



Red Hat Advanced Cluster Management for Kubernetes 2.5

Web console

Read more to learn how to use console components.

Red Hat Advanced Cluster Management for Kubernetes 2.5 Web console

Read more to learn how to use console components.

Legal Notice

Copyright © 2023 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

Read more to learn how to use console components.

Table of Contents

CHAPTER 1. WEB CONSOLE	3
1.1. ACCESSING YOUR CONSOLE	3
1.1.1. From the OpenShift Container Platform web console	3
1.1.2. From the Red Hat OpenShift CLI:	3
1.2. CONSOLE OVERVIEW	3
1.2.1. Console components	4
1.2.2. OpenShift Container Platform console plug-ins	4
1.2.3. Home	5
1.2.4. Infrastructure	5
1.2.5. Applications	6
1.2.6. Governance	6
1.2.7. Credentials	6
1.3. SEARCH IN THE CONSOLE	6
1.3.1. Search components	6
1.3.2. Search customization	7
1.3.2.1. Options to increase the redisgraph memory	7
1.3.3. Queries in the console	7
1.3.3.1. Query ArgoCD applications	8

CHAPTER 1. WEB CONSOLE

Learn how to access the Red Hat Advanced Cluster Management for Kubernetes console for your Red Hat Advanced Cluster Management console and learn how to use console components from the following documentation:

- [Accessing the console](#)
- [Console overview](#)
- [Search in the console](#)

1.1. ACCESSING YOUR CONSOLE

From the Red Hat OpenShift Container Platform web console, you can access the Red Hat Advanced Cluster Management for Kubernetes console. You can also access the console from your terminal after installation. See all the options for accessing your console.

1.1.1. From the OpenShift Container Platform web console

- **OpenShift Container Platform 4.10 Technology Preview:** See the OpenShift Container Platform documentation about [Adding a dynamic plug-in to the OpenShift Container Platform web console](#) to create and deploy a dynamic plug-in on your cluster that is loaded at run-time.
- For OpenShift Container Platform 4.8 and later, you can launch Red Hat Advanced Cluster Management from the perspective switcher near the navigation menu. **Note:** Only for *earlier* versions of OpenShift Container Platform, you can click the **Application launcher** in the header and choose the *Red Hat Advanced Cluster Management for Kubernetes* option.
- You can also access the console from the left-hand navigation:
 - Click **Networking > Routes**, then click on the URL from the *Location* column.
 - From the *Project* menu, select the namespace where Red Hat Advanced Cluster Management for Kubernetes is installed. The default namespace is **open-cluster-management**.
 - Click the **multicloud-console** URL from the *Location* column.

1.1.2. From the Red Hat OpenShift CLI:

1. While logged in to Red Hat OpenShift Container Platform and after you install Red Hat Advanced Cluster Management for Kubernetes, run the following command to find the route. Here, **<namespace-from-install>** is the namespace where you installed the product:

```
oc get routes -n <namespace-from-install>
```

