

---

# **MiCADO Documentation**

**Attila Farkas**

**Sep 30, 2019**



<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	Deployment . . . . .	3
1.2	Dashboard . . . . .	9
1.3	REST API . . . . .	9
1.4	Application Description Template (ADT) . . . . .	10
1.5	Tutorials . . . . .	37
1.6	Release Notes . . . . .	40



This software is developed by the [COLA project](#) and is hosted at the [MiCADO-scale github repository](#). Please, visit the [MiCADO homepage](#) for general information about the product.



MiCADO is an auto-scaling framework for Docker containers, orchestrated by Kubernetes. It supports autoscaling at two levels. At virtual machine (VM) level, a built-in Kubernetes cluster is dynamically extended or reduced by adding/removing cloud virtual machines. At Kubernetes level, the number of replicas tied to a specific Kubernetes Deployment can be increased/decreased.

MiCADO requires a TOSCA-based Application Description Template to be submitted containing three sections:

1. the definition of the individual applications making up a Kubernetes Deployment,
2. the specification of the virtual machine and
3. the implementation of policies for scaling and monitoring both levels of the application.

The format of the Application Description Template for MiCADO is detailed later.

To use MiCADO, first the MiCADO core services must be deployed on a virtual machine (called MiCADO Master) by an Ansible playbook. MiCADO Master is configured as the Kubernetes Master Node and has installed the Docker Engine, Occopus (to scale VMs), Prometheus (for monitoring), Policy Keeper (to perform decision on scaling) and Submitter (to provide submission endpoint) microservices to realize the autoscaling control loops. During operation MiCADO workers (realised on new VMs) are instantiated on demand which deploy Prometheus Node Exporter and CAdvisor as Kubernetes DaemonSets and the Docker engine through contextualisation. The newly instantiated MiCADO workers join the Kubernetes cluster managed by the MiCADO Master.

In the current release, the status of the system can be inspected through the following ways: REST API provides interface for submission, update and list functionalities over applications. Dashboard provides three graphical view to inspect the VMs and Kubernetes Deployments. They are the Kubernetes Dashboard, Grafana and Prometheus. Finally, advanced users may find the logs of the MiCADO core services useful in the Kubernetes Dashboard under the `micado-system` and `micado-worker` namespaces, or directly on the MiCADO master.

## 1.1 Deployment

To deploy MiCADO you need a (separate) virtual machine, called MiCADO master. There are two ways of deployment:

- remote: download the Ansible playbook on your local machine, configure the MiCADO master as target machine and run the playbook to perform the deployment remotely.
- local: login to the MiCADO master, download the Ansible playbook, configure the localhost as target machine and run the playbook to perform the deployment locally.

We recommend to perform the installation remotely as all your configuration files are preserved on your machine, i.e. it is easier to repeat the deployment if needed.

### 1.1.1 Prerequisites

For cloud interfaces supported by MiCADO:

- EC2 (tested on Amazon and OpenNebula)
- Nova (tested on OpenStack)
- CloudSigma
- CloudBroker

For the MiCADO master:

- Ubuntu 16.04, 18.04 (the worker image **must be** the same)
- (Minimum) 2GHz CPU & 3GB RAM & 15GB DISK
- (Recommended) 2GHz CPU & 4GB RAM & 20GB DISK

For the host where the Ansible playbook is executed (differs depending on local or remote):

- Ansible 2.8 or greater
- curl
- jq (to pretty-format API responses)
- wrk (to load test nginx & wordpress demonstrators)

### Ansible

Note: Ansible in the Ubuntu APT repository is outdated and insufficient (at the time of writing this document)

To install Ansible on Ubuntu, use these commands:

```
sudo apt-get update
sudo apt-get install software-properties-common
sudo apt-add-repository ppa:ansible/ansible
sudo apt-get update
sudo apt-get install ansible
```

To install Ansible on other operation systems follow the [official installation guide](#).

### curl

To install curl on Ubuntu, use this command:

```
sudo apt-get install curl
```

To install curl on other operating systems follow the [official installation guide](#).



## jq

To install jq on Ubuntu, use this command:

```
sudo apt-get install jq
```

To install jq on other operating systems follow the [official installation guide](#).

## wrk

To install wrk on Ubuntu, use this command:

```
sudo apt-get install wrk
```

To install wrk on other operating systems check the sidebar on the [github wiki](#).

### 1.1.2 Installation

Perform the following steps either on your local machine or on MiCADO master depending on the installation method.

#### Step 1: Download the ansible playbook.

```
curl --output ansible-micado-0.8.0.tar.gz -L https://github.com/micado-scale/ansible-  
↪micado/releases/download/v0.8.0/ansible-micado-0.8.0.tar.gz  
tar -zxvf ansible-micado-0.8.0.tar.gz  
cd ansible-micado-0.8.0/
```

#### Step 2: Specify cloud credential for instantiating MiCADO workers.

MiCADO master will use this credential against the cloud API to start/stop VM instances (MiCADO workers) to host the application and to realize scaling. Credentials here should belong to the same cloud as where MiCADO master is running. We recommend making a copy of our predefined template and edit it. MiCADO expects the credential in a file, called `credentials-cloud-api.yml` before deployment. Please, do not modify the structure of the template!

```
cp sample-credentials-cloud-api.yml credentials-cloud-api.yml  
edit credentials-cloud-api.yml
```

Edit `credentials-cloud-api.yml` to add cloud credentials. You will find predefined sections in the template for each cloud interface type MiCADO supports. Fill only the section belonging to your target cloud.

Optionally you can use the [Ansible Vault](#) mechanism to keep the credential data in an encrypted format. To achieve this, create the above file using Vault with the command

```
ansible-vault create credentials-cloud-api.yml
```

This will launch the editor defined in the `$EDITOR` environment variable to make changes to the file. If you wish to make any changes to the previously encrypted file, you can use the command

```
ansible-vault edit credentials-cloud-api.yml
```

### Step 3a: Specify security settings and credentials to access MiCADO.

MiCADO master will use these security-related settings and credentials to authenticate its users for accessing the REST API and Dashboard.

```
cp sample-credentials-micado.yml credentials-micado.yml
edit credentials-micado.yml
```

Specify the provisioning method for the x509 keypair used for TLS encryption of the management interface in the `tls` subtree:

- The ‘self-signed’ option generates a new keypair with the specified hostname as subject (or ‘micado-master’ if omitted).
- The ‘user-supplied’ option lets the user add the keypair as plain multiline strings (in unencrypted format) in the `ansible_user_data.yml` file under the ‘cert’ and ‘key’ subkeys respectively.

Specify the default username and password for the administrative user in the `authentication` subtree.

Optionally you may use the Ansible Vault mechanism as described in Step 2 to protect the confidentiality and integrity of this file as well.

### Step 3b: (Optional) Specify credentials to use private Docker registries.

Set the Docker login credentials of your private Docker registry in which your private containers are stored. We recommend making a copy of our predefined template and edit it. MiCADO expects the docker registry credentials in a file, called `credentials-docker-registry.yml`. Please, do not modify the structure of the template!

```
cp sample-credentials-docker-registry.yml credentials-docker-registry.yml
edit credentials-docker-registry.yml
```

Edit `credentials-docker-registry.yml` and add username, password, and registry url. To login to the default `docker_hub`, leave `DOCKER_REPO` as is (<https://index.docker.io/v1/>).

Optionally you may use the Ansible Vault mechanism as described in Step 2 to protect the confidentiality and integrity of this file as well.

### Step 4: Launch an empty cloud VM instance for MiCADO master.

This new VM will host the MiCADO core services.

**a)** Default port number for MiCADO service is 443. Optionally, you can modify the port number stored by the variable called `web_listening_port` defined in the ansible playbook file called `micado-master.yml`.

**b)** Configure a cloud firewall settings which opens the following ports on the MiCADO master virtual machine:

Protocol	Port(s)	Service
TCP	443*	web listening port (configurable*)
TCP	22	SSH
TCP	2379-2380	etcd server
TCP	6443	kube-apiserver
TCP	10250-10252	kubelet, kube-controller, kube-scheduler
UDP	8285 & 8472	flannel overlay network

**NOTE:** [`web_listening_port`] should match with the actual value specified in Step 4a.

**NOTE:** MiCADO master has built-in firewall, therefore you can leave all ports open at cloud level.

c) Finally, launch the virtual machine with the proper settings (capacity, ssh keys, firewall): use any of aws, ec2, nova, etc command-line tools or web interface of your target cloud to launch a new VM. We recommend a VM with 2 cores, 4GB RAM, 20GB disk. Make sure you can ssh to it (password-free i.e. ssh public key is deployed) and your user is able to sudo (to install MiCADO as root). Store its IP address which will be referred as `IP` in the following steps.

### Step 5: Customize the inventory file for the MiCADO master.

We recommend making a copy of our predefined template and edit it. Use the template inventory file, called `sample-hosts.yml` for customisation.

```
cp sample-hosts.yml hosts.yml
edit hosts.yml
```

Edit the `hosts.yml` file to set the variables. The following parameters under the key **micado-target** can be updated:

- **ansible\_host**: specifies the publicly reachable ip address of the target machine where you intend to build/deploy a MiCADO Master or build a MiCADO Worker. Set the public or floating `IP` of the master regardless the deployment method is remote or local. The ip specified here is used by the Dashboard for webpage redirection as well
- **ansible\_connection**: specifies how the target host can be reached. Use “ssh” for remote or “local” for local installation. In case of remote installation, make sure you can authenticate yourself against MiCADO master. We recommend to deploy your public ssh key on MiCADO master before starting the deployment
- **ansible\_user**: specifies the name of your sudoer account, defaults to “ubuntu”
- **ansible\_become**: specifies if account change is needed to become root, defaults to “True”
- **ansible\_become\_method**: specifies which command to use to become superuser, defaults to “sudo”
- **ansible\_python\_interpreter**: specifies the interpreter to be used for ansible on the target host, defaults to “/usr/bin/python3”

Please, revise all the parameters, however in most cases the default values are correct.

### Step 6: Customize the deployment

A few parameters can be fine tuned before deployment. They are as follows:

- **disable\_optimizer**: Setting this parameter to False enables the deployment of the Optimizer module, to perform more advanced scaling. Default is True.
- **disable\_worker\_updates**: Setting this parameter to False enables periodic software updates of the worker nodes. Default is True.
- **grafana\_admin\_pwd**: The string defined here will be the password for Grafana administrator.
- **web\_listening\_port**: Port number of the dashboard on MiCADO master. Default is 443.
- **web\_session\_timeout**: Timeout value in seconds for the Dashboard. Default is 600.

### Step 7: Start the installation of MiCADO master.

Run the following command to build and initialise a MiCADO master node on the empty VM you launched in Step 4 and pointed to in `hosts.yml` Step 5.

```
ansible-playbook -i hosts.yml micado-master.yml
```

If you have used Vault to encrypt your credentials, you have to add the path to your vault credentials to the command line as described in the [Ansible Vault documentation](#) or provide it via command line using the command

```
ansible-playbook -i hosts.yml micado-master.yml --ask-vault-pass
```

Optionally, you can split the deployment of your MiCADO Master in two. The `build` tags prepare the node will all the necessary dependencies, libraries and images necessary for operation. The `start` tags initialise the cluster and all the MiCADO core components.

