Using OpenStack to reduce HPC service complexity

... no, that is not an oxymoron!

John Garbutt, Principal Engineer, StackHPC
5th February 2022

StackHPC

Why build a Supercomputer with OpenStack?

StackHPC

Green500 Data

Rank	TOP500 Rank	System	Cores	Rmax (TFlop/s)	Power (kW)	Power Efficiency (GFlops/watts)
1	301	**MN-3** - MN-Core Server, Xeon Platinum 8260M 24C 2.4GHz, Preferred Networks MN-Core, MN-Core DirectConnect, Preferred Networks Preferred Networks Japan	1,664	2,181.2	55	39.379
2	291	**SSC-21 Scalable Module** - Apollo 6500 Gen10 plus, AMD EPYC 7543 32C 2.8GHz, NVIDIA A100 80GB, Infiniband HDR200, HPE Samsung Electronics South Korea	16,704	2,274.1	103	33.983
3	295	**Tethys** - NVIDIA DGX A100 Liquid Cooled Prototype, AMD EPYC 7742 64C 2.25GHz, NVIDIA A100 80GB, Infiniband HDR, Nvidia NVIDIA Corporation United States	19,840	2,255.0	72	31.538
4	280	**Wilkes-3** - PowerEdge XE8545, AMD EPYC 7763 64C 2.45GHz, NVIDIA A100 80GB, Infiniband HDR200 dual rail, DELL EMC University of Cambridge United Kingdom	26,880	2,287.0	74	30.797
5	30	**HiPerGator AI** - NVIDIA DGX A100, AMD EPYC 7742 64C 2.25GHz, NVIDIA A100, Infiniband HDR, Nvidia University of Florida United States	138,880	17,200.0	583	29.521

Traditional HPC

StackHPC

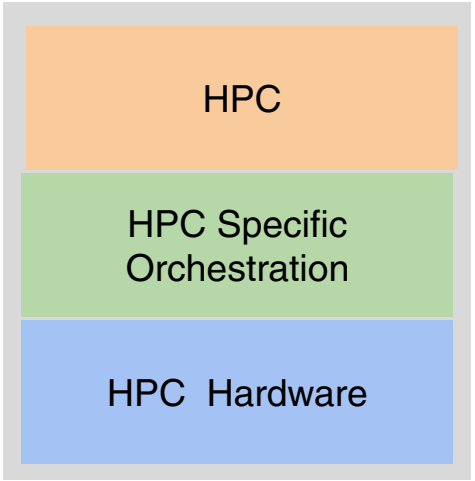

Traditional HPC and TREs

StackHPC

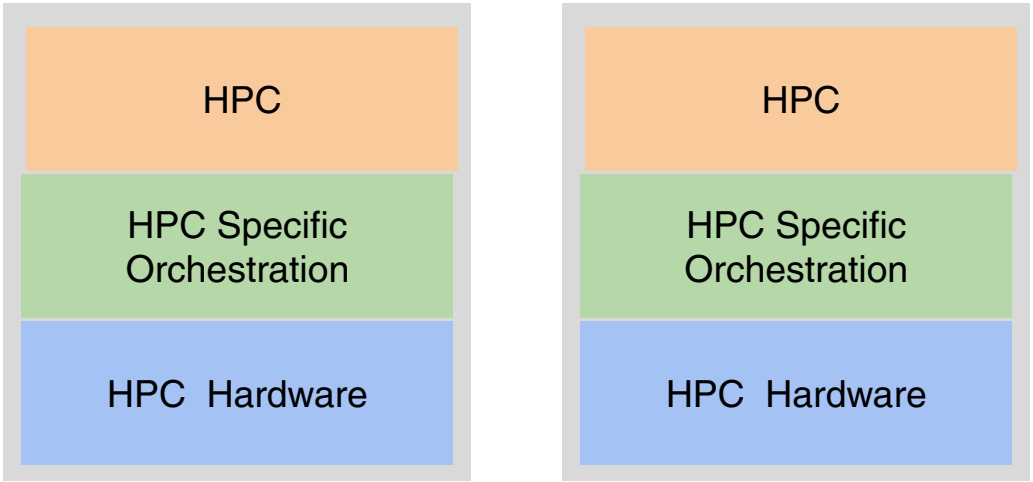

Traditional HPC and HPDA

StackHPC

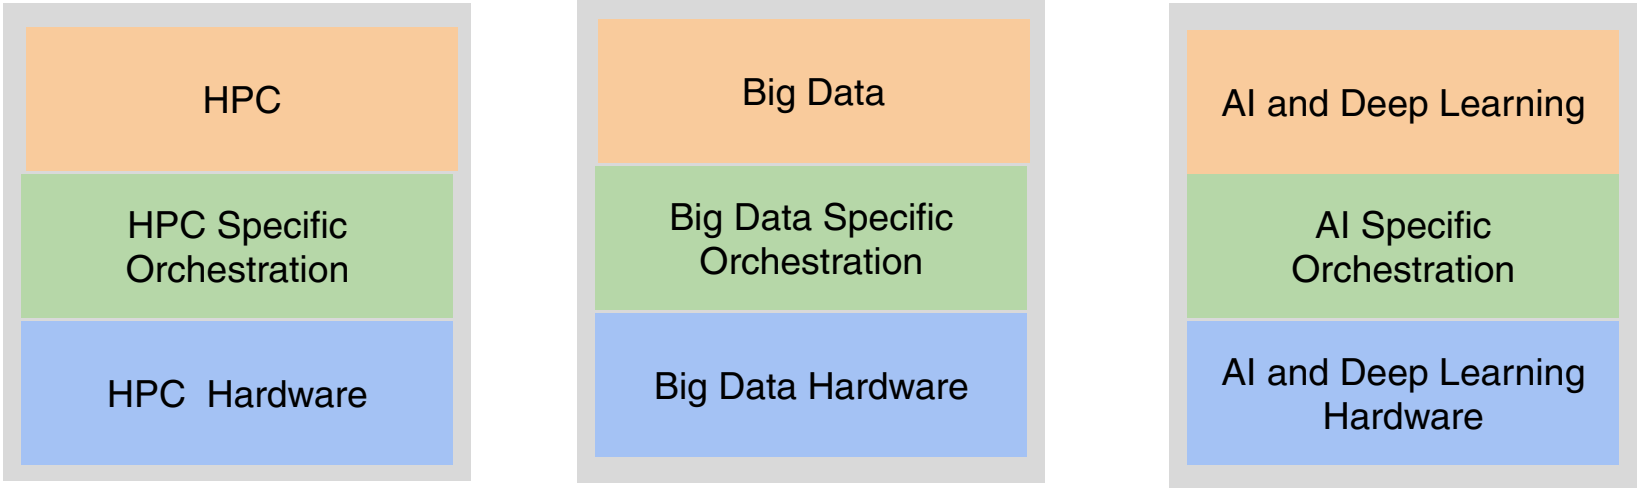

Evolving User Requirements

StackHPC

Evolution away from a single cluster:

- Trusted Research Environments (TREs)
- GPU enabled Kubernetes Cluster?
- Partitions with specialist hardware types
- Hadoop based Big Data processing
- Dedicated AI platforms
- Large data sets
- “Bring your own” interactive data platforms
- …

Key Challenges

StackHPC

- Managing Complexity
- Sharing Knowledge
- Performance vs Flexibility
- Maintaining High Utilization
- Enable new use cases, while keeping the lights on

What does Success look like?