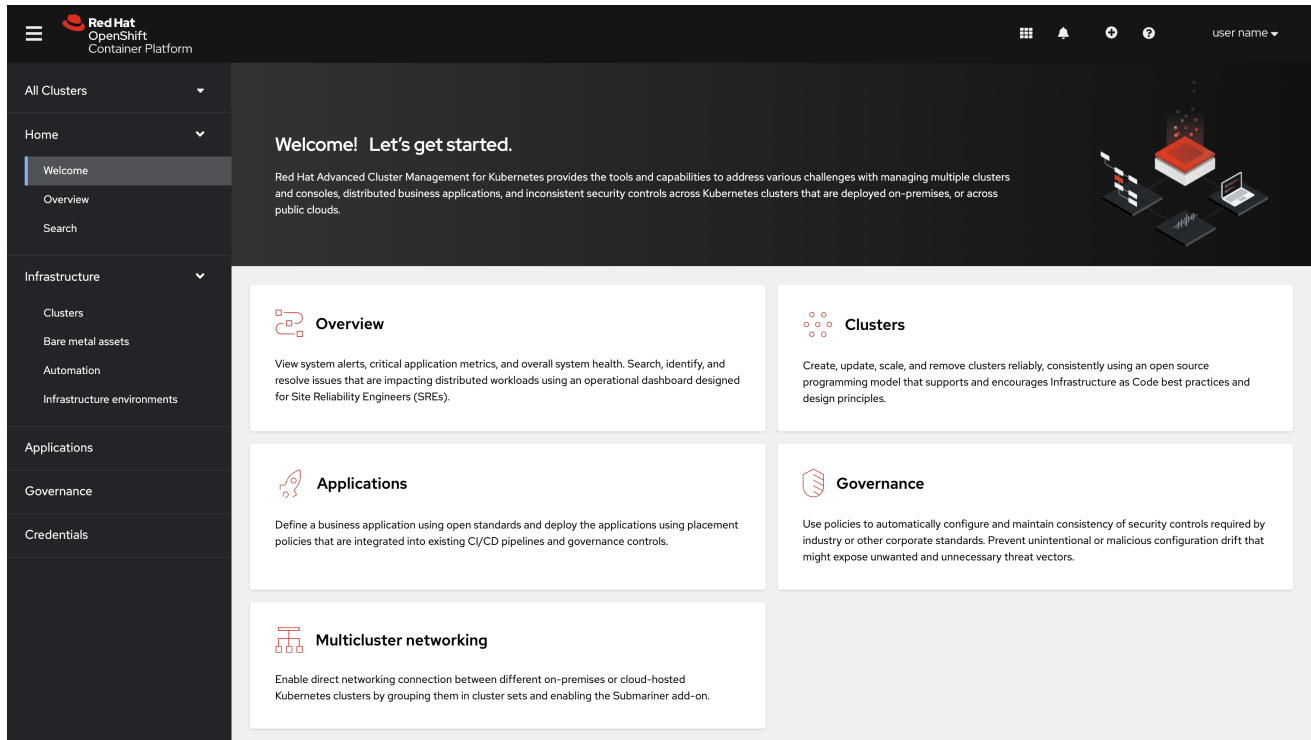
1. Find the **open-cluster-management** name and the **Host/Port** columns to get your URL.

Learn more about the Red Hat Advanced Cluster Management for Kubernetes console, see [Web console](#).

1.2. CONSOLE OVERVIEW

Learn more about console components that you can use to view, manage, or customize your console.

See the following image of the *Navigation* from the Red Hat Advanced Cluster Management for Kubernetes console, which is described in more detail later in each section. See that the navigation represents major production function.



1.2.1. Console components

- [OpenShift Container Platform console plug-ins](#)
- [Home](#)
- [Infrastructure](#)
- [Applications](#)
- [Governance](#)
- [Credentials](#)

1.2.2. OpenShift Container Platform console plug-ins

With OpenShift Container Platform 4.10 and Red Hat Advanced Cluster Management 2.5, Red Hat Advanced Cluster Management web console can be integrated into the OpenShift Container Platform console. To use this feature, the console plug-ins must remain enabled and accessed from the *All Clusters* view.

For new installations, the console plug-ins are enabled by default. If you upgraded from a previous version of Red Hat Advanced Cluster Management and want to enable the plug-ins, or if you want to disable the plug-ins, see the following instructions:

1. To disable the plug-in, be sure you are in the *Administrator* perspective in the OpenShift Container Platform console.

2. Find **Administration** in the navigation and click **Cluster Settings**, then click *Configuration* tab.
3. From the list of *Configuration resources*, click the **Console** resource with the **operator.openshift.io** API group, which contains cluster-wide configuration for the web console.
4. Click on the *Console plug-ins* tab. Both the **acm** and **mce** plug-ins are listed.
5. Modify plug-in status from the table. In a few moments, you are prompted to refresh the console.

