
KubeNow Documentation

Release 0.3.1

mcapuccini

Mar 06, 2018

1	Prerequisites	3
1.1	Install Docker	3
1.2	Get KubeNow	3
2	Deploy KubeNow on a host cloud	5
2.1	Deploy on OpenStack	5
2.2	Deploy on Google Cloud (GCE)	7
2.3	Deploy on Amazon Web Services (AWS)	8
2.4	Deploy on Microsoft Azure	10
3	Deploy your first application	13
3.1	Traefik reverse proxy	13
4	Clean after yourself	15
5	Terraform troubleshooting	17
5.1	Corrupted Terraform state	17
6	OpenStack troubleshooting	19
6.1	Console logs on OpenStack	19
6.2	Missing DomainID or DomainName to authenticate by Username	19
7	Kubernetes troubleshooting	21
7.1	List kubernetes pods	21
7.2	Describe status of a specific pod	21
7.3	Get the kubelet service log	22
8	More troubleshooting	23
8.1	SSH connection errors	23
8.2	Figure out hostnames and IP numbers	23
9	Image building	25

Welcome to KubeNow's documentation! This is a place where we aim to help you to provision Kubernetes, the KubeNow's way. If you are new to Kubernetes, and to cloud computing, this is going to take a while to grasp the first time. Luckily, once you get the procedure, it's going to be very quick to spawn your clusters.

CHAPTER 1

Prerequisites

1.1 Install Docker


KubeNow provisioners are distributed via [Docker](#). Please start by installing Docker on your workstation: [Install Docker](#).

1.2 Get KubeNow

In order to get the provisioning tools please run:

```
docker pull kubenow/provisioners:0.3.1
```

We wrote up a handy CLI that wraps around the Docker container above, you can install it with a one-liner:

```
curl -Lo kn https://github.com/kubenow/KubeNow/releases/download/0.3.1/kn &&   
↪ chmod +x kn && sudo mv kn /usr/local/bin/
```

Deploy KubeNow on a host cloud

The following steps are slightly different for each host cloud. Here you find a section for each of the supported providers.

Sections

- *Deploy KubeNow on a host cloud*
 - *Deploy on OpenStack*
 - *Deploy on Google Cloud (GCE)*
 - *Deploy on Amazon Web Services (AWS)*
 - *Deploy on Microsoft Azure*

2.1 Deploy on OpenStack

2.1.1 Prerequisites

In this section we assume that:

- You have downloaded and sourced the OpenStack RC file for your tenancy: `source project-openrc.sh` (<https://docs.openstack.org/user-guide/common/cli-set-environment-variables-using-openstack-rc.html#download-and-source-the-openstack-rc-file>)
- You have a floating IP quota that allows to allocate at least one public IP

2.1.2 Deployment configuration

First we need to initialize a deploy configuration directory by running:

```
kn init openstack my_deployment
```

The configuration directory contains a new SSH key pair for your deployments, and a [Terraform](#) configuration template that we need to fill in.

Locate into `my_deployment` :

```
cd my_deployment
```

In the configuration file `terraform.tfvars` you will need to set at least:

Cluster configuration

- **cluster_prefix**: every resource in your tenancy will be named with this prefix
- **floating_ip_pool**: a floating IP pool label
- **external_network_uuid**: the uuid of the external network in the OpenStack tenancy

If you are wondering where you can get `floating_ip_pool` and `external_network_uuid`, then one way is to inquiry your OpenStack settings by running:

```
kn openstack network list
```

Depending on your tenancy settings you should get a similar output:

```
+-----+-----+...
| ID              | Label          |...
+-----+-----+...
| 5f274562-89b6-4ab2-a18f-94b159b0b85d | internal       |...
| d9384930-baa5-422b-8657-1d42fb54f89c | net_external   |...
+-----+-----+...
```

Thus in this specific case the above mentioned fields will be set as it follows:

```
floating_ip_pool = "net_external"
external_network_uuid = "d9384930-baa5-422b-8657-1d42fb54f89c"
```

Master configuration

- **master_flavor**: an instance flavor for the master

Node configuration

- **node_count**: number of Kubernetes nodes to be created (no floating IP is needed for these nodes)
- **node_flavor**: an instance flavor name for the Kubernetes nodes

If you are wondering yet again where you can fetch correct flavor label names then no worries, you are not being a stranger here. The openstack command-line interface will come in handy. Just run the following command:

```
kn openstack flavor list
```

Depending on your tenancy settings you should get a similar output:

```
+-----+-----+...
| ID      | Name        |...
+-----+-----+...
| 8c7ef1  | ssc.tiny    |...
| 8d7ef2  | ssc.small   |...
| 8e7ef3  | ssc.medium  |...
| 8f7ef4  | ssc.large   |...
```

```
| 8g7ef5 | ssc.xlarge | ...
+-----+-----+...
```

You may want to select the flavor according to the number of resources you'd like to allocate. E.g.:

```
master_flavor = "ssc.medium"
node_flavor = "ssc.large"
```

2.1.3 Deploy KubeNow

Once you are done with your settings you are ready to deploy your cluster running:

```
kn apply
```

The first time you are going to deploy it will take longer, since the KubeNow image needs to be imported. Future deployments will be considerably faster, since the image will be already present in your user space.

