JupyterHub on Scientific OpenStack

John Garbutt - Feb 2020

StackHPC

Jupyter Ecosystem

- “Classic” Jupyter Notebook
- JupyterLab Notebook (replaces classic)
- JupyterHub (multi-user, creates both notebooks)
- Others: BinderHub, Pangeo, Conda, Voila

Jupyter Notebooks @ DiRAC

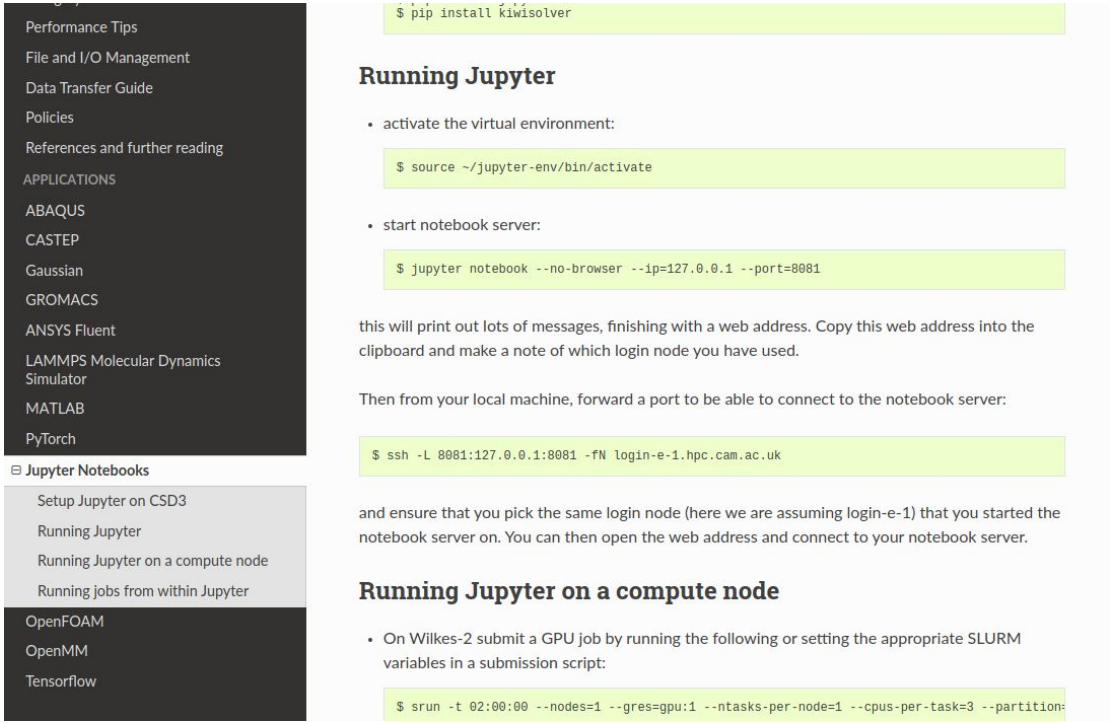

https://docs.hpc.cam.ac.uk/hpc/software-packages/jupyter.html

Pangeo

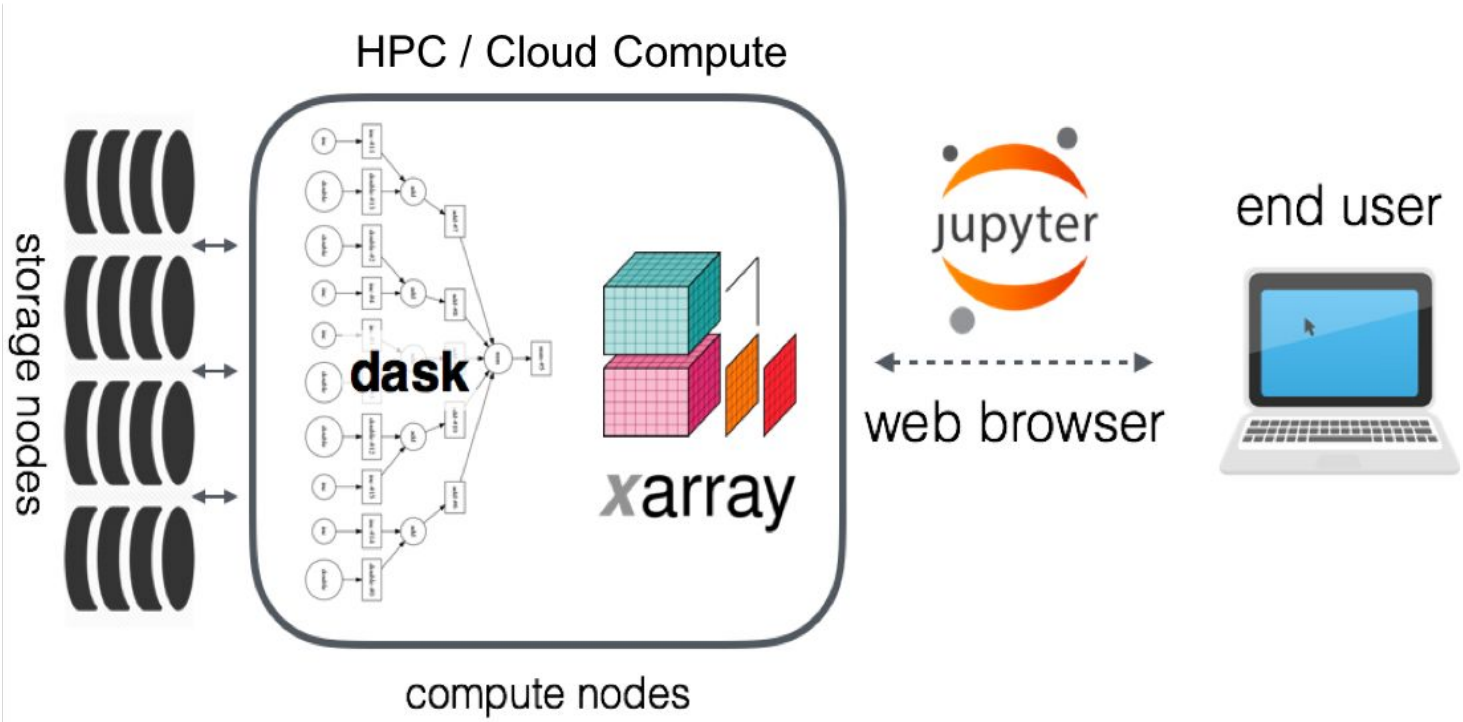
HPC / Cloud Compute
storage nodes
dask
xarray
compute nodes
jupyter
web browser
end user

