster is running on using the AWS resource identifier. The OpenShift reports in cost management (generated from Prometheus data) include this identifier for nodes. Similarly in Microsoft Azure, each virtual machine instance ID is included in the OpenShift reports for cost management.

3.1.4. Special OpenShift tags

There are three special-case AWS tags you can use to associate cost with OpenShift:

- **openshift_cluster**
- **openshift_node**
- **openshift_project**

These tags have matching priority over custom tags, and are especially useful in differentiating the costs of different OpenShift clusters running on the same AWS instance.

To use this tagging method to identify an OpenShift cluster, tag your AWS instance with the key **openshift_cluster**, and provide the OpenShift integration name as the value. In the following example, the name of OpenShift integration in the cost management application is **dev-cluster**:

Table 3.2. Example: Special OpenShift tags

Source and resource type	Key	Value
AWS resource (RDS)	openshift_cluster	dev-cluster
OpenShift cluster	No tags needed. This will match if the name of the OpenShift integration in cost management is dev-cluster .	No tags needed.

3.1.5. Custom tags

You can use any key:value combination as tags, and cost management will associate identical tag key and values together. You can then group costs by tag key, account, service, region, and more to view your costs and charge for that tag.

Table 3.3. Example: Custom tags

Source and resource type	Key	Value
AWS resource (RDS)	team	engineering
OpenShift pod	team	engineering

3.2. CONFIGURING TAGS ON YOUR INTEGRATIONS

To control the tags cost management imports, activate or enable the tags you want to view by integration:

- AWS tags must be activated, and are then selected and exported to cost management in the Cost and Usage report. See [Activating AWS tags for cost management](#) in the *Adding an Amazon Web Services (AWS) source* guide for instructions.
- Microsoft Azure tags are exported to cost management in the cost export report configured in [Configuring a daily Azure data export schedule](#).
- OpenShift Container Platform labels are exported by Metering and included in the metrics reports that cost management uses as input.

3.2.1. Adding tags to an AWS resource

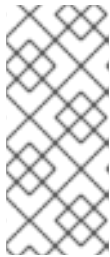
Amazon creates certain identifiers automatically, like the EC2 instance resource identifier (a number such as **i-123456789**), which cost management uses similarly to a tag.

You can also add your own tags at the individual resource level. These tags must be activated for Cost and Usage Reporting to export them to the cost management application.

Configure AWS tags for cost management using the following steps:

Procedure

1. Create and apply tags to your AWS resources.



NOTE

If you converted from a compatible third-party Linux distribution to Red Hat Enterprise Linux (RHEL) and purchased the RHEL for third party migration listing in AWS, activate the AWS Cost Allocation Tags tags for your RHEL systems. Enable the **com_redhat_rhel** tag on the cost allocation tags page. Enter **RHEL 7 ELS**, **RHEL 8**, or **RHEL 8 ELS** as the tag value to match your version of RHEL.

For instructions in the AWS documentation, see [User-Defined Cost Allocation Tags](#).

2. Activate the tags you want to be collected by the cost management application through the Cost and Usage Report. In the AWS Billing console, select the tags you want to activate from the Cost Allocation Tags area.

For instructions in the AWS documentation, see [Activating the AWS-Generated Cost Allocation Tags](#).

3.2.2. Adding tags to a Microsoft Azure resource

Adding an Microsoft Azure integration creates identifiers for virtual machine instances automatically, which cost management uses similarly to tags to associate Azure resources to related OpenShift resources.

You can also add your own tags in Microsoft Azure at the individual resource level.

Create and apply Microsoft Azure tags for cost management using the instructions in the Microsoft Azure documentation: [Use tags to organize your Azure resources and management hierarchy](#).

3.2.3. Adding tags to a Google Cloud resource

You can apply custom labels to Google cloud resources, such as virtual machine instances, images, and persistent disks. These labels are automatically added to your BigQuery export and sent to cost management.

Procedure

- Create and apply labels to your Google Cloud resources.
See [Creating and managing labels](#) in the Google Cloud documentation for instructions.

3.2.4. Viewing labels in an OpenShift namespace

The AWS or Microsoft Azure tag equivalent in OpenShift is a label, which also consists of a key:value pair. The cost management service collects OpenShift tag data from nodes, pods, and persistent volumes (or persistent volume claims) using Prometheus metrics and Cost Management Metrics Operator.

To view the available tags, navigate to a resource in the OpenShift web console. Any assigned labels are listed under the **Labels** heading, for example: `openshift.io/cluster-monitoring=true`.

3.2.5. Enabling and Disabling tags in cost management

All cloud provider tags are activated in cost management by default. Sometimes having too many resource tags can affect cost management performance. Enabled tags are limited to 200 per account. Unnecessary tags can also make managing your costs more complicated when grouping tags and matching key:value pairs. Disable tags you are not actively using to reduce these potential issues.