To check that your cluster is up and running you can run:

```
kn kubectl get nodes
```

As long as you are in the `my_deployment` directory you can use `kubectl` over SSH to control Kubernetes. If you want to open an interactive SSH terminal onto the master then you can use the `kn ssh` command:

```
kn ssh
```

If everything went well, now you are ready to *deploy your first application*.

2.2 Deploy on Google Cloud (GCE)

2.2.1 Prerequisites

In this section we assume that:

- You have enabled the Google Compute Engine API: API Manager > Library > Compute Engine API > Enable
- You have created and downloaded a service account file for your GCE project: Api manager > Credentials > Create credentials > Service account key

2.2.2 Deployment configuration

First we need to initialize a deployment configuration directory by running:

```
kn init gce my_deployment
```

The configuration directory contains a new SSH key pair for your deployments, and a [Terraform](#) configuration template that we need to fill in.

Locate into `my_deployment` :

```
cd my_deployment
```

In the configuration file `terraform.tfvars` you will need to set at least:

Cluster configuration

- **cluster_prefix**: every resource in your project will be named with this prefix (the name must match `(?:[a-z](?:[a-z0-9]{0,61}[a-z0-9])?)`, e.g. “kubenow”)

Google credentials

- **gce_project**: your project id
- **gce_zone**: some GCE zone (e.g. europe-west1-b)

Master configuration

- **master_flavor**: an instance flavor for the master (e.g. n1-standard-2)
- **master_disk_size**: master disk size in GB

Node configuration

- **node_count**: number of Kubernetes nodes to be created
- **node_flavor**: an instance flavor for the Kubernetes nodes (e.g. n1-standard-2)
- **node_disk_size**: nodes disk size in GB

In addition, when deploying on GCE you need to copy your service account file in the deployment configuration directory:

```
# assuming that you are in my_deployment
cp /path/to/service-account.json ./
```

2.2.3 Deploy KubeNow

Once you are done with your settings you are ready deploy your cluster running:

```
kn apply
```

The first time you are going to deploy it will take longer, since the KubeNow image needs to be imported. Future deployments will be considerably faster, since the image will be already present in your user space.

To check that your cluster is up and running you can run:

```
kn kubectl get nodes
```

As long as you are in the `my_deployment` directory you can use `kubectl` over SSH to control Kubernetes. If you want to open an interactive SSH terminal onto the master then you can use the `kn ssh` command:

```
kn ssh
```

If everything went well, now you are ready to *deploy your first application*.

2.3 Deploy on Amazon Web Services (AWS)

2.3.1 Prerequisites

In this section we assume that:

- You have an IAM user along with its `access key` and `security credentials` (http://docs.aws.amazon.com/IAM/latest/UserGuide/id_users_create.html)

2.3.2 Deployment configuration

First we need to initialize a deploy configuration directory by running:

```
kn init aws my_deployment
```

The configuration directory contains a new SSH key pair for your deployments, and a [Terraform](#) configuration template that we need to fill in.

Locate into `my_deployment` :

```
cd my_deployment
```

In the configuration file `terraform.tfvars` you will need to set at least:

Cluster configuration

- **cluster_prefix**: every resource in your tenancy will be named with this prefix
- **aws_region**: the region where your cluster will be bootstrapped (e.g. `eu-west-1`)
- **availability_zone**: an availability zone for your cluster (e.g. `eu-west-1a`)

Credentials

- **aws_access_key_id**: your access key id
- **aws_secret_access_key**: your secret access key

Master configuration

- **master_instance_type**: an instance type for the master (e.g. `t2.medium`)
- **master_disk_size**: edges disk size in GB

Node configuration

- **node_count**: number of Kubernetes nodes to be created
- **node_instance_type**: an instance type for the Kubernetes nodes (e.g. `t2.medium`)
- **node_disk_size**: edges disk size in GB

2.3.3 Deploy KubeNow

Once you are done with your settings you are ready deploy your cluster running:

```
kn apply
```

To check that your cluster is up and running you can run:

```
kn kubectl get nodes
```

As long as you are in the `my_deployment` directory you can use `kubectl` over SSH to control Kubernetes. If you want to open an interactive SSH terminal onto the master then you can use the `kn ssh` command:

```
kn ssh
```

If everything went well, now you are ready to *deploy your first application*.

