

Scientific OpenStack FY21

OpenStack Blazar and
Azimuth Cloud Portal's On-Demand Platforms


February 2022

StackHPC

StackHPC Company Overview

StackHPC


- Formed 2016, based in Bristol, UK
 - Based in Bristol with presence in Cambridge, France and Poland
 - Currently around 20 people
- Founded on HPC expertise
 - Software Defined Networking
 - Systems Integration
 - OpenStack Development and Operations
- Motivation to transfer this expertise into Cloud to address HPC & HPDA (AI)
- “Open” Modus Operandi
 - Upstream development of OpenStack capability
 - Consultancy/Support to end-user organizations in managing HPC service transition
 - Scientific-WG engagement for the Open Infrastructure Foundation
- Hybrid Cloud Enablement



Scientific OpenStack IRIS Assets FY21

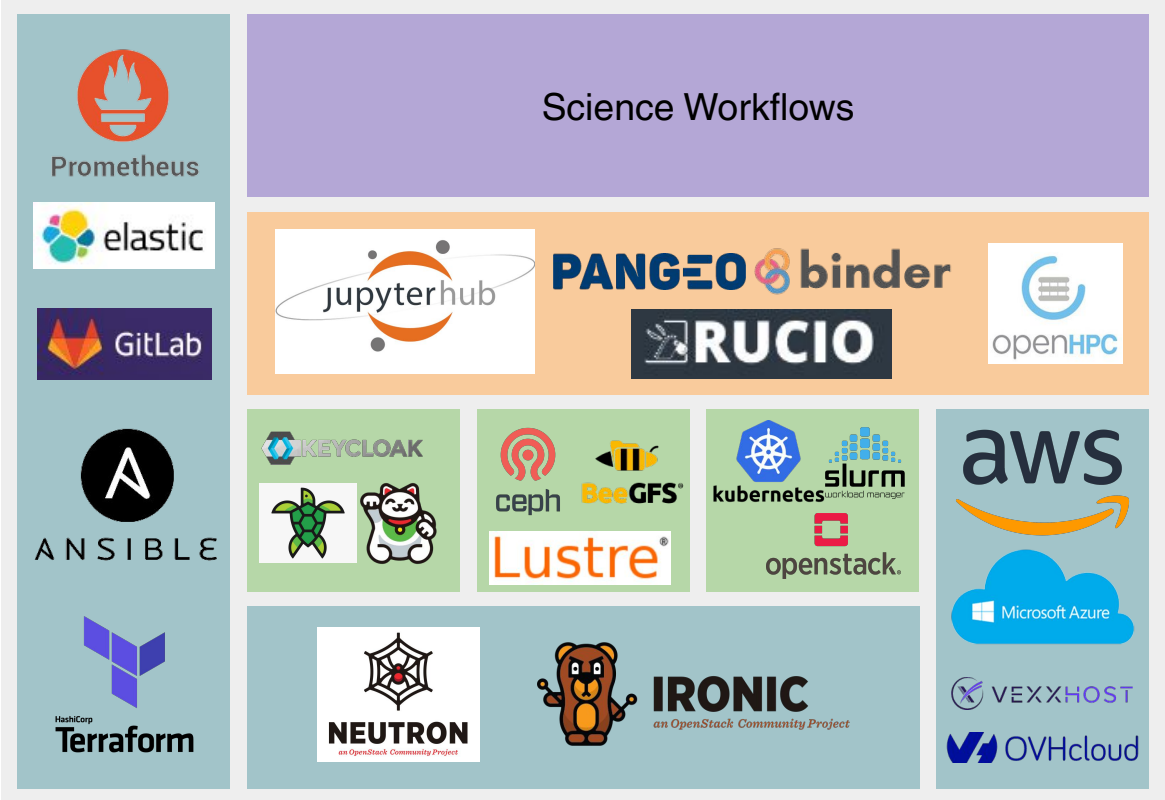
Scientific OpenStack

StackHPC



Integration can be a Challenge!

StackHPC



Previous IRIS Assets

- IRIS Production is largely based on FY19 assets (and Arcus a bit)
- IRIS Stage largely based on FY20 assets
 - Blazar Preemptible Instances, and easier to use Reservations
 - Multiple configurations in a single kayobe config branch
 - CI/CD Improvements, Monitoring Improvements (alerts and dashboards)
 - Ansible based server bootstrapping from UM6P
 - Scripts to deploy hypervisors on Ironic overcloud instead of the seed
 - Creation of Slurm appliance, supporting Updated platforms and Magnum templates
- Some parts of FY21 assets are ready for testing

IRIS Assets for FY2021

StackHPC

- IRIS-FY21A
 - OpenStack Configuration and Deployment (Wallaby on CentOS Stream 8, cephadm, OVN)
 - Better Resource Sharing (Monitor utilization, Improve Blazar operability and utilization)
 - Accelerated Infrastructure Management (RDMA and GPUs within K8s on VMs, NVMe in K8s)
 - Cloud Portal & K8s (OIDC, Portal CI, Slurm appliance autoscale, Magnum K8s, Grafana)
- IRIS-FY21B
 - WP1: Cloud Portal Enhancements (40%)
 - K8s Applications, Storage Import/Export, Secure Application Proxy, AI/ML Benchmarks
 - WP2: Lustre and NVMe powered Ceph file system and Ceph object store via Jupyter (20%)
 - WP3: Port Scientific OpenStack to Rocky Linux 8 / RHEL 8 (30%)
 - WP4: Centralised open access Scientific OpenStack repository and CI/CD platform (10%)

OpenStack Blazar

Before creating a Science Platform...

... you need **access** to infrastructure and a **resource allocation**

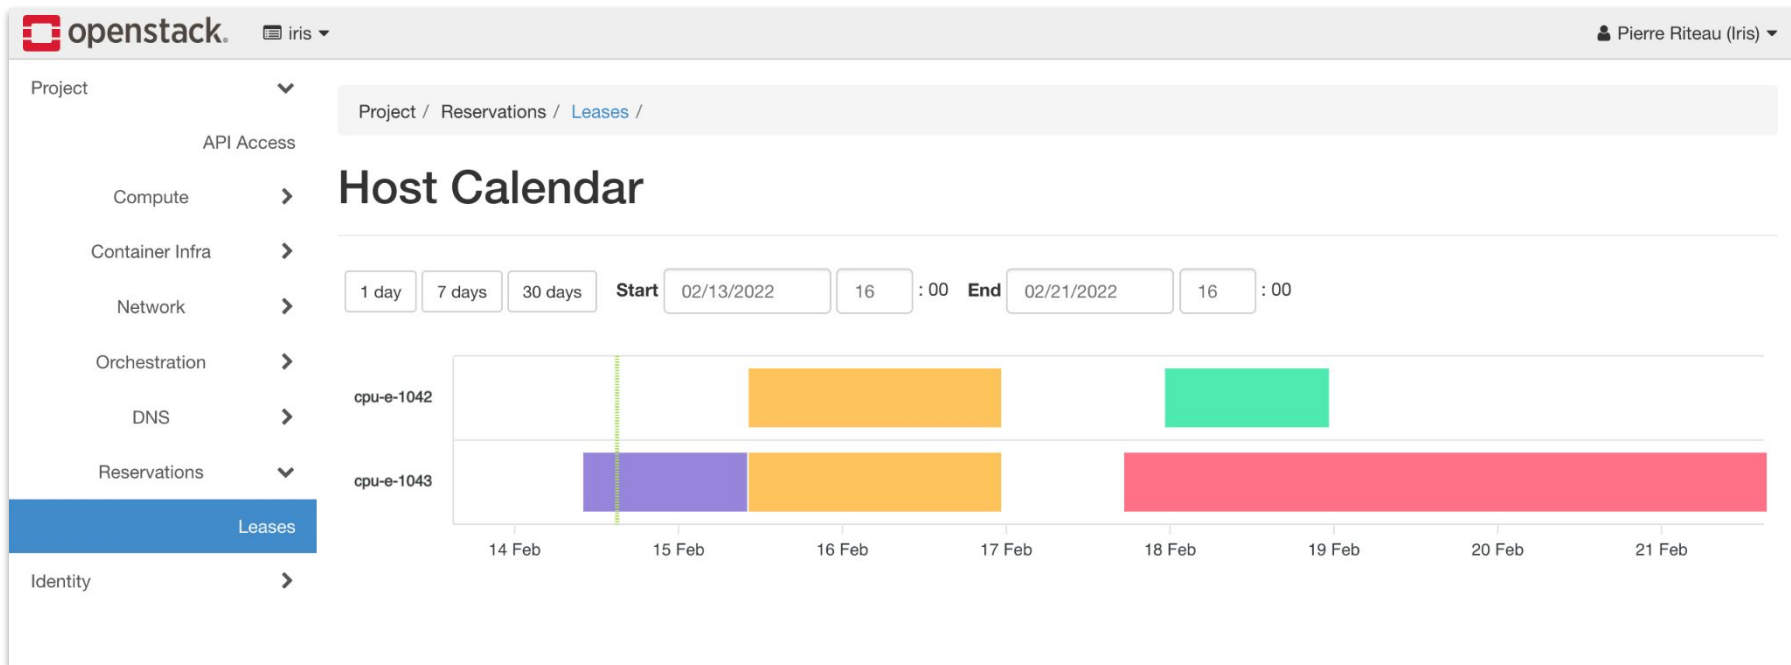
- Pick access method
 - GridPP, Slurm, IaaS (e.g. OpenStack)
- Understand what resources to request
- Request IRIS allocation
- IRIS IAM group membership
- Onboarded by allocated Infrastructure provider
 - Submit your job, or create your VM
 - ... “no valid host”!?

A detailed diorama of a coral reef ecosystem. The scene is set on a sandy seabed with various types of coral and shells. On the left, there is a large, green and white branching coral structure. In the center, there are several pink and orange branching corals. To the right, there are large, flat, green table corals. In the foreground, there are several shells, including a large, colorful scallop with red, green, and white edges, and a smaller, dark shell with a white spot. The background shows a blue water surface with more coral and a small, yellowish object. A dark, semi-transparent banner is overlaid across the center of the image, containing the text "The Coral Reef Cloud" in white.

The Coral Reef Cloud

OpenStack Blazar

StackHPC



OpenStack Blazar

StackHPC

openstack. Iris

Pierre Riteau (Iris)

Create Lease

Lease Name *

Description:
Create a lease with the provided values.

Start Date ⓘ
YYYY-MM-DD HH:MM (blank for now)

End Date ⓘ
YYYY-MM-DD HH:MM (blank for Start Date + 24h)

Resource Type *

Physical Host

Minimum Number of Hosts ⓘ

Maximum Number of Hosts ⓘ

Hypervisor Properties ⓘ

Resource Properties ⓘ

Cancel Create Lease

Host Calendar + Create Lease

Host Calendar + Create Lease


Actions

Using Blazar


- Allocation mapped to “cloud credit hours”, not Quota
- User creates a Reservation
 - Pick when and how much resources, depending on availability
 - Cloud credits are consumed when you create the reservation
 - You can't create a reservation without sufficient cloud credits
- Typically reserve a number of hosts
 - Scheduling implemented similarly to tenant isolation filters
 - ... does not require additional scheduler hints
- GPU reservation, ideally reserve 1/8th of an A100
 - Need more granular reservation than a full host
 - ... but, must create resources using a reservation specific private flavor
- Preemptible instances can use unreserved space
 - with minimum lifetime