You can clone the drive of a “**built**” MiCADO Master (or otherwise make an image from it) to be reused again and again. This will greatly speed up the deployment of future instances of MiCADO.

Running the following command will `build` a MiCADO Master node on an empty Ubuntu 16.04 VM.

```
ansible-playbook -i hosts.yml micado-master.yml --tags 'build'
```

You can then run the following command to `start` any “**built**” MiCADO Master node which will initialise and launch the core components for operation.

```
ansible-playbook -i hosts.yml micado-master.yml --tags 'start'
```

As a last measure of increasing efficiency, you can also `build` a MiCADO Worker node. You can then clone/snapshot/image the drive of this VM and point to it in your ADT descriptions. Before running this operation, Make sure the `hosts.yml` points to the empty VM where you intend to build the worker image. Adjust the values under the key `micado-target` as needed. The following command will `build` a MiCADO Worker node on an empty Ubuntu 16.04 VM.

```
ansible-playbook -i hosts.yml build-micado-worker.yml
```

### 1.1.3 After deployment

Once the deployment has successfully finished, you can proceed with

- visiting the *Dashboard*
- using the *REST API*
- playing with the *Tutorials*
- creating your *Application Description Template (ADT)*

### 1.1.4 Check the logs

All logs are now available via the Kubernetes Dashboard on the MiCADO Dashboard. You can navigate to them by changing the `namespace` to `micado-system` or `micado-worker` and then accessing the logs in the **Pods** section. You can also SSH into MiCADO master and check the logs at any point after MiCADO is successfully deployed. All logs are kept under `/var/log/micado` and are organised by components. Scaling decisions, for example, can be inspected under `/var/log/micado/policykeeper`

### 1.1.5 Accessing user-defined service

In case your application contains a container exposing a service, you will have to ensure the following to access it.

- First set **nodePort: xxxxx** (where xxxxx is a port in range 30000-32767) in the **properties: ports:** TOSCA description of your docker container. More information on this in the *Application Description Template (ADT)*

- The container will be accessible at `<IP>:<port>` . Both, the IP and the port values can be extracted from the Kubernetes Dashboard (in case you forget it). The **IP** can be found under *Nodes > my\_micado\_vm > Addresses* menu, while the **port** can be found under *Discovery and load balancing > Services > my\_app > Internal endpoints* menu.

## 1.2 Dashboard

MiCADO has a simple dashboard that collects web-based user interfaces into a single view. To access the Dashboard, visit `https://[IP]:[PORT]`, where

- [IP] is the ip address of MiCADO master, the virtual machine you have launched in Step 4 of *Deployment*
- [PORT] is the port number configured during Step 4 of *Deployment*, its value is held by the `web_listening_port` variable specified in the `micado-master.yml` ansible file.

The following webpages are currently exposed:

- Kubernetes Dashboard: A read-only instance of the Kubernetes WebUI providing a full overview of the infrastructure. Simply *SKIP* the authentication pop-up to gain read-only access to the dashboard.
- Grafana: graphically visualize the resources (nodes, containers) in time. After deploying your application, you can select the service whose metrics you want using the ‘Service’ drop down running above the graphs area.
- Prometheus: monitoring subsystem. Recommended for developers, experts.

## 1.3 REST API

MiCADO has a TOSCA compliant submitter which enables submitting, updating, listing and removing MiCADO applications. MiCADO offers two modes for running applications:

- **Normal mode:** The application is executed by MiCADO as described in the TOSCA ADT (default).
- **Dryrun mode:** The application is not executed by MiCADO, submission is only simulated and tested. The *dryrun* mode is activated by setting the parameter `dryrun=True` when launching a new application.

### 1.3.1 The submitter exposes the following REST API:

- To **launch** an application specified by an Application Description Template (ADT) using a **local file**. You can optionally set an *ID* or run the application in *dryrun* mode:

```
curl --insecure -F file=@<Path_to_ADT> [-F id=<APP_ID>] [-F dryrun=True] -X POST https://<username>:<password>@<IP>:<port>/toscasubmitter/v1.0/app/launch/
```

- To **launch** an application specified by an ADT using a **URL**. You can optionally set an *ID* or run the application in *dryrun* mode:

```
curl --insecure -d input="<URL_to_ADT>" [-d id=<APP_ID>] [-d dryrun=True] -X POST https://<username>:<password>@<IP>:<port>/toscasubmitter/v1.0/app/launch/
```

- To **validate** an application specified by an ADT using a **local file**:

```
curl --insecure -F file=@<Path_to_ADT> -X POST https://<username>:<password>@<IP>:<port>/toscasubmitter/v1.0/app/validate/
```

- To **validate** an application specified by an ADT using a **URL**:

```
curl --insecure -d input="<URL_to_ADT>" -X POST https://<username>:<password>@<IP>:  
↳<port>/toscasubmitter/v1.0/app/validate/
```

- To **update** a running MiCADO application using a **local file**:

```
curl --insecure -F file=@<Path_to_ADT> -X PUT https://<username>:<password>@<IP>:  
↳<port>/toscasubmitter/v1.0/app/update/<APP_ID>
```

- To **update** a running MiCADO application using a **URL**:

```
curl --insecure -d input="<URL_to_ADT>" -X PUT https://<username>:<password>@<IP>:  
↳<port>/toscasubmitter/v1.0/app/update/[APP_ID]
```

- To **undeploy** a running MiCADO application:

```
curl --insecure -X DELETE https://<username>:<password>@<IP>:<port>/toscasubmitter/v1.  
↳0/app/undeploy/[APP_ID]
```

- To **list all** the running MiCADO applications:

```
curl --insecure -X GET https://<username>:<password>@<IP>:<port>/toscasubmitter/v1.0/  
↳list_app/
```

- To **query** a running MiCADO application using the application's ID:

```
curl --insecure -X GET https://<username>:<password>@<IP>:<port>/toscasubmitter/v1.0/  
↳app/[APP_ID]/status
```

- To **query the full execution status** of MiCADO:

```
curl --insecure -X GET https://<username>:<password>@<IP>:<port>/toscasubmitter/v1.0/  
↳info_threads
```

- To **query the services** of a running MiCADO application, use this command:

```
curl --insecure -d query='services' -X GET https://<username>:<password>@<IP>:<port>/  
↳toscasubmitter/v1.0/app/query/[APP_ID]
```

- To **query the nodes** hosting a running MiCADO application:

```
curl --insecure -d query='nodes' -X GET https://<username>:<password>@<IP>:<port>/  
↳toscasubmitter/v1.0/app/query/[APP_ID]
```

## 1.4 Application Description Template (ADT)

### 1.4.1 Overview

MiCADO executes applications described by Application Description Template. The ADT follows the [TOSCA Specification](#) and is described in detail in this section.

#### Main sections of the ADT

##### Top-level definitions

- **tosca\_definitions\_version:** `tosca_simple_yaml_1_1_0`.
- **imports:** List of urls pointing to custom TOSCA types. The default url points to the custom types defined for MiCADO. Please, do not modify this url.
- **repositories:** Docker repositories with their addresses.

### Topology template section

- **node\_templates:** Definitions of the application containers (see **Specification of the Application**) and auxiliary components such as a volume (see **Specification of Volumes**) and virtual machines (see **Specification of the Virtual Machine**)
- **policies:** Scaling & metric policies (see **Specification of Policies**)

### Types section (optional)

This section is used to optionally define additional detailed types which can be referenced in the **topology\_template** section to benefit from abstraction. Under **policy\_types:** for example, complex scaling logic can be defined here, then referenced in the **policies** section above

### Example of the overall structure of an ADT

```
tosca_definitions_version: tosca_simple_yaml_1_1_0

imports:
  - https://raw.githubusercontent.com/micado-scale/tosca/v0.8.0/micado_types.yaml

repositories:
  docker_hub: https://hub.docker.com/
  custom_registry: https://my-registry.mydomain.eu/

topology_template:
  node_templates:
    YOUR-KUBERNETES-APP:
      type: tosca.nodes.MiCADO.Container.Application.Docker
      properties:
        ...
      artifacts:
        ...
      interfaces:
        ...
      requirements:
        ...

    YOUR-VOLUME:
      type: tosca.nodes.MiCADO.Container.Volume
      properties:
        ...
      interfaces:
        ...

    YOUR-VIRTUAL-MACHINE:
      type: tosca.nodes.MiCADO.<CLOUD_API_TYPE>.Compute
      properties:
        ...
      interfaces:
        ...
      capabilities:
```

(continues on next page)

```

...
policies:
- monitoring:
  type: tosca.policies.Monitoring.MiCADO
  properties:
    enable_container_metrics: true
    enable_node_metrics: false
- scalability:
  type: tosca.policies.Scaling.MiCADO
  targets: [ YOUR-VIRTUAL-MACHINE ]
  properties:
    ...
- scalability:
  type: tosca.policies.Scaling.MiCADO
  targets: [ YOUR-KUBERNETES-APP ]
  properties:
    ...
- network:
  type: tosca.policies.Security.MiCADO.Network.HttpProxy
  properties:
    encryption: true
    encryption_key: |
      -----BEGIN PRIVATE KEY-----
    ...
- secret:
  type: tosca.policies.Security.MiCADO.Secret.KubernetesSecretDistribution
  properties:
    ...

```

## 1.4.2 Application

Under the `node_templates` section you can define one or more Docker containers and choose to orchestrate them with Kubernetes (see **YOUR-KUBERNETES-APP**). Each container is described as a separate named node which references a **type** (more on types below). The definition of the most basic container consists of the following:

**NOTE** Kubernetes does not allow for underscores in any resource names (ie TOSCA node names). Names must also begin and end with an alphanumeric.

### Properties

The fields under the **properties** section of the Kubernetes app are a collection of options specific to all iterations of Docker containers. The translator understands both Docker-Compose style naming and Kubernetes style naming, though the Kubernetes style is recommended. You can find additional information about properties in the [translator documentation](#). These properties will be translated into Kubernetes manifests on deployment.

Under the **properties** section of an app (see **YOUR-KUBERNETES-APP**) here are a few common keywords:

- **name**: name for the container (defaults to the TOSCA node name)
- **command**: override the default command line of the container (*list*)
- **args**: override the default entrypoint of container (*list*)
- **env**: list of required environment variables in format:
  - **name**:



- **value:**
- **valueFrom: !! see note below**
- **envFrom: !! see note below**
- **resource:**
  - **requests:**
    - \* **cpu:** CPU reservation, core components usually require 100m so assume 900m as a maximum
- **ports:** list of published ports to the host machine, you can specify these keywords in the style of a flattened (*Service*, *ServiceSpec* and *ServicePort* can all be defined at the same level - see [Kubernetes Service](#))
  - **targetPort:** the port to target (assumes port if not specified)
  - **port:** the port to publish (assumes targetPort if not specified)
  - **name:** the name of this port in the service (generated if not specified)
  - **protocol:** the protocol for the port (defaults to: TCP)
  - **nodePort:** the port (30000-32767) to expose on the host (will create a nodePort Service unless type is explicitly set below)
  - **type:** the type of service for this port (defaults to: ClusterIP unless nodePort is defined above)
  - **clusterIP:** the desired (internal) IP (10.0.0.0/24) for this service (defaults to next available)
  - **metadata:** service metadata, giving the option to set a name for the service. Explicit naming can be used to group different ports together (default grouping is by type)

Environment variables can be loaded in from configuration data in Kubernetes ConfigMaps. This can be accomplished by using **envFrom:** with a list of **configMapRef:** to load all data from a ConfigMap into environment variables as seen [here](#) , or by using **env:** and **valueFrom:** with **configMapKeyRef:** to load specific values into environment variables as seen [here](#) .

