



Cost Management Service 1-latest

Using cost models

Configuring cost models to reflect your cloud costs

Cost Management Service 1-latest Using cost models

Configuring cost models to reflect your cloud 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

This guide explains the cost model feature, configuration and related terminology. 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. USING COST MODELS AND ANALYZING YOUR USAGE	3
1.1. WHAT IS A COST MODEL?	3
1.2. COST MODEL CONCEPTS	3
1.2.1. Cost model terminology	3
1.3. THE COST MODEL WORKFLOW	4
1.4. ANALYZING UNUSED CLUSTER AND NODE CAPACITY	6
1.5. UNDERSTANDING COST DISTRIBUTION IN COST MANAGEMENT	7
1.5.1. Distributing costs	7
1.5.2. Calculating costs	8
1.5.3. Adding OpenShift projects	8
CHAPTER 2. SETTING UP A COST MODEL	9
2.1. APPLYING A MARKUP OR DISCOUNT TO CLOUD INTEGRATIONS	9
2.2. CREATING A COST MODEL FOR AN OPENSIFT CONTAINER PLATFORM CLUSTER	10
2.3. CALCULATING EFFECTIVE USAGE WITH COST MODELS	13
PROVIDING FEEDBACK ON RED HAT DOCUMENTATION	15

CHAPTER 1. USING COST MODELS AND ANALYZING YOUR USAGE

You can use cost models in cost management to apply a price to usage in your hybrid cloud environment, then distribute the costs to resources.

1.1. WHAT IS A COST MODEL?

It can be difficult to find out the real cost of cloud-based IT systems. Different integrations provide a variety of cost data and metrics, which can be complicated to calculate together and accurately distribute the costs.

A cost model is a framework used to define the calculations for the costs stored in cost management. Using a cost model, you can associate a price to metrics provided by your integrations to charge for utilization of resources.

In some cases, those costs are related to the raw costs of the infrastructure, while others will be based on a price list that associates usage to costs. The data must be normalized before you can add a markup to it to cover your overhead, and then distribute the charges to your resources or end customers. It helps to align costs to utilization: those customers that use a resource more will be charged more.

1.2. COST MODEL CONCEPTS

These terms and concepts are important for understanding cost management's cost model workflow.

Cost model

A cost model is a framework used to define the calculations for the costs stored in cost management, using raw costs and metrics. The cost model is used for budgeting and accounting, and also the visualization and analysis of costs – in cost management, it provides the basis for the cost information that the user will view. Costs generated by a cost model can be recorded, categorized and allocated to specific customers, business units or projects.

1.2.1. Cost model terminology

Markup

The portion of cost calculated by applying markup or discount to infrastructure raw cost in the cost management application.

Example: For a raw cost of \$100 and a markup of 10%, the markup would be \$10, and the cost would be \$110 (sum of both).

Request

The pod resources requested, as reported by OpenShift.

Usage

The pod resources used, as reported by OpenShift.

Usage cost

The portion of cost calculated by applying hourly or monthly price list rates to metrics.

Example: For a metric of 100 core-hours and a rate of \$1/core-hour, the usage cost would be \$100.


Effective usage

The pod resources used or requested each hour, whichever is higher.

Monthly cost

The portion of cost calculated by applying monthly price list rates to metrics returned as part of usage cost. Monthly cost can be configured for OpenShift nodes or clusters in a cost model to account for subscription costs, for example.

Monthly costs are currently shown as part of the usage costs in the cost management API and interface.



NOTE

These costs are not currently amortized over the reported period in the API, so when viewing a daily breakdown of cost, monthly costs will only show on the first of the month.

Example: For an OpenShift cluster with 10 nodes at a rate of \$10,000 per node per month, the monthly cost would be \$100,000.

Price list

A list of rates used within the application's cost model to calculate the usage cost of resources.

Distributed costs

The costs calculated by the cost model are distributed to higher level application concepts, such as project, account, service, and so on. How the costs are distributed depends on the cost type (infrastructure or supplementary) assigned in the cost model.

Count (cluster)

The distinct number of clusters identified during the month.

Count (nodes)

The distinct number of nodes identified during the month.


Count (persistent volume claims)

The distinct number of volume claims identified during the month.

1.3. THE COST MODEL WORKFLOW

The following diagram shows the cost model workflow that cost management uses to apply a price to metrics and inventory, normalize cost data from different integrations, apply a markup (or discount), then distribute the costs across the relevant resources.

The cost model also helps differentiate raw costs from the costs used in cost management.