Reservations and Preemptibles

StackHPC



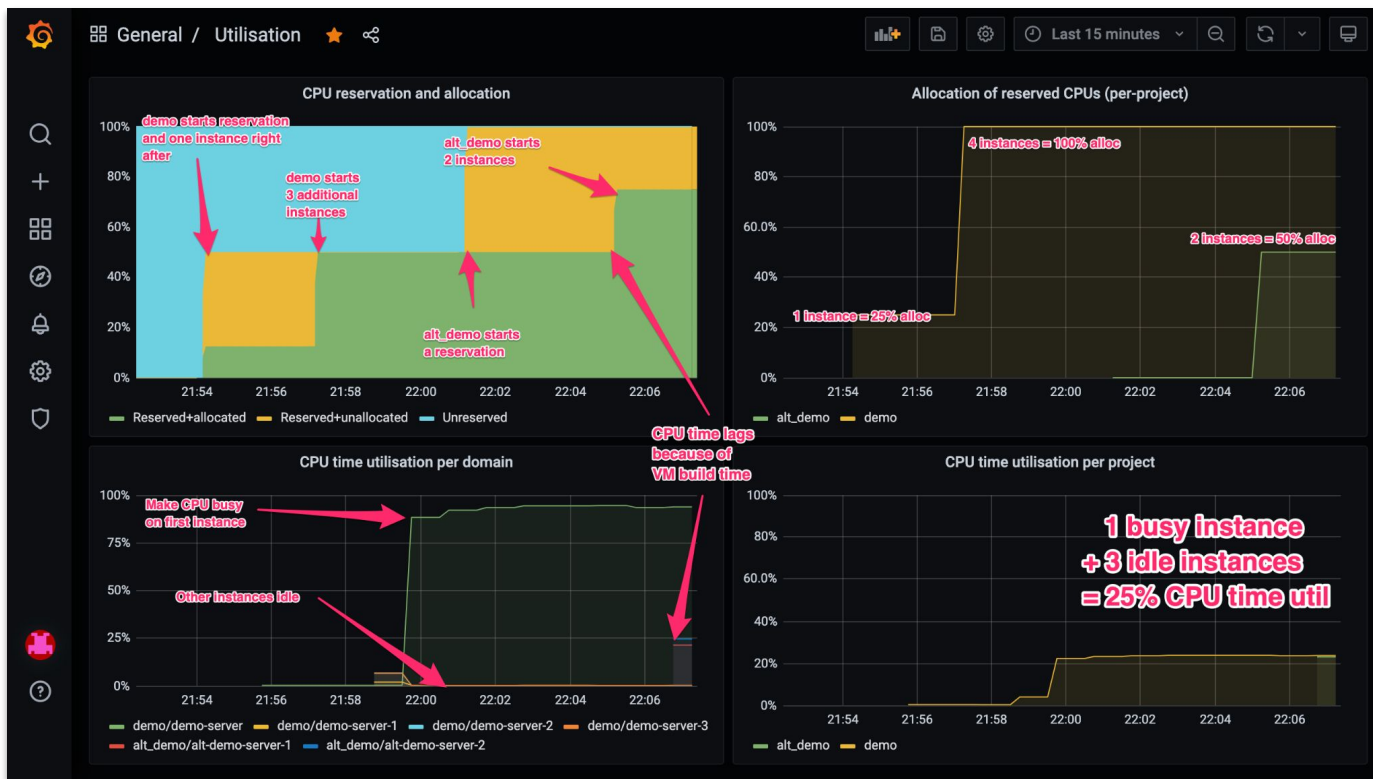
Reservations and Preemptibles



Planned Blazar testing at Cambridge


- Slurm with Preemptible VM partition
 - Large reservation for shared Slurm (or non-blazar managed servers)
 - Some usage of preemptibles, as backfill for unused Blazar resource
 - ... in a way this creates space for possible reservations
 - Minimum lifetime of preemptible == max job time
- Cloud Portal Resources via OpenStack
 - Resource allocation maps to cloud credits (e.g. CPU hours)
 - When your reservation starts, use OpenStack as normal
 - Reservations ending trigger deleting all VMs in the reservation
 - ... option of preemptibles if you have no credit

Tracking Reservation Efficiency



Tracking VM Efficiency

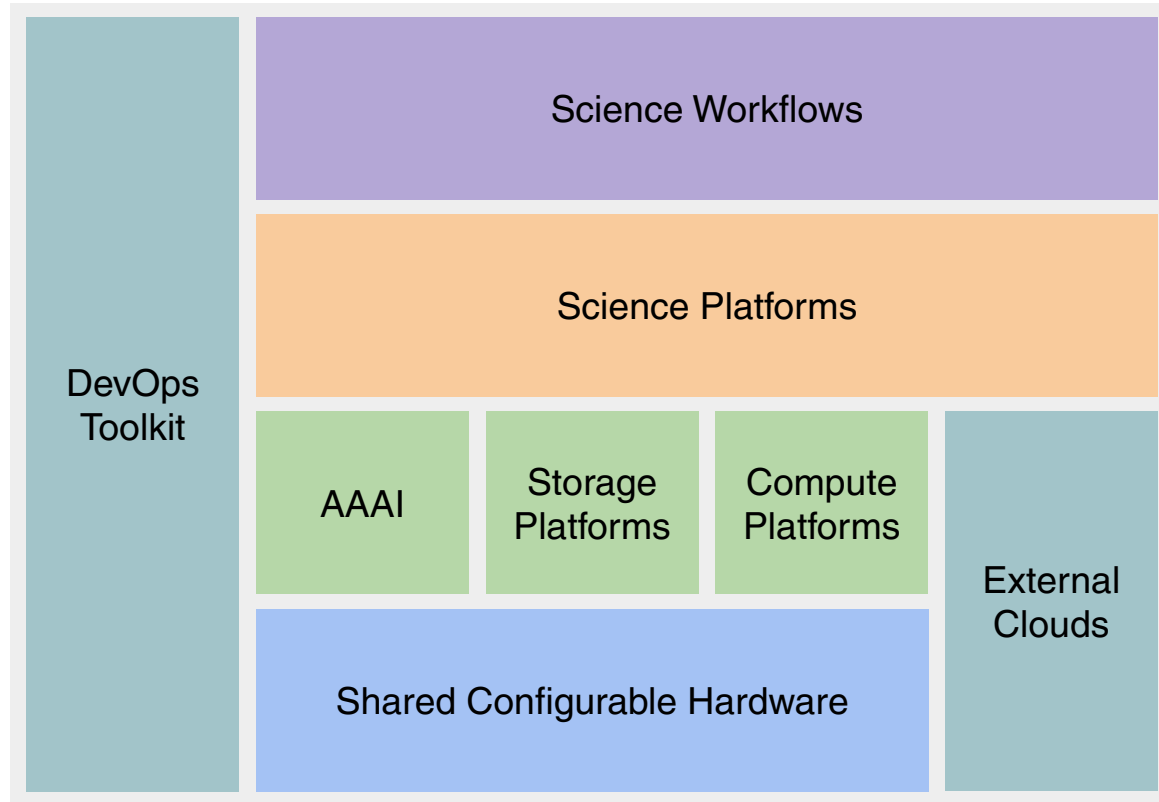
- Already collected in Prometheus via libvirt exporter
 - Example workload: stress-ng (with 75% allocation)
- WIP: Summary per project via CloudKitty
- WIP: OpenStack Grafana and Azimuth integration



Azimuth Cloud Portal

Scientific OpenStack

StackHPC




Azimuth Cloud Portal

StackHPC

- Self-service portal for managing cloud resources
 - Focus on ease of use for Scientists
 - While being optimized for HPC and AI use cases
- Make it easier to **find** and **reuse** common lego bricks
 - Could be a consistent interface across all IRIS OpenStack clouds
 - Links to: IRIS Docs, IRIS accounting portal, IRIS status pages, etc.
- Reduce time to science
- Reduce operational effort of onboarding new communities

Builds on the STFC funded work done by JASMIN



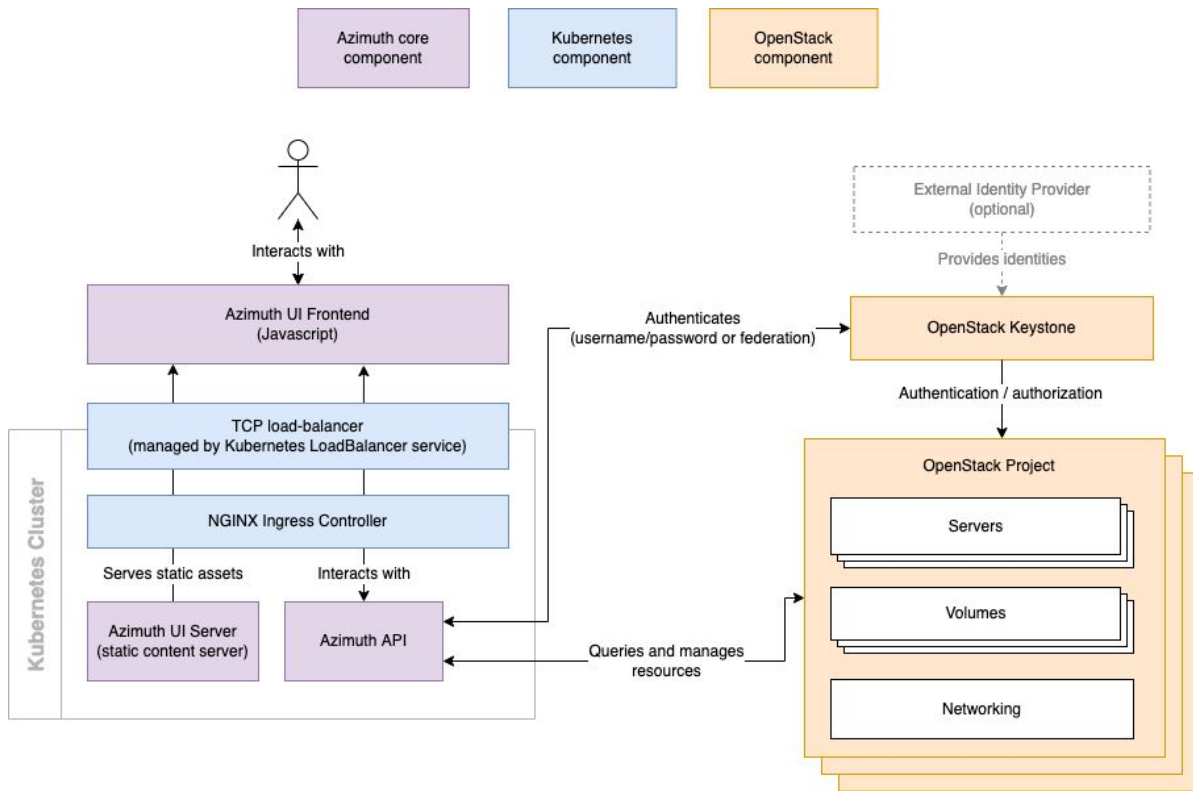
Key Azimuth Changes in FY21

- JupyterHub on Kubernetes
 - Built on Kubernetes Cluster API (not OpenStack Magnum, not Kubespray)
 - Batteries included: Monitoring, KubeApps, Dashboard, Ingress, Cinder CSI, and more
 - KubeApps pre-populated with helm charts (e.g. DB as a Service)
 - Simplify use of RDMA and GPUs, when available
 - NVMe optionally exposed as RWX PVCs (using Rook.io)
 - WIP: Tooling for Lustre shared between JupyterHub and CSD3
- Zenith Proxy
 - Streamline user experience
 - More efficient use of external IPs (only for the routers)
 - (optionally) share authn/z with Azimuth
- Slurm
 - AWX uses terraform to create Rocky Linux based OpenHPC cluster, with Open OnDemand

