



OpenShift Container Platform 4.6

Getting started with cost management

Learn about and configure cost management

OpenShift Container Platform 4.6 Getting started with cost management

Learn about and configure cost management

Legal Notice

Copyright © 2021 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 describes the initial steps to begin using cost management.

Table of Contents

CHAPTER 1. INTRODUCTION TO COST MANAGEMENT	3
1.1. ABOUT COST MANAGEMENT	3
1.1.1. Terminology	3
1.2. PLANNING FOR COST MANAGEMENT	3
CHAPTER 2. LIMITING ACCESS TO COST MANAGEMENT RESOURCES	5
2.1. DEFAULT USER ROLES IN COST MANAGEMENT	5
2.2. ADDING A ROLE	5
2.3. ADDING A ROLE TO A GROUP	7
CHAPTER 3. ADDING AN OPENSIFT CONTAINER PLATFORM SOURCE TO COST MANAGEMENT	9
3.1. INSTALLATION TASKS SUMMARY	9
3.2. INSTALLING THE KOKU METRICS OPERATOR	9
3.3. CONFIGURING THE OPERATOR INSTANCE FOR A NEW INSTALLATION	10
3.4. REPLACING THE PRIOR OPERATOR INSTANCE WITH THE KOKU METRICS OPERATOR	11
3.5. REMOVING A PRIOR COST OPERATOR	12
3.6. VERIFYING KOKU-METRICS-OPERATOR	13
3.7. CONFIGURING BASIC AUTHENTICATION FOR COST OPERATOR	14
3.7.1. Creating the secret key/value pair for basic authentication	14
3.7.2. Modifying the YAML file	15
3.8. CREATING AN OPENSIFT CONTAINER PLATFORM SOURCE MANUALLY	16
3.9. ADDING A RESTRICTED NETWORK SOURCE	17
3.9.1. Installing the Koku Metrics Operator on a restricted network	17
3.9.2. Configuring Koku Metrics Operator on a restricted network	18
3.9.3. Downloading cost reports	19
3.9.4. Uploading cost reports to cloud.redhat.com	20
CHAPTER 4. ADDING AN AMAZON WEB SERVICES (AWS) SOURCE TO COST MANAGEMENT	22
4.1. CREATING AN S3 BUCKET FOR REPORTING	22
4.2. ACTIVATING AWS TAGS FOR COST MANAGEMENT	23
4.3. ENABLING MINIMAL ACCOUNT ACCESS FOR COST AND USAGE CONSUMPTION	23
4.3.1. Enabling additional account access for cost and usage consumption	25
CHAPTER 5. ADDING A MICROSOFT AZURE SOURCE TO COST MANAGEMENT	27
5.1. CREATING AN AZURE RESOURCE GROUP AND STORAGE ACCOUNT	27
5.2. CONFIGURING AZURE ROLES	28
5.3. CONFIGURING A DAILY AZURE DATA EXPORT SCHEDULE	28
CHAPTER 6. ADDING AN OPENSIFT CONTAINER PLATFORM 4.3 AND 4.4 SOURCE TO COST MANAGEMENT	30
6.1. INSTALLING THE COST MANAGEMENT OPERATOR	30
6.2. CONFIGURING THE COST MANAGEMENT OPERATOR	31
CHAPTER 7. NEXT STEPS FOR MANAGING YOUR COSTS	35
7.1. CONFIGURE TAGGING FOR YOUR SOURCES	35
7.2. CONFIGURE COST MODELS TO ACCURATELY REPORT COSTS	35

CHAPTER 1. INTRODUCTION TO COST MANAGEMENT

This document provides instructions to begin using cost management, including prerequisites and instructions for connecting your cloud environments, and configuring users and permissions.

After completing the setup described in this guide, you will be able to track cost and usage data for your Amazon Web Services (AWS), Microsoft Azure, and OpenShift Container Platform environments.

If you have a suggestion for improving this guide or have found an error, please submit a Bugzilla report at <http://bugzilla.redhat.com> against *Cloud Software Services (cloud.redhat.com)* for the **Cost Management** component.

1.1. ABOUT COST MANAGEMENT

Cost management is an OpenShift Container Platform service that enables you to better understand and track costs for clouds and containers. It is based on the upstream project Koku.

You can access the cost management application from <https://cloud.redhat.com/cost-management/>.

Cost management allows you to simplify management of resources and costs across various environments, including:

- Container platforms such as OpenShift Container Platform
- Public clouds such as Amazon Web Services (AWS) and Microsoft Azure

The cost management application allows you to:

- Visualize, understand and analyze the use of resources and costs
- Forecast your future consumption and compare them with budgets
- Optimize resources and consumption
- Identify patterns of usage that should be investigated
- Integrate with third party tools that can benefit from cost and resourcing data

1.1.1. Terminology

Source

A cloud provider account that is connected to cost management to be monitored, for example, an OpenShift Container Platform deployment, or an AWS or Azure account.

Organization Administrator

The highest permission level for Red Hat accounts, with full access to content and features. This is the only role that can manage users and control their access and permissions on an account. An account may have multiple Organization Administrators.

See [Roles and Permissions for Red Hat Subscription Management](#) for more details.

1.2. PLANNING FOR COST MANAGEMENT

When configuring cost management for your needs, consider the scope of your environments that you want to manage costs for, and the users who will have access to the data.

Some considerations in creating a new Red Hat organization and users for different customer types include:

Scope:

- Customer company wide
- Customer division or organization wide
- Partner company managing several tenants

Data:

- How does your business need the data? Do you want information about projects, or users, for example?
- Planning AWS tags to reflect previous use cases.
- Enforcement: Is there any way for you to ensure that the proper tags and metadata are included in each item of the inventory?

User access:

- What level of access do you want your users to have?
- Do you want some users to have access to all cost data, while other users can view only a portion of the environment or certain sources?

CHAPTER 2. LIMITING ACCESS TO COST MANAGEMENT RESOURCES

You may not want users to have access to all cost data, but instead only data specific to their projects or organization. Using role-based access control, you can limit the visibility of resources involved in cost management reports. For example, you may want to restrict a user's view to only AWS sources, instead of the entire environment.

Role-based access control works by organizing users into groups, which can be associated with one or more roles. A role defines a permission and a set of resource definitions.

By default, a user who is not an account administrator will not have access to data, but instead must be granted access to resources. Account administrators can view all data without any further role-based access control configuration.



NOTE

A Red Hat account user with Organization Administrator entitlements is required to configure Red Hat account users. This Red Hat login allows you to look up users, add them to groups, and to assign roles that control visibility to resources.

