

---

# **OpenStack-Ansible Documentation:**

## **os\_magnum role**

*Release 18.1.0.dev237*

**OpenStack-Ansible Contributors**

**May 11, 2024**



## CONTENTS

<b>1</b>	<b>Default variables</b>	<b>3</b>
<b>2</b>	<b>Dependencies</b>	<b>9</b>
<b>3</b>	<b>Wiring docker with cinder</b>	<b>11</b>
<b>4</b>	<b>Example playbook</b>	<b>13</b>
<b>5</b>	<b>Tags</b>	<b>15</b>
<b>6</b>	<b>Post-deployment configuration</b>	<b>17</b>
<b>7</b>	<b>Post-deployment debugging</b>	<b>21</b>



Ansible role that installs and configures OpenStack Magnum. Magnum is installed behind the Apache webserver listening on port 9511 by default.

To clone or view the source code for this repository, visit the role repository for [os\\_magnum](#).



## DEFAULT VARIABLES

```
## Verbosity Options
debug: False

#python venv executable
magnum_venv_python_executable: "{{ openstack_venv_python_executable | default(
↳ 'python3') }}"

# Enable/Disable Ceilometer
magnum_ceilometer_enabled: "{{ (groups['ceilometer_all'] is defined) and
↳ (groups['ceilometer_all'] | length > 0) }}"

# Set the host which will execute the shade modules
# for the service setup. The host must already have
# clouds.yaml properly configured.
magnum_service_setup_host: "{{ openstack_service_setup_host | default(
↳ 'localhost') }}"
magnum_service_setup_host_python_interpreter: "{{ openstack_service_setup_
↳ host_python_interpreter | default((magnum_service_setup_host == 'localhost
↳ ') | ternary(ansible_playbook_python, ansible_facts['python']['executable
↳ '])) }}"

# Set the package install state for distribution packages
# Options are 'present' and 'latest'
magnum_package_state: "{{ package_state | default('latest') }}"

magnum_system_group_name: magnum
magnum_system_user_name: magnum
magnum_system_user_comment: Magnum System User
magnum_system_user_shell: /bin/false
magnum_system_user_home: "/var/lib/{{ magnum_system_user_name }}"
magnum_etc_directory: /etc/magnum

magnum_service_name: magnum
magnum_service_user_name: magnum
magnum_service_type: container-infra
magnum_service_description: "OpenStack Containers (Magnum)"
magnum_service_project_name: service
magnum_service_role_names:
```

(continues on next page)

(continued from previous page)

```

- admin
- service
magnum_service_token_roles:
- service
magnum_service_token_roles_required: "{{ openstack_service_token_roles_
↳required | default(True) }}"
magnum_service_region: "{{ service_region | default('RegionOne') }}"
magnum_barbican_service_region: "{{ magnum_service_region }}"
magnum_cinder_service_region: "{{ magnum_service_region }}"
magnum_glance_service_region: "{{ magnum_service_region }}"
magnum_heat_service_region: "{{ magnum_service_region }}"
magnum_neutron_service_region: "{{ magnum_service_region }}"
magnum_nova_service_region: "{{ magnum_service_region }}"
magnum_keystone_service_region: "{{ magnum_service_region }}"
magnum_octavia_service_region: "{{ magnum_service_region }}"
magnum_bind_port: 9511
magnum_service_proto: http
magnum_service_publicuri_proto: "{{ openstack_service_publicuri_proto |
↳default(magnum_service_proto) }}"
magnum_service_publicurl: "{{ magnum_service_publicuri_proto }}://{{ external_
↳lb_vip_address }}:{{ magnum_bind_port }}"
magnum_service_internaluri_proto: "{{ openstack_service_internaluri_proto |
↳default(magnum_service_proto) }}"
magnum_service_internalurl: "{{ magnum_service_internaluri_proto }}://{{
↳internal_lb_vip_address }}:{{ magnum_bind_port }}"
magnum_service_adminuri_proto: "{{ openstack_service_adminuri_proto |
↳default(magnum_service_proto) }}"
magnum_service_adminurl: "{{ magnum_service_adminuri_proto }}://{{ internal_
↳lb_vip_address }}:{{ magnum_bind_port }}"
magnum_service_in_ldap: "{{ service_ldap_backend_enabled | default(False) }}"

magnum_config_overrides: {}
magnum_policy_overrides: {}
magnum_api_paste_ini_overrides: {}
magnum_keystone_auth_default_policy: []

magnum_pip_install_args: "{{ pip_install_options | default('') }}"

# Name of the virtual env to deploy into
magnum_venv_tag: "{{ venv_tag | default('untagged') }}"
magnum_venv_path: "/openstack/venvs/magnum-{{ magnum_venv_tag }}"
magnum_bin: "{{ magnum_venv_path }}/bin"

magnum_git_repo: "https://opendev.org/openstack/magnum"
magnum_git_install_branch: master
magnum_upper_constraints_url: "{{ requirements_git_url | default('https://
↳releases.openstack.org/constraints/upper/' ~ requirements_git_install_
↳branch | default('master')) }}"
magnum_git_constraints:

```

(continues on next page)



(continued from previous page)

```

- "--constraint {{ magnum_upper_constraints_url }}"

# Database vars
magnum_db_setup_host: "{{ openstack_db_setup_host | default('localhost') }}"
magnum_db_setup_python_interpreter: "{{ openstack_db_setup_python_interpreter |
↳| default((magnum_db_setup_host == 'localhost') | ternary(ansible_playbook_
↳python, ansible_facts['python']['executable'])) }}"
magnum_galera_address: "{{ galera_address | default('127.0.0.1') }}"
magnum_galera_database_name: magnum_service
magnum_galera_user: magnum
magnum_galera_use_ssl: "{{ galera_use_ssl | default(False) }}"
magnum_galera_ssl_ca_cert: "{{ galera_ssl_ca_cert | default('') }}"
magnum_galera_port: "{{ galera_port | default('3306') }}"
magnum_db_max_overflow: "{{ openstack_db_max_overflow | default('50') }}"
magnum_db_max_pool_size: "{{ openstack_db_max_pool_size | default('5') }}"
magnum_db_pool_timeout: "{{ openstack_db_pool_timeout | default('30') }}"
magnum_db_connection_recycle_time: "{{ openstack_db_connection_recycle_time |
↳default('600') }}"

# Oslo Messaging vars
# RPC
magnum_oslomsg_rpc_host_group: "{{ oslomsg_rpc_host_group | default('rabbitmq_
↳all') }}"
magnum_oslomsg_rpc_setup_host: "{{ (magnum_oslomsg_rpc_host_group in groups) |
↳ternary(groups[magnum_oslomsg_rpc_host_group][0], 'localhost') }}"
magnum_oslomsg_rpc_transport: "{{ oslomsg_rpc_transport | default('rabbit') }}"
↳
magnum_oslomsg_rpc_servers: "{{ oslomsg_rpc_servers | default('127.0.0.1') }}"
magnum_oslomsg_rpc_port: "{{ oslomsg_rpc_port | default('5672') }}"
magnum_oslomsg_rpc_use_ssl: "{{ oslomsg_rpc_use_ssl | default(False) }}"
magnum_oslomsg_rpc_userid: magnum
magnum_oslomsg_rpc_vhost: /magnum
magnum_oslomsg_rpc_ssl_version: "{{ oslomsg_rpc_ssl_version | default('TLSv1.2
↳') }}"
magnum_oslomsg_rpc_ssl_ca_file: "{{ oslomsg_rpc_ssl_ca_file | default('') }}"

# Notify
magnum_oslomsg_notify_host_group: "{{ oslomsg_notify_host_group | default(
↳'rabbitmq_all') }}"
magnum_oslomsg_notify_setup_host: "{{ (magnum_oslomsg_notify_host_group in
↳groups) | ternary(groups[magnum_oslomsg_notify_host_group][0], 'localhost')
↳}}"
magnum_oslomsg_notify_transport: "{{ oslomsg_notify_transport | default(
↳'rabbit') }}"
magnum_oslomsg_notify_servers: "{{ oslomsg_notify_servers | default('127.0.0.1
↳') }}"
magnum_oslomsg_notify_port: "{{ oslomsg_notify_port | default('5672') }}"
magnum_oslomsg_notify_use_ssl: "{{ oslomsg_notify_use_ssl | default(False) }}"
magnum_oslomsg_notify_userid: "{{ magnum_oslomsg_rpc_userid }}"