StackHPC

Adopting DevOps & ResOps in HPC

StackHPC

Four key measures of Software Delivery Performance:

- **Lead Time:**
  from customer request to being satisfied
- **Mean Time to Restore (MTTR):**
  failure will happen, get good recovery
- **Change Fail Percentage:**
  a proxy for quality throughout the process
- **Deployment Frequency:**
  a proxy for small batch size

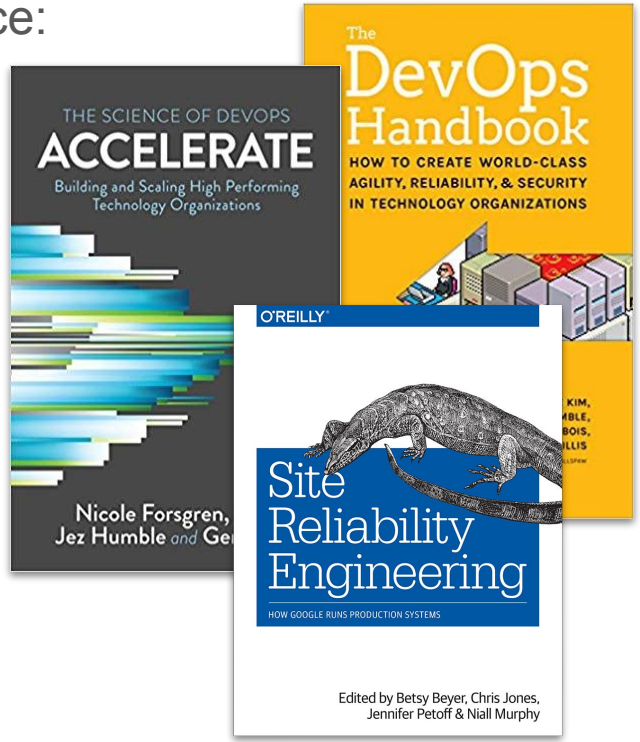

Traditional HPC and HPDA

StackHPC