For more information about Red Hat account roles, see [Roles and Permissions for Red Hat Subscription Management](#) and [How To Create and Manage Users](#).

2.1. DEFAULT USER ROLES IN COST MANAGEMENT

You can configure custom user access roles for cost management, or assign each user a predefined role.

To use a default role, determine the required level of access to permit your users based on the following predefined roles in cost management:

Administrator roles

- Cost Administrator: has read and write permissions to all resources in cost management
- Cost Price List Administrator: has read and write permissions on price list rates

Viewer roles

- Cost Cloud Viewer: has read permissions on cost reports related to cloud sources
- Cost OpenShift Viewer: has read permissions on cost reports related to OpenShift sources
- Cost Price List Viewer: has read permissions on price list rates


2.2. ADDING A ROLE

Create a new role to manage and limit the scope of information that users can see within cost management.

Prerequisites

- You must be an Account Administrator or a member of a group with the *RBAC Administrator* role to create a role.

Procedure

1. From cost management, click  (**Settings**) to navigate to **User Access**.
2. Click the **Roles** tab.
3. Click **Create Role** to open the **Add role** wizard.
4. In the **Name and Description** screen, enter a name for the new role, and optionally, a description. Click **Next**.
5. In the **Permission** screen, specify the Red Hat Cloud Services application you are creating the role for (in this case, cost management) as well as the resource and permission type:
 - a. For **Application**, enter *cost-management*.
 - b. For **Resource type**, specify the resource this permission will be used to access from the following list:
 - *aws.account*
 - *aws.organizational_unit*
 - *azure.subscription_guid*
 - *openshift.cluster*
 - *openshift.node*
 - *openshift.project*

NOTE

When you add an AWS organizational unit as a **Resource Type**, any user who has access to the parent node also has access to all children and sub-children of the parent node.

- c. For **Permission**, specify *read* as all cost resource data is read-only.
For example, to create a role with read-only permissions to AWS account data, set *aws.account* as the **Resource type** and *read* as the **Permission**. In the next step, you can specify the AWS account to apply this role to.
6. In the **Resource definitions** screen, you can provide more details about the resources the permission will be used for. For example, to grant this role access to a specific AWS account, enter the following and click **Add to definitions**:
 - **Key:** *aws.account*
 - Options for **Key** are: *aws.account*, *aws.organizational_unit*, *azure.subscription_guid*, *openshift.cluster*, *openshift.node*, *openshift.project*
 - **Operation:** *equal*
 - Use *equal* if you know the exact value, or *list* to see a list of values that will work for this role.

- **Value:** Your AWS account number or account alias.
 - This is specific to the resource defined in the **Key** field. Examples include the AWS account ID or alias, AWS organizational unit, Azure subscription ID, OpenShift cluster ID, OpenShift node name, or OpenShift project name.
 - You can also enter * in this field as a wildcard to create a role that matches everything of the resource type defined in **Key**.
7. Add more resource definitions if desired and click **Next** when finished.
 8. Review the details for this role and click **Confirm** to create the role.

Your new role will be listed in the **Roles** tab on the **User Access Management** screen.

Next steps

- Add this role to a group to provide the role with access to resources.


2.3. ADDING A ROLE TO A GROUP

Add your role to a group to manage and limit the scope of information that users in that group can see within cost management.

Prerequisites

- You must be an Account Administrator or a member of a group with the *RBAC Administrator* role to create a role.

Procedure

1. From cost management, click  (**Settings**) to navigate to **User Access**.
2. Click the **Groups** tab.
3. Click **Create group**.
4. In the **General information** screen, enter a name for the new group, and optionally, a description. Click **Next**.
5. In the **Add members** screen, select the user(s) in your organization to add to the new group. Click **Next**.
6. (Optional) In the **Select roles** screen, select one or more role(s) to add to the group. Default roles available for cost management are:
 - *Cost Administrator* : grants read and write permissions
 - *Cost Cloud Viewer* : grants read permissions on cost reports related to cloud sources
 - *Cost OpenShift Viewer* : grants read permissions on cost reports related to OpenShift sources
 - *Cost Price List Administrator* : grants read and write permissions on price list rates
7. Review the details for this group and click **Confirm** to create the group.

Your new group will be listed in the **Groups** list on the **User Access** screen.

To verify your configuration, log out of the cost management application and log back in as a user added to the group.

CHAPTER 3. ADDING AN OPENSIFT CONTAINER PLATFORM SOURCE TO COST MANAGEMENT

To support cost management in Red Hat OpenShift Container Platform 4.5 and more recent versions, a new community operator is introduced, **koku-metrics-operator**.

The community **cost-mgmt-operator** supported in OpenShift Container Platform 4.3 and 4.4 is deprecated. Though this cost operator functions with cost management in OpenShift Container Platform 4.5 or later, it is no longer supported. For installing this deprecated operator, see [Chapter 6, Adding an OpenShift Container Platform 4.3 and 4.4 source to cost management](#) .


To avoid gaps in your cost management data, you can wait 24 to 48 hours before removing **cost-mgmt-operator** while you verify that **koku-metrics-operator** provides cost management reports.

3.1. INSTALLATION TASKS SUMMARY

Whether you are replacing a prior cost management operator with the Koku metrics operator or installing it for the first time, the basic tasks are the same.

Operator installation, configuration, and source management can all be performed from the OpenShift Container Platform web console.

You will perform the following tasks to install the **koku-metrics-operator** and begin using the cost management application in OpenShift Container Platform 4.5 or later.



NOTE

To install and configure **koku-metrics-operator** from the OpenShift Container Platform web console, you must use an account with cluster administrator privileges.

Prerequisites

- The OpenShift Container Platform cluster is installed.
- You can access the OpenShift Container Platform web console using an account that has cluster administrator privileges.

Task summary

- Install the Koku metrics operator (**koku-metrics-operator**) and use the default token authentication.
- Create a KokuMetricsConfig YAML file that configures **koku-metrics-operator**.
- Create a cost management OpenShift Container Platform source with a new installation, or confirm an existing source with an replacement installation.
- After installing the Koku metrics operator, delete the old cost operator. This task is required only when the **cost-mgmt-operator** was previously installed.

If you use Basic authentication, additional steps are required to configure the Secret that holds username and password credentials.

3.2. INSTALLING THE KOKU METRICS OPERATOR

Install the Koku operator **koku-metrics-operator** from the OpenShift Container Platform web console.

