


IRIS Scientific OpenStack Digital Asset Update

John Garbutt
December 2019

StackHPC

Scientific OpenStack

StackHPC



Scientists



Platform Ops



Infrastructure Ops

Scientific OpenStack FY19 Digital Assets

- Building on FY18 assets
- Driven by Science Communities needs
- Provides Reference Platforms
- Reference OpenStack Architecture and Configuration tuned for Scientific Computing
- Tooling to help Operate OpenStack

StackHPC



ANSIBLE



openstack®



ceph



kubernetes

Development Approach

StackHPC



FY19 Deliverables

StackHPC

1. OpenStack

- WP1.1 CI/CD of OpenStack Train
- WP1.2 Improved Monitoring

2. Efficient Resource Usage

- WP2.1 Blazar, WP2.2 Backfill best practices
- WP2.6 (stretch) Low Latency Ethernet

3. Platforms

- WP3.1 Octavia
- WP3.2 Slurm as a Service for Euclid
- WP3.3 Kubernetes, (Stretch) Spark on Kubernetes

4. Federation

- WP4.1 APEL
- WP4.2 IRIS IAM



openstack®



ceph



kubernetes

OpenStack Digital Assets

StackHPC

OpenStack Progress Highlights

StackHPC

- Being used for IRIS Science, since April 2019
- Operations by Cambridge HPCS
- User onboarding and automation scripts
- DiRAC Slurm on OpenStack Ironic

- Updates: Octavia and Magnum, Stein soon, Train planned
- Federation: APEL, IRIS IAM (in pre-prod)
- CI/CD: Tempest, (OnBoarding, Containers, Config, Kayobe)
- Monitoring and Alert improvements (in development)

Additional site volunteers?



openstack®

Baremetal via Ironic

StackHPC

- Maximum Performance
- Latency sensitive, e.g. MPI
 - RDMA Ethernet, RoCEv2 or iWARP
 - Dataset larger than single node's memory
 - SR-IOV is a possible alternative
- Trust issues around direct access to hardware
 - Cleaning is already supported
 - Can be avoided by providing a “Managed” service
- Optionally used by Kayobe for Server Lifecycle Management



IRONIC

an OpenStack Community Project

Platform Types

StackHPC

OpenStack Server

- Terraform creates Infrastructure
 - Use base OS image
 - No difference for Baremetal vs VMs
- Ansible modifies base OS to deploy Platform stack and Monitoring stack



Kubernetes

- Terraform creates K8s cluster
 - Manila CSI, cluster-autoscaler, Octavia Ingress, Prometheus, Grafana
- Ansible deploys apps via Helm, Kustomize



Platforms:
OpenHPC Slurm

StackHPC

OpenHPC Slurm for Euclid

- Powering real Science
- Operations in the hands of Euclid team
- Software VPN: RAL and Cambridge via Edinburgh
- Multi-Site Hyperconverged CephFS

Future Ideas:

- Access to Lustre File System
- Hardware offload of VPN

StackHPC



Generic OpenHPC Slurm

StackHPC

- Open OnDemand via IRIS IAM
- Hybrid: Terraform instead of OpenStack Heat
- Monitoring with Prometheus

- Lustre inside OpenStack, designed
- Exploring AutoScale (image based deploy) and Updates
- Exploring DiRAC via Ironic convergence



Example: Terraform for OpenHPC

StackHPC

```
provider "openstack" {  
  cloud = "cumulus"  
}
```

```
resource "openstack_compute_instance_v2" "login" {  
  name           = "ohpc-login"  
  image_name     = "CentOS7-1907"  
  flavor_name    = "general.v1.tiny"  
  key_pair       = "johng"  
  security_groups = ["default"]  
  
  network {  
    name = "cumulus-internal"  
  }  
}
```

```
resource "openstack_compute_instance_v2" "comp" {  
  name           = "ohpc-compute-${count.index}"  
  image_name     = "CentOS7-1907"  
  flavor_name    = "general.v1.medium"  
  key_pair       = "johng"  
  security_groups = ["default"]  
  count          = 5  
  
  network {  
    name = "cumulus-internal"  
  }  
}
```



Demo:
IRIS IAM and Slurm



Welcome to **IRIS IAM**

Sign in with your IRIS IAM credentials

[Forgot your password?](#)

Or sign in with

Not a member?



