



Red Hat Hybrid Cloud Console 2022

Planning for Red Hat Hybrid Cloud Console

Planning Red Hat Hybrid Cloud Console Administration

Red Hat Hybrid Cloud Console 2022 Planning for Red Hat Hybrid Cloud Console

Planning Red Hat Hybrid Cloud Console Administration

Red Hat Customer Content Services

Legal Notice

Copyright © 2022 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 provides an overview of Red Hat Hybrid Cloud Console features and can help Red Hat Enterprise Linux administrators with the planning and usage of the many services included in the Hybrid Cloud Console.

Table of Contents

| | |
|--|-----------|
| CHAPTER 1. WHAT IS USER ACCESS | 3 |
| 1.1. WHO CAN USE USER ACCESS | 3 |
| 1.2. ADDITIVE ACCESS | 3 |
| 1.3. THE USER ACCESS GROUPS, ROLES, AND PERMISSIONS | 3 |
| CHAPTER 2. SETTING USER PREFERENCES | 4 |
| CHAPTER 3. SYSTEM TAGS AND GROUPS | 5 |
| 3.1. TAG STRUCTURE | 5 |
| CHAPTER 4. SYSTEM STALENESS AND DELETION | 7 |
| 4.1. RULES FOR SYSTEM STALENESS AND DELETION | 7 |
| 4.2. VIEWING STALE SYSTEM SUMMARY | 7 |
| 4.3. FILTERING SYSTEM INVENTORY | 7 |
| CHAPTER 5. CONFIGURING SOURCES FOR RED HAT SERVICES | 9 |
| 5.1. ABOUT SOURCES | 9 |
| 5.2. ADDING CLOUD SOURCES | 9 |
| 5.3. ADDING RED HAT SOURCES | 12 |
| 5.4. SOURCES REFERENCE MATERIAL | 12 |
| MAKING OPEN SOURCE MORE INCLUSIVE | 14 |
| PROVIDING FEEDBACK ON RED HAT DOCUMENTATION | 15 |

CHAPTER 1. WHAT IS USER ACCESS

The User Access feature is an implementation of role-based access control (RBAC) that controls user access to various services hosted on the [Red Hat Hybrid Cloud Console](#). You configure the User Access feature to grant user access to services hosted on Hybrid Cloud Console.

1.1. WHO CAN USE USER ACCESS

To initially view and manage User Access on [Red Hat Hybrid Cloud Console](#), you must be an Organization Administrator. This is because User Access requires user management capabilities that are designated from the Red Hat Customer Portal at [Customer Portal](#). Those capabilities belong solely to the Organization Administrator.

The **User Access administrator** role is a special role that the Organization Administrator can assign. This role allows users who are not Organization Administrator users to manage User Access on [Red Hat Hybrid Cloud Console](#).

1.2. ADDITIVE ACCESS

User access on [Red Hat Hybrid Cloud Console](#) uses an additive model, which means that there are no **deny** roles. In other words, actions are only permitted. You control access by assigning the appropriate roles with the desired permissions to groups then adding users to those groups. The access permitted to any individual user is a sum of all roles assigned to all groups to which that user belongs.

1.3. THE USER ACCESS GROUPS, ROLES, AND PERMISSIONS

User Access uses the following categories to determine the level of user access that an Organization Administrator can grant to the supported [Red Hat Hybrid Cloud Console](#) services. The access provided to any authorized user depends on the group that the user belongs to and the roles assigned to that group.

- **Group:** A collection of users belonging to an account which provides the mapping of roles to users. An Organization Administrator can use groups to assign one or more roles to a group and to include one or more users in a group. You can create a group with no roles and no users.
- **Roles:** A set of permissions that provide access to a given service, such as Insights. The permissions to perform certain operations are assigned to specific roles. Roles are assigned to groups. For example, you might have a **read** role and a **write** role for a service. Adding both roles to a group grants all members of that group read and write permissions to that service.
- **Permissions:** A discrete action that can be requested of a service. Permissions are assigned to roles.