Prerequisites

- You are logged into the OpenShift Container Platform web console with cluster administrator privileges.

Procedure

1. Login to the OpenShift Container Platform web console and click on the **Operators > OperatorHub** tab.
2. Search for and locate **Koku Metrics Operator**.
3. Click on the displayed **Koku Metrics Operator** tile.
4. If an information panel appears with the Community Operators message, click **Continue**.
5. When the **Install Operator** window appears, you must select the **koku-metrics-operator** namespace for installation. If the namespace does not yet exist, it is created for you.
6. Click on the **Install** button.
7. After a short wait, **Koku Metrics Operator** appears in the **Installed Operators** tab under **Project: all projects** or **Project: koku-metrics-operator**.

3.3. CONFIGURING THE OPERATOR INSTANCE FOR A NEW INSTALLATION

Use the OpenShift Container Platform web console to configure the **koku-metrics-operator** instance after it is installed.

Prerequisites

- You are logged into the OpenShift Container Platform web console with cluster administrator privileges.
- The **Koku Metrics Operator** appears in the **Installed Operators** tab.

Procedure

1. Under the **Name** heading in the list of installed operators, click the **KokuMetricsOperator** link. The **Installed Operators > Operator Details** window appears for **Koku Metrics Operator**.
2. In the **Details** window, click **+ Create Instance**.
3. A **Koku Metrics Operator > Create KokuMetricsConfig** window appears.
4. Click the **YAML view** radio button to view and modify the contents of the YAML configuration file.
5. When you create a new cost management instance for the Koku metrics operator, make the following modifications in the YAML configuration file.
6. Locate the following two lines in the YAML file.

```
create_source: false
name: INSERT-SOURCE-NAME
```

- a. Change **false** to **true**.
- b. Change **INSERT-SOURCE-NAME** to the new name of your source.


Example

```
create_source: true
name: koku-cost-source
```

7. Click the **Create** button. These actions create a new source for cost management that will appear in the cloud.redhat.com cost management application.

3.4. REPLACING THE PRIOR OPERATOR INSTANCE WITH THE KOKU METRICS OPERATOR

If you are replacing a prior cost management operator with the Koku metrics operator, make certain your existing cost management source is properly configured in the YAML configuration file.




IMPORTANT

When you are replacing a prior cost management operator with the **Koku Metrics Operator** and you want to use an existing source, you must make certain that the **name: INSERT-SOURCE-NAME** in the YAML file matches your existing source.

Prerequisites

- You are logged into the OpenShift Container Platform web console with cluster administrator privileges.
- You can access cloud.redhat.com and view existing cost management sources.

Procedure

1. Under the **Name** heading in the list of installed operators, click the **KokuMetricsOperator** link. The **Installed Operators > Operator Details** window appears for **Koku Metrics Operator**.
2. In the **Details** window, click **+ Create Instance**.
3. A **Koku Metrics Operator > Create KokuMetricsConfig** window appears.
4. Click the **YAML view** radio button to view and modify the contents of the **KokuMetricsConfig.yaml** file.
5. Open cloud.redhat.com and log in using your Organization Administrator account.
6. Click on  (Settings).
7. Click on the **Sources** tab to display existing sources.
8. Select an existing cost management source and copy its name.

- In the **KokuMetricsConfig.yaml** file, replace **INSERT-SOURCE-NAME** with the source name that you copied from the cost management source list for your organization.

```
create_source: false
name: INSERT-SOURCE-NAME <<<< replace this string
```

The **create_source: false** does not change because you are matching an existing source, not creating a new source.

- Click the **Create** button. No further actions are needed to configure the operator instance.

3.5. REMOVING A PRIOR COST OPERATOR

After installing the **koku-metrics-operator**, uninstall the prior **cost-mgmt-operator** operator.

To avoid gaps in your cost management data, you can wait 24 to 48 hours before removing **cost-mgmt-operator** while you verify that **koku-metrics-operator** provides cost management reports.



NOTE

If you mistakenly remove the **Koku Metrics Operator**, reinstall it.

Prerequisites

- The **cost-mgmt-operator** is installed.





NOTE

The **cost-mgmt-operator** is deprecated in OpenShift Container Platform 4.5 and later. Beginning with OpenShift Container Platform 4.5 or later, you must install **koku-metrics-operator**.

- You installed the **Koku Metrics Operator**.
- You are logged into the OpenShift Container Platform web console with cluster administrator privileges.
- You can view the operators in the **Installed Operators** tab.

Procedure

- In the Installed Operators list, select the operator you want to remove.
- Click on the  (More options) icon in that row.
- Click on the **Uninstall Operator** option. Confirm the action to remove the operator.
- In the OpenShift Container Platform web console, click the **Administration > Custom Resource Definitions** tab.
- In the window that displays the custom resource definitions (CRD), locate the **CostManagement** CRD and the **CostManagementData** CRD.

6. For each CRD, click on the  (More options) icon and click on **Delete Custom Resource Definition**. Confirm the delete action.
7. When these CRDs are deleted, the **cost-mgmt-operator** is fully uninstalled.



NOTE

When you install **Koku Metrics Operator**, a **KokuMetricsConfig** CRD appears in **Administration > Custom Resource Definitions** list.

3.6. VERIFYING KOKU-METRICS-OPERATOR

View the Koku configuration YAML file to verify the cost management operator is functioning.

Prerequisites

- You can access the OpenShift Container Platform web console.
- You can locate and view the **Installed Operators** tab.

Procedure

1. Click on the **Installed Operators** tab.
2. In the list of installed operators, click on the **Koku Metrics Operator** entry.
3. When the metrics operator window opens, click on the **KokuMetricsConfig** tab to show a list of the configuration file names.
4. In the name list, click on the configuration file. In the default installation, the file name is **kokumetricscfg-sample**.
5. When the Details window opens, click on the **YAML** tab. and visually check the following items.
 - a. Prometheus configuration and connection are **true**.

```
prometheus:
  last_query_start_time: '2021-01-25T20:59:06Z'
  last_query_success_time: '2021-01-25T20:59:06Z'
  prometheus_configured: true
  prometheus_connected: true
  service_address: 'https://thanos-querier.openshift-monitoring.svc:9091'
  skip_tls_verification: false
```

- b. Upload information shows the ingress path, successful upload and time, and accepted status.

```
upload:
  ingress_path: /api/ingress/v1/upload
  last_successful_upload_time: '2021-01-25T20:59:35Z'
  last_upload_status: 202 Accepted
  last_upload_time: '2021-01-25T20:59:35Z'
  upload: true
```

```
upload_cycle: 360
upload_wait: 28
validate_cert: true
```

3.7. CONFIGURING BASIC AUTHENTICATION FOR COST OPERATOR

You can configure the cost operator to use basic authentication. By default, the cost operator uses token authentication.