2.4 Deploy on Microsoft Azure

2.4.1 Prerequisites

In this section we assume that:

- You have created an application API key (Service Principal) in your Microsoft Azure subscription: (https://www.terraform.io/docs/providers/azurerm/authenticating_via_service_principal.html#creating-a-service-principal)

2.4.2 Deployment configuration

First we need to initialize a deploy configuration directory by running:

```
kn init azure my_deployment
```

The configuration directory contains a new SSH key pair for your deployments, and a [Terraform](#) configuration template that we need to fill in.

Locate into `my_deployment` :

```
cd my_deployment
```

In the configuration file `terraform.tfvars` you will need to set at least:

Cluster configuration

- **cluster_prefix**: every resource in your tenancy will be named with this prefix
- **location**: some Azure location (e.g. `West Europe`)

Azure credentials

- **subscription_id**: your subscription id
- **client_id**: your client id (also called `appId`)
- **client_secret**: your client secret (also called `password`)
- **tenant_id**: your tenant id

Master configuration

- **master_vm_size**: the vm size for the master (e.g. `Standard_DS2_v2`)

Node configuration

- **node_count**: number of Kubernetes nodes to be created
- **node_vm_size**: the vm size for the Kubernetes nodes (e.g. `Standard_DS2_v2`)

2.4.3 Deploy KubeNow

Once you are done with your settings you are ready deploy your cluster running:

```
kn apply
```

The first time you are going to deploy it will take longer, since the KubeNow image needs to be imported. Future deployments will be considerably faster, since the image will be already present in your user space.

To check that your cluster is up and running you can run:

```
kn kubectl get nodes
```

As long as you are in the `my_deployment` directory you can use `kubectl` over SSH to control Kubernetes. If you want to open an interactive SSH terminal onto the master then you can use the `kn ssh` command:

```
kn ssh
```

If everything went well, now you are ready to *deploy your first application*.

Deploy your first application

In this guide we are going to deploy a simple application: [cheese](#). We will deploy 3 web pages with a 2 replication factor. The master node will act as a reverse proxy, load balancing the requests among the replicas in the Kubernetes nodes.

The simple cluster that we just deployed uses [nip.io](#) as base domain for incoming HTTP traffic. First, we need to figure out our cluster domain by running:

```
grep domain inventory
```

The command will return something like `domain=37.153.138.137.nip.io`, meaning that our cluster domain name in this case would be `37.153.138.137.nip.io`.

In KubeNow we encourage to deploy and define SaaS-layer applications using [Helm](#). The KubeNow community maintain a Helm repository that contains applications that are developed and tested for KubeNow: <https://github.com/kubenow/helm-charts>. To deploy the cheese application you can run the following command, substituting `<your-domain>` with the domain name that you got above:

```
kn helm install --name cheese --set domain=<your-domain> kubenow/cheese
```

If everything goes well you should be able to access the web pages at:

- <http://stilton.<your-domain>>
- <http://cheddar.<your-domain>>
- <http://wensleydale.<your-domain>>

3.1 Traefik reverse proxy

KubeNow uses the [Traefik](#) reverse proxy as ingress controller for your cluster. Traefik is installed on one or more nodes, namely edge nodes, which have a public IP associated. In this way, we can access services with a few floating IP without needing LBaaS, which may not be available on certain cloud providers.

In the default setting KubeNow doesn't deploy any edge node, and it runs Traefik in the master node.

3.1.1 Accessing the Traefik UI

One simple and quick way to access the Traefik UI is to tunnel via SSH to one of the edge nodes with the following command:

```
ssh -N -f -L localhost:8081:localhost:8081 ubuntu@<your-domain>
```

Using SSH tunnelling, the Traefik UI should be reachable at <http://localhost:8081>, and it should look something like this:

The screenshot displays the Traefik UI interface. At the top, there are tabs for 'Providers', 'Health', 'Documentation', and 'traefik.io'. Below the tabs, a search bar shows 'kubernetes'. The main content area is divided into two columns. The left column shows the configuration for three services: 'galaxy.carmat.phenomenal.cloud/', 'luigi.carmat.phenomenal.cloud/', and 'notebook.carmat.phenomenal.cloud/'. Each service configuration includes a table with 'Route' and 'Rule' columns. The right column shows the configuration for 'galaxy.carmat.phenomenal.cloud/', 'luigi.carmat.phenomenal.cloud/', and 'notebook.carmat.phenomenal.cloud/'. Each service configuration includes a table with 'Server', 'URL', and 'Weight' columns. Below each table is a 'Load Balancer: wrr' button.