An Organization Administrator adds or deletes roles and users to groups. The group can be a new group created by an Organization Administrator or the group can be an existing group. By creating a group that has one or more specific roles and then adding users to that group, you control how that group and its members interact with the [Red Hat Hybrid Cloud Console](#) services.


When you add users to a group, they become members of that group. A group member inherits the roles of all other groups they belong to. The user interface lists users in the **Members** tab.

CHAPTER 2. SETTING USER PREFERENCES

Use the following procedure to set or update your email preferences.

Procedure

1. Click the user menu located on the upper-right side, then go to **User preferences → Email preferences**. The Email preferences screen opens.
Alternatively, on the Red Hat Hybrid Cloud Console dashboard, in the left-side navigation panel at the top, click **Red Hat Enterprise Linux**, and then click **User Preferences**. The Email preferences screen opens.
2. Depending on your email notification preference, you can subscribe to **Instant notification** emails for each system with triggered policies and/or **Daily digest** (summary) of all systems with triggered policies. On this page, you can also select your preference for other [Red Hat Hybrid Cloud Console](#) emails you want to receive.



NOTE


Subscribing to instant notification can result in receiving many emails on large inventories, that is, one email per system checking in.

3. Click **Submit**.

CHAPTER 3. SYSTEM TAGS AND GROUPS

The Red Hat Hybrid Cloud Console enables administrators to filter systems in inventory and in individual services using group tags. Groups are identified by the method of system data ingestion to the console. The Hybrid Cloud Console enables filtering groups of systems by those running SAP workloads, by Satellite host group, and by custom tags that are defined by system administrators with root access to configure the Insights client on the system.

Use the global, **Filter by status** box to filter by SAP workloads, Satellite host groups, or custom tags added to the Insights client configuration file.



Prerequisites

The following prerequisites and conditions must be met to use the tagging features in Red Hat Hybrid Cloud Console:

- The Insights client is installed and registered on each system.
- To create custom tags, root permissions, or their equivalent, are required to add to or change the `/etc/insights-client/tags.yaml` file.

3.1. TAG STRUCTURE

Tags use a **namespace/key=value** paired structure.

- **Namespace.** The namespace is the name of the ingestion point, *insights-client*, and cannot be changed. The **tags.yaml** file is abstracted from the namespace, which is injected by the client before upload.
- **Key.** The key can be a user-chosen key or a predefined key from the system. You can use a mix of capitalization, letters, numbers, symbols and whitespace.
- **Value.** Define your own descriptive string value. You can use a mix of capitalization, letters, numbers, symbols and whitespace.

CHAPTER 4. SYSTEM STALENESS AND DELETION

System deletion is the automated removal of systems from the Red Hat Insights inventory after all sources stop reporting information about it for a defined period of time.

System staleness is reporting when a system has missed check-ins for a defined period of time but is not yet deleted.

4.1. RULES FOR SYSTEM STALENESS AND DELETION

The inventory reporting service, as part of their messaging, includes a timestamp for when the report about the host is considered stale. This timestamp is determined by the reporting service and defaults to a value set by the user account.

When various reporters contribute data to the host in the host inventory, staleness states are recalculated.

Systems in the inventory have the following three fields related to staleness and deletion:

- **"stale_timestamp": "2019-12-13T19:36:30.979Z"**
- **"stale_warning_timestamp": "2019-12-13T19:36:30.979Z"**
- **"culled_timestamp": "2019-12-13T19:36:30.978Z"**

Rules:

- Before the `stale_timestamp` is reached, a system is considered fresh
- Between the `stale_timestamp` and `stale_warning_timestamp`, a system is considered stale
- Between the `stale_warning_timestamp` and `culled_timestamp`, a system is considered in the "stale warning" state and is scheduled for deletion
- After the `culled_timestamp` is reached, a system and all associated data are automatically deleted

4.2. VIEWING STALE SYSTEM SUMMARY

You can see stale systems and the ones scheduled for deletion in the following places in the Red Hat Insights user interface:

Dashboard: In the Red Hat Insights dashboard summary, you can see the number of stale systems and the number of systems scheduled for deletion displayed under System inventory. Click on the respective links to see the list of stale systems and the ones marked for deletion.

