



Red Hat Ceph Storage 5

Monitoring Ceph with Datadog Guide

Guide on Monitoring Ceph with Datadog

Red Hat Ceph Storage 5 Monitoring Ceph with Datadog Guide

Guide on Monitoring Ceph with Datadog

Legal Notice

Copyright © 2022 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

XFS[®] is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

This document provides information on monitoring the status of the Ceph Storage cluster with the Datadog monitoring tool. Red Hat is committed to replacing problematic language in our code, documentation, and web properties. We are beginning with these four terms: master, slave, blacklist, and whitelist. Because of the enormity of this endeavor, these changes will be implemented gradually over several upcoming releases. For more details, see our CTO Chris Wright's message.

Table of Contents

CHAPTER 1. MONITORING DATADOG AND CEPH	3
CHAPTER 2. INSTALLING DATADOG FOR CEPH INTEGRATION	4
CHAPTER 3. INSTALLING AND CONFIGURING THE DATADOG AGENT FOR CEPH	7
CHAPTER 4. VIEWING THE CEPH OVERVIEW WITH DATADOG	8
CHAPTER 5. CEPH METRICS FOR DATADOG	9
CHAPTER 6. CREATING ALERTS IN DATADOG	11

CHAPTER 1. MONITORING DATADOG AND CEPH

The [Datadog](#) integration with Ceph enables Datadog to execute and process the output from:

- **ceph status**
- **ceph health detail**
- **ceph df detail**
- **ceph osd perf**; and,
- **ceph osd pool stats**.

The integration enables Datadog to:

- Monitor the status and health of the Red Hat Ceph Storage cluster.
- Monitor I/O and performance metrics.
- Track disk usage across storage pools.

Using Datadog

Using [Datadog](#) to monitor Ceph requires installing a Datadog agent on at least one Ceph monitor node. When monitoring Ceph, the Datadog agent will execute Ceph command line arguments. Consequently, each Ceph node must have an appropriate Ceph key providing access to the cluster, usually in `/etc/ceph`. Once the agent executes the Ceph command, it sends Red Hat Ceph Storage cluster status and statistics back to Datadog. Then, Datadog will present the status and statistics in the Datadog user interface.

Since Datadog uses an agent, the Red Hat Ceph Storage cluster must be able to reach the internet. However, the Red Hat Ceph Storage cluster does not have to be reachable from the internet.



NOTE

Datadog supports retrieving **ceph status** with Red Hat Ceph Storage version 2 or higher. Datadog will provide an update to support **ceph status** for Red Hat Ceph Storage cluster 3 in a subsequent release of its **dd-agent**.



IMPORTANT

Red Hat works with our technology partners to provide this documentation as a service to our customers. However, Red Hat does not provide support for this product. If you need technical assistance for this product, then contact Datadog for support.

CHAPTER 2. INSTALLING DATADOG FOR CEPH INTEGRATION

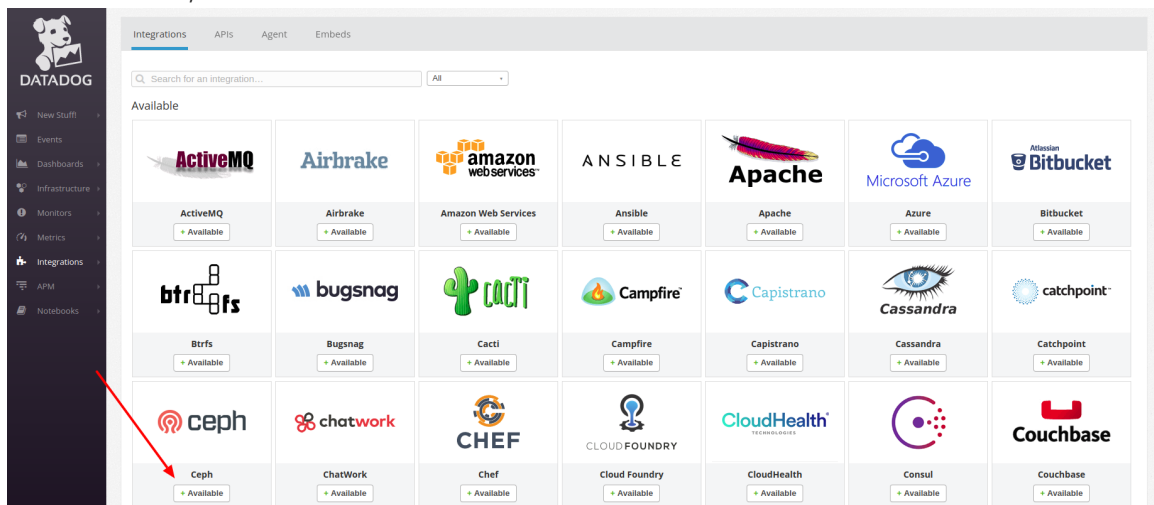
After installing the Datadog agent, configure the Datadog agent to report Ceph metrics to Datadog.

