



Red Hat Advanced Cluster Management for Kubernetes 2.2

Web console

Read more to learn how to use console components.

Red Hat Advanced Cluster Management for Kubernetes 2.2 Web console

Read more to learn how to use console components.

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

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. OBSERVABILITY IN THE CONSOLE	3
1.2.1. Console components	4
1.2.2. Access your console	4
1.2.3. Welcome page	5
1.2.4. Observe environments Overview page	5
1.2.5. Observe environment details	5
1.2.6. Automate infrastructure	5
1.2.7. Manage applications	5
1.2.7.1. Viewing your pod health	5
1.2.8. Governance and risk dashboard	6
1.3. SEARCH IN THE CONSOLE	6
1.3.1. Search customization	6
1.3.2. Queries in the console	6
1.4. VISUAL WEB TERMINAL	7
1.4.1. Starting a session	8
1.4.2. Supported command list	8
1.4.3. Searching with Visual Web Terminal	10
1.5. MANAGING CLUSTER LABELS	11

CHAPTER 1. WEB CONSOLE

Learn how to access the console for your Red Hat Advanced Cluster Management console and learn how to use some of the console components.

- [Accessing the console](#)
- [Observability in the console](#)
- [Search in the console](#)
- [Searching with Visual Web Terminal](#)
- [Managing cluster labels](#)