OnDemand provides an integrated, single access point for all of your HPC resources.

Message of the Day

powered by



OnDemand version: v1.6.20 | Dashboard version: v1.35.3

Last login: Tue Dec 10 17:46:09 2019 from 128-232-227-155.vss.cloud.cam.ac.uk

[jmfg2@ohpc-login ~]\$ sinfo

PARTITION	AVAIL	TIMELIMIT	NODES	STATE	NODELIST
compute*	up	1-00:00:00	1	idle	ohpc-compute-0

[jmfg2@ohpc-login ~]\$ squeue

JOBID	PARTITION	NAME	USER	ST	TIME	NODES	NODELIST(Reason)
-------	-----------	------	------	----	------	-------	------------------

[jmfg2@ohpc-login ~]\$ cat ~/.ssh/authorized_keys

```
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQACxvbnNzxtmEQszJEfKKLcLh2Q/HNAb1jcc0Za7N8HVh0LJhwn15SFcPzDC1e/0njo0pKa/3F+Wedz  
0/wra0Z1WLxX7FMvYX40Fgb10xWc6VuniAb+Fil1mysjW3ft8nhMQ6Q8TNY50q+lgtx4rxsr8fIH8AEkyX0lqAcgr2J5dSH4bmPGMaHdegiaKMuDsl  
vf6BynshHbwbmIY+1cWzQAVVZI3oipi2d3a7zZ2oCU+zHbzj7vFCT0PaYzjmk5PqKABAb9bR1Np6pGMH7GvCLcK17kpaG+0XgmSD/TV3LBmkVMWMTu/  
h+gldMCEm/wN+L/iPYJWwxWgM7ZdeUIfLed ansible-generated on ohpc-login.novalocal
```

```
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQCa1SHgtJax630Bvy8UUEIVUgh91PSIH5wmngI1prY16+rV3ydQg8xY1AqdjYc8PBjw+UXbdCuwUUd  
APDAtZ+1nehHj33/anfi+TxG7cdJf1J5ItwvwlkwmkqkyxJ/iQu2yDyYmqxiyk1J05nJYBcLn24V2yqn8MzkSk8VIUZ65vKJmVQ/lz2fHNbXhDIx2sB  
3D5p1qhu0J9GjUcPcIcA2aeWTGNjQoPsS4TD75aimxiv/Xzw20Y+QTsmsfw3XVZJwlaCijg8H1gG5Hx/0XKIVU8011ZP+aHoy/7GiBHosMvV8X7Pnq3  
jJp+DSDou5CMaBZmDGLDHn9xdtGaUmmfz3T ansible-generated on ohpc-login.novalocal
```

[jmfg2@ohpc-login ~]\$

Interactive Apps

Desktops

Secure Remote Desktop

Secure Remote Desktop (3)

1 node | 1 core | Running

Host: >_ohpc-compute-0

Delete

Created at: 2019-12-16 12:17:42 UTC

Time Remaining: 7 hours and 59 minutes

Session ID: 3c0302ae-5057-46be-b63d-f62f60697716

Compression



0 (low) to 9 (high)


Image Quality



0 (low) to 9 (high)

Launch Secure Remote Desktop

View Only (Share-able Link)



Your Jobs ▾

All Clusters ▾

Active Jobs

Show entriesFilter:

ID	Name	User	Account	Time Used	Queue	Status	Cluster
> 3	sys/dashboard/sys/bc_desktop/ohpc	jmfg2	root	00:02:30	compute	Running	ohpc

Showing 1 to 1 of 1 entries

Previous **1** Next

Platforms: Kubernetes

OpenStack Magnum

- Magnum and Kubernetes updates
- OpenStack cloud provider
 - Storage: Cinder, Manila
 - Load Balancer: Octavia
 - Ingress to reduce Public IP usage
- Cluster Autoscaler
- Hybrid: Terraform to create Kubernetes cluster
- Working on Spark for GAIA
- Jupyter Hub and Dask on Kubernetes via Pangeo
- Planned: in-place upgrade, more updates and testing