Prerequisites

- Root-level access to the Ceph monitor node.
- Appropriate Ceph key providing access to the Red Hat Ceph Storage cluster.
- Internet access.

Procedure

1. Install the Ceph integration.
 - a. Log in to the [Datadog App](#). The user interface will present navigation on the left side of the screen.
 - b. Click **Integrations**.
 - c. Either enter **ceph** into the search field or scroll to find the Ceph integration. The user interface will present whether the Ceph integration is *available* or already *installed*.
 - d. If it is *available*, click the button to install it.



2. Configuring the Datadog agent for Ceph
 - a. Navigate to the Datadog Agent configuration directory:

```
[root@mon ~]# cd /etc/dd-agent/conf.d
```

- b. Create a **ceph.yaml** file from the **ceph.yml.sample** file:

```
[root@mon ~]# cp ceph.yaml.example ceph.yaml
```

- c. Modify the **ceph.yaml** file:

```
[root@mon ~]# vim ceph.yaml
```


Example

The following is a sample of what the modified **ceph.yaml** file looks like.

```
init_config:

instances:
# - tags:
#   - name:mars_cluster
#
#   ceph_cmd: /usr/bin/ceph
#   ceph_cluster: ceph
#
# If your environment requires sudo, please add a line like:
#   dd-agent ALL=(ALL) NOPASSWD:/usr/bin/ceph
# to your sudoers file, and uncomment the below option.
#
#   use_sudo: True
```

Uncomment the **-tags**, **-name**, **ceph_command**, **ceph_cluster**, and **use_sudo: True** lines. The default values for **ceph_command** and **ceph_cluster** are **/usr/bin/ceph** and **ceph** respectively.

When complete, it will look like this:

```
init_config:

instances:
  - tags:
    - name:ceph-RHEL
#
  ceph_cmd: /usr/bin/ceph
  ceph_cluster: ceph
#
# If your environment requires sudo, please add a line like:
#   dd-agent ALL=(ALL) NOPASSWD:/usr/bin/ceph
# to your sudoers file, and uncomment the below option.
#
  use_sudo: True
```

d. Modify the sudoers file:

```
[root@mon ~]# visudo
```

e. Add the following line:

```
dd-agent ALL=(ALL) NOPASSWD:/usr/bin/ceph
```

f. Enable the Datadog agent so that it will restart if the Ceph host reboots:

```
[root@mon ~]# systemctl enable datadog-agent
```

g. Restart the Datadog agent:

```
| [root@mon ~]# systemctl status datadog-agent
```

CHAPTER 3. INSTALLING AND CONFIGURING THE DATADOG AGENT FOR CEPH

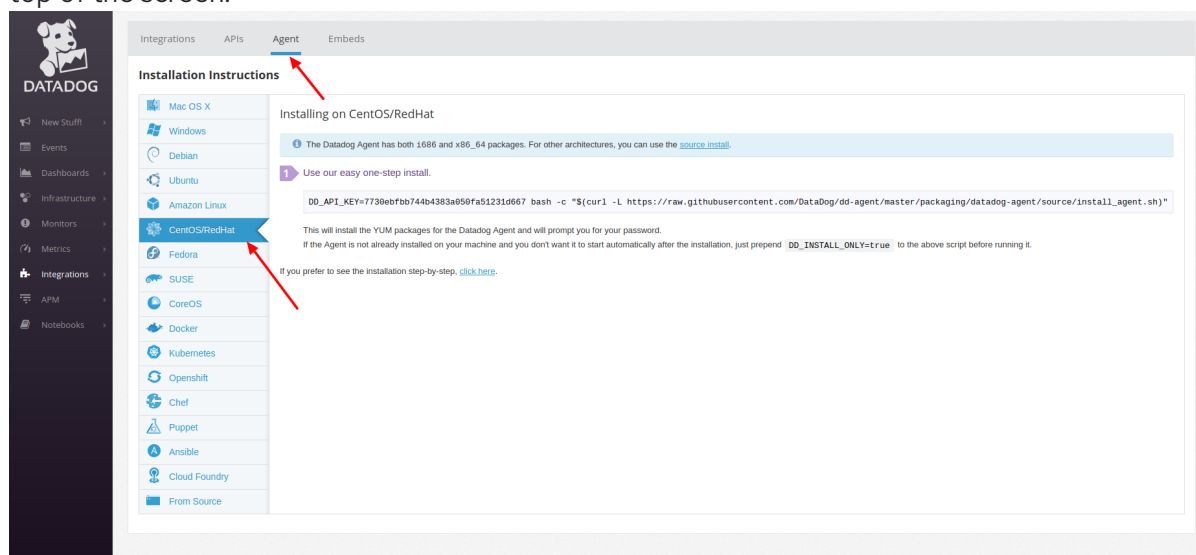
Install the Datadog agent for Ceph and configure it to report back the Ceph data to the Datadog App.

