



Cost Management Service 1-latest

Getting started with cost management

Learn how you can track your OpenShift costs

Cost Management Service 1-latest Getting started with cost management

Learn how you can track your OpenShift costs

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

Use cost management to track the cost of your clouds and containers. Start by setting up an integration, configuring cost management to process your data, and then learn how to use cost management to analyze your data. 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. This getting started guide will teach you everything you need to know to begin using cost management.

Table of Contents

PREFACE	3
CHAPTER 1. WHY USE COST MANAGEMENT	4
1.1. WHAT CAN YOU ACCOMPLISH WITH COST MANAGEMENT?	4
1.2. HOW DOES COST MANAGEMENT WORK?	4
CHAPTER 2. HOW TO SET UP COST MANAGEMENT	6
2.1. SIGN UP	6
2.2. CONNECT AND VIEW YOUR COST DATA	6
2.2.1. Option 1: On-premise	7
2.2.1.1. Install the Metrics Operator	7
2.2.1.2. Setting up OpenShift tags	7
2.2.1.3. Create a Red Hat OpenShift cost model	7
2.2.2. Option 2: Cloud	7
2.2.2.1. Install the Metrics Operator	7
2.2.2.2. Add an integration for your cloud provider	7
2.2.2.3. Setting up OpenShift tags	8
2.2.2.4. Create a cloud cost model	8
2.3. CONTROL YOUR PERMISSIONS	8
2.4. ANALYZE YOUR RESULTS	8
CHAPTER 3. KEY CONCEPTS IN COST MANAGEMENT	9
3.1. COST LAYERS	9
3.2. ADDITIONAL RESOURCES	10
CHAPTER 4. YOUR PRIVACY AND DATA IN COST MANAGEMENT	11
PROVIDING FEEDBACK ON RED HAT DOCUMENTATION	12

PREFACE

Cost management helps you monitor and analyze your OpenShift Container Platform and Public cloud costs to improve the management of your business. It is based on the upstream project Koku.

To get started, learn about the following topics:

- What you can do with cost management and why your organization might want to use it
- How to set up and configure cost management
- How to adjust your settings after setup
- How to use cost management

You can use cost management to track cost and usage data for your Amazon Web Services (AWS), Microsoft Azure, Google Cloud, Oracle Cloud, and OpenShift Container Platform environments.

CHAPTER 1. WHY USE COST MANAGEMENT

Cost management is a free offering as part of your subscription to the Red Hat Insights portfolio of services. With cost management, you can monitor and analyze your costs to improve the management of your business.

Cost management helps you simplify the management of your resources and costs across container platforms like OpenShift Container Platform, as well as public clouds like Amazon Web Services (AWS), Google Cloud, Oracle Cloud, and Microsoft Azure.

1.1. WHAT CAN YOU ACCOMPLISH WITH COST MANAGEMENT?

With the expanding scale and performance of containerized business applications, you need aggregated and meaningful data so that you can quickly analyze your cluster spending and align with business priorities.

To overcome business challenges, cost management gives your organization visibility into your costs down to the project level for on-premise and public cloud environments. This visibility gives IT and financial stakeholders a unique snapshot into the costs associated with applications.

With cost management, you can achieve some of the following goals:

- Visualize, understand, and analyze how you use your resources and costs across hybrid cloud infrastructure
- Track cost trends
- Map charges to projects and organizations
- Use cost models to normalize data and add markups
- Forecast your future consumption and compare it with your budgets
- Optimize your resources and usage
- Identify patterns of usage that you might want to investigate
- Integrate with third party tools that can use your cost and resourcing data

These preceding goals can ultimately help your organization optimize costs, increase efficiency, and save money.

1.2. HOW DOES COST MANAGEMENT WORK?

It's important to understand some key OpenShift concepts:

Cluster

a group of servers that are managed together and participate in workload management.

Node

a worker machine that is either virtual or physical, depending on the cluster.

- **Master node:** The master node hosts the control plane and manages the cluster, including scheduling and scaling applications and maintaining the state of the cluster.

- **Worker node:** Worker nodes are responsible for running the containers and executing the workloads.

Pod

a collection of one or more containers. It is the smallest unit possible.

Persistent volume claim (PVC)

Persistent volume (PV) framework enables cluster administrators to provision persistent storage for a cluster. Developers can use persistent volume claims (PVCs) to request PV resources.

