OpenShift Container Platform 4.4

Support

Getting support for OpenShift Container Platform

Last Updated: 2021-02-14

OpenShift Container Platform 4.4 Support

Getting support for OpenShift Container Platform

Legal Notice

Abstract

This document provides information on getting support from Red Hat for OpenShift Container Platform. It also contains information on remote health monitoring through Telemetry and the Insights Operator.

Table of Contents

CHAPTER 1. GETTING SUPPORT 3
1.1. GETTING SUPPORT 3

CHAPTER 2. GATHERING DATA ABOUT YOUR CLUSTER 4
2.1. ABOUT THE MUST-GATHER TOOL 4
2.2. GATHERING DATA ABOUT YOUR CLUSTER FOR RED HAT SUPPORT 4
2.3. GATHERING DATA ABOUT SPECIFIC FEATURES 5
2.4. OBTAINING YOUR CLUSTER ID 9

CHAPTER 3. REMOTE HEALTH MONITORING WITH CONNECTED CLUSTERS 10
3.1. ABOUT REMOTE HEALTH MONITORING 10
3.1.1. About Telemetry 10
3.1.1.1. Information collected by Telemetry 10
3.1.2. About the Insights Operator 11
3.1.2.1. Information collected by the Insights Operator 11
3.2. SHOWING DATA COLLECTED BY REMOTE HEALTH MONITORING 12
3.2.1. Showing data collected by Telemetry 12
3.2.2. Showing data collected by the Insights Operator 13
3.3. OPTING OUT OF REMOTE HEALTH REPORTING 13
3.3.1. Consequences of disabling remote health reporting 13
3.3.2. Modifying the global cluster pull secret to disable remote health reporting 14
3.3.3. Updating the global cluster pull secret 14

CHAPTER 1. GETTING SUPPORT

1.1. GETTING SUPPORT

If you experience difficulty with a procedure described in this documentation, visit the Red Hat Customer Portal. Through the Customer Portal, you can:

- Search or browse through the Red Hat Knowledgebase of technical support articles about Red Hat products.
- Submit a support case to Red Hat Support.

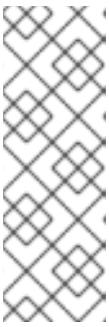

NOTE

When submitting a support case, it is recommended to provide the following information about your cluster to Red Hat Support to aid in troubleshooting:

  - Data gathered using the **oc adm must-gather** command
  - The unique cluster ID. Navigate to **(?) Help → Open Support Case** to have the cluster ID autofilled when you submit the case.

- Access other product documentation.

If you have a suggestion for improving this documentation or have found an error, please submit a Bugzilla report against the **OpenShift Container Platform** product for the **Documentation** component. Please provide specific details, such as the section name and OpenShift Container Platform version.

CHAPTER 2. GATHERING DATA ABOUT YOUR CLUSTER

When opening a support case, it is helpful to provide debugging information about your cluster to Red Hat Support.

It is recommended to provide:

- Data gathered using the **oc adm must-gather** command
- The unique cluster ID

2.1. ABOUT THE MUST-GATHER TOOL

The **oc adm must-gather** CLI command collects the information from your cluster that is most likely needed for debugging issues, such as:

- Resource definitions
- Audit logs
- Service logs

You can specify one or more images when you run the command by including the **--image** argument. When you specify an image, the tool collects data related to that feature or product.

When you run **oc adm must-gather**, a new pod is created on the cluster. The data is collected on that pod and saved in a new directory that starts with **must-gather.local**. This directory is created in the current working directory.

2.2. GATHERING DATA ABOUT YOUR CLUSTER FOR RED HAT SUPPORT

You can gather debugging information about your cluster by using the **oc adm must-gather** CLI command.

Prerequisites

- Access to the cluster as a user with the **cluster-admin** role.
- The OpenShift Container Platform CLI (**oc**) installed.