Prerequisites

- You must have Organization Administrator or Cost Price List Administrator privileges to change these settings in cost management. See [Limiting access to cost management resources](#) in *Getting started with cost management* for more information about user roles and access.

Procedure

1. From [cost management](#), click **Cost Management → Settings**.
2. Click the **Tags and labels** tab.
3. Select the tags you want to disable.
4. Click **Disable tags**.
This tag is now deactivated for the cost management application. You can enable tags you have disabled on the same page by selecting the tags you want to enable and clicking **Enable tags**.

3.2.6. Configuring Amazon Web Services cost categories in cost management

You can enable or disable Amazon Web Services (AWS) cost categories in the cost management service. AWS cost categories allow your organization to group meaningful cost and usage information in addition to tags. In order to use cost categories in cost management, they must first be configured in the AWS Console. The following procedure describes how to enable cost categories in the cost management service.

Prerequisites

- You must have Organization Administrator or Cost Price List Administrator privileges to change these settings in cost management. See [Limiting access to cost management resources](#) in *Getting started with cost management* for more information about user roles and access.
- You have an Amazon Web Services integration added to cost management with cost categories enabled through the Amazon Web Services Console.

Procedure

1. From [cost management](#), click **Cost Management → Settings**.
2. Click the **Cost categories** tab.
3. Select the cost category to enable. You can select more than one.
4. Click **Enable categories**.
The selected cost categories are now enabled for the cost management application. You can also disable cost categories by selecting the cost categories you want to disable and clicking **Disable categories**.

CHAPTER 4. VIEWING AND EXPORTING YOUR COST DATA

4.1. FILTERING YOUR COST DATA VIEW


Tags allow you to customize your view of cost data. You can view resources by type (for example, project, node, cluster) or tag or label to investigate why certain resources show an increase in cost, or when data looks abnormal.

This example shows how to see how much each OpenShift project within a cluster is costing.

Prerequisites

- Your OpenShift cluster added as a cost management data integration. See [Integrating OpenShift Container Platform data into cost management](#) in *Getting started with cost management* for instructions.
- Your cloud infrastructure account added as a cost management data integration. See [Getting started with cost management](#) for instructions for your cloud provider type.
- Configure tags on your integrations. For tips and configuration instructions, see [Section 3.2, "Configuring tags on your integrations"](#).

Procedure

1. From the **OpenShift details** menu, click the filter button and select **Tag**.
2. In the **Choose key** dropdown list, select the key to filter by. For example, select `environment` to view clusters with the `environment` tag. Selecting a tag key reveals another dropdown to choose the value to filter by.
3. In the **Choose value** dropdown list, select one or more values to filter by. For example, select `qe` and `dev` to show cost data for OpenShift projects with these tags.
4. To view more information about each project:
 - Click the arrow icon for each resource to see more information such as the cluster(s) the resource belongs to, and CPU and memory usage, limits, and requests.
 - Click **More options**  to reveal more viewing options:
 - Click **View price lists** to see rates applied to the OpenShift metrics to calculate the costs.
 - Click **View historical data** to open the daily usage comparison view, which compares usage, request, and limits by day between months for that resource.
 - Click **View all projects** or **View all tags** to see related resources and metadata.
5. Click **Clear all filters** to reset your OpenShift details view.

4.2. GROUPING COST DATA BY TAG CATEGORY

You can group resources by tag category to further investigate your cost data.

Grouping and filtering is useful for finding the root cause of a cost or problem, or investigating part of the environment that acts independently of the rest, such as a cost center or an specific environment.


This allows you to hide information about the rest of the environment to help avoid unneeded complexity in cost data results and allow you to find the desired information, which can otherwise be hidden among other data.

This example shows how an educational course provider who is running a lab environment on OpenShift Container Platform can use tag grouping to filter cost information by student and course.

Prerequisites

- Your OpenShift cluster added as a cost management data integration. See [Integrating OpenShift Container Platform data into cost management](#) in *Getting started with cost management* for instructions.
- Your cloud infrastructure account added as a cost management data integration. See [Getting started with cost management](#) for instructions for your cloud provider type.
- Configure tags on your integrations. For tips and configuration instructions, see [Section 3.2, "Configuring tags on your integrations"](#).

Procedure

1. From the **OpenShift details** page in the **Group cost by:** field, select the tag key to group the cost by. In this case, select *Tag Key:user* to show results grouped by student user.
2. In the filter area, select **Tag**.
3. In the **Choose key** list, select the tag key *user*.
4. In the **Choose value** dropdown list, check the values *course_id* and *course_type* to identify how many students have taken the course X and the costs for that course.
5. To view more information about each resource, for example, how much course X has cost:
 - Click the arrow icon for each resource to see more information such as the cluster(s) the resource belongs to, and CPU and memory usage, limits, and requests.
 - Click **View Historical Data** to open the daily usage comparison view, which compares by month the usage, request, and limits per day for that resource.
 - Click **More options**  to reveal more viewing options:
 - Click **View historical data** to open the daily usage comparison view, which compares usage, request, and limits by day between months for that resource.
 - Click **Export data** to create a .CSV file for reporting. Specify a daily or monthly aggregate and click **Generate and download**.
6. Click **Clear all filters** to reset your OpenShift details view.