Azimuth Cloud Portal Architecture

StackHPC

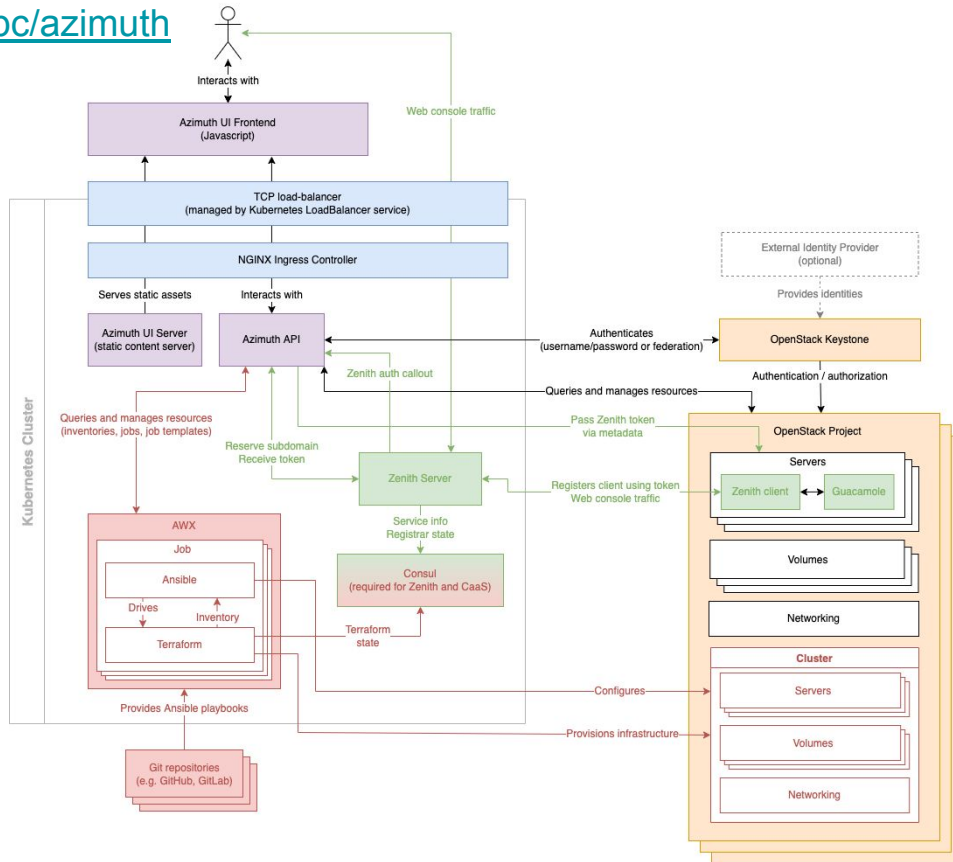
<https://github.com/stackhpc/azimuth>



Azimuth with AWX and Zenith

<https://github.com/stackhpc/azimuth>

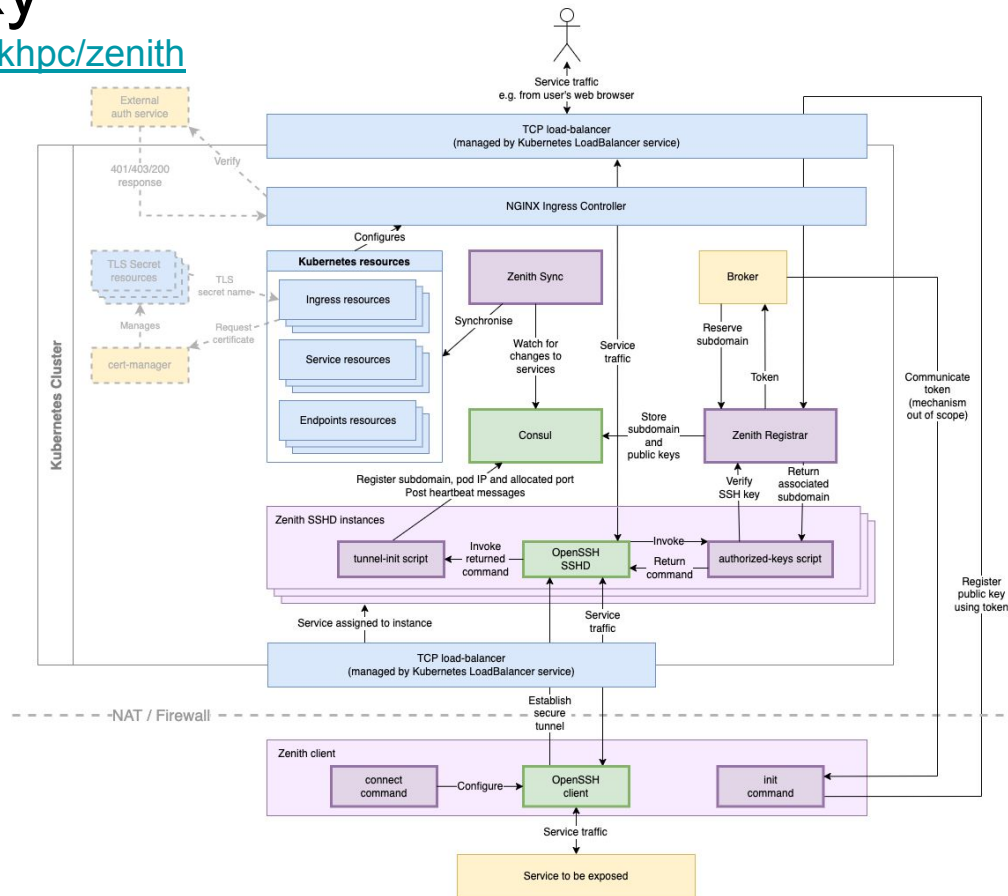
StackHPC



Zenith Proxy

<https://github.com/stackhpc/zenith>

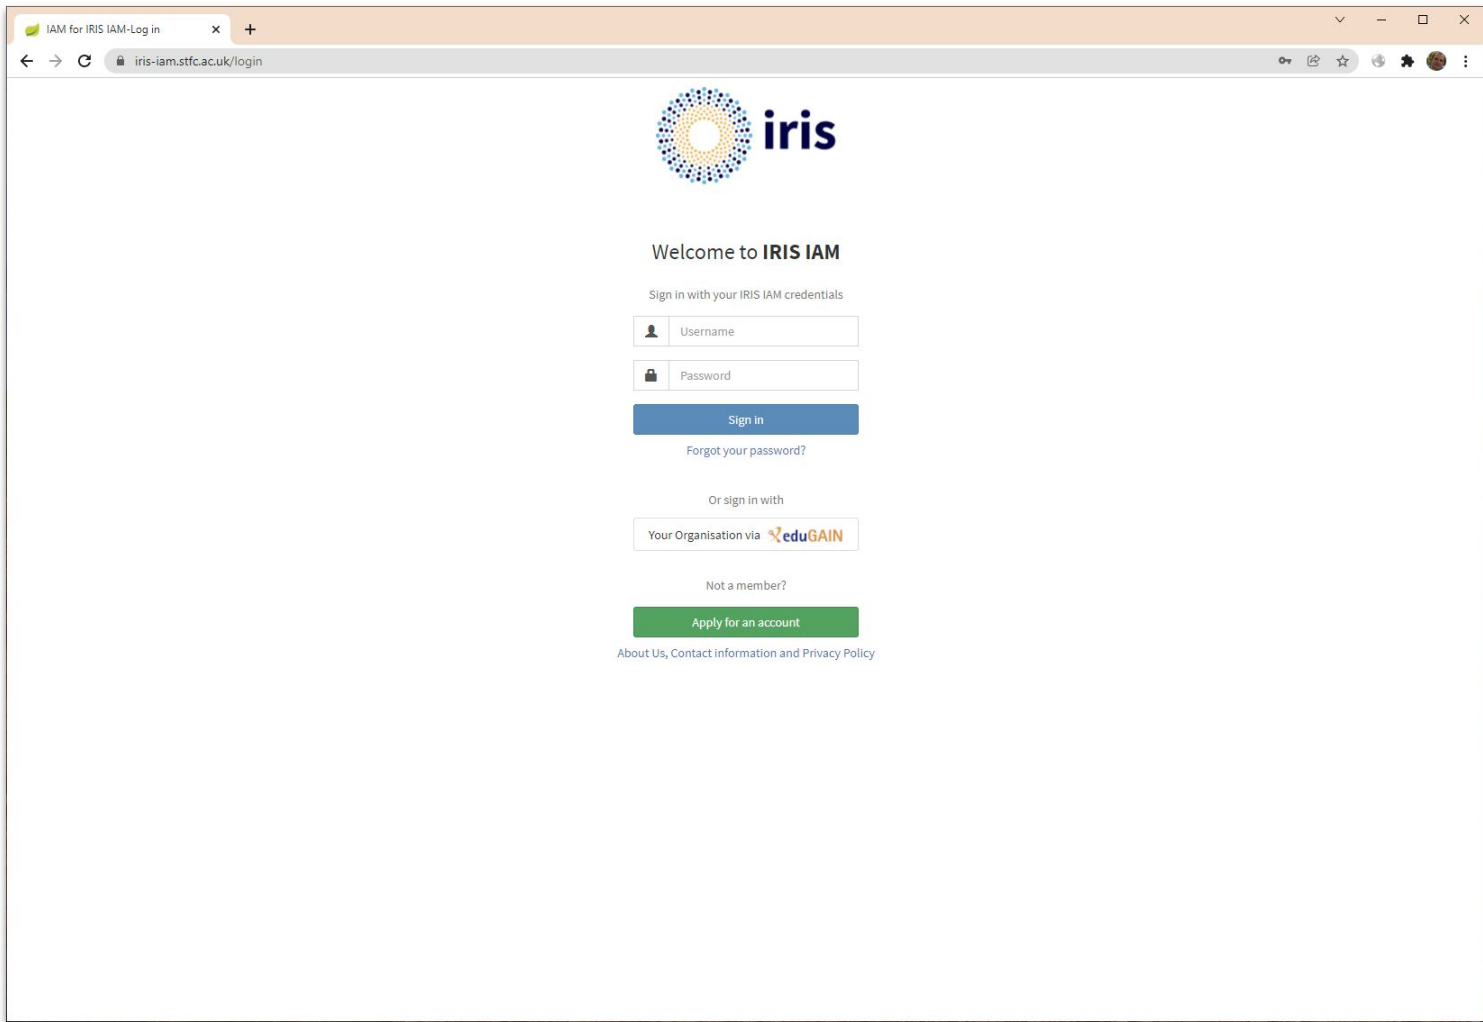
StackHPC



Future Ideas

- Current focus on short lived platforms
- More work to support creating platforms for your group
 - ... rather than only matching existing OpenStack project
- Lots of ideas from interviewing IRIS communities
- Includes:
 - More single machine apps: VirtualGL, Jupyter repo2docker, Matlab, R-Studio...
 - Automatically registering independent IRIS IAM client, via helm
 - Automatically using RDMA network for Slurm clusters
 - Storage for Slurm surviving the cluster
 - In place updates for Slurm
 - Slurm autoscale via the portal
 - ...

Cloud Portal: Get me a Slurm Cluster



Welcome to IRIS IAM

Sign in with your IRIS IAM credentials

Sign in

[Forgot your password?](#)


Or sign in with

Your Organisation via  eduGAIN

Not a member?

Apply for an account

[About Us, Contact information and Privacy Policy](#)



Dashboard

Available tenancies

iris-alaska-prod

iris-euclid

rcp-cloud-portal-demo

rcp-cloud-portal-dev

rcp-cloud-portal-prod

rcp-rds-lustre-dev

rcp-scientific-openstack-ci-cd

Clusters | rcp-cloud-portal-demo x +

portalapps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/clusters

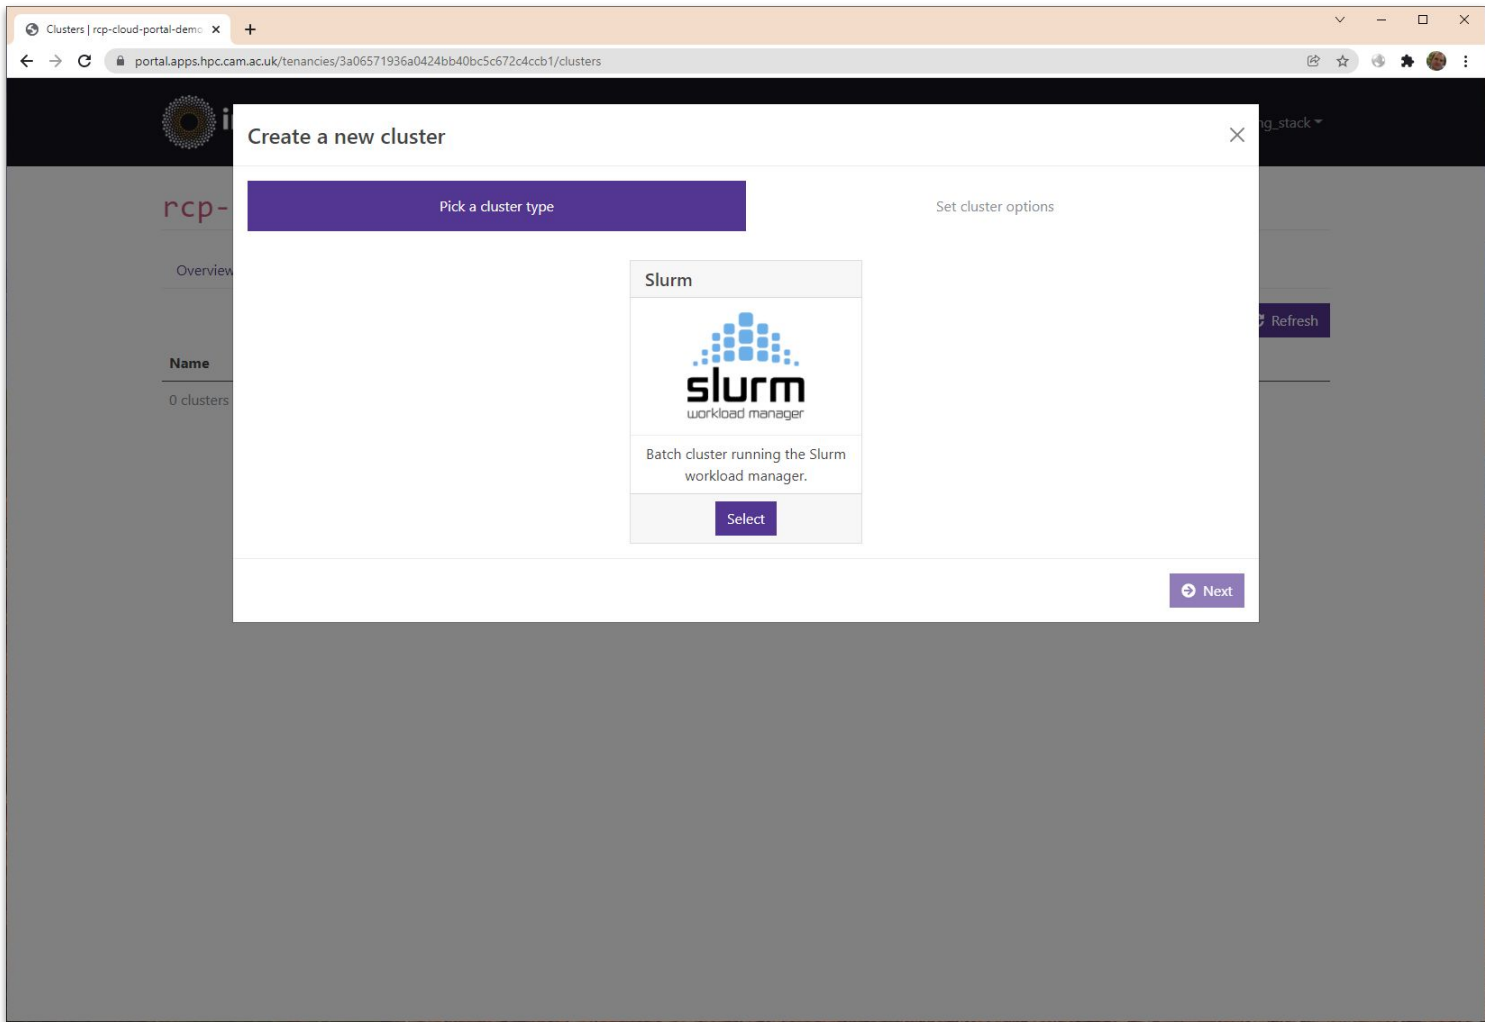
iris Select Tenancy johng_stack

rcp-cloud-portal-demo

Overview Machines Volumes Kubernetes Clusters

New cluster Refresh

Name	Cluster Type	Status	Task	Created	Updated	Patched
0 clusters						



Create a new cluster

Pick a cluster type

Set cluster options

Slurm



Batch cluster running the Slurm workload manager.