For OpenShift Container Platform 4.10 with the plug-ins enabled, you can access Red Hat Advanced Cluster Management within the OpenShift Container Platform console from the cluster switcher by selecting **All Clusters** from the drop-down menu. Without plug-ins enabled on OpenShift Container Platform versions 4.8 through 4.10, Red Hat Advanced Cluster Management is available in the perspective switcher.

You can also click **Networking > Routes** to access the standalone or separate Red Hat Advanced Cluster Management web console.

1.2.3. Home

From the Red Hat Advanced Cluster Management for Kubernetes *Home* page, you get more information about the product and you can access header features, as well as the pages for the major components of the product.

- Access the *Welcome* page and the *Overview*, which gives you visibility into your clusters. You can view the following information about your clusters on the *Overview* dashboard:
 - Metric data from your managed clusters by selecting the Grafana link
 - Cluster, node, and pod counts across all clusters and for each provider
 - Cluster status
 - Cluster compliance
 - Pod status
- Select **Grafana** to access the Grafana dashboard.
- Click **Add provider connections** to access the *Clusters* page.

Search is also available from the *Home* tab. To learn about Search, see [Search in the console](#).

1.2.4. Infrastructure

From *Clusters*, you can create new clusters or import existing clusters. From *Bare metal assets*, you can create and import assets. From *Automation*, you can create an Ansible template.

For more information about managing clusters, see [Managing your clusters with Red Hat Advanced Cluster Management for Kubernetes](#).

Additionally, see specific information on these cluster types at [Creating and modifying bare metal assets](#) and [Configuring Ansible Tower tasks to run on managed clusters](#).

1.2.5. Applications

Create an application and edit a **.yaml** file. Access an overview or more advanced information about each application. For more information about application resources, see [Managing applications](#).

1.2.6. Governance

Create and edit a **.yaml** file to create a policy. Use the *Governance* dashboard to manage policies and policy controllers.

For more information, see [Governance](#).

1.2.7. Credentials

The credential stores the access information for a cloud provider. Each provider account requires its own credential, as does each domain on a single provider.

Review your credentials or add a credential.

See [Managing credentials overview](#) for more specific information about providers and credentials.

1.3. SEARCH IN THE CONSOLE

For Red Hat Advanced Cluster Management for Kubernetes, search provides visibility into your Kubernetes resources across all of your clusters. Search indexes the Kubernetes resources and the relationships to other resources. You can create a **searchcustomization** custom resource (CR) to define the storage settings for search persistence if you want to change the storage class and storage size.

- [Search components](#)
- [Search customization](#)
 - [Options to increase the redisgraph memory](#)
- [Queries in the console](#)
 - [Query ArgoCD applications](#)

1.3.1. Search components

The search architecture is composed of the following components:

- **Collector:** Watches the Kubernetes resources and creates an index. The **search-collector** computes relationships for resources within the managed cluster.
- **Aggregator:** Receives data from the collector and writes to a database. The **search-aggregator** watches resources in the hub cluster, computes multicluster relationships, and tracks the activity from connected collectors.
- **Search API:** Provides access to the data in the search index and enforces role-based access control.

Search is enabled by default. Search is also enabled when you provision or manually import a managed cluster. If you want to disable search on your managed cluster, see [Modifying the kubernetes add-ons settings of your cluster](#) for more information.

1.3.2. Search customization

When you install Red Hat Advanced Cluster Management, the product is configured to persist the data in-memory to a file system. The StatefulSet **search-redisgraph** deploys the Redisgraph pod, which mounts the persistent volume named **persist**. If your cluster has a defined default storage class, the search component creates a Persistent Volume Claim (PVC) of 10Gi on the default storage class. If a default storage class does not exist in your cluster, search saves the index in an empty directory (**emptyDir**).