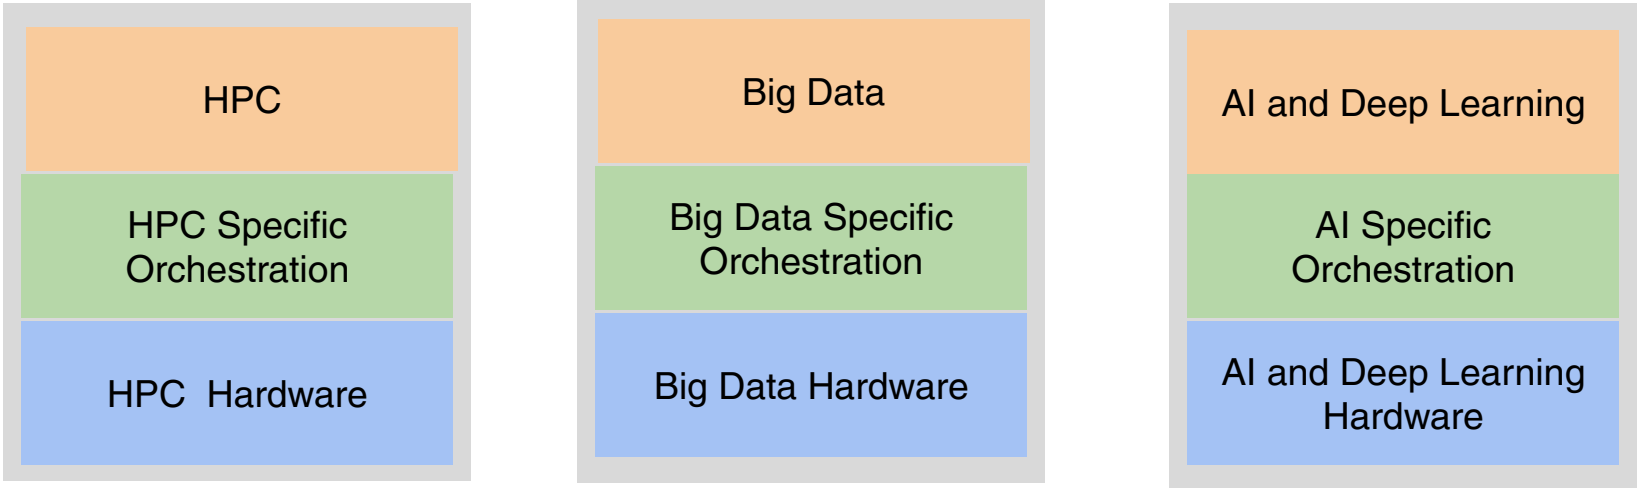

Supercomputing Clouds

StackHPC

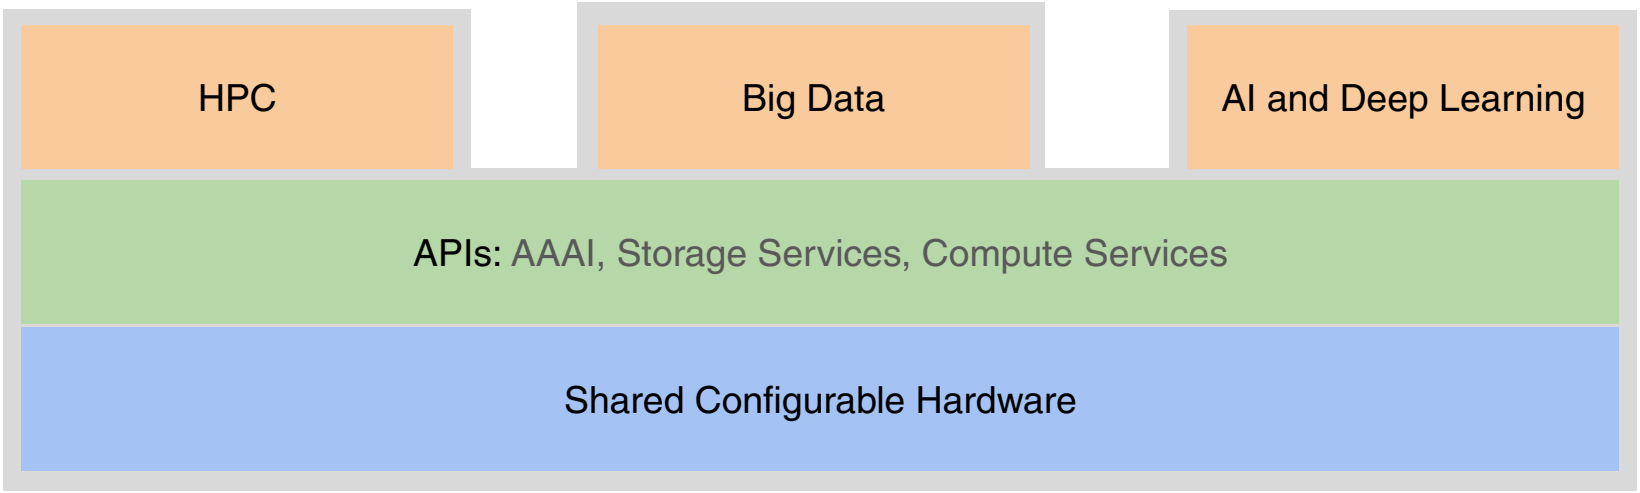

HPC 2.0

StackHPC

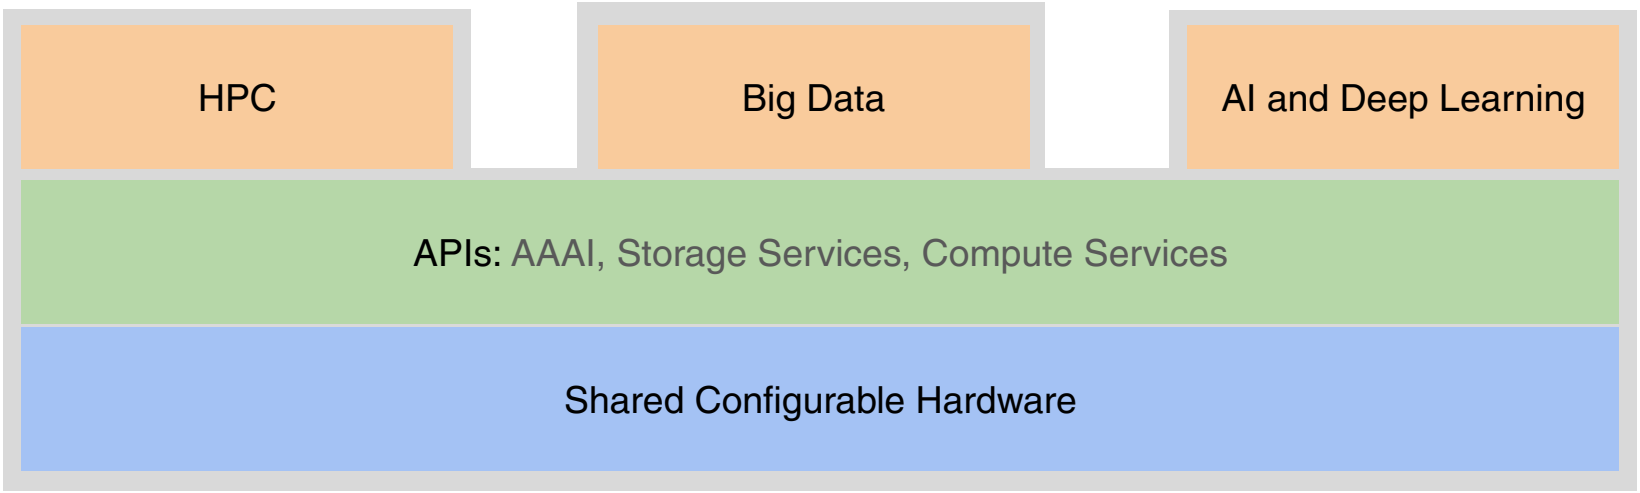

Supercomputing Cloud

StackHPC

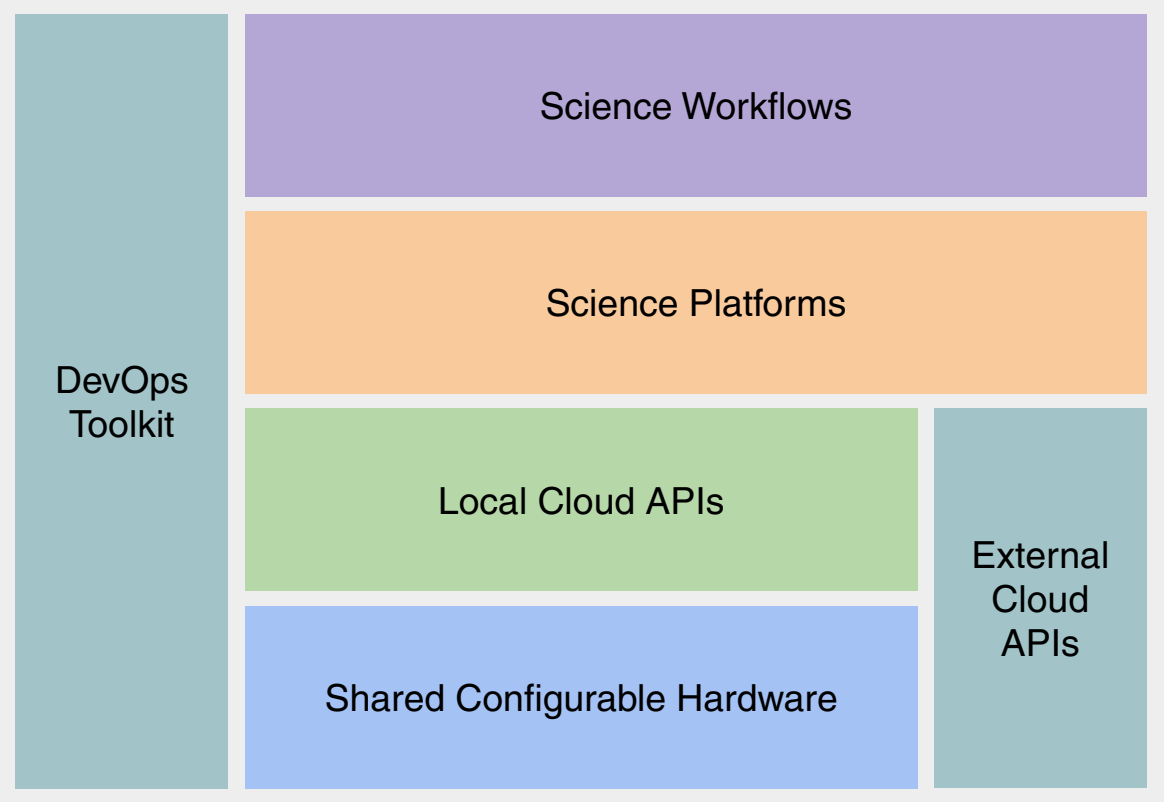

Supercomputing Cloud

StackHPC

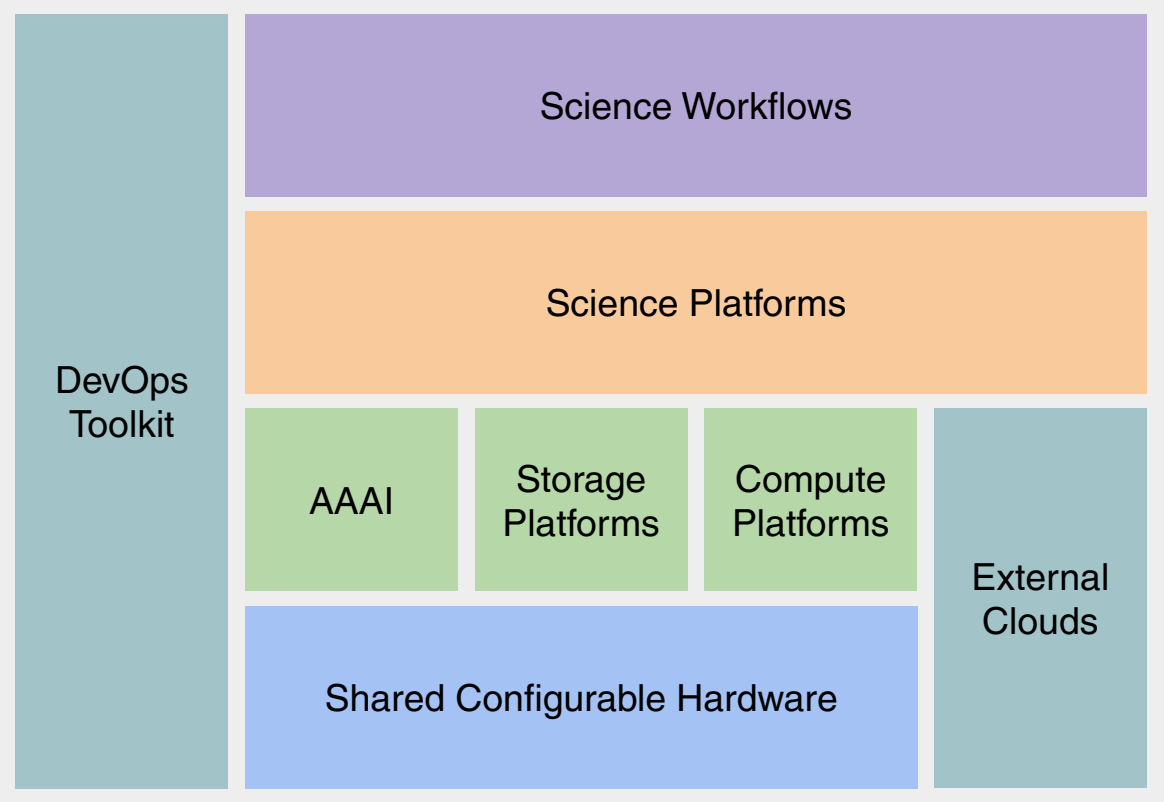

Supercomputing Cloud

StackHPC

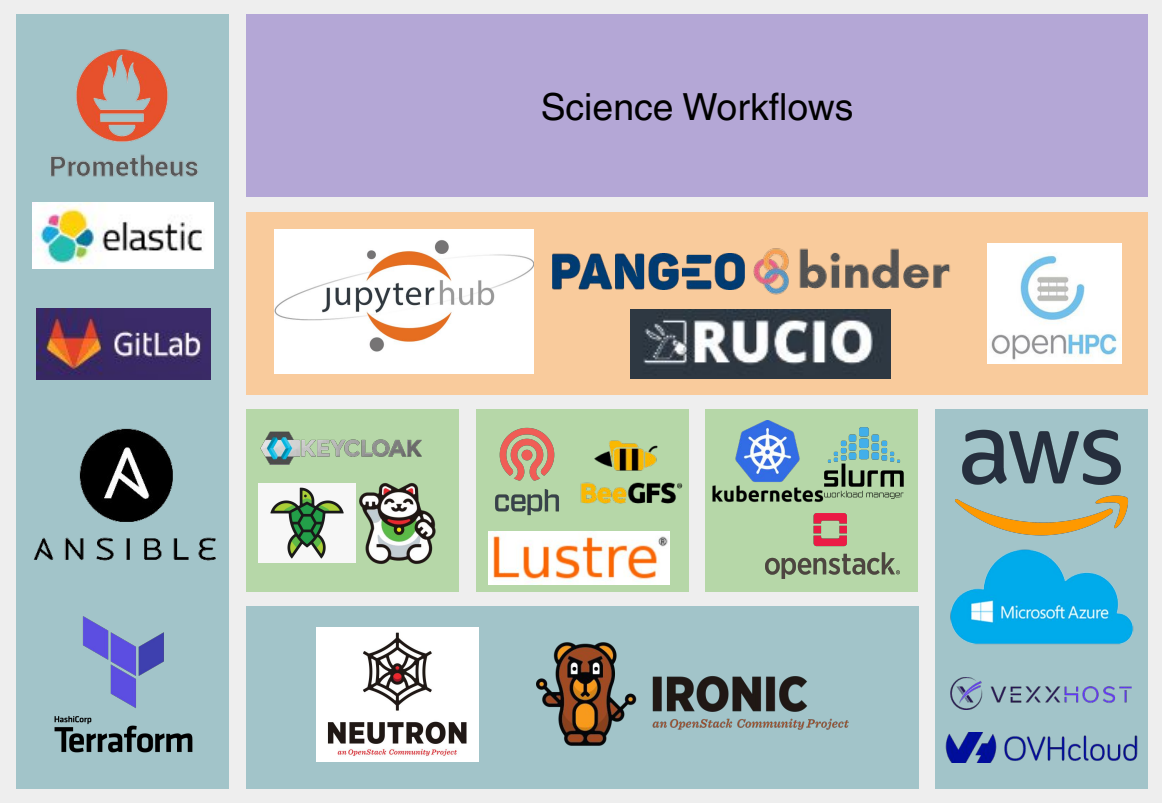

Cambridge HPCS OpenStack Journey

StackHPC

In the Beginning (c. 2015)

