Red Hat Ceph Storage 3

Monitoring Ceph with Datadog

Guide on Monitoring Ceph with Datadog

Red Hat Ceph Storage 3 Monitoring Ceph with Datadog

Guide on Monitoring Ceph with Datadog

Red Hat Ceph Storage 3 Monitoring Ceph with Datadog

Guide on Monitoring Ceph with Datadog

Legal Notice

Abstract

This document provides information on monitoring the status of the Ceph Storage cluster with the Datadog monitoring tool.

Table of Contents

CHAPTER 1. INTRODUCTION

The Datadog integration with Ceph allows 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 Ceph Storage cluster

- Monitor I/O and performance metrics; and,

- Track disk usage across storage pools.

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 Ceph 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 Ceph cluster must be able to reach the internet; however, the Ceph cluster does not have to be reachable from the internet.

NOTE

Datadog supports retrieving **ceph status** with RHCS 2. Datadog will provide an update to support **ceph status** for RHCS 3 in a subsequent release of its **dd-agent**.
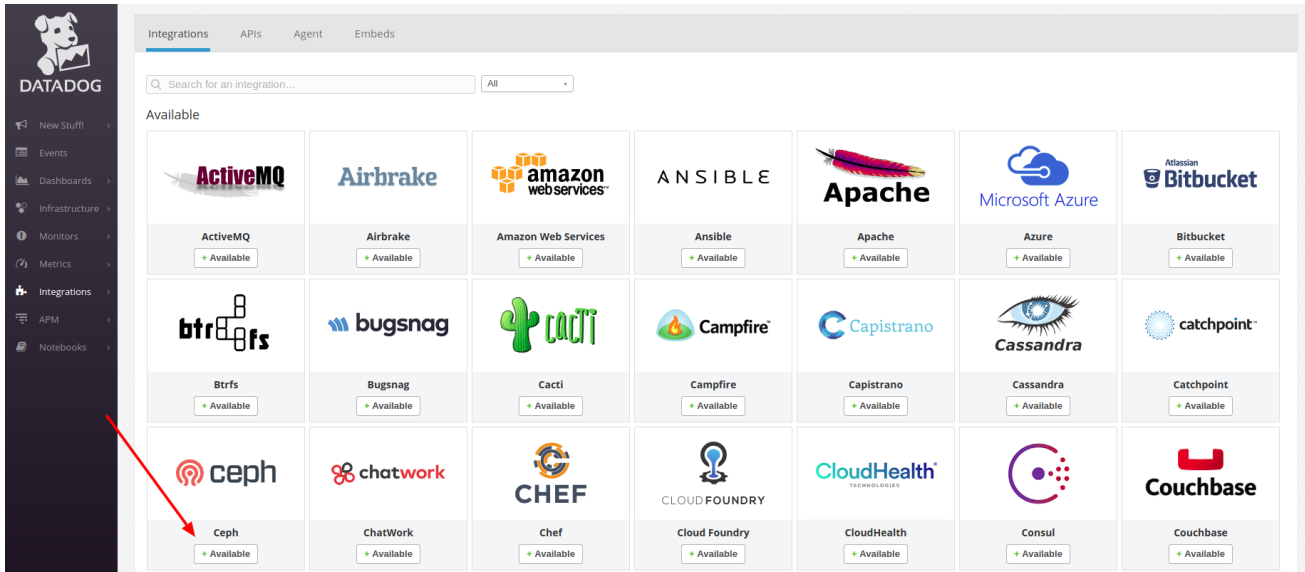
IMPORTANT