Alternatively, ConfigMaps can be mounted as volumes as discussed [here](#) , in the same way other volumes are attached to a container, using the **requirements:** notation below. Also see the examples in **Specification of Configuration Data** below.

## Artifacts

Under the **artifacts** section you can define the docker image for the kubernetes app. Three fields must be defined:

- **type:** `tosca.artifacts.Deployment.Image.Container.Docker`
- **file:** docker image for the kubernetes app (e.g. `sztakilpds/cqueue_frontend:latest` )
- **repository:** name of the repository where the image is located. The name used here (e.g. `docker_hub`), must be defined at the top of the description under the **repositories** section.

## Requirements

Under the **requirements** section you can define the virtual machine you want to host this particular app, restricting the container to run **only** on that VM. If you do not provide a host requirement, the container will run on any possible virtual machine. You can also attach a volume or ConfigMap to this app - the definition of volumes can be found in the next section. Requirements takes a list of map objects:

- **host:** name of your virtual machine as defined under `node_templates`
- **volume:**

- **node:** name of your volume (or ConfigMap) as defined under `node_templates`
- **relationship: !!**
  - \* **type:** `tosca.relationships.AttachesTo`
  - \* **properties:**
    - **location:** path in container
  - **container:** name of a sidecar container defined as a `tosca.nodes.MiCADO.Container.Application.Docker` type under `node_templates`. The sidecar will share the Kubernetes Pod with the main container (the sidecar should not be given an interface)

If a relationship is not defined for a volume the path on container will be the same as the path defined in the volume (see Specification of Volumes). If no path is defined in the volume, the path defaults to `/etc/micado/volumes` for a Volume or `/etc/micado/configs` for a ConfigMap

## Interfaces

Under the **interfaces** section you can define orchestrator specific options, to instruct MiCADO to use Kubernetes, we use the key **Kubernetes**. Fields under **inputs:** will be translated directly to a Kubernetes manifest so it is possible to use the full range of properties which Kubernetes offers as long as field names and syntax follow [the Kubernetes documentation](#) If **inputs:** is omitted a set of defaults will be used to create a Deployment

- **create:** *this key tells MiCADO to create a workload (Deployment/DaemonSet/Job/Pod etc. . . ) for this container*
  - **inputs:** *top-level workload and workload spec options go here. . . two examples, for more see [translator documentation](#)*
    - \* **kind:** overwrite the workload type (defaults to Deployment)
    - \* **spec:**
      - **strategy:**
      - **type:** Recreate (kill pods then update instead of RollingUpdate)
- **configure:** *this key configures the Pod for this workload*
  - **inputs:** `PodTemplateSpec` options follow here. . . For example
    - \* **spec:**
      - **restartPolicy:** change the restart policy (defaults to Always)

## Types

Through abstraction, it is possible to reference a pre-defined parent type and simplify the description of a container. These parent types can hide or reduce the complexity of more complex TOSCA constructs such as **artifacts** and **interfaces** by enforcing defaults or moving them to a simpler construct such as **properties**. Currently MiCADO supports the following types:

- **tosca.nodes.MiCADO.Container.Application.Docker** - The base and most common type for Docker containers in MiCADO. If the desired Docker container image is stored in DockerHub, the property **image:** can be used instead of defining **artifacts:**
- **tosca.nodes.MiCADO.Container.Application.Docker.Deployment** - As above, but orchestrated as a Kubernetes Deployment so that **interfaces:** is not required
- **tosca.nodes.MiCADO.Container.Application.Docker.DaemonSet** - As above, but for a Kubernetes DaemonSet

- **tosca.nodes.MiCADO.Container.Pod.Kubernetes** - Creates an empty Pod. No properties are available, so to use this type a container must be defined and **assigned no interface** as type `tosca.nodes.MiCADO.Container.Application.Docker` and referenced under **requirements:** (more than one container can be referenced to run multiple containers in a single Pod)
- **tosca.nodes.MiCADO.Container.Pod.Kubernetes.Deployment** - As above, but a Kubernetes Deployment

### Examples of the definition of a basic application

#### With *tosca.nodes.MiCADO.Container.Application.Docker* and the Docker image in a custom repository

```
YOUR-KUBERNETES-APP:
  type: toasca.nodes.MiCADO.Container.Application.Docker
  properties:
    name:
    command:
    args:
    env:
    ...
  artifacts:
    image:
      type: toasca.artifacts.Deployment.Image.Container.Docker
      file: YOUR_DOCKER_IMAGE
      repository: custom_registry
  requirements:
  - host: YOUR-VIRTUAL-MACHINE
  interfaces:
    Kubernetes:
      create:
      inputs:
      ...
```

#### With *tosca.nodes.MiCADO.Container.Application.Docker* and the Docker image in DockerHub

```
YOUR-KUBERNETES-APP:
  type: toasca.nodes.MiCADO.Container.Application.Docker
  properties:
    image: YOUR_DOCKER_IMAGE
    name:
    command:
    args:
    env:
    ...
  requirements:
  - host: YOUR-VIRTUAL-MACHINE
  interfaces:
    Kubernetes:
      create:
      inputs:
      ...
```

#### With *tosca.nodes.MiCADO.Container.Application.Docker.Deployment* and the Docker image in DockerHub

```
YOUR-KUBERNETES-APP:
  type: toasca.nodes.MiCADO.Container.Application.Docker.Deployment
  properties:
    image: YOUR_DOCKER_IMAGE
```

(continues on next page)

(continued from previous page)

```
name:
command:
args:
env:
...
requirements:
- host: YOUR-VIRTUAL-MACHINE
```

### Multiple containers in a single Pod, images in DockerHub

```
YOUR-KUBERNETES-APP:
  type: toasca.nodes.MiCADO.Container.Application.Docker
  properties:
    image: YOUR_DOCKER_IMAGE
    name:
    command:
    ...

YOUR-OTHER-KUBERNETES-APP:
  type: toasca.nodes.MiCADO.Container.Application.Docker
  properties:
    image: YOUR_OTHER_DOCKER_IMAGE
    name:
    command:
    ...

YOUR-KUBERNETES-POD:
  type: toasca.nodes.MiCADO.Container.Pod.Kubernetes
  requirements:
  - container: YOUR-KUBERNETES-APP
  - container: YOUR-OTHER-KUBERNETES-APP
```

### Networking in Kubernetes

Kubernetes networking is inherently different to the approach taken by Docker/Swarm. This is a complex subject which is worth a [read here](#) . Since every pod gets its own IP, which any pod can by default use to communicate with any other pod, this means there is no network to explicitly define. If the **ports** keyword is defined in the definition above, pods can reach each other over CoreDNS via their hostname (container name).

Under the **outputs** section (this key is nested within *topology\_template*) you can define an output to retrieve from Kubernetes via the adaptor. Currently, only port info is obtainable.

```
outputs:
  ports:
    value: { get_attribute: [ YOUR-KUBERNETES-APP, port ] }
```

### 1.4.3 Volume

Volumes are defined at the same level as virtual machines and containers, and are then connected to containers using the **requirements:** notation discussed above in the container spec. Some examples of attaching volumes will follow.

## Interfaces

Under the **interfaces** section you should define orchestrator specific options, here we again use the key **Kubernetes**:

- **create**: *this key tells MiCADO to create a persistent volume and claim*
  - **inputs**: persistent volume specific spec options... here are two popular examples, see [Kubernetes volumes](#) for more
    - \* **nfs**:
      - **server**: IP of NFS server
      - **path**: path on NFS share
    - \* **hostPath**:
      - **path**: path on host
- **configure**:
  - **inputs**: using this key, options can be overwritten in the claim

## Types

**NEW in v0.8.0** Through abstraction, it is possible to reference a pre-defined parent type and simplify the description of a volume. These parent types can hide or reduce the complexity of more complex TOSCA constructs such as **interfaces** by enforcing defaults or moving them to a simpler construct such as **properties**. Currently MiCADO supports the following volume types:

- **tosca.nodes.MiCADO.Container.Volume** - The base and most common type for volumes in MiCADO. It is necessary to define further fields under **interfaces**:
- **tosca.nodes.MiCADO.Container.Volume.EmptyDir** - Creates a [EmptyDir](#) persistent volume (PV) and claim (PVC) in Kubernetes
- **tosca.nodes.MiCADO.Container.Volume.HostPath** - Creates a [HostPath](#) PV and PVC. Define the path on host as **path**: under **properties**:
- **tosca.nodes.MiCADO.Container.Volume.NFS** - Creates an [NFS](#) PV and PVC. Define the path and server IP as **path**: and **server**: under **properties**:
- **tosca.nodes.MiCADO.Container.Volume.GlusterFS** - Creates a [GlusterFS](#) PV and PVC. Define path, endpoint and readOnly flag as **path**:, **endpoints**:, and **readOnly**: under **properties**:

## Examples of the definition of a basic volume

With *tosca.nodes.MiCADO.Container.Volume*

```
YOUR-VOLUME:
  type: toasca.nodes.MiCADO.Container.Volume
  interfaces:
    Kubernetes:
      create:
        inputs:
          nfs:
            path: /exports
            server: 10.96.0.1
YOUR-KUBERNETES-APP:
```

(continues on next page)

(continued from previous page)

```

type: tosca.nodes.MiCADO.Container.Application.Docker.Deployment
properties:
  ...
requirements:
- volume:
  node: YOUR-VOLUME
  relationship:
    type: tosca.relationships.AttachesTo
  properties:
    location: /tmp/container/mount/point

```

**Another example with *tosca.nodes.MiCADO.Container.Volume***

Here, no **relationship** is defined under **requirements** so the path defined by the volume */etc/mypath* will be used as the container mount point

```

YOUR-VOLUME:
  type: tosca.nodes.MiCADO.Container.Volume
  interfaces:
    Kubernetes:
      create:
        inputs:
          hostPath:
            path: /etc/mypath

YOUR-KUBERNETES-APP:
  type: tosca.nodes.MiCADO.Container.Application.Docker.Deployment
  properties:
    ...
  requirements:
  - volume: YOUR-VOLUME

```

**With *tosca.nodes.MiCADO.Container.Volume.EmptyDir***

```

YOUR-VOLUME:
  type: tosca.nodes.MiCADO.Container.Volume.EmptyDir

YOUR-KUBERNETES-APP:
  type: tosca.nodes.MiCADO.Container.Application.Docker.Deployment
  properties:
    ...
  requirements:
  - volume:
    node: YOUR-VOLUME
    relationship:
      type: tosca.relationships.AttachesTo
    properties:
      location: /tmp/container/mount/point

```

**With *tosca.nodes.MiCADO.Container.Volume.NFS***

```

YOUR-VOLUME:
  type: tosca.nodes.MiCADO.Container.Volume.NFS
  properties:
    path: /exports
    server: 10.96.0.1

```

(continues on next page)

(continued from previous page)

```
YOUR-KUBERNETES-APP:
  type: tosca.nodes.MiCADO.Container.Application.Docker.Deployment
  properties:
    ...
  requirements:
  - volume:
    node: YOUR-VOLUME
    relationship:
      type: tosca.relationships.AttachesTo
    properties:
      location: /tmp/container/mount/point
```

## 1.4.4 Configuration Data

**NEW in v0.8.0** Configuration data (a Kubernetes **ConfigMap**) are to be defined at the same level as virtual machines, containers and volumes and then loaded into environment variables, or mounted as volumes in the definition of containers as discussed in **Specification of the Application**. Some examples of using configurations will follow at the end of this section.

