



# **Red Hat Virtualization 4.2**

## **Metrics Store Installation Guide**

Installing Metrics Store for Red Hat Virtualization



# Red Hat Virtualization 4.2 Metrics Store Installation Guide

---

Installing Metrics Store for Red Hat Virtualization

Red Hat Virtualization Documentation Team

Red Hat Customer Content Services

[rhev-docs@redhat.com](mailto:rhev-docs@redhat.com)

## Legal Notice

Copyright © 2018 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, 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

A comprehensive guide to installing and configuring Metrics Store for Red Hat Virtualization.

---

## Table of Contents

<b>CHAPTER 1. INTRODUCTION .....</b>	<b>3</b>
1.1. SYSTEM REQUIREMENTS .....	3
1.2. PREREQUISITES .....	3
<b>CHAPTER 2. SETTING UP THE RED HAT VIRTUALIZATION MANAGER AND HOSTS .....</b>	<b>4</b>
2.1. COPYING OPENSIFT ANSIBLE FILES .....	4
<b>CHAPTER 3. SETTING UP OPENSIFT AGGREGATED LOGGING .....</b>	<b>6</b>
3.1. CONFIGURING ANSIBLE PREREQUISITES .....	6
3.2. OPENING PORTS .....	6
3.3. CONFIGURING SUDO .....	7
3.4. ATTACHING SUBSCRIPTIONS AND ENABLING REPOSITORIES .....	7
3.5. INSTALLING OPENSIFT AGGREGATED LOGGING PACKAGES .....	8
3.6. CONFIGURING PERSISTENT STORAGE FOR ELASTICSEARCH .....	8
3.7. RUNNING ANSIBLE .....	8
3.8. ENABLING ELASTICSEARCH TO MOUNT THE DIRECTORY .....	9
3.9. VERIFYING THE OPENSIFT AGGREGATED LOGGING INSTALLATION .....	10
3.10. CONFIGURING COLLECTD AND FLUENTD .....	10
<b>CHAPTER 4. VERIFYING THE INSTALLATION .....</b>	<b>12</b>



# CHAPTER 1. INTRODUCTION

OpenShift Aggregated Logging is based on the OpenShift Logging stack running on OpenShift Container Platform (OCP). Ansible is used to install OpenShift Aggregated Logging using OpenShift Ansible logging roles.

## 1.1. SYSTEM REQUIREMENTS

- 4 cores, 16GB RAM, and 500GB disk for an environment with 50 hosts.
- Red Hat highly recommends using SSD disks.
- OpenShift Aggregated Logging requires RHEL 7.5.

## 1.2. PREREQUISITES

### Metrics Store Machine Prerequisites


1. Add the hostname of the OpenShift Aggregated Logging machine to your enterprise hostname resolution system, for example, DNS.
2. Add the following aliases:
  - `es.FQDN` for Elasticsearch
  - `kibana.FQDN` for Kibana  
where `FQDN` is the hostname and domain of the OpenShift Aggregated Logging machine.
3. The machine must meet all [Minimum Hardware Requirements](#) detailed in the **Masters** section.
4. Ensure that `libvirt` is not installed on the machine:

```
# rpm -qa | grep libvirt
```

If `libvirt` is installed, remove it from the machine:

```
# yum remove libvirt*
```

5. Create a preallocated 500GB partition called `/var`, which will be used for persistent storage. Do not use `root (/)`.



### IMPORTANT

XFS is the only supported file system for persistent storage.

### Manager Machine Prerequisites

Ensure that the time stamp in the `/var/log/ovirt-engine/engine.log` file contains a UTC offset suffix, rather than a letter such as Z or A. For example: 2018-03-27 13:35:06,720+01

## CHAPTER 2. SETTING UP THE RED HAT VIRTUALIZATION MANAGER AND HOSTS

### Prerequisites

Install a 4.2 environment as described in the *Installation Guide* or *Self-Hosted Installation Guide*, depending on your environment. Alternatively, upgrade your 4.x environment to 4.2.