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.
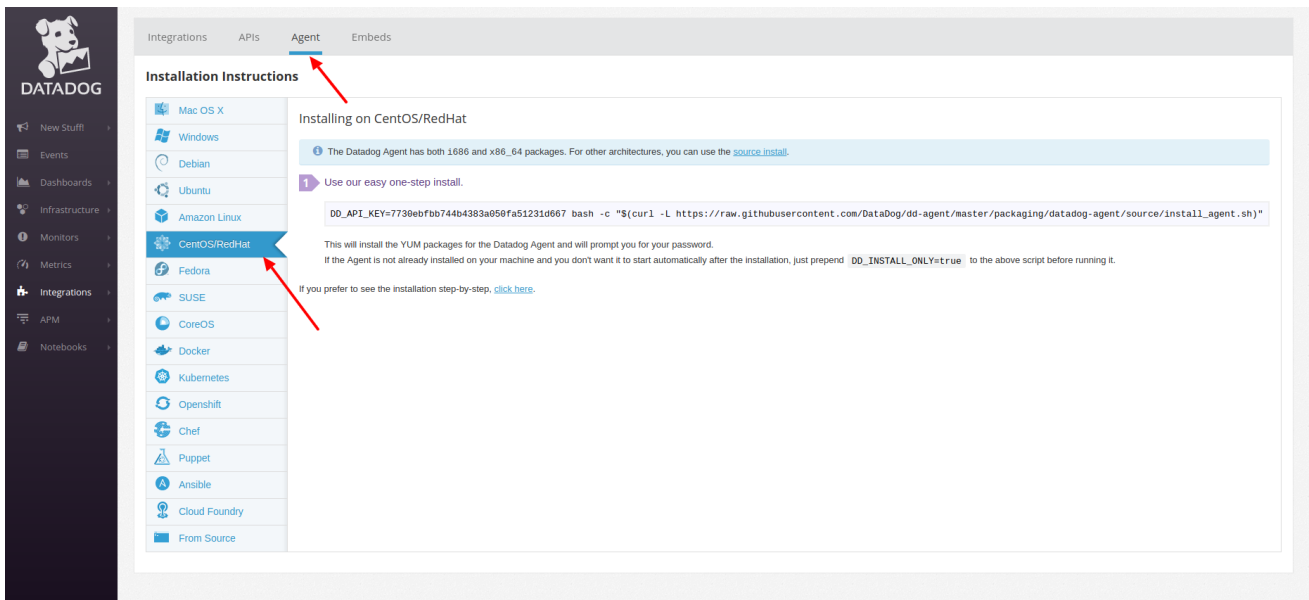
CHAPTER 2. INSTALLING THE CEPH INTEGRATION

To install the Ceph integration, log in to the Datadog App. The user interface will present navigation on the left side of the screen. Click **Integrations**. 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*. If it is *available*, click the button to install it.

Installing the Datadog Agent for Ceph

To install the Datadog agent for Red Hat Ceph Storage, log in to the Datadog App. The user interface will present navigation on the left side of the screen. Click **Integrations**. To install the agent from the command line, click on the **Agent** tab at the top of the screen.

For Red Hat Ceph Storage on RHEL 7 or Ubuntu 16.04, open a command line. Then, enter the one-step command line agent installation. For example:

```
# 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 3. CONFIGURING THE DATADOG AGENT FOR CEPH

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

1. Navigate to the Datadog Agent configuration directory.

   ```
   # cd /etc/dd-agent/conf.d
   ```

2. Create a **ceph.yaml** file from the **ceph.yml.sample** file.

   ```
   # cp ceph.yaml.example ceph.yaml
   ```

3. Modify the **ceph.yaml** file.

   ```
   # vim ceph.yaml
   ```

   It will look like this:

   ```
   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 and ceph_cluster lines. The default values for ceph_command and ceph_cluster are /usr/bin/ceph and ceph respectively. For RHEL 7, uncomment use_sudo: True; however, this step is optional for Ubuntu, since Ubuntu disables the root user and gives the initial admin user root permissions.

   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
```

4. For RHEL 7, modify the sudoers file.

```
# visudo
```

Add the following line.

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

> **NOTE**
>
> For Ubuntu, if **ceph.yml** enables **use_sudo: True**, perform this step, too.

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

```
# systemctl enable datadog-agent
```

6. Finally, restart the Datadog agent.