There are two procedures required when you configure basic authentication.

- [Creating the secret key/value pair for basic authentication](#)
- [Modifying the YAML file](#)

3.7.1. Creating the secret key/value pair for basic authentication

Prerequisites

- You are logged into the OpenShift Container Platform web console with cluster administrator privileges.
- The **Koku Metrics Operator** appears in the **Installed Operators** tab.
- You have a username and password for your cloud.redhat.com account.

Procedure

This procedure describes setting up basic authentication using the OpenShift Container Platform web console.

1. In the OpenShift Container Platform web console, click on the **Workloads > Secrets** tab.
2. In the **Secrets** window, select **Project:koku-metrics-operator** from the dropdown.
3. Click the **Create > Key/Value Secret** selection.
4. In the **Create Key/Value Secret** window enter the following information to create a new secret that contains a username key and a password key and a value for each key.
 - a. Enter a name for your secret in the **Secret Name** field.

Secret name example

```
basic-auth-secret
```

- b. In the **Key** field, enter **username**. This is your first key of the key pair.

Key name example for username

```
username
```

- c. In the **Value** field for the username key, enter the actual username for your authorized cloud.redhat.com user account.

Key value example for username

```
RedHatUser
```

- d. Click the **Add Key/Value** link to add the required password key name and value.
- e. In the **Key** field, enter **password**. This is your second key of the key pair.

Key name example for password

```
password
```

- f. In the **Value** field for the password key, enter the actual password for your authorized cloud.redhat.com user account.

Key value example for password

```
my.User!password
```

- g. Click the **Create** button to complete the creation of your basic authorization secret.
- h. After you click the **Create** button, you can verify the key information details for the secret.

**NOTE**

Do *not* add the secret to the workload.

3.7.2. Modifying the YAML file

Modify the Koku metrics operator API YAML file to use basic authentication with a secret username and password key/value pair.

Prerequisites

- You are logged into the OpenShift Container Platform web console with cluster administrator privileges.
- You created a secret name for the username and password key/value pair.
- The Koku metrics operator is installed.

Procedure

1. Click on the **Operators > Installed Operators** tab.
2. Locate the row that contains **KokuMetricsOperator** and click on the **KokuMetricsConfig** link that is under the **Provided APIs** heading.
3. When the **KokuMetricsConfigs** window appears, click on Koku configuration file listed in the **Name** column.
The default name is **kokumetricscfg-sample**.
4. When the **kokumetricscfg-sample** window appears, click in the **YAML** tab to open an edit and view window.

5. Locate the following lines in the YAML view.

```
authentication:  
  type: token
```

6. Change **type: token** to **type: basic**.
7. Insert a new line for **secret_name**. Enter the value for **secret_name**, which is the name you previously created.

Example

```
authentication:  
  secret_name: basic-auth-secret  
  type: basic
```

8. Click the **Save** button. A confirmation message appears.


3.8. CREATING AN OPENSIFT CONTAINER PLATFORM SOURCE MANUALLY

If you follow the previous steps, your OpenShift Container Platform source should be created automatically. However, there are situations, such as restricted network installations, when an OpenShift Container Platform source must be created manually on cloud.redhat.com.

Prerequisites

- OpenShift Container Platform cluster installed.
- Red Hat account user with Organization Administrator entitlements.
- You are logged into the OpenShift Container Platform web console.

Procedure

1. From cost management, click  (**Settings**).
2. Click **Sources**.
3. Click **Red Hat Sources**.
4. Click **Add source** to open the dialog.
5. Enter a name for the source and click **Next**.
6. Select the **Red Hat OpenShift Container Platform** tile as the source type.
7. Select **Cost Management** as the application and click **Next**.
8. Copy your **Cluster Identifier** from the OpenShift Container Platform web console **Home > Overview** tab and click **Next**.
9. Review the details and click **Add** to create the source.

3.9. ADDING A RESTRICTED NETWORK SOURCE

You can install OpenShift Container Platform on a restricted network that does not have access to the internet.

The procedure to add an OpenShift Container Platform cluster operating on a restricted network as a cost management source is different in the following ways:

1. Operator Lifecycle Manager is configured to install and run local sources.
2. The **koku-metrics-operator** is configured to store cost report CSV files locally using a persistent volume claim (PVC).
3. Cost reports stored in the PVC are downloaded to a workstation.
4. An OpenShift Container Platform source is created manually.
5. Cost reports are uploaded to cloud.redhat.com from your workstation.

3.9.1. Installing the Koku Metrics Operator on a restricted network

For OpenShift Container Platform clusters that are installed on restricted networks, Operator Lifecycle Manager (OLM) by default cannot access the koku-metrics-operator hosted remotely because those remote sources require full Internet connectivity. Therefore, OLM must be configured to install and run local sources.

Prerequisites

- OpenShift Container Platform cluster installed.
- Workstation with unrestricted network access.
- You are logged into the OpenShift Container Platform web console with cluster administrator privileges.

Procedure

1. Complete the following OpenShift Container Platform procedure to create a local mirror of the **koku-metrics-operator**: [Using Operator Lifecycle Manager on restricted networks](#)



NOTE

The **koku-metrics-operator** is found in the **community-operators** Catalog in the **registry.redhat.io/redhat/community-operator-index:latest** index.

Red Hat recommends pruning unwanted objects from the index before pushing to the mirrored registry. Make sure you keep the **koku-metrics-operator** package.

2. Log in to the OpenShift Container Platform web console and click **Operators > OperatorHub**.
3. Search for and locate **Koku Metrics Operator**.
4. Click the **Koku Metrics Operator** tile.

5. If an information panel appears with the **Community Operators** message, click **Continue**.
6. When the **Install Operator** window appears, you must select the **koku-metrics-operator** namespace for installation. If the namespace does not yet exist, it is created for you.
7. Click **Install**.

Verification steps

- After a short wait, **Koku Metrics Operator** appears in the **Installed Operators** tab under **Project: all projects** or **Project: koku-metrics-operator**.

Additional resources

- For more details on the Operator Lifecycle Manager, see [What is Operator Lifecycle Manager?](#)

3.9.2. Configuring Koku Metrics Operator on a restricted network

After the **koku-metrics-operator** is installed, you must configure it to run on a restricted network.