Procedure

1. Navigate to the directory where you want to store the **must-gather** data.
2. Run the **oc adm must-gather** command:

```
$ oc adm must-gather
```

NOTE

If this command fails, for example if you cannot schedule a pod on your cluster, then use the **oc adm inspect** command to gather information for particular resources. Contact Red Hat Support for the recommended resources to gather.

NOTE

If your cluster is using a restricted network you must import the default must-gather image before running the **oc adm must-gather** command.

```
$ oc import-image is/must-gather -n openshift
```

3. Create a compressed file from the **must-gather** directory that was just created in your working directory. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar cvaf must-gather.tar.gz must-gather.local.5421342344627712289/ 1
```

1 Make sure to replace **must-gather-local.5421342344627712289/** with the actual directory name.

4. Attach the compressed file to your support case on the Red Hat Customer Portal .

2.3. GATHERING DATA ABOUT SPECIFIC FEATURES

You can gather debugging information about specific features by using the **oc adm must-gather** CLI command with the **--image** or **--image-stream** argument. The **must-gather** tool supports multiple images, so you can gather data about more than one feature by running a single command.

Table 2.1. Supported must-gather images

Image	Purpose
registry.redhat.io/container-native-virtualization/cnv-must-gather-rhel8	Data collection for container-native virtualization.
registry.redhat.io/openshift-serverless-1/svls-must-gather-rhel8	Data collection for OpenShift Serverless.
registry.redhat.io/openshift-service-mesh/istio-must-gather-rhel7	Data collection for Red Hat OpenShift Service Mesh.
registry.redhat.io/rhcam-1-2/openshift-migration-must-gather-rhel8	Data collection for migration-related information.
registry.redhat.io/ocs4/ocs-must-gather-rhel8	Data collection for Red Hat OpenShift Container Storage.
registry.redhat.io/openshift4/ose-cluster-logging-operator	Data collection for Red Hat OpenShift cluster logging.

NOTE

To collect the default **must-gather** data in addition to specific feature data, add the **--image-stream=openshift/must-gather** argument.

Prerequisites

- Access to the cluster as a user with the **cluster-admin** role.
- The OpenShift Container Platform CLI (**oc**) installed.

Procedure

1. Navigate to the directory where you want to store the **must-gather** data.
2. Run the **oc adm must-gather** command with one or more **--image** or **--image-stream** arguments. For example, the following command gathers both the default cluster data and information specific to {VirtProductName}:

```
$ oc adm must-gather \
 --image-stream=openshift/must-gather \ 1
 --image=registry.redhat.io/container-native-virtualization/cnv-must-gather-rhel8 2
```

1. The default OpenShift Container Platform **must-gather** image
2. The must-gather image for {VirtProductName}

You can use the **must-gather** tool with additional arguments to gather data that is specifically related to cluster logging and the Cluster Logging Operator in your cluster. For cluster logging, run the following command:

```
$ oc adm must-gather --image=$(oc -n openshift-logging get deployment.apps/cluster-logging-operator \
 -o jsonpath='{.spec.template.spec.containers[?(@.name == "cluster-logging-operator")].image}')
```

Example 2.1. Example must-gather output for cluster logging

```
├── cluster-logging
│   ├── clo
│   │   ├── cluster-logging-operator-74dd5994f-6ttgt
│   │   ├── clusterlogforwarder_cr
│   │   ├── cr
│   │   ├── csv
│   │   ├── deployment
│   │   └── logforwarding_cr
│   ├── collector
│   │   ├── fluentd-2tr64
│   ├── curator
│   │   └── curator-1596028500-zkz4s
│   ├── eo
│   │   ├── csv
│   │   ├── deployment
│   │   └── elasticsearch-operator-7dc7d97b9d-jb4r4
│   ├── es
│   │   ├── cluster-elasticsearch
│   │   │   ├── aliases
│   │   │   ├── health
│   │   │   ├── indices
```