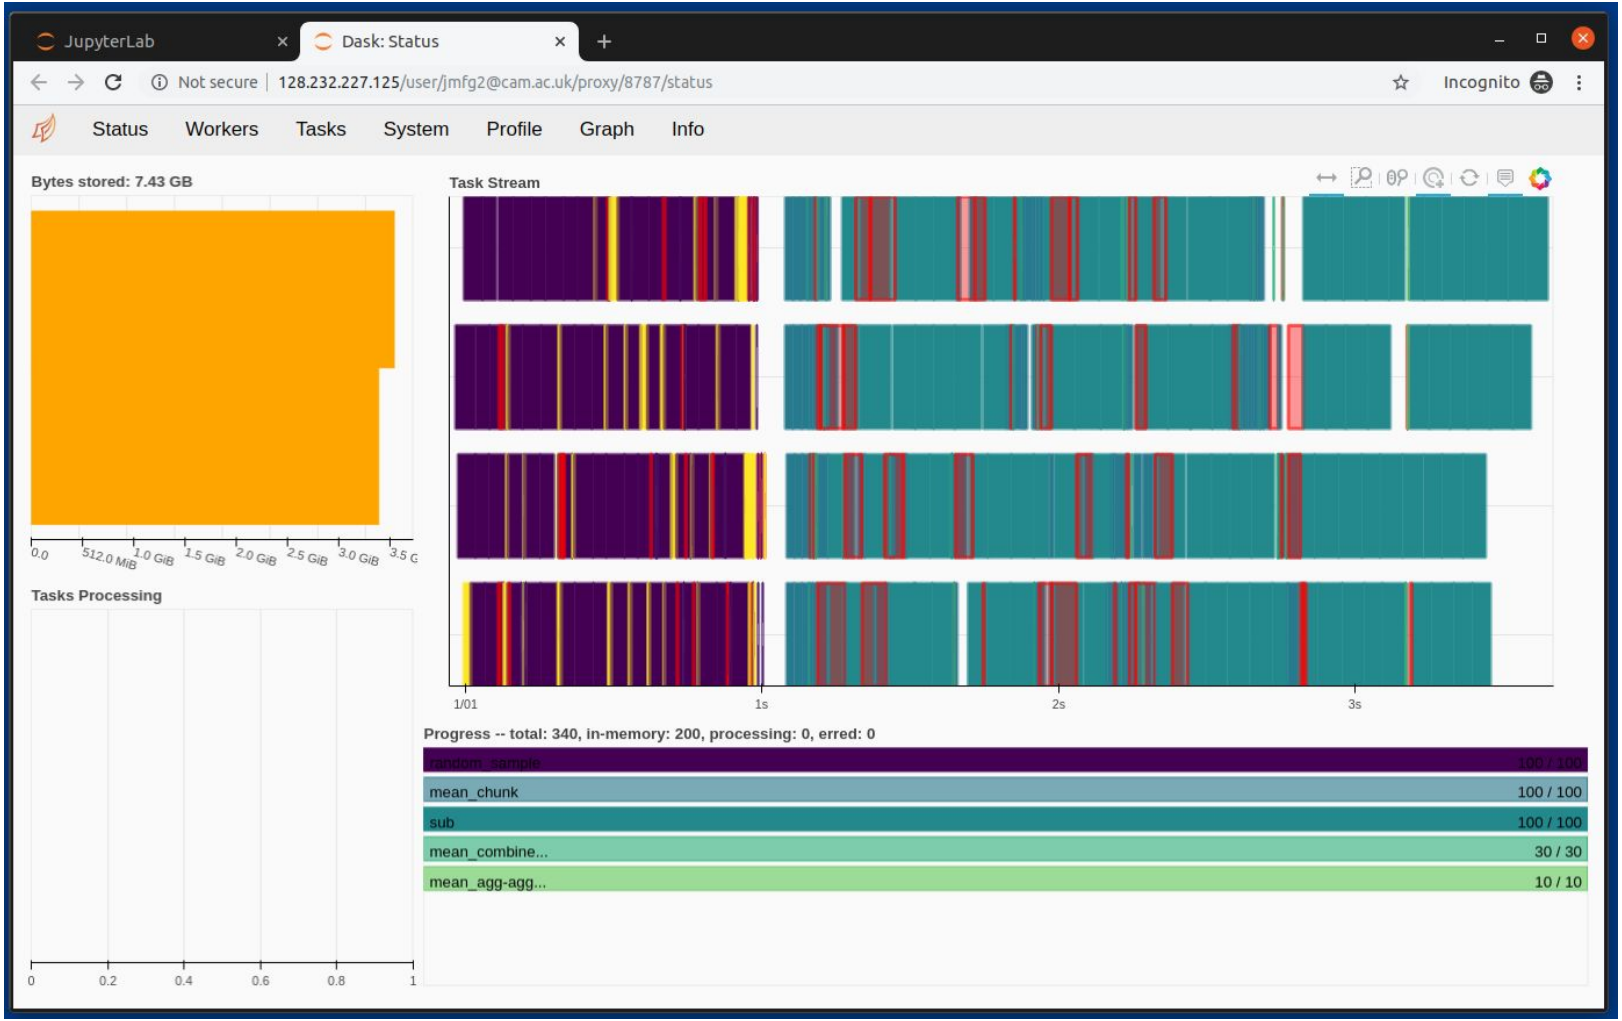
JupyterLab
Dask: Status
Not secure | 128.232.227.125/user/jmfg2@cam.ac.uk/proxy/8787/status
Incognito
Status
Workers
Tasks
System
Profile
Graph
Info
Bytes stored: 7.43 GB
0.0
512.0 MiB
1.0 GiB
1.5 GiB
2.0 GiB
2.5 GiB
3.0 GiB
Task Stream
1/01
1s
2s
3s
Tasks Processing
0
0.2
0.4
0.6
0.8
1
Progress -- total: 340, in-memory: 200, processing: 0, erred: 0
random_sample
100 / 100
mean_chunk
100 / 100
sub
100 / 100
mean_combine...
30 / 30
mean_agg-agg...
10 / 10

Deploying JupyterHub

- Ground up install
  https://jupyterhub.readthedocs.io
- The Littlest JupyterHub (single server)
  https://tljh.jupyter.org
- Zero to JupyterHub (for Kubernetes)
  https://z2jh.jupyter.org

Search the docs ...

Kubernetes on Google Cloud (GKE)

Kubernetes on Microsoft Azure Kubernetes Service (AKS)

Kubernetes on Microsoft Azure Kubernetes Service (AKS) with Autoscaling

Kubernetes on Amazon Web Services (AWS)

Kubernetes on Amazon Web Services (AWS) with Elastic Container with Kubernetes (EKS)

Kubernetes on Red Hat OpenShift

Kubernetes on IBM Cloud

Kubernetes on Digital Ocean

Setup Kubernetes

Kubernetes' documentation describes the many ways to set up a cluster. We attempt to provide quick instructions for the most painless and popular ways of setting up a Kubernetes cluster on various cloud providers and on other infrastructure.

Choose one option and proceed.

- Kubernetes on Google Cloud (GKE)
- Kubernetes on Microsoft Azure Kubernetes Service (AKS)
- Kubernetes on Microsoft Azure Kubernetes Service (AKS) with Autoscaling
- Kubernetes on Amazon Web Services (AWS)
- Kubernetes on Amazon Web Services (AWS) with Elastic Container with Kubernetes (EKS)
- Kubernetes on Red Hat OpenShift
- Kubernetes on IBM Cloud
- Kubernetes on Digital Ocean

Note

- During the process of setting up JupyterHub, you'll be creating some files for configuration purposes. It may be helpful to create a folder for your JuypterHub deployment to keep track of these files.

FY19 Digital Assets

StackHPC

OpenStack Magnum

- Kubernetes Cloud Provider OpenStack
  - Storage: Cinder, Manila
  - Load Balancer: Octavia
  - Cluster Autoscaler and Auto-healing
- Monitoring via Prometheus and Grafana

- Improved documentation:
  - Terraform to create/re-size Kubernetes
  - Ingress via Nginx, Terraform to add a Router to CephFS
  - JupyterHub on Kubernetes, with IRIS IAM (using Helm 3)
  - Apache Spark (using Helm 3)

https://github.com/RSE-Cambridge/iris-magnum

Example: Terraform for K8s

```
terraform {
 required_version = ">= 0.12, < 0.13"
}

module "cluster" {
 source = "../../modules/cluster"

 cluster_name = "my-test"
 cluster_template_name = "kubernetes-1.15.9-20200205"
 master_flavor_name = "general.v1.tiny"
 flavor_name = "general.v1.tiny"
 master_count = 1
 node_count = 1
 max_node_count = 2
}
```

Cinder Volumes in K8s

Assuming your magnum templated enabled the built in cinder volumes, you can use them for PVCs by create a new storage class:

```
cat <<END | kubectl apply -f -
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: cinder
  annotations:
provisioner: kubernetes.io/cinder
END
```

You can make it the default by applying:

```
kubectl patch storageclass cinder -p \
  '{"metadata": {"annotations":{"storageclass.kubernetes.io/is-default-class":"true"}}}'
```

Note: the above uses the in-tree cinder driver that is deprecated. We are looking to change OpenStack Magnum to rely on the more supported setup of using the cloud-provider-openstack based driver. But for now, this seems to be the best way to use cinder volumes.

Zero to Juypterhub

At this point you will have Helm installed and kubectl pointing at your kubernetes cluster created by OpenStack Magnum.