```
# systemctl status datadog-agent
```

CHAPTER 4. MONITORING CEPH WITH DATADOG

After installing and configuring the Datadog integration with Ceph, return to the Datadog App . The user interface will present navigation on the left side of the screen. Hover over **Dashboards** to expose the submenu; then, click **Ceph Overview**.

Datadog presents an overview of the Ceph Storage Cluster. Click **Dashboards→New Dashboard** to create a custom Ceph dashboard.

CHAPTER 5. CEPH METRICS

The Datadog agent collects the following metrics from Ceph. These metrics may be included in custom dashboards and in alerts.
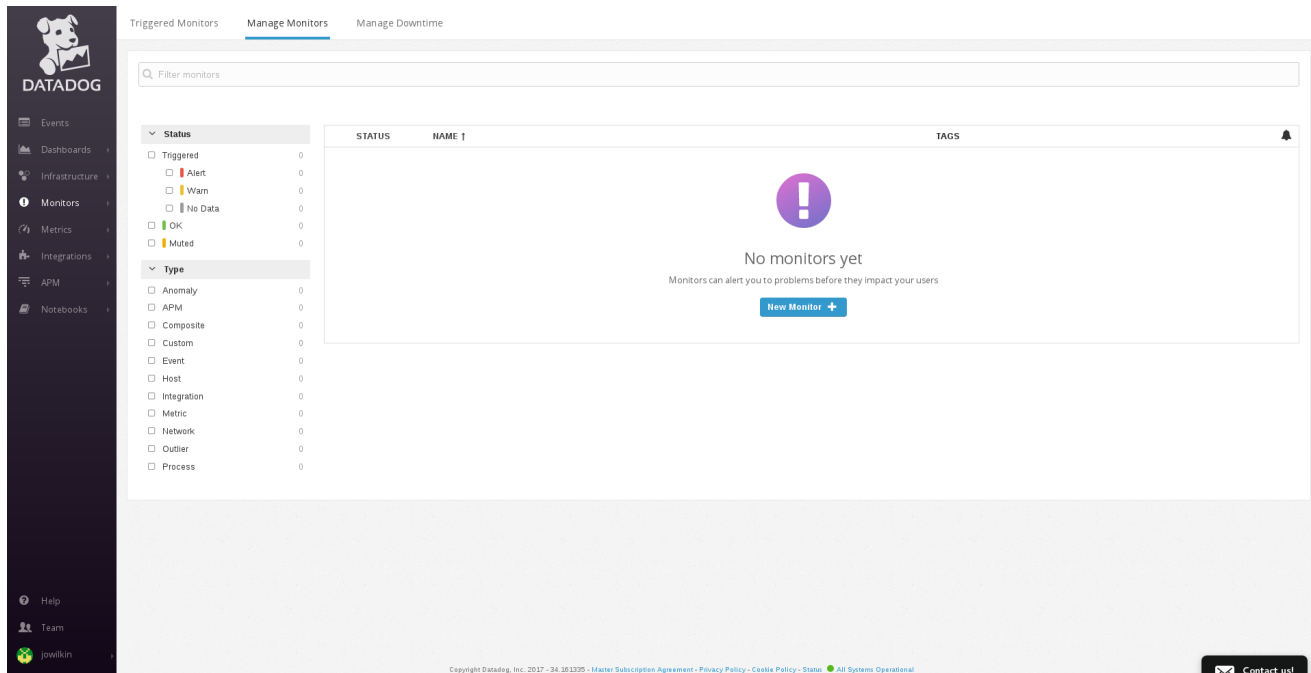
Metric Name	Description
ceph.commit_latency_ms	The time taken to commit an operation to the journal.
ceph.apply_latency_ms	Time taken to flush an update to disks.
