



OpenShift Container Platform 4.20

Power Monitoring

Configuring and using power monitoring for OpenShift Container Platform

OpenShift Container Platform 4.20 Power Monitoring

Configuring and using power monitoring for OpenShift Container Platform

Legal Notice

Copyright © Red Hat.

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, 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 Software Collections 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

Use power monitoring to monitor the power consumption for various components, such as CPU and DRAM, for each container running in an OpenShift Container Platform cluster.

Table of Contents

CHAPTER 1. POWER MONITORING FOR RED HAT OPENSIFT 0.5 (TECHNOLOGY PREVIEW) RELEASE NOTES	4
1.1. POWER MONITORING 0.5 (TECHNOLOGY PREVIEW) RELEASE NOTES OVERVIEW	4
1.1.1. Power monitoring Technology Preview 0.5 new features	4
1.1.2. Power monitoring Technology Preview 0.5 enhancements	4
1.1.2.1. Updated dashboards	5
1.1.3. Power monitoring 0.5 (Technology Preview) deprecated and removed features	5
1.1.4. Power monitoring 0.5 (Technology Preview) support tables	5
CHAPTER 2. POWER MONITORING FOR RED HAT OPENSIFT RELEASE NOTES	7
2.1. POWER MONITORING 0.4 (TECHNOLOGY PREVIEW)	7
2.1.1. Features	7
2.1.2. Bug fixes	7
2.1.3. CVEs	7
2.2. POWER MONITORING 0.3 (TECHNOLOGY PREVIEW)	8
2.2.1. Bug fixes	8
2.2.2. CVEs	8
2.3. POWER MONITORING 0.2 (TECHNOLOGY PREVIEW)	8
2.3.1. Features	8
2.3.2. CVEs	9
2.4. POWER MONITORING 0.1 (TECHNOLOGY PREVIEW)	9
2.4.1. Features	9
CHAPTER 3. POWER MONITORING OVERVIEW	10
3.1. ABOUT POWER MONITORING	10
3.2. POWER MONITORING ARCHITECTURE	10
3.3. KEPLER HARDWARE SUPPORT	10
3.4. ABOUT FIPS COMPLIANCE FOR POWER MONITORING OPERATOR	11
3.5. ADDITIONAL RESOURCES	11
CHAPTER 4. INSTALLING POWER MONITORING FOR RED HAT OPENSIFT	12
4.1. INSTALLING THE POWER MONITORING OPERATOR	12
4.2. DEPLOYING POWERMONITOR CUSTOM RESOURCE	13
CHAPTER 5. CONFIGURING POWER MONITORING	14
5.1. THE KEPLER CONFIGURATION	14
5.2. MONITORING THE KEPLER STATUS	16
CHAPTER 6. VISUALIZING POWER MONITORING METRICS	17
6.1. POWER MONITORING DASHBOARDS OVERVIEW	17
6.1.1. Power Monitor / Overview dashboard	17
6.1.2. Power Monitor / Namespace (Pods) dashboard	17
6.2. ACCESSING POWER MONITORING DASHBOARDS AS A CLUSTER ADMINISTRATOR	18
6.3. ACCESSING POWER MONITORING DASHBOARDS AS A DEVELOPER	18
6.4. ADDITIONAL RESOURCES	19
CHAPTER 7. UNINSTALLING POWER MONITORING	20
7.1. DELETING KEPLER	20
7.2. DELETING THE POWERMONITOR CUSTOM RESOURCE	20
7.3. UNINSTALLING THE POWER MONITORING OPERATOR	21
CHAPTER 8. POWER MONITORING REFERENCES	22
8.1. POWER MONITORING KEPLER POWER ATTRIBUTION GUIDE	22

- CHAPTER 9. POWER MONITORING API REFERENCE 23
 - 9.1. POWERMONITORING API SPECIFICATIONS 23
 - 9.1.1. PowerMonitor.spec 23
 - 9.1.2. PowerMonitor.status.conditions 24
 - 9.1.3. PowerMonitor.status.kepler 24


CHAPTER 1. POWER MONITORING FOR RED HAT OPENSIFT

0.5 (TECHNOLOGY PREVIEW) RELEASE NOTES

These release notes track the development of power monitoring for Red Hat OpenShift in the OpenShift Container Platform.