Prerequisites

- Root-level access to the Ceph monitor node.
- Appropriate Ceph key providing access to the Red Hat Ceph Storage cluster.
- Internet access.

Procedure

1. Log in to the [Datadog App](#). The user interface will present navigation on the left side of the screen.
2. Click **Integrations**. To install the agent from the command line, click on the **Agent** tab at the top of the screen.



3. Open a command line and enter the one-step command line agent installation.

Example

```
[root@mon ~]# DD_API_KEY=KEY-STRING bash -c "$(curl -L
https://raw.githubusercontent.com/DataDog/dd-agent/master/packaging/datadog-
agent/source/install_agent.sh)"
```



NOTE

Copy the example from the Datadog user interface, as the key differs from the example above and with each user account.

CHAPTER 5. CEPH METRICS FOR DATADOG

The Datadog agent collects the following metrics from Ceph. These metrics may be included in custom dashboards and in alerts.

Metric Name	Description
ceph.commit_latency_ms	The time taken to commit an operation to the journal.
ceph.apply_latency_ms	Time taken to flush an update to disks.
ceph.op_per_sec	The number of I/O operations per second for given pool.
ceph.read_bytes_sec	The bytes per second being read.
ceph.write_bytes_sec	The bytes per second being written.
ceph.num_osds	The number of known storage daemons.
ceph.num_in_osds	The number of participating storage daemons.
ceph.num_up_osds	The number of online storage daemons.
ceph.num_pgs	The number of placement groups available.
ceph.num_mons	The number of monitor daemons.
ceph.aggregate_pct_used	The overall capacity usage metric.
ceph.total_objects	The object count from the underlying object store.
ceph.num_objects	The object count for a given pool.
ceph.read_bytes	The per-pool read bytes.
ceph.write_bytes	The per-pool write bytes.
ceph.num_pools	The number of pools.
ceph.pgstate.active_clean	The number of active+clean placement groups.
ceph.read_op_per_sec	The per-pool read operations per second.
ceph.write_op_per_sec	The per-pool write operations per second.
ceph.num_near_full_osds	The number of nearly full OSDs.

Metric Name	Description
ceph.num_full_osds	The number of full OSDs.
ceph.osd.pct_used	The percentage used of full or near-full OSDs.

