



Red Hat OpenStack Platform 16.1

Monitoring Tools Configuration Guide

A guide to OpenStack logging and monitoring tools

Red Hat OpenStack Platform 16.1 Monitoring Tools Configuration Guide

A guide to OpenStack logging and monitoring tools

OpenStack Team
rhos-docs@redhat.com

Legal Notice

Copyright © 2020 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 guide provides information on configuring logging and monitoring for a Red Hat OpenStack Platform environment.

Table of Contents

CHAPTER 1. INTRODUCTION	3
CHAPTER 2. MONITORING ARCHITECTURE	4
2.1. CENTRALIZED LOGGING	4
2.2. AVAILABILITY MONITORING	4
CHAPTER 3. INSTALLING THE CLIENT-SIDE TOOLS	8
3.1. SETTING CENTRALIZED LOGGING CLIENT PARAMETERS	8
3.2. SETTING MONITORING CLIENT PARAMETERS	8
3.3. COLLECTD PLUG-IN CONFIGURATIONS	9
3.3.1. Configuring the pcie_errors plug-in	9
3.3.2. Configuring the mcelog plug-in	9
3.4. YAML FILES	10
CHAPTER 4. MONITOR THE OPENSTACK PLATFORM	11
CHAPTER 5. VALIDATE THE SENSU CLIENT INSTALLATION	12
CHAPTER 6. REVIEW THE STATE OF A NODE	13
CHAPTER 7. REVIEW THE STATE OF AN OPENSTACK SERVICE	14

CHAPTER 1. INTRODUCTION

Monitoring tools are an optional suite of tools designed to help operators maintain an OpenStack environment. The tools perform the following functions:

- **Centralized logging:** Allows you gather logs from all components in the OpenStack environment in one central location. You can identify problems across all nodes and services, and optionally, export the log data to Red Hat for assistance in diagnosing problems.
- **Availability monitoring:** Allows you to monitor all components in the OpenStack environment and determine if any components are currently experiencing outages or are otherwise not functional. You can also configure the system to alert you when problems are identified.

CHAPTER 2. MONITORING ARCHITECTURE

Monitoring tools use a client-server model with the client deployed onto the Red Hat OpenStack Platform overcloud nodes. The Rsyslog service provides client-side centralized logging (CL) and the collectd with enabled sensubility plugin provides client-side availability monitoring (AM).

2.1. CENTRALIZED LOGGING

In your Red Hat OpenStack environment, collecting the logs from all services in one central location simplifies debugging and administration. These logs come from the operating system, such as syslog and audit log files, infrastructure components such as RabbitMQ and MariaDB, and OpenStack services such as Identity, Compute, and others.

The centralized logging toolchain consists of the following components:

- Log Collection Agent (Rsyslog)
- Data Store (Elasticsearch)
- API/Presentation Layer (Kibana)



NOTE

Red Hat OpenStack Platform director does not deploy the server-side components for centralized logging. Red Hat does not support the server-side components, including the Elasticsearch database and Kibana.

2.2. AVAILABILITY MONITORING

With availability monitoring, you have one central place to monitor the high-level functionality of all components across your entire OpenStack environment.

The availability monitoring toolchain consists of several components:

- Monitoring Agent (collectd with enabled sensubility plugin)
- Monitoring Relay/Proxy (RabbitMQ)
- Monitoring Controller/Server (Sensu server)
- API/Presentation Layer (Uchiwa)



NOTE

Red Hat OpenStack Platform director does not deploy the server-side components for availability monitoring. Red Hat does not support the server-side components, including Uchiwa, Sensu Server, the Sensu API plus RabbitMQ, and a Redis instance running on a monitoring node.


The availability monitoring components and their interactions are laid out in the following diagrams:



NOTE

Items shown in blue denote Red Hat-supported components.