Select

Next


Clusters | rcp-cloud-portal-demo x Web console not ready x | +

portalapps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/clusters

iris Select Tenancy johng_stack

Create a new cluster

Pick a cluster type [Set cluster options](#)

 **Slurm** Batch cluster running the Slurm workload manager.

Cluster name

Must contain lower-case alphanumeric characters and dash (-) only.

External IP

 [+](#)

The external IP to use for the login node.

Compute node count

The number of compute nodes in the cluster.

Login node size

 [v](#)

The size to use for the login node.

Control node size

 [v](#)

The size to use for the control node.

Compute node size

 [v](#)

The size to use for the compute node.

Cluster monitoring

Enable cluster monitoring?

If selected, a monitoring stack will be deployed allowing you to track and visualise the state of the cluster.
WARNING: This can take a significant amount of time to deploy and configure.

[New cluster](#) [Refresh](#)

Patched

Name	Cluster Type
0 clusters	

Clusters | rcp-cloud-portal-demo x +

portalapps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/clusters

iris Select Tenancy

slurm Batch cluster running the Slurm workload manager.

Cluster name

test-slurm

Must contain lower-case alphanumeric characters and dash (-) only.

External IP

128.232.222.183

The external IP to use for the login node.

Compute node count

2

The number of compute nodes in the cluster.

Login node size

vm.alaska.cpu.general.small 4 cpus, 4GB RAM, 40GB disk

The size to use for the login node.

Control node size

vm.alaska.cpu.general.small 4 cpus, 4GB RAM, 40GB disk

The size to use for the control node.

Compute node size

vm.iris.cpu.dac.quarter 14 cpus, 38GB RAM, 2100GB disk

The size to use for the compute node.

Cluster monitoring

Enable cluster monitoring?

If selected, a monitoring stack will be deployed allowing you to track and visualise the state of the cluster.
WARNING: This can take a significant amount of time to deploy and configure.

Post-configuration validation

Run post-configuration validation?

If selected, post-configuration jobs will be executed to validate the core functionality of the cluster when it is re-configured.

[Back](#) [+ Create cluster](#)

Overview Machines Volume

Name	Cluster Type
0 clusters	

johnh_stack

New cluster Refresh

Patched

Clusters | rcp-cloud-portal-demo x +

portalapps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/clusters

iris Select Tenancy johng_stack

rcp-cloud-portal-demo

Overview Machines Volumes Kubernetes Clusters

New cluster Refresh

Name	Cluster Type	Status	Task	Created	Updated	Patched
test-slurm	Slurm	CONFIGURING	Waiting to be scheduled	a few seconds ago	-	-

1 cluster

Clusters | rcp-cloud-portal-demo x +

portalapps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/clusters

iris Select Tenancy johng_stack

rcp-cloud-portal-demo

Overview Machines Volumes Kubernetes Clusters

New cluster Refresh

Name	Cluster Type	Status	Task	Created	Updated	Patched	
test-slurm	Slurm	READY	-	19 minutes ago	4 minutes ago	4 minutes ago	Actions

1 cluster

Machines | rcp-cloud-portal-demo x +

portalapps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/machines

iris Select Tenancy johng_stack

rcp-cloud-portal-demo

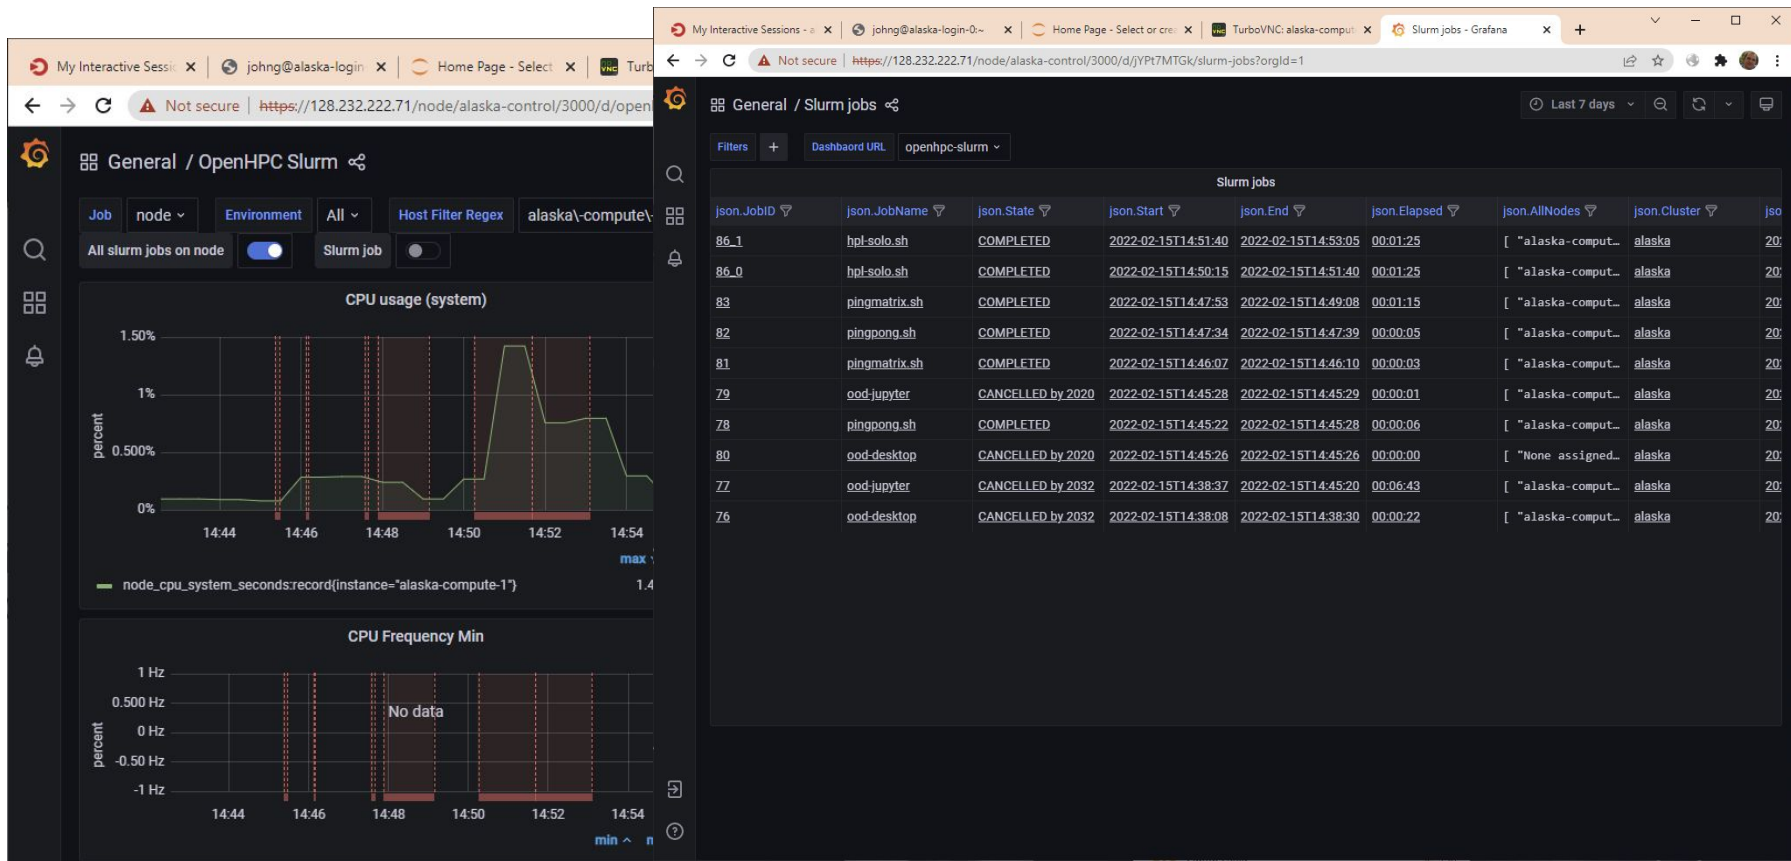
Overview Machines Volumes Kubernetes Clusters

New machine Refresh

Name	Image	Size	Status	Power State	Task	Internal IP	External IP	Created	
test-slurm-compute-0	CentOS8-2105	vm.iris.cpu.dac.quarter	ACTIVE	Running	-	192.168.3.126	-	19 minutes ago	Actions
test-slurm-compute-1	CentOS8-2105	vm.iris.cpu.dac.quarter	ACTIVE	Running	-	192.168.3.128	-	19 minutes ago	Actions
test-slurm-control-0	CentOS8-2105	vm.alaska.cpu.general.small	ACTIVE	Running	-	192.168.3.140	-	19 minutes ago	Actions
test-slurm-login-0	CentOS8-2105	vm.alaska.cpu.general.small	ACTIVE	Running	-	192.168.3.158	128.232.222.183	19 minutes ago	Actions
test-web-console	Ubuntu-Focal-20.04-20210624	vm.alaska.cpu.general.tiny	ACTIVE	Running	-	192.168.3.161	-	an hour ago	Actions
johng-dactest	CentOS8.4-OFED-5.4-1.0.3.0	vm.iris.cpu.dac.quarter	ACTIVE	Running	-	192.168.3.176	-	an hour ago	Actions

6 machines

Slurm integrated Grafana



WIP: Open OnDemand via Zenith

StackHPC

The screenshot shows the 'Jupyter Notebook' configuration page. The browser address bar is `login-web.hpc.cam.ac.uk/pun/sys/dashboard/batch_connect/sys/jupyter/csd3/session_contexts/new`. The navigation bar includes 'Research Computing Services HPC', 'Files', 'Jobs', 'Clusters', 'Interactive Apps', and 'My Interactive Sessions'. The left sidebar lists 'Interactive Apps' with 'Jupyter Notebook' selected. The main content area is titled 'Jupyter Notebook version: v1.0.1-3-g94d29b4'. It contains several form fields: 'Project Account', 'Partition', 'Reservation', 'Number of hours' (set to 1), 'Number of GPUs' (set to 0), and 'Modules' (set to 'jupyterlab/3'). A checkbox at the bottom is labeled 'I would like to receive an email when the session starts'.