CHAPTER 6. CREATING ALERTS IN DATADOG

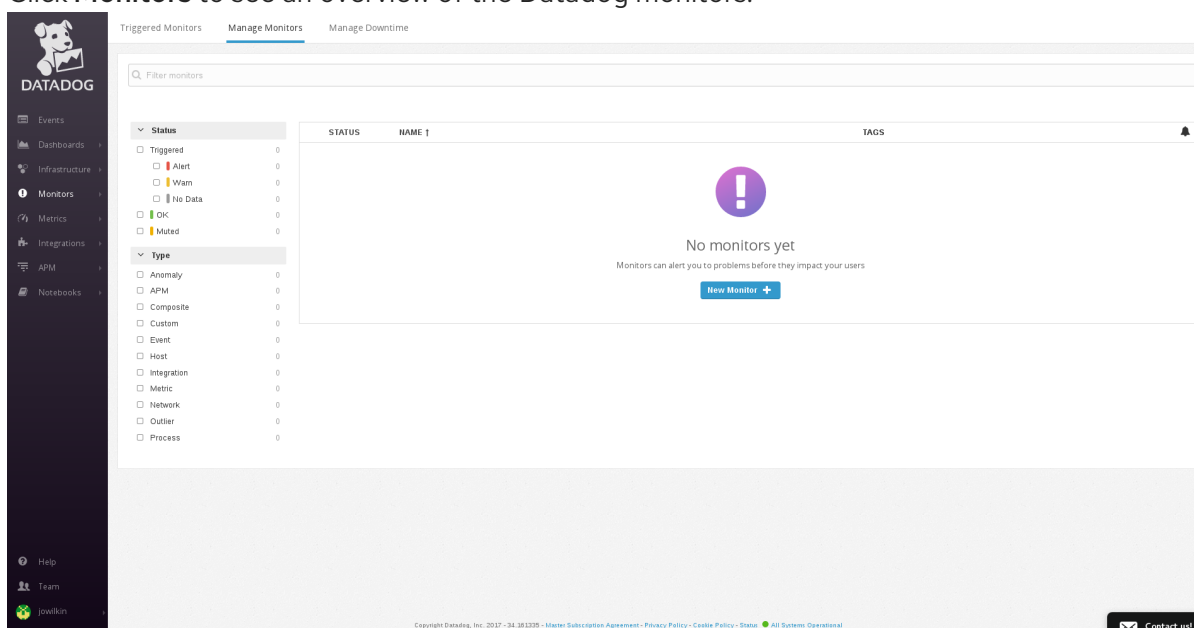
Administrators can create monitors that track the metrics of the Red Hat Ceph Storage cluster and generate alerts. For example, if an OSD is down, Datadog can alert an administrator that one or more OSDs are down.

Prerequisites

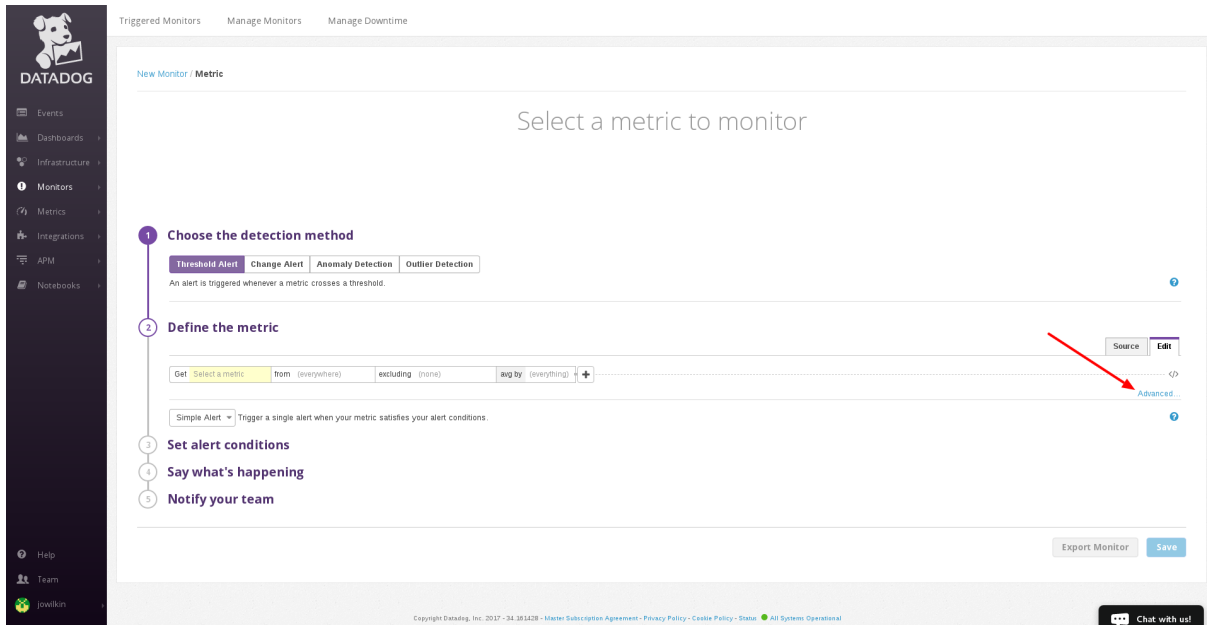
- Root-level access to the Ceph Monitor node.
- Appropriate Ceph key providing access to the Red Hat Ceph Storage cluster.
- Internet access.