To make it easy to tidy up and separate your work, we create a new namespace to contain all the work on JupyterHub:

```
kubectl create namespace jhub
```

We can now follow this tutorial: https://zero-to-jupyterhub.readthedocs.io/en/latest/setup-jupyterhub/setup-jupyterhub.html

```
helm repo add jupyterhub https://jupyterhub.github.io/helm-chart/
helm repo update

cat >config.yml <<END
proxy:
  secretToken: "`openssl rand -hex 32`"
END

helm install jhub jupyterhub/jupyterhub --values config.yml --namespace jhub

kubectl --namespace jhub get service
```

Once the external IP appears you can now login as any username and password. This is clearly an unsafe default for a

Example: JupyterHub with IRIS IAM

```
jupyterhub:
 hub:
   extraEnv:
     OAUTH2_AUTHORIZE_URL: https://iris-iam.stfc.ac.uk/authorize
     OAUTH2_TOKEN_URL: https://iris-iam.stfc.ac.uk/token
     OAUTH_CALLBACK_URL: http://128.232.227.125/hub/oauth_callback
 auth:
   type: custom
   custom:
     className: oauthenticator.generic.GenericOAuthenticator
     config:
       login_service: "IRIS IAM"
       client_id: "75dc9e2a-3c08-436e-b110-d710d1e3c530"
       client_secret: "xxxx"
       token_url: https://iris-iam.stfc.ac.uk/token
       userdata_url: https://iris-iam.stfc.ac.uk/userinfo
       userdata_method: GET
       userdata_params: {'state': 'state'}
       username_key: preferred_username
       oauth_callback_url: "http://128.232.227.125/hub/oauth_callback"
```

```
proxy:
   service:
     type: LoadBalancer
     loadBalancerIP: 128.232.227.125
```

Demo: *JupyterHub via Pangeo*

StackHPC

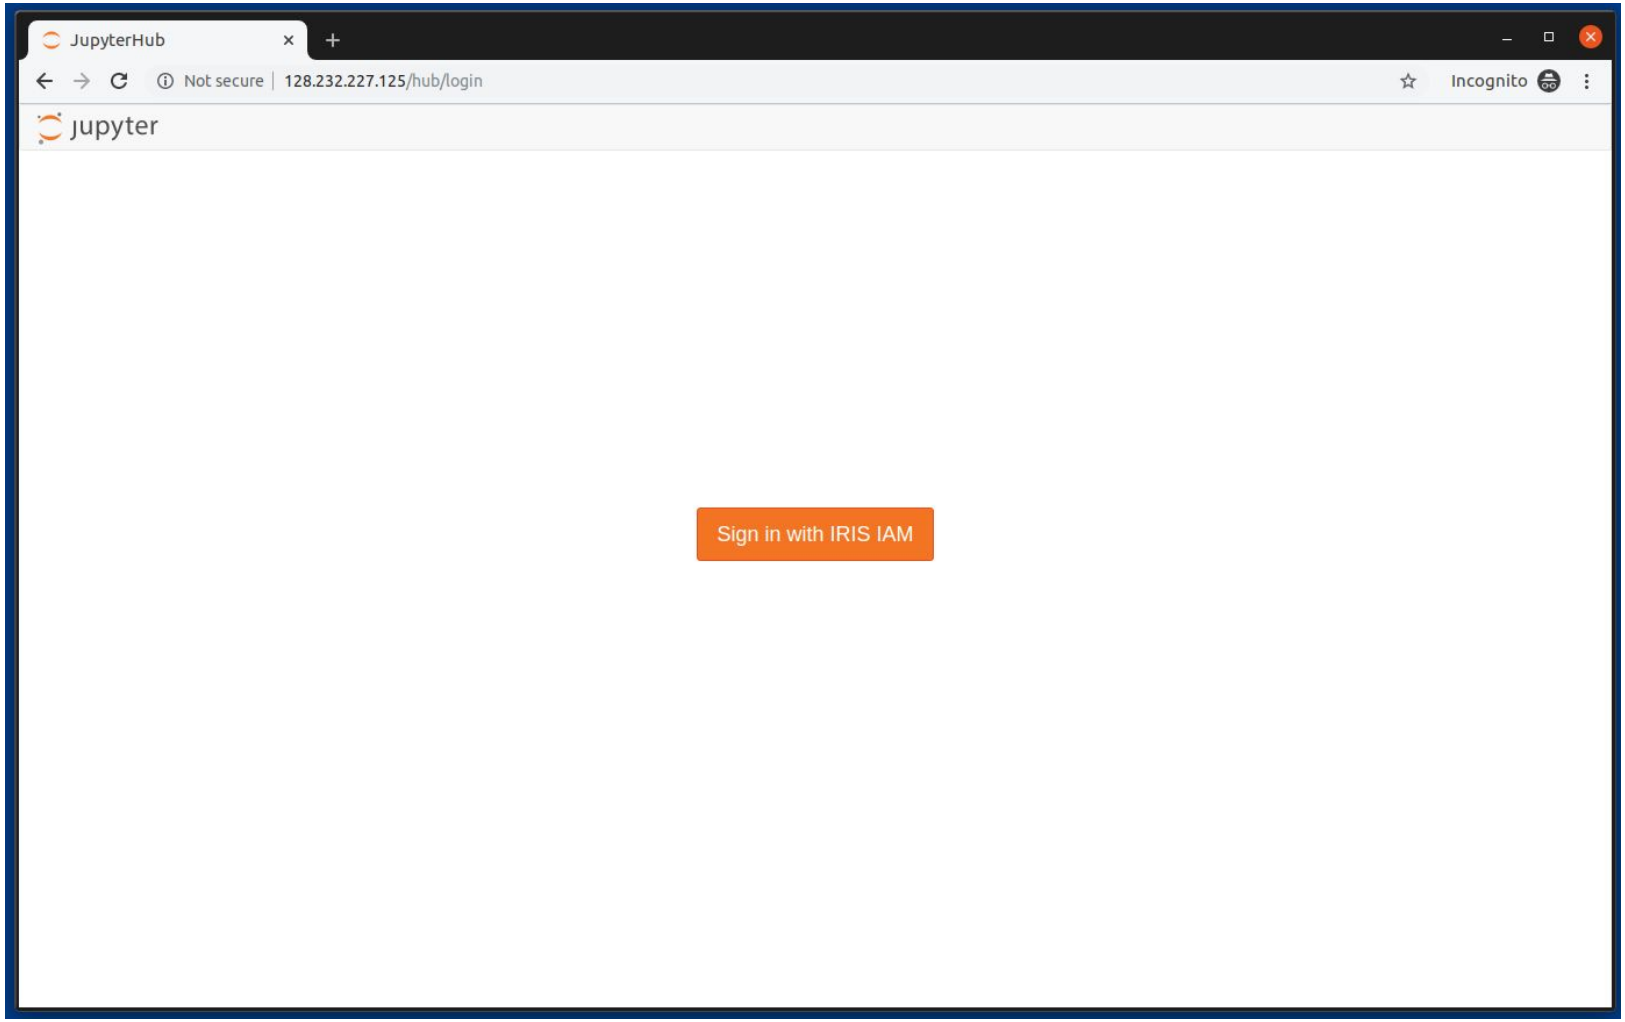
JupyterHub
Not secure | 128.232.227.125/hub/login
Incognito
jupyter
Sign in with IRIS IAM