For an overview of the Power Monitoring Operator, see [About power monitoring](#).

1.1. POWER MONITORING 0.5 (TECHNOLOGY PREVIEW) RELEASE NOTES OVERVIEW



IMPORTANT

Power monitoring is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

Power monitoring for Red Hat OpenShift enables you to monitor the power usage of workloads and identify the most power-consuming namespaces running in an OpenShift Container Platform cluster with key power consumption metrics, such as CPU or DRAM, measured at container level.

This release of power monitoring and the Power Monitoring Operator provides more accurate data, includes new dashboards, and removes some features and functionality.

This release of power monitoring and the Power Monitoring Operator is supported on:

- OpenShift Container Platform 4.17+
- Bare metal deployments

1.1.1. Power monitoring Technology Preview 0.5 new features

This release of power monitoring for Red Hat OpenShift and the Power Monitoring Operator, based on the Kepler Project, includes the following new feature:

- Deployment and deletion of **PowerMonitor** custom resource definition (CRD).

1.1.2. Power monitoring Technology Preview 0.5 enhancements

This release of power monitoring for Red Hat OpenShift and the Power Monitoring Operator, based on the Kepler Project, includes the following enhancements:

- Dynamic detection of Nodes Running Average Power Limit (RAPL) zones
- More accurate power measurement based on active CPU usage
- Improved Virtual Machine (VM), container, and pod detection

- More relevant label values for processes, containers, VMs, and pods
- Requires only **readonly** access to host: **/proc** and **/sys**
 - No more **CAP_SYSADMIN** and **CAP_BPF**
- Significantly reduced resource usage compared to earlier Kepler implementations
- Multi-level energy tracking for the following levels:
 - node
 - process
 - container
 - VM
 - pod
- Terminated workload tracking with configurable retention policies
- Energy-based prioritization for terminated resources
- Real-time data collection with configurable intervals and staleness detection

1.1.2.1. Updated dashboards

With this update, power monitoring for Red Hat OpenShift has the following dashboard changes:

- Updated **Power Monitor / Overview** dashboard.
- Updated **Power Monitor / Namespace (Pods)** dashboard.



IMPORTANT

The older metrics and dashboards are no longer supported. If you are managing your own custom dashboard or queries, you need to update to the newer versions.

1.1.3. Power monitoring 0.5 (Technology Preview) deprecated and removed features

- In the Red Hat OpenShift power monitoring technology preview 0.5 release, the **Kepler** custom resource has been deprecated, and will be removed in a future release. Use the **PowerMonitor** custom resource instead.
- In the Red Hat OpenShift power monitoring technology preview 0.5 release, the Redfish configuration has been removed. It is no longer supported in previous versions of power monitoring.

1.1.4. Power monitoring 0.5 (Technology Preview) support tables

This release includes the following support updates:

Table 1.1. Power Monitoring Operator supported version table

Kepler	0.10.2
Power Monitoring Operator	0.20.0

Table 1.2. Power monitoring supported platforms


OpenShift Container Platform	4.17+
Bare metal	X



IMPORTANT

Installations in virtual machines are not supported and will not function.

CHAPTER 2. POWER MONITORING FOR RED HAT OPENSIFT RELEASE NOTES



IMPORTANT

Power monitoring is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

Power monitoring for Red Hat OpenShift enables you to monitor the power usage of workloads and identify the most power-consuming namespaces running in an OpenShift Container Platform cluster with key power consumption metrics, such as CPU or DRAM, measured at container level.

These release notes track the development of power monitoring for Red Hat OpenShift in the OpenShift Container Platform.

For an overview of the Power Monitoring Operator, see [About power monitoring](#).

2.1. POWER MONITORING 0.4 (TECHNOLOGY PREVIEW)

This release includes the following version updates:

- Kepler 0.7.12
- Power Monitoring Operator 0.15.0

2.1.1. Features

- With this release, FIPS mode is now enabled for power monitoring for Red Hat OpenShift. When installed on an OpenShift Container Platform cluster in FIPS mode, Power Monitoring Operator ensures compatibility without affecting the FIPS support status of the cluster.