85_OpenShift_0420

1. Cost management collects cost data from several integrations:

- Inventory** - All the resources that have ever run in your integration, including those that are no longer in use. For example, if your OpenShift Container Platform environment contains a node that is not in use, that node still costs \$x amount per month. There are several ways to collect inventory data into cost management: cost management can generate inventory from the AWS data export, Azure or Google Cloud export, or OpenShift Metering Operator reports.
- Metrics** - A subset of the OpenShift inventory showing usage and consumption for each resource.
- Cloud raw costs**- AWS, Azure, and Google Cloud provide regular reports to cost management listing consumption of resources and their costs, which cost management uses for calculations. As a result, configuring a custom price list is not necessary for cloud integrations.

2. The **cost model** allows you to apply a markup or discount of your choice to account for other costs and overhead, and provides options for assigning a cost layer (infrastructure or supplementary) to costs:
 - a. For OpenShift Container Platform integrations - Since the metrics and inventory data do not have a price assigned to usage, you must create and assign a **price list** to your integrations to determine the **usage cost** of these resources. A price list includes rates for storage, memory, CPU usage and requests, and clusters and nodes.
 - b. For AWS, Azure, and Google Cloud integrations - You can create cost models for these integrations to account for any extra costs or overhead in your environment by applying a **markup** percentage, or a negative percentage to calculate a discount.
3. Costs from integrations are then collected together and allocated as **infrastructure cost** and **supplementary cost**.
4. The cost is then **distributed** to resources across the environment. Depending on your organization, these resources may include OpenShift projects, nodes, or clusters, and cloud integration services or accounts. You can also use tagging to distribute the costs by team, project, or group within your organization.



NOTE

For more information about configuring tagging for cost management, see [Managing cost data using tagging](#).

1.4. ANALYZING UNUSED CLUSTER AND NODE CAPACITY

You can analyze your cluster usage with cost management by examining the unrequested capacity and unused capacity of your cluster. The unrequested capacity identifies how much of the requested resources are being used in the cluster. When this value is high, there are nodes in your cluster that are requesting more resources than it uses. You can find the nodes that are responsible and adjust your requests to make your cluster usage more efficient. However, the usage can exceed the amount requested, so the unused capacity can help you to understand if you should adjust the overall capacity.

Request

The pod resources requested, which OpenShift reports.

Unrequested capacity

is the requests minus the usage.

Effective usage

The pod resources used or requested each hour, whichever is higher.

Unused capacity

is the capacity minus the effective usage.

For more information, see [Section 1.2, "Cost model concepts"](#).

Prerequisites

- You created a integration on [Red Hat Hybrid Cloud Console](#).

Procedure

1. Log in to the [Red Hat Hybrid Cloud Console](#).

2. From the **Services** menu, click **Spend Management → Cost Management**.
3. To view unused cluster capacity and unrequested cluster capacity:
 - a. In the Global Navigation, click **Cost Management → OpenShift**.
 - b. On the OpenShift details page, in the **Group By** menu, select **Cluster**. Filter by **Cluster**, then select a cluster from the results.
 - c. On the Cost overview tab, you can see your unused capacity and unrequested capacity in the CPU tile. If the unrequested capacity is higher than your unused capacity, or the unrequested core-hours are too high as a percentage of capacity, you can search for the nodes in your cluster that are responsible.
4. To view unused node capacity and unrequested node capacity:
 - a. In the Global Navigation, click **Cost Management → OpenShift**.
 - b. On the OpenShift details page, in the **Group By** menu, select **Node**. Filter by **Node**, then select a node from the results.
5. If the unrequested capacity is higher than your unused capacity, or the unrequested core-hours are too high as a percentage of capacity for your node, you can adjust them in your cloud service provider to optimize your cloud spending.

1.5. UNDERSTANDING COST DISTRIBUTION IN COST MANAGEMENT

Costs can fall into two different groups:

Platform costs

Costs that incur from running the OpenShift Container Platform. Platform costs include the cost of all projects with the label **Default**. These namespaces and projects contain **openshift-** or **kube-** in the name. These projects were not created by the user but are required for OpenShift to run. You can optionally add namespaces and projects to platform costs. For more information, see [Section 1.5.3, "Adding OpenShift projects"](#).

Worker unallocated costs

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

1.5.1. Distributing costs

To configure the distribution of platform and worker unallocated costs into your projects, you can set costs to **Distribute** or **Do not distribute**. When you create a cost model, the costs are set to **Distribute** by default. This default setting means that the cost of the Platform projects are set to zero.