StackHPC



Example: Terraform for K8s

StackHPC

```
provider "openstack" {
  cloud = "cumulus"
}

resource "openstack_containerinfra_clustertemplate_v1" "kubernetes_template" {
  name = "kubernetes-1.15.3"
}

resource "openstack_containerinfra_cluster_v1" "cluster" {
  name                = "my_test_k8s"
  cluster_template_id = "${openstack_containerinfra_clustertemplate_v1.kubernetes_template.id}"
  master_count       = 2
  node_count         = 2
  keypair             = "johng"
  labels              = {
    min_node_count = 1
    max_node_count = 4
  }
}
```



Demo:
Kubernetes AutoScale

Sign in with IRIS IAM



Welcome to IRIS IAM

Sign in with your IRIS IAM credentials

[Forgot your password?](#)

Or sign in with

Not a member?

[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/

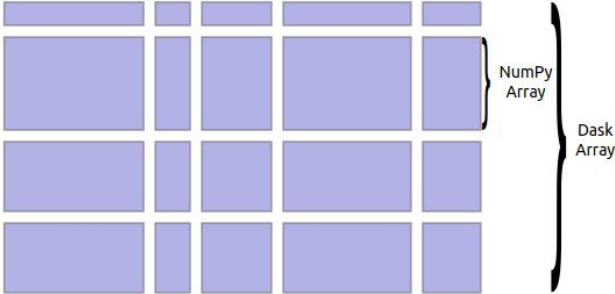
▶ Event log

Waiting for 128.232.227.125...

JupyterLab interface showing a file browser on the left and a code editor on the right. The file browser lists files: lost+found, dask-array.ipynb (13 minutes ago), and dask.yaml (16 minutes ago). The code editor displays the following content:

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


Bytes stored: 7.43 GB



Tasks Processing



Task Stream



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 x +

cumulus.openstack.hpc.cam.ac.uk/ngdetails/OS::Magnum::Cluster/29f11ac7-f5d9-4139-bbca-9382a5b68c6e Incognito johng-admin

openstack demo

Project / Container Infra / Clusters

API Access

Compute > Back

Volumes >

Container Infra >

k8s-atomic

Show Certificate

Clusters

Cluster Templates

Network >

Orchestration >

Object Store >

Share >

Admin >

Identity >

Cluster Template

Name	k8s-atomic
ID	04245965-3dea-4202-9e58-70b389985059
COE	kubernetes
Image ID	Fedora-AtomicHost-29-20191126.0

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

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

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 ledelection.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 bae86ecf-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 Dashboard

cumulus.openstack.hpc.cam.ac.uk/project/instances/

openstack demo johng-admin

Project / Compute / Instances

Instances


Instance Name = k8s* Filter Launch Instance Delete Instances More Actions

Displaying 4 items


<input type="checkbox"/>	Instance Name	Image Name	IP Address	Flavor	Key Pair	Status	Availability Zone	Task	Power State	Time since created	Actions
<input type="checkbox"/>	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
<input type="checkbox"/>	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
<input type="checkbox"/>	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
<input type="checkbox"/>	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


Bytes stored: 8.80 GB



Tasks Processing



Task Stream



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

Better OpenStack Resource Management

OpenStack Resource Management

StackHPC

- Quota = max concurrent usage
- Allocation = fair share of usage, averaged over one year
- Quota can divide resources between projects
- Resources per project, not per group of projects
- OpenStack generally hides current capacity
- First to take resources wins




openstack®

Resource Management

StackHPC

- Aims
 - Maximise Science done on given resources
 - While maintaining a fare share inline with allocation
- Use Cases
 - GridPP Backfill
 - Autoscaling Platforms (Kubernetes, Slurm)
 - Schedule re-processing runs
 - Known time slot resource needs



Problems with Backfilling?

- Making space for Backfill
 - Isolated platforms help loose coupling, but can be costly
 - Stop VMs you don't need
 - Kubernetes cluster autoscaler, Slurm elastic compute
- Taking back resources from Backfill
 - Retry until you get lucky?
 - External system manage project quotas and wait for drain?
 - External system delete “preemptible” instances?
 - ... or something else? Reservations?