2.1.2. Bug fixes

- Before this update, the **Install** screen in the OperatorHub page for power monitoring for Red Hat OpenShift displayed incorrect documentation links. With this update, the links now direct to the correct paths.

2.1.3. CVEs

- [CVE-2024-24790](#)
- [CVE-2024-24791](#)
- [CVE-2024-34155](#)
- [CVE-2024-34158](#)

2.2. POWER MONITORING 0.3 (TECHNOLOGY PREVIEW)

This release includes the following version updates:

- Kepler 0.7.12
- Power Monitoring Operator 0.15.0

The following advisory is available for power monitoring 0.3:

- [RHEA-2024:9961](#)

2.2.1. Bug fixes

- Before this update, the Power Monitoring Operator dashboard used an invalid Prometheus rule, which caused the panel for **OTHER Power Consumption(W) by Pods** to display incorrect data. With this update, the rule is corrected, ensuring the dashboard now shows accurate power consumption data.

2.2.2. CVEs

- [CVE-2023-37920](#)
- [CVE-2024-2236](#)
- [CVE-2024-2511](#)
- [CVE-2024-3596](#)
- [CVE-2024-4603](#)
- [CVE-2024-4741](#)
- [CVE-2024-5535](#)
- [CVE-2024-34397](#)

2.3. POWER MONITORING 0.2 (TECHNOLOGY PREVIEW)

This release includes the following version updates:

- Kepler 0.7.10
- Power Monitoring Operator 0.13.0

The following advisory is available for power monitoring 0.2:

- [RHEA-2024:2923](#)

2.3.1. Features

- With this update, you can enable the Redfish API in Kepler. Kepler uses Redfish to get the power consumed by nodes.

- With this update, you can install the Power Monitoring Operator in the namespace of your choice.
- With this update, you can access the power monitoring **Overview** dashboard from the **Developer** perspective.

2.3.2. CVEs

- [CVE-2022-48554](#)
- [CVE-2023-2975](#)
- [CVE-2023-3446](#)
- [CVE-2023-3817](#)
- [CVE-2023-5678](#)
- [CVE-2023-6129](#)
- [CVE-2023-6237](#)
- [CVE-2023-7008](#)
- [CVE-2024-0727](#)
- [CVE-2024-25062](#)
- [CVE-2024-28834](#)
- [CVE-2024-28835](#)

2.4. POWER MONITORING 0.1 (TECHNOLOGY PREVIEW)

This release introduces a Technology Preview version of power monitoring for Red Hat OpenShift. The following advisory is available for power monitoring 0.1:

- [RHEA-2024:0078](#)

2.4.1. Features

- Deployment and deletion of Kepler
- Power usage metrics from Intel-based bare-metal deployments
- Dashboards for plotting power usage

CHAPTER 3. POWER MONITORING OVERVIEW



IMPORTANT

Power monitoring is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

3.1. ABOUT POWER MONITORING

You can use power monitoring for Red Hat OpenShift to monitor the power usage and identify power-consuming containers running in an OpenShift Container Platform cluster. Power monitoring collects and exports energy-related system statistics from various components, such as CPU and DRAM. It provides estimates and granular power consumption data for Kubernetes pods and namespaces, and reads the power consumption of nodes.



WARNING

Power monitoring Technology Preview works only in bare-metal deployments. Most public cloud vendors do not expose Kernel Power Management Subsystems to virtual machines.

3.2. POWER MONITORING ARCHITECTURE

Power monitoring is made up of the following major components:

The Power Monitoring Operator

For administrators, the Power Monitoring Operator streamlines the monitoring of power usage for workloads by simplifying the deployment and management of Kepler in an OpenShift Container Platform cluster. The setup and configuration for the Power Monitoring Operator are simplified by adding a **PowerMonitor** custom resource definition (CRD). The Operator also manages operations, such as upgrading, removing, configuring, and redeploying Kepler.

Kepler

Kepler is a key component of power monitoring. It is responsible for monitoring the power usage of containers running in OpenShift Container Platform. It generates metrics related to the power usage of both nodes and containers.