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

- 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**
 - From the *Project* menu, select the namespace where Red Hat Advanced Cluster Management for Kubernetes is installed.
 - Find **multicloud-console** and click the 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>
```

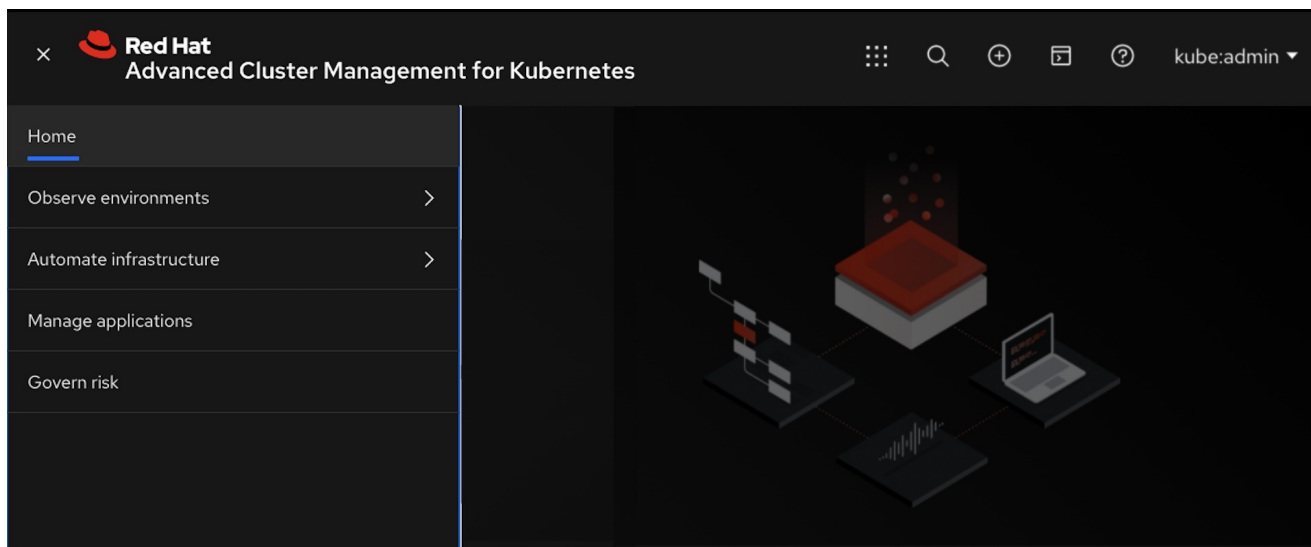
1. Find the **multicloud-console** name and the **Host/Port** columns to get your URL.

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

1.2. OBSERVABILITY IN THE CONSOLE

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

- [Access your console](#)
- [Welcome page](#)
- [Observe environments Overview page](#)
- [Observe environment details](#)
- [Automate Infrastructure](#)
- [Manage applications](#)
- [Governance and risk dashboard](#)

To learn about Search, see [Search in the console](#)

1.2.2. Access your console

- Access from the Red Hat OpenShift Container Platform console:
 - From the left-hand navigation, click **Networking** > **Routes**
 - From the *Project* menu, select the namespace where Red Hat Advanced Cluster Management for Kubernetes is installed.
 - Find **multicloud-console** and click the URL from the **Location** column.
- Access the console from the Red Hat OpenShift Container Platform CLI:
 - While logged in to OpenShift, run the following command to find the route, where **<namespace-from-install>** is the namespace where you installed Red Hat Advanced Cluster Management for Kubernetes:

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


- a. Find the **multicloud-console** name and the **Host/Port** columns to get your URL.

1.2.3. Welcome page

From the Red Hat Advanced Cluster Management for Kubernetes Welcome page, you get more information about the product and you can access header features, such as the *Applications* menu, *Search*, *Create resource*, *Visual Web Terminal*, and *Info*.

When you select the *Applications* menu, you can navigate to different managed clusters. Select the **Search** icon to query and observe cluster resources from the *Search* page. Additionally, search capabilities are also available in the *Visual Web Terminal* from the header. Use the Visual Web Terminal to run commands across your cluster. For more details, see [Visual Web Terminal](#).

Select **Create resource** to create a YAML file or JSON file for a deployment. Also from the *Header*, you can click the **Info** (?) icon to view the *About* page and the documentation. From the *User* menu, you can access the *Configure Client* page.

1.2.4. Observe environments Overview page

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

Additionally, you can view many clickable elements on the dashboard open a search for related resources. Click on a provider card to view information for clusters from a single provider.

1.2.5. Observe environment details

You can gain insight to optimize your managed clusters by viewing cluster and pod information from Grafana dashboards. For more details, see [Observing environments](#).

1.2.6. Automate infrastructure

Create clusters or import existing clusters. Scale up or down and delete clusters as needed. For more information, see [Managing your clusters with Red Hat Advanced Cluster Management for Kubernetes](#).

1.2.7. Manage applications

Click **New application** to edit a **.yaml** file and create your application. Click **Overview** and **Advanced configuration** to view the information of each application. For more information about application resources, see [Managing applications](#).

1.2.7.1. Viewing your pod health

View the pod health for all of your clusters by expanding the Heat map.

Click **Show details** to view the map. The size of the color-coordinated boxes represents the number of nodes on your cluster. Hover your cursor over the box to view the response time of your cluster.

1.2.8. Governance and risk dashboard

Use the Governance and risk dashboard to create and manage policies and policy controllers. For more information, see [Governance and risk](#).

1.3. SEARCH IN THE CONSOLE

For Red Hat Advanced Cluster Management for Kubernetes, search provides visibility into your resources across all of your clusters. Search uses Redisgraph, an in-memory database, to store the Kubernetes resources and the resource relationships. 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.

1.3.1. 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
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. 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.

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

1.4. VISUAL WEB TERMINAL

The Red Hat Advanced Cluster Management for Kubernetes Visual Web Terminal is a **tech preview** feature. You can use the Visual Web Terminal to run many commands across your cluster. Run commands and get outputs without leaving the entire console. You can run multiple commands, similar to what you can do in a standard terminal.

With Visual Web Terminal, commands that return cluster resource data are displayed visually in interactive tabular format, rather than plain text that is displayed in a standard terminal. You can see the data, click on a row, and then see more details about that resource.

By default, Visual Web Terminal includes the versions of the **oc**, **kubectl**, and **helm** CLI binaries that are equivalent to the latest OpenShift Container Platform version that is supported by Red Hat Advanced Cluster Management for the hub cluster.

You can check the versions by running the following commands:

- **oc version**
- **kubectl version**
- **helm version**

Notes:

The information that is provided in the Visual Web Terminal is limited by the permissions of the user. When you run a command, only the items that you have permission to view are displayed.

If there are over ten sessions that are running at the same time, there might be a slower response time when running commands.

The Visual Web Terminal connects to the cloud that you are already logged into with the hub cluster console.

- [Starting a session](#)
- [Supported command list](#)
- [Searching with Visual Web Terminal](#)

1.4.1. Starting a session

You can access the Visual Web Terminal by completing the following steps:

1. Log in to your console.
2. From the console *Header*, select the **Visual Web Terminal** icon.
3. Choose to open in a new tab or open in a current tab.
4. Start running commands in the input field. For example, you can start running any **oc** or **kubectl** commands and get a visual output.