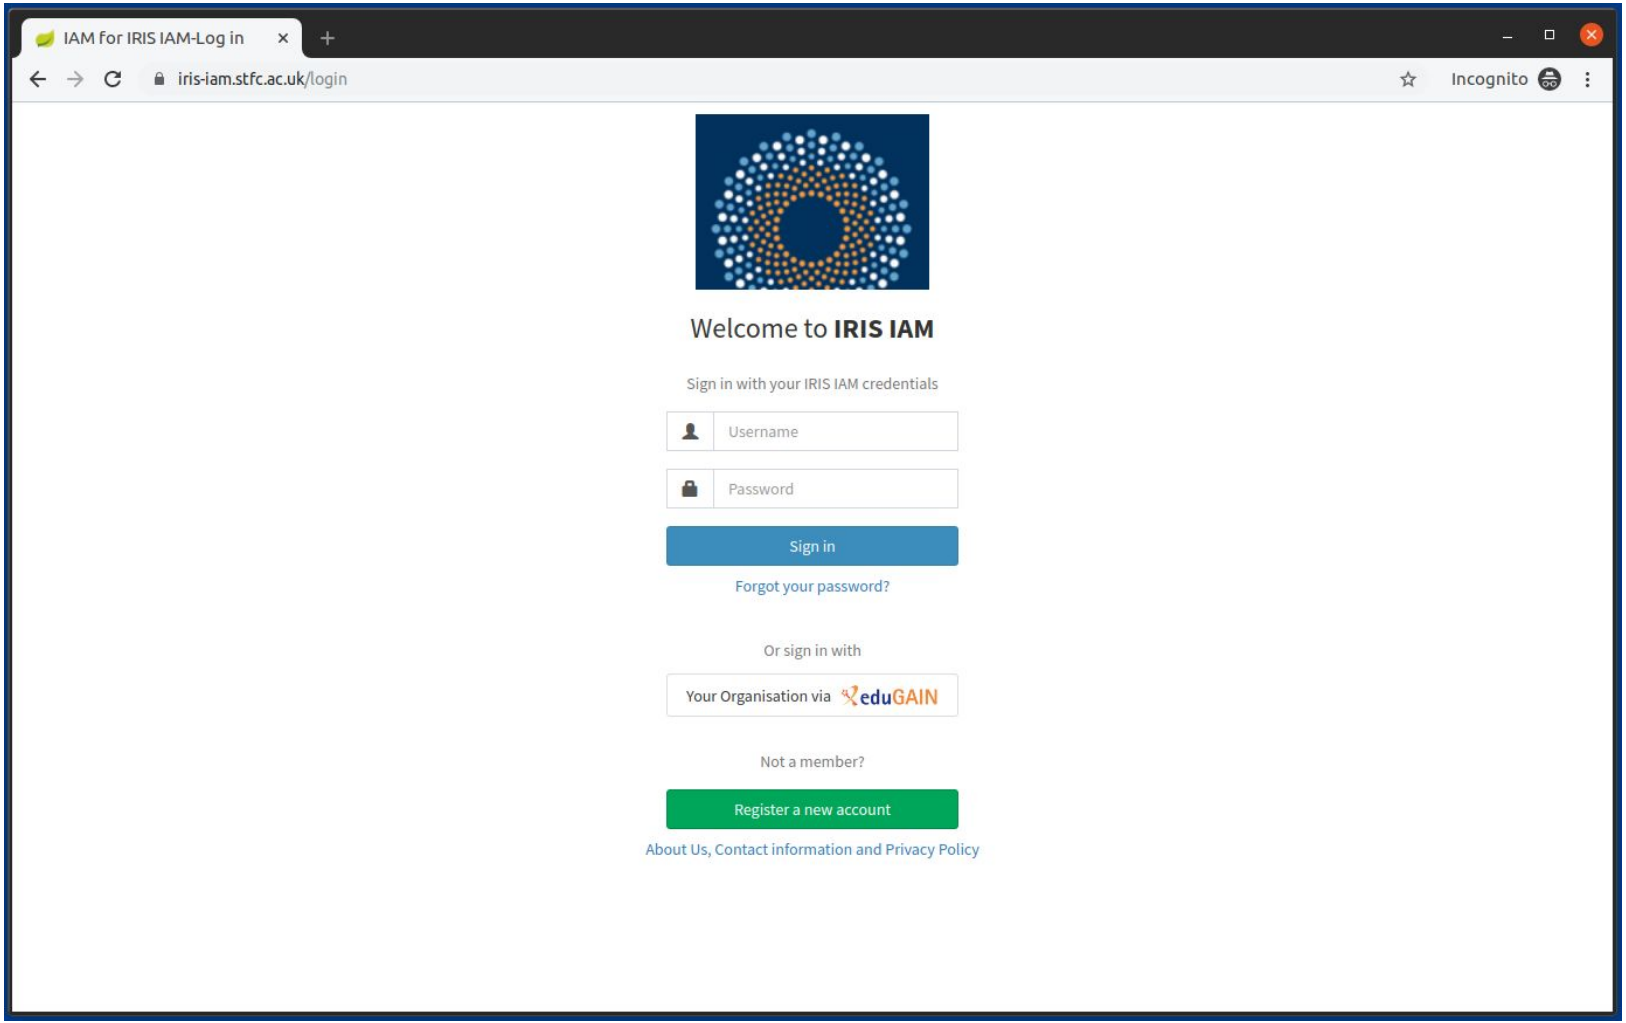
IAM for IRIS IAM-Log in
iris-iam.stfc.ac.uk/login
Incognito
Welcome to IRIS IAM
Sign in with your IRIS IAM credentials
Username
Password
Sign in
Forgot your password?
Or sign in with
Your Organisation via eduGAIN
Not a member?
Register a new account
About Us, Contact information and Privacy Policy

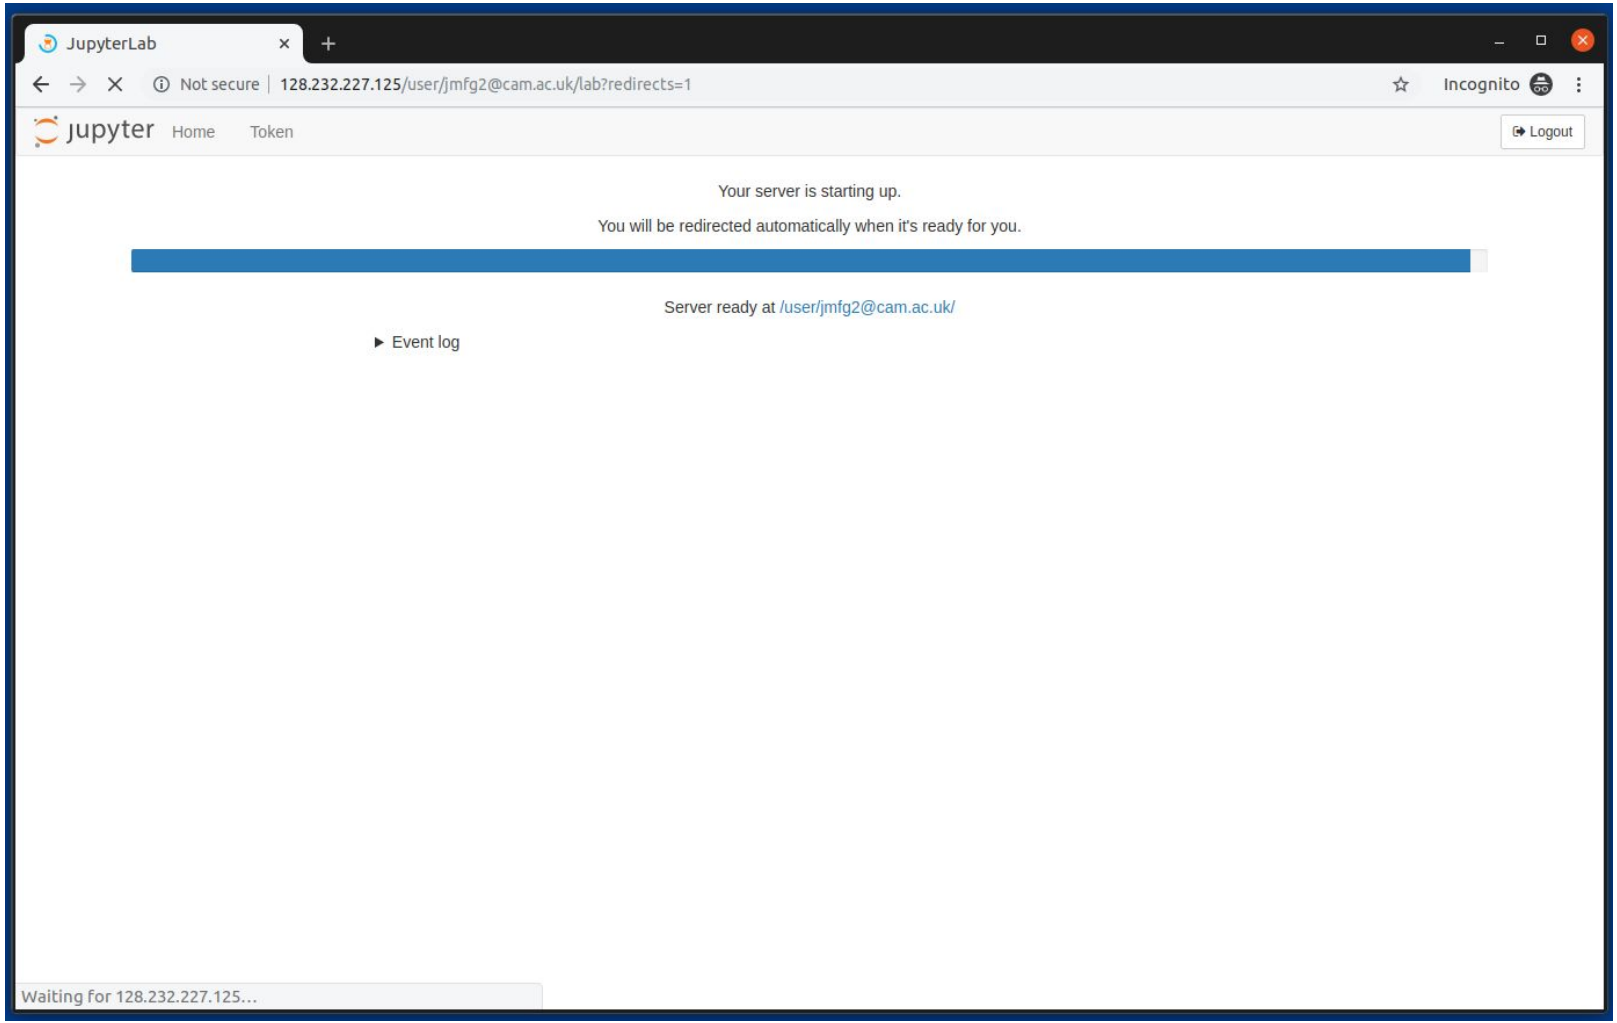
JupyterLab
Not secure | 128.232.227.125/user/jmfg2@cam.ac.uk/lab?redirects=1
Incognito
jupyter
Home
Token
Logout
Your server is starting up.
You will be redirected automatically when it's ready for you.
Server ready at /user/jmfg2@cam.ac.uk/
Event log
Waiting for 128.232.227.125...