3.3. KEPLER HARDWARE SUPPORT

Kepler is the key component of power monitoring that collects real-time CPU power consumption data from a node through the RAPL Subsystem. By understanding the total power consumption of the node and calculating the percent of CPU time each process is using, it is able to estimate the power consumption at a per process and container level.

Kernel Power Management Subsystem

- **rapl-sysfs:** This requires access to the `/sys/class/powercap/intel-rapl` directory.

3.4. ABOUT FIPS COMPLIANCE FOR POWER MONITORING OPERATOR

Starting with version 0.4, Power Monitoring Operator for Red Hat OpenShift is FIPS compliant. When deployed on an OpenShift Container Platform cluster in FIPS mode, it uses Red Hat Enterprise Linux (RHEL) cryptographic libraries validated by National Institute of Standards and Technology (NIST).

For details on the NIST validation program, see [Cryptographic module validation program](#). For the latest NIST status of RHEL cryptographic libraries, see [Compliance activities and government standards](#).

To enable FIPS mode, you must install Power Monitoring Operator for Red Hat OpenShift on an OpenShift Container Platform cluster. For more information, see "Do you need extra security for your cluster?".

3.5. ADDITIONAL RESOURCES

- [Power monitoring dashboards overview](#)
- [Do you need extra security for your cluster?](#)

CHAPTER 4. INSTALLING POWER MONITORING FOR RED HAT OPENSIFT



IMPORTANT

Power monitoring is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

You can install power monitoring for Red Hat OpenShift by deploying the Power Monitoring Operator in the OpenShift Container Platform web console.

4.1. INSTALLING THE POWER MONITORING OPERATOR

As a cluster administrator, you can install the Power Monitoring Operator from the software catalog by using the OpenShift Container Platform web console.



WARNING

You must remove any previously installed versions of the Power Monitoring Operator before installation.

Prerequisites

- You have access to the OpenShift Container Platform web console.
- You are logged in as a user with the **cluster-admin** role.

Procedure


1. In the web console, go to **Ecosystem → Software Catalog**.
2. Search for **power monitoring**, click the **Power monitoring for Red Hat OpenShift** tile, and then click **Install**.
3. Click **Install** again to install the Power Monitoring Operator.
Power monitoring for Red Hat OpenShift is now available in all namespaces of the OpenShift Container Platform cluster.

Verification

1. Verify that the Power Monitoring Operator is listed in **Ecosystem → Installed Operators**. The **Status** should resolve to **Succeeded**.

4.2. DEPLOYING POWERMONITOR CUSTOM RESOURCE

You can deploy Kepler by creating an instance of the **PowerMonitor** custom resource (CR) using the Power Monitoring Operator.



IMPORTANT


The **Kepler** custom resource definition (CRD) has been deprecated and will be removed in a future release. Use the **PowerMonitor** custom resource instead.

Prerequisites

- You have access to the OpenShift Container Platform web console.
- You are logged in as a user with the **cluster-admin** role.
- You have installed the Power Monitoring Operator.

Procedure

1. In the web console, go to **Ecosystem → Installed Operators**.
2. Click **Power monitoring for Red Hat OpenShift** from the **Installed Operators** list and go to the **PowerMonitor** tab.
3. Click **Create PowerMonitor**.
4. On the **Create PowerMonitor** page, ensure the **Name** is set to **power-monitor**.



IMPORTANT

The name of your **PowerMonitor** instance must be set to **power-monitor**. All other instances are ignored by the Power Monitoring Operator.

5. Click **Create** to deploy the PowerMonitor and power monitoring dashboards.

CHAPTER 5. CONFIGURING POWER MONITORING



IMPORTANT

Power monitoring is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

The **PowerMonitor** resource is a Kubernetes custom resource definition (CRD) that enables you to configure the deployment and monitor the status of the **PowerMonitor** resource.

5.1. THE KEPLER CONFIGURATION

You can configure Kepler with the **spec** field of the **PowerMonitor** resource.



IMPORTANT

Ensure that the name of your **PowerMonitor** instance is **power-monitor**. All other instances are rejected by the Power Monitoring Operator Webhook.