### Interfaces

Currently MiCADO only supports the definition of configuration data as Kubernetes ConfigMaps. Under the **interfaces** section of this type use the key **Kubernetes:** to instruct MiCADO to create a ConfigMap.

- **create:** *this key tells MiCADO to create a ConfigMap*
  - **inputs:** ConfigMap fields to be overwritten, for more detail see [ConfigMap](#)
    - \* **data:** for UTF-8 byte values
    - \* **binaryData:** for byte values outside of the UTF-8 range

### Types

Through abstraction, it is possible to reference a pre-defined parent type and simplify the description of a ConfigMap. These parent types can hide or reduce the complexity of more complex TOSCA constructs such as **interfaces** by enforcing defaults or moving them to a simpler construct such as **properties**. Currently MiCADO supports the following ConfigMap types:

- **tosca.nodes.MiCADO.Container.Config** - The base and most common type for configuration data in MiCADO. It is necessary to define further fields under **interfaces:** as indicated above
- **tosca.nodes.MiCADO.Container.Config.Kubernetes** - Defaults to a Kubernetes interface and abstracts the inputs to properties. Define the data or binary data fields as **data:** and **binaryData:** under **properties:**

### Examples of the definition of a simple ConfigMap

#### Single ENV var with *tosca.nodes.MiCADO.Container.Config*

Here the environment variable MY\_COLOUR is assigned a value from the ConfigMap

```
YOUR-CONFIG:
  type: tosca.nodes.MiCADO.Container.Config
  interfaces:
    Kubernetes:
      create:
        inputs:
          data:
            color: purple
            how: fairlyNice
            textmode: "true"

YOUR-KUBERNETES-APP:
  type: tosca.nodes.MiCADO.Container.Application.Docker.Deployment
  properties:
    env:
      - name: MY_COLOUR
        valueFrom:
          configMapKeyRef:
            name: YOUR-CONFIG
            key: color
```

### All ENV vars with *tosca.nodes.MiCADO.Container.Config.Kubernetes*

Here an environment variable is created for each key (this becomes the variable name) and value pair in the ConfigMap

```
YOUR-CONFIG:
  type: tosca.nodes.MiCADO.Container.Config.Kubernetes
  properties:
    data:
      color: purple
      how: fairlyNice
      textmode: "true"

YOUR-KUBERNETES-APP:
  type: tosca.nodes.MiCADO.Container.Application.Docker.Deployment
  properties:
    envFrom:
      - configMapRef:
          name: YOUR-CONFIG
```

### A volume with *tosca.nodes.MiCADO.Container.Config.Kubernetes*

Here a volume at /etc/config is populated with three files named after the ConfigMap key names and containing the matching values

```
YOUR-CONFIG:
  type: tosca.nodes.MiCADO.Container.Config.Kubernetes
  properties:
    data:
      color: purple
      how: fairlyNice
      textmode: "true"

YOUR-KUBERNETES-APP:
  type: tosca.nodes.MiCADO.Container.Application.Docker.Deployment
  requirements:
    - volume:
```

(continues on next page)



(continued from previous page)

```
node: YOUR-CONFIG
relationship:
  type: tosca.relationships.AttachesTo
properties:
  location: /etc/config
```

### 1.4.5 Virtual Machine

The collection of docker containers (kubernetes applications) specified in the previous section is orchestrated by Kubernetes. This section introduces how the parameters of the virtual machine can be configured which will host the Kubernetes worker node. During operation MiCADO will instantiate as many virtual machines with the parameters defined here as required during scaling. MiCADO currently supports four different cloud interfaces: CloudSigma, CloudBroker, EC2, Nova. MiCADO supports multiple virtual machine “sets” which can be restricted and host only specific containers (defined in the requirements section of the container specification). At the moment multi-cloud support is in alpha stage, so only certain combinations of different cloud service providers will work.

**NOTE** Underscores are not permitted in virtual machine names (ie TOSCA node names). Names should also begin and end with an alphanumeric.

The following ports and protocols should be enabled on the virtual machine acting as MiCADO worker, replacing [exposed\_application\_ports] with ports you wish to expose on the host:

Protocol	Port(s)	Service
TCP	30000-32767*	exposed application node ports (configurable*)
TCP	22	SSH
TCP	10250	kubelet
UDP	8285 & 8472	flannel overlay network

The following subsections details how to configure them.

#### General

The **capabilities** sections for all virtual machine definitions that follow are identical and are **ENTIRELY OPTIONAL**. They are filled with metadata to support human readability.:

- **num\_cpus** under *host* is a readable string specifying clock speed of the instance type
- **mem\_size** under *host* is a readable string specifying RAM of the instance type
- **type** under *os* is a readable string specifying the operating system type of the image
- **distribution** under *os* is a readable string specifying the OS distro of the image
- **version** under *os* is a readable string specifying the OS version of the image

The **interfaces** section of all virtual machine definitions that follow are **REQUIRED**, and allow you to provide orchestrator specific inputs, in the examples below we use **Occopus**.

- **create**: *this key tells MiCADO to create the VM using Occopus*
  - **inputs**: Specific settings for Occopus follow here
    - \* **interface\_cloud**: tells Occopus which cloud type to interface with
    - \* **endpoint\_cloud**: tells Occopus the endpoint API of the cloud

### CloudSigma

To instantiate MiCADO workers on CloudSigma, please use the template below. MiCADO **requires** `num_cpus`, `mem_size`, `vnc_password`, `libdrive_id`, `public_key_id` and `firewall_policy` to instantiate VM on *CloudSigma*.

```
YOUR-VIRTUAL-MACHINE:
  type: tosca.nodes.MiCADO.CloudSigma.Compute
  properties:
    num_cpus: ADD_NUM_CPUS_FREQ (e.g. 4096)
    mem_size: ADD_MEM_SIZE (e.g. 4294967296)
    vnc_password: ADD_YOUR_PW (e.g. secret)
    libdrive_id: ADD_YOUR_ID_HERE (eg. 87ce928e-e0bc-4cab-9502-514e523783e3)
    public_key_id: ADD_YOUR_ID_HERE (e.g. d7c0f1ee-40df-4029-8d95-ec35b34dae1e)
  nics:
    - firewall_policy: ADD_YOUR_FIREWALL_POLICY_ID_HERE (e.g. fd97e326-83c8-44d8-
↪90f7-0a19110f3c9d)
      ip_v4_conf:
        conf: dhcp
  capabilities:
    # OPTIONAL METADATA
  host:
    properties:
      num_cpus: 2GHz
      mem_size: 2GB
  os:
    properties:
      type: linux
      distribution: ubuntu
      version: 16.04
  interfaces:
    Occopus:
      create:
        inputs:
          interface_cloud: cloudsigma
          endpoint_cloud: ADD_YOUR_ENDPOINT (e.g for cloudsigma https://zrh.
↪cloudsigma.com/api/2.0 )
```

Under the **properties** section of a CloudSigma virtual machine definition these inputs are available.:

- **num\_cpus** is the speed of CPU (e.g. 4096) in terms of MHz of your VM to be instantiated. The CPU frequency required to be between 250 and 100000
- **mem\_size** is the amount of RAM (e.g. 4294967296) in terms of bytes to be allocated for your VM. The memory required to be between 268435456 and 137438953472
- **vnc\_password** set the password for your VNC session (e.g. secret).
- **libdrive\_id** is the image id (e.g. 87ce928e-e0bc-4cab-9502-514e523783e3) on your CloudSigma cloud. Select an image containing a base os installation with cloud-init support!
- **public\_key\_id** specifies the keypairs (e.g. d7c0f1ee-40df-4029-8d95-ec35b34dae1e) to be assigned to your VM.
- **nics[.firewall\_policy && .ip\_v4\_conf.conf]** specifies network policies (you can define multiple security groups in the form of a list for your VM).

### CloudBroker

To instantiate MiCADO workers on CloudBroker, please use the template below. MiCADO **requires** `deployment_id` and `instance_type_id` to instantiate a VM on *CloudBroker*.

```
YOUR-VIRTUAL-MACHINE:
  type: tosca.nodes.MiCADO.CloudBroker.Compute
  properties:
    deployment_id: ADD_YOUR_ID_HERE (e.g. e7491688-599d-4344-95ef-aff79a60890e)
    instance_type_id: ADD_YOUR_ID_HERE (e.g. 9b2028be-9287-4bf6-bbfe-bcbc92f065c0)
    key_pair_id: ADD_YOUR_ID_HERE (e.g. d865f75f-d32b-4444-9fbb-3332bcedeb75)
    opened_port: ADD_YOUR_PORTS_HERE (e.g. '22,2377,7946,8300,8301,8302,8500,8600,
↪9100,9200,4789')
  capabilities:
    # OPTIONAL METADATA
    host:
      properties:
        num_cpus: 2GHz
        mem_size: 2GB
    os:
      properties:
        type: linux
        distribution: ubuntu
        version: 16.04
  interfaces:
    Occopus:
      create:
        inputs:
          interface_cloud: cloudbroker
          endpoint_cloud: ADD_YOUR_ENDPOINT (e.g https://cola-prototype.cloudbroker.
↪com )
```

Under the **properties** section of a CloudBroker virtual machine definition these inputs are available.:

- **deployment\_id** is the id of a preregistered deployment in CloudBroker referring to a cloud, image, region, etc. Make sure the image contains a base OS (preferably Ubuntu) installation with cloud-init support! The id is the UUID of the deployment which can be seen in the address bar of your browser when inspecting the details of the deployment.
- **instance\_type\_id** is the id of a preregistered instance type in CloudBroker referring to the capacity of the virtual machine to be deployed. The id is the UUID of the instance type which can be seen in the address bar of your browser when inspecting the details of the instance type.
- **key\_pair\_id** is the id of a preregistered ssh public key in CloudBroker which will be deployed on the virtual machine. The id is the UUID of the key pair which can be seen in the address bar of your browser when inspecting the details of the key pair.
- **opened\_port** is one or more ports to be opened to the world. This is a string containing numbers separated by a comma.