Inventory: On the last seen column, you will see systems marked for deletion with a warning icon. Hover on the last seen information to see if a system will be removed from the inventory in the next x days or if it is already scheduled for deletion.

4.3. FILTERING SYSTEM INVENTORY

You can filter the inventory by system status: Fresh, Stale, and Stale warning. Note that by default fresh and stale systems are listed, but stale warning systems are not shown in the user interface. You can also select the source, which is Insights by default, and filter by system name. You can remove a particular

filter in effect, or clear all at once.

CHAPTER 5. CONFIGURING SOURCES FOR RED HAT SERVICES

A data source is a service, application, or provider that supplies data to a [Red Hat Hybrid Cloud Console](#) application or service. Sources comprise cloud sources and Red Hat sources.

5.1. ABOUT SOURCES

Sources are how the services and applications on the [Red Hat Hybrid Cloud Console](#) connect with public cloud providers and other services or tools to collect information for the service or application. You add and manage sources in the **Sources** application located within the **Settings** bundle. You can access **Settings** by clicking the gear icon in the masthead on the [Red Hat Hybrid Cloud Console](#).

The **Sources** menu uses a wizard to help you add cloud sources and Red Hat sources. For cloud sources, you can associate the provider with Red Hat applications, such as Cost Management and the RHEL management bundle. For Red Hat sources, you can add Red Hat OpenShift Container Platform. Adding applications is optional for cloud sources but is required for Red Hat sources.

5.2. ADDING CLOUD SOURCES

The **Add a cloud source** wizard steps you through creating a source. You can add Amazon Web Services (AWS), Google Cloud, and Microsoft Azure. The wizard provides detailed information for each public cloud provider.

Amazon Web Services

The workflow for adding AWS as a cloud source includes the following high-level steps:

1. Selecting the source type
2. Naming the source
3. Selecting the configuration
4. Selecting applications
5. Configuring cost and usage reporting (for cost management)
6. Identifying tags, aliases, and organizational units (for cost management)
7. Enabling account access
8. Reviewing details
9. Adding the source

You have two choices for the configuration mode: * **Account authorization** (recommended) * **Manual configuration**

If you select **Account authorization**, you provide your AWS account credentials (Access key ID and Secret key ID) and Red Hat configures and manages your source for you. This option automatically selects the **Cost Management** and the **RHEL management** bundle applications. You can deselect these applications.

If you select **Manual configuration**, you choose **Cost Management**, **RHEL management** bundle, or **No application**.

The **Cost Management** application allows you to perform financially related tasks, such as:

- Visualizing, understanding, and analyzing the use of resources and costs
- Forecasting your future consumption and comparing them with budgets
- Optimizing resources and consumption
- Identifying patterns of usage for further analysis
- Integrating with third-party tools that can benefit from cost and resourcing data

The **RHEL management** bundle includes the following items:

- Red Hat gold images
- High precision subscription watch data
- Autoregistration

The **Cost Management** and the **RHEL management** bundle applications require you to enable account access. You accomplish this by creating an IAM policy, an IAM role, and entering your Amazon Resource Name (ARN). An ARN is a generic name for an Amazon resource and has a common format depending on the service involved. In this case, it is the identity and access management (IAM) service and Role resource-type.

If you select **No application**, you choose which credentials to supply:

- AWS Secret key
- Cost Management ARN
- Subscription Watch ARN

Google Cloud

The workflow for adding Google Cloud as a cloud source includes the following high-level steps:

1. Selecting the source type
2. Naming the source
3. Selecting applications
4. Adding a project
5. Enabling account access
6. Creating a dataset
7. Setting up billing export information
8. Reviewing details
9. Adding the source

Cost Management is the only application choice. You must create an IAM role and assign access.

