



Red Hat OpenStack Platform 11

Auto Scaling for Instances

Configure Auto Scaling in Red Hat OpenStack Platform

Red Hat OpenStack Platform 11 Auto Scaling for Instances

Configure Auto Scaling in Red Hat OpenStack Platform

OpenStack Team
rhos-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

Automatically scale out your Compute instances in response to system usage.

Table of Contents

CHAPTER 1. ABOUT THIS GUIDE	3
CHAPTER 2. CONFIGURE AUTO SCALING FOR COMPUTE INSTANCES	4
2.1. ARCHITECTURAL OVERVIEW	4
2.1.1. Orchestration	4
2.1.2. Telemetry	4
2.1.3. Key Terms	4
2.2. EXAMPLE: AUTO SCALING BASED ON CPU USAGE	4
2.2.1. Test Automatic Scaling Up Instances	9
2.2.2. Automatically Scaling Down Instances	10
2.2.3. Troubleshooting the setup	11
2.3. EXAMPLE: AUTO SCALING APPLICATIONS	13
2.3.1. Test Auto Scaling Applications	20
2.3.2. Automatically Scaling Down Applications	21

CHAPTER 1. ABOUT THIS GUIDE



WARNING

Red Hat is currently reviewing the information and procedures provided in this guide for this release.

This document is based on the Red Hat OpenStack Platform 10 document, available at https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/?version=10.

If you require assistance for Red Hat OpenStack Platform 11, please contact Red Hat support.

CHAPTER 2. CONFIGURE AUTO SCALING FOR COMPUTE INSTANCES

This guide describes how to automatically scale out your Compute instances in response to heavy system usage. By using pre-defined rules that consider factors such as CPU or memory usage, you can configure Orchestration (heat) to add and remove additional instances automatically, when they are needed.

2.1. ARCHITECTURAL OVERVIEW

2.1.1. Orchestration

The core component providing automatic scaling is Orchestration (heat). Orchestration allows you to define rules using human-readable YAML templates. These rules are applied to evaluate system load based on Telemetry data to find out whether there is need to add more instances into the stack. Once the load has dropped, Orchestration can automatically remove the unused instances again.

2.1.2. Telemetry

Telemetry does performance monitoring of your OpenStack environment, collecting data on CPU, storage, and memory utilization for instances and physical hosts. Orchestration templates examine Telemetry data to assess whether any pre-defined action should start.

2.1.3. Key Terms

- **Stack** - A stack stands for all the resources necessary to operate an application. It can be as simple as a single instance and its resources, or as complex as multiple instances with all the resource dependencies that comprise a multi-tier application.
- **Templates** - YAML scripts that define a series of tasks for Heat to execute. For example, it is preferable to use separate templates for certain functions:
 - **Template File** - This is where you define thresholds that Telemetry should respond to, and define the auto scaling group.
 - **Environment File** - Defines the build information for your environment: which flavor and image to use, how the virtual network should be configured, and what software should be installed.

2.2. EXAMPLE: AUTO SCALING BASED ON CPU USAGE

In this example, Orchestration examines Telemetry data, and automatically increases the number of instances in response to high CPU usage. A stack template and environment template are created to define the needed rules and subsequent configuration. This example makes use of existing resources (such as networks), and uses names that are likely to differ in your own environment.

1. Create the environment template, describing the instance flavor, networking configuration, and image type and save it in the template `/home/<user>/stacks/example1/cirros.yaml` file. Please, replace the `<user>` variable with a real user name:

```
heat_template_version: 2016-10-14
description: Template to spawn an cirros instance.
```



```

parameters:
  metadata:
    type: json
  image:
    type: string
    description: image used to create instance
    default: cirros
  flavor:
    type: string
    description: instance flavor to be used
    default: m1.tiny
  key_name:
    type: string
    description: keypair to be used
    default: mykeypair
  network:
    type: string
    description: project network to attach instance to
    default: internal1
  external_network:
    type: string
    description: network used for floating IPs
    default: external_network

resources:
  server:
    type: OS::Nova::Server
    properties:
      block_device_mapping:
        - device_name: vda
          delete_on_termination: true
          volume_id: { get_resource: volume }
      flavor: {get_param: flavor}
      key_name: {get_param: key_name}
      metadata: {get_param: metadata}
      networks:
        - port: { get_resource: port }

  port:
    type: OS::Neutron::Port
    properties:
      network: {get_param: network}
      security_groups:
        - default

  floating_ip:
    type: OS::Neutron::FloatingIP
    properties:
      floating_network: {get_param: external_network}

  floating_ip_assoc:
    type: OS::Neutron::FloatingIPAssociation
    properties:
      floatingip_id: { get_resource: floating_ip }
      port_id: { get_resource: port }

```

```

volume:
  type: OS::Cinder::Volume
  properties:
    image: {get_param: cirros}
    size: 1

```

2. Register the Orchestration resource in `~/stacks/example1/environment.yaml`:


```

resource_registry:

  "OS::Nova::Server::Cirros": ~/stacks/example1/cirros.yaml

```

3. Create the stack template, describing the CPU thresholds to watch for, and how many instances should be added. An instance group is also created, defining the minimum and maximum number of instances that can participate in this template.