Prerequisites

- **koku-metrics-operator** installed.
- You are logged into the OpenShift Container Platform web console with cluster administrator privileges.

Procedure

1. From the OpenShift Container Platform web console, select **Operators > Installed Operators > koku-metrics-operator > KokuMetricsConfig > Create Instance**.
2. Specify the desired storage. If not specified, the operator will create a default persistent volume claim called **koku-metrics-operator-data** with **10Gi** of storage.




NOTE

To configure the **koku-metrics-operator** to use or create a different PVC, update the **volume_claim_template** configuration in **YAML view**.

3. Select **YAML view**.
4. Specify the maximum number of reports to store using **max_reports_to_store**, and time between report generation in minutes using **upload_cycle**.

```
packaging:  
  max_reports_to_store: 30  
  max_size_MB: 100
```

```
upload:  
  upload_cycle: 360
```



IMPORTANT

The **koku-metrics-operator** creates one report every **360 minutes** by default. Therefore, the default value of **30** reports and **360 minutes** gives you **7.5 days** of reports.

Any report generated after the total number specified will replace the oldest report in storage. Make sure to download generated reports from your PVC before they are lost.

5. Set **upload_toggle** to **false**.

```
upload:
  upload_cycle: 360
  upload_toggle: false
```

6. Replace the configuration in the **source** section with empty brackets.

```
source: {}
```

7. Replace the configuration in the **authentication** section with empty brackets.

```
authentication: {}
```

8. Click **Create**.

Verification steps

1. Select the **KokuMetricsConfig** you created.
2. Select **YAML view**.
3. Verify that a report has been created in the **packaging** section.

```
packaging:
  last_successful_packaging_time: `current date and time`
  max_reports_to_store: 30
  max_size_MB: 100
  number_of_reports_stored: 1
  packaged_files:
    - >-
      /tmp/koku-metrics-operator-reports/upload/YYYYMMDDTHHMMSS-cost-mgmt.tar.gz
```



NOTE

koku-metrics-operator will generate an initial report after it is configured. Generated reports will be listed under **packaged_files**.

3.9.3. Downloading cost reports

If the **koku-metrics-operator** is configured to run in a restricted network, copy the reports from the PVC where they are temporarily stored to a workstation with unrestricted network access for upload to cloud.redhat.com.

**NOTE**

The default configuration saves one week of reports. Therefore, download the reports locally and upload them to cloud.redhat.com weekly to prevent loss of metrics data.

Prerequisites

- Workstation with unrestricted network access.
- **koku-metrics-operator** reports in your PVC.

Procedure

1. Create the following pod with **claimName** matching the PVC containing the report data:

```
kind: Pod
apiVersion: v1
metadata:
  name: volume-shell
  namespace: koku-metrics-operator
spec:
  volumes:
  - name: koku-metrics-operator-reports
    persistentVolumeClaim:
      claimName: koku-metrics-operator-data
  containers:
  - name: volume-shell
    image: busybox
    command: ['sleep', '3600']
    volumeMounts:
    - name: koku-metrics-operator-reports
      mountPath: /tmp/koku-metrics-operator-reports
```

2. Use **rsync** to copy all of the files from the PVC to a local folder.

```
$ oc rsync volume-shell:/tmp/koku-metrics-operator-reports/upload local/path/to/save/folder
```

3. Confirm that the files have been copied.
4. Connect to the pod and delete the contents of the upload folder.

```
$ oc rsh volume-shell
$ rm /tmp/koku-metrics-operator-reports/upload/*
```

5. (Optional) Delete the pod that was used to connect to the PVC.

```
$ oc delete -f volume-shell.yaml
```

Additional resources

- For more information on persistent volume claims, see [Understanding persistent storage](#).

3.9.4. Uploading cost reports to cloud.redhat.com

You must manually upload locally stored cost reports from a restricted network to cloud.redhat.com.



NOTE

The default configuration saves one week of reports. Therefore, download the reports locally and upload them to cloud.redhat.com weekly to prevent loss of metrics data.

Prerequisites

- **koku-metrics-operator** reports downloaded locally.
- Source created on cloud.redhat.com, see [Section 3.8, "Creating an Openshift Container Platform source manually"](#).
- Red Hat account user with Organization Administrator entitlements.
- Workstation with unrestricted network access.

Procedure

- Upload your reports to cloud.redhat.com, replacing USERNAME and PASSWORD with your cloud.redhat.com login credentials, and FILE_NAME with the report to upload:


```
$ curl -vvvv -F "file=@FILE_NAME.tar.gz;type=application/vnd.redhat.hccm.tar+tgz"
https://cloud.redhat.com/api/ingress/v1/upload -u USERNAME:PASS
```

Verification steps

1. From cloud.redhat.com/cost-management, click **OpenShift**.
2. Verify you have OpenShift usage data for your cluster on the **OpenShift details** page.

CHAPTER 4. ADDING AN AMAZON WEB SERVICES (AWS) SOURCE TO COST MANAGEMENT


To add an AWS account to cost management, you must configure your AWS account to provide metrics, then add your AWS account as a source from the cost management user interface.



NOTE

You must have a Red Hat account user with Organization Administrator entitlements before you can add sources to cost management.

When you add your AWS account as a source, this creates a read-only connection to AWS in order to collect cost information hourly in cost management, but does not make any changes to the AWS account.



IMPORTANT

You must use an AWS master account for this procedure, as a linked AWS account does not have sufficient access to create billing accounts. After you add the master account as a source, cost management will collect data from any linked accounts as well.

Before you can add your AWS account to cost management as a data source, you must configure the following services on your AWS account to allow cost management access to metrics:

1. An S3 bucket to store cost and usage data reporting for cost management
2. An Identity Access Management (IAM) policy and role for cost management to process the cost and usage data

As you will complete some of the following steps in the AWS console, and some steps in the cost management user interface, keep both applications open in a web browser.

Add your AWS source to cost management from the settings area at <https://cloud.redhat.com/settings/sources/>.



NOTE

As non-Red Hat products and documentation can change without notice, instructions for configuring the third-party sources provided in this guide are general and correct at the time of publishing. See the [AWS documentation](#) for the most up-to-date and accurate information.


4.1. CREATING AN S3 BUCKET FOR REPORTING

Cost management requires an Amazon S3 bucket with permissions configured to store billing reports.