The following is the list of configuration options:

Table 5.1. PowerMonitor configuration options

Name	Description	Default Behavior
deployment.no deSelector	The nodes on which Kepler (created by PowerMonitor) pods are scheduled.	kubernetes.io/os: linux
deployment.tol erations	The tolerations for Power Monitor that allow the pods to be scheduled on nodes with specific characteristics.	- operator: "Exists"
deployment.se curity.mode	Security mode can be set to either none , allowing unrestricted access to Kepler's metrics by any entity, or rbac , securing the metrics endpoint with TLS encryption and restricting access to authorized service accounts listed in allowedSANames .	Set to rbac by default and only user workload prometheus is allowed access.
deployment.se curity.allowedS ANames	A list of Service Account Names that can access Kepler's metrics endpoint when security mode is rbac .	In OpenShift, set to openshift-user-workload-monitoring:prometheus-user-workload to allow user workload monitoring to scrape Kepler.

Name	Description	Default Behavior
config.logLevel	The level of logs to expose by Kepler.	Set to info.
config.metricLevels	A list of energy metric levels to expose. Possible values include node , process , container , vm , and pod .	The default list includes node , pod , and vm .
config.staleness	Specifies how long to wait before considering calculated power values as stale.	500ms (500 milliseconds).
config.sampleRate	Specifies the interval for monitoring resources such as processes, containers, and VMs.	5s (5 seconds).
config.maxTerminated	Controls terminated workload tracking. A negative value tracks unlimited workloads, zero disables tracking, and a positive value tracks the top N terminated workloads by energy consumption.	500.

Example PowerMonitor resource with default configuration

```

apiVersion: v1alpha1
kind: PowerMonitor
metadata:
  labels:
    app.kubernetes.io/name: powermonitor
    app.kubernetes.io/instance: powermonitor
    app.kubernetes.io/part-of: kepler-operator
  name: power-monitor
spec:
  kepler:
    deployment:
      nodeSelector:
        kubernetes.io/os: linux

    tolerations:
      - key: key1
        operator: Equal
        value: value1
        effect: NoSchedule

    security:
      mode: rbac
      allowedSANames:
        - openshift-user-workload-monitoring:prometheus-user-workload

    config:
      logLevel: info
      metricLevels: [node, pod, vm]
```

```

staleness: 1s
sampleRate: 10s
maxTerminated: 1000

```

5.2. MONITORING THE KEPLER STATUS

You can monitor the state of the Kepler exporter with the **status** field of the **PowerMonitor** resource.

The **status** field includes information, such as the following:

- The number of nodes currently running the Kepler pods
- The number of nodes that should be running the Kepler pods
- Conditions representing the health of the Kepler resource

This provides you with valuable insights into the changes made through the **spec** field.

Example state of the **PowerMonitor** resource


```

apiVersion: kepler.system.sustainable.computing.io/v1alpha1
kind: PowerMonitor
metadata:
  name: power-monitor
status:
  conditions: 1
  - lastTransitionTime: '2024-01-11T11:07:39Z'
    message: Reconcile succeeded
    observedGeneration: 1
    reason: ReconcileSuccess
    status: 'True'
    type: Reconciled
  - lastTransitionTime: '2024-01-11T11:07:39Z'
    message: >-
      power-monitor daemonset "openshift-power-monitoring/power-monitor" is deployed to all nodes
    and
      available; ready 2/2
    observedGeneration: 1
    reason: DaemonSetReady
    status: 'True'
    type: Available
  currentNumberScheduled: 2 2
  desiredNumberScheduled: 2 3

```

- 1 The health of the **PowerMonitor** resource. In this example, the **PowerMonitor** resource is successfully reconciled and ready.
- 2 The number of nodes currently running the Kepler pods is 2.
- 3 The wanted number of nodes to run the Kepler pods is 2.

CHAPTER 6. VISUALIZING POWER MONITORING METRICS



IMPORTANT

Power monitoring is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

You can visualize power monitoring metrics in the OpenShift Container Platform web console by accessing power monitoring dashboards or by exploring **Metrics** under the **Observe** tab.