NOTE

The **granularity** parameter needs to be set according to **gnocchi cpu_util** metric granularity. For more information, refer to this [solution article](#).

Save the following values in `~/stacks/example1/template.yaml`:

```

heat_template_version: 2016-10-14
description: Example auto scale group, policy and alarm
resources:
  scaleup_group:
    type: OS::Heat::AutoScalingGroup
    properties:
      cooldown: 300
      desired_capacity: 1
      max_size: 3
      min_size: 1
      resource:
        type: OS::Nova::Server::Cirros
        properties:
          metadata: {"metering.server_group": {get_param:
"OS::stack_id"}}

  scaleup_policy:
    type: OS::Heat::ScalingPolicy
    properties:
      adjustment_type: change_in_capacity
      auto_scaling_group_id: { get_resource: scaleup_group }
      cooldown: 300
      scaling_adjustment: 1

  scaledown_policy:
    type: OS::Heat::ScalingPolicy
    properties:
      adjustment_type: change_in_capacity
      auto_scaling_group_id: { get_resource: scaleup_group }
      cooldown: 300
      scaling_adjustment: -1

```

```

cpu_alarm_high:
  type: OS::Aodh::GnocchiAggregationByResourcesAlarm
  properties:
    description: Scale up if CPU > 80%
    metric: cpu_util
    aggregation_method: mean
    granularity: 300
    evaluation_periods: 1
    threshold: 80
    resource_type: instance
    comparison_operator: gt
    alarm_actions:
      - str_replace:
          template: trust+url
          params:
            url: {get_attr: [scaleup_policy, signal_url]}
    query:
      str_replace:
        template: '{"=": {"server_group": "stack_id"}}'
        params:
          stack_id: {get_param: "OS::stack_id"}

cpu_alarm_low:
  type: OS::Aodh::GnocchiAggregationByResourcesAlarm
  properties:
    metric: cpu_util
    aggregation_method: mean
    granularity: 300
    evaluation_periods: 1
    threshold: 5
    resource_type: instance
    comparison_operator: lt
    alarm_actions:
      - str_replace:
          template: trust+url
          params:
            url: {get_attr: [scaledown_policy, signal_url]}
    query:
      str_replace:
        template: '{"=": {"server_group": "stack_id"}}'
        params:
          stack_id: {get_param: "OS::stack_id"}

outputs:
  scaleup_policy_signal_url:
    value: {get_attr: [scaleup_policy, signal_url]}

  scaledown_policy_signal_url:
    value: {get_attr: [scaledown_policy, signal_url]}

```

4. Run the following OpenStack command to build the environment and deploy the instance:

```

$ openstack stack create -t template.yaml -e environment.yaml
example
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

```

--+
| Field                | Value
|
+-----+-----+
--+
| id                   | 248a98bb-f56e-4934-a281-fffde62d78d8
|
| stack_name          | example
|
| description          | Example auto scale group, policy and alarm |
| creation_time       | 2017-03-06T15:00:29Z
|
| updated_time        | None
|
| stack_status         | CREATE_IN_PROGRESS
|
| stack_status_reason | Stack CREATE started
|
+-----+-----+
--+

```

- Orchestration will create the stack and launch a defined minimum number of cirros instances, as defined in the `min_size` parameter of the `scaleup_group` definition. Verify that the instances were created successfully:

```

$ openstack server list
+-----+-----+-----+-----+-----+-----+
| ID                | Name
| Status | Task State | Power State | Networks
|
+-----+-----+-----+-----+-----+-----+
| e1524f65-5be6-49e4-8501-e5e5d812c612 | ex-3gax-5f3a4og5cwn2-
png47w3u2vjd-server-vaajhuv4mj3j | ACTIVE | -          | Running
| internal1=10.10.10.9, 192.168.122.8 |
+-----+-----+-----+-----+-----+-----+
--+

```

- Orchestration also creates two cpu alarms which are used to trigger scale-up or scale-down events, as defined in `cpu_alarm_high` and `cpu_alarm_low`. Verify that the triggers exist:

```

$ openstack alarm list
+-----+-----+-----+-----+-----+-----+
| alarm_id          | type
| name              | state
| severity | enabled |
+-----+-----+-----+-----+-----+-----+
--+

```

```

| 022f707d-46cc-4d39-a0b2-afd2fc7ab86a |
gnocchi_aggregation_by_resources_threshold | example-cpu_alarm_high-
odj77qpbld7j | insufficient data | low      | True      |
| 46ed2c50-e05a-44d8-b6f6-f1ebd83af913 |
gnocchi_aggregation_by_resources_threshold | example-cpu_alarm_low-
m37jvnm56x2t | insufficient data | low      | True      |
+-----+-----+-----+-----+
-----+-----+-----+-----+
-----+-----+-----+-----+

```

2.2.1. Test Automatic Scaling Up Instances

Orchestration can scale instances automatically based on the **cpu_alarm_high** threshold definition. Once the CPU utilization reaches a value defined in the **threshold** parameter, another instance is started to balance the load. The **threshold** value in the above **template.yaml** file is set to 80%.

1. Login to the instance and run several **dd** commands to generate the load:

```

$ ssh -i ~/mykey.pem cirros@192.168.122.8
$ sudo dd if=/dev/zero of=/dev/null &
$ sudo dd if=/dev/zero of=/dev/null &
$ sudo dd if=/dev/zero of=/dev/null &

```

2. Having run the **dd** commands, you can expect to have 100% CPU utilization in the cirros instance. Verify that the alarm has been triggered:

```

$ openstack alarm list
+-----+-----+-----+-----+
-----+-----+-----+-----+
+-----+-----+
| alarm_id                               | type
| name                                   | state | severity | enabled
|
+-----+-----+-----+-----+
-----+-----+-----+-----+
+-----+-----+
| 022f707d-46cc-4d39-a0b2-afd2fc7ab86a |
gnocchi_aggregation_by_resources_threshold | example-cpu_alarm_high-
odj77qpbld7j | alarm | low      | True      |
| 46ed2c50-e05a-44d8-b6f6-f1ebd83af913 |
gnocchi_aggregation_by_resources_threshold | example-cpu_alarm_low-
m37jvnm56x2t | ok    | low      | True      |
+-----+-----+-----+-----+
-----+-----+-----+-----+
+-----+-----+

```

3. After some time (approximately 60 seconds), Orchestration will start another instance and add it into the group. You can verify this with the **nova list** command:

```

$ openstack server list
+-----+-----+-----+-----+
-----+-----+-----+-----+
+-----+-----+
| ID                               | Name

```

```

| Status | Task State | Power State | Networks
|
+-----+-----+-----+-----+
-----+-----+-----+-----+
-+-----+
| 477ee1af-096c-477c-9a3f-b95b0e2d4ab5 | ex-3gax-4urpikl5koff-
yrxk3zxzfmpf-server-2hde4tp4trnk | ACTIVE | -           | Running
| internal1=10.10.10.13, 192.168.122.17 |
| e1524f65-5be6-49e4-8501-e5e5d812c612 | ex-3gax-5f3a4og5cwn2-
png47w3u2vjd-server-vaajhuv4mj3j | ACTIVE | -           | Running
| internal1=10.10.10.9, 192.168.122.8   |
+-----+-----+-----+-----+
-----+-----+-----+-----+
-+-----+

```

- After another short period, you will observe that Orchestration has auto scaled again to three instances. The configuration is set to three instances maximally, so it will not scale any higher (the **scaleup_group** definition: **max_size**). Again, you can verify that with the above mentioned command:

```

$ openstack server list
+-----+-----+-----+-----+
-----+-----+-----+-----+
-+-----+
| ID                               | Name
| Status | Task State | Power State | Networks
|
+-----+-----+-----+-----+
-----+-----+-----+-----+
-+-----+
| 477ee1af-096c-477c-9a3f-b95b0e2d4ab5 | ex-3gax-4urpikl5koff-
yrxk3zxzfmpf-server-2hde4tp4trnk | ACTIVE | -           | Running
| internal1=10.10.10.13, 192.168.122.17 |
| e1524f65-5be6-49e4-8501-e5e5d812c612 | ex-3gax-5f3a4og5cwn2-
png47w3u2vjd-server-vaajhuv4mj3j | ACTIVE | -           | Running
| internal1=10.10.10.9, 192.168.122.8   |
| 6c88179e-c368-453d-a01a-555eae8cd77a | ex-3gax-fvxz3tr63j4o-
36fhftuja3bw-server-rhl4sqkjuy5p | ACTIVE | -           | Running
| internal1=10.10.10.5, 192.168.122.5   |
+-----+-----+-----+-----+
-----+-----+-----+-----+
-+-----+

```

2.2.2. Automatically Scaling Down Instances

Orchestration can also automatically scale down instances based on the **cpu_alarm_low** threshold. In this example, the instances are scaled down once CPU utilization is below 5%.

- Terminate the running **dd** processes and you will observe Orchestration begin to scale the instances back down.