## EC2

To instantiate MiCADO workers on a cloud through EC2 interface, please use the template below. MiCADO **requires** `region_name`, `image_id` and `instance_type` to instantiate a VM through *EC2*.

```
YOUR-VIRTUAL-MACHINE:
  type: tosca.nodes.MiCADO.EC2.Compute
  properties:
    region_name: ADD_YOUR_REGION_NAME_HERE (e.g. eu-west-1)
    image_id: ADD_YOUR_ID_HERE (e.g. ami-12345678)
    instance_type: ADD_YOUR_INSTANCE_TYPE_HERE (e.g. t1.small)
  capabilities:
```

(continues on next page)

(continued from previous page)

```
# OPTIONAL METADATA
host:
  properties:
    num_cpus: 2GHz
    mem_size: 2GB
os:
  properties:
    type: linux
    distribution: ubuntu
    version: 16.04
interfaces:
  Occopus:
    create:
      inputs:
        interface_cloud: ec2
        endpoint_cloud: ADD_YOUR_ENDPOINT (e.g https://ec2.eu-west-1.amazonaws.com)
```

Under the **properties** section of an EC2 virtual machine definition these inputs are available.:

- **region\_name** is the region name within an EC2 cloud (e.g. eu-west-1).
- **image\_id** is the image id (e.g. ami-12345678) on your EC2 cloud. Select an image containing a base os installation with cloud-init support!
- **instance\_type** is the instance type (e.g. t1.small) of your VM to be instantiated.
- **key\_name** optionally specifies the keypair (e.g. my\_ssh\_keypair) to be deployed on your VM.
- **security\_group\_ids** optionally specify security settings (you can define multiple security groups or just one, but this property must be formatted as a list, e.g. [sg-93d46bf7]) of your VM.
- **subnet\_id** optionally specifies subnet identifier (e.g. subnet-644e1e13) to be attached to the VM.

## Nova

To instantiate MiCADO workers on a cloud through Nova interface, please use the template below. MiCADO **requires** `image_id` `flavor_name`, `project_id` and `network_id` to instantiate a VM through *Nova*.

```
YOUR-VIRTUAL-MACHINE:
  type: toasca.nodes.MiCADO.Nova.Compute
  properties:
    image_id: ADD_YOUR_ID_HERE (e.g. d4f4e496-031a-4f49-b034-f8dafe28e01c)
    flavor_name: ADD_YOUR_ID_HERE (e.g. 3)
    project_id: ADD_YOUR_ID_HERE (e.g. a678d20e71cb4b9f812a31e5f3eb63b0)
    network_id: ADD_YOUR_ID_HERE (e.g. 3fd4c62d-5fbe-4bd9-9a9f-c161dabeefde)
    key_name: ADD_YOUR_KEY_HERE (e.g. keyname)
    security_groups:
      - ADD_YOUR_ID_HERE (e.g. d509348f-21f1-4723-9475-0cf749e05c33)
  capabilities:
    # OPTIONAL METADATA
    host:
      properties:
        num_cpus: 2GHz
        mem_size: 2GB
    os:
      properties:
        type: linux
```

(continues on next page)

(continued from previous page)

```

distribution: ubuntu
version: 16.04
interfaces:
  Occopus:
    create:
      inputs:
        interface_cloud: nova
        endpoint_cloud: ADD_YOUR_ENDPOINT (e.g https://sztaki.cloud.mta.hu:5000/v3)

```

Under the **properties** section of a Nova virtual machine definition these inputs are available.:

- **project\_id** is the id of project you would like to use on your target Nova cloud.
- **image\_id** is the image id on your Nova cloud. Select an image containing a base os installation with cloud-init support!
- **flavor\_name** is the name of flavor to be instantiated on your Nova cloud.
- **server\_name** optionally defines the hostname of VM (e.g.: "helloworld").
- **key\_name** optionally sets the name of the keypair to be associated to the instance. Keypair name must be defined on the target nova cloud before launching the VM.
- **security\_groups** optionally specify security settings (you can define multiple security groups in the form of a list) for your VM.
- **network\_id** is the id of the network you would like to use on your target Nova cloud.

## Types

**NEW in v0.8.0** Through abstraction, it is possible to reference a pre-defined type and simplify the description of a virtual machine. Currently MiCADO supports these additional types for CloudSigma, but more can be written:

- **tosca.nodes.MiCADO.CloudSigma.Compute.Occo** - Automatically orchestrates on Zurich with Occopus. There is no need to define further fields under **interfaces:** but Zurich can be changed by overwriting **endpoint** under **properties:**
- **tosca.nodes.MiCADO.CloudSigma.Compute.Occo.small** - As above but creates a 2GHz/2GB node by default
- **tosca.nodes.MiCADO.CloudSigma.Compute.Occo.big** - As above but creates a 4GHz/4GB node by default
- **tosca.nodes.MiCADO.CloudSigma.Compute.Occo.small.NFS** - As *small* above but installs NFS dependencies by default

## Example definition of a VM using abstraction

**With *tosca.nodes.MiCADO.CloudSigma.Compute.Occo.small* and omitting capabilities metadata**

```

YOUR-VIRTUAL-MACHINE:
  type: toska.nodes.MiCADO.CloudSigma.Compute.Occo.small
  properties:
    vnc_password: ADD_YOUR_PW (e.g. secret)
    libdrive_id: ADD_YOUR_ID_HERE (eg. 87ce928e-e0bc-4cab-9502-514e523783e3)
    public_key_id: ADD_YOUR_ID_HERE (e.g. d7c0f1ee-40df-4029-8d95-ec35b34dae1e)
    nics:
      - firewall_policy: ADD_YOUR_FIREWALL_POLICY_ID_HERE (e.g. fd97e326-83c8-44d8-
↪90f7-0a19110f3c9d)

```

(continues on next page)

```
ip_v4_conf:
  conf: dhcp
```

## 1.4.6 Monitoring Policy

**NEW in v0.8.0** Metric collection is now disabled by default. The basic exporters from previous MiCADO versions can be enabled through the monitoring policy below. If the policy is omitted, or if one property is left undefined, then the relevant metric collection will be disabled.

```
policies:
- monitoring:
  type: tosca.policies.Monitoring.MiCADO
  properties:
    enable_container_metrics: true
    enable_node_metrics: true
```

## 1.4.7 Scaling Policy

### Basic scaling

To utilize the autoscaling functionality of MiCADO, scaling policies can be defined on virtual machine and on the application level. Scaling policies can be listed under the **policies** section. Each **scalability** subsection must have the **type** set to the value of `tosca.policies.Scaling.MiCADO` and must be linked to a node defined under **node\_template**. The link can be implemented by specifying the name of the node under the **targets** subsection. You can attach different policies to different containers or virtual machines, though a new policy should exist for each. The details of the scaling policy can be defined under the **properties** subsection. The structure of the **policies** section can be seen below.

```
topology_template:
  node_templates:
    YOUR-VIRTUAL-MACHINE:
      type: tosca.nodes.MiCADO.<CLOUD_API_TYPE>.Compute
      ...
    YOUR-OTHER-VIRTUAL-MACHINE:
      type: tosca.nodes.MiCADO.<CLOUD_API_TYPE>.Compute
      ...
    YOUR-KUBERNETES-APP:
      type: tosca.nodes.MiCADO.Container.Application.Docker
      ...
    YOUR-OTHER-KUBERNETES-APP:
      type: tosca.nodes.MiCADO.Container.Application.Docker
      ...

  policies:
  - scalability:
    type: tosca.policies.Scaling.MiCADO
    targets: [ YOUR-VIRTUAL-MACHINE ]
    properties:
      ...
  - scalability:
    type: tosca.policies.Scaling.MiCADO
    targets: [ YOUR-OTHER-VIRTUAL-MACHINE ]
```

(continues on next page)

(continued from previous page)

```

properties:
  ...
- scalability:
  type: toasca.policies.Scaling.MiCADO
  targets: [ YOUR-KUBERNETES-APP ]
  properties:
    ...
- scalability:
  type: toasca.policies.Scaling.MiCADO
  targets: [ YOUR-OTHER-KUBERNETES-APP ]
  properties:
    ...

```

The scaling policies are evaluated periodically. In every turn, the virtual machine level scaling policies are evaluated, followed by the evaluation of each scaling policies belonging to kubernetes-deployed applications.

The **properties** subsection defines the scaling policy itself. For monitoring purposes, MiCADO integrates the Prometheus monitoring tool with two built-in exporters on each worker node: Node exporter (to collect data on nodes) and CAdvisor (to collect data on containers). Based on Prometheus, any monitored information can be extracted using the Prometheus query language and the returned value can be associated to a user-defined variable. Once variables are updated, scaling rule is evaluated. Scaling rule is specified by (a short) Python code. The code can refer to/use the variables. The structure of the scaling policy can be seen below.

```

- scalability:
  ...
  properties:
    sources:
      - 'myprometheus.exporter.ip.address:portnumber'
    constants:
      LOWER_THRESHOLD: 50
      UPPER_THRESHOLD: 90
      MYCONST: 'any string'
    queries:
      THELOAD: 'Prometheus query expression returning a number'
      MYLISTOFSTRING: ['Prometheus query returning a list of strings as tags',
↪ 'tagname as filter']
      MYEXPR: 'something refering to {{MYCONST}}'
    alerts:
      - alert: myalert
        expr: 'Prometheus expression for an event important for scaling'
        for: 1m
    min_instances: 1
    max_instances: 5
    scaling_rule: |
      if myalert:
        m_node_count=5
      if THELOAD>UPPER_THRESHOLD:
        m_node_count+=1
      if THELOAD<LOWER_THRESHOLD:
        m_node_count-=1

```

The subsections have the following roles:

- **sources** supports the dynamic attachment of an external exporter by specifying a list endpoints of exporters (see example above). Each item found under this subsection is configured under Prometheus to start collecting the information provided/exported by the exporters. Once done, the values of the parameters provided by the exporters become available. MiCADO supports Kubernetes service discovery to define such a source, simply

pass the name of the app as defined in TOSCA and do not specify any port number

- **constants** subsection is used to predefined fixed parameters. Values associated to the parameters can be referred by the scaling rule as variable (see `LOWER_THRESHOLD` above) or in any other sections referred as Jinja2 variable (see `MYEXPR` above).
- **queries** contains the list of Prometheus query expressions to be executed and their variable name associated (see `THELOAD` or `MYLISTOFSTRING` above)
- **alerts** subsection enables the utilization of the alerting system of Prometheus. Each alert defined here is registered under Prometheus and fired alerts are represented with a variable of their name set to True during the evaluation of the scaling rule (see `myalert` above).
- **min\_instances** keyword specifies the lowest number of instances valid for the node.
- **max\_instances** keyword specifies the highest number of instances valid for the node.
- **scaling\_rule** specifies Python code to be evaluated periodically to decide on the number of instances. The Python expression must be formalized with the following conditions:
  - Each constant defined under the ‘constants’ section can be referred; its value is the one defined by the user.
  - Each variable defined under the ‘queries’ section can be referred; its value is the result returned by Prometheus in response to the query string.
  - Each alert name defined under the ‘alerts’ section can be referred, its value is a logical True in case the alert is firing, False otherwise
  - Expression must follow the syntax of the Python language
  - Expression can be multiline
  - The following predefined variables can be referred; their values are defined and updated before the evaluation of the scaling rule
    - \* `m_nodes`: python list of nodes belonging to the kubernetes cluster
    - \* `m_node_count`: the target number of nodes
    - \* `m_nodes_todrop`: the ids or ip addresses of the nodes to be dropped in case of downscaling
    - \* `m_container_count`: the target number of containers for the service the evaluation belongs to
    - \* `m_time_since_node_count_changed`: time in seconds elapsed since the number of nodes changed
  - In a scaling rule belonging to the virtual machine, the name of the variable to be updated is `m_node_count`; as an effect the number stored in this variable will be set as target instance number for the virtual machines.
  - In a scaling rule belonging to the virtual machine, the name of the variable to be updated is `m_nodes_todrop`; the variable must be filled with list of ids or ip addresses and as an effect the valid nodes will be dropped. The variable `m_node_count` should not be modified in case of node dropping, MiCADO will update it automatically.
  - In a scaling rule belonging to a kubernetes deployment, the name of the variable to be set is `m_container_count`; as an effect the number stored in this variable will be set as target instance number for the kubernetes service.