- Capture use-cases for Software Defined HPC
- Help establish a Research Computing SIG within OpenStack Foundation
- Establish Cambridge as thought-leader in Scientific OpenStack
- Leverage SKA use-case and develop OpenStack Kayobe

StackHPC

OpenStack in Cambridge (c. 2020)

StackHPC

- Clinical Cloud (2016+)
  - Virtualised Cloud for Brain Imaging, etc
- OpenCB (2016+)
  - Genomics Platform
- AlaSKA (2016+)
  - Bare metal platform for SKA-SDP prototyping
- STFC-IRIS (2018+)
  - Virtual and bare metal cloud for medium scale HPC
- Secure Research Computing Platform (2018+)
  - Virtualised resource for medical informatics
- Arcus (2020+)
  - Unified OpenStack, including Large-scale HPC

Green500 Data

Rank	TOP500 Rank	System	Cores	Rmax (TFlop/s)	Power (kW)	Power Efficiency (GFlops/watts)
1	301	**MN-3** - MN-Core Server, Xeon Platinum 8260M 24C 2.4GHz, Preferred Networks MN-Core, MN-Core DirectConnect, Preferred Networks Preferred Networks Japan	1,664	2,181.2	55	39.379
2	291	**SSC-21 Scalable Module** - Apollo 6500 Gen10 plus, AMD EPYC 7543 32C 2.8GHz, NVIDIA A100 80GB, Infiniband HDR200, HPE Samsung Electronics South Korea	16,704	2,274.1	103	33.983
3	295	**Tethys** - NVIDIA DGX A100 Liquid Cooled Prototype, AMD EPYC 7742 64C 2.25GHz, NVIDIA A100 80GB, Infiniband HDR, Nvidia NVIDIA Corporation United States	19,840	2,255.0	72	31.538
4	280	**Wilkes-3** - PowerEdge XE8545, AMD EPYC 7763 64C 2.45GHz, NVIDIA A100 80GB, Infiniband HDR200 dual rail, DELL EMC University of Cambridge United Kingdom	26,880	2,287.0	74	30.797
5	30	**HiPerGator AI** - NVIDIA DGX A100, AMD EPYC 7742 64C 2.25GHz, NVIDIA A100, Infiniband HDR, Nvidia University of Florida United States	138,880	17,200.0	583	29.521