The costs distribute *into* your project costs according to the sum of the effective CPU or the Memory usage of your cost model. Most users use the default **Distribute** setting to track platform and worker unallocated costs for their organizations.

If you instead set the costs to **Do not distribute**, the costs of each Platform project are displayed individually instead of spread across all of the projects. The worker unallocated cost still is calculated, but it appears as an individual project in the **OpenShift details** page. With this option, you cannot see how the costs would distribute to user projects.

You can always distribute platform or worker unallocated costs independently of each other, or you can choose to distribute none of them.

1.5.2. Calculating costs

Cost management uses effective usage to calculate both platform and worker unallocated costs, in addition to project costs.

To distribute platform costs, cost management uses the following formula:

(individual user project effective usage) / (sum of usage for all user project's effective usage) * (platform cost)

To distribute worker unallocated costs, cost management uses the following formula:

(individual user project effective usage) / (sum of usage for all user project's effective usage) * (worker unallocated cost)

1.5.3. Adding OpenShift projects

In cost management, the **Group** named **Platform** has default projects that you cannot remove. These projects start with the prefixes **openshift** or **kube** and have a **Default** label in the **OpenShift details** page. You can add your own projects to the **Platform** group so that you have control over what is considered a platform cost. Any projects that have the cost of some Platform projects have the **Overhead** label.

For example, you might have a cost that you consider overhead and that you want to show up as a platform cost. You can add the cost to your Platform projects to distribute the costs according to your cost model.

Prerequisites

- You must have a cluster that has a cost model set to **Distribute**.

Procedure

To add OpenShift projects to the **Platform** group, complete the following steps:

1. In **Settings** in cost management, click the **Platform projects** tab.
2. Select a project to add to the **Platform** group.
3. Click **Add projects**.

The project now has the label **Platform**, but **not** the label **Default**.

Verification

Complete the following steps to verify that your costs are distributing properly:

1. In cost management, click **OpenShift** to open the **OpenShift Details** page.
2. Select the cluster whose project you edited in the previous steps.

The project should display a cost of **\$0** because you set the cost to distribute across all other projects. A project with the label **Overhead** includes the cost of that project plus the default project costs.

CHAPTER 2. SETTING UP A COST MODEL

Configure cost models in cost management based on your requirements.

2.1. APPLYING A MARKUP OR DISCOUNT TO CLOUD INTEGRATIONS

Create a AWS, Azure, or Google Cloud cost model to apply a markup or discount to your cloud infrastructure costs.

As cloud infrastructure integrations (AWS, Azure, or Google Cloud) collect cost and usage data into cost management with prices already assigned, you only need to assign a markup or discount (if desired) to accurately reflect your costs.

Adding a markup to your *raw costs* can allow you to account for your overhead costs, such as the cost of administering your AWS account, Azure subscription, or other support costs. A markup is an estimation to cover your costs not shown by metrics or usage.

This example shows how to add a 10% markup to the information collected from the AWS Cost and Usage Reports. The same method can be used to apply a markup or discount to your Azure or Google Cloud costs.



NOTE

Creating, editing, or deleting a cost model only updates calculations starting from the first day of the current month.

Prerequisites

- A user with Cost Administrator or Cost Price List Administrator permissions. See [Limiting access to cost management resources](#) in *Getting started with cost management* for instructions on configuring user roles.
- Your AWS account added to cost management as a data integration. See [Integrating Amazon Web Services \(AWS\) data into cost management](#) in *Getting started with cost management* for instructions.

Procedure

1. From [cost management](#), click **Cost Management** → **Settings**.
2. In the **Cost models** tab, click **Create cost model** to open the cost model wizard.
3. Enter a name and description for the cost model, and select **Amazon Web Services (AWS)** as the integration type to apply the cost model to. Click **Next**.
4. To apply a 10% markup to the base costs of your resources, enter *10* in the as the **Markup rate** and click **Next**.




NOTE

- To apply a discount rather than a markup, enter a minus sign with the value (for example, *-15*).
- If you do not want to apply a markup or discount, enter *0*.

5. Select one or multiple AWS integrations to assign the cost model to and click **Next**. Selecting a integration with a cost model already assigned will override its previous cost model. You can also assign integrations to your cost model at a later time.
6. Review the cost model details, and click **Create**.
7. Click **Close** to exit the cost model wizard.

Your new cost model will appear in the list on the **Cost models** page.