ceph.op_per_sec	The number of I/O operations per second for given pool.
ceph.read_bytes_sec	The bytes per second being read.
ceph.write_bytes_sec	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.

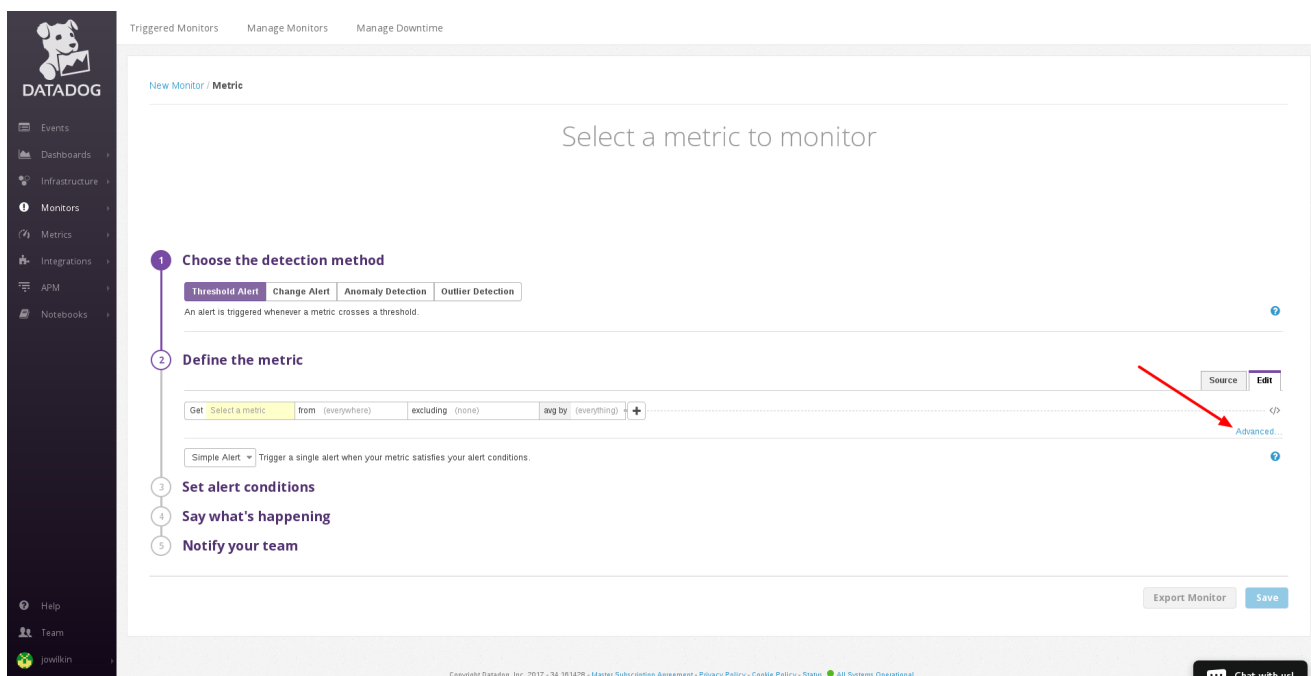| Metric Name	Description
ceph.num_full_osds	The number of full OSDs.
ceph.osd.pct_used	The percentage used of full or near-full OSDs.

CHAPTER 6. CREATE AN ALERT

Administrators can create monitors that track the metrics of the Ceph cluster and generate alerts. For example, if an OSD is down, Datadog can alert an administrator that one or more OSDs are down.
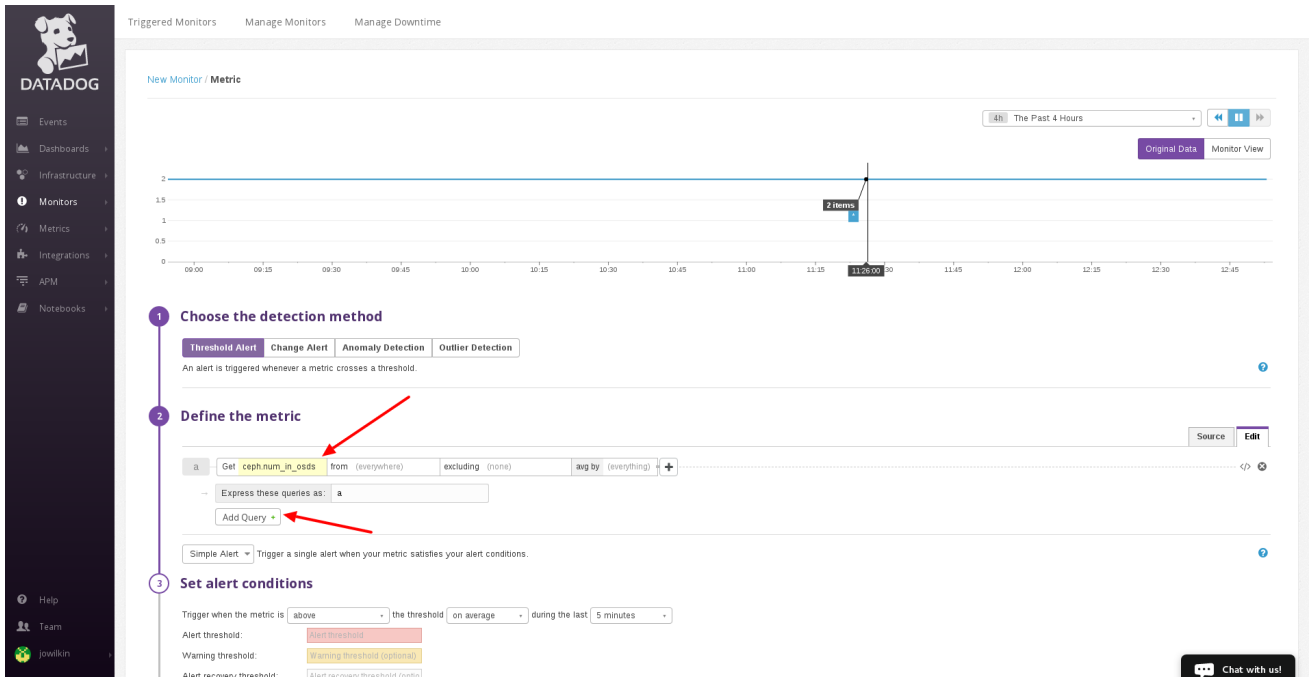
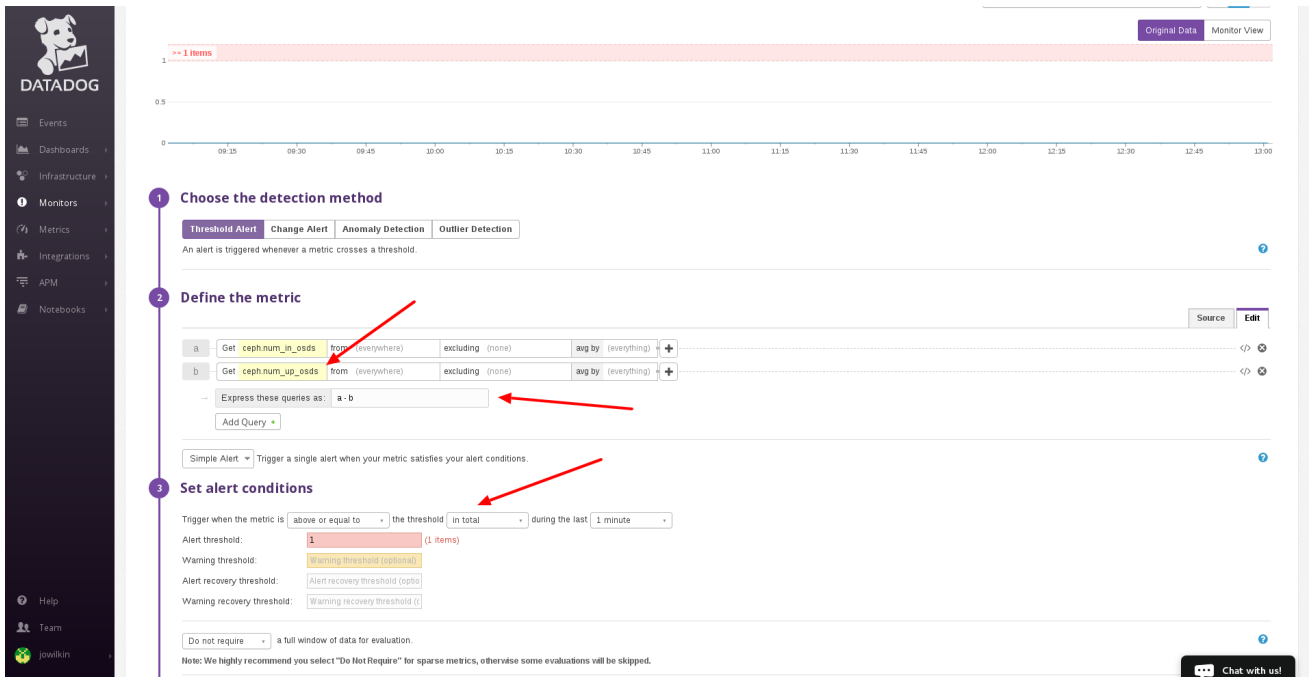Click **Monitors** to see an overview of the Datadog monitors.

To create a monitor, select **Monitors→New Monitor**. At step 1, select the detection method. For example, "Threshold Alert."

At step 2, 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. Then, click **Add Query+** to add another query.

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

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.

At step 3, 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**. Then, 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.

At step 4, give the monitor a title in the input field below **Preview** and **Edit**. This is required to save the monitor. Enter a description of the alert in the text field.

NOTE

The text field supports metric variables and Markdown syntax.

At step 5, add the recipients of the alert. This will add an email address to the text field of step 4. When the alert gets triggered, the recipients will receive the alert.