Shared Configurable Hardware

StackHPC

Supercomputing Cloud

StackHPC

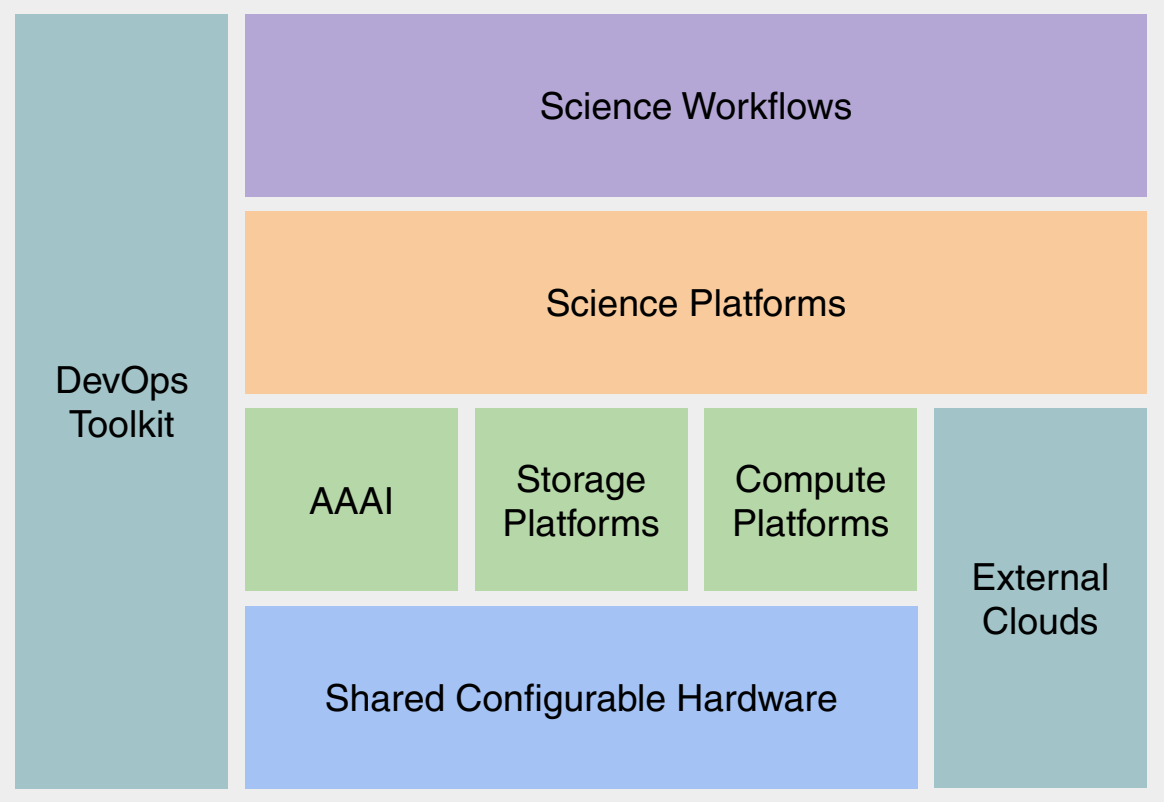

OpenStack Kayobe Architecture

StackHPC

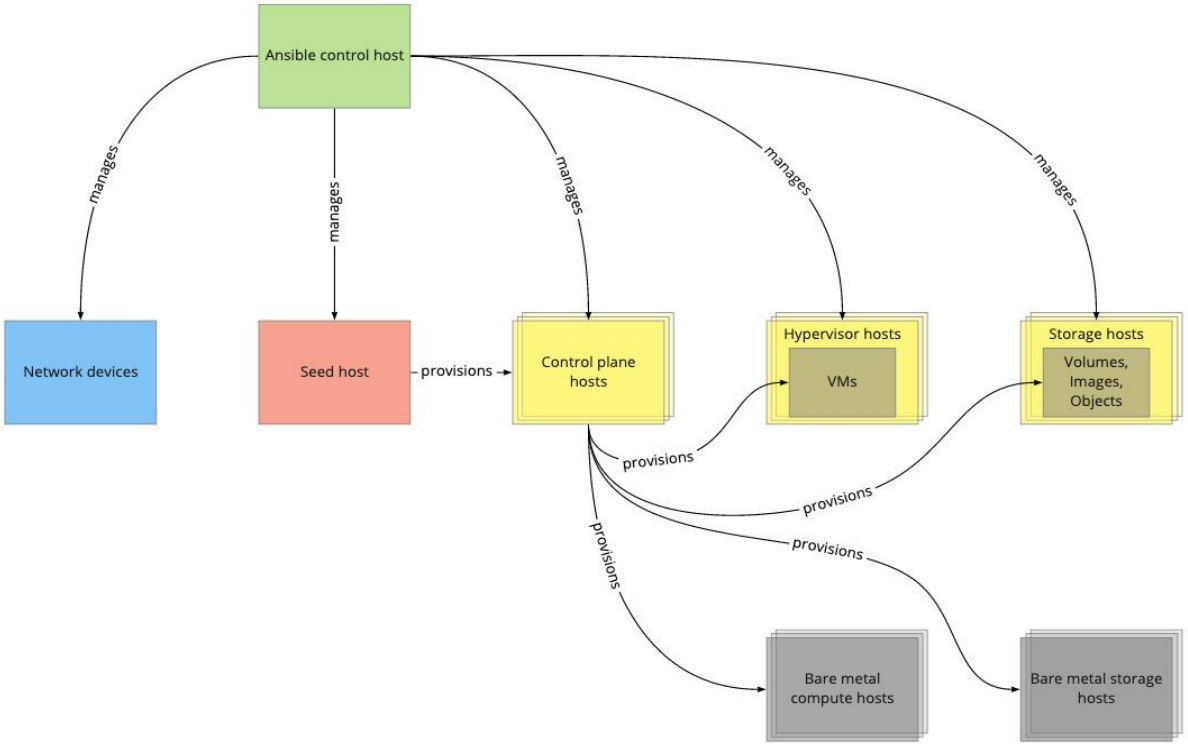

IRONIC
an OpenStack Community Project

OpenStack Kayobe Deployment

StackHPC

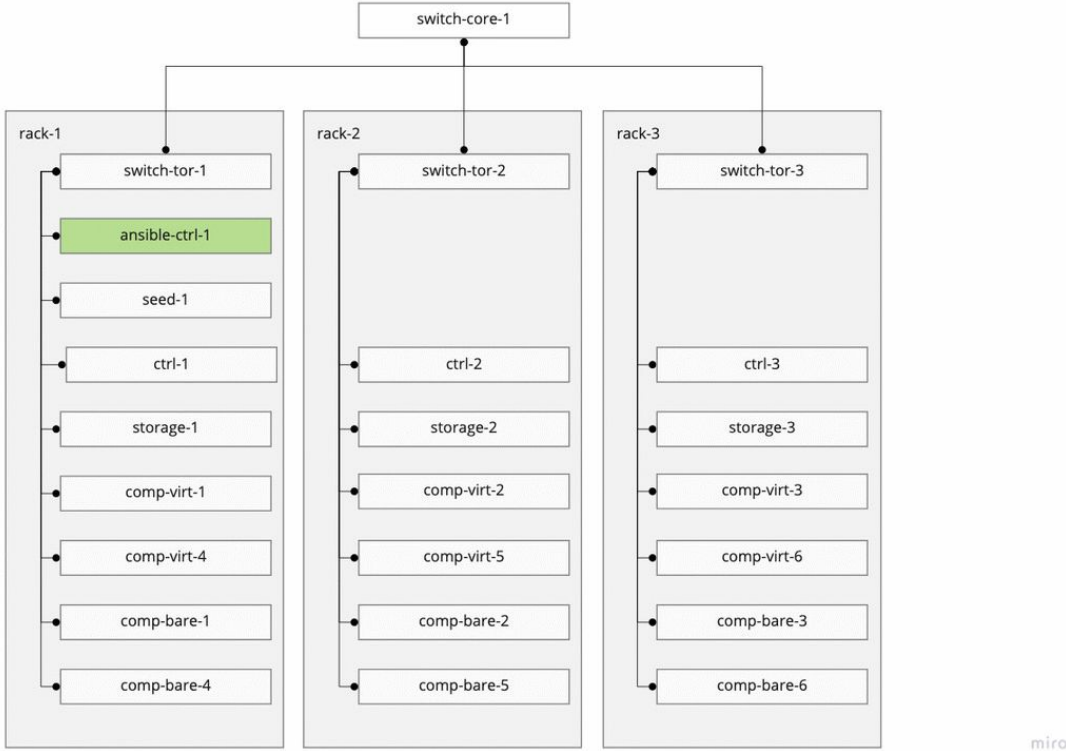

IRONIC
an OpenStack Community Project

OpenStack Kayobe Deployment

StackHPC

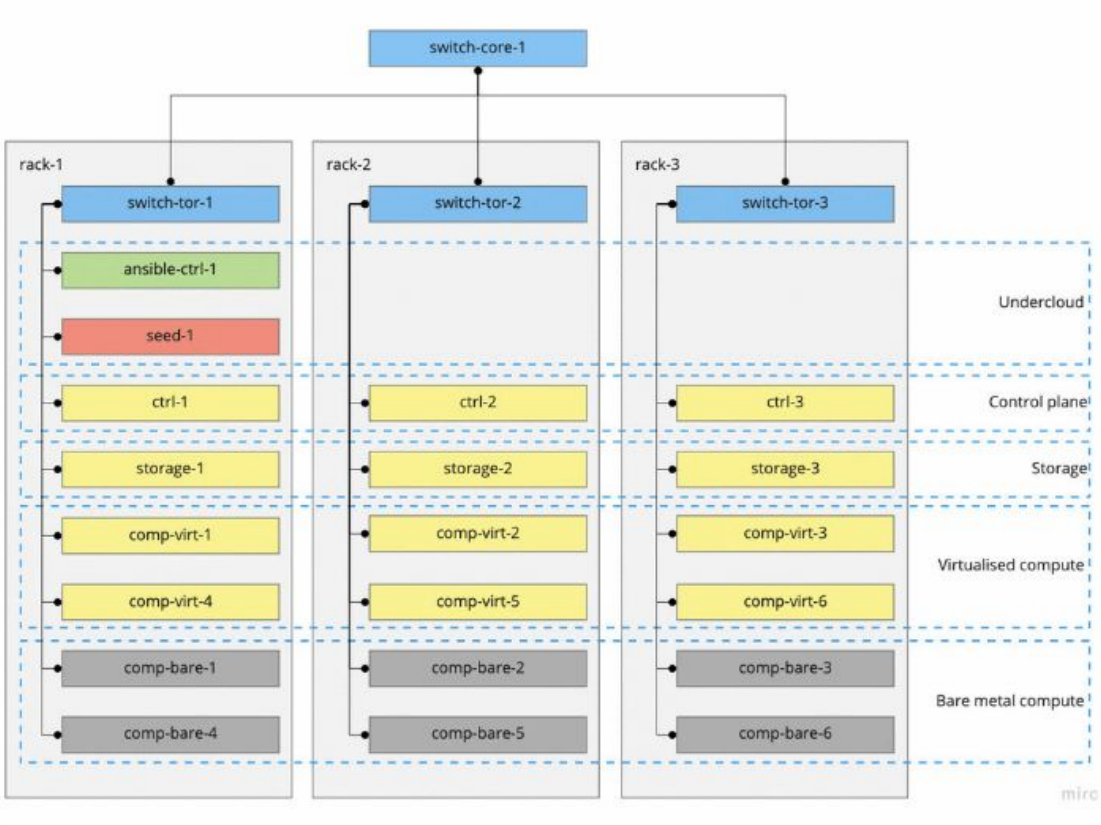

IRONIC
an OpenStack Community Project

Baremetal Node <-> Hypervisor

StackHPC

- Kayobe and Kolla-Ansible, OpenStack
  - Network switches configured via ansible in Kayobe
- Ironic to deploy Controllers and Hypervisors
  - Controllers via Bifrost Ironic
  - Hypervisors are a baremetal workload
  - Science Platforms can use baremetal and/or VM servers
- Storage
  - Lustre and Ceph are applications running on baremetal
- Dynamic Networking
  - OpenStack Neutron used to configure physical switches
  - networking-generic-switch can change the access VLAN
  - Dedicated networks for inspecting, provisioning and cleaning

KOLLA

an OpenStack Community Project

Bootstrapping Physical Infrastructure

StackHPC

- Barcode scan servers in racks
- iDRAC DHCP
  - MAC from scanning, Neutron for DHCP
- Ansible driven Ironic Enrollment
  - Ansible enables IPMI, updates Firmware
  - State machine stored in Ironic
- Ironic Inspector
  - First on 1GbE, and update ConnectX-6
  - Re-inspection on 50GbE, LLDP to get switch info
- Testing: Burn-in with HPL, MPI ping tests
- https://github.com/stackhpc/arcus-terraform-idrac