### 2.1. COPYING OPENSIFT ANSIBLE FILES

1. On the Manager machine, copy `/etc/ovirt-engine-metrics/config.yml.example` to `config.yml`:

```
# cp /etc/ovirt-engine-metrics/config.yml.example /etc/ovirt-engine-metrics/config.yml
```

2. Update the values of `/etc/ovirt-engine-metrics/config.yml` to match the details of your specific environment:

```
# vi /etc/ovirt-engine-metrics/config.yml
```



#### IMPORTANT

All parameters are mandatory.

Table 2.1. `config.yml` Parameters

Name	Default Value	Description
<code>ovirt_env_name</code>	Yes	<p>The environment name. This is used to identify data collected from the Manager for this Red Hat Virtualization environment.</p> <p>Use the following conventions:</p> <ul style="list-style-type: none"><li>• Include only alphanumeric characters and hyphens ( "-" ).</li><li>• The name cannot begin with a hyphen or a number, or end with a hyphen.</li><li>• A maximum of 49 characters can be used.</li><li>• Wildcard patterns (for example, <code>ovirt-metrics</code>) cannot be used.</li></ul>
<code>fluentd_elasticsearch_host</code>	No	The address or FQDN of the Elasticsearch server host.

3. Copy the Manager's public key to your Metrics Store machine:



```
# mytemp=$(mktemp -d)

# cp /etc/pki/ovirt-engine/keys/engine_id_rsa $mytemp

# ssh-keygen -y -f $mytemp/engine_id_rsa > $mytemp/engine_id_rsa.pub

# ssh-copy-id -i $mytemp/engine_id_rsa.pub
root@fluentd_elasticsearch_host
```

It should ask for root password (on first attempt), supply it. After that, run:

```
# rm -rf $mytemp
```

To test that you are able to log into the metrics store machine from the engine, run:

```
# ssh -i /etc/pki/ovirt-engine/keys/engine_id_rsa
root@fluentd_elasticsearch_host
```

4. As the root user, run the Ansible script that generates the Ansible inventory and vars.yaml files and copies them to the Metrics Store machine (by default to / (root)):

```
# /usr/share/ovirt-engine-
metrics/setup/ansible/configure_ovirt_machines_for_metrics.sh \
--playbook=ovirt-metrics-store-installation.yml
```

## CHAPTER 3. SETTING UP OPENSIFT AGGREGATED LOGGING

### 3.1. CONFIGURING ANSIBLE PREREQUISITES

You must be able to log into the machine using an SSH keypair. The following instructions assume you are running Ansible on the same machine that you will be running OpenShift Aggregated Logging.

#### Configure Ansible Prerequisites

1. Assign the machine an FQDN and IP address so that it can be reached from another machine. These are the `public_hostname` and `public_ip` parameters.
2. Use the root user or create a user account. This user will be referred to below as `$USER`. If you do not use the root user, you must update `ansible_ssh_user` and `ansible_become` in `vars.yaml`, which is saved to the `/root` directory on the Metrics Store machine by default.
3. Create an SSH public key for this user account using the `ssh-keygen` command.

```
# ssh-keygen
```

4. Add the SSH public key to the user account `$HOME/.ssh/authorized_keys`:

```
# cat $HOME/.ssh/id_rsa.pub >> $HOME/.ssh/authorized_keys
```

5. Add the SSH hostkey for localhost to your SSH `known_hosts`:

```
# ssh-keyscan -H localhost >> $HOME/.ssh/known_hosts
```

6. Add the SSH hostkey for `public_hostname` to your SSH `known_hosts`:

```
# ssh-keyscan -H public_hostname >> $HOME/.ssh/known_hosts
```

7. If you *not* using the root user, enable passwordless sudo by adding `$USER ALL=(ALL) NOPASSWD: ALL` to `/etc/sudoers`.

8. Verify that passwordless SSH works, and that you do not get prompted to accept host verification, by running:

```
# ssh localhost 'ls -al'
# ssh public_hostname 'ls -al'
```

Ensure that you are not prompted to provide a password or to accept host verification.



#### NOTE

`openshift-ansible` may attempt to SSH to localhost. This is the expected behavior.

### 3.2. OPENING PORTS

The TCP ports listed below are required by OpenShift Container Platform. Ensure that they are open on your network and configured to allow access between hosts.