```
$ killall dd
```

- Stopping the **dd** processes causes the **cpu_alarm_low event** to trigger. As a result, Orchestration begins to automatically scale down and remove the instances. Verify, that the corresponding alarm has been triggered.

```
$ openstack alarm list
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| alarm_id | type | state | severity | enabled |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 022f707d-46cc-4d39-a0b2-afd2fc7ab86a | gnocchi_aggregation_by_resources_threshold | example-cpu_alarm_high-odj77qpbld7j | ok | low | True |
| 46ed2c50-e05a-44d8-b6f6-f1ebd83af913 | gnocchi_aggregation_by_resources_threshold | example-cpu_alarm_low-m37jvnm56x2t | alarm | low | True |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
```

After a few minutes, Orchestration continually reduce the number of instances to the minimum value defined in the **min_size** parameter of the **scaleup_group** definition. In this scenario, the **min_size** parameter is set to **1**.

2.2.3. Troubleshooting the setup

If your environment is not working properly, you can look for errors in the log files and history records.

- To get information on state transitions, you can list the stack event records:

```
$ openstack stack event list example
2017-03-06 11:12:43Z [example]: CREATE_IN_PROGRESS Stack CREATE
started
2017-03-06 11:12:43Z [example.scaleup_group]: CREATE_IN_PROGRESS
state changed
2017-03-06 11:13:04Z [example.scaleup_group]: CREATE_COMPLETE state
changed
2017-03-06 11:13:04Z [example.scaledown_policy]: CREATE_IN_PROGRESS
state changed
2017-03-06 11:13:05Z [example.scaleup_policy]: CREATE_IN_PROGRESS
state changed
2017-03-06 11:13:05Z [example.scaledown_policy]: CREATE_COMPLETE
state changed
2017-03-06 11:13:05Z [example.scaleup_policy]: CREATE_COMPLETE
state changed
2017-03-06 11:13:05Z [example.cpu_alarm_low]: CREATE_IN_PROGRESS
state changed
2017-03-06 11:13:05Z [example.cpu_alarm_high]: CREATE_IN_PROGRESS
state changed
2017-03-06 11:13:06Z [example.cpu_alarm_low]: CREATE_COMPLETE state
changed
```

```

2017-03-06 11:13:07Z [example.cpu_alarm_high]: CREATE_COMPLETE
state changed
2017-03-06 11:13:07Z [example]: CREATE_COMPLETE Stack CREATE
completed successfully
2017-03-06 11:19:34Z [example.scaleup_policy]: SIGNAL_COMPLETE
alarm state changed from alarm to alarm (Remaining as alarm due to 1
samples outside threshold, most recent: 95.4080102993)
2017-03-06 11:25:43Z [example.scaleup_policy]: SIGNAL_COMPLETE
alarm state changed from alarm to alarm (Remaining as alarm due to 1
samples outside threshold, most recent: 95.8869217299)
2017-03-06 11:33:25Z [example.scaledown_policy]: SIGNAL_COMPLETE
alarm state changed from ok to alarm (Transition to alarm due to 1
samples outside threshold, most recent: 2.73931707966)
2017-03-06 11:39:15Z [example.scaledown_policy]: SIGNAL_COMPLETE
alarm state changed from alarm to alarm (Remaining as alarm due to 1
samples outside threshold, most recent: 2.78110858552)

```

2. To read the alarm history log:

```

$ openstack alarm-history show 022f707d-46cc-4d39-a0b2-afd2fc7ab86a
+-----+-----+-----+
| timestamp                | type                | detail                |
| event_id                  |                     |                       |
+-----+-----+-----+
| 2017-03-06T11:32:35.510000 | state transition | {"transition_reason": "Transition to ok due to 1 samples inside threshold, most recent: 25e0e70b-3eda-466e-abac-42d9cf67e704", "state": "ok"} |
| 2017-03-06T11:17:35.403000 | state transition | {"transition_reason": "Transition to alarm due to 1 samples outside threshold, most recent: 8322f62c-0d0a-4dc0-9279-435510f81039", "state": "alarm"} |
| 2017-03-06T11:15:35.723000 | state transition | {"transition_reason": "Transition to ok due to 1 samples inside threshold, most recent: 1503bd81-7eba-474e-b74e-ded8a7b630a1", "state": "ok"} |
| 2017-03-06T11:13:06.413000 | creation          | {"alarm_actions": ["trust+http://fca6e27e3d524ed68abdc0fd576aa848:delete@192.168.122.126:8004/v1/fd/224f15c0-b6f1-4690-9a22-0c1d236e65f6/1c345135be4ee587fef424c241719d/stacks/example/d9ef59ed-b8f8-4e90-bd9b-"] |

```