JupyterLab

Not secure | 128.232.227.125/user/jmfg2@cam.ac.uk/lab?

File Edit View Run Kernel Hub Tabs Settings Help

Name	Last Modified
lost+found	an hour ago
dask-array.ipynb	13 minutes ago
dask.yaml	16 minutes ago

dask-array.ipynb

Markdown

Python 3

Dask Array scales NumPy

Dask array implements the NumPy API. It is composed of many small NumPy arrays

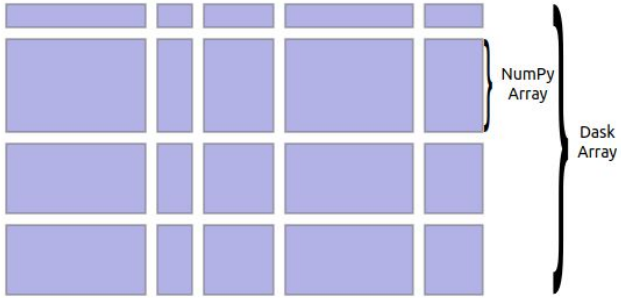

[1]:
```
from dask_kubernetes import KubeCluster
cluster = KubeCluster(n_workers=2)
cluster
```

Error creating widget: could not find model

[2]:
```
from dask.distributed import Client
client = Client(cluster)
client
```

[2]:

Client	Cluster
Scheduler: tcp://192.168.3.6:37753	**Workers:** 0
Dashboard: /user/jmfg2@cam.ac.uk/proxy/8787/status	**Cores:** 0
	Memory: 0 B

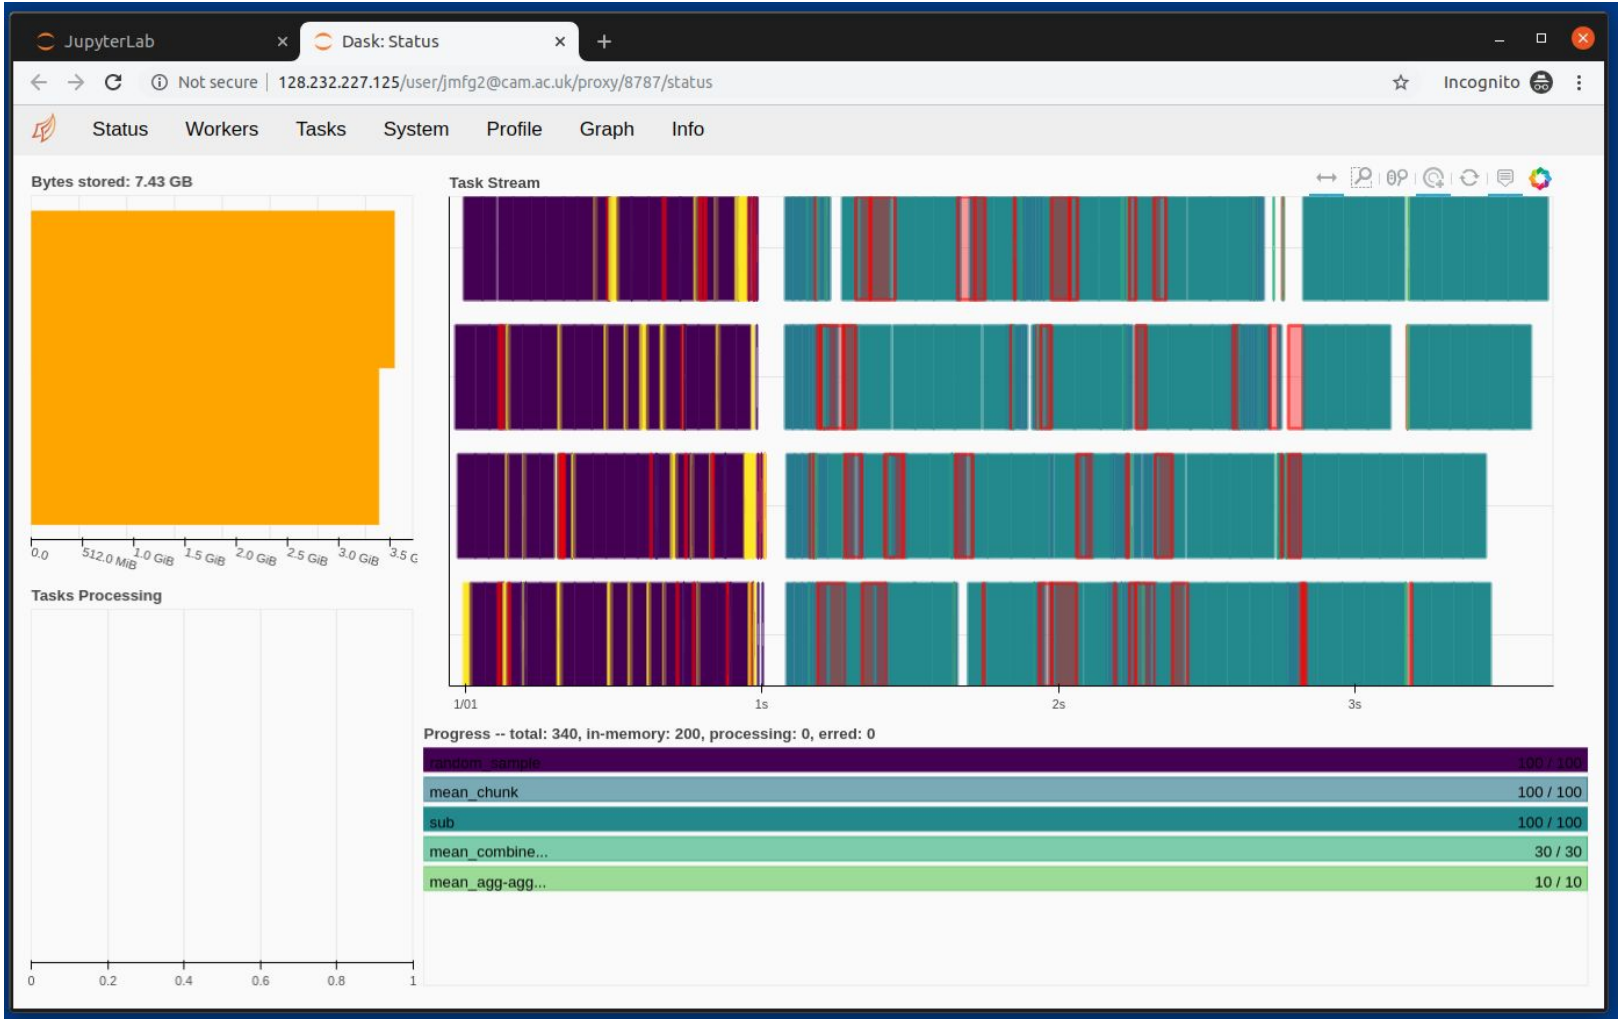
JupyterLab
Dask: Status
Not secure | 128.232.227.125/user/jmfg2@cam.ac.uk/proxy/8787/status
Incognito
Status
Workers
Tasks
System
Profile
Graph
Info
Bytes stored: 7.43 GB
0.0
512.0 MiB
1.0 GiB
1.5 GiB
2.0 GiB
2.5 GiB
3.0 GiB
Task Stream
1/01
1s
2s
3s
Tasks Processing
0
0.2
0.4
0.6
0.8
1
Progress -- total: 340, in-memory: 200, processing: 0, erred: 0
random_sample
100 / 100
mean_chunk
100 / 100
sub
100 / 100
mean_combine...
30 / 30
mean_agg-agg...
10 / 10

