

# Scalable JupyterHub deployment on bare-metal Kubernetes

Rohini Joshi

[r.joshi@skatelescope.org](mailto:r.joshi@skatelescope.org)

# JupyterHub Prototype

- Prototype running on <http://130.246.212.44/escape>
- Openstack running bare-metal Kubernetes cluster
- Helm based JupyterHub deployment
- Dynamic persistent storage with NFS
- ESCAPE IAM auth integration example [here](#)
- Custom user environments
- DEMO

<input type="checkbox"/>	Instance Name	Image Name	IP Address	Flavor	Key Pair	Status	Availability Zone	Task	Power State	Age	Actions
<input type="checkbox"/>	kubernetes-worker-3	Ubuntu-Bionic-NoGui	192.168.50.78	c2.large	RohiniAENEASUIMachine	Active	<input type="checkbox"/> ceph	None	Running	3 weeks, 3 days	CREATE SNAPSHOT ▾
<input type="checkbox"/>	kubernetes-worker-2	Ubuntu-Bionic-NoGui	192.168.50.29	c3.medium	RohiniAENEASUIMachine	Active	<input type="checkbox"/> ceph	None	Running	1 month, 1 week	CREATE SNAPSHOT ▾
<input type="checkbox"/>	kubernetes-worker-1	Ubuntu-Bionic-NoGui	192.168.50.61	c3.medium	RohiniAENEASUIMachine	Active	<input type="checkbox"/> ceph	None	Running	1 month, 1 week	CREATE SNAPSHOT ▾
<input type="checkbox"/>	kubernetes-head-01	Ubuntu-Bionic-NoGui	192.168.50.44, 130.246.212.44	c3.large	RohiniLaptop	Active	<input type="checkbox"/> ceph	None	Running	1 month, 1 week	CREATE SNAPSHOT ▾

# Lessons learnt

Start small, adding a variable at a time

Persistent volumes are hard!

Magnum clusters or another cluster orchestration service will simplify things

Kubernetes

- Double check yaml file for misconfigs
- Lots of things can change between versions
- K8s documentation, blog posts are your friend

More useful monitoring strategy probably after testing the prototype

Custom user environments: Be careful of permissions of the default jovyan user

ssh tunnels used initially. Exposed port 80 needed for IAM auth to work

# Next steps

Documenting how to query the JupyterHub REST API. [Gitlab issue](#)


Custom environment for Rucio to interact with the data lake from a notebook

Using auth tokens and propagation of auth tokens

Find and document how to query JHub API

☰ | Collinson, James ESCAPEES

Identified SPO-793



Reproduce and develop a JupyterHub custom environment for Rucio

☰ | Collinson, James ESCAPEES

Identified SPO-786



# Useful Links

- Small self-contained deployment  
<https://medium.com/@georgepaw/jupyterhub-with-kubernetes-on-single-bare-metal-instance-tutorial-67cbd5ec0b00>
- Zero to JupyterHub  
<https://zero-to-jupyterhub.readthedocs.io/en/latest/>
- Customisation Guide  
<https://zero-to-jupyterhub.readthedocs.io/en/latest/customizing/index.html>
- Current deployment documentation  
<https://github.com/rohinijoshi06/jupyterhub-on-k8s>
- Stelios' previous tech talk  
<https://github.com/stvoutsin/escape-wp5-dev/blob/master/doc/20200608-ESCAPEWP5.pdf>