Tracking Bootstrapping

StackHPC

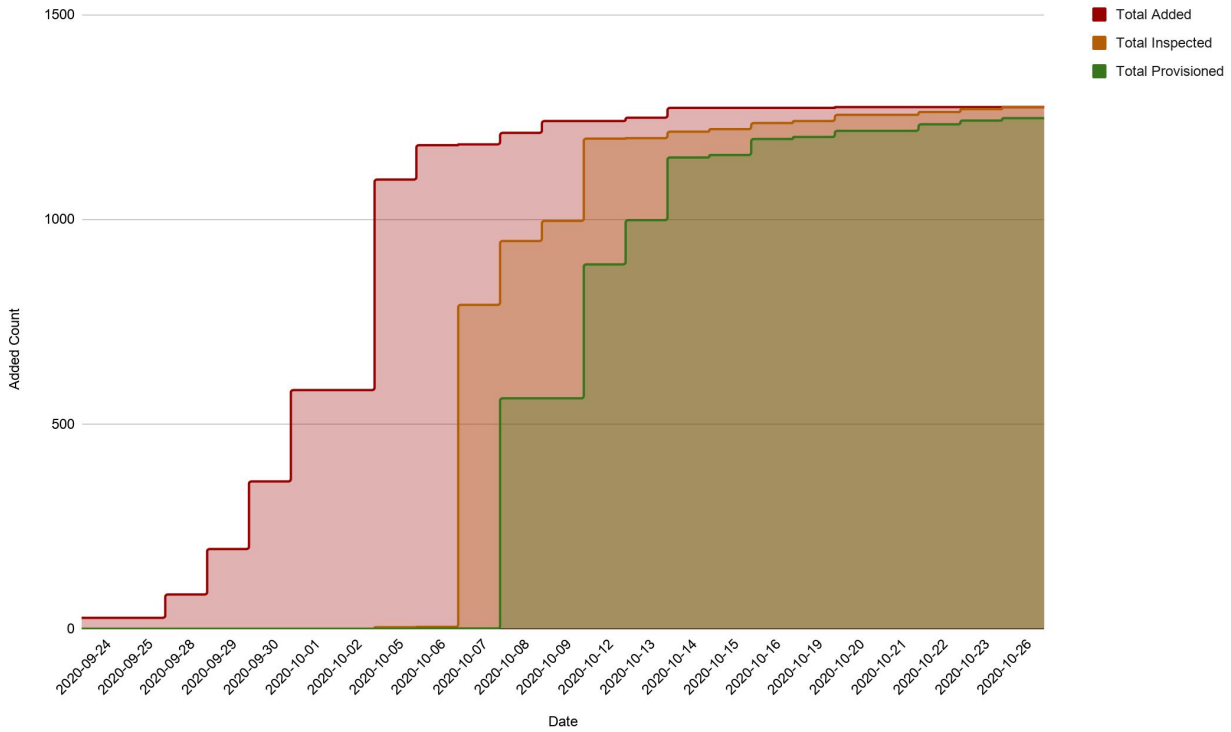

Tracking Bootstrapping

StackHPC

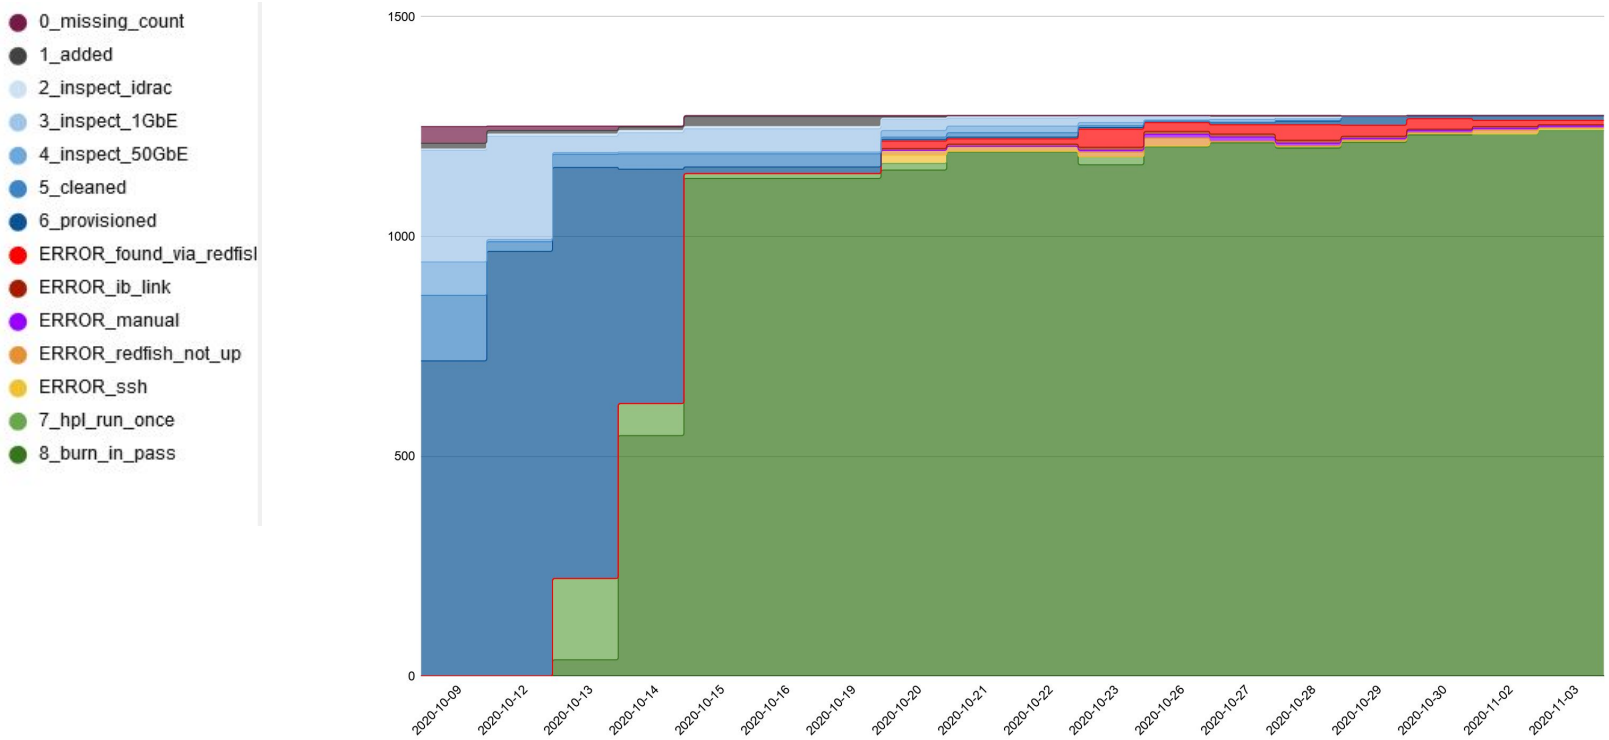

Tracking Bootstrapping

StackHPC

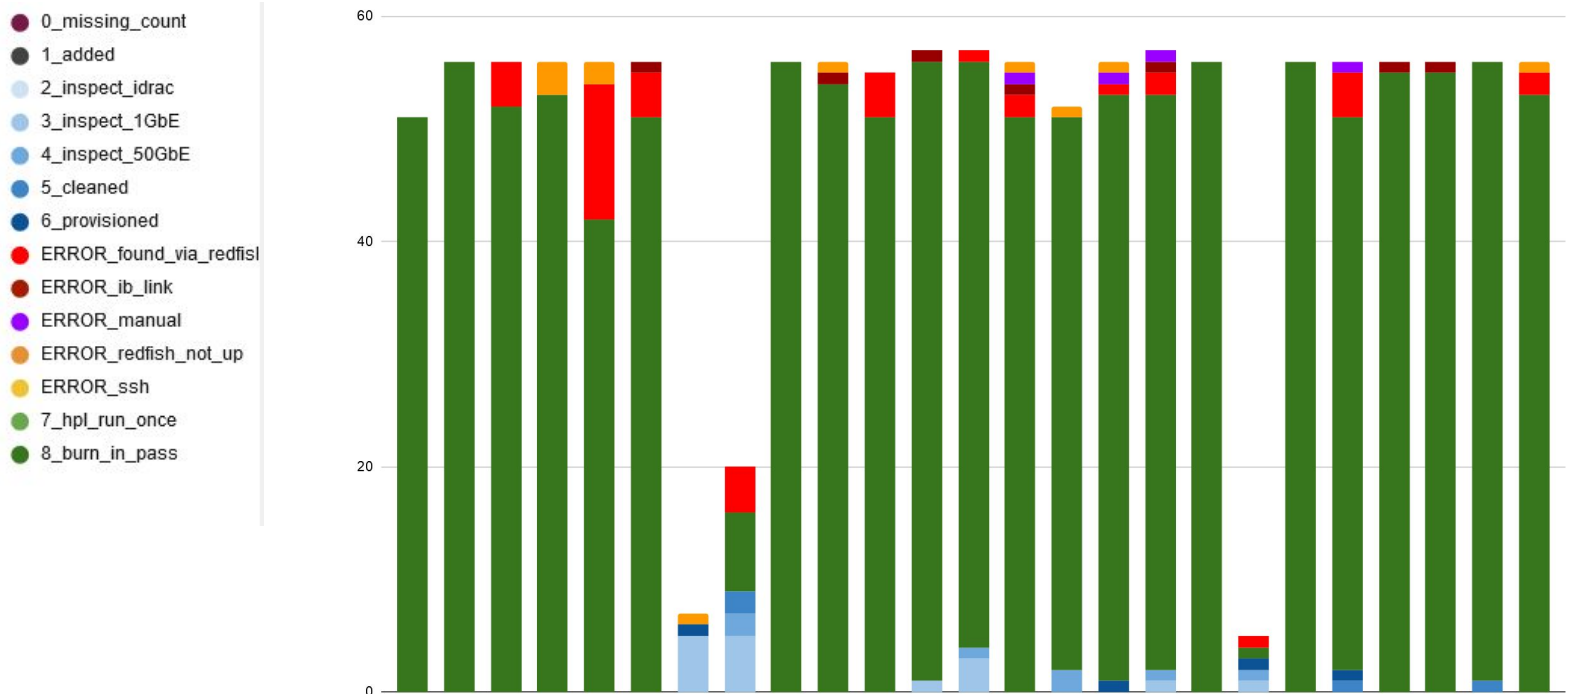

Science Platform:
OpenHPC Slurm

StackHPC

Supercomputing Cloud

StackHPC

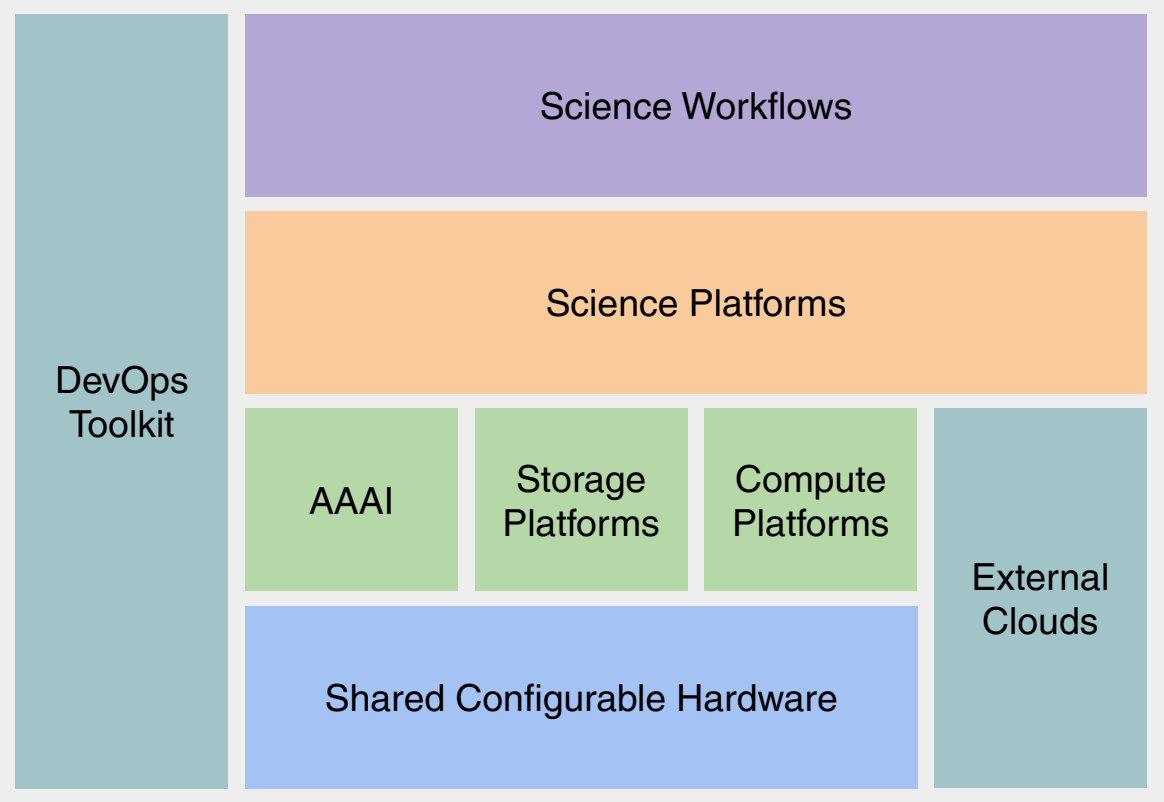

OpenHPC Slurm on OpenStack

StackHPC

- Ansible install and configure OpenHPC packages
  - Terraform can manage infrastructure
  - Open OnDemand web interface
  - Monitoring using Grafana and Prometheus
  - Self-tests: HPL, MPI ping pong
- (Optional) Image build pipeline, using Packer and Ansible
- Non-impacting upgrades
  - Slurm rebuild scripts
  - OpenStack rebuild
- Autoscaling (coming soon)
- https://github.com/stackhpc/ansible-slurm-appliance

Tuning Ironic for Scale

StackHPC

- Target: rebuild all hosts to apply new kernel
- networking-generic-switch
  - Added cumulus driver to networking-generic-switch
  - Added experimental “batching” of config
- Ironic-conductor configuration
  - IMPI driver, avoids reboot in iDRAC driver
  - Direct deploy, not iSCSI
    - force_raw_images = false
    - Avoid neutron router bottleneck
  - Server Delete: increased rpc_timeout
- connect_timeout in HAproxy and MariaDB

Redfish exporter

2020-11-03 12:16:11 to 2020-11-05 07:52:11

Overview

Power Consumption

Power Consumption

	max	avg	current
h22d5	29.72 kW	24.21 kW	14.41 kW
h22d8	29.59 kW	24.18 kW	14.37 kW
h22b5	29.21 kW	23.68 kW	13.84 kW