Figure 2.1. Availability monitoring architecture at a high level



OPENSTACK_435795_0117

Figure 2.2. Single-node deployment for Red Hat OpenStack Platform





Figure 2.3. HA deployment for Red Hat OpenStack Platform



OPENSTACK_435795_0117

CHAPTER 3. INSTALLING THE CLIENT-SIDE TOOLS

Before you deploy the overcloud, you need to determine the configuration settings to apply to each client. Copy the example environment files from the heat template collection and modify the files to suit your environment.

3.1. SETTING CENTRALIZED LOGGING CLIENT PARAMETERS

For more information, see https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/16.1/html/logging_monitoring_and_troubleshooting_guide/logging#enabl

3.2. SETTING MONITORING CLIENT PARAMETERS

The monitoring solution collects system information periodically and provides a mechanism to store and monitor the values in a variety of ways using a data collecting agent. Red Hat supports collectd as a collection agent. You can use Service Telemetry Framework (STF) to store the data, and in turn, monitor systems, find performance bottlenecks, and predict future system load.

To configure collectd, complete the following steps:

1. Create **config.yaml** in your home directory, for example, **/home/templates/custom**, and configure the **MetricsQdrConnectors** parameter to point to the Service Telemetry Framework server side:

```
MetricsQdrConnectors:
  - host: qdr-normal-sa-telemetry.apps.remote.tld
    port: 443
    role: inter-router
    sslProfile: sslProfile
    verifyHostname: false
MetricsQdrSSLProfiles:
  - name: sslProfile
```

2. In the **config.yaml** file, list the plug-ins you want under **CollectdExtraPlugins**. You can also provide parameters in the **ExtraConfig** section. By default, collectd comes with the **cpu**, **df**, **disk**, **hugepages**, **interface**, **load**, **memory**, **processes**, **tcpconns**, **unixsock**, and **uptime** plug-ins. You can add additional plug-ins using the **CollectdExtraPlugins** parameter. You can also provide additional configuration information for the **CollectdExtraPlugins** using the **ExtraConfig** option as shown. For example, to enable the **virt** plug-in, and configure the connection string and the hostname format, use the following syntax:

```
parameter_defaults:
  CollectdExtraPlugins:
    - disk
    - df
    - virt

  ExtraConfig:
    collectd::plugin::virt::connection: "qemu:///system"
    collectd::plugin::virt::hostname_format: "hostname uuid"
```

**NOTE**

Do not remove the **unixsock** plug-in. Removal results in the permanent marking of the collectd container as unhealthy.

3. Deploy the overcloud. Include **config.yaml**, **collectd-write-qdr.yaml**, and one of the **qdr-*.yaml** files in your **overcloud deploy** command. For example:

```
$ openstack overcloud deploy
-e /home/templates/custom/config.yaml
-e tripleo-heat-templates/environments/metrics/collectd-write-qdr.yaml
-e tripleo-heat-templates/environments/metrics/qdr-form-controller-mesh.yaml
```

4. Optional: To enable overcloud RabbitMQ monitoring, include the **collectd-read-rabbitmq.yaml** file in your **overcloud deploy** command. For more information about the YAML files, see [Section 3.4, “YAML files”](#).

Additional resources

- For more information, see the [Service Telemetry Framework](#) guide.
- For more information about collectd plug-ins, see [Section 3.3, “Collectd plug-in configurations”](#).

3.3. COLLECTD PLUG-IN CONFIGURATIONS

There are multiple configuration possibilities of Red Hat OpenStack Platform director. You can configure multiple collectd plug-ins to suit your environment.

Additional resources

To view the complete list of collectd plug-ins, see [collectd plug-ins](#) in the *Service Telemetry Framework* guide.

3.3.1. Configuring the pcie_errors plug-in

Use the `pcie_errors` plug-in to poll PCI config space for baseline and Advanced Error Reporting (AER) errors, and to parse syslog for AER events. Errors are reported through notifications.