```

|
| ae87e73ef6e2/resources/scaleup_policy/signal"], "user_id":
| "a85f83b7f7784025b6acdc06ef0a8fd8",
|
|
| "example-cpu_alarm_high-odj77qpbld7j", "state": "insufficient data",
| "timestamp":
| "2017-03-
| 06T11:13:06.413455", "description": "Scale up if CPU > 80%",
| "enabled": true,
|
|
| "state_timestamp": "2017-03-06T11:13:06.413455", "rule":
| {"evaluation_periods": 1, "metric":
|
| "cpu_util",
| "aggregation_method": "mean", "granularity": 300, "threshold": 80.0,
| "query": "{\\\"=\\\":
|
| {\\\"server_group\\\": \\\"d9ef59ed-b8f8-4e90-bd9b-ae87e73ef6e2\\\"}}",
| "comparison_operator": "gt",
|
|
| "instance"}, "alarm_id": "022f707d-46cc-4d39-a0b2-afd2fc7ab86a",
|
|
| "time_constraints": [], "insufficient_data_actions": null,
| "repeat_actions": true, "ok_actions":
|
|
| "project_id": "fd1c345135be4ee587fef424c241719d", "type":
|
|
| "gnocchi_aggregation_by_resources_threshold", "severity": "low"}
|
|
+-----+-----+-----+-----+-----+-----+-----+-----+
|
+-----+-----+-----+-----+-----+-----+-----+-----+

```

3. To see the records of scale-out or scale-down operations that heat collects for the existing stack, you can use **awk** to parse the **heat-engine.log**:

```

$ awk '/Stack UPDATE started/,/Stack CREATE completed successfully/
{print $0}' /var/log/heat/heat-engine.log

```

4. To see the **aodh** related information, examine the **evaluator.log**:

```

$ grep -i alarm /var/log/aodh/evaluator.log | grep -i transition

```

2.3. EXAMPLE: AUTO SCALING APPLICATIONS

The functionality described earlier can also be used to scale up applications; for example, a dynamic web page that be served by one of multiple instances running at a time. In this case, *neutron* can be configured to provide *Load Balancing-as-a-Service*, which works to evenly distribute traffic among

instances.

In the following example, Orchestration again examines Telemetry data and increases the number of instances if high CPU usage is detected, or decreases the number of instances if CPU usage returns below a set value.

1. Create the template describing the properties of the *load-balancer* environment. Enter the following values in `~/stacks/example2/lb-env.yaml`:

```

heat_template_version: 2014-10-16
description: A load-balancer server
parameters:
  image:
    type: string
    description: Image used for servers
  key_name:
    type: string
    description: SSH key to connect to the servers
  flavor:
    type: string
    description: flavor used by the servers
  pool_id:
    type: string
    description: Pool to contact
  user_data:
    type: string
    description: Server user_data
  metadata:
    type: json
  network:
    type: string
    description: Network used by the server

resources:
  server:
    type: OS::Nova::Server
    properties:
      flavor: {get_param: flavor}
      image: {get_param: image}
      key_name: {get_param: key_name}
      metadata: {get_param: metadata}
      user_data: {get_param: user_data}
      networks:
        - port: { get_resource: port }

  member:
    type: OS::Neutron::PoolMember
    properties:
      pool_id: {get_param: pool_id}
      address: {get_attr: [server, first_address]}
      protocol_port: 80

  port:
    type: OS::Neutron::Port
    properties:
      network: {get_param: network}

```

```

    security_groups:
      - base

outputs:
  server_ip:
    description: IP Address of the load-balanced server.
    value: { get_attr: [server, first_address] }
  lb_member:
    description: LB member details.
    value: { get_attr: [member, show] }

```

2. Create another template for the instances that will be running the web application. The following template creates a load balancer and uses the existing networks. Be sure to replace the parameters according to your environment, and save the template in a file such as `~/stacks/example2/lb-webserver-rhel7.yaml`:

```

heat_template_version: 2014-10-16
description: AutoScaling RHEL 7 Web Application
parameters:
  image:
    type: string
    description: Image used for servers
    default: RHEL 7
  key_name:
    type: string
    description: SSH key to connect to the servers
    default: admin
  flavor:
    type: string
    description: flavor used by the web servers
    default: m2.tiny
  network:
    type: string
    description: Network used by the server
    default: private
  subnet_id:
    type: string
    description: subnet on which the load balancer will be located
    default: 9daa6b7d-e647-482a-b387-dd5f855b88ef
  external_network_id:
    type: string
    description: UUID of a Neutron external network
    default: db17c885-77fa-45e8-8647-dbb132517960

resources:
  webserver:
    type: OS::Heat::AutoScalingGroup
    properties:
      min_size: 1
      max_size: 3
      cooldown: 60
      desired_capacity: 1
    resource:
      type: file:///etc/heat/templates/lb-env.yaml
      properties:
        flavor: {get_param: flavor}

```

```

        image: {get_param: image}
        key_name: {get_param: key_name}
        network: {get_param: network}
        pool_id: {get_resource: pool}
        metadata: {"metering.stack": {get_param: "OS::stack_id"}}
        user_data:
            str_replace:
                template: |
                    #!/bin/bash -v

                    yum -y install httpd php
                    systemctl enable httpd
                    systemctl start httpd
                    cat <<EOF > /var/www/html/hostname.php
                    <?php echo "Hello, My name is " . php_uname('n'); ?
>
                EOF
            params:
                hostip: 192.168.122.70
                fqdn: sat6.example.com
                shortname: sat6

web_server_scaleup_policy:
    type: OS::Heat::ScalingPolicy
    properties:
        adjustment_type: change_in_capacity
        auto_scaling_group_id: {get_resource: webserver}
        cooldown: 60
        scaling_adjustment: 1

web_server_scaledown_policy:
    type: OS::Heat::ScalingPolicy
    properties:
        adjustment_type: change_in_capacity
        auto_scaling_group_id: {get_resource: webserver}
        cooldown: 60
        scaling_adjustment: -1

cpu_alarm_high:
    type: OS::Aodh::GnocchiAggregationByResourcesAlarm
    properties:
        description: Scale-up if the average CPU > 95% for 1 minute
        meter_name: cpu_util
        statistic: avg
        period: 60
        evaluation_periods: 1
        threshold: 95
        alarm_actions:
            - str_replace:
                template: trust+url
                params:
                    url: {get_attr: [server_scaleup_policy,
signal_url]}
                matching_metadata: {'metadata.user_metadata.stack':
{get_param: "OS::stack_id"}}
                comparison_operator: gt

```

```

cpu_alarm_low:
  type: OS::Aodh::GnocchiAggregationByResourcesAlarm
  properties:
    description: Scale-down if the average CPU < 15% for 1 minute
    meter_name: cpu_util
    statistic: avg
    period: 60
    evaluation_periods: 1
    threshold: 15
    alarm_actions:
      - str_replace:
          template: trust+url
          params:
            url: {get_attr: [server_scaleup_policy,
signal_url]}
            matching_metadata: {'metadata.user_metadata.stack':
{get_param: "OS::stack_id"}}
            comparison_operator: lt

monitor:
  type: OS::Neutron::HealthMonitor
  properties:
    type: TCP
    delay: 5
    max_retries: 5
    timeout: 5

pool:
  type: OS::Neutron::Pool
  properties:
    protocol: HTTP
    monitors: [{get_resource: monitor}]
    subnet_id: {get_param: subnet_id}
    lb_method: ROUND_ROBIN
    vip:
      protocol_port: 80

lb:
  type: OS::Neutron::LoadBalancer
  properties:
    protocol_port: 80
    pool_id: {get_resource: pool}

lb_floating:
  type: OS::Neutron::FloatingIP
  properties:
    floating_network_id: {get_param: external_network_id}
    port_id: {get_attr: [pool, vip, port_id]}

outputs:
  scale_up_url:
    description: >
      This URL is the webhook to scale up the autoscaling group.
You
      can invoke the scale-up operation by doing an HTTP POST to

```