You can customize the storage settings for search by creating the **searchcustomization** CR. Search customization is namespace-scoped and located where search is installed in the hub cluster. View the following example of the search customization CR:

```
apiVersion: search.open-cluster-management.io/v1alpha1
kind: SearchCustomization
metadata:
  name: searchcustomization
  namespace: open-cluster-management
  labels:
    cluster.open-cluster-management.io/backup: ""
spec:
  persistence: true
  storageClass: gp2
  storageSize: 12Gi
```

Run the following command to view search customization CRD:

```
oc get crd searchcustomizations.search.open-cluster-management.io -o yaml
```

You can disable persistence by updating the **persistence** flag to **false** in the customization CR, which turns off saving search index to the file system. A status for persistence can be retrieved from the search operator (**searchoperator**) CR. Run the following command to view search operator CR: **oc get searchoperator searchoperator -o yaml**.

1.3.2.1. Options to increase the redisgraph memory

Redisgraph is an in-memory database that needs linear growth of memory as the number of objects are cached. A Red Hat Advanced Cluster Management cluster with many managed clusters, or with a large number of Kubernetes objects require limit memory updates for the redisgraph pod (**search-redisgraph-0**).

By default, the redisgraph pod (**search-redisgraph-0**) is deployed with a memory limit of **4Gi**. If you are managing larger clusters, you might need to increase this limit by editing the **redisgraph_resource.limit_memory** for the **searchoperator** in the hub cluster namespace. For example, you can update the limit to **8Gi** with the following command:

```
oc patch searchoperator searchoperator --type='merge' -p '{"spec":{"redisgraph_resource":{"limit_memory":"8Gi"}}}'
```

After the change is made, the **search-redisgraph** pod automatically restarts with the updated configuration.

1.3.3. Queries in the console

You can type any text value in the *Search box* and results include anything with that value from any property, such as a name or namespace. Users are unable to search for values that contain an empty space.

For more specific search results, include the property selector in your search. You can combine related values for the property, for a more precise scope of your search. For example, search for **cluster:dev red** to receive results that match the string "red" in the **dev** cluster.

View the following steps to make queries with search:

1. Click **Search** in the navigation menu.
2. Type a word in the *Search box*, then Search finds your resources that contain that value.
 - As you search for resources, you receive other resources that are related to your original search result, which help you visualize how the resources interact with other resources in the system.
 - Search returns and lists each cluster with the resource that you search. For resources in the *hub* cluster, the cluster name is displayed as *local-cluster*.
 - Your search results are grouped by **kind**, and each resource **kind** is grouped in a table.
 - Your search options depend on your cluster objects. You can refine your results with specific labels. Search is case-sensitive when you query labels. See the following examples: name, namespace, status, and other resource fields. Auto-complete provides suggestions to refine your search. See the following example:
 - Search for a single field, such as **kind:pod** to find all pod resources.
 - Search for multiple fields, such as **kind:pod namespace:default** to find the pods in the default namespace.

Notes:

- You can also search with conditions by using characters, such as **>**, **>=**, **<**, **<=**, **!=**.
 - When you search for more than one property selector with multiple values, the search returns either of the values that were queried. View the following examples:
 - When you search for **kind:pod name:a**, any pod named **a** is returned.
 - When you search for **kind:pod name:a,b**, any pod named **a** or **b** are returned.
 - Search for **kind:pod status:!Running** to find all pod resources where the status is not **Running**.
 - Search for **kind:pod restarts:>1** to find all pods that restarted at least twice.
3. If you want to save your search, click the **Save search** icon.

1.3.3.1. Query ArgoCD applications

When you search for an ArgoCD application, you are directed to the *Applications* page. Complete the following steps to access the ArgoCD application from the *Search* page:

1. Log in to your Red Hat Advanced Cluster Management hub cluster.

2. From the console header, select the *Search* icon.
3. Filter your query with the following values: **kind:application** and **apigroup:argoproj.io**
4. Select an application to view. The *Application* page displays an overview of information for the application.

Learn more about the Red Hat Advanced Cluster Management for Kubernetes console, see [Web console](#).