Log into your AWS master account to begin configuring cost and usage reporting:

1. In the AWS S3 console, create a new S3 bucket or use an existing bucket. If you are configuring a new S3 bucket, accept the default settings.
2. In the AWS Billing console, create a Cost and Usage Report that will be delivered to your S3 bucket. Specify the following values (and accept the defaults for any other values):

- Report name: *<any-name>* (note this name as you will use it later)
- Additional report details: Include resource IDs
- S3 bucket: *<the S3 bucket you configured previously>*
- Time granularity: Hourly
- Enable report data integration for: Amazon Redshift, Amazon QuickSight (do not enable report data integration for Amazon Athena)
- Compression type: GZIP
- Report path prefix: cost

**NOTE**

See the *AWS Billing and Cost Management* documentation for more details on configuration.

3. In the cloud.redhat.com platform, open the **Sources** menu (<https://cloud.redhat.com/settings/sources/>) to begin adding an AWS source to cost management:
 - a. Navigate to **Sources** and click **Add a source** to open the Sources wizard.
 - b. Enter a name for your source and click **Next**.
 - c. Select **Cost Management** as the application and **Amazon Web Services (AWS)** as the source type. Click **Next**.
 - d. Paste the name of your S3 bucket and click **Next**.

4.2. ACTIVATING AWS TAGS FOR COST MANAGEMENT

To use tags to organize your AWS resources in the cost management application, activate your tags in AWS to allow them to be imported automatically.

Procedure

1. In the AWS Billing console:
 - a. Open the *Cost Allocation Tags* section.
 - b. Select the tags you want to use in the cost management application, and click **Activate**.
2. In the cloud.redhat.com **Sources** wizard, click **Next** to move to the next screen.

4.3. ENABLING MINIMAL ACCOUNT ACCESS FOR COST AND USAGE CONSUMPTION

To provide data within the web interface and API, cost management needs to consume the Cost and Usage Reports produced by AWS. For cost management to obtain this data with a minimal amount of access, create an IAM policy and role for cost management to use. This configuration provides access to the stored information and nothing else.

Procedure

1. From the AWS Identity and Access Management (IAM) console, create a new IAM policy for the S3 bucket you configured previously.
 - a. Select the JSON tab and paste the following content in the JSON policy text box:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "s3:Get*",
        "s3:List*"
      ],
      "Resource": [
        "arn:aws:s3:::bucket_name",
        "arn:aws:s3:::bucket_name/*"
      ]
    },
    {
      "Sid": "VisualEditor1",
      "Effect": "Allow",
      "Action": [
        "s3:HeadBucket",
        "cur:DescribeReportDefinitions"
      ],
      "Resource": "*"
    }
  ]
}
```

- a. Provide a name for the policy and complete the creation of the policy. Keep the AWS IAM console open as you will need it for the next step.
 1. In the cloud.redhat.com **Sources** wizard, click **Next** to move to the next screen.
 2. In the AWS IAM console, create a new IAM role:
- b. For the type of trusted entity, select **Another AWS account**
- c. Enter *589173575009* as the Account ID to provide the cost management application with read access to the AWS account cost data.
- d. Attach the IAM policy you just configured.
- e. Enter a role name (and description if desired) and finish creating the policy.
 1. In the cloud.redhat.com **Sources** wizard, click **Next** to move to the next screen.
 2. In the AWS IAM console under Roles, open the summary screen for the role you just created and copy the Role ARN (a string beginning with *arn:aws:*).
 3. In the cloud.redhat.com **Sources** wizard, paste your Role ARN and click **Next**.

4. Review the details and click **Finish** to add the AWS account to cost management.

Cost management will begin collecting cost and usage data from your master AWS account and any linked AWS accounts.

The data can take a few days to populate before it shows on the cost management dashboard (<https://cloud.redhat.com/cost-management/>).

4.3.1. Enabling additional account access for cost and usage consumption

Cost management can display additional data that might be useful. For example:

- Include the Action `iam:ListAccountAliases` to display an AWS account alias rather than an account number in cost management.
- Include the Actions `organization:List*` and `organizations:Describe*` to obtain the display names of AWS member accounts if you are using consolidated billing rather than the account ID.

The following configuration provides access to additional stored information and nothing else.

Procedure

1. From the AWS Identity and Access Management (IAM) console, create a new IAM policy for the S3 bucket you configured previously.
2. Select the JSON tab and paste the following content in the JSON policy text box:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "s3:Get*",
        "s3:List*"
      ],
      "Resource": [
        "arn:aws:s3:::bucket",
        "arn:aws:s3:::bucket/*"
      ]
    },
    {
      "Sid": "VisualEditor1",
      "Effect": "Allow",
      "Action": [
        "iam:ListAccountAliases",
        "s3:HeadBucket",
        "cur:DescribeReportDefinitions",
        "organizations:List*",
        "organizations:Describe*"
      ],
      "Resource": "*"
    }
  ]
}
```

The remainder of the configuration steps are the same as in [Section 4.3.1, “Enabling additional account access for cost and usage consumption”](#).

You have completed adding your AWS account as a source.

CHAPTER 5. ADDING A MICROSOFT AZURE SOURCE TO COST MANAGEMENT

This section describes how to configure your Microsoft Azure account to allow cost management access.

Configuring your Azure account to be a cost management source requires:

1. Creating a storage account and resource group
2. Configuring a Storage Account Contributor and Reader roles for access
3. Scheduling daily cost exports



NOTE

As non-Red Hat products and documentation can change without notice, instructions for configuring the third-party sources provided in this guide are general and correct at the time of publishing. See the [Microsoft Azure documentation](#) for the most up-to-date and accurate information.

Add your Azure source to cost management from <https://cloud.redhat.com/settings/sources/>.

5.1. CREATING AN AZURE RESOURCE GROUP AND STORAGE ACCOUNT

Cost export data is written to a storage account, which exists within a resource group. The resource group must be accessible by cost management in order to read the Azure cost data.

Create a new storage account in Azure to contain the cost data and metrics that cost management will collect. This requires a resource group; Red Hat recommends creating a dedicated resource group for this storage account.



NOTE

You must have a Red Hat account user with Organization Administrator entitlements before you can add sources to cost management.

1. In the cloud.redhat.com platform, open the **Sources** menu (<https://cloud.redhat.com/settings/sources/>) to begin adding an Azure source to cost management:
 - a. Navigate to **Sources** and click **Add a source** to open the Sources wizard.
 - b. Enter a name for your source and click **Next**.
 - c. Select **Cost Management** as the application and **Microsoft Azure** as the source type. Click **Next**.
2. Create a resource group and storage account in your Azure account using the instructions in the See [Azure documentation](#) *Create a storage account*.