```

(continues on next page)

(continued from previous page)

```

magnum_oslmsg_notify_password: "{{ magnum_oslmsg_rpc_password }}"
magnum_oslmsg_notify_vhost: "{{ magnum_oslmsg_rpc_vhost }}"
magnum_oslmsg_notify_ssl_version: "{{ oslmsg_notify_ssl_version | default(
  ↪ 'TLSv1_2') }}"
magnum_oslmsg_notify_ssl_ca_file: "{{ oslmsg_notify_ssl_ca_file | default(
  ↪ ') }}"

## (Qdrouterd) integration
# TODO(ansmith): Change structure when more backends will be supported
magnum_oslmsg_amqp1_enabled: "{{ magnum_oslmsg_rpc_transport == 'amqp' }}"

# Keystone AuthToken/Middleware
magnum_keystone_auth_plugin: password
magnum_service_project_domain_name: Default
magnum_service_user_domain_name: Default

# Trustee User
magnum_trustee_domain_admin_name: trustee_domain_admin
magnum_trustee_domain_name: magnum
magnum_trustee_domain_admin_roles:
  - admin
magnum_cluster_user_trust: True

#Glance images
## Example Glance Image - Fedora Atomic
# - name: fedora-atomic-latest           #Name of the image in Glance
#   disk_format: qcow2                 #Disk format (e.g. qcow2)
#   image_format: bare                 #Image format
#   public: true                       #Boolean - is the image public
#   file: https://builds.coreos.fedoraproject.org/prod/streams/stable/builds/
  ↪ 31.20200210.3.0/x86_64/fedora-coreos-31.20200210.3.0-openstack.x86_64.qcow2.
  ↪ xz
#   distro: fedora-atomic              #Value for the os_distro metadata
#   checksum:
  ↪ "sha256:9a5252e24b82a5edb1ce75b05653f59895685b0f1028112462e908a12deae518"
magnum_glance_images: []

# Define cluster templates to create. It should be list of
# dictionaries with keys that are supported by os_coe_cluster_template
# module (https://docs.ansible.com/ansible/latest/modules/os\_coe\_cluster\_
  ↪ template\_module.html)
#magnum_cluster_templates:
# - name: k8s
#   cloud: default
#   coe: kubernetes
#   docker_volume_size: 50
#   external_network_id: public
#   network_driver: flannel
magnum_cluster_templates: []

```

(continues on next page)

(continued from previous page)

```

# Create extra flavors to be used by magnum cluster template. It should be_
↪list
# of dictionaries with keys that are supported by os_nova_flavor module
# (https://docs.ansible.com/ansible/latest/modules/os_nova_flavor_module.html)
#magnum_flavors:
# - name: k8s-pod
#   cloud: default
#   ram: 256
#   vcpus: 1
#   disk: 5
magnum_flavors: []

# Set the directory where the downloaded images will be stored
# on the magnum_service_setup_host host. If the host is localhost,
# then the user running the playbook must have access to it.
magnum_image_path: "{{ lookup('env', 'HOME') }}/openstack-ansible/magnum"
magnum_image_path_owner: "{{ lookup('env', 'USER') }}"

magnum_pip_packages:
- "git+{{ magnum_git_repo }}@{{ magnum_git_install_branch }}#egg=magnum"
- osprofiler
- PyMySQL
- pymemcache
- python-memcached
- systemd-python

# Memcached override
magnum_memcached_servers: "{{ memcached_servers }}"

# Specific pip packages provided by the user
magnum_user_pip_packages: []

magnum_optional_oslomsg_amqp1_pip_packages:
- oslo.messaging[amqp1]

# Store certificates in DB by default (x509keypair)
# Other valid values are: barbican, local
magnum_cert_manager_type: x509keypair

magnum_api_init_config_overrides: {}
magnum_conductor_init_config_overrides: {}

magnum_services:
  magnum-conductor:
    group: magnum_all
    service_name: magnum-conductor
    execstarts: "{{ magnum_bin }}/magnum-conductor"
    init_config_overrides: "{{ magnum_conductor_init_config_overrides }}"