4.3. EXPORTING COST DATA TO A REPORTING TOOL

Tags allow you to customize your view of cost data. This is useful when you want to investigate further why certain resources show an increase in cost, or data looks abnormal.

This example shows how to view data for specific OpenShift resources, and export the data to your desired reporting tool.

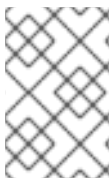
Prerequisites

- Your OpenShift cluster added as a cost management data integration. See [Integrating OpenShift Container Platform data into cost management](#) in *Getting started with cost management* for instructions.
- Your cloud infrastructure account added as a cost management data integration. See [Getting started with cost management](#) for instructions for your cloud provider type.
- Configure tags on your integrations. For tips and configuration instructions, see [Section 3.2, "Configuring tags on your integrations"](#).


Procedure

1. From the **OpenShift details** menu, click the filter button and select **Tag**.
2. In the **Choose key** dropdown list, select the key to filter by. For example, select `version`. Selecting a tag key reveals another dropdown to choose the value to filter by.
3. In the **Choose value** dropdown list, select one or more values to filter by. For example, select `qe` and `dev` to show cost data for OpenShift resources with these tags.
4. To export data about your resources, check the boxes next to each resource you want to export data for. Click **Export** to open the export options dialog.
5. Specify a daily or monthly aggregate and click **Generate and download**.

The CSV file will download to your local system, and you can use it in your desired reporting tool.



NOTE

You can also export your data as a .CSV file from the **More options**  **> Export data** menu on each resource.

Click **Clear all filters** to reset your OpenShift details view.

CHAPTER 5. ADDITIONAL RESOURCES

5.1. TAG SPECIFICATIONS BY INTEGRATION TYPE

Tagging standards differ between integration types. To use the same tags/labels across integrations, you must use the most common of all the restrictions across the different providers.

The following table summarizes tagging and labelling criteria across AWS, Azure and OpenShift Container Platform 4:

Table 5.1. Tagging specifications by integration

Criteria	AWS	Azure	Google Cloud	Red Hat OpenShift
Name	Tags	Tags	Labels	Labels
Format	Key & value	Name & value	Key & value	Key & value Keys: [prefix/]name Prefix: must be a DNS subdomain
Allows empty value	Yes	Yes	Yes	Yes
Unique label per key	Yes	Yes	Yes	Yes
Case sensitive	Yes	No	Only lowercase letters	Yes
Limit per resource	50	50 (15 for storage)	64	N/A
Length of key	128	512 (128 for storage)	63	253(prefix) / 63(name)
Length of value	256	256	63	63
Allowed characters	Letters, numbers, and spaces representable in UTF-8, and the following characters: + - = . _ : / @	Tag names cannot contain these characters: <, >, %, &, \, ?, /	Only lowercase letters, numeric characters, underscores, and dashes.	The name segment is required and must be 63 characters or less, beginning and ending with an alphanumeric character ([a-z0-9A-Z]) with dashes (-), underscores (_), dots (.), and alphanumerics between

Criteria	AWS	Azure	Google Cloud	Red Hat OpenShift
Restrictions	The prefix "aws:" is reserved. Tags applied to EC2 can use any character. Not all resource types support tags.	Not all resource types support tags. Generalized VMs do not support tags. Tags applied to the resource group are not inherited by the resources.	Keys must start with a lowercase letter or international character.	Prefixes kubernetes.io/ and k8s.io/ are reserved. Not all resource types support tags.
Notes	You need to select the tag keys that will be included in cost and usage files and billing reports.	You can use a JSON string to go over the limit of keys.	There is no limit on how many labels you can apply across all resources within a project.	If the prefix is omitted, the label Key is presumed to be private to the user.

5.2. FURTHER READING

The following links provide further guidance on tagging for each integration type.

AWS:

- [AWS tagging strategies](#)
- [IAM: Add a specific tag with specific values](#)

OpenShift:

- [Kubernetes labels and selectors](#)
- [Kubernetes user guide: labels](#)

Microsoft Azure:

- [Azure resource naming and tagging decision guide](#)
- [Azure recommended naming and tagging conventions](#)
- [Use tags to organize your Azure resources and management hierarchy](#)
- [Enforce tags in Azure resource groups](#)

Google Cloud:

- [Creating and managing labels](#)

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

If you found an error or have a suggestion on how to improve these guidelines, open an issue in the [cost management Jira board](#) and add the **Documentation** label.

We appreciate your feedback!