```
│   │   │   ├── latest_documents.json
│   │   │   ├── nodes
│   │   │   ├── nodes_stats.json
│   │   │   └── thread_pool
│   │   ├── cr
│   │   ├── elasticsearch-cdm-lp8l38m0-1-794d6dd989-4jxms
│   │   └── logs
│   │       ├── elasticsearch-cdm-lp8l38m0-1-794d6dd989-4jxms
│   ├── install
│   │   ├── co_logs
│   │   ├── install_plan
│   │   ├── olmo_logs
│   │   └── subscription
│   └── kibana
│       ├── cr
│       ├── kibana-9d69668d4-2rkvz
├── cluster-scoped-resources
│   └── core
│       ├── nodes
│       │   ├── ip-10-0-146-180.eu-west-1.compute.internal.yaml
│       └── persistentvolumes
│           ├── pvc-0a8d65d9-54aa-4c44-9ecc-33d9381e41c1.yaml
├── event-filter.html
├── gather-debug.log
└── namespaces
    ├── openshift-logging
    │   ├── apps
    │   │   ├── daemonsets.yaml
    │   │   ├── deployments.yaml
    │   │   ├── replicasets.yaml
    │   │   └── statefulsets.yaml
    │   ├── batch
    │   │   ├── cronjobs.yaml
    │   │   └── jobs.yaml
    │   ├── core
    │   │   ├── configmaps.yaml
    │   │   ├── endpoints.yaml
    │   │   ├── events
    │   │   │   ├── curator-1596021300-wn2ks.162634ebf0055a94.yaml
    │   │   │   ├── curator.162638330681bee2.yaml
    │   │   │   ├── elasticsearch-delete-app-1596020400-gm6nl.1626341a296c16a1.yaml
    │   │   │   ├── elasticsearch-delete-audit-1596020400-9l9n4.1626341a2af81bbd.yaml
    │   │   │   ├── elasticsearch-delete-infra-1596020400-v98tk.1626341a2d821069.yaml
    │   │   │   ├── elasticsearch-rollover-app-1596020400-cc5vc.1626341a3019b238.yaml
    │   │   │   ├── elasticsearch-rollover-audit-1596020400-s8d5s.1626341a31f7b315.yaml
    │   │   │   ├── elasticsearch-rollover-infra-1596020400-7mgv8.1626341a35ea59ed.yaml
    │   │   ├── events.yaml
    │   │   ├── persistentvolumeclaims.yaml
    │   │   ├── pods.yaml
    │   │   ├── replicationcontrollers.yaml
    │   │   ├── secrets.yaml
    │   │   └── services.yaml
    │   ├── openshift-logging.yaml
    │   ├── pods
    │   │   ├── cluster-logging-operator-74dd5994f-6ttgt
    │   │   │   ├── cluster-logging-operator
```

```
│   │   │   │   └── cluster-logging-operator
│   │   │   │       └── logs
│   │   │   │           ├── current.log
│   │   │   │           ├── previous.insecure.log
│   │   │   │           └── previous.log
│   │   │   └── cluster-logging-operator-74dd5994f-6ttgt.yaml
│   │   ├── cluster-logging-operator-registry-6df49d7d4-mxxff
│   │   │   ├── cluster-logging-operator-registry
│   │   │   │   └── cluster-logging-operator-registry
│   │   │   │       └── logs
│   │   │   │           ├── current.log
│   │   │   │           ├── previous.insecure.log
│   │   │   │           └── previous.log
│   │   │   ├── cluster-logging-operator-registry-6df49d7d4-mxxff.yaml
│   │   │   └── mutate-csv-and-generate-sqlite-db
│   │   │       └── mutate-csv-and-generate-sqlite-db
│   │   │           └── logs
│   │   │               ├── current.log
│   │   │               ├── previous.insecure.log
│   │   │               └── previous.log
│   │   ├── curator-1596028500-zkz4s
│   │   ├── elasticsearch-cdm-lp8l38m0-1-794d6dd989-4jxms
│   │   ├── elasticsearch-delete-app-1596030300-bpgcx
│   │   │   ├── elasticsearch-delete-app-1596030300-bpgcx.yaml
│   │   │   └── indexmanagement
│   │   │       └── indexmanagement
│   │   │           └── logs
│   │   │               ├── current.log
│   │   │               ├── previous.insecure.log
│   │   │               └── previous.log
│   │   ├── fluentd-2tr64
│   │   │   ├── fluentd
│   │   │   │   └── fluentd
│   │   │   │       └── logs
│   │   │   │           ├── current.log
│   │   │   │           ├── previous.insecure.log
│   │   │   │           └── previous.log
│   │   │   ├── fluentd-2tr64.yaml
│   │   │   └── fluentd-init
│   │   │       └── fluentd-init
│   │   │           └── logs
│   │   │               ├── current.log
│   │   │               ├── previous.insecure.log
│   │   │               └── previous.log
│   │   ├── kibana-9d69668d4-2rkvz
│   │   │   ├── kibana
│   │   │   │   └── kibana
│   │   │   │       └── logs
│   │   │   │           ├── current.log
│   │   │   │           ├── previous.insecure.log
│   │   │   │           └── previous.log
│   │   │   ├── kibana-9d69668d4-2rkvz.yaml
│   │   │   └── kibana-proxy
│   │   │       └── kibana-proxy
│   │   │           └── logs
│   │   │               ├── current.log
```