h24c5	29.14 kW	23.60 kW	13.84 kW
h23c5	29.12 kW	23.69 kW	13.90 kW
h21d5	29.09 kW	23.56 kW	13.87 kW
h22b8	29.08 kW	23.62 kW	13.86 kW

Power states

server	Power state
svn3-h21a5-u28	ON
svn2-h21a5-u27	ON
svn3-h21a5-u36	ON
svn1-h21a5-u28	ON
svn4-h21a5-u25	ON
svn3-h21a5-u26	ON
svn2-h21a5-u25	ON
svn1-h21a5-u26	ON
svn4-h21a5-u23	ON
svn3-h21a5-u24	ON
svn2-h21a5-u23	ON
svn1-h21a5-u24	ON
svn2-h21a5-u35	ON
svn4-h21a5-u13	ON

1 2 3 4 5 6 7 8 9

Chassis status

server	Status
svn2-h23c8-u35	ERROR
svn4-h22c5-u27	ERROR
svn3-h22c5-u28	ERROR
svn2-h22c5-u27	ERROR
svn1-h22c5-u28	ERROR
svn4-h23c5-u33	ERROR
svn3-h23c5-u34	ERROR
svn2-h23c5-u33	ERROR
svn1-h23c5-u34	ERROR
svn4-h24d8-u31	HEALTHY
svn3-h24d8-u4	HEALTHY
svn2-h22d8-u35	HEALTHY
svn2-h24d8-u5	HEALTHY
svn3-h24d8-u6	HEALTHY

1 2 3 4 5 6 7 8 9

Max Inlet Temp

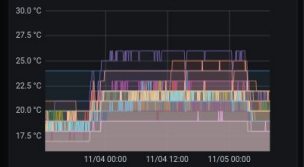

	max	avg	current
max inlet h22b8	26.00 °C	24.20 °C	21.00 °C
max inlet h24c5	25.00 °C	22.59 °C	21.00 °C
max inlet h21c5	25.00 °C	22.74 °C	22.00 °C
max inlet h22d8	24.00 °C	24.00 °C	24.00 °C
max inlet h24b8	23.00 °C	21.36 °C	22.00 °C

Max CPU1 Temp

	max	avg	current
h23c8	82.00 °C	67.03 °C	52.00 °C
h23d5	77.00 °C	67.33 °C	51.00 °C
h24b8	80.00 °C	68.56 °C	50.00 °C
h24d8	77.00 °C	65.81 °C	47.00 °C
h21c8	79.00 °C	67.78 °C	47.00 °C

Max CPU2 Temp

	max	avg	current
h23c8	72.00 °C	61.78 °C	50.00 °C
h24d8	72.00 °C	61.22 °C	48.00 °C
h23d5	72.00 °C	63.61 °C	47.00 °C
h22a5	72.00 °C	63.23 °C	44.00 °C
h21c8	75.00 °C	62.97 °C	44.00 °C