OpenStack Blazar Reservation

StackHPC

- OpenStack Blazar manages a subset of nodes
 - Adds a time axis to OpenStack placement
- Users can request a reservation of specific size and time
- Negotiate when to do your “big run” with other projects

- Added support for Preemptibles in Blazar
- Independent to CERN efforts around preemptibles in Nova
- Future ideas:
 - Limit max CPU hours of reservations, track usages vs limit
 - Auto-shrink unused reservations, or require explicit claim
 - Specify constraints not specific size and time
 - Webhooks to automate usage on start



Demo:
Blazar Preemptibles

Blazar Preemptibles Demo

- Run on pre-prod at Cambridge
- One hypervisor is a Blazar host resource
 - Normally only used by Blazar reservations
- Successfully create 4 preemptibles
- Successfully create a Blazar reservation
- Triggers delete of preemptible instances
- Unsuccessfully try to create more preemptible instances
- Successfully create instance to use Blazar reservation

openstack. admin

Admin / Compute / All Hypervisors

All Hypervisors

Hypervisor Summary

VCPU Usage
Used 5 of 56

Memory Usage
Used 31GB of 381.4GB

Local Disk Usage
Used 122GB of 208GB

Hypervisor Compute Host

Displaying 2 items

Hostname	Type	VCPUs (used)	VCPUs (total)	RAM (used)	RAM (total)	Local Storage (used)	Local Storage (total)	Instances
cpu-e-1042	QEMU	0	28	9GB	190.7GB	39GB	104GB	0
cpu-e-1043	QEMU	5	28	22GB	190.7GB	83GB	104GB	3

Displaying 2 items

Hypervisor allocated to Blazar →

Currently unused →

Hypervisor allocated to Blazar

Currently unused

VCPUs (used)

- Project >
- Admin >
- Overview
- Compute >
- Hypervisors
- Host Aggregates
- Instances**
- Flavours
- Images
- Volume >
- Network >
- Container Infra >
- System >
- Identity >

Admin / Compute / Instances

Instances

**Preemptibles instances
are scheduled on the
Blazar host**

Project Name = Filter

Displaying 7 items

<input type="checkbox"/>	Project	Host	Name	Image Name	IP Address	Flavour	Status	Task	Power State	Time since created	Actions
<input type="checkbox"/>	preemptible	cpu-e-1042	preemptible-instance	CentOS7-1907	10.2.0.15	general.v1.tiny	Build	Spawning	No State	0 minutes	Edit Instance
<input type="checkbox"/>	preemptibles	cpu-e-1042	preemptible-instance	CentOS7-1907	10.2.0.9	general.v1.tiny	Build	Spawning	No State	0 minutes	Edit Instance
<input type="checkbox"/>	preemptibles	cpu-e-1042	preemptible-instance	CentOS7-1907	10.2.0.17	general.v1.tiny	Build	Spawning	No State	0 minutes	Edit Instance
<input type="checkbox"/>	preemptibles	cpu-e-1042	preemptible-instance	CentOS7-1907	10.2.0.6	general.v1.tiny	Build	Spawning	No State	0 minutes	Edit Instance

Host reservation

- Create a reservation
- Set it to start immediately

Create Lease

Lease Name *

Start Date ?

End Date ?

Resource Type *

Minimum Number of Hosts ?

Maximum Number of Hosts ?

Hypervisor Properties ?

Resource Properties ?

Description:
Create a lease with the provided values.

Reservation Request Accepted

StackHPC

Leases - OpenStack Dashboard

cumulus-dev.openstack.hpc.cam.ac.uk/project/leases/

openstack demo

Project / Reservations /

Success: Lease test-lease was successfully created.