Next steps

- From the **Cost models** summary page, you can review information about your cost models, including:
 - The integration type the cost model was created for
 - The number of integrations assigned to the cost model
 - The date the cost model was last modified
- To see more information, assign, or edit a cost model:
 - From the **Cost models** summary page, click a cost model name to bring up the details page where you can edit the cost model, including integration assignments, markup, and other Settings.
- To delete a cost model:
 - From the **Cost models** summary page, click **More options** , then **Delete**.
- Review your tags and tagging strategy to ensure that costs are being distributed to the correct resources, cost centers, or teams. See [Managing cost data using tagging](#) for more information.

2.2. CREATING A COST MODEL FOR AN OPENSIFT CONTAINER PLATFORM CLUSTER

As metrics and inventory from OpenShift Container Platform integrations do not have costs associated, you must create a cost model to associate a price to resources.

Creating a cost model for an OpenShift integration includes assigning prices for usage and requests using **CPU, memory, node, cluster, storage, or persistent volume claim** metrics, and applying a markup or discount to determine the total costs of your OpenShift infrastructure. You can also use tags to measure costs for specific parts of your infrastructure, such as various storage methods, giving you greater control over your individual requirements. See [Managing cost data using tagging](#) for more information about tagging in cost management.

This example shows how to design and apply a cost model for an OpenShift Container Platform cluster on cloud infrastructure (for example, AWS or Azure). The cloud infrastructure costs are seen in cost management as part of the cluster cost, so creating a cost model allows you to distribute the underlying infrastructure cost to accurately reflect the costs of running the cluster.



NOTE

Creating, editing, or deleting a cost model only updates calculations starting from the first day of the current month.

Prerequisites

- A user with Cost Administrator or Cost Price List Administrator permissions. See [Limiting access to cost management resources](#) in *Getting started with cost management* for instructions on configuring user roles.
- 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.

Procedure

1. From [cost management](#), click **Cost Management → Settings**.
2. In the **Cost models** tab, click **Create cost model** to open the cost model wizard.
3. Enter a name and description for the cost model, and select **OpenShift Container Platform** as the integration type to apply the cost model to. Click **Next**.
4. Create a price list so you can assign rates to usage or requests. The cost management service collects these metrics from OpenShift but there is no cost attached to them in cost management until you apply a cost model. You can also create your price list later.
 - a. To apply a monthly node cost of \$1,000:
 - i. Click **Create rate**.
 - ii. For **Metric**, select **Node**.
 - iii. For **Measurement**, select **Count (node-month)**.
 - iv. The cost management service classifies node and cluster costs as *infrastructure costs* by default. To change the cost type you are calculating, select it under **Calculation type**. See [Key concepts](#) to learn more.
 - v. In the **Rate** field, enter *1000* to assign a \$1,000 per price to each node in the integrations you select (in a later step).
 - vi. Click **Create rate** to save the node rate.
 - b. To apply a rate to CPU requests:
 - i. Click **Create rate**.
 - ii. For **Metric**, select **CPU**.
 - iii. For **Measurement**, select **Request (core-hours)**.
 - iv. The cost management service classifies metrics-based data such as CPU requests as *supplementary costs* by default. To change the cost type you are calculating, select it under **Calculation type**. See [Key concepts](#) to learn more.


- v. In the **Rate** field, enter *0.09* to apply a \$0.09 cost to each core-hour of CPU requests.
 - vi. Click **Create rate** to save the CPU request rate.
- c. To apply tag-based rates to persistent volume claims:
- i. Click **Create rate**.
 - ii. For **Metric**, select **Persistent volume claims**.
 - iii. For **Measurement**, select **Count (pvc-month)**.
 - iv. The cost management service classifies persistent volume claim rates as *infrastructure* costs by default. To change the cost type you are calculating, select it under **Calculation type**. See [Key concepts](#) to learn more.
 - v. Select **Enter rate by tag**.
 - vi. Enter a tag key in the **Filter by tag key** field.
 - vii. Enter a **Tag value** and **Rate**. You can select **Default** to specify a rate as default for that tag key.



NOTE

- Specifying a default rate will apply that rate to all tag values for the corresponding tag key that are not defined. You can enter a rate of *0* for any tag value that you do not want the default rate applied to.

- viii. Click **Add more tag values** to add as many tag values as you require.
 - ix. Click **Create rate** to save the persistent volume claim rates.
- d. Add additional rates by repeating the previous steps, or click **Next** to finish configuring your price list.
5. If desired, set a markup or discount. To apply a 10% markup to the base costs of your resources, enter *10* in the as the **Markup rate** and click **Next**.



NOTE

- To apply a discount rather than a markup, enter a minus sign with the value (for example, *-15*).
- If you do not want to apply a markup or discount, enter *0*.