Procedure

1. Click **Monitors** to see an overview of the Datadog monitors.

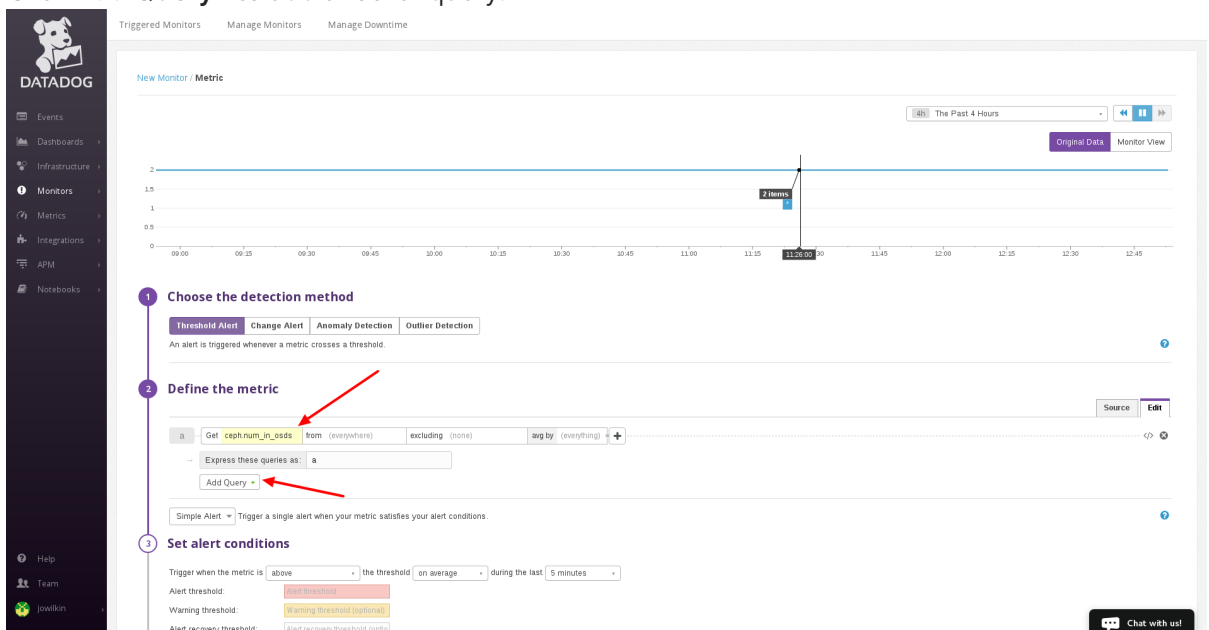


2. To create a monitor, select **Monitors**→**New Monitor**.
3. Select the detection method. For example, "Threshold Alert."



4. Define the metric. To create an advanced alert, click on the *Advanced...* link. Then, select a metric from the combo box. For example, select the **ceph.num_in_osds** Ceph metric.

5. Click **Add Query+** to add another query.



6. Select another metric from the combo box. For example, select the **ceph.num_up_osds** Ceph metric.

The screenshot shows the 'Define the metric' step in the Datadog alert configuration. It includes two queries: 'a' (Get ceph.num_in_osds) and 'b' (Get ceph.num_up_osds). The 'Express these queries as' field is set to 'a-b'. The 'Set alert conditions' section shows the trigger set to 'above or equal to' the threshold 'in total' during the last '1 minute', with an alert threshold of '1'.

7. In the **Express these queries as:** field, enter **a-b**, where **a** is the value of **ceph.num_in_osds** and **b** is the value of **ceph.num_up_osds**. When the difference is **1** or greater, there is at least one OSD down.
8. Set the alert conditions. For example, set the trigger to be **above or equal to**, the threshold to **in total** and the time elapsed to **1 minute**.
9. Set the **Alert threshold** field to **1**. When at least one OSD is in the cluster and it is not up and running, the monitor will alert the user.
10. Give the monitor a title in the input field below **Preview** and **Edit**. This is required to save the monitor.
11. Enter a description of the alert in the text field.

The screenshot shows the 'Say what's happening' step in the Datadog alert configuration. It includes a preview and edit field with the text: 'At least one OSD is down! @jowilkin@redhat.com At least one OSD is down. Restart it, or Ceph will begin backfilling shortly.' The 'Notify your team' section shows 'John Wilkins' as the contact.



NOTE

The text field supports metric variables and Markdown syntax.

12. Add the recipients of the alert. This will add an email address to the text field. When the alert gets triggered, the recipients will receive the alert.