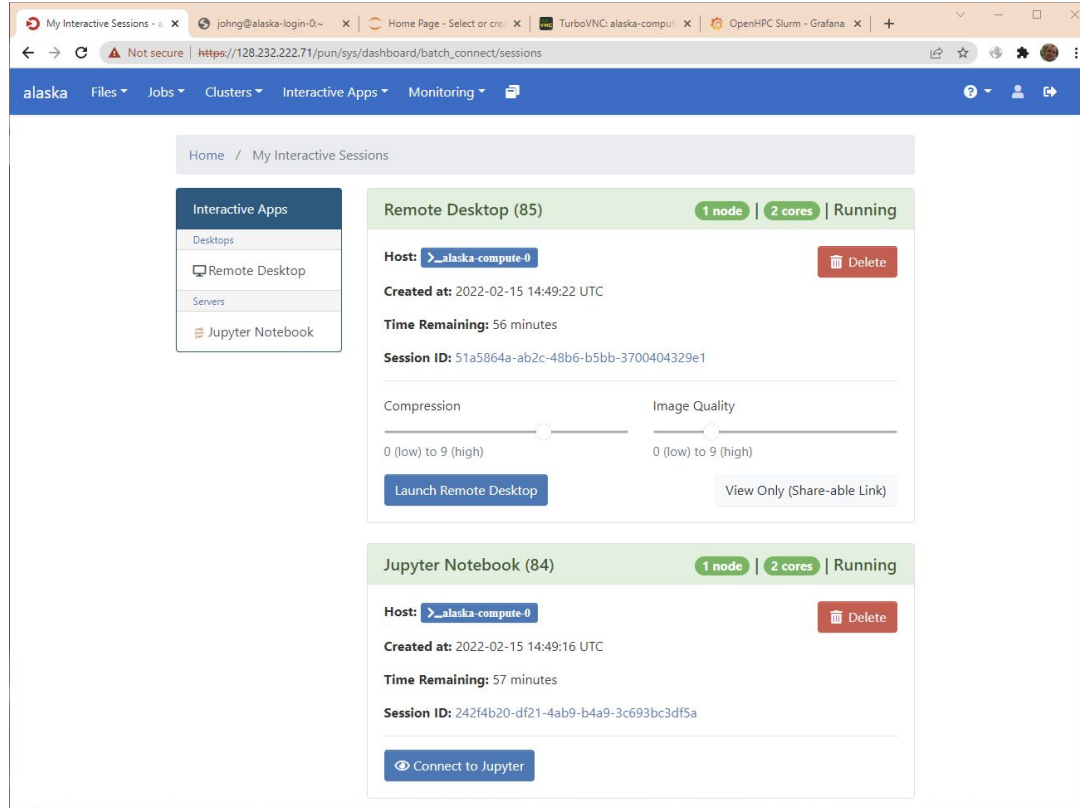
The screenshot shows the 'Templates' page in the Open OnDemand interface. The browser address bar is `login-web.hpc.cam.ac.uk/pun/sys/myjobs/workflows/new`. The navigation bar includes 'Open OnDemand', 'Job Composer', 'Jobs', and 'Templates'. The page title is 'Templates' with the subtitle 'To create a new job, select a template to copy, fill out the form to the right, and click "Create New Job".' Below the title is a table of templates:

Name	Cluster	Source
(default) Simple CPU Job	Csd3	System Templates
Simple GPU Job	Csd3	System Templates
Simple KNL Job	Csd3	System Templates

The 'Simple GPU Job' template is selected. To the right of the table is a form for creating a new job. The form includes: 'Job Name' (Simple GPU Job), 'Cluster' (CSD3), 'Script Name' (slurm_submit_wilkes2), and a 'Create New Job' button. Below the form is a 'Selected Template Details' section with fields for 'Template location' (`/etc/ood/config/apps/myjobs/templates/gpu`), 'Folder Contents' (manifest.yml, slurm_submit_wilkes2), and 'Script Contents'.

WIP: Open OnDemand via Zenith

StackHPC



The screenshot displays the Open OnDemand web interface. The browser's address bar shows the URL `https://128.232.222.71/pun/sys/dashboard/batch_connect/sessions`. The navigation menu includes "alaska", "Files", "Jobs", "Clusters", "Interactive Apps", and "Monitoring". The main content area is titled "My Interactive Sessions" and features a sidebar with "Interactive Apps" (Desktops, Remote Desktop, Servers, Jupyter Notebook). Two session cards are visible:

- Remote Desktop (85)**: 1 node, 2 cores, Running. Host: `_alaska-compute-0`. Created at: 2022-02-15 14:49:22 UTC. Time Remaining: 56 minutes. Session ID: 51a5864a-ab2c-48b6-b5bb-3700404329e1. Includes sliders for Compression and Image Quality, and buttons for "Launch Remote Desktop" and "View Only (Share-able Link)".
- Jupyter Notebook (84)**: 1 node, 2 cores, Running. Host: `_alaska-compute-0`. Created at: 2022-02-15 14:49:16 UTC. Time Remaining: 57 minutes. Session ID: 242f4b20-df21-4ab9-b4a9-3c693bc3df5a. Includes a "Connect to Jupyter" button.

WIP: Open OnDemand via Zenith

StackHPC

The image displays three overlapping screenshots related to an Open OnDemand environment:

- Terminal Screenshot (Left):** Shows a shell session on a host named 'alaska-login-0'. The user 'johnhg' runs 'sqlite3 squeue' multiple times, which outputs a table of job information. The table columns are JOBID, PARTITION, and NAME. The output shows jobs with IDs 78, 79, and 77, all in the 'small' partition, with names 'pingpong', 'ood-jupy', and 'ood-jupy' respectively.
- JupyterLab Screenshot (Top Right):** Shows the JupyterLab web interface. The browser address bar indicates the URL 'https://128.232.222.71/node/alaska-compute-0/60156/tree?'. The interface includes a 'Quit' and 'Logout' button, and a file browser showing a tree view with 'Desktop' and 'ondemand' folders, both last modified 4 months ago.
- Desktop Screenshot (Bottom):** Shows a desktop environment with a dark theme. The desktop contains icons for 'Filesystem root', 'home', 'scratch', and 'File System'. The system tray at the bottom shows the date and time as 'Tue 15 Feb, 14:54 John Garbutt'.


Live Demo of Azimuth

Live Demo of Azimuth

StackHPC

- Bigger Laptop
- Kubernetes
 - Zenith proxy to Grafana and KubeApps
 - JupyterHub SSO
- Customizing JupyterHub, and its Storage

Cloud Portal:
Get me a bigger laptop,
via IRIS IAM



Welcome to IRIS IAM

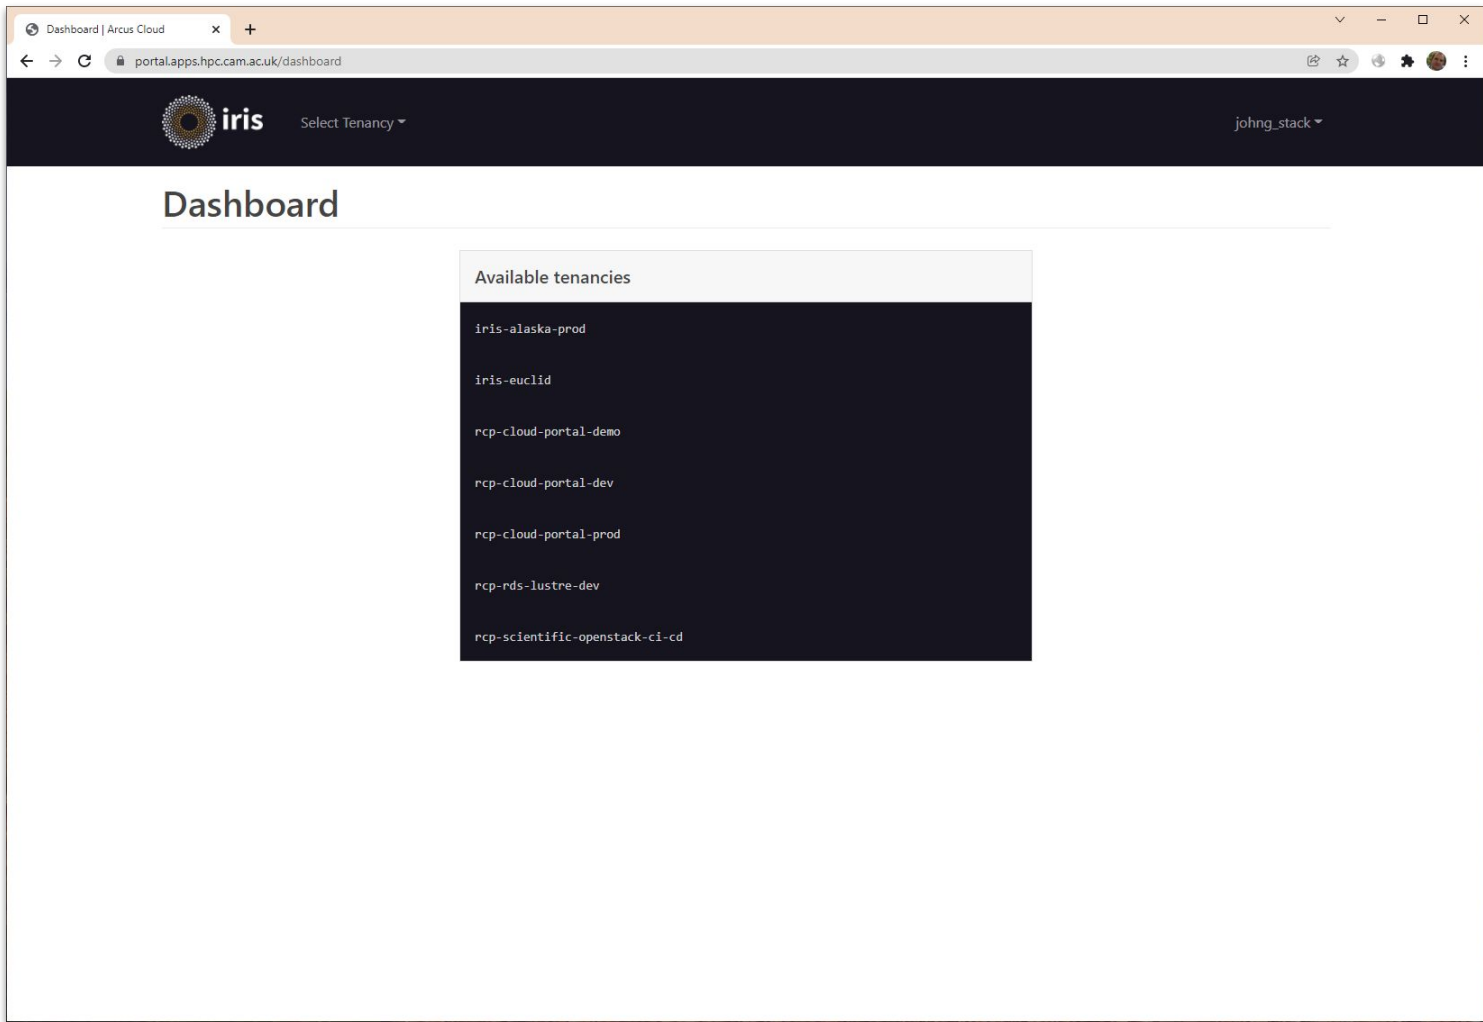
Sign in with your IRIS IAM credentials

[Forgot your password?](#)

Or sign in with

Not a member?

[About Us, Contact information and Privacy Policy](#)



Dashboard

Available tenancies

iris-alaska-prod

iris-euclid

rcp-cloud-portal-demo

rcp-cloud-portal-dev

rcp-cloud-portal-prod

rcp-rds-lustre-dev

rcp-scientific-openstack-ci-cd

Machines | rcp-cloud-portal-den x +

portalapps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/machines

iris Select Tenancy johng_stack

Create a new machine

Machine name

Must contain alphanumeric characters, dot (.) and dash (-) only.

Image

Size

Enable web console?
Installs [Apache Guacamole](#) to provide access to the machine via a web browser.

Enable remote desktop for web console?
WARNING: The remote desktop can take a long time to install and configure.

+ Create machine

rcp-cloud-portal

Overview Machines Volume

Name	Image
johng-dactest	CentOS8.4-OPEL

1 machine

New machine Refresh

Created 12 minutes ago Actions

Machines | rcp-cloud-portal-demo x +

portalapps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/machines

iris Select Tenancy johng_stack

rcp-cloud-portal-demo

Overview Machines Volumes Kubernetes Clusters

[New machine](#) [Refresh](#)

Name	Image	Size	Status	Power State	Task	Internal IP	External IP	Created
test-web-console	Ubuntu-Focal-20.04-20210624	vm.alaska.cpu.general.tiny	BUILD	Unknown	Spawning	-	-	a few seconds ago
johng-dactest	CentOS8.4-OFED-5.4-1.0.3.0	vm.iris.cpu.dac.quarter	ACTIVE	Running	-	192.168.3.176	-	14 minutes ago

2 machines

Actions

Machines | rcp-cloud-portal-demo x +

portalapps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/machines

iris Select Tenancy johng_stack


rcp-cloud-portal-demo

Overview Machines Volumes Kubernetes Clusters

New machine Refresh

Name	Image	Size	Status	Power State	Task	Internal IP	External IP	Created	Actions
test-web-console	Ubuntu-Focal-20.04-20210624	vm.alaska.cpu.general.tiny	ACTIVE	Running	-	192.168.3.161	-	a minute ago	Actions
johng-dactest	CentOS8.4-OFED-5.4-1.0.3.0	vm.iris.cpu.dac.quarter	ACTIVE	Running	-	192.168.3.176	-	14 min	Access web console Attach external IP Detach external IP Firewall rules Start machine Stop machine Restart machine View machine logs Delete machine