```

this
  URL; no body nor extra headers are needed.
  value: {get_attr: [web_server_scaleup_policy, alarm_url]}
scale_dn_url:
  description: >
    This URL is the webhook to scale down the autoscaling group.
    You can invoke the scale-down operation by doing an HTTP POST
to
  this URL; no body nor extra headers are needed.
  value: {get_attr: [web_server_scaledown_policy, alarm_url]}
pool_ip_address:
  value: {get_attr: [pool, vip, address]}
  description: The IP address of the load balancing pool
website_url:
  value:
    str_replace:
      template: http://serviceip/hostname.php
      params:
        serviceip: { get_attr: [lb_floating, floating_ip_address]
}
  description: >
    This URL is the "external" URL that can be used to access the
    website.
ceilometer_query:
  value:
    str_replace:
      template: >
        ceilometer statistics -m cpu_util
        -q metadata.user_metadata.stack=stackval -p 60 -a avg
      params:
        stackval: { get_param: "OS::stack_id" }
  description: >
    This is a Ceilometer query for statistics on the cpu_util
meter
  Samples about OS::Nova::Server instances in this stack. The -
q
  parameter selects Samples according to the subject's metadata.
  When a VM's metadata includes an item of the form
metering.X=Y,
  the corresponding Ceilometer resource has a metadata item of
the
  form user_metadata.X=Y and samples about resources so tagged
can
  be queried with a Ceilometer query term of the form
  metadata.user_metadata.X=Y. In this case the nested stacks
give
  their VMs metadata that is passed as a nested stack parameter,
  and this stack passes a metadata of the form metering.stack=Y,
  where Y is this stack's ID.

```

- Update the Telemetry collection interval. By default, Telemetry polls instances every 10 minutes for CPU data. For this example, change the interval to 60 seconds in **/etc/ceilometer/pipeline.yaml**:

```

- name: cpu_source
  interval: 60

```

```

meters:
- "cpu"
sinks:
- cpu_sink

```



NOTE

A polling period of 60 seconds is not recommended for production environments, as a higher polling interval can result in increased load on the control plane.

- Restart all OpenStack ceilometer services to apply the updated Telemetry setting:

```
# systemctl restart openstack-ceilometer*
```

- Run the Orchestration scripts. This will build the environment and use the template to deploy the instance:

```
# heat stack-create webfarm -f /root/lb-webserver-rhel7.yaml
```

Replace `/root/lb-webserver-rhel7.yaml` with the actual path and file name.

You can monitor the creation of the stack in Dashboard under *Orchestration* → *Stacks* → *Webfarm*. Once the stack has been created, you are presented with multiple useful pieces of information, notably:

- URLs that you can use to trigger manual scale-up or scale-down events.
- The floating IP address, which is the IP address of the website.
- The Telemetry command which shows the CPU load for the whole stack, and which you can use to check whether the scaling is working as expected.

This is what the page looks like in Dashboard:

The screenshot shows the OpenStack Dashboard interface. The top navigation bar includes 'RED HAT ENTERPRISE LINUX OPENSTACK PLATFORM' and various menu items like 'Project', 'Admin', 'Identity', 'Compute', 'Network', 'Object Store', 'Orchestration', 'Database', and 'Data Processing'. The main content area is titled 'Stacks' and 'Resource Types'. Below this, there are tabs for 'Topology', 'Overview', 'Resources', 'Events', and 'Template'. The 'Stack Overview' section is active, displaying the following information:

- Information:** Name: webfarm, ID: 8f8c3d515c14a64-b9e8-70215498c046, Description: AutoScaling RHEL7 Web Application.
- Status:** Created: 1 minute, Last Updated: Never, Status: Create_Complete: Stack CREATE completed successfully.
- Outputs:**
 - pool_ip_address:** The IP address of the load balancing pool, 10.10.1.175.
 - scale_dn_url:** This URL is the webhook to scale down the autoscaling group. You can invoke the scale-down operation by doing an HTTP POST to this URL; no body nor extra headers are needed. http://192.168.122.80:8000/v1/signal/arm%3Aopenstack%3Aheat%3A%3A8cfd9e086c54ad96cc3d40181ea284%3Astack%2Fwebfarm%2F8f8c3d515c14a64-b9e8-70215498c046%2Fresources%2Fweb_server_scaledown_policy?TImestamp=2015-10-05T09%3A46%3A32Z&SignatureMethod=HmacSHA256&AWSAccessKeyId=b2bbbeba4b475492be77b28c3c9ef8&SignatureVersion=2&Signature=BfVtqU9%2FYm12hymmpc411440lbnK00K1398c%3D
 - scale_up_url:** This URL is the webhook to scale up the autoscaling group. You can invoke the scale-up operation by doing an HTTP POST to this URL; no body nor extra headers are needed. http://192.168.122.80:8000/v1/signal/arm%3Aopenstack%3Aheat%3A%3A8cfd9e086c54ad96cc3d40181ea284%3Astack%2Fwebfarm%2F8f8c3d515c14a64-b9e8-70215498c046%2Fresources%2Fweb_server_scaleup_policy?TImestamp=2015-10-05T09%3A46%3A32Z&SignatureMethod=HmacSHA256&AWSAccessKeyId=85a6a7e495a844839f0978951efatb1&SignatureVersion=2&Signature=dR1oVKEIST%2FLKIKL0M9ovZPP3IC%2Bz8KUGq6WTDNg%3D
 - website_url:** This URL is the "external" URL that can be used to access the website. <http://192.168.122.179/hostname.php>
 - ceilometer_query:** This is a Ceilometer query for statistics on the cpu_util meter. Samples about OS::Nova::Server instances in this stack. The -q parameter selects Samples according to the subject's metadata. When a VM's metadata includes an item of the form metering.X:Y, the corresponding Ceilometer resource has a metadata item of the form user_metadata.X:Y and samples about resources so tagged can be queried with a Ceilometer query term of the form metadata.user_metadata.X:Y. In this case the nested stacks give their VMs metadata that is passed as a nested stack parameter, and this stack passes a metadata of the form metering.stack=Y, where Y is this stack's ID. `ceilometer statistics -m cpu_util -q metadata.user_metadata.stack=8f8c3d515c14a64-b9e8-70215498c046 -p 600 -a avg`

Open *Network* → *Load Balancers* to view the load balancer:

Name	Description	Provider	Subnet	Protocol	Status	VIP	Actions
webfarm-pool-hsmjmgaa2d7		haproxy	10.10.1.0/24	HTTP	Active	pool.vip	Edit Pool

Click *Members*. This page displays the members of the load balancing pool; these are the instances to which the website traffic can be distributed. Note that a member will not have the *Active* status until the corresponding instance has been created, and Apache has been installed and configured.

When the web server has started, the instance is visible as an active member of the load balancer:

IP Address	Protocol Port	Weight	Pool	Status	Actions
10.10.1.176	80	1	webfarm-pool-hsmjmgaa2d7	Active	Edit Member

You are now able to access the web application at `http://IP/hostname.php`. You can expect to see output similar to the following:

```
Hello, My name is we-zrwm-t4ezkpx34gxu-qbg5d7dqbc4j-server-mzdvigk2jugl
```

You can now view the stack's CPU performance data by running the Telemetry command from the stack overview in Dashboard. The command looks like the following:

```
# ceilometer statistics -m cpu_util -q
metadata.user_metadata.stack=8f86c3d5-15cf-4a64-b9e8-70215498c046 -p 60 -a
avg
```

2.3.1. Test Auto Scaling Applications

To manually trigger application scaling, use the REST *scale-up URL* from the stack overview in Dashboard, or generate load by running a resource-intensive command on the initially deployed instance.


- To use the REST API, you need a tool which can perform **HTTP POST** requests, such as the [REST Easy Firefox add on](#) or **curl**. Copy the *scale-up URL* and either paste it into the *REST Easy* form:

Or use it as a parameter on the **curl** command line:

```
$ curl -X POST "scale-up URL"
```


- To artificially generate load, allocate a floating IP to the instance, log in to it with SSH, and run a command which will keep the CPU busy. For example:


```
$ dd if=/dev/zero of=/dev/null &
```



IMPORTANT

Check whether CPU usage is above 95%, for example, using the `top` command. If the CPU usage is not sufficiently high, run the `dd` command multiple times in parallel, or use another method to keep the CPU busy.

The next time Telemetry collects CPU data from the stack, the scale-up event will trigger and appear at *Orchestration* → *Stacks* → *Webfarm* → *Events*. A new web server instance will be created and added to the load balancer. When this is done, the instance becomes active, and you will notice that the website URL is routed through the load balancer to both instances in the stack.




NOTE

The creation can take several minutes because the instance must be initialized, Apache installed and configured, and the application deployed. This is monitored by HAProxy, which ensures that the website is available on the instance before it is marked as active.

This is what the list of members of the load balancing pool looks like in the Dashboard while the new instance is being created:

IP Address	Protocol Port	Weight	Pool	Status	Actions
10.10.1.176	80	1	webfarm-pool-hsmjmgaa217	Active	Edit Member
10.10.1.177	80	1	webfarm-pool-hsmjmgaa217	Inactive	Edit Member



IMPORTANT

The average CPU usage of the instances in the *heat* stack is taken into account when deciding whether or not an additional instance gets created. Because the second instance will most likely have normal CPU usage, it will balance out the first instance. However, if the second instance becomes busy as well and the average CPU usage of the first and second instance exceeds 95%, another (third) instance will be created.

2.3.2. Automatically Scaling Down Applications

This is similar to [Section 2.2.2, “Automatically Scaling Down Instances”](#) in that the scale-down policy is triggered when the average CPU usage for the stack drops below a predefined value, which is 15% in the example described in [Section 2.3.1, “Test Auto Scaling Applications”](#). In addition, when an instance is removed from the stack this way, it is also automatically removed from the load balancer. The website traffic is then automatically distributed among the rest of the instances.