For debugging purposes, the following support is provided:

- `m_dryrun` can be specified in the **constant** as list of components towards which the communication is disabled. It has the following syntax: `m_dryrun: ["prometheus","occpus","k8s","optimizer"]` Use this feature with caution!



- the standard output of the python code defined by the user under the scaling rule section is collected in a separate log file stored under the policy keeper log directory. It can also be used for debugging purposes.

For further examples, inspect the scaling policies of the demo examples detailed in the next section.

## Optimiser-based scaling

For implementing more advanced scaling policies, it is possible to utilize the built-in Optimiser in MiCADO. The role of the Optimiser is to support decision making in calculating the number of worker nodes (virtual machines) i.e. to scale the nodes to the optimal level. Optimiser is implemented using machine learning algorithm aiming to learn the relation between various metrics and the effect of scaling events. Based on this learning, the Optimiser is able to calculate and advise on the necessary number of virtual machines.

### Current limitations

- only web based applications are supported
- only one of the node sets can be supported
- no container scaling is supported

### Optimiser can be utilised based on the following principles

- User specifies a so-called target metric with its associated minimum and maximum thresholds. The target metric is a monitored Prometheus expression for which the value is tried to be kept between the two thresholds by the Optimiser with scaling advices.
- User specifies several so-called input metrics which represent the state of the system correlating to the target variable
- User specifies several initial settings (see later) for the Optimiser
- User submits the application activating the Optimiser through the ADT
- Optimiser starts with the ‘training’ phase during which the correlations are learned. During the training phase artificial load must be generated for the web application and scaling activities must be performed (including extreme values) in order to present all possible situations for the Optimiser. During the phase, Optimiser continuously monitors the input/target metrics and learns the correlations.
- When correlations are learnt, Optimiser turns to ‘production’ phase during which advice can be requested from the Optimiser. During this phase, Optimiser returns advice on request, where the advice contains the number of virtual machines (nodes) to be scaled to. During the production phase, the Optimiser continues its learning activity to adapt to the new situations.

**Activation of the Optimiser** Optimiser must be enabled at deployment time. By default it is disabled. Once it is enabled and deployed, it can be driven through the scaling policy in subsections “constants” and “queries”. Each parameter relating to the Optimiser must start with the “m\_opt\_” string. In case no variable name with this prefix is found in any sections, Optimiser is not activated.

**Initial settings for the Optimiser** Parameters for initial settings are defined under the “constants” section and their name must start with the “m\_opt\_init\_” prefix. These parameters are as follows:

- **m\_opt\_init\_knowledge\_base** is a parameter which specifies the way how the knowledge base must be built under the Optimiser. When defined as “build\_new”, Optimiser empties its knowledge base and starts building a new knowledge i.e. starts learning the correlations. When using the “use\_existing” value, the knowledge is kept and continued building further. Default is “use\_existing”.
- **m\_opt\_init\_training\_samples\_required** defines how many sample of the metrics must be collected by the Optimiser before start learning the correlations. Default is 300.
- **m\_opt\_init\_max\_upscale\_delta** specifies the maximum change in number of node for an upscaling advice. Default is 6.

- **m\_opt\_init\_max\_downscale\_delta** specifies the maximum change in number of node for a downscaling advice. Default is 6.
- **m\_opt\_init\_advice\_freeze\_interval** specifies how many seconds must elapse before the Optimiser advises a different number of node. Can be used to mitigate the frequency of scaling. Defaults to 0.

**Definition of input metrics for the Optimizer** Input metrics must be specified for the Optimiser under the “queries” subsection to perform the training i.e. learning the correlations. Each parameter must start with the “m\_opt\_input\_” prefix, e.g. m\_opt\_input\_CPU. The following two pieces of variable must be specified for the web application:

- **m\_opt\_input\_AVG\_RR** should specify the average request rate of the web server.
- **m\_opt\_input\_SUM\_RR** should specify the summary of request rate of the web server.

**Definition of the target metric for the Optimizer** Target metric is a continuously monitored parameter that must be kept between thresholds. To specify it, together with the thresholds, “m\_opt\_target\_” prefix must be used. These three parameter must be defined under the “queries” sections. They are as follows:

- **m\_opt\_target\_query\_MYTARGET** specifies the prometheus query for the target metric called MYTARGET.
- **m\_opt\_target\_minth\_MYTARGET** specifies the value above which the target metric must be kept.
- **m\_opt\_target\_maxth\_MYTARGET** specifies the value below which the target metric must be kept.

**Requesting scaling advice from the Optimizer** In order to receive a scaling advice from the Optimiser, the method **m\_opt\_advice()** must be invoked in the scaling\_rule section of the node.

**IMPORTANT! Minimum and maximum one node must contain this method invocation in its scaling\_rule section for proper operation!**

The **m\_opt\_advice()** method returns a python dictionary containing the following fields:

- **valid** stores True/False value indicating whether the advise can be considered or not.
- **phase** indicates whether the Optimiser is in “training” or “production” phase.
- **vm\_number** represents the advise for the target number of nodes to scale to.
- **reliability** represents the goodness of the advice with a number between 0 and 100. The bigger the number is the better/more reliable the advice is.
- **error\_msg** contains the error occured in the Optimiser. Filled when valid is False.

## 1.4.8 Network policy

There are six types of MiCADO network security policy.

- `tosca.policies.Security.MiCADO.Network.Passthrough`: Pass through network policy. Specifies no additional filtering, no application-level firewall on the nodes.
- `tosca.policies.Security.MiCADO.Network.L7Proxy`: Apply application-level firewall; can provide TLS control. No protocol enforcement.

```
properties:
  encryption:
    type: boolean
    description: Specifies if encryption should be used
    required: true
  encryption_key:
    type: string
```

(continues on next page)

(continued from previous page)

```

description: The key file for TLS encryption as unencrypted .PEM
required: false
encryption_cert:
  type: string
  description: The cert file for TLS encryption as .PEM
  required: false
encryption_offload:
  type: string
  description: Controls whether connection should be re-encrypted server side
  required: false
encryption_cipher:
  type: string
  description: Specifies allowed ciphers client side during TLS handshake
  required: false

```

- `tosca.policies.Security.MiCADO.Network.SmtpProxy`: Enforce SMTP protocol, can provide TLS control.

```

properties:
  relay_check:
    type: boolean
    description: Toggle relay checking
    required: true
  permit_percent_hack:
    type: boolean
    description: Allow the % symbol in the local part of an email address
    required: false
  error_soft:
    type: boolean
    description: Return a soft error when recipient filter does not match
    required: false
  relay_domains:
    type: list
    description: Domain mails are accepted for use postfix style lists
    required: false
  permit_exclamation_mark:
    type: boolean
    description: Allow the ! symbol in the local part of an email address
    required: false
  relay_domains_matcher_whitelist:
    type: list
    description: Domains mails accepted based on list of regex (precedence)
    required: false
  relay_domains_matcher_blacklist:
    type: list
    description: Domain mails rejected based on list of regular expressions
    required: false
  sender_matcher_whitelist:
    type: list
    description: Sender addresses accepted based on list of regex (precedence)
    required: false
  sender_matcher_blacklist:
    type: list
    description: Sender addresses rejected based on list of regex
    required: false
  recipient_matcher_whitelist:
    type: list
    description: Recipient addresses accepted based on list of regex (precedence)

```

(continues on next page)

```
    required: false
recipient_matcher_blacklist:
  type: list
  description: Recipient addresses rejected based on list of regex
  required: false
autodetect_domain_from:
  type: string
  description: Let Zorp autodetect firewall domain name and write to received line
  constraints:
    - valid_values: ["mailname", "fqdn"]
  required: false
append_domain:
  type: string
  description: Domain to append to email addresses which do not specify a domain
  required: false
permit_omission_of_angle_brackets:
  type: boolean
  description: Permit MAIL From and RCPT To params without normally required_
↳brackets
  required: false
interval_transfer_noop:
  type: integer
  description: Interval between two NOOP commands sent to server while waiting for_
↳stack proxy results
  required: false
resolve_host:
  type: boolean
  description: Resolve client host from IP address and write to received line
  required: false
permit_long_responses:
  type: boolean
  description: Permit overly long responses as some MTAs include variable parts in_
↳responses
  required: false
max_auth_request_length:
  type: integer
  description: Maximum allowed length of a request during SASL style authentication
  required: false
max_response_length:
  type: integer
  description: Maximum allowed line length of server response
  required: false
unconnected_response_code:
  type: integer
  description: Error code sent to client if connecting to server fails
  required: false
add_received_header:
  type: boolean
  description: Add a received header into the email messages transferred by proxy
  required: false
domain_name:
  type: string
  description: Fix a domain name into added receive line. add_received_header must_
↳be true
  required: false
tls_passthrough:
  type: boolean
```

(continues on next page)

(continued from previous page)

```

description: Change to passthrough mode
required: false
extensions:
  type: list
  description: Allowed ESMTP extensions, indexed by extension verb
  required: false
require_crlf:
  type: boolean
  description: Specify whether proxy should enforce valid CRLF line terminations
  required: false
timeout:
  type: integer
  description: Timeout in ms - if no packet arrives, connection is dropped
  required: false
max_request_length:
  type: integer
  description: Maximum allowed line length of client requests
  required: false
permit_unknown_command:
  type: boolean
  description: Enable unknown commands
  required: false

```

- `tosca.policies.Security.MiCADO.Network.HttpProxy`: Enforce HTTP protocol, can provide TLS control.

```

properties:
  max_keepalive_requests:
    type: integer
    description: Max number of requests allowed in a single session
    required: false
  permit_proxy_requests:
    type: boolean
    description: Allow proxy type requests in transparent mode
    required: false
  reset_on_close:
    type: boolean
    description: If connection is terminated without a proxy generated error, send an ↔RST instead of a normal close
    required: false
  permit_unicode_url:
    type: boolean
    description: Allow unicode characters in URLs encoded as u'
    required: false
  permit_server_requests:
    type: boolean
    description: Allow server type requests in non transparent mode
    required: false
  max_hostname_length:
    type: integer
    description: Maximum allowed length of hostname field in URLs
    required: false
  parent_proxy:
    type: string
    description: Address or hostname of parent proxy to be connected
    required: false
  permit_ftp_over_http:
    type: boolean

```

(continues on next page)

(continued from previous page)

```
description: Allow processing FTP URLs in non transparent mode
required: false
parent_proxy_port:
  type: integer
  description: Port of parent proxy to be connected
  required: false
permit_http09_responses:
  type: boolean
  description: Allow server responses to use limited HTTP 0 9 protocol
  required: false
rewrite_host_header:
  type: boolean
  description: Rewrite host header in requests when URL redirection occurs
  required: false
max_line_length:
  type: integer
  description: Maximum allowed length of lines in requests and responses
  required: false
max_chunk_length:
  type: integer
  description: Maximum allowed length of a single chunk when using chunked transfer
  required: false
strict_header_checking_action:
  type: string
  description: Specify Zorp action if non rfc or unknown header in communication
  constraints:
    - valid_values: ["accept", "drop", "abort"]
  required: false
non_transparent_ports:
  type: list
  description: List of ports that non transparent requests may use
  required: false
strict_header_checking:
  type: boolean
  description: Require RFC conformant HTTP headers
  required: false
max_auth_time:
  type: integer
  description: Force new auth request from client browser after time in seconds
  required: false
max_url_length:
  type: integer
  description: Maximum allowed length of URL in a request
  required: false
timeout_request:
  type: integer
  description: Time to wait for a request to arrive from client
  required: false
rerequest_attempts:
  type: integer
  description: Control number of attempts proxy takes to send request to server
  required: false
error_status:
  type: integer
  description: On error, Zorp uses this as status code of HTTP response
  required: false
```