6.1. POWER MONITORING DASHBOARDS OVERVIEW

There are two types of power monitoring dashboards. Both provide different levels of details around power consumption metrics for a single cluster:

6.1.1. Power Monitor / Overview dashboard

This dashboard allows you to view the following information:

Cluster-wide power consumption

View current total, active, and idle CPU power consumption, grouped by zones.

Node-level power details

Analyze historical and current power consumption (total, active, and idle) for individual nodes.

Hardware information

Display CPU model and core counts for each node in the cluster.

Time-series analysis

Track power consumption trends over time with graphs that can be filtered by node and zone. This provides a comprehensive view of your cluster's energy usage.

6.1.2. Power Monitor / Namespace (Pods) dashboard

This dashboard allows you to monitor and analyze power consumption for Kubernetes namespaces and pods. It provides the following information:

Top ten power consuming namespaces

A real-time table showing the top ten namespaces based on their current power usage. This helps you quickly identify the most resource-intensive workloads.

Total namespace power consumption

A historical graph showing the total power consumption of pods within a selected namespace over time, grouped by zone. This helps you see trends and understand an application's or service's total power use.

Individual pod power consumption

A detailed graph showing the power consumption of individual pods, so you can analyze them in detail.

6.2. ACCESSING POWER MONITORING DASHBOARDS AS A CLUSTER ADMINISTRATOR

You can access power monitoring dashboards of the OpenShift Container Platform web console.

Prerequisites

- You have access to the OpenShift Container Platform web console.
- You are logged in as a user with the **cluster-admin** role.
- You have installed the Power Monitoring Operator.
- You have deployed Kepler in your cluster.
- You have enabled monitoring for user-defined projects.

Procedure

1. In the web console, go to **Observe → Dashboards**.
2. From the **Dashboard** drop-down list, select the power monitoring dashboard you want to see:
 - **Power Monitor / Overview**
 - **Power Monitor / Namespace (Pods)**

6.3. ACCESSING POWER MONITORING DASHBOARDS AS A DEVELOPER

You can access power monitoring dashboards from OpenShift Container Platform web console.

Prerequisites

- You have access to the OpenShift Container Platform web console.
- You have access to the cluster as a developer or as a user.
- You have installed the Power Monitoring Operator.
- You have deployed Kepler in your cluster.
- You have enabled monitoring for user-defined projects.
- You have **view** permissions for the namespace **openshift-power-monitoring**, the namespace where Kepler is deployed to.

Procedure

1. In the web console, go to **Observe → Dashboard**.


2. From the **Dashboard** drop-down list, select the power monitoring dashboard you want to see:

- **Power Monitor / Overview**

6.4. ADDITIONAL RESOURCES

- [Enabling monitoring for user-defined projects](#)

CHAPTER 7. UNINSTALLING POWER MONITORING



IMPORTANT


Power monitoring is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

You can uninstall power monitoring by deleting the Kepler instance and then the Power Monitoring Operator in the OpenShift Container Platform web console.

7.1. DELETING KEPLER

You can delete Kepler by removing the Kepler instance of the **Kepler** custom resource definition (CRD) from the OpenShift Container Platform web console.



IMPORTANT

Starting with power monitoring for Red Hat OpenShift 0.5 (Technology Preview), use the **PowerMonitor** CRD, and remove all instances of the **Kepler** CRD.

Prerequisites

- You have access to the OpenShift Container Platform web console.
- You are logged in as a user with the **cluster-admin** role.

Procedure

1. In the web console, go to **Ecosystem → Installed Operators**.
2. Click **Power monitoring for Red Hat OpenShift** from the **Installed Operators** list and go to the **Kepler** tab.
3. Locate the Kepler instance entry in the list.
4. Click  for this entry and select **Delete Kepler**.
5. In the **Delete Kepler?** dialog, click **Delete** to delete the Kepler instance.

7.2. DELETING THE POWERMONITOR CUSTOM RESOURCE

You can delete the **PowerMonitor** custom resource (CR) by removing the **power-monitor** instance of the **PowerMonitor** CR from the OpenShift Container Platform web console.

Prerequisites

- You have access to the OpenShift Container Platform web console.
- You are logged in as a user with the **cluster-admin** role.