Leases

+ Create Lease Delete Leases

Displaying 1 item

<input type="checkbox"/>	Lease name	Start date	End date	Status	Degraded	Actions
<input type="checkbox"/>	test-lease	2019-12-13 17:41 UTC	2019-12-14 17:41 UTC	PENDING	No	Update Lease

Displaying 1 item

Leases

Identity

Reservation is STARTING

StackHPC

The screenshot shows the OpenStack Dashboard interface. The browser address bar displays the URL `cumulus-dev.openstack.hpc.cam.ac.uk/project/leases/`. The dashboard header includes the OpenStack logo, a user profile for 'priteau', and a search bar. The left sidebar contains a navigation menu with categories like Project, API Access, Compute, Volumes, Container Infra, Network, Orchestration, DNS, Reservations, and Identity. The 'Reservations' category is expanded, and the 'Leases' sub-page is selected. The main content area shows the breadcrumb 'Project / Reservations / Leases' and the title 'Leases'. There are two buttons: '+ Create Lease' and 'Delete Leases'. Below the buttons, it says 'Displaying 1 item'. A table lists the lease details:

<input type="checkbox"/>	Lease name	Start date	End date	Status	Degraded	Actions
<input type="checkbox"/>	test-lease	2019-12-13 17:41 UTC	2019-12-14 17:41 UTC	STARTING	No	Update Lease ▾

Preemptible Instances Terminated





Instances

Project Name =

Filter

Delete Instances

Displaying 7 items

<input type="checkbox"/>	Project	Host	Name	Image Name	IP Address	Flavour	Status	Task	Power State	Time since created	Actions
<input type="checkbox"/>	preemptibles	cpu-e-1042	preemptible-instance	CentOS7-1907	10.2.0.15	general.v1.tiny	Active	 Deleting	Running	9 minutes	Soft Reboot Instance
<input type="checkbox"/>	preemptibles	cpu-e-1042	preemptible-instance	CentOS7-1907	10.2.0.9	general.v1.tiny	Active	 Deleting	Running	9 minutes	Soft Reboot Instance
<input type="checkbox"/>	preemptibles	cpu-e-1042	preemptible-instance	CentOS7-1907	10.2.0.17	general.v1.tiny	Active	 Deleting	Running	9 minutes	Soft Reboot Instance
<input type="checkbox"/>	preemptibles	cpu-e-1042	preemptible-instance	CentOS7-1907	10.2.0.6	general.v1.tiny	Active	 Deleting	Running	9 minutes	Soft Reboot Instance

No More Preemptible Instances

```
openstack_compute_instance_v2.preemptible[2]: Creating...
openstack_compute_instance_v2.preemptible[1]: Creating...
openstack_compute_instance_v2.preemptible[3]: Creating...
openstack_compute_instance_v2.preemptible[0]: Still creating... [10s elapsed]
openstack_compute_instance_v2.preemptible[1]: Still creating... [10s elapsed]
openstack_compute_instance_v2.preemptible[2]: Still creating... [10s elapsed]
openstack_compute_instance_v2.preemptible[3]: Still creating... [10s elapsed]

Error: Error waiting for instance (36217b21-a2b8-47b0-879b-f331c3f6a259) to become ready: unexpected state 'ERROR', wanted target 'ACTIVE'. 1
ast error: %!s(<nil>)

on preemptibles.tf line 6, in resource "openstack_compute_instance_v2" "preemptible":
  6: resource "openstack_compute_instance_v2" "preemptible" {

Error: Error waiting for instance (c668cc6f-b90c-47e6-a2a9-97ebdf3a4134) to become ready: unexpected state 'ERROR', wanted target 'ACTIVE'. 1
ast error: %!s(<nil>)

on preemptibles.tf line 6, in resource "openstack_compute_instance_v2" "preemptible":
  6: resource "openstack_compute_instance_v2" "preemptible" {

Error: Error waiting for instance (fc5850b0-87c6-451b-b0c4-537abd513143) to become ready: unexpected state 'ERROR', wanted target 'ACTIVE'. 1
ast error: %!s(<nil>)

on preemptibles.tf line 6, in resource "openstack_compute_instance_v2" "preemptible":
  6: resource "openstack_compute_instance_v2" "preemptible" {

Error: Error waiting for instance (5469d333-8e76-41f0-bd48-e1d001a4f07e) to become ready: unexpected state 'ERROR', wanted target 'ACTIVE'. 1
ast error: %!s(<nil>)

on preemptibles.tf line 6, in resource "openstack_compute_instance_v2" "preemptible":
  6: resource "openstack_compute_instance_v2" "preemptible" {

[0] 0:priteau@cumulus-seed:~/src/cumulus-kayobe-config- 1:priteau@cumulus-seed:~/preemptibles-demo*Z "cumulus-seed" 18:06 13-Dec-19
```