(continues on next page)

(continued from previous page)

```

keep_persistent:
  type: boolean
  description: Try to keep connection to client persistent, even if unsupported
  required: false
error_files_directory:
  type: string
  description: Location of HTTP error messages
  required: false
max_header_lines:
  type: integer
  description: Maximum number of eader lines allowed in requests and responses
  required: false
use_canonicalized_urls:
  type: boolean
  description: Enable canonicalization - converts URLs to canonical form
  required: false
max_body_length:
  type: integer
  description: Maximum allowed length of HTTP request or response body
  required: false
require_host_header:
  type: boolean
  description: Require presence of host header
  required: false
buffer_size:
  type: integer
  description: Size of I O buffer used to transfer entity bodies
  required: false
permitted_responses:
  type: list
  description: Normative policy hash for HTTP responses indexed by HTTP method and
↳ response code
  entry_schema:
    description: dictionary (string/int)
    type: map
    required: false
transparent_mode:
  type: boolean
  description: Enable transparent mode for the proxy
  required: false
permit_null_response:
  type: boolean
  description: Permit RFC incompilant responses with headers not terminated by CRLF,
↳ and not containing entity body
  required: false
language:
  type: string
  description: Specify language of HTTP error pages displayed to client
  required: false
  default: English
error_silent:
  type: boolean
  description: Turns off verbose error reporting to HTTP client, making firewall_
↳ fingerprinting more difficult
  required: false
permitted_requests:
  type: list

```

(continues on next page)

(continued from previous page)

```

description: List of permitted HTTP methods indexed by verb
required: false
use_default_port_in_transparent_mode:
  type: boolean
  description: Enable use of default port in transparent mode
  required: false
timeout_response:
  type: integer
  description: Time to wait for the HTTP status line to arrive from the server
  required: false
permit_invalid_hex_escape:
  type: boolean
  description: Allow invalid hexadecimal escaping in URLs
  required: false
auth_cache_time:
  type: integer
  description: Caching authentication information time in seconds
  required: false
timeout:
  type: integer
  description: General I O timeout in ms
  required: false
default_port:
  type: integer
  description: Used in non transparent mode when URL does not contain a port number
  required: false
  default: 80

```

- `tosca.policies.Security.MiCADO.Network.HttpURIFilterProxy`: Enforce HTTP protocol with regex URL filtering capabilities

```

properties:
  matcher_whitelist:
    type: list
    description: List of regex determining permitted access to a URL (precedence)
    required: true
  matcher_blacklist:
    type: list
    description: List of regex determining prohibited access to a URL
    required: true

```

- `tosca.policies.Security.MiCADO.Network.HttpWebdavProxy`: Enforce HTTP protocol with request methods for WebDAV.

This proxy has no additional properties.

### 1.4.9 Secret policy

There is a way to define application-level secrets in the MiCADO application description. These secrets are distributed by kubernetes.

```

tosca.policies.Security.MiCADO.Secret.KubernetesSecretDistribution:
  derived_from: toasca.policies.Root
  description: distributes secrets to services
  properties:

```

(continues on next page)



(continued from previous page)

```
file_secrets:
  type: map
text_secrets:
  type: map
```

## 1.5 Tutorials

You can find some demo applications under the subdirectories of the ‘testing’ directory in the downloaded (and unzipped) installation package of MiCADO.

### 1.5.1 stressng

This application contains a single service, performing a constant CPU load. The policy defined for this application scales up/down both nodes and the stressng service based on cpu consumption. Helper scripts have been added to the directory to ease application handling.

**Note:** make sure you have the `jq` tool installed required by the helper scripts.

- Step1: make a copy of the TOSCA file which is appropriate for your cloud - `stressng_<your_cloud>.yaml` - and name it `stressng.yaml` (ie. by issuing the command `cp stressng_cloudsigma.yaml stressng.yaml`)
- Step2: fill in the requested fields beginning with `ADD_YOUR_...`. These will differ depending on which cloud you are using.

**Important:** Make sure you create the appropriate firewall policy for the MiCADO workers as described [here!](#)

- In CloudSigma, for example, the `libdrive_id`, `public_key_id` and `firewall_policy` fields must be completed. Without these, CloudSigma does not have enough information to launch your worker nodes. All information is found on the CloudSigma Web UI. `libdrive_id` is the long alphanumeric string in the URL when a drive is selected under “Storage/Library”. `public_key_id` is under the “Access & Security/Keys Management” menu as **Uuid**. `firewall_policy` can be found when selecting a rule defined under the “Networking/Policies” menu. The following ports must be opened for MiCADO workers: *all inbound connections from MiCADO master*
- Step3: Update the parameter file, called `_settings`. You need the ip address for the MiCADO master and should name the application by setting the `APP_ID` **\*the application ID can not contain any underscores ( \_ )** You should also change the SSL user/password/port information if they are different from the default.
- Step4: run `1-submit-tosca-stressng.sh` to create the minimum number of MiCADO worker nodes and to deploy the Kubernetes Deployment including the stressng app defined in the `stressng.yaml` TOSCA description.
- Step4a: run `2-list-apps.sh` to see currently running applications and their IDs
- Step5: run `3-stress-cpu-stressng.sh 85` to stress the service and increase the CPU load. After a few minutes, you will see the system respond by scaling up virtual machines and containers to the maximum specified.
- Step6: run `3-stress-cpu-stressng.sh 10` to update the service and decrease the CPU load. After a few moments the system should respond by scaling down virtual machines and containers to the minimum specified.
- Step7: run `4-undeploy-stressng.sh` to remove the stressng stack and all the MiCADO worker nodes

## 1.5.2 cqueue

This application demonstrates a deadline policy with CQueue. CQueue provides a lightweight queuing service for executing containers. The entire infrastructure will be deployed by a single ADT as a microservices architecture. CQueue server (implemented by RabbitMQ, Redis and a web-based frontend) will be run on a static VM in the cluster. The server stores items in a queue where each item represents a container execution. CQueue workers will run on a separate set of scalable VMs, and are responsible for fetching items and performing the execution of the container locally. The demonstration below shows that the items can be consumed before a deadline using MiCADO for scaling the CQueue worker and its VM nodes.

If you prefer to launch your own cQueue server externally, use the docker-compose file `docker-compose-cqueue-server.yaml` and edit the relevant shell scripts to point to your server and to launch `micado-cqworker.yaml` instead.

**Note:** make sure you have the `jq` tool installed required by the helper scripts.

- Step1: Update the file `cq-microservice.yaml` with the CloudSigma ID details necessary to launch your **two** VM sets
  - Update each ‘ADD\_YOUR\_ID\_HERE’ string with the proper value retrieved under your CloudSigma account.
  - Make sure port 30888 is open on the `cq-server` virtual machine set
- Step2: Update the parameter file, called `_settings`. You need the ip address for the MiCADO master and, once your worker nodes are running, you should enter the IP for the CQueue server which is about to be deployed. Setting the IP of the CQueue server is a required step if your MiCADO Master does not have the appropriate port open.
- Step3: Run `./1-deploy-cq-microservices.sh` to deploy the cQueue server and worker components to separate virtual machine nodes
- Step4: Use your Cloud WebUI and find the public IP of the VM hosting the cQueue server (in fact, this can be **any** VM in your cluster with port 30888 open)
- Step5: Run `./3-get_date_in_epoch_plus_seconds.sh 600` to calculate the unix timestamp representing the deadline by which the items (containers) must be finished. Take the value from the last line of the output produced by the script. The value is 600 seconds from now.
- Step6: Run `./4-update-cqueue-deadline.sh xxxxxxxx` where **xxxxxxx** is the unix timestamp taken from the previous step.
- Step7: Run `./5-submit-jobs.sh 50` to generate and send 50 jobs to CQueue server. Each item will be a simple Hello World app (combined with some sleep) realized in a container. You can later override this with your own container.
- Step8a: You can run `./2-list-running-apps.sh` to list the apps running under MiCADO.
- Step8b: You can run `query-services.sh` to see the details of docker services of your application
- Step8c: You can run `query-nodes.sh` to see the details of docker nodes hosting your application
- Step9: Run `./6-undeploy-cq-microservices.sh` to remove your application from MiCADO when all items are consumed.
- Step10: You can have a look at the state `./cqueue-get-job-status.sh <task_id>` or stdout of container executions `./cqueue-get-job-status.sh <task_id>` using one of the task id values printed during Step 3.

### 1.5.3 nginx

This application deploys a http server with nginx. The container features a built-in prometheus exporter for HTTP request metrics. The policy defined for this application scales up/down both nodes and the nginx service based on active http connections. wrk (`apt-get install wrk` | <https://github.com/wg/wrk>) is recommended for HTTP load testing.

**Note:** make sure you have the jq tool and wrk benchmarking app installed as these are required by the helper scripts. Best results for wrk are seen on multi-core systems.

- Step1: make a copy of the TOSCA file which is appropriate for your cloud - `nginx_<your_cloud>.yaml` - and name it `nginx.yaml`
- Step2: fill in the requested fields beginning with `ADD_YOUR_...`. These will differ depending on which cloud you are using.

**Important:** Make sure you create the appropriate firewall policy for the MiCADO workers as described [here!](#)

- In CloudSigma, for example, the `libdrive_id`, `public_key_id` and `firewall_policy` (port 30012 must be open) fields must be completed. Without these, CloudSigma does not have enough information to launch your worker nodes. All information is found on the CloudSigma Web UI. `libdrive_id` is the long alphanumeric string in the URL when a drive is selected under “Storage/Library”. `public_key_id` is under the “Access & Security/Keys Management” menu as **Uaid**. `firewall_policy` can be found when selecting a rule defined under the “Networking/Policies” menu. The following ports must be opened for MiCADO workers: *all inbound connections from MiCADO master*
- Step3: Update the parameter file, called `_settings`. You need the ip address for the MiCADO master and should name the deployment by setting the `APP_ID`. **\*the application ID can not contain any underscores ( \_ )** The `APP_NAME` must match the name given to the application in TOSCA (default: **nginxapp**) You should also change the SSL user/password/port information if they are different from the default.
- Step4: run `1-submit-tosca-nginx.sh` to create the minimum number of MiCADO worker nodes and to deploy the Kubernetes Deployment including the nginx app defined in the `nginx.yaml` TOSCA description.
- Step4a: run `2-list-apps.sh` to see currently running applications and their IDs, as well as the ports forwarded to 8080 for accessing the HTTP service, which should now be accessible on `<micado_worker_ip>:30012`
- Step5: run `3-generate-traffic.sh` to generate some HTTP traffic. After thirty seconds or so, you will see the system respond by scaling up containers, and eventually virtual machines to the maximum specified. **NOTE:** In some cases, depending on your cloud, the pre-configured load test may be too weak to trigger a scaling response from MiCADO. If this is the case, edit the file `3-generate-traffic.sh` and increase the load options in the command on the very last line, for example `wrk -t4 -c40 -d8m http://.....`. On the other hand, a load test too powerful will be like launching a denial-of-service attack on yourself.
- Step5a: the load test will finish after 10 minutes and the infrastructure will scale back down
- Step6: run `4-undeploy-nginx.sh` to remove the nginx deployment and all the MiCADO worker nodes

### 1.5.4 wordpress

This application deploys a wordpress blog, complete with MySQL server and a Network File Share for persistent data storage. It is a proof-of-concept and is **NOT** production ready. The policy defined for this application scales up/down both nodes and the wordpress frontend container based on network load. wrk (`apt-get install wrk` | <https://github.com/wg/wrk>) is recommended for HTTP load testing, but you can use any load generator you wish.