At a high level, cost management calculates your costs by processing data from your integrations in the following ways:

- From your cloud bill, cost management takes the cost of all of your nodes and determines what nodes belong to what cluster and which nodes are worker or master nodes.
- Cost management then determines what pods are running on what cluster and namespace and calculates how much central processing units (CPU), memory, disk space, and PVCs each one uses.
- Cost management multiplies the cost from the cloud bill by the established usage metrics to calculate the amount of money that each pod is costing you.
- If you have a cost model, it distributes the cost of the platform or the cost of unallocated capacity.
 - If you do not create a Red Hat OpenShift Container Platform cost model, we use the implicit cost model. This method distributes the cost from the cloud bill based on CPU effective use.

Cost management does not use public prices. Rather, it reads your cloud bill to process the savings plans, reserved instances, discounts, or other costs that you have. Cost management also tracks which nodes run on which pods. If you have different instance types, or same instance types but with different prices, cost management can still attribute the correct cost to each pod.

CHAPTER 2. HOW TO SET UP COST MANAGEMENT

To get started with cost management, complete the following four steps, plus any applicable substeps:

1. [Access and log in to cost management.](#)
2. [Connect and view your cost data:](#)

Cost management can analyze cost data from on-premise instances of OpenShift or cloud-based instances of OpenShift. If your organization has an on-premise instance of OpenShift, complete the following additional steps:

- a. [Install the Metrics Operator.](#)
- b. [Set up OpenShift tags.](#)
- c. [Create a Red Hat OpenShift cost model.](#)

Cost management also supports AWS, Google Cloud, Oracle Cloud, and Microsoft Azure. To set up cost management for OpenShift that is running on a cloud provider, complete the following additional steps:

- a. [Install the Metrics Operator.](#)
- b. [Add an integration for your cloud provider.](#)
- c. [Set up tags.](#)
- d. [Create a cloud cost model.](#)

Finally, finish getting set up for both on-premise and cloud with the following steps:

3. [Control your permissions](#)
4. [Analyze your results](#)

2.1. SIGN UP

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.

You can access cost management in the [Hybrid Cloud Console](#). To get started, click [OpenShift → Cost Management](#).

After you sign up, configure a user with Cloud Administrator access that can add cloud or OpenShift integrations to your cost management. For more information, see [Configuring cloud integrations for Red Hat services](#).

2.2. CONNECT AND VIEW YOUR COST DATA

To begin analyzing your cost data, you need to enter information about your costs.

The steps you will take depend on if your organization set up an on-premise cluster with OpenShift, or integrated with a cloud provider.

2.2.1. Option 1: On-premise

To get started with an on-premise cluster, complete the following steps:

2.2.1.1. Install the Metrics Operator

Red Hat® OpenShift® Operators automate the creation, configuration, and management of instances of Kubernetes-native applications. Your OpenShift cluster should already be set up, but you additionally need to set up the Metrics Operator.

To install the Metrics Operator, follow the instructions in [Installing a cost operator](#).

2.2.1.2. Setting up OpenShift tags

Tags, also called labels, are strings of custom metadata that you assign to resources. You can use tags to differentiate and allocate costs between various parts of your environment.

To learn about the different use cases for tags and how to set them up, see [Managing cost data using tagging](#).

2.2.1.3. Create a Red Hat OpenShift cost model

Finally, you must add a cost model to accurately analyze your costs.

A cost model is a framework that uses raw costs and metrics to define calculations for your costs. You can record, categorize, and distribute the costs that the cost model generates to specific customers, business units, or projects.

To learn how to set up a cost model, see [Using cost models](#).

2.2.2. Option 2: Cloud

Cost management supports AWS, Google Cloud, Oracle Cloud, and Microsoft Azure. Unlike an on-premise cluster, you need to set up an integration to connect to your cloud provider.

To integrate cost management with your cloud provider, complete the following steps:

2.2.2.1. Install the Metrics Operator

Red Hat® OpenShift® Operators automate the creation, configuration, and management of instances of Kubernetes-native applications. Your OpenShift cluster should already be set up, but you additionally need to set up the Metrics Operator.

To install the Metrics Operator, follow the instructions in [Installing a cost operator](#).

2.2.2.2. Add an integration for your cloud provider

To enable cost management to monitor your costs with a cloud provider such as AWS, Google, Oracle, or Azure, you need to set up an integration. An integration is a provider account that cost management connects to and monitors.