```
│ │ │ ├── previous.insecure.log
│ │ │ └── previous.log
│ └── route.openshift.io
│ └── routes.yaml
└── openshift-operators-redhat
  ├── ...
```

3. Create a compressed file from the **must-gather** directory that was just created in your working directory. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar cvaf must-gather.tar.gz must-gather.local.5421342344627712289/ 1
```

1 Make sure to replace **must-gather-local.5421342344627712289/** with the actual directory name.

4. Attach the compressed file to your support case on the Red Hat Customer Portal .

2.4. OBTAINING YOUR CLUSTER ID

When providing information to Red Hat Support, it is helpful to provide the unique identifier for your cluster. You can have your cluster ID autofilled by using the OpenShift Container Platform web console. You can also manually obtain your cluster ID by using the web console or the OpenShift CLI (**oc**).

Prerequisites

- Access to the cluster as a user with the **cluster-admin** role.
- Access to the web console or the OpenShift CLI (**oc**) installed.

Procedure

- To open a support case and have your cluster ID autofilled using the web console:
  a. From the toolbar, navigate to **(?) Help → Open Support Case**.
  b. The 'Cluster ID' value is autofilled.
- To manually obtain your cluster ID using the web console:
  a. Navigate to **Home → Dashboards → Overview**.
  b. The value is available in the **Cluster ID** field of the **Details** section.
- To obtain your cluster ID using the OpenShift CLI (**oc**), run the following command:

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

CHAPTER 3. REMOTE HEALTH MONITORING WITH CONNECTED CLUSTERS

3.1. ABOUT REMOTE HEALTH MONITORING

OpenShift Container Platform collects anonymized aggregated information about the health, usage, and size of clusters and reports it to Red Hat via two integrated components: Telemetry and the Insights Operator. This information allows Red Hat to improve OpenShift Container Platform and to react to issues that impact customers more quickly. This also simplifies the subscription and entitlement process for Red Hat customers and enables the Red Hat OpenShift Cluster Manager service to provide an overview of your clusters and their health and subscription status.

A cluster that reports data to Red Hat via Telemetry and the Insights Operator is considered a *connected cluster*.

3.1.1. About Telemetry

Telemetry sends a carefully chosen subset of the cluster monitoring metrics to Red Hat. These metrics are sent continuously and describe:

- The size of an OpenShift Container Platform cluster
- The health and status of OpenShift Container Platform components
- The health and status of any upgrade being performed
- Limited usage information about OpenShift Container Platform components and features
- Summary info about alerts reported by the cluster monitoring component

This continuous stream of data is used by Red Hat to monitor the health of clusters in real time and to react as necessary to problems that impact our customers. It also allows Red Hat to roll out OpenShift Container Platform upgrades to customers so as to minimize service impact and continuously improve the upgrade experience.