Use iptables to open ports. The following example opens port **22**:

```
# iptables OS_FIREWALL_ALLOW -p tcp -m state --state NEW -m tcp \
--dport 22 -j ACCEPT
```

### Required Ports

- **22** Required for SSH by the installer or system administrator.
- **443** For use by Kibana.
- **8443** For use by the OpenShift Container Platform web console, shared with the API server. This enables Metrics users to access the OpenShift Management user interface.
- **9200** For Elasticsearch API use. Required to be internally open on any infrastructure nodes to enable Kibana to retrieve logs. It can be externally opened for direct access to Elasticsearch by means of a route. The route can be created using `oc expose`.

## 3.3. CONFIGURING SUDO

### Configure sudo not to require a tty

Create a file under `/etc/sudoers.d/`, for example `999-cloud-init-requiretty`, and add **Defaults !requiretty** to the file.

For example:

```
# cat /etc/sudoers.d/999-cloud-init-requiretty
Defaults !requiretty
```

## 3.4. ATTACHING SUBSCRIPTIONS AND ENABLING REPOSITORIES

OpenShift Aggregated Logging requires RHEL 7.5 and OpenShift 3.9 subscriptions.

1. Register your system with the Content Delivery Network, entering your Customer Portal user name and password when prompted:

```
# subscription-manager register
```

2. Pull the latest subscription data from Red Hat Subscription Manager:

```
# subscription-manager refresh
```

3. Find the the **OpenShift Container Platform** subscription pool and note down the pool ID:

```
# subscription-manager list --available
```

4. Use the pool IDs to attach the subscriptions to the system:

```
# subscription-manager attach --pool=pool_id
```

5. Enable the required repositories:

■

```
# subscription-manager repos --enable="rhel-7-server-rpms" \
--enable="rhel-7-server-extras-rpms" \
--enable="rhel-7-server-ose-3.9-rpms" \
--enable="rhel-7-fast-datapath-rpms" \
--enable="rhel-7-server-ansible-2.4-rpms"
```

### 3.5. INSTALLING OPENSIFT AGGREGATED LOGGING PACKAGES


The installer for OpenShift Container Platform is provided by the `atomic-openshift-utils` package.

Install the OpenShift Container Platform package:

```
# yum -y install wget git net-tools bind-utils iptables-services bridge-
utils bash-completion kexec-tools sos psacct
# yum -y update
# yum -y install atomic-openshift-utils
# yum -y install docker
```

### 3.6. CONFIGURING PERSISTENT STORAGE FOR ELASTICSEARCH

Elasticsearch requires persistent storage for the database. By default, Elasticsearch uses ephemeral storage, and therefore you need to manually configure persistent storage.



#### IMPORTANT

Before proceeding, ensure you have set up the storage according to the instructions in [Section 1.2, “Prerequisites”](#).

#### Configuring Persistent Storage for Elasticsearch

1. Create the `/lib/elasticsearch` directory that will be used for persistent storage using the `/var` mounted storage partition you created in [Section 1.2, “Prerequisites”](#):

```
# mkdir -p /var/lib/elasticsearch
```

2. Change the group ownership of the directory to 65534:

```
# chgrp 65534 /var/lib/elasticsearch
```

3. Make this directory writable by the group:

```
# chmod -R 0770 /var/lib/elasticsearch
```

4. Run the following commands:

```
# semanage fcontext -a -t container_file_t
"/var/lib/elasticsearch(/.*)?"
# restorecon -R -v /var/lib/elasticsearch
```

### 3.7. RUNNING ANSIBLE

Prior to running Ansible, verify that the value for hostname and IP address that you defined in the DNS matches the values Ansible will use.