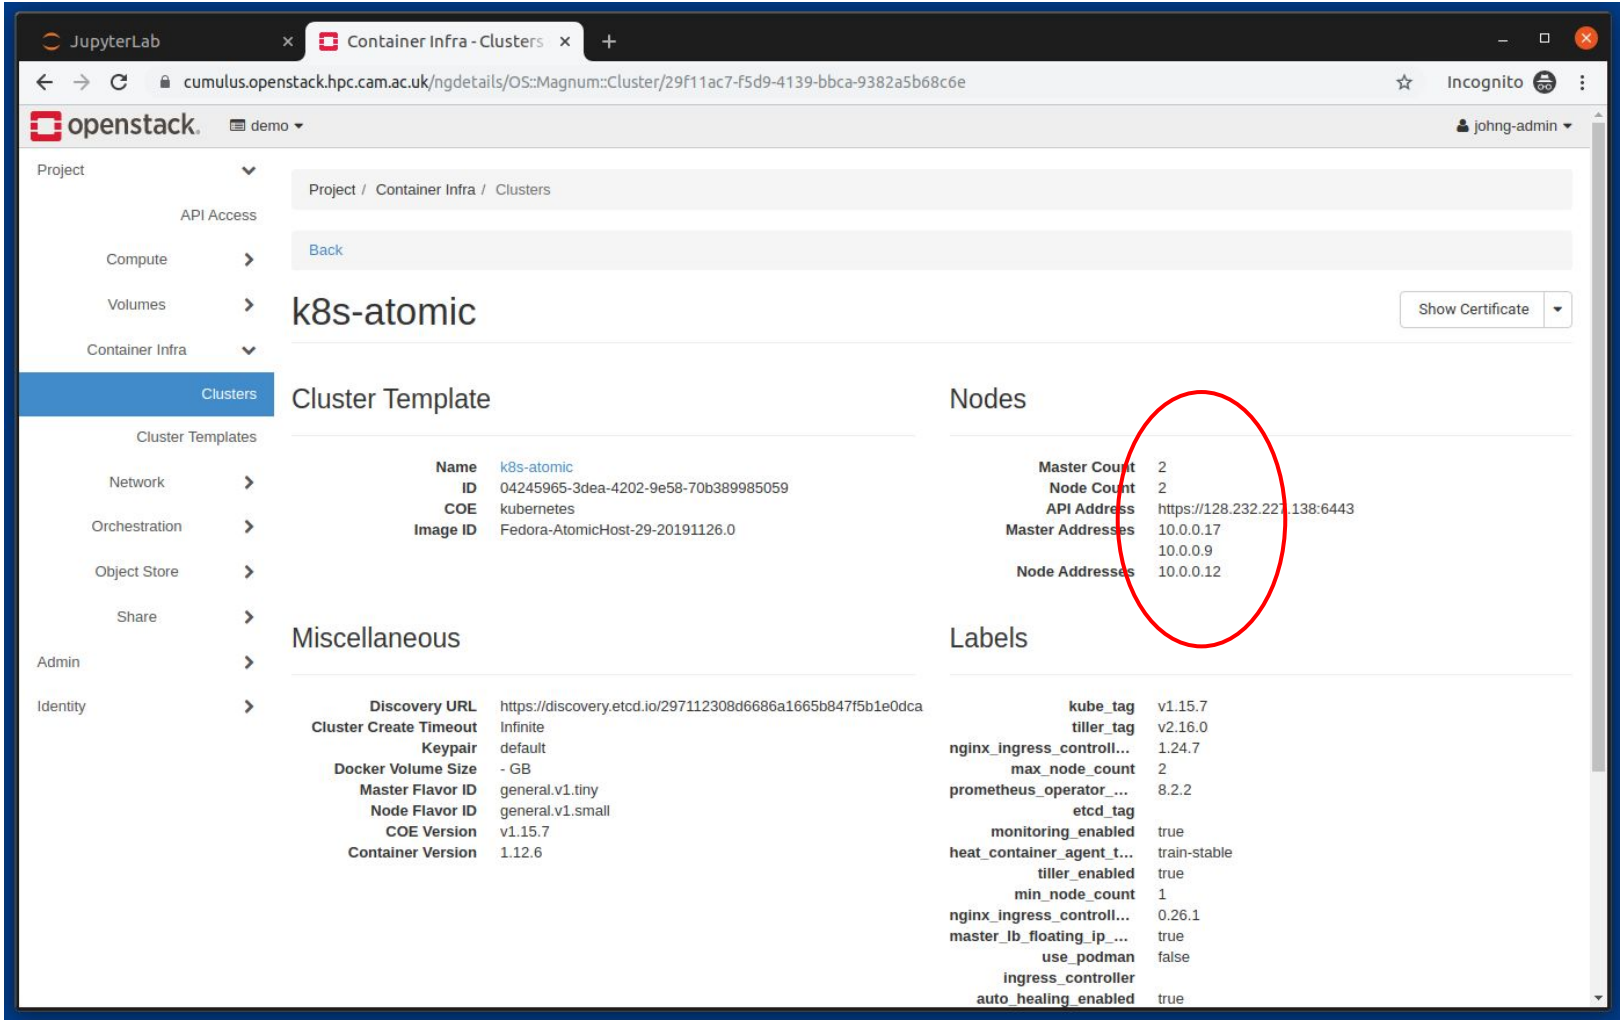
JupyterLab
Container Infra - Clusters
cumulus.openstack.hpc.cam.ac.uk/ngdetails/OS::Magnum::Cluster/29f11ac7-f5d9-4139-bbca-9382a5b68c6e
Incognito
openstack.
demo
johng-admin
Project
API Access
Compute
Volumes
Container Infra
Clusters
Cluster Templates
Network
Orchestration
Object Store
Share
Admin
Identity
Project / Container Infra / Clusters
Back
k8s-atomic
Show Certificate
Cluster Template
Name k8s-atomic
ID 04245965-3dea-4202-9e58-70b389985059
COE kubernetes
Image ID Fedora-AtomicHost-29-20191126.0
Nodes
Master Count 2
Node Count 2
API Address https://128.232.221.138:6443
Master Addresses 10.0.0.17
10.0.0.9
Node Addresses 10.0.0.12
Miscellaneous
Discovery URL https://discovery.etcd.io/297112308d6686a1665b847f5b1e0dca
Cluster Create Timeout Infinite
Keypair default
Docker Volume Size - GB
Master Flavor ID general.v1.tiny
Node Flavor ID general.v1.small
COE Version v1.15.7
Container Version 1.12.6
Labels
kube_tag v1.15.7
tiller_tag v2.16.0
nginx_ingress_controll... 1.24.7
max_node_count 2
prometheus_operator_... 8.2.2
etcd_tag
monitoring_enabled true
heat_container_agent_t... train-stable
tiller_enabled true
min_node_count 1
nginx_ingress_controll... 0.26.1
master_lb_floating_ip_... true
use_podman false
ingress_controller
auto_healing_enabled true

```
I1216 15:53:27.616349       1 leaderelection.go:227] successfully acquired lease kube-system/cluster-autoscaler
I1216 15:53:30.912732       1 magnum_manager_heat.go:293] For stack ID f864efb5-ceec-45b1-a660-a987f482d0ac, stack name is k8s-atomic-j7g2o2g55jsz
I1216 15:53:31.111256       1 magnum_manager_heat.go:310] Found nested kube_minions stack: name k8s-atomic-j7g2o2g55jsz-kube_minions-pdtyqdas7jho, ID bae8
6ecf-7b61-446f-830e-a30cf9b0e01d
I1216 16:29:07.348067       1 scale_up.go:689] Scale-up: setting group default-worker size to 2
I1216 16:29:12.250370       1 magnum_nodegroup.go:101] Increasing size by 1, 1->2
I1216 16:29:48.614285       1 magnum_nodegroup.go:67] Waited for cluster UPDATE_IN_PROGRESS status
I1216 16:33:34.949178       1 magnum_nodegroup.go:67] Waited for cluster UPDATE_COMPLETE status
I1216 17:16:27.410946       1 scale_down.go:882] Scale-down: removing empty node k8s-atomic-j7g2o2g55jsz-node-1
I1216 17:16:34.387762       1 magnum_manager_heat.go:344] Resolved node k8s-atomic-j7g2o2g55jsz-node-1 to stack index 1
I1216 17:16:35.804520       1 magnum_manager_heat.go:280] Waited for stack UPDATE_IN_PROGRESS status
I1216 17:17:06.289528       1 magnum_manager_heat.go:280] Waited for stack UPDATE_COMPLETE status
```