Procedure

- Add the **pcie_errors** plug-in configuration to **config.yaml**:

```
parameter_defaults:
  CollectdExtraPlugins:
    - pcie_errors
```

Additional resources

- For more information about configuring collectd plug-ins, see [Setting monitoring client parameters](#).

3.3.2. Configuring the mcelog plug-in

Use the **mcelog** plug-in to send notifications and statistics relevant to Machine Check Exceptions when they occur. Configure **mcelog** to run on the platform in daemon mode and ensure that logging capabilities are enabled.

Procedure

- Add the **mcelog** plug-in configuration to **config.yaml**:

```
parameter_defaults:
  CollectdExtraPlugins: mcelog
  CollectdEnableMcelog: true
```

Additional resources

- For more information about configuring collectd plug-ins, see [Setting monitoring client parameters](#).

3.4. YAML FILES

When you configure collectd, you can include the following YAML files in your overcloud deploy command:

- **collectd-read-rabbitmq.yaml**: Enables python-collect-rabbitmq and configures it to monitor overcloud RabbitMQ instance.
- **collectd-write-qdr.yaml**: Enables collectd to send telemetry and notification data through QPID dispatch routers (QDR).
- **qdr-edge-only.yaml**: Enables deployment of QPID dispatch routers. Each overcloud node will have one local qdrouterd service running and operating in edge mode, for example, sending received data straight to defined MetricsQdrConnectors.
- **qdr-form-controller-mesh.yaml**: Enables deployment of QPID dispatch routers. Each overcloud node will have one local qdrouterd service running and forming a mesh topology. For example, QDRs running on controllers operate in interior router mode, with connections to defined MetricsQdrConnectors, and QDRs running on other node types connect in edge mode to the interior routers running on the controllers.

CHAPTER 4. MONITOR THE OPENSTACK PLATFORM

See the Sensu documentation for further details about the Ssensu stack infrastructure:

<https://docs.sensu.io/sensu-core/1.7/overview/architecture/>

Red Hat supplies a set of check scripts in the **osops-tools-monitoring-oschecks** package. The majority of the check scripts only check the API connection to the OpenStack component. However, certain scripts also perform additional OpenStack resource tests for OpenStack Compute (nova), OpenStack Block Storage (cinder), OpenStack Image (glance), and OpenStack Networking (neutron). For example, the OpenStack Identity (keystone) API check returns the following result when **keystone** is running:

OK: Got a token, Keystone API is working.

CHAPTER 5. VALIDATE THE SENSU CLIENT INSTALLATION

1. Check the status of the **sensu-client** on each overcloud node:

```
# podman ps | grep sensu-client
```

2. Review the error log for any issues: **/var/log/containers/sensu/sensu-client.log**
3. Verify that each overcloud node has the **/var/lib/config-data/puppet-generated/sensu/etc/sensu/conf.d/rabbitmq.json** file that sets the IP address of the monitoring server.

CHAPTER 6. REVIEW THE STATE OF A NODE

If you have a deployment of the Uchiwa dashboard, you can use it with the Sensu server to review the state of your nodes:

1. Login to the Uchiwa dashboard and click the **Data Center** tab to confirm that the Data Center is operational.

 `http://<SERVER_IP_ADDRESS>/uchiwa`

2. Check that all overcloud nodes are in a **Connected** state.
3. At a suitable time, reboot one of the overcloud nodes and review the rebooted node's status in the Uchiwa dashboard. After the reboot completes, verify that the node successfully re-connects to the Sensu server and starts executing checks.

CHAPTER 7. REVIEW THE STATE OF AN OPENSTACK SERVICE

This example tests the monitoring of the **openstack-ceilometer-central** service.

1. Confirm that the **openstack-ceilometer-central** service is running:

```
docker ps -a | grep ceilometer
```

2. Connect to the Uchiwa dashboard and confirm that a successful **ceilometer** check is present and running as defined in the **ceilometer** JSON file.