```

(continues on next page)

(continued from previous page)

```
    start_order: 1
  magnum-api:
    group: magnum_all
    service_name: magnum-api
    init_config_overrides: "{{ magnum_api_init_config_overrides }}"
    start_order: 2
    wsgi_app: True
    wsgi_path: "{{ magnum_bin }}/magnum-api-wsgi"
    uwsgi_overrides: "{{ magnum_api_uwsgi_ini_overrides }}"
    uwsgi_port: "{{ magnum_bind_port }}"
    uwsgi_bind_address: "{{ magnum_api_uwsgi_bind_address }}"

# uWSGI Settings
magnum_api_uwsgi_ini_overrides: {}
magnum_wsgi_processes_max: 16
magnum_wsgi_processes: "{{ [(ansible_facts['processor_vcpus']//ansible_facts[
↪ 'processor_threads_per_core'])|default(1), 1] | max * 2, magnum_wsgi_
↪ processes_max] | min }}"
magnum_wsgi_threads: 1
magnum_api_uwsgi_bind_address: "{{ openstack_service_bind_address | default(
↪ '0.0.0.0') }}"

# conductor settings
magnum_conductor_workers_max: 16
magnum_conductor_workers: "{{ [(ansible_facts['processor_vcpus']//ansible_
↪ facts['processor_threads_per_core'])|default(1), 1] | max * 2, magnum_
↪ conductor_workers_max] | min }}"
```

## DEPENDENCIES

This role needs pip >= 7.1 installed on the target host.

To use this role, define the following variables:

```
# Magnum TCP listening port
magnum_service_port: 9511

# Magnum service protocol http or https
magnum_service_proto: http

# Magnum Galera address of internal load balancer
magnum_galera_address: "{{ internal_lb_vip_address }}"

# Magnum Galera database name
magnum_galera_database_name: magnum_service

# Magnum Galera username
magnum_galera_user: magnum

# Magnum rpc userid
magnum_oslmsg_rpc_userid: magnum

# Magnum rpc vhost
magnum_oslmsg_rpc_vhost: /magnum

# Magnum notify userid
magnum_oslmsg_notify_userid: magnum

# Magnum notify vhost
magnum_oslmsg_notify_vhost: /magnum
```

This list is not exhaustive. See role internals for further details.



## WIRING DOCKER WITH CINDER

If you need to use volumes, `default_docker_volume_type` should be set. By default, Magnum doesn't need one.

To deploy Magnum with cinder integration, please set the following in your `/etc/openstack_deploy/user_variables.yml`:

```
magnum_config_overrides:  
  cinder:  
    default_docker_volume_type: lvm
```

If you have defined `cinder_default_volume_type` for all your nodes, by defining it in your `user_variables`, you can re-use it directly:

```
magnum_config_overrides:  
  cinder:  
    default_docker_volume_type: "{{ cinder_default_volume_type }}"
```





## EXAMPLE PLAYBOOK

```
---
- name: Install Magnum server
  hosts: magnum_all
  user: root
  roles:
    - role: "os_magnum"
      tags:
        - os-magnum"
  vars:
    magnum_galera_address: "{{ internal_lb_vip_address }}"
    magnum_galera_password: secrete
    magnum_service_password: secrete
    magnum_oslmsg_rpc_password: secrete
    magnum_trustee_password: secrete
```



**TAGS**

This role supports two tags: `magnum-install` and `magnum-config`. The `magnum-install` tag can be used to install and upgrade. The `magnum-config` tag can be used to maintain configuration of the service.



---

## POST-DEPLOYMENT CONFIGURATION

---

Deploying the magnum service makes the API components available to use. Additional configuration is required to make a working Kubernetes cluster, including loading the correct Image and setting up a suitable Cluster Template

This example is intended to show the steps required and should be updated as needed for the version of k8s and associated components. The example has been tested by a deployer with magnum SHA `fe35af8ef5d9e65a4074aa3ba3ed3116b7322415`.

First, upload the coreos image. this can be done either manually or using the `os_magnum` playbooks.