**Note:** make sure you have the jq tool and wrk benchmarking app installed as these are required by the helper scripts to force scaling. Best results for wrk are seen on multi-core systems.

- Step1: make a copy of the TOSCA file which is appropriate for your cloud - `wordpress_<your_cloud>.yaml` - and name it `wordpress.yaml`

- Step2: fill in the requested fields beginning with `ADD_YOUR_...`. These will differ depending on which cloud you are using.

**Important:** Make sure you create the appropriate firewall policy (port 30010 must be open) for the MiCADO workers as described [here!](#)

- In CloudSigma, for example, the `libdrive_id`, `public_key_id` and `firewall_policy` fields must be completed. Without these, CloudSigma does not have enough information to launch your worker nodes. All information is found on the CloudSigma Web UI. `libdrive_id` is the long alphanumeric string in the URL when a drive is selected under “Storage/Library”. `public_key_id` is under the “Access & Security/Keys Management” menu as **Uuid**. `firewall_policy` can be found when selecting a rule defined under the “Networking/Policies” menu. The following ports must be opened for MiCADO workers: *all inbound connections from MiCADO master*
- Step3: Update the parameter file, called `_settings`. You need the ip address for the MiCADO master and should name the deployment by setting the `APP_ID`. **\*the application ID can not contain any underscores ( \_ )** The `FRONTEND_NAME`: must match the name given to the application in TOSCA (default: **wordpress**) You should also change the SSL user/password/port information if they are different from the default.
- Step4: run `1-submit-tosca-wordpress.sh` to create the minimum number of MiCADO worker nodes and to deploy the Kubernetes Deployments for the NFS and MySQL servers and the Wordpress frontend.
- Step4a: run `2-list-apps.sh` to see currently running applications and their IDs, as well as the nodePort open on the host for accessing the HTTP service (defaults to 30010)
- Step5: navigate to your wordpress blog (generally at `<worker_node_ip>:30010`) and go through the setup tasks until you can see the front page of your blog
- Step6: run `3-generate-traffic.sh` to generate some HTTP traffic. After thirty seconds or so, you will see the system respond by scaling up a VM and containers to the maximum specified. **NOTE:** In some cases, depending on your cloud, the pre-configured load test may be too weak to trigger a scaling response from MiCADO. If this is the case, edit the file `3-generate-traffic.sh` and increase the load options in the command on the very last line, for example `wrk -t4 -c40 -d8m http://.....`. On the other hand, a load test too powerful will be like launching a denial-of-service attack on yourself.
- Step6a: the load test will stop after 10minutes and the infrastructure will scale back down
- Step7: run `4-undeploy-wordpress.sh` to remove the wordpress deployment and all the MiCADO worker nodes

## 1.6 Release Notes

### v0.8.0 (30 September 2019)

- simplify ADTs by introducing pre-defined TOSCA node types
- add support for Kubernetes ConfigMaps, Namespaces and multi-container Pods
- metric collection (disabled by default) is now enabled with “monitoring” policy
- upgrade all components (Docker, Kubernetes, Grafana, Prometheus, etc...)
- introduce new Optimizer supported scaling
- add MiCADO version on dashboard and Grafana
- introduce log rotate for Docker and components
- introduce node downscale mechanism with node selection
- redirect stdout of `scaling_rule` usercode to different log file

- add support of keystone V3 for OpenStack in Occopus
- improve cloud API handling in Occopus
- make the master node web authentication timeout configurable
- make master-worker node VPN connection more restrictive
- implement ADT-based application secret distribution
- push cloud secrets to Credential Store at deploy time
- implement Security Policy Manager adaptor in the TOSCA Submitter
- add support for configuring application-level firewalling rules for the application through the ADT (FWaaS)
- generate node certificate with the right common name for the master node
- make the micadoctl command line utility to work after the transition to Kubernetes pods
- fix keypair distribution to worker nodes
- update TOSCA template for Kubernetes application-level secret distribution
- refactor Kubernetes translation
- fix Policy Keeper Kubernetes node maintenance
- propagate Kubelet configuration to worker nodes
- support system cGroup driver by Docker & Kubernetes
- fix Kubernetes node objects to be deleted on “undeploy”
- fix Occopus create & import actions to correctly raise exceptions
- fix Occopus updates not to kill unrelated nodes
- support updates of an ADT with no Occopus nodes
- support updates of an ADT with no Kubernetes nodes
- add a timeout to Kubernetes undeploy
- simplify hosts.yml file

#### **v0.7.3 (14 Jun 2019)**

- update MiCADO internal core services to run in Kubernetes pods
- remove Consul and replace it with Prometheus’ Kubernetes Service Discovery
- update cAdvisor and NodeExporter to run as Kubernetes DaemonSets
- introduce the support for creating prepared image for the MiCADO master and the MiCADO worker
- introduce the support for deploying unique “sets” of virtual machines scaling independently
- update Grafana to track the independently scaling VMs from the drop-down Node ID
- update scrape interval between Prometheus and cAdvisor to be less frequent
- fix the Occopus Adaptor to correctly raise exceptions for the submitter
- update Kubernetes Dashboard to improve RBAC permissions
- update the Flannel Overlay deployment
- update the Kubernetes eviction thresholds on the Master node to be lowered
- remove Docker-Compose from Master & Workers

- fix dependencies and vulnerabilities
- add dry-run support for the Submitter upon launch of TOSCA ADT
- add new api call for the Submitter to validate TOSCA template
- improve Submitter logs
- improve Submitter responses to users
- improve handling of wrong template by Submitter
- add support for hv\_relaxed and hv\_tsc CloudSigma specific properties
- add support for tagging EC2 type resources
- add disk and free space checking to the deployment playbook
- update the Wordpress demo to demonstrate “virtual machine sets”
- update the cQueue demo to demonstrate “virtual machine sets”
- fix and improve the NGINX demo

### **v0.7.2-rev1 (01 Apr 2019)**

- fix dependency issue for Kubernetes 1.13.1 ([kubernetes/kubernetes#75683](https://github.com/kubernetes/kubernetes/issues/75683))

### **v0.7.2 (25 Feb 2019)**

- add checking for minimal memory on micado master at deployment
- support private networks on cloudsigma
- support user-defined contextualisation
- support re-use across other container & cloud orchestrators in ADT
- new TOSCA to Kubernetes Manifest Adaptor
- add support for creating DaemonSets, Jobs, StatefulSets (with limited functionality) and standalone Pods
- add support for creating PersistentVolumes & PVClaims
- add support for specifying custom service details (NodePort, ClusterIP, etc.)
- minor improvements to Grafana dashboard
- support asynchronous calls through TOSCASubmitter API
- fix kubectl error on MiCADO Master restart
- fix TOSCASubmitter rollback on errors
- fix TOSCASubmitter status & output display
- add support for encrypting master-worker communication
- automatically provision and revoke security credentials for worker nodes
- update default MTU to 1400 to ensure compatibility with OpenStack and AWS
- add Credential Store security enabler
- add Security Policy Manager security enabler
- add Image Integrity Verifier Security enabler
- add Crypto Engine security enabler
- add support for kubernetes secrets

- reimplement Credential Manager using the flask-users library

**v0.7.1 (10 Jan 2019)**

- Fix: Add SKIP back to Dashboard (defaults changed in v1.13.1)
- Fix: URL not found for Kubernetes manifest files
- Fix: Make sure worker node sets hostname correctly
- Fix: Don't update Kubernetes if template not changed
- Fix: Make playbook more idempotent
- Add Support for outputs via TOSCA ADT
- Add Kubernetes service discovery support to Prometheus
- Add new demo: nginx (HTTP request scaling)

**v0.7.0 (12 Dec 2018)**

- Introduce Kubernetes as the primary container orchestration engine
- Replace the swarm-visualiser with the Kubernetes Dashboard

**v0.6.1 (15 Oct 2018)**

- enable VM-only deployments
- add support for special characters in SSL credentials
- fix missing vm instance number reset at undeployment
- add option to disable auto-updates on worker nodes
- modify default launch-order of TOSCA adaptors
- add cloud-specific TOSCA templates and improve helper scripts for stressng
- flatten CPU scaling policies
- improve virtual machine build time
- fix Zorp starting dependency
- fix Docker login timing issue
- remove unnecessary port from docker compose file
- enable Prometheus DB export

**v0.6.0 (10 Sept 2018)**

- introduce documentation repository and host its content at <http://micado-scale.readthedocs.io>
- improve MiCADO master containers restart policy
- fix MTU issue in relation to Docker
- fix Occopus restart issue
- fix health-checking for Cloudbroker-AWS platform
- update host naming convention for worker and master nodes
- make wait-update task idempotent in ansible playbook
- fix issue with worker node deployment in EC2 clouds
- fix issue with user-defined Docker networks in OpenStack clouds

- make Submitter response message structure uniform
- add 'nodes' and 'services' query methods to REST API
- improve 'stressng' and 'cqueue' test helper scripts
- add more compose properties to custom TOSCA definition
- fix floating ip issues in the Dashboard component
- add new links to Dashboard to reflect the changes introduced by reverse proxying
- fix Dashboard to generate links based on the contents of the Host header to find the frontend URL automatically
- make consul security encryption based on generated random key instead of static key
- add reverse proxy, TLS encryption and application-level firewalling capabilities to the web interfaces exposed by the MiCADO master node
- add packet filtering for closing down non-public ports
- add systemd unit for MiCADO services
- update the ansible playbook to use the built-in service module for installing and handling MiCADO services
- update the documentation to reflect the changes after the introduction of reverse proxying
- add support for form-based authentication of exposed web services
- add COLA-themed login page
- add the Credential Manager component to store and handle web service users and passwords securely
- add support for provisioning a user to the Credential Manager via Ansible
- add support for user and admin roles in the Credential Manager
- add support for authorization of the web services based on user role
- add documentation about the Ansible Vault mechanism to protect sensitive deployment details
- add support for HTTP basic authentication for APIs
- add support for making the web interface's listening port configurable
- update the documentation of API calls in terms of authentication, encryption and reverse proxying
- add micadoctl tool for user and service management
- add HTTP method filter to firewall in order to control requests directed to containers
- add support for IPv6 exposure of services
- add IPv6 packet filtering

### **v0.5.0 (12 July 2018)**

- introduce supporting TOSCA
- introduce supporting user-defined scaling policy
- dashboard added with Docker Visualizer, Grafana, Prometheus
- deployment with Ansible playbook
- support private docker registry
- improve persistence of MiCADO master services