### Running Ansible

1. To check the host's FQDN:

```
# ansible -m setup localhost -a 'filter=ansible_fqdn'
```

2. To check the host's IP address:

```
# ansible -m setup localhost -a 'filter=ansible_default_ipv4'
```

3. Run Ansible using the *prerequisites.yml* playbook to ensure the machine is configured correctly:

```
# cd /usr/share/ansible/openshift-ansible
# ANSIBLE_LOG_PATH=/tmp/ansible-prereq.log ansible-playbook -vvv -e
@/root/vars.yaml -i /root/ansible-inventory-ocp-39-aio
playbooks/prerequisites.yml
```

4. Run Ansible using the *openshift-node/network\_manager.yml* playbook to ensure that the networking and the NetworkManager are configured correctly:

```
# cd /usr/share/ansible/openshift-ansible
# ANSIBLE_LOG_PATH=/tmp/ansible-network.log ansible-playbook -vvv -e
@/root/vars.yaml -i /root/ansible-inventory-ocp-39-aio
playbooks/openshift-node/network_manager.yml
```

5. Run Ansible using the *deploy\_cluster.yml* playbook to install both OpenShift and the OpenShift Logging components:

```
# cd /usr/share/ansible/openshift-ansible
# ANSIBLE_LOG_PATH=/tmp/ansible.log ansible-playbook -vvv -e
@/root/vars.yaml -i /root/ansible-inventory-ocp-39-aio
playbooks/deploy_cluster.yml
```

6. Check */tmp/ansible.log* to ensure that no errors occurred. If there are errors, fix the machine's definitions and/or *vars.yaml* and run Ansible again.



#### NOTE

If the installation fails, inspect the Ansible log files in */var/log/ovirt-engine/ansible/*, fix the issue, and run the installation again.

## 3.8. ENABLING ELASTICSEARCH TO MOUNT THE DIRECTORY

After the installation, the Elasticsearch service will not be able to run until granted permission to mount that directory.

### Enabling Elasticsearch to Mount the Directory

Run the following:

```
# oc project logging
# oadm policy add-scc-to-user hostmount-anyuid \
  system:serviceaccount:logging:aggregated-logging-elasticsearch

# oc rollout cancel $( oc get -n logging dc -l component=es -o name )
# oc rollout latest $( oc get -n logging dc -l component=es -o name )
# oc rollout status -w $( oc get -n logging dc -l component=es -o name )
```

## 3.9. VERIFYING THE OPENSIFT AGGREGATED LOGGING INSTALLATION

The following procedures verify that all pods and services are running, and that the hostname, IPs, and routes are correctly configured.

### Verifying the OpenShift Aggregated Logging Installation

1. Log into the project:

```
# oc project logging
```

2. To confirm that Elasticsearch, Curator, and Kibana pods are running, run:

```
# oc get pods
```

3. Check that the **STATUS** is **Running**.

4. To confirm that the Elasticsearch and Kibana services are running, run:

```
# oc get svc
```

5. Ensure that the **EXTERNAL-IP** and **PORT(S)** fields are correct.

6. To confirm that there are routes for Elasticsearch and Kibana, run:

```
# oc get routes
```

7. Ensure that the value of **HOST/PORT** is correct.

## 3.10. CONFIGURING COLLECTD AND FLUENTD

Deploy and configure collectd and fluentd to send the metrics and logs to OpenShift Aggregated Logging.

### Configuring Collectd and Fluentd

On the Manager machine, run the following:

```
# /usr/share/ovirt-engine-
metrics/setup/ansible/configure_ovirt_machines_for_metrics.sh
```

**NOTE**

Deploying additional hosts after running this script does *not* require running the script again; the Manager configures the hosts automatically.

## CHAPTER 4. VERIFYING THE INSTALLATION

Access the Kibana console to view the logs and statistics about clusters, hosts, virtual machines, and the Manager.

### Verifying the Installation

1. Access Kibana at <https://kibana.<FQDN>>.
2. In the **Discover** tab, check that you can view the `project.ovirt-logs-ovirt_env_name-uuid*` index, where `ovirt_env_name` is the name you defined in [Configuring Collectd and Fluentd](#). See the [Discover](#) section in the *Kibana documentation* for more information about working with logs.
3. Use the **Visualization** tab to build visualization for the `project.ovirt-metrics-ovirt_env_name-uuid*` and the `project.ovirt-logs-ovirt_env_name-uuid*` indexes.

See the [Metrics User Guide](#) for the available parameters. See the [Visualize](#) section of the *Kibana documentation* for more information about visualizing logs and metrics.



### NOTE

You can access the OpenShift portal at <https://FQDN:8443>.