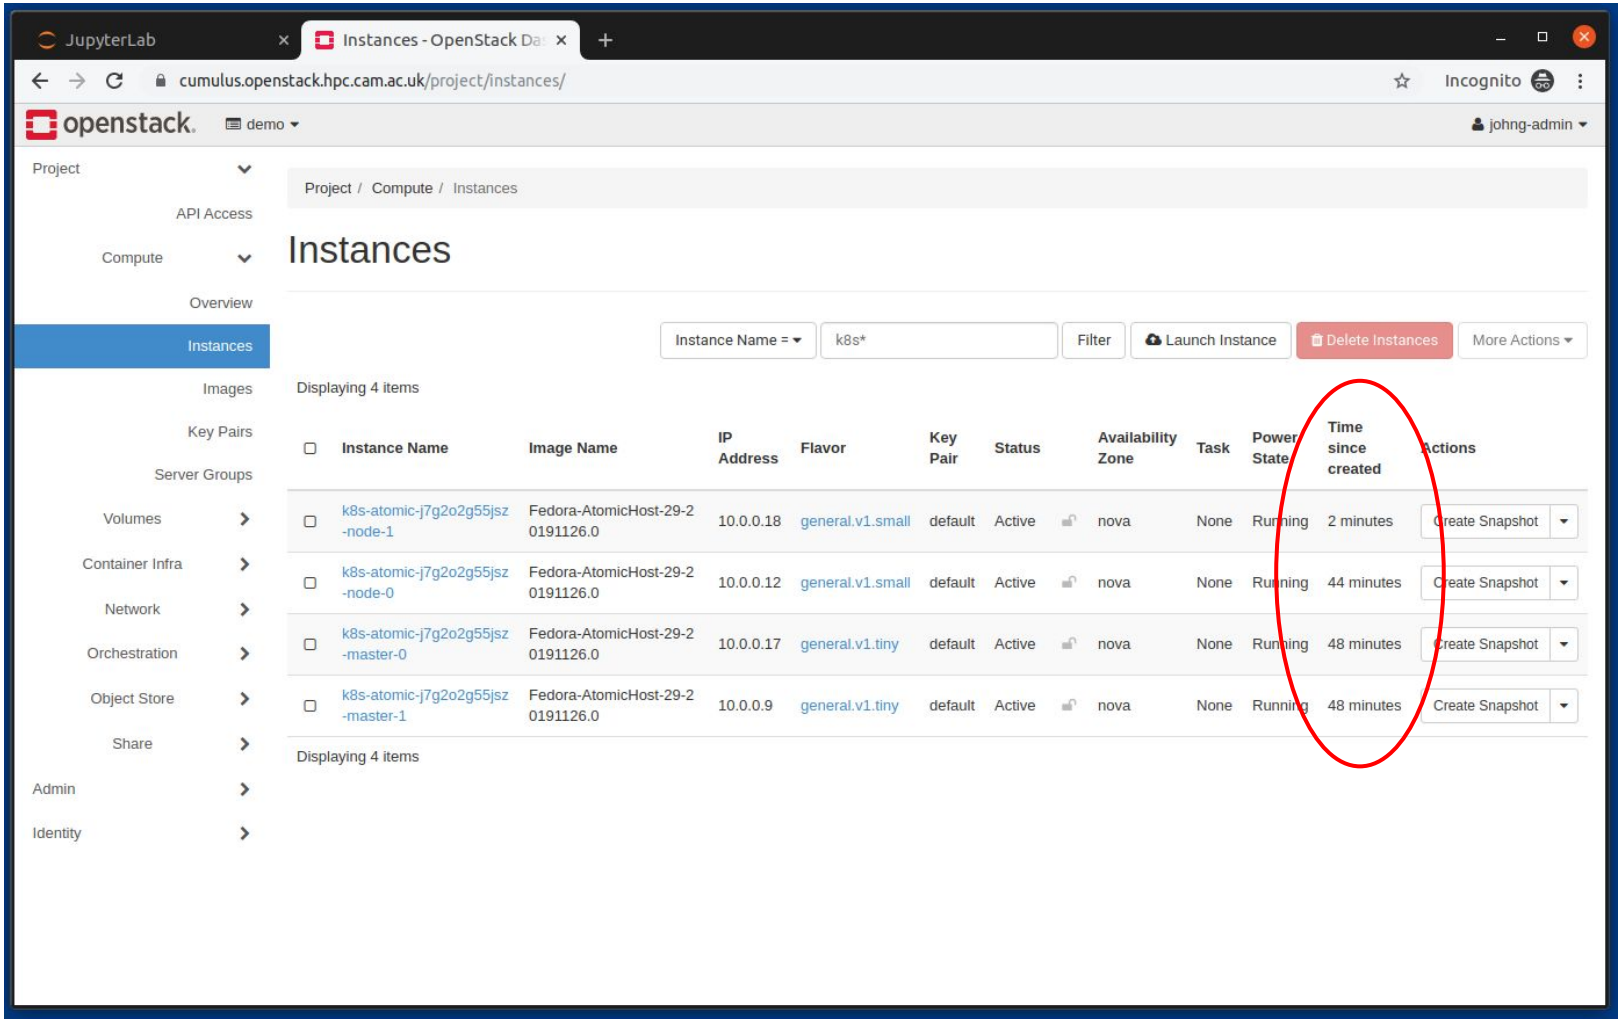
JupyterLab
Instances - OpenStack Da
cumulus.openstack.hpc.cam.ac.uk/project/instances/
Incognito
openstack.
demo
johng-admin
Project
API Access
Compute
Overview
Instances
Images
Key Pairs
Server Groups
Volumes
Container Infra
Network
Orchestration
Object Store
Share
Admin
Identity
Project / Compute / Instances
Instances
Instance Name =
k8s*
Filter
Launch Instance
Delete Instances
More Actions
Displaying 4 items
Instance Name
Image Name
IP Address
Flavor
Key Pair
Status
Availability Zone
Task
Power State
Time since created
Actions
k8s-atomic-j7g2o2g55jsz-node-1
Fedora-AtomicHost-29-20191126.0
10.0.0.18
general.v1.small
default
Active
nova
None
Running
2 minutes
Create Snapshot
k8s-atomic-j7g2o2g55jsz-node-0
Fedora-AtomicHost-29-20191126.0
10.0.0.12
general.v1.small
default
Active
nova
None
Running
44 minutes
Create Snapshot
k8s-atomic-j7g2o2g55jsz-master-0
Fedora-AtomicHost-29-20191126.0
10.0.0.17
general.v1.tiny
default
Active
nova
None
Running
48 minutes
Create Snapshot
k8s-atomic-j7g2o2g55jsz-master-1
Fedora-AtomicHost-29-20191126.0
10.0.0.9
general.v1.tiny
default
Active
nova
None
Running
48 minutes
Create Snapshot
Displaying 4 items