Manual configuration:

```
wget https://builds.coreos.fedoraproject.org/prod/streams/stable/builds/32.
↪20201004.3.0/x86_64/fedora-coreos-32.20201004.3.0-openstack.x86_64.qcow2.xz

(convert to raw if necessary here for ceph backed storage)

openstack image create "fedora-coreos-latest" --disk-format raw --container-
↪format bare \
--file fedora-coreos-32.20201004.3.0-openstack.x86_64.raw --property os_
↪distro='fedora-coreos'
```

Via `os_magnum` playbooks and data in `user_variables.yml`

```
magnum_glance_images:
- name: fedora-coreos-latest
  disk_format: qcow2
  image_format: bare
  public: true
  file: https://builds.coreos.fedoraproject.org/prod/streams/stable/builds/
↪31.20200210.3.0/x86_64/fedora-coreos-31.20200210.3.0-openstack.x86_64.qcow2.
↪xz
  distro: "coreos"
  checksum:
↪"sha256:9a5252e24b82a5edb1ce75b05653f59895685b0f1028112462e908a12deae518"
```

Second, create the cluster template.

Manual configuration:

```

openstack coe cluster template create <name> --coe kubernetes --external-
↪network <ext-net> \
--image "fedora-coreos-latest" --master-flavor <flavor> --flavor <flavor> --
↪master-lb-enabled \
--docker-volume-size 50 --network-driver calico --docker-storage-driver_
↪overlay2 \
--volume-driver cinder \
--labels boot_volume_type=<your volume type>,boot_volume_size=50,kube_tag=v1.
↪18.6,availability_zone=nova,helm_client_url="https://get.helm.sh/helm-v3.4.
↪0-linux-amd64.tar.gz",helm_client_sha256=
↪"270acb0f085b72ec28aee894c7443739271758010323d72ced0e92cd2c96ffdb",helm_
↪client_tag="v3.4.0",etcd_volume_size=50,auto_scaling_enabled=true,auto_
↪healing_enabled=true,auto_healing_controller=magnum-auto-healer,etcd_volume_
↪type=<your volume type>,kube_dashboard_enabled=True,monitoring_enabled=True,
↪ingress_controller=nginx,cloud_provider_tag=v1.19.0,magnum_auto_healer_
↪tag=v1.19.0,container_infra_prefix<docker-registry-without-rate-limit> -f_
↪yaml -c uuid

```

The equivalent Cluster Template configuration through os\_magnum and data in user\_variables.yml

```

magnum_cluster_templates:
- name: <name>
  coe: kubernetes
  external_network_id: <network-id>
  image_id: <image-id>
  master_flavor_id: <master-flavor-id>
  flavor_id: <minion-flavor-id>
  master_lb_enabled: true
  docker_volume_size: 50
  network_driver: calico
  docker_storage_driver: overlay2
  volume_driver: cinder
  labels:
    boot_volume_type: <your volume type>
    boot_volume_size: 50
    kube_tag: v1.18.6
    availability_zone: nova
    helm_client_url: "https://get.helm.sh/helm-v3.4.0-linux-amd64.tar.gz"
    helm_client_sha256:
↪"270acb0f085b72ec28aee894c7443739271758010323d72ced0e92cd2c96ffdb"
    helm_client_tag: v3.4.0
    etcd_volume_size: 50
    auto_scaling_enabled: true
    auto_healing_enabled: true
    auto_healing_controller: magnum-auto-healer
    etcd_volume_type: <your volume type>
    kube_dashboard_enabled: True
    monitoring_enabled: True
    ingress_controller: nginx
    cloud_provider_tag: v1.19.0

```

(continues on next page)

(continued from previous page)

```
magnum_auto_healer_tag: v1.19.0
container_infra_prefix: <docker-registry-without-rate-limit>
```

Note that openstack-ansible deploys the Magnum API service. It is not in scope for openstack-ansible to maintain a guaranteed working cluster template as this will vary depending on the precise version of Magnum deployed and the required version of k8s and its dependencies.

It will be necessary to specify a docker registry (potentially hosting your own mirror or cache) which does not enforce rate limits when deploying Magnum in a production environment.





## POST-DEPLOYMENT DEBUGGING

If the k8s cluster does not create properly, or times out during creation, then the cloud-init logs in the master/minion nodes should be examined, also check the heat-config log and heat-container-agent status.