The process to set up an integration for each provider varies. To learn how to add your specific integration to cost management, see the following guides:

- [Integrating OpenShift Container Platform data into cost management](#)

- [Integrating Amazon Web Services \(AWS\) data into cost management](#)
- [Integrating Google Cloud data into cost management](#)
- [Integrating Microsoft Azure data into cost management](#)
- [Integrating Oracle Cloud data into cost management](#)

2.2.2.3. Setting up OpenShift tags

Tags, also called labels, are strings of custom metadata that you assign to resources. You can use tags to differentiate and allocate costs between various parts of your environment.

To learn about the different use cases for tags and how to set them up, see [Managing cost data using tagging](#).

2.2.2.4. Create a cloud cost model

Finally, depending on your cloud provider, you need to add either an AWS, Google, or Azure cost model to accurately analyze your costs.

A cost model is a framework that uses raw costs and metrics to define calculations for your costs. You can record, categorize, and distribute the costs that the cost model generates to specific customers, business units, or projects.

To learn how to set up a cloud cost model, see [Using cost models](#).

2.3. CONTROL YOUR PERMISSIONS

You might want to limit access to your data to only specific users or organizations. To learn how to control permissions, see [Limiting access to cost management resources](#).

2.4. ANALYZE YOUR RESULTS

Now that your cost data is generated, you can analyze your results and make changes in your business. To learn more about cost analysis, go to [Visualizing your costs using Cost Explorer](#).

=

CHAPTER 3. KEY CONCEPTS IN COST MANAGEMENT

Review the following terms to ensure you understand important definitions in cost management.

Cost

The sum of infrastructure and supplementary costs or the total cost for the filtered data. This sum can include the cost for a particular account, region, or service.

Raw cost

The costs that a cloud provider reports before any cost model calculations.

Unallocated costs

The costs that you can view when you group OpenShift cost data by project. These costs are in the rows **Platform unallocated** and **Worker unallocated** and are available when viewing OpenShift cost data by project.

Platform unallocated costs

The costs for parts of primary and infrastructure nodes that are not allocated to run workloads. In this case, those workloads are the OpenShift platform or control plane.

Network unattributed costs

Costs associated with ingress and egress network traffic for individual nodes.

Worker unallocated costs

The costs that represent any unused part of your worker node's usage and request capacity.

Unattributed cost

The costs associated with a cloud integration that is not attributed to an OpenShift project.

Overhead costs

The costs that you outlay to administer your cloud account.

3.1. COST LAYERS

You can create costs at different layers of the environment. In cost management, all costs can be classified as either infrastructure or supplementary costs. To customize what cost management shows on the Details page, you can allocate and present your costs as any of the cost layers in your cost model.

Infrastructure cost

In cost management, all cloud costs are considered infrastructure costs. They can include costs that cloud providers directly report through an Amazon (AWS), Microsoft Azure, or Google Cloud export. However, you can also designate infrastructure costs calculated from a price list for rates specifically marked as infrastructure.



NOTE

OpenShift Container Platform node and cluster cost per month rates are classified as infrastructure cost by default. You can choose to classify them as supplementary.

Supplementary cost

All costs not directly attributed to infrastructure are identified as supplementary costs. In cost management, these costs come from applying a price list to the metrics taken from an OpenShift cluster. The rates in the price list must be specified as supplementary. This is the default for all OpenShift costs except node and cluster cost per month.

3.2. ADDITIONAL RESOURCES

For more information about cost models, see [Cost model terminology](#).

For information about distributing costs, see [Understanding cost distribution in cost management](#)

CHAPTER 4. YOUR PRIVACY AND DATA IN COST MANAGEMENT

To run cost management, we gather your cost and usage data, but do not collect identifying information such as user names, passwords, or certificates.

For more information about your privacy and data, log in to the [Customer Portal](#) and see [Red Hat's Privacy Policy](#) and [our FAQ page](#).

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

We appreciate and prioritize your feedback regarding our documentation. Provide as much detail as possible, so that your request can be quickly addressed.

Prerequisites

- You are logged in to the Red Hat Customer Portal.

Procedure

To provide feedback, perform the following steps:

1. Click the following link: [Create Issue](#).
2. Describe the issue or enhancement in the **Summary** text box.
3. Provide details about the issue or requested enhancement in the **Description** text box.
4. Type your name in the **Reporter** text box.
5. Click the **Create** button.

This action creates a documentation ticket and routes it to the appropriate documentation team. Thank you for taking the time to provide feedback.