The **Cost Management** application allows you to perform financially related tasks, such as:

- Visualizing, understanding, and analyzing the use of resources and costs
- Forecasting your future consumption and comparing them with budgets
- Optimizing resources and consumption
- Identifying patterns of usage for further analysis
- Integrating with third-party tools that can benefit from cost and resourcing data

If you select **No application**, you provide the Project ID and the Service Account JSON as credentials.

Microsoft Azure

The workflow for adding Microsoft Azure as a cloud source includes the following high-level steps:

1. Selecting the source type
2. Naming the source
3. Selecting applications
4. Creating a resource group and a storage account (for cost management)
5. Entering a Subscription ID (for cost management)
6. Creating roles (for cost management)
7. Setting up daily exports (for cost management)
8. Providing credentials
9. Reviewing details
10. Adding the source

Application choices include **Cost Management**, **RHEL management** bundle, or **No application**.

The **Cost Management** application allows you to perform financially related tasks, such as:


- Visualizing, understanding, and analyzing the use of resources and costs
- Forecasting your future consumption and comparing them with budgets
- Optimizing resources and consumption
- Identifying patterns of usage for further analysis
- Integrating with third-party tools that can benefit from cost and resourcing data

The **RHEL management** bundle includes the following items:

- Red Hat gold images

- Autoregistration

You create a dedicated resource group and a storage account in the Azure Portal so you can collect cost data and metrics for cost management. You then use your subscription ID to create a Cost Management Storage Account Contributor role in the Cloud Shell. Using the subscription ID to run a second command in the Cloud Shell gives you the tenant (directory) ID, client (application) ID, and client secret necessary to complete setting up that role.



NOTE

Configure dedicated credentials to grant cost management read-only access to Azure cost data.

5.3. ADDING RED HAT SOURCES

The workflow for adding Red Hat OpenShift Container Platform as a cloud source includes the following high-level steps:

1. Selecting the source type and application
2. Naming the source
3. Installing and configuring the operator
4. Reviewing the details
5. Adding the source

Cost Management is the only application choice. The **Cost Management** application allows you to perform financially related tasks, such as:

- Visualizing, understanding, and analyzing the use of resources and costs
- Forecasting your future consumption and comparing them with budgets
- Optimizing resources and consumption
- Identifying patterns of usage for further analysis
- Integrating with third-party tools that can benefit from cost and resourcing data



NOTE

For Red Hat OpenShift Container Platform 4.6 and later, install the **costmanagement-metrics-operator** from the OpenShift Container Platform web console. For more information, see [Adding an OpenShift Container Platform source to cost management](#) .

5.4. SOURCES REFERENCE MATERIAL

[Getting started with cost management](#)

[Getting Started with the Subscriptions Service](#)

[Adding sources for public cloud metering](#)

[Getting started with Automation Services Catalog](#)

[Bucket restrictions and limitations](#)

[Bucket naming rules](#)

MAKING OPEN SOURCE MORE INCLUSIVE

Red Hat is committed to replacing problematic language in our code, documentation, and web properties. We are beginning with these four terms: master, slave, blacklist, and whitelist. Because of the enormity of this endeavor, these changes will be implemented gradually over several upcoming releases. For more details, see [our CTO Chris Wright's message](#).

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

We appreciate your feedback on our documentation. To provide feedback, highlight text in a document and add comments.

Prerequisites

- You are logged in to the Red Hat Customer Portal.
- In the Red Hat Customer Portal, the document is in the **Multi-page HTML** viewing format.

Procedure

To provide your feedback, perform the following steps:

1. Click the **Feedback** button in the top-right corner of the document to see existing feedback.



NOTE

The feedback feature is enabled only in the **Multi-page HTML** format.

2. Highlight the section of the document where you want to provide feedback.
3. Click the **Add Feedback** pop-up that appears near the highlighted text.
A text box appears in the feedback section on the right side of the page.
4. Enter your feedback in the text box and click **Submit**.
A documentation issue is created.
5. To view the issue, click the issue link in the feedback view.