2 machines



```
Machines | rcp-cloud-portal-den x shell x +
cfhykb9nz6fvov3e5wwqj58m5tidjier.apps.hpc.cam.ac.uk/guacamole/#/client/c2hlgWwAYwBkZwZhdWx0
Welcome to Ubuntu 22.04.1 LTS (GNU/Linux 5.15.0-46-generic x86_64)
...8m5tidjier.apps.hpc.cam.ac.uk wants to
* Document: See text and images copied to the clipboard
* Manage: cal.com
* Support: tage
Allow Block
System information as of the date of the update:
System load: 1.16          Processes:              148
Usage of /: 17.1% of 19.21GB Users logged in:          0
Memory usage: 62%         IPv4 address for ens3: 192.168.3.161
Swap usage: 0%

126 updates can be applied immediately.
51 of these updates are standard security updates.
To see these additional updates run: apt list --upgradable

The programs included with the Ubuntu system are free software;
the exact distribution terms for each program are described in the
individual files in /usr/share/doc/*/copyright.

Ubuntu comes with ABSOLUTELY NO WARRANTY, to the extent permitted by
applicable law.

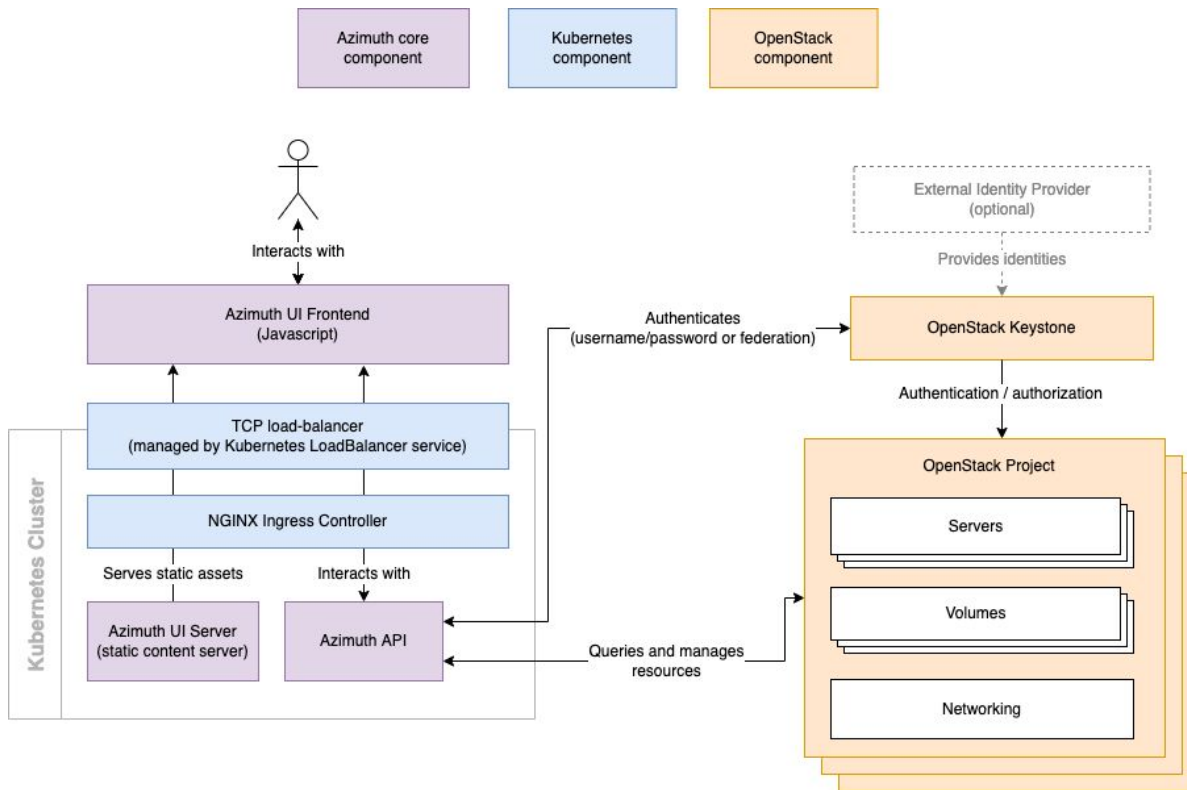
To run a command as administrator (user "root"), use "sudo <command>".
See "man sudo_root" for details.

ubuntu@test-web-console:~$
```

Azimuth Cloud Portal Architecture

<https://github.com/stackhpc/azimuth>

StackHPC



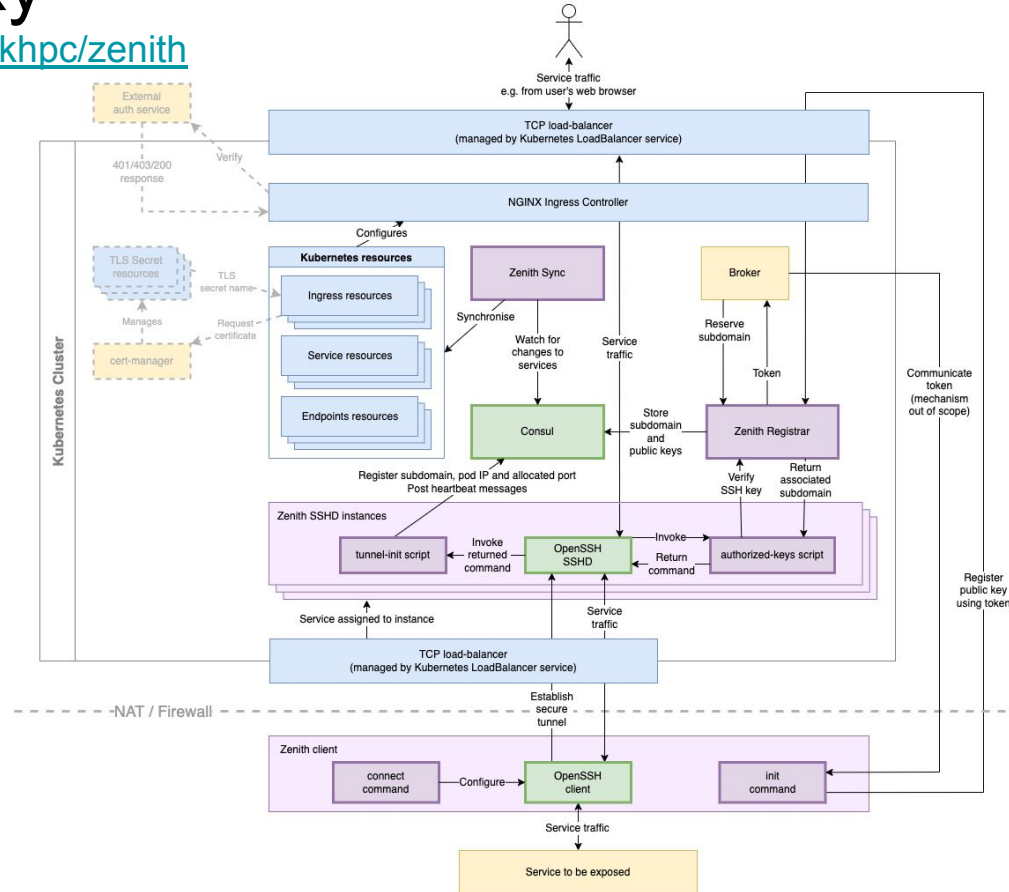
How did you get access to that VM?

- There is nothing hiding the OpenStack API here, similar to Exposhere
- IRIS IAM login to OpenStack Keystone
 - Get keystone token to access the API
 - No credentials ever go through the Cloud Portal
- Create OpenStack server via API
 - Cloud-init configures guacamole
 - And starts an ssh session to a proxy
- Azimuth Proxy
 - HA possible via use of Nginx, consul, consul-template, ssh
 - <https://github.com/stackhpc/azimuth>

Zenith Proxy

<https://github.com/stackhpc/zenith>

StackHPC



Cloud Portal:
Get me a K8s Cluster,
via 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](#)



Dashboard

Available tenancies

iris-alaska-prod

iris-euclid

rcp-cloud-portal-demo

rcp-cloud-portal-dev

rcp-cloud-portal-prod

rcp-rds-lustre-dev

rcp-scientific-openstack-ci-cd

Kubernetes | rcp-cloud-portal-ci x +

portal.apps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/kubernetes

iris Select Tenancy johng_stack

Create a new Kubernetes cluster

Cluster name

Must contain lower-case alphanumeric characters and dash (-) only.

Cluster template

The template determines the Kubernetes version for the cluster.

Control Plane Size

The size to use for the Kubernetes control plane node(s).

Enable auto-healing?
If enabled, the cluster will try to remediate unhealthy nodes automatically.

Node Groups

Name	Node Size	Node Count
No node groups configured yet.		

[+ Add node group](#)

Cluster Addons

- Enable cert-manager?
- Enable Kubernetes Dashboard?
- Enable Kubernetes Ingress?
- Enable cluster monitoring?
- Enable applications dashboard?

[+ Create cluster](#)

rcp-cloud-portal-ci

Overview Machines Volume

Name	Status
data-demo	Reconciling
jhub-demo	Reconciling

2 clusters

New cluster Refresh

minutes ago Actions

minutes ago Actions

Kubernetes | rcp-cloud-portal-cl... x

portal.apps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/kubernetes

iris Select Tenancy johng_stack

rcp-cloud-portal-demo

Overview Machines Volumes **Kubernetes** Clusters

New cluster Refresh

Name	Status	Template	Control Plane	Workers	Addons	Created	Actions
jhub-demo	Ready	v1.22.6	Ready	2 (2 ready)	14 (14 ready)	an hour ago	Generate kubeconfig Cluster details Modify cluster Upgrade cluster Delete cluster

1 cluster

Kubernetes | rcp-cloud-portal-... x +

portal.apps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/kubernetes

iris Select Tenancy johng_stack

Kubeconfig for jhub-demo

Copy to clipboard Download Regenerate

Use this configuration file with the `kubectl` command-line tool to access your cluster.

```
apiVersion: v1
clusters:
- cluster:
  certificate-authority-data: LS0tLS1CRUdJTTI8DRVJU5UZJ08FURS0tLS0tcck1JSUM2akNDQWRRLZ0F3SU1B70I=
  server: https://k9rt75xegzvj56gpowfu031ut7uyqnvcm0jd47ccy5q.apps.hpc.cam.ac.uk:443
  name: jhub-demo
contexts:
- context:
  cluster: jhub-demo
  user: jhub-demo-admin
  name: jhub-demo-admin@jhub-demo
current-context: jhub-demo-admin@jhub-demo
kind: Config
preferences: {}
users:
- name: jhub-demo-admin
  user:
```

Close

Overview Machines Volume

Name	Status
data-demo	Reconciling
jhub-demo	Ready

2 clusters

New cluster Refresh

minutes ago Actions

minutes ago Actions

Kubernetes | rcp-cloud-portal-ci x +

portal.apps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/kubernetes

iris Select Tenancy johng_stack

Upgrade Kubernetes cluster jhub-demo

Upgrading a Kubernetes cluster is a long-running and potentially disruptive operation that may affect workloads running on the cluster.

Once started, an upgrade cannot be stopped.

Cluster template

1.23.3 Kubernetes version: 1.23.3

The template determines the Kubernetes version for the cluster.

[New cluster](#) [Refresh](#)

Name	Status
data-demo	Reconciling
jhub-demo	Ready

2 clusters

[Upgrade cluster](#)

minutes ago Actions

minutes ago Actions

Kubernetes | rcp-cloud-portal-
portal.apps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/kubernetes

iris Select Tenancy johng_stack

Update Kubernetes cluster jhub-demo



Cluster name
jhub-demo
Must contain lower-case alphanumeric characters and dash (-) only.

Cluster template
v1.22.6 Kubernetes version: 1.22.6
The template determines the Kubernetes version for the cluster.

Control Plane Size
vm.alaska.cpu.general.small 4 cpus, 4GB RAM, 40GB disk
The size to use for the Kubernetes control plane node(s).

Enable auto-healing?
If enabled, the cluster will try to remediate unhealthy nodes automatically.

Node Groups

Name	Node Size	Node Count	
md-0	vm.alaska.cpu.general.small	2	 

[+ Add node group](#)

Cluster Addons

- Enable cert-manager?
- Enable Kubernetes Dashboard?
- Enable Kubernetes Ingress?
- Enable cluster monitoring?
- Enable applications dashboard?