Make a note of the resource group and storage account. They will be needed in subsequent steps.

3. In the cloud.redhat.com **Sources** wizard, enter the **Resource group name** and **Storage account name** and click **Next**.

5.2. CONFIGURING AZURE ROLES

Red Hat recommends configuring dedicated credentials to grant cost management read-only access to Azure cost data. Configure a Storage Account Contributor and Reader role in Azure to provide this access to cost management.

1. In Azure Cloud Shell, run the following command to obtain your Subscription ID:

```
$ az account show --query "{subscription_id: id}"
```

2. In the cloud.redhat.com **Sources** wizard, enter your **Subscription ID**. Click **Next** to move to the next screen.
3. In Azure Cloud Shell, run the following command to create a Cost Management Storage Account Contributor role, and obtain your tenant ID, client (application) ID, and client secret:

```
$ az ad sp create-for-rbac -n "CostManagement" --role "Storage Account Contributor" --query '{"tenant": tenant, "client_id": appld, "secret": password}'
```

4. In the cloud.redhat.com **Sources** wizard, enter your Azure **Tenant ID**, **Client ID**, and **Client Secret**.
5. In Azure Cloud Shell, run the following command to create a Cost Management Reader role with your subscription ID. Copy the full command from the cloud.redhat.com **Sources** wizard, which will automatically substitute your Azure subscription ID obtained earlier for `<SubscriptionID>`:

```
$ az role assignment create --role "Cost Management Reader" --assignee http://CostManagement --subscription <SubscriptionID>
```

6. Click **Next**.

5.3. CONFIGURING A DAILY AZURE DATA EXPORT SCHEDULE

Create a recurring task to export your cost data on a daily basis automatically to your Azure storage account, where cost management will retrieve the data.

1. In Azure, add a new export as described in the instructions in the [Azure article](#) *Create and manage exported data*.
 - For **Export type**, select **Daily export of billing-period-to-date costs**
 - For **Storage account**, select the account you created earlier.
 - Enter any value for the container name and directory path for the export. These values provide the tree structure in the storage account where report files are stored.
 - Click **Run now** to start exporting data to the Azure storage container.
2. In the cloud.redhat.com **Sources** wizard, click **Next** when you have created the export schedule and review the source details.
3. Click **Finish** to complete adding the Azure source to cost management.

After the schedule is created, cost management will begin polling Azure for cost data, which will appear on the cost management dashboard (<https://cloud.redhat.com/cost-management/>).

You have completed adding your Azure account as a source.

CHAPTER 6. ADDING AN OPENSIFT CONTAINER PLATFORM 4.3 AND 4.4 SOURCE TO COST MANAGEMENT



NOTE

Use these instructions with OpenShift Container Platform 4.3 or 4.4 for use with the **cost-mgmt-operator** operator. This operator is deprecated for OpenShift Container Platform 4.5 and later.

For OCP 4.5 and later, see [Chapter 3, Adding an OpenShift Container Platform source to cost management](#).

To add an OpenShift Container Platform cluster as a source to cost management, you must first configure your cluster to provide usage data (metrics) using the Cost Management Operator.



NOTE

You must have a Red Hat account user with Organization Administrator entitlements before you can add sources to cost management.

The Cost Management Operator (**cost-mgmt-operator**) collects the metrics required for cost management by:

- Using Operator Metering to create usage reports specific to cost management.
- Collecting and packaging these reports to a tarball which is uploaded to cost management through cloud.redhat.com.

To add your OpenShift Container Platform cluster as a cost management source:

1. Install the Cost Management Operator in OpenShift from OperatorHub
2. Configure the Cost Management Operator to collect OpenShift usage data (metrics) using Operator Metering
3. Provide the cluster identifier to cost management

As you will complete some of the following steps in OpenShift Container Platform, and some steps in the cloud.redhat.com platform (<https://cloud.redhat.com/settings/sources/>), have both applications open in a web browser, as well as a terminal to access the command line interface (CLI).

6.1. INSTALLING THE COST MANAGEMENT OPERATOR

The Cost Management Operator collects the metrics required for cost management.

Begin adding your OpenShift Container Platform cluster as a source to cost management, then install the Cost Management Operator from OperatorHub.



NOTE

See [Operators](#) in the OpenShift documentation for more information about Operators and OperatorHub.

Prerequisites

- OpenShift Container Platform 4.3 or 4.4

Procedure

1. In the cloud.redhat.com platform, open the **Sources** menu (<https://cloud.redhat.com/settings/sources/>) to begin adding an OpenShift source to cost management:
 - a. Navigate to **Sources** and click **Add source** to open the Sources wizard.
 - b. Enter a name for your source and click **Next**.
 - c. Select **Cost Management** as the application and **OpenShift Container Platform** as the source type. Click **Next**.
2. In OpenShift, create a namespace called **openshift-metering** if one does not exist, and label the namespace with **openshift.io/cluster-monitoring=true**.
3. In OpenShift, install the Cost Management Operator in the **openshift-metering** namespace, using either the OpenShift web console (search for *cost management* in OperatorHub) or the CLI.



IMPORTANT

You must install the Cost Management Operator in the **openshift-metering** namespace. Other namespaces are not supported for installation.

See [Adding Operators to a cluster](#) in the OpenShift documentation for instructions for installing an Operator.

Additional resources

- See [Metering](#) in the OpenShift documentation for more information about installing Metering.

6.2. CONFIGURING THE COST MANAGEMENT OPERATOR

The Cost Management Operator (**cost-mgmt-operator**) collects the metrics required for cost management.


After installing the Cost Management Operator, configure authentication and the **operator-metering** namespace, then configure the Cost Management Operator.

Prerequisites

- OpenShift Container Platform 4.3 or 4.4
- The Cost Management Operator installed in the **openshift-metering** namespace
- A user with access to the **openshift-config** namespace

Procedure

1. Configure authentication inside the **openshift-metering** project. This allows you to upload OpenShift data to cloud.redhat.com.



NOTE

For most installations you can use token authentication or basic authentication to upload the usage reports (metrics) to cost management. Except for Azure RedHat OpenShift installation, the default and recommended method is token authentication.



NOTE

If you are performing an Azure RedHat OpenShift installation managed by Azure, you must use basic authentication. Token authentication is not supported for Azure-managed installations.