6. On the **Cost distribution** page, select the **CPU** or **Memory** distribution type. The distribution type distributes costs based on CPU or memory metrics in project cost breakdowns. If your cluster has high memory usage, select **Memory**. If your cluster has high CPU usage, select **CPU**.
7. On the **Cost distribution** page, select whether or not to distribute your **Platform** or **Worker unallocated** costs into your projects and click **Next**.
8. Select one or multiple OpenShift integrations to assign the cost model to and click **Next**. Selecting a integration with a cost model already assigned will override its previous cost model. You can also assign integrations to your cost model later.

9. Review the cost model details, and click **Create**.

10. Click **Close** to exit the cost model wizard.

Your new cost model will appear in the list on the **Cost models** page.

Next steps

- From the **Cost models** summary page, you can review information about your cost models, including:
 - The integration type the cost model was created for
 - The number of integrations assigned to the cost model
 - The date the cost model was last modified
- To see more information, assign, or edit a cost model:
 - From the **Cost models** summary page, click a cost model name to bring up the details page where you can edit the cost model, including integration assignments, price list, markup, and other Settings.
- To delete a cost model:
 - From the **Cost models** summary page, click **More options**  , then **Delete**.
- Review your tags and tagging strategy to ensure that costs are being distributed to the correct resources, cost centers, or teams. See [Managing cost data using tagging](#) for more information.

2.3. CALCULATING EFFECTIVE USAGE WITH COST MODELS

Cloud providers charge for the infrastructure costs of running a cluster, regardless of your overall usage. By calculating the effective usage in cost management, you can more accurately correlate cloud costs with a pod or namespace by considering their direct utilization.

A pod typically requests resources, such as CPU or memory, from a cluster. The cluster then reserves these requested resources as a minimum, but the pod might use more or less than it initially requested. The effective usage metric in cost management uses whichever kind of usage, CPU or memory, is greater per hour.

You can create a cost model in cost management to estimate your effective usage. Ultimately, you can use this data to understand how infrastructure cost is distributed to your OpenShift project.

Prerequisites

- You must be a user with Cost Administrator or Cost Price List Administrator permissions. To learn how to configure user roles, see [Limiting access to cost management resources](#) in *Getting started with cost management*
- You must add your OpenShift cluster as a cost management data integration. For more details, see [Integrating OpenShift Container Platform data into cost management](#) in *Getting started with cost management*.

Procedure

1. Log in to [Red Hat Hybrid Cloud Console](#).
2. From the **Services** menu, click **Spend Management** → **Cost Management**.
3. In the Global Navigation, click **Cost Management** → **Settings**.
4. In the **Cost Models** tab, click **Create cost model** to open the cost model wizard.
5. Enter a name and description for the cost model and select **OpenShift Container Platform** as the integration type. Click **Next**.
6. Create a price list so that you can assign rates to usage or requests. The cost management service collects these metrics from OpenShift but there is no cost attached to them in cost management until you apply a cost model.
 - a. To create a price list to calculate effective CPU usage, click **Create rate**.
 - i. Add a description. In this example, enter **effective cpu usage**.
 - ii. In the **Metric** field, select **CPU**.
 - iii. In the **Measurement** field, select **Effective-usage (core-hours)**.
 - iv. In the **Rate** field, enter the rate you pay for CPU usage. In this example, enter **2**. Click **Create rate**.
 - b. To create a price list to calculate effective memory usage, click **Create rate**.
 - i. Add a description. In this example, enter **effective memory usage**.
 - ii. In the **Metric** field, select **Memory**.
 - iii. In the **Measurement** field, select **Effective-usage (GB-hours)**.
 - iv. In the **Rate** field, enter the rate you pay for memory usage. In this example, enter **1**. Click **Create rate**.
7. Click **Next**.
8. (Optional) On the **Cost calculations** page, apply a markup or discount to change how raw costs are calculated for your integrations. Adding a markup to your *raw costs* can allow you to account for your overhead costs, such as the cost of administering your AWS account, Azure subscription, or other support costs. A markup is an estimation to cover your costs not shown by metrics or usage.
9. On the **Cost distribution** page, select the **CPU** or **Memory** distribution type. The distribution type distributes costs based on CPU or memory metrics in project cost breakdowns. If your cluster has high memory usage, select **Memory**. If your cluster has high CPU usage, select **CPU**. Click **Next**.
10. Assign an integration to your cost model and then click **Next**.
11. Review the details and then click **Create**.
12. To review the results of your cost model on a integration, in the Global Navigation, click **Cost Management** → **OpenShift**.
13. Select a project and view the results.

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!