Reservation Owner Creates Instances

Instances

Instance ID = Filter Launch Instance Delete Instances More Actions

Displaying 3 items

<input type="checkbox"/>	Instance Name	Image Name	IP Address	Flavour	Key Pair	Status	Availability Zone	Task	Power State	Time since created	Actions
<input type="checkbox"/>	reserved-instances	CentOS7-1907	10.1.0.16	general.v1.tiny	cumulus-seed	Active	blazar_74bbcc9d-a65f-47fc-9d10-30ff152aeb6d	None	Running	0 minutes	Create Snapshot

OpenStack Blazar Future Ideas

StackHPC

- Enhancements to Preemptible support
- Option to use Reservations by default, with Magnum
- Limit max CPU hours of reservations
 - Track actual usages vs reservations vs limit
- Webhooks to automate usage on start and clean up on finish
- Specify constraints, rather than specific time and size
 - Shuffle future reservations based on new requests
- Auto-shrink unused reservations
 - Consider explicit claim, followed by return, followed by later reclaim
- API to query current preemptible capacity
 - Communicate preemptible capacity via quota



Summary:
Scientific OpenStack
FY19 Digital Asset

FY19 Deliverables

1. OpenStack

- WP1.1 CI/CD of OpenStack Train, WP1.2 Improved Monitoring

2. Efficient Resource Usage

- WP2.1 Blazar, WP2.2 Backfill best practices
- WP2.6 (stretch) Low Latency Ethernet

3. Platforms

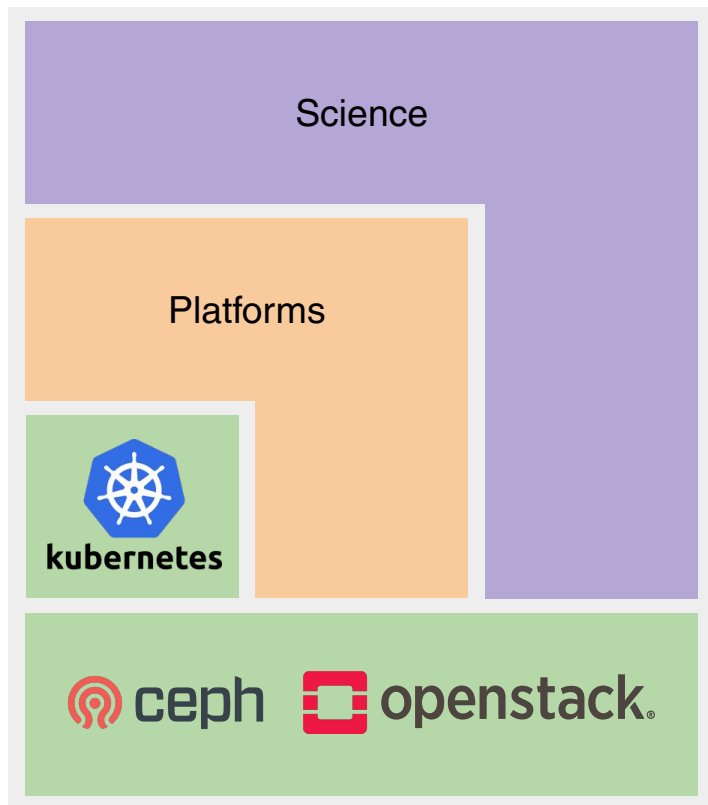
- WP3.1 Octavia
- WP3.2 Slurm as a Service for Euclid
- WP3.3 Kubernetes, (Stretch) Spark on Kubernetes

4. Federation

- WP4.1 APEL
- WP4.2 IRIS IAM

Scientific OpenStack

StackHPC



Scientists




Platform Ops



Infrastructure Ops

Scientific OpenStack

StackHPC



Scientists



Platform Ops



Infrastructure Ops

@stackhpc
@johnthetubaguy

StackHPC