



# Cost Management Service 1-latest

## Getting started with resource optimization for OpenShift

Learn about resource optimization for OpenShift



# Cost Management Service 1-latest Getting started with resource optimization for OpenShift

---

Learn about resource optimization for OpenShift

## 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<sup>®</sup> is the registered trademark of Linus Torvalds in the United States and other countries.

Java<sup>®</sup> is a registered trademark of Oracle and/or its affiliates.

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

MySQL<sup>®</sup> is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js<sup>®</sup> 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<sup>®</sup> 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 you can use resource optimization for OpenShift to optimize the size of containers and pods within projects in your OpenShift clusters and improve the sizing and capacity management of nodes within the cluster.

---

## Table of Contents

<b>PREFACE</b> .....	<b>3</b>
<b>CHAPTER 1. WHAT IS RESOURCE OPTIMIZATION FOR OPENSIFT?</b> .....	<b>4</b>
<b>CHAPTER 2. PREREQUISITES</b> .....	<b>5</b>
<b>CHAPTER 3. RESOURCE OPTIMIZATION FOR OPENSIFT OPTIMIZATION REPORTS</b> .....	<b>6</b>
3.1. ENABLING OPTIMIZATION	6
3.2. VIEWING OPTIMIZATION REPORTS	6
3.3. OPTIMIZING FOR COST OR FOR PERFORMANCE	6
3.4. UNDERSTANDING BOX PLOTS	6



## PREFACE

As part of cost management, resource optimization for OpenShift assesses and monitors your usage across clusters to optimize your Red Hat OpenShift resources.

## CHAPTER 1. WHAT IS RESOURCE OPTIMIZATION FOR OPENSIFT?

Resource optimization for OpenShift uses current and historical data from OpenShift to recommend actions to take:

- Shows metrics for CPU and memory usage and analyzes them
- Compares defined container requests and limits
- Analyzes the historical usage patterns to return optimization recommendations
- Reports usage of applications and deployments
- Optimizes the size of your pods
- Manages costs

The data that resource optimization for OpenShift provides can improve your resource allocation and help you save money on your OpenShift cluster deployment.



## CHAPTER 2. PREREQUISITES

To use resource optimization for OpenShift, you must have completed the following steps:

- You installed the Cost Management Metrics Operator version 3.3.0 or newer.
- You added your OpenShift integration to cost management. For more information, see [Integrating OpenShift Container Platform data into cost management](#).

## CHAPTER 3. RESOURCE OPTIMIZATION FOR OPENSIFT

### OPTIMIZATION REPORTS

Access resource optimization for OpenShift from the [Red Hat Hybrid Cloud Console](#) to see detailed recommendations for how to optimize your Red Hat OpenShift clusters.

#### 3.1. ENABLING OPTIMIZATION

To receive resource optimization recommendations for your namespaces, you must first enable each namespace. To enable a namespace, label it with `insights_cost_management_optimizations='true'`.

In the CLI, run: `oc label namespace NAMESPACE insights_cost_management_optimizations="true" --overwrite=true`

#### 3.2. VIEWING OPTIMIZATION REPORTS

##### Prerequisites

- You added an OpenShift integration to [Red Hat Hybrid Cloud Console](#).
- You uploaded at least 24 hours of data from the operator.
- You logged in to the [Red Hat Hybrid Cloud Console](#).

##### Procedure

1. In cost management, click the tab [Optimizations](#).
2. Search for an optimization or use the filter. Click the link to the optimization that you selected.
3. View details about the recommendation and toggle between **Cost optimizations** and **Performance optimizations**. For more information, see [Optimizing for cost or for performance](#).

#### 3.3. OPTIMIZING FOR COST OR FOR PERFORMANCE

After you select an optimization, you can toggle between two tabs called **Cost optimizations** and **Performance optimizations**. Optimizing for cost uses less resources and is useful when you are performing tests where there is no impact to users. Optimizing for performance provides all the resources possible and is helpful for apps running in a production cluster.

In **Cost optimizations**, recommendations get generated when CPU usage is at or above the 60th percentile and memory usage is at the 100th percentile. In **Performance optimizations**, recommendations get generated when CPU usage is at or above the 98th percentile and when memory usage is at the 100th percentile.

#### 3.4. UNDERSTANDING BOX PLOTS

On the [Optimizations](#) page, there are two box plots for your **Current CPU utilization** and your **Current memory utilization**. These visualizations can help you understand resource distribution and identify outliers in your data. You can export the data in CSV and JSON format.

The box plots display the following data points in millicpu (m) for CPU and in mebibytes (Mi) for memory:

- Minimum
- Maximum
- Median
- First quartile (Q1): value where 25% of data points are when they are arranged in increasing order
- Third quartile (Q3): value where 75% of data points are when they are arranged in increasing order
- Recommended limit
- Recommended request

The data points are based on the time period that you selected:

- 1 day = 4 data points
- 7 days = 7 data points
- 14 days = 14 data points