Procedure

1. In the web console, go to **Ecosystem → Installed Operators**.
2. Click **Power monitoring for Red Hat OpenShift** from the **Installed Operators** list and go to the **PowerMonitor** tab.
3. Locate the **PowerMonitor** instance entry in the list.
4. Click the  for this entry and select **Delete PowerMonitor**.
5. In the **Delete PowerMonitor?** dialog, click **Delete** to delete the **PowerMonitor** instance.

7.3. UNINSTALLING THE POWER MONITORING OPERATOR

If you installed the Power Monitoring Operator by using the software catalog, you can uninstall it from the OpenShift Container Platform web console.

Prerequisites

- You have access to the OpenShift Container Platform web console.
- You are logged in as a user with the **cluster-admin** role.

Procedure

1. Delete the Kepler instance.



WARNING

Ensure that you have deleted the Kepler instance before uninstalling the Power Monitoring Operator.

2. Go to **Ecosystem → Installed Operators**.
3. Locate the **Power monitoring for Red Hat OpenShift** entry in the list.
4. Click  for this entry and select **Uninstall Operator**.
5. In the **Uninstall Operator?** dialog, click **Uninstall** to uninstall the Power Monitoring Operator.

CHAPTER 8. POWER MONITORING REFERENCES



IMPORTANT

Power monitoring is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

8.1. POWER MONITORING KEPLER POWER ATTRIBUTION GUIDE

Kepler's power attribution system provides practical, proportional distribution of hardware energy consumption to individual workloads. While CPU-time-based attribution has inherent limitations due to modern CPU complexity, it offers a good balance between accuracy, simplicity, and performance overhead for most monitoring and optimization use cases.

For more information about power attribution, see [Kepler Power Attribution Guide](#).

CHAPTER 9. POWER MONITORING API REFERENCE



IMPORTANT

Power monitoring is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

PowerMonitor is the Schema for the PowerMonitor API.

9.1. POWERMONITORING API SPECIFICATIONS

PowerMonitor

PowerMonitor is the schema for the PowerMonitor API.

Name	Type	Description	Required
apiVersion	string	kepler.system.sustainable.computing.io/v1alpha1	true
kind	string	PowerMonitor	true
object	Refer to the Kubernetes API documentation for the fields of the metadata field.	true	spec
object	PowerMonitorSpec defines the desired state of Power Monitor	false	status

9.1.1. PowerMonitor.spec

PowerMonitorSpec defines the desired state of Power Monitor

Name	Type	Description	Required
kepler	object		true

9.1.2. PowerMonitor.status.conditions

Name	Type	Description	Required
lastTransitionTime	string	The last time the condition transitioned from one status to another. This should be when the underlying condition changed. If that is not known, then using the time when the API field changed is acceptable. Format: date-time	true
message	string	A human-readable message indicating details about the transition. This may be an empty string.	true
reason	string	Contains a programmatic identifier indicating the reason for the condition's last transition.	true
status	string	The status of the condition, which can be one of True, False, or Unknown.	true
type	string	The type of Kepler Condition, such as Reconciled or Available.	true
observedGeneration	integer	Represents the .metadata.generation that the condition was set based upon. For instance, if .metadata.generation is currently 12, but the .status.conditions[x].observedGeneration is 9, the condition is out of date. Format: int64 Minimum: 0	false

9.1.3. PowerMonitor.status.kepler

Name	Type	Description	Required
currentNumberScheduled	integer	The number of nodes that are running at least one power-monitor pod and are supposed to run it. Format: int32	true
desiredNumberScheduled	integer	The total number of nodes that should be running the power-monitor pod. Format: int32	true
numberMisscheduled	integer	The number of nodes running the power-monitor pod that are not supposed to. Format: int32	true

Name	Type	Description	Required
numberReady	integer	The number of nodes that should be running the power-monitor pod and have at least one pod with a Ready condition. Format: int32	true
numberAvailable	integer	The number of nodes that should be running the power-monitor pod and have at least one pod running and available. Format: int32	false
numberUnavailable	integer	The number of nodes that should be running the power-monitor pod but have no pods running and available. Format: int32	false
updatedNumberScheduled	integer	The total number of nodes that are running an updated power-monitor pod. Format: int32	false