5. Optional: Select the **getting started** command to pull up the *About* page and more, or type the command in the input field.
6. Optional: You can change your default namespace. Click **@default** in the status bar at the bottom of the page to run the command that displays your namespaces. View the output, then click on a row in the table to set a different namespace.
7. Use the *Supported command list* to learn about commands that you can run using **kubectl** or **oc**.

1.4.2. Supported command list

The following table provides information and examples for the command types that are supported by the Visual Web Terminal:

Command	Description	Example
!!	Run the previous command again	!!
bash commands	Run the following basic bash commands: base64, basename, cat, cksum, cp, cut, date, dirname, echo, grep, head, ls, mkdir, mv, printf, readlink, rm, sed, tail, touch, uname, vi, and vim	mkdir new_directory
cd	Change the current working directory for local filesystem changes	cd dir_name
clear	Clear the screen of previous commands and output	clear
export	Export the output of the command to an external file	export search search_criteria
helm	Run supported helm commands. Tip: See the helm commands topic in the Helm documentation for more information about the Helm commands.	helm list --chart-name
help (alias for getting started)	Display additional information about a command that is supported in the Visual Web Terminal in a side panel.	help
history	Displays the previous 25 commands that were run during the current session. Tip: You can limit the search criteria by adding it to the end of the command. For example, history help displays the last 25 commands that contained the string <i>help</i> .	history
kubectl or k	Run supported kubectl (or shorthand k) commands. Tip: See the Kubernetes CLI reference for more information about the kubectl commands.	kubectl get pods

Command	Description	Example
oc	Runs Red Hat OpenShift oc commands. Tip: See Red Hat OpenShift CLI Operations for more information about the oc commands.	oc get svc
search	Searches your Kubernetes resources for those that match the specified search filters and parameters	search pod
savedsearches	Displays your saved searches that you created and saved in the <i>Search</i> page.	savedsearches Tip: You can specify a string after the savedsearches command to limit the data that is returned. For example, savedsearches projectA displays the saved searches that contain the string <i>projectA</i> .
sleep	Pauses activity during a session. Note: Requires a sleepTime parameter that identifies the number of seconds that it pauses	sleep 20
themes	Displays the contrast theme options of the interface, which you can set to either light or dark	themes

1.4.3. Searching with Visual Web Terminal

The Visual Web Terminal search function provides visibility into your resources across all your clusters.

You can only search for resources based on your role-based access control level assignment. If you save and share a search query with another user, returned results for that user depend on his or her access level.

1. Start a Visual Web Terminal session.
2. In the command entry field of the Visual Web Terminal, type: **search**. When you run a **search** command, the Visual Web Terminal verifies that the search function is available. If it is not available, a message indicates that either the search function is not installed, or that it is just not available. If it is installed, but not available, it might be a network issue.
3. Add a space after the **search** command. The list of filters that are available for the search is displayed. The list of filters might be empty because it is dependent on the resources that are available in your environment and your role permissions.

4. Select one of the filters from the list. The selected filter is added to the search criteria on your command line, and the next level of filters for that selection are displayed. **Tip:** You can also enter a string after the **search** command, rather than selecting a filter from the list.
5. Optional: Add filters by entering a single space after each filter entry until your command contains all of the required filters.
6. When you are finished adding your search criteria, press **Enter** to run the search.

You can refine your results with specific fields. 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.

You can also search with conditions using characters, such as (>, >=, <, <=, !=).

See the following examples:

- 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 have restarted at least twice.

The search returns resources that meet the criteria of your search request, as well as links to view related resources. Viewing all of them helps you visualize how the resources interact with other resources in the system.

Your search results are grouped by **kind** and each resource **kind** is grouped in a table. You can reorder the data in the tables by selecting the row and column headings.

You can use a search that was saved in the console by entering the **savedsearches** command and selecting the search that you want to run.

1.5. MANAGING CLUSTER LABELS

Add a label to your cluster to select group resources. See [Labels and Selectors](#) for more information.

You can add new labels, remove existing labels, and edit existing labels for your clusters. Complete the following steps to manage your labels:

1. From the navigation menu, click **Automate infrastructure > Clusters**.
2. Find your cluster in the *Clusters* table.
3. Click the **Actions** icon to see more.
4. Click **Edit labels**.
5. From the *Edit labels* dialog box, enter a label. Your entry might resemble the following content: `Key=Value``.
6. Hit enter or space to add additional labels.
7. Click **Add**. Labels are saved when the **Add** button is clicked.

8. Click **Done** to close the Edit labels dialog, when you are finished.

- If you want to remove an existing label, click the **Remove** icon, next to the label in the list.
- If you want to undo a **Remove** action, click on the **Deleted** icon, which appears after **Remove** is clicked.
- If you want to update an existing label, you can reassign its key to a new value by adding a new label using the same key with a different value.

Note: You can also edit a cluster label from a cluster details page. From the navigation menu, click **Automate infrastructure > Clusters**. From the **Clusters** view, access a cluster's details page by clicking the cluster's name, then **Actions > Edit labels**.

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