- a. Copy the following into a file called **auth_secret.yaml**:

```
kind: Secret
apiVersion: v1
metadata:
  name: auth-secret-name
  namespace: openshift-metering
  annotations:
    kubernetes.io/service-account.name: cost-mgmt-operator
data:
  username: >-
    Y2xvdWQu cmVkaGF0LmNvbSB1c2VybmFtZQ==
  password: >-
    Y2xvdWQu cmVkaGF0LmNvbSBwYXNzd29yZA==
  token: >-
    Y2xvdWQu cmVkaGF0LmNvbSB0b2t1bG==
```

- b. Choose a name for your authentication secret and replace the **metadata.name** value with it.
- c. To configure token authentication (the default method), obtain the correct *auth* token and then edit the secret to replace the token value:

- i. Install the **jq** JSON processor.
- ii. Change to the **openshift-config** namespace:

```
$ oc project openshift-config
```

- iii. Replace the token value in **auth_secret.yaml** with the authentication token for *cloud.openshift.com*. Obtain the token by running the following command, and copy only the *tokenvalue* to **auth_secret.yaml** (excluding the quotation marks):

```
$ oc get secret pull-secret -o "jsonpath={.data.\.dockerconfigjson}" | base64 --decode
| jq '.auths."cloud.openshift.com".auth'
```

**NOTE**

To use basic authentication, edit the secret to replace the username and password values with your base64-encoded username and password for connecting to cloud.redhat.com.

- d. Deploy the secret to your OpenShift cluster in the **openshift-metering** namespace:

```
$ oc create -f auth-secret.yaml
```

For both methods of authentication, the name of the secret should match the *authentication_secret_name* set in the **CostManagement** custom resource configured in the next steps.

2. Configure the Metering Operator.
Cost management uses the Metering Operator to create, collect, package, and upload metrics to cost management. In order for metering to work properly, configure **operator-metering** using the [OpenShift documentation](#) to create a MeteringConfig resource.
3. Configure the Cost Management Operator by creating the **CostManagement** and **CostManagementData** custom resources.
Creating these resources also starts the roles that create the resources to obtain the usage reports (metrics). This takes about an hour to run and the reports are collected, packaged, and uploaded every six hours.

**NOTE**

The Cost Management Operator requires the **clusterID**, **reporting_operator_token_name**, and **authentication_secret_name** to be specified in a **CostManagement** custom resource.

- a. Copy the following **CostManagement** resource template and save it to a file called **cost-mgmt-resource.yaml**:

```
apiVersion: cost-mgmt.openshift.io/v1alpha1
kind: CostManagement
metadata:
  name: cost-mgmt-setup
spec:
  clusterID: '123a45b6-cd8e-9101-112f-g131415hi1jk'
  reporting_operator_token_name: 'reporting-operator-token-123ab'
  validate_cert: 'false'
  authentication: 'basic'
  authentication_secret_name: 'basic_auth_creds-123ab'
```

- b. Edit the following values in your **cost-mgmt-resource.yaml** file:

- The **clusterID** value to your cluster ID. Obtain your cluster ID by running:

```
$ oc get clusterversion -o jsonpath='{.items[].spec.clusterID}'
```

- The **reporting_operator_token_name** to the **reporting-operator-token** secret name inside the **openshift-metering** namespace. Obtain this by running:

```
$ oc get secret -n openshift-metering | grep reporting-operator-token
```

**NOTE**

Depending on your configuration, this command can return two token names. You can use either to configure the Cost Management Operator.

- Specify the authentication type you are using (**token** or **basic**). If you are using **token** authentication, you can remove the authentication field as token authentication is the default.
- Change the **authentication_secret_name** to the name of your authentication secret you created earlier.

c. Deploy the **CostManagement** resource:

```
$ oc create -f cost-mgmt-resource.yaml
```

d. Create a **CostManagementData** resource to start the collection. Copy the following template and save it as **cost-mgmt-data-resource.yaml**:


```
apiVersion: cost-mgmt-data.openshift.io/v1alpha1
kind: CostManagementData
metadata:
  name: cost-mgmt-data-example
```

e. Deploy the **CostManagementData** resource:

```
$ oc create -f cost-mgmt-data-resource.yaml
```

The Cost Management Operator will now create, collect, package, and upload your OpenShift usage reports to cost management.

4. When configuration is complete, enter the cluster identifier into the cloud.redhat.com **Sources** wizard, click **Next**.

**NOTE**

The cluster identifier can be found in **Help > About** in OpenShift.

5. In the cloud.redhat.com **Sources** wizard, review the details and click **Finish** to add the OpenShift Container Platform cluster to cost management.

Additional resources

- See [Operators](#) in the OpenShift documentation for more information about Operators and OperatorHub.

Cost management will begin collecting usage data (metrics) from your OpenShift Container Platform cluster. The data can take a few days to populate before it shows on the cost management dashboard.

You have completed adding your OpenShift Container Platform cluster as a source.

CHAPTER 7. NEXT STEPS FOR MANAGING YOUR COSTS

After adding your infrastructure and cloud sources, in addition to showing cost data by source, cost management will automatically show AWS and Azure cost and usage related to running your OpenShift Container Platform clusters on AWS or Azure.

On the cost management **Overview** page, use the **Perspective** options for different views of your cost data.

Use the **Details** menu to look more closely at your costs.

7.1. CONFIGURE TAGGING FOR YOUR SOURCES

The cost management application tracks cloud and infrastructure costs using tags (called labels in OpenShift), which you can refine to filter and attribute to resources. Tags in cost management allow you to organize your resources by cost and to allocate the costs to different parts of your cloud infrastructure.



IMPORTANT

Tags and labels can only be configured directly on a source. You cannot edit tags and labels in the cost management application.

See [Managing cost data using tagging](#) to learn more about:

- Planning your tagging strategy to organize your view of cost data
- Understanding how cost management associates tags
- Configuring tags and labels on your sources

7.2. CONFIGURE COST MODELS TO ACCURATELY REPORT COSTS

Now that you have configured your sources to collect cost and usage data into cost management, you can configure cost models to associate prices to metrics and usage, and fine-tune the costs of running your cloud.

A cost model is a framework used to define the calculations for the costs stored in cost management, using raw costs and metrics. Costs generated by a cost model can be recorded, categorized and distributed to specific customers, business units or projects.

From the [Cost models](#) area of cost management, you can:

- Classify your costs as infrastructure or supplementary costs
- Capture monthly costs for OpenShift nodes and clusters
- Apply a markup to account for additional support costs
- Learn how to configure a cost model in [Using cost models](#).