This debugging information is available to Red Hat Support and engineering teams with the same restrictions as accessing data reported via support cases. All connected cluster information is used by Red Hat to help make OpenShift Container Platform better and more intuitive to use. None of the information is shared with third parties.

3.1.1.1. Information collected by Telemetry

Primary information collected by Telemetry includes:

- The number of updates available per cluster
- Channel and image repository used for an update
- The number of errors that occurred during an update
- Progress information of running updates
- The number of machines per cluster
- The number of CPU cores and size of RAM of the machines

- The number of members in the etcd cluster and number of objects currently stored in the etcd cluster
- The number of CPU cores and RAM used per machine type - infra or master
- The number of CPU cores and RAM used per cluster
- Use of OpenShift Container Platform framework components per cluster
- The version of the OpenShift Container Platform cluster
- Health, condition, and status for any OpenShift Container Platform framework component that is installed on the cluster, for example Cluster Version Operator, Cluster Monitoring, Image Registry, and Elasticsearch for Logging
- A unique random identifier that is generated during installation
- The name of the platform that OpenShift Container Platform is deployed on, such as Amazon Web Services

Telemetry does not collect identifying information such as user names, passwords, or the names or addresses of user resources.

3.1.2. About the Insights Operator

The Insights Operator periodically gathers anonymized configuration and component failure status and reports that to Red Hat. This is a subset of the information captured by the **must-gather** tool and allows Red Hat to assess important configuration and deeper failure data than is reported via Telemetry. This data is sent several times a day and describes:

- Important configuration information about the environment that the cluster runs in
- Details about the state of the cluster and its major components
- Debugging information about infrastructure pods or nodes that are reporting failures

This debugging information is available to Red Hat Support and engineering teams with the same restrictions as accessing data reported via support cases. All connected cluster information is used by Red Hat to help make OpenShift Container Platform better and more intuitive to use. None of the information is shared with third parties.

3.1.2.1. Information collected by the Insights Operator

Primary information collected by the Insights Operator includes:

- The version of the cluster and its components, as well as the unique cluster identifier
- Channel and image repository used for an update
- Details about errors that have occurred in the cluster components
- Progress and health information of running updates and the status of any component upgrades
- Anonymized details about the cluster configuration that is relevant to Red Hat Support
- Details about any Technology Preview or unsupported configurations that might impact Red Hat Support

- Details of the platform that OpenShift Container Platform is deployed on, such as Amazon Web Services, and the region that the cluster is located in
- Information about pods of degraded OpenShift Container Platform cluster Operators
- Information about nodes marked as **NotReady**
- Events for all namespaces listed as "related objects" for Degraded operator
- Anonymized certificate signing requests (CSR) and information about the validity of certificates

The Insights Operator does not collect identifying information such as user names, passwords, or the names or addresses of user resources.

3.2. SHOWING DATA COLLECTED BY REMOTE HEALTH MONITORING

As an administrator, you can review the metrics collected by Telemetry and the Insights Operator.

3.2.1. Showing data collected by Telemetry

You can see the cluster and components time series data captured by Telemetry.

Prerequisites

- Install the OpenShift CLI (**oc**).
- You must log in to the cluster with a user that has the **cluster-admin** role.

Procedure

1. Find the URL for the Prometheus service that runs in the OpenShift Container Platform cluster:

```
$ oc get route prometheus-k8s -n openshift-monitoring -o jsonpath="{.spec.host}"
```

2. Navigate to the URL.
3. Enter this query in the **Expression** input box and press **Execute**:

```
{__name__="up"} or {__name__="cluster_version"} or
{__name__="cluster_version_available_updates"} or {__name__="cluster_operator_up"} or
{__name__="cluster_operator_conditions"} or {__name__="cluster_version_payload"} or
{__name__="cluster_version_payload_errors"} or
{__name__="instance:etcd_object_counts:sum"} or
{__name__="ALERTS",alertstate="firing"} or
{__name__="code:apiserver_request_count:rate:sum"} or
{__name__="kube_pod_status_ready:etcd:sum"} or
{__name__="kube_pod_status_ready:image_registry:sum"} or
{__name__="cluster:capacity_cpu_cores:sum"} or
{__name__="cluster:capacity_memory_bytes:sum"} or
{__name__="cluster:cpu_usage_cores:sum"} or
{__name__="cluster:memory_usage_bytes:sum"} or
{__name__="openshift:cpu_usage_cores:sum"} or
{__name__="openshift:memory_usage_bytes:sum"} or
{__name__="cluster:node_instance_type_count:sum"}
```

This query replicates the request that Telemetry makes against a running OpenShift Container Platform cluster's Prometheus service and returns the full set of time series captured by Telemetry.