[Update cluster](#)

rcp-cloud-portal-
Overview Machines Volume
Name Status
data-demo Reconciling
jhub-demo ✔ Ready
2 clusters

New cluster Refresh
minutes ago Actions
minutes ago Actions

Kubernetes | rcp-cloud-portal-
portal.apps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/kubernetes

Cluster details for jhub-demo

Overview

Cluster details	
Name	jhub-demo
Template	v1.22.6
Kubernetes version	1.22.6
Status	✓ Ready
Autohealing?	✓ Enabled
Node groups	1
Created	an hour ago

Control plane

Status	✓ Ready
Size	vm.alaska.cpu.general.small
Node Count	3

Nodes

Services	
JupyterHub	↗
Applications	↗
Kubernetes Dashboard	↗
Monitoring	↗

Cluster addons

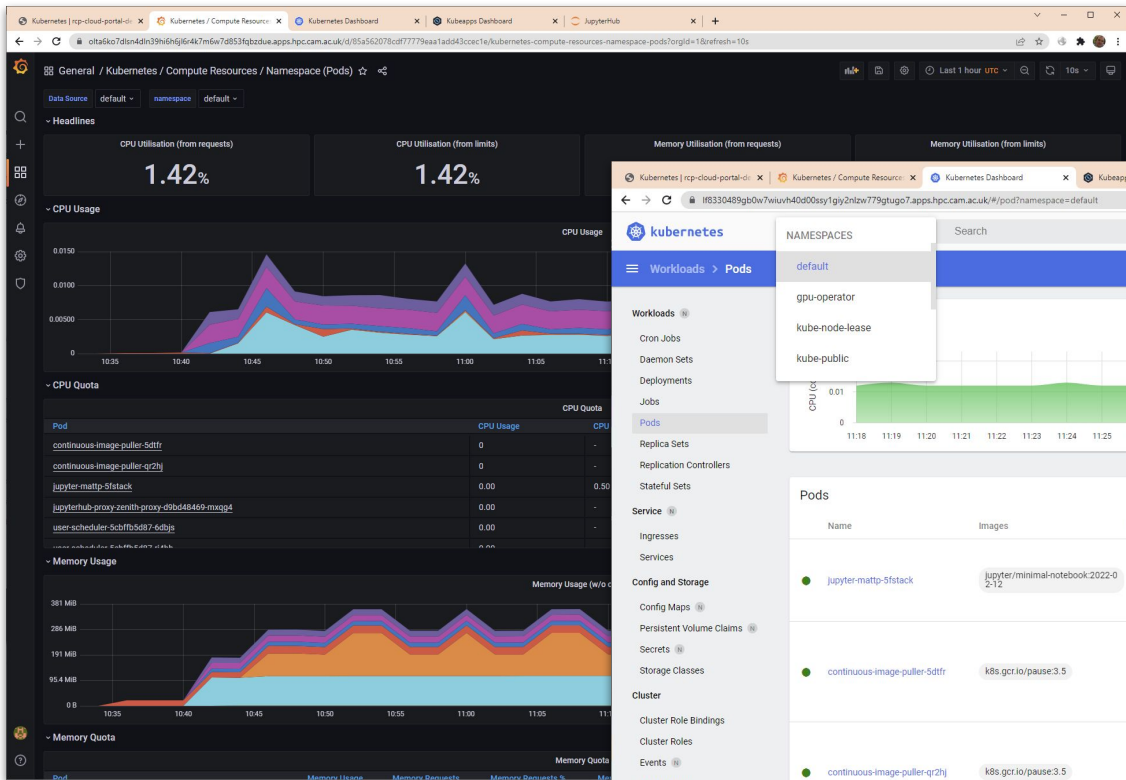
ccm-openstack	✓ Ready
cloud-config	✓ Ready
cni-cilium	✓ Ready
csi-cinder	✓ Ready
jupyterhub-proxy	✓ Ready
kubeapps	✓ Ready
kubeapps-proxy	✓ Ready
kubernetes-dashboard	✓ Ready
kubernetes-dashboard-proxy	✓ Ready
metrics-server	✓ Ready
monitoring	✓ Ready
monitoring-proxy	✓ Ready
node-feature-discovery	✓ Ready
nvidia-gpu-operator	✓ Ready

Refresh

actions

ig_stack

Close



kubernetes

Workloads > **Pods**

Namespaces

- default
- gpu-operator
- kube-node-lease
- kube-public

Pod List

Name	Images	Labels	Node	Status	Restarts	CPU Usage (cores)	Memory Usage (bytes)	Created
jupyter-matp-5fstack	jupyter/minimal-notebook:2022-02-12	app: jupyterhub chart: jupyterhub-1.2.0 component: singleuser-server	jhub-demo-md-0-8d3a6fca-m255t	Running	0	1.00m	80.45Mi	20 minutes ago
continuous-image-puller-5dtfr	k8s.gcr.io/pause:3.5	component: continuous-image-puller controller-revision-hash: 597d478fd7	jhub-demo-md-0-8d3a6fca-wgpc	Running	0	0.00m	836.00Ki	53 minutes ago
continuous-image-puller-qz2h	k8s.gcr.io/pause:3.5	component: continuous-image-puller controller-revision-hash: 597d478fd7	jhub-demo-md-0-8d3a6fca-m255t	Running	0	0.00m	956.00Ki	53 minutes ago
hub-5b769897b5-h5sfh	jupyterhub/k8s-hub:1.2.0	app: jupyterhub component: hub hub.jupyter.org/network-access-proxy-api: true	jhub-demo-md-0-8d3a6fca-m255t	Running	0	3.00m	110.27Mi	53 minutes ago
proxy-766c8f1cc9-hplw7	jupyterhub/configurable-http-proxy:4.5.0	app: jupyterhub component: proxy hub.jupyter.org/network-access-proxy-api: true	jhub-demo-md-0-8d3a6fca-m255t	Running	0	1.00m	17.34Mi	53 minutes ago

Kubernetes | rcp-cloud-portal-... x +

portal.apps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/kubernetes

Cluster details for jhub-demo

Overview **Nodes**

Name	Status	Size	Kubelet Version	IP address
jhub-demo-control-plane-bxxx* control-plane	✓ Ready	vm.alaska.cpu.general.small	1.22.6	192.168.3.234
jhub-demo-control-plane-1c9d6 control-plane	✓ Ready	vm.alaska.cpu.general.small	1.22.6	192.168.3.152
jhub-demo-control-plane-n7vt2 control-plane	✓ Ready	vm.alaska.cpu.general.small	1.22.6	192.168.3.172
jhub-demo-md-0-5c796b8697-cc859 worker md-0	✓ Ready	vm.alaska.cpu.general.small	1.22.6	192.168.3.176
jhub-demo-md-0-5c796b8697-f4994 worker md-0	✓ Ready	vm.alaska.cpu.general.small	1.22.6	192.168.3.190

5 nodes

Close

Machines | rcp-cloud-portal-de... x +

portal.apps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/machines

iris Select Tenancy johng_stack

rcp-cloud-portal-demo

Overview Machines Volumes Kubernetes Clusters

[New machine](#) [Refresh](#)

Name	Image	Size	Status	Power State	Task	Internal IP	External IP	Created	Actions
jhub-demo-control-plane-8d3a6fca-d9hg	ubuntu-2004-kube-v1.22.6	vm.alaska.cpu.general.small	ACTIVE	Running	-	192.168.3.152	-	an hour ago	Actions
jhub-demo-control-plane-8d3a6fca-vqv7h	ubuntu-2004-kube-v1.22.6	vm.alaska.cpu.general.small	ACTIVE	Running	-	192.168.3.234	-	an hour ago	Actions
jhub-demo-md-0-8d3a6fca-m255t	ubuntu-2004-kube-v1.22.6	vm.alaska.cpu.general.small	ACTIVE	Running	-	192.168.3.176	-	an hour ago	Actions
jhub-demo-md-0-8d3a6fca-wggpc	ubuntu-2004-kube-v1.22.6	vm.alaska.cpu.general.small	ACTIVE	Running	-	192.168.3.190	-	an hour ago	Actions
jhub-demo-control-plane-8d3a6fca-26jwh	ubuntu-2004-kube-v1.22.6	vm.alaska.cpu.general.small	ACTIVE	Running	-	192.168.3.172	-	an hour ago	Actions
lustre-exporter	RockyLinux-8.5-20211114.2	vm.alaska.cpu.general.small	ACTIVE	Running	-	192.168.3.4	-	2 hours ago	Actions
lustre-admin	RockyLinux-8.5-20211114.2	vm.alaska.cpu.general.small	ACTIVE	Running	-	10.60.102.244	-	2 hours ago	Actions
vg-slurm-test-compute-0	RockyLinux-8.5-20211114.2	vm.alaska.cpu.general.small	ACTIVE	Running	-	192.168.3.254	-	13 hours ago	Actions
vg-slurm-test-login-0	RockyLinux-8.5-20211114.2	vm.alaska.cpu.general.small	ACTIVE	Running	-	192.168.3.48	128.232.222.147	13 hours ago	Actions
lustre-csd-client	RockyLinux-8.5-20211114.2	vm.alaska.cpu.general.small	ACTIVE	Running	-	10.60.102.193	-	19 hours ago	Actions

How did you do create the K8s cluster?

StackHPC

- Kubernetes Cluster API
 - Machines provided by Cluster API Provider OpenStack
- Operator and helm charts on top of Cluster API
 - Install addons, including Monitoring and Networking Drivers
- Future funded work: Azimuth Proxy support:
 - Grafana, Kubernetes Dashboard, and kubeapps are the first targets

Cloud Portal: Get me a JupyterHub, with Azimuth SSO

Kubernetes | rcp-cloud-portal-... portal.apps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/kubernetes

Cluster details for jhub-demo

Overview

Cluster details	
Name	jhub-demo
Template	v1.22.6
Kubernetes version	1.22.6
Status	✓ Ready
Autohealing?	✓ Enabled
Node groups	1
Created	an hour ago

Control plane

Status	✓ Ready
Size	vm.alaska.cpu.general.small
Node Count	3

Nodes

Services	
JupyterHub	🔗
Applications	🔗
Kubernetes Dashboard	🔗
Monitoring	🔗

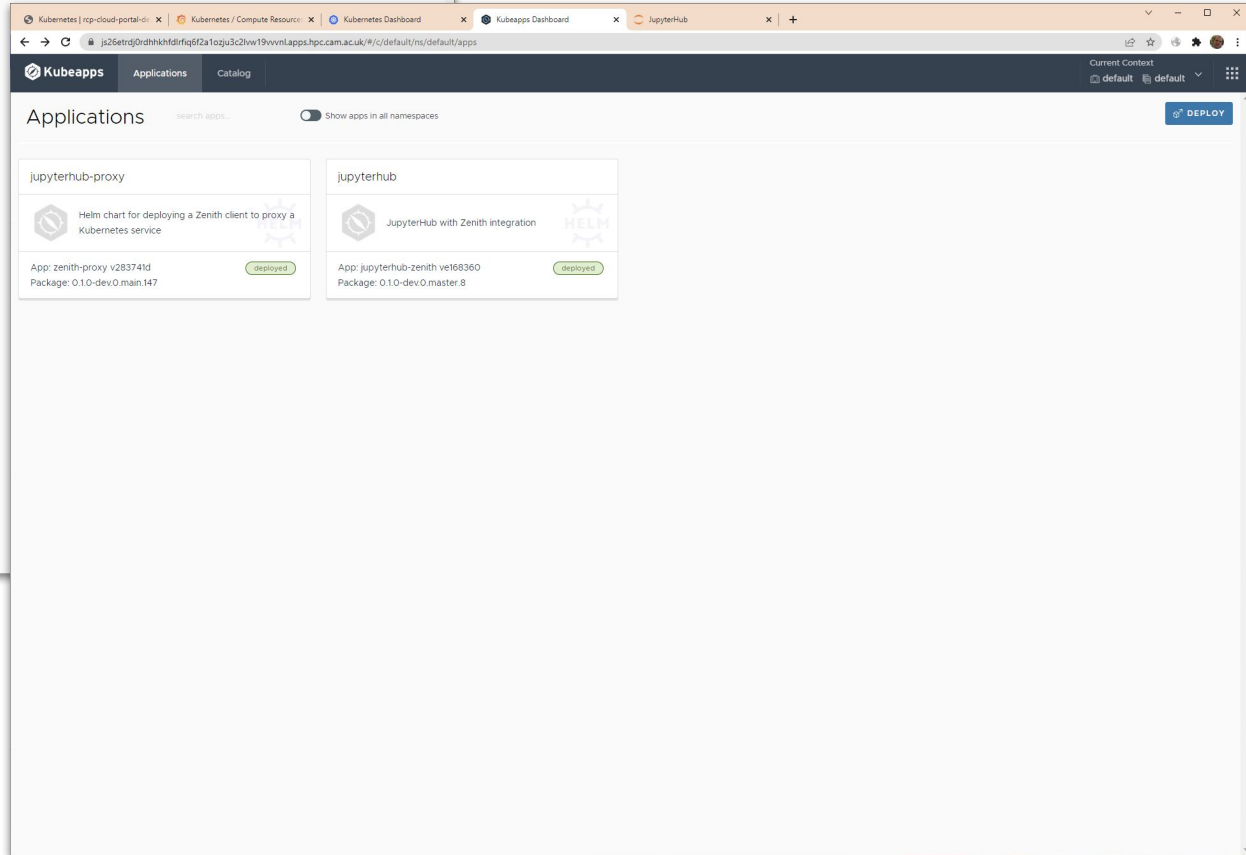

Cluster addons

ccm-openstack	✓ Ready
cloud-config	✓ Ready
cni-cilium	✓ Ready
csi-cinder	✓ Ready
jupyterhub-proxy	✓ Ready
kubeapps	✓ Ready
kubeapps-proxy	✓ Ready
kubernetes-dashboard	✓ Ready
kubernetes-dashboard-proxy	✓ Ready
metrics-server	✓ Ready
monitoring	✓ Ready
monitoring-proxy	✓ Ready
node-feature-discovery	✓ Ready
nvidia-gpu-operator	✓ Ready