Route	Rule
/	PathPrefix:/
galaxy.carmat.phenomenal.cloud	Host:galaxy.carmat.phenomenal.cloud

Server	URL	Weight
http://10.44.0.2:8080	http://10.44.0.2:8080	1

Route	Rule
/	PathPrefix:/
luigi.carmat.phenomenal.cloud	Host:luigi.carmat.phenomenal.cloud

Server	URL	Weight
http://10.32.0.2:8082	http://10.32.0.2:8082	1

Route	Rule
/	PathPrefix:/
notebook.carmat.phenomenal.cloud	Host:notebook.carmat.phenomenal.cloud

Server	URL	Weight
http://10.44.0.1:8888	http://10.44.0.1:8888	1

In the left side you can find your deployed frontends URLs, whereas on the right side the backend services.

CHAPTER 4

Clean after yourself

Cloud resources are typically pay-per-use, hence it is good to release them when they are not used. Here we show how to destroy a KubeNow cluster.

To release the resources, please run:

```
kn destroy
```

Warning: if you delete the cluster configuration directory (`my_deployment`) the cluster status will be lost, and you'll have to delete the resources manually.

Terraform troubleshooting

Since Terraform applies changes incrementally, when there is a minor issue (e.g. network timeout) it's sufficient to rerun the command. However, here we try to collect some tips that can be useful when rerunning the command doesn't help.

Contents

- *Terraform troubleshooting*
 - *Corrupted Terraform state*

5.1 Corrupted Terraform state

Due to network issues, Terraform state files can get out of synch with your infrastructure, and cause problems. Since Terraform apply changes incrementally. A possible way to fix the issue is to destroy your nodes manually, and remove all state files and cached modules:

```
rm -R .terraform/  
rm terraform.tfstate  
rm terraform.tfstate.backup
```


OpenStack troubleshooting

Contents

- *OpenStack troubleshooting*
 - *Console logs on OpenStack*
 - *Missing DomainID or DomainName to authenticate by Username*

6.1 Console logs on OpenStack

Can't get the status from the nodes with `ansible master -a "kubectl get nodes"`? The nodes might not have started all right. Checking the console logs with nova could help.

List node IDs, floating IPs etc.:

```
nova list
```

Show console output from node of interest:

```
nova console-log <node-id>
```

6.2 Missing DomainID or DomainName to authenticate by Username

When running terraform and/or packer you may be prompted with the following error:

```
You must provide exactly one of DomainID or DomainName to authenticate by Username
```

If this is the case, then setting either `OS_DOMAIN_ID` or `OS_DOMAIN_NAME` in your environment should fix the issue. For further information, please refer to this document: <https://www.terraform.io/docs/providers/openstack/index.html>.

Kubernetes troubleshooting

Here you can find some frequently used commands to list the status and logs of kubernetes. If this doesn't help, please refer to <http://kubernetes.io/docs>.

Contents

- *Kubernetes troubleshooting*
 - *List kubernetes pods*
 - *Describe status of a specific pod*
 - *Get the kubelet service log*

7.1 List kubernetes pods

```
# If you are logged into the node via SSH:
kubectl get pods --all-namespaces

# With Ansible from your local computer:
kn ansible master -a "kubectl get pods --all-namespaces"
```

7.2 Describe status of a specific pod

```
# If you are logged into the node via SSH:
kubectl describe pod <pod id> --all-namespaces

# With Ansible from your local computer:
kn ansible master -a "kubectl describe pod <pod id> --all-namespaces"
```

7.3 Get the kubelet service log

```
# If you are logged into the node via SSH:
sudo journalctl -r -u kubelet

# With Ansible from your local computer:
kn ansible master -a "journalctl -r -u kubelet"
```

More troubleshooting

Contents

- *More troubleshooting*
 - *SSH connection errors*
 - *Figure out hostnames and IP numbers*

8.1 SSH connection errors

In case of SSH connection errors:

- Make sure to add your private SSH key to your local keyring:

```
ssh-add ~/.ssh/id_rsa
```

- Make sure that port 22 is allowed in your cloud provider security settings.

If you still experience problems, checking out the console logs from your cloud provider could help.

8.2 Figure out hostnames and IP numbers

The bootstrap step should create an Ansible inventory list, which contains hostnames and IP addresses:

```
cat inventory
```


CHAPTER 9

Image building

KubeNow uses prebuilt images to speed up the deployment. Image continuous integration is defined in this repository: <https://github.com/kubeflow/image>.

The images are exported on AWS, GCE and Azure:

- `https://storage.googleapis.com/kubeflow-images/kubeflow-<version-without-dots>.tar.gz`
- `https://s3.amazonaws.com/kubeflow-us-east-1/kubeflow-<version-without-dots>.qcow2`
- `https://kubeflow.blob.core.windows.net/system?restype=container&comp=list`

Please refer to this page to figure out the image version: <https://github.com/kubeflow/image/releases>. It is important to point out that the image versioning is now disjoint from the main KubeNow repository versioning. The main reason lies in the fact that pre-built images require less revisions and updates compared to the main KubeNow package.