3.2.2. Showing data collected by the Insights Operator

You can review the data that is collected by the Insights Operator.

Prerequisites

- Access to the cluster as a user with the **cluster-admin** role.

Procedure

1. Find the name of the currently running pod for the Insights Operator:

```
$ INSIGHTS_OPERATOR_POD=$(oc get pods --namespace=openshift-insights -o custom-columns=:metadata.name --no-headers  --field-selector=status.phase=Running)
```

2. Copy the recent data archives collected by the Insights Operator:

```
$ oc cp openshift-insights/$INSIGHTS_OPERATOR_POD:/var/lib/insights-operator ./insights-data
```

The recent Insights Operator archives are now available in the **insights-data** directory.

3.3. OPTING OUT OF REMOTE HEALTH REPORTING

You might need to opt out of reporting health and usage data for your cluster. For example, you might need to comply with privacy laws or standards governing how your organization reports monitoring data.

To opt out of remote health reporting, you must:

1. Modify the global cluster pull secret to disable remote health reporting.
2. Update the cluster to use this modified pull secret.

3.3.1. Consequences of disabling remote health reporting

In OpenShift Container Platform, customers can opt out of reporting health and usage information. However, connected clusters allow Red Hat to react more quickly to problems and better support our customers, as well as better understand how product upgrades impact clusters.

Red Hat strongly recommends leaving health and usage reporting enabled for pre-production and test clusters even if it is necessary to opt out for production clusters. This allows Red Hat to be a participant in qualifying OpenShift Container Platform in your environments and react more rapidly to product issues.

Some of the consequences of opting out of having a connected cluster are:

- Red Hat will not be able to monitor the success of product upgrades or the health of your clusters without a support case being opened.

- Red Hat will not be able to use anonymized configuration data to better triage customer support cases and identify which configurations our customers find important.
- The Red Hat OpenShift Cluster Manager will not show data about your clusters including health and usage information.
- Your subscription entitlement information must be manually entered via cloud.redhat.com without the benefit of automatic usage reporting.

In restricted networks, Telemetry and Insights data can still be reported through appropriate configuration of your proxy.

3.3.2. Modifying the global cluster pull secret to disable remote health reporting

You can modify your existing global cluster pull secret to disable remote health reporting. This disables both Telemetry and the Insights Operator.

Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.

Procedure

1. Download the global cluster pull secret to your local file system.

```
$ oc extract secret/pull-secret -n openshift-config --to=.
```

2. In a text editor, edit the **.dockerconfigjson** file that was downloaded.
3. Remove the **cloud.openshift.com** JSON entry, for example:

```
"cloud.openshift.com":{"auth":"<hash>","email":"<email_address>"}
```

4. Save the file.

You can now update your cluster to use this modified pull secret.

3.3.3. Updating the global cluster pull secret

You can update the global pull secret for your cluster.

WARNING

Cluster resources must adjust to the new pull secret, which can temporarily limit the usability of the cluster.

Prerequisites

- You have a new or modified pull secret file to upload.

- You have access to the cluster as a user with the **cluster-admin** role.

Procedure

- Run the following command to update the global pull secret for your cluster:

```
$ oc set data secret/pull-secret -n openshift-config --from-file=.dockerconfigjson=<pull-secret-location> 1
```

1 Provide the path to the new pull secret file.

This update is rolled out to all nodes, which can take some time depending on the size of your cluster. During this time, nodes are drained and pods are rescheduled on the remaining nodes.