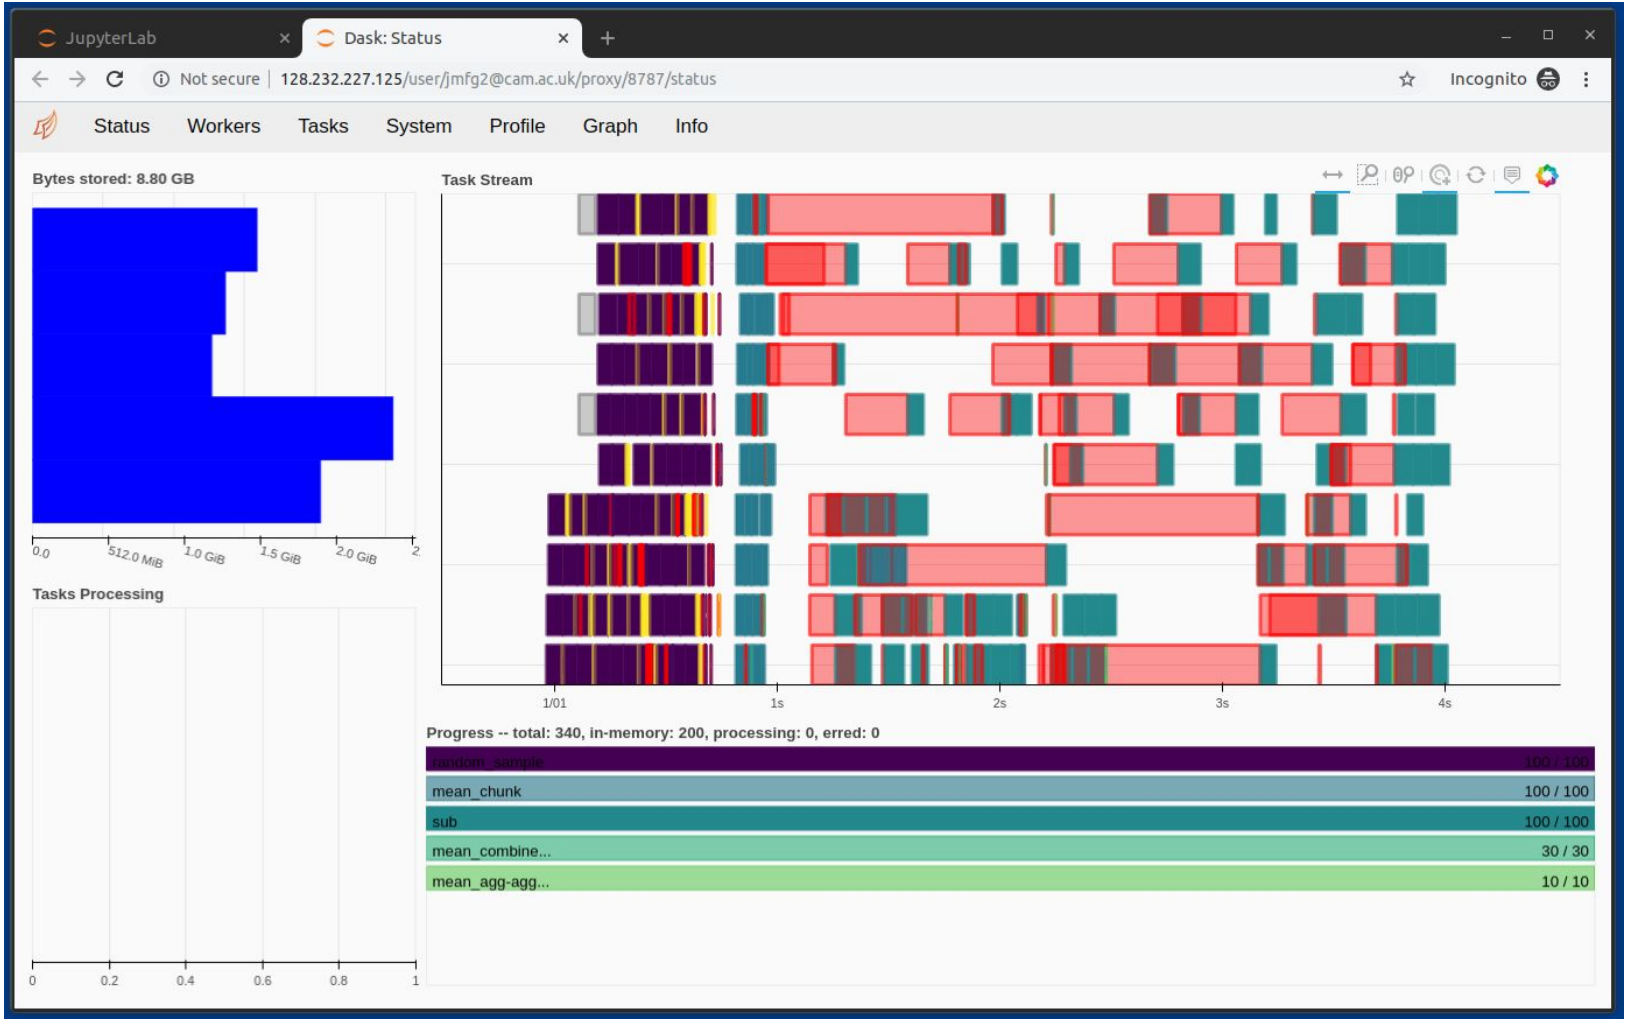
JupyterLab
Dask: Status
Not secure
128.232.227.125/user/jmfg2@cam.ac.uk/proxy/8787/status
Incognito
Status
Workers
Tasks
System
Profile
Graph
Info
Bytes stored: 8.80 GB
0.0
512.0 MiB
1.0 GiB
1.5 GiB
2.0 GiB
Task Stream
1/01
1s
2s
3s
4s
Tasks Processing
0
0.2
0.4
0.6
0.8
1
Progress -- total: 340, in-memory: 200, processing: 0, erred: 0
random_sample
100 / 100
mean_chunk
100 / 100
sub
100 / 100
mean_combine...
30 / 30
mean_agg-agg...
10 / 10

Plans for FY20

StackHPC

Plans for FY2020

- Clusters with multiple node types
- GPU support
- Upgrading kubernetes in-place
- More Magnum automation of common setups

@johnthetubaguy
@stackhpc
johng@stackhpc.com

StackHPC