Refresh

actions

Close



Kubernetes | rcp-cloud-portal-... portal.apps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/kubernetes

Cluster details for jhub-demo

Overview

Cluster details	
Name	jhub-demo
Template	v1.22.6
Kubernetes version	1.22.6
Status	✔ Ready
Autohealing?	✔ Enabled
Node groups	1
Created	an hour ago

Control plane

Status	✔ Ready
Size	vm.alaska.cpu.general.small
Node Count	3

Nodes

Services	
JupyterHub	🔗
Applications	🔗
Kubernetes Dashboard	🔗
Monitoring	🔗

Cluster addons

ccm-openstack	✔ Ready
cloud-config	✔ Ready
cni-cilium	✔ Ready
csi-cinder	✔ Ready
jupyterhub-proxy	✔ Ready
kubeapps	✔ Ready
kubeapps-proxy	✔ Ready
kubernetes-dashboard	✔ Ready
kubernetes-dashboard-proxy	✔ Ready
metrics-server	✔ Ready
monitoring	✔ Ready
monitoring-proxy	✔ Ready
node-feature-discovery	✔ Ready
nvidia-gpu-operator	✔ Ready

Refresh

actions

Close

Kubernetes | rcp-cloud-portal-... x | Kubernetes / Compute Resource: x | Kubernetes Dashboard x | Kubeapps Dashboard x | JupyterHub x +

x9mczce306rkqypf14nqz4j9dvme7xzl9o9xjs376ee.apps.hpc.cam.ac.uk/hub/spawn

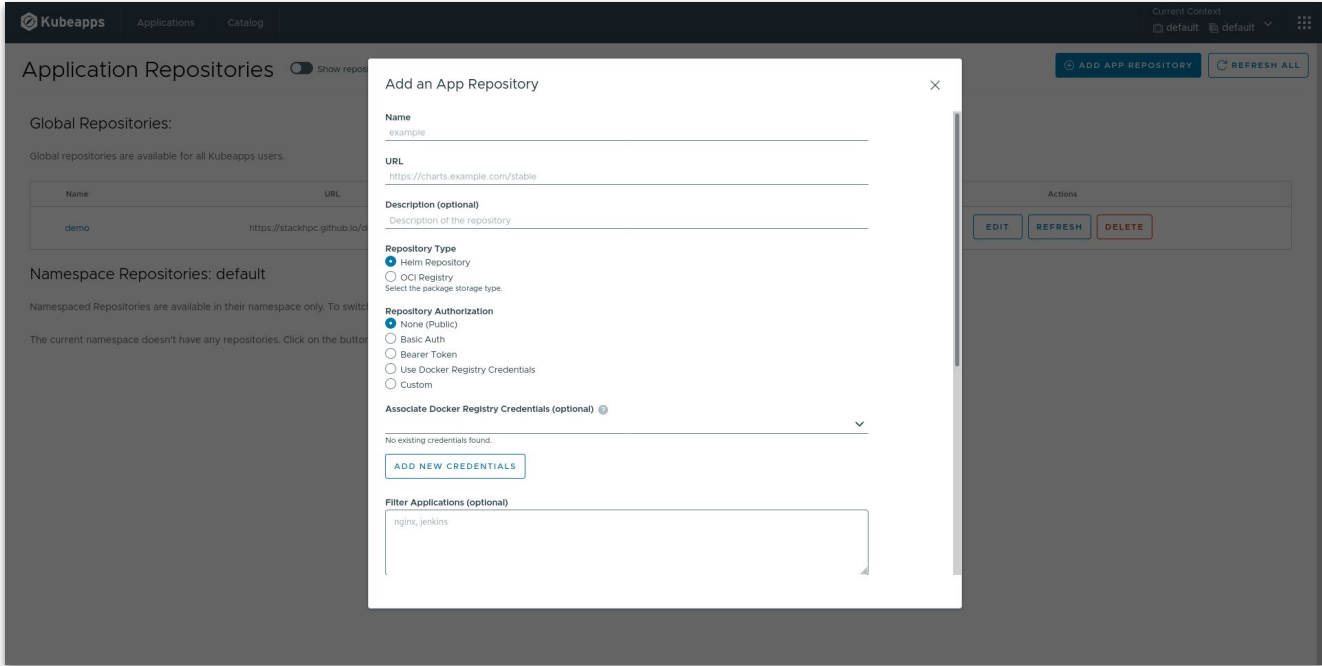
jupyterhub Home Token johng_stack Logout

Server Options

- Python environment (minimal)**
Minimal Python environment
- Spark environment**
Python, R, and Scala support for Apache Spark.

Start

Cloud Portal: Customize my JupyterHub



Catalog

search packages...

FILTERS [CLEAR ALL](#)

Category

Unknown

Application Repository

demo

Repository: demo X

jupyterhub-zenith



JupyterHub with Zenith integration



0.1.0-dev.0.master.13

demo

Helm

jupyterhub-zenith-nfs



JupyterHub with Zenith integration and NFS



0.1.0-dev.0.master.13

demo

Helm

jupyterhub-zenith-nfs-customrepo



JupyterHub with Zenith integration, NFS and a custom git repository



0.1.0-dev.0.master.13

demo

Helm



demo/jupyterhub-zenith-nfs-customrepo

Helm Chart

Package Version 0.1.0-dev.0.master.13 / App Version ca596ec

App Version
ca596ecPackage Version
0.1.0-dev.0.master.13

Name

jupyterhub

Form **YAML** Changes

```
56 -   - display name: "Spark environment"
57 -     description: "Python, R, and Scala support for Apache Spark."
58 -     kubespawner_override:
59 -       image: jupyter/all-spark-notebook:2022-02-12
60 -   memory:
61 -     limit: "1G"
62 -     guarantee: "1G"
63 -   cpu:
64 -     limit: .5
65 -     guarantee: .5
66 -   storage:
67 -     capacity: "10Gi"
68 -     extraVolumes:
69 -       - name: dac-nvme
70 -         persistentVolumeClaim:
71 -           claimName: rook-nfs-pv-claim
72 -     extraVolumeMounts:
73 -       - name: dac-nvme
74 -         mountPath: /nvme-nfs
75 -   lifecycleHooks:
76 -     postStart:
77 -       exec:
78 -         command:
79 -           - 'bash'
80 -           - '-c'
81 -           - >
82 -             pip install nbgitpuller && gitpuller https://github.com/m-bull/spark-toy-notebook main spark-toy-notebook
83 -
```

Note: Only comments from the original package values will be preserved.

DEPLOY 0.1.0-DEV.0.MASTER.13

RESTORE DEFAULTS

Applications

search apps...

Show apps in all namespaces

DEPLOY

jupyterhub-proxy



Helm chart for deploying a Zenith client to proxy a Kubernetes service

deployed

App: zenith-proxy vc1589d5
Package: 0.1.0-dev.0.main.149

jupyterhub



JupyterHub with Zenith integration, NFS and a custom git repository

deployed

App: jupyterhub-zenith-nfs-customrepo vca596ec
Package: 0.1.0-dev.0.master.13

Kubernetes | rcp-cloud-portal-... portal.apps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/kubernetes

Cluster details for jhub-demo

Overview

Cluster details	
Name	jhub-demo
Template	v1.22.6
Kubernetes version	1.22.6
Status	✓ Ready
Autohealing?	✓ Enabled
Node groups	1
Created	an hour ago

Control plane

Status	✓ Ready
Size	vm.alaska.cpu.general.small
Node Count	3

Nodes

Services	
JupyterHub	🔗
Applications	🔗
Kubernetes Dashboard	🔗
Monitoring	🔗

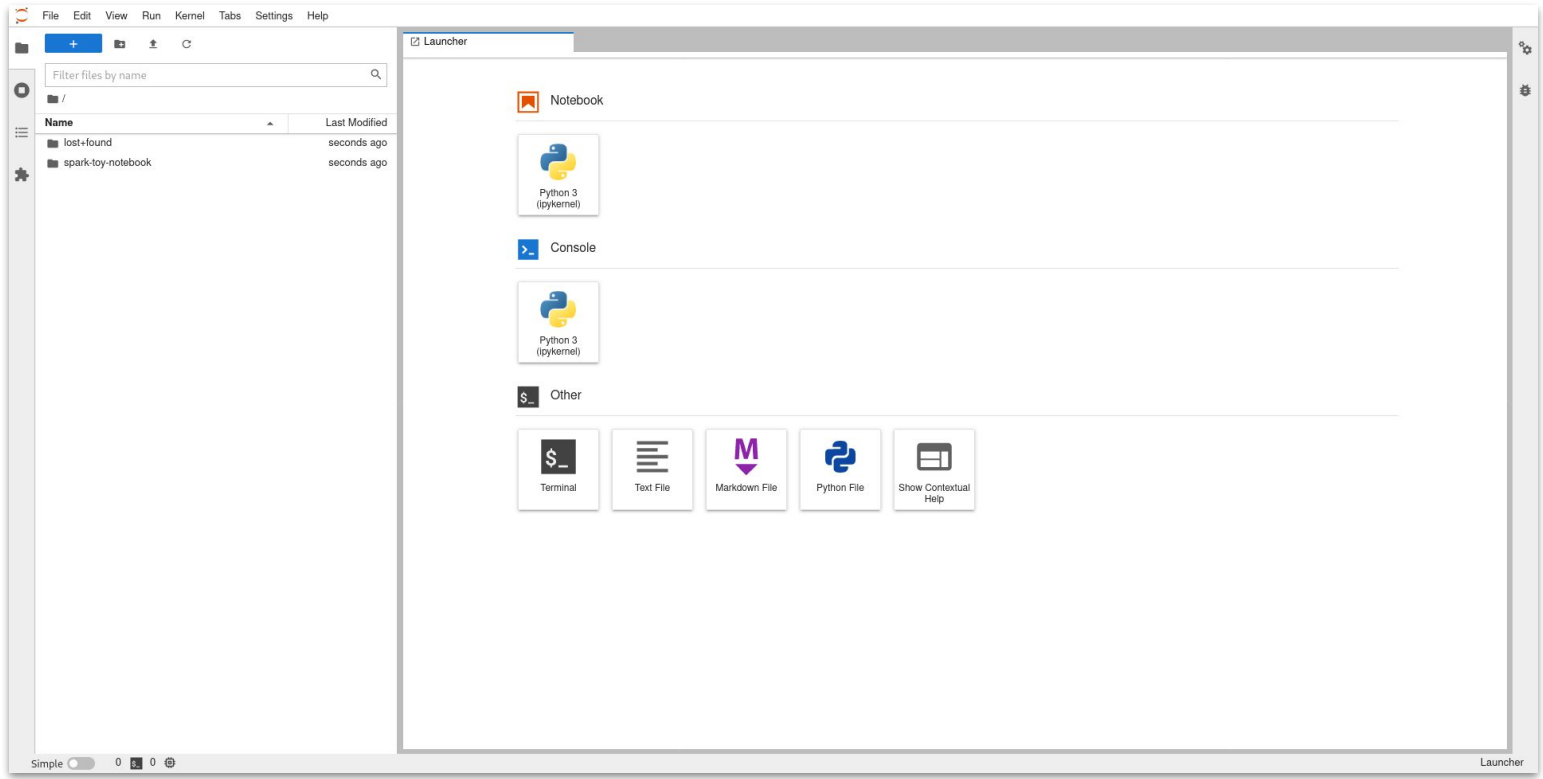
Cluster addons

ccm-openstack	✓ Ready
cloud-config	✓ Ready
cni-cilium	✓ Ready
csi-cinder	✓ Ready
jupyterhub-proxy	✓ Ready
kubeapps	✓ Ready
kubeapps-proxy	✓ Ready
kubernetes-dashboard	✓ Ready
kubernetes-dashboard-proxy	✓ Ready
metrics-server	✓ Ready
monitoring	✓ Ready
monitoring-proxy	✓ Ready
node-feature-discovery	✓ Ready
nvidia-gpu-operator	✓ Ready

Refresh

actions

Close



File Edit View Run Kernel Tabs Settings Help

Terminal 1

```
joyan@jupyter-matta-5fstack:~$ df -h
Filesystem                Size      Used Avail Use% Mounted on
overlay                   50G        28G   28G  41% /
tmpfs                     64M         0   64M   0% /dev
tmpfs                     19G         0   19G   0% /sys/fs/cgroup
172.24.231.125:/nfs-local-claim/default-rook-nfs-pv-claim-pvc-3e2cea9-b0f2-49b2-aeaf-771bee31e405 2.0T      80M   1.9T   1% /nvme-nfs
/dev/vde                  9.8G       37M   9.8G   1% /home/joyan
/dev/vda1                 50G        28G   28G  41% /etc/hosts
shm                       64M         0   64M   0% /dev/shm
tmpfs                     19G         0   19G   0% /proc/acpi
tmpfs                     19G         0   19G   0% /proc/scsi
tmpfs                     19G         0   19G   0% /sys/firmware
```

joyan@jupyter-matta-5fstack:~\$

Simple 1 0

Terminal 1

Questions?

StackHPC