Powered ON by Rack

	max	current
Powered up h22d8	57.00	57.00
Powered up h22d5	57.00	57.00
Powered up h24d5	56.00	56.00
Powered up h24c5	56.00	56.00
Powered up h24b8	56.00	56.00

Powered OFF by Rack

	max	current
Powered down h24d8	1.000	
Powered down h23e8	1.000	
Powered down h22a8	4.000	
Powered down h21d5	1.000	

Max Inlet Temp

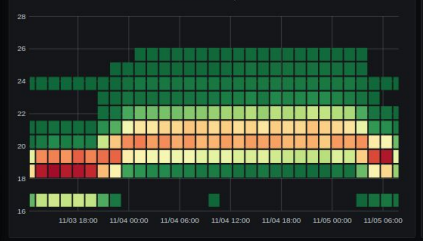

Max server fan speed

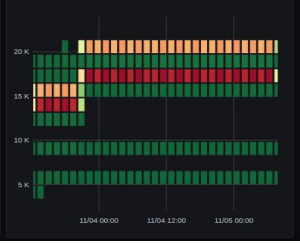

CPU1 Temp

CPU2 Temp

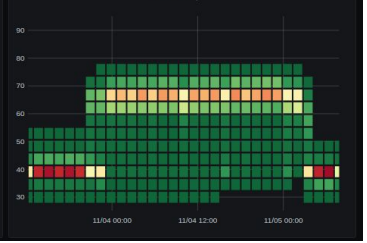

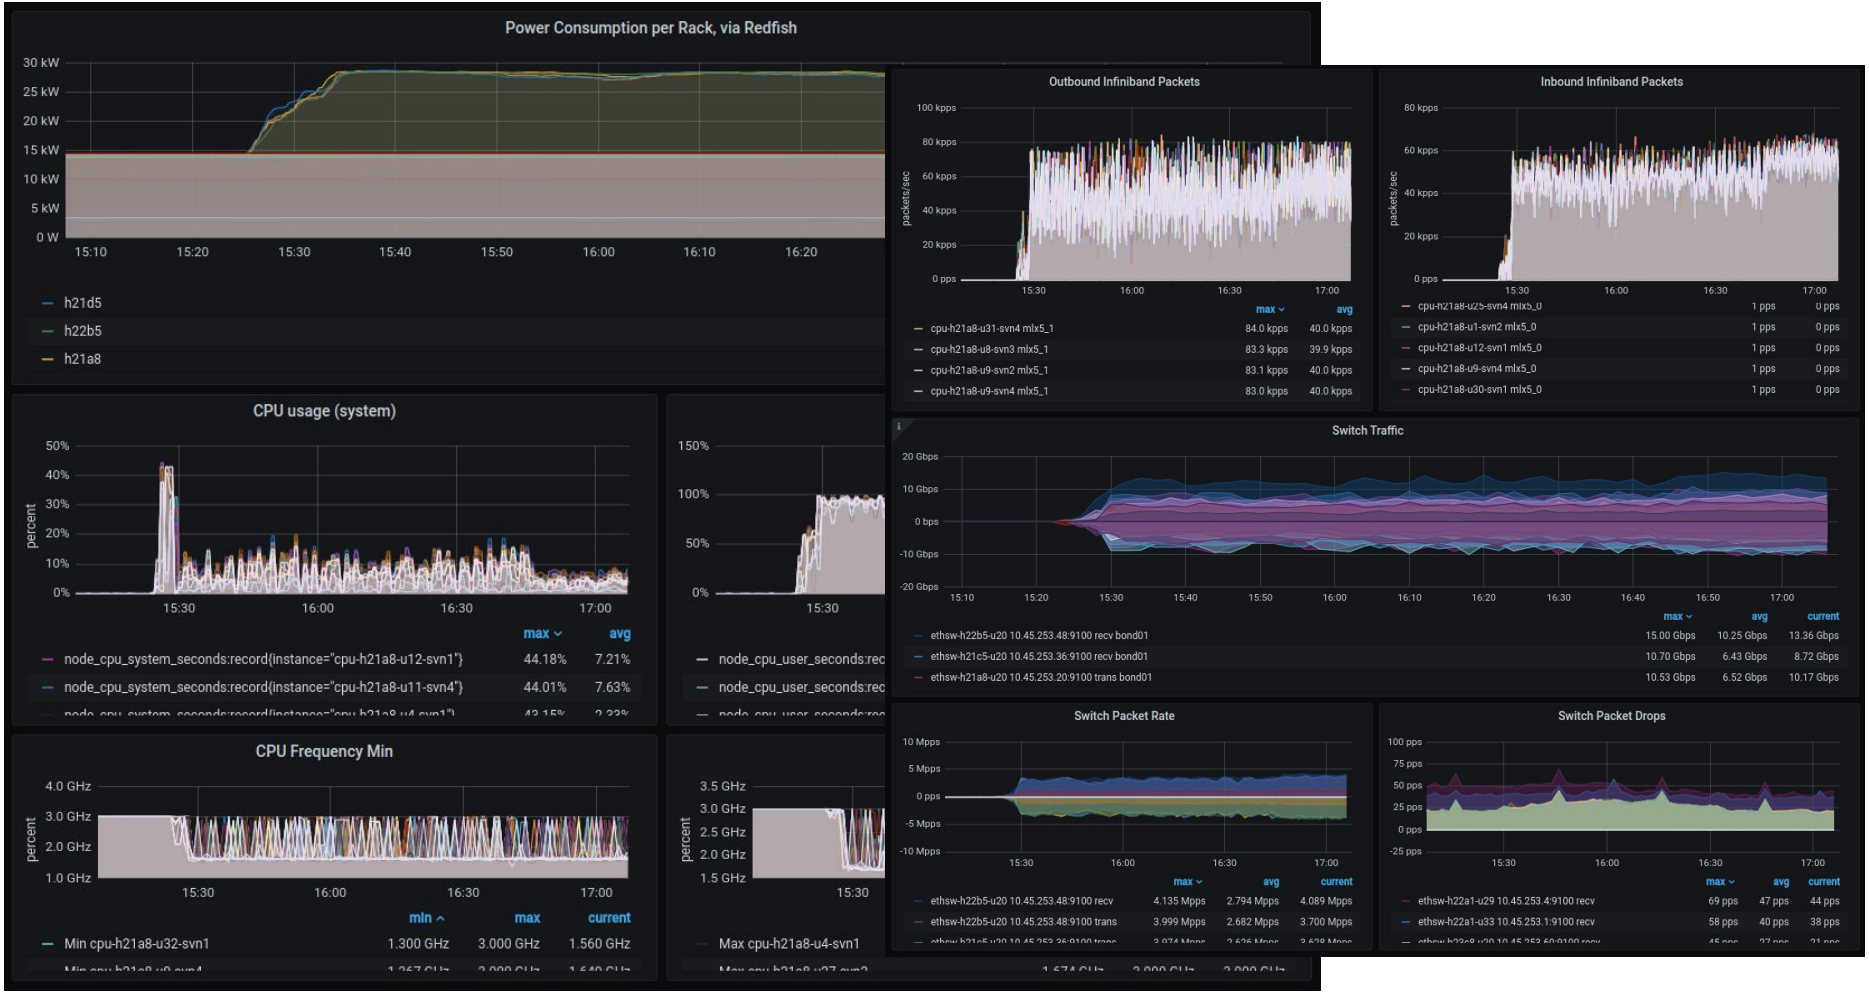
Power Consumption per Rack, via Redfish
30 kW
25 kW
20 kW
15 kW
10 kW
5 kW
0 W
15:10
15:20
15:30
15:40
15:50
16:00
16:10
16:20
h21d5
h22b5
h21a8
Outbound Infiniband Packets
100 kpps
80 kpps
60 kpps
40 kpps
20 kpps
0 pps
packets/sec
15:30
16:00
16:30
17:00
max
avg
cpu-h21a8-u31-svn4 mlx5_1 84.0 kpps 40.0 kpps
cpu-h21a8-u8-svn3 mlx5_1 83.3 kpps 39.9 kpps
cpu-h21a8-u9-svn2 mlx5_1 83.1 kpps 40.0 kpps
cpu-h21a8-u9-svn4 mlx5_1 83.0 kpps 40.0 kpps
Inbound Infiniband Packets
80 kpps
60 kpps
40 kpps
20 kpps
0 pps
packets/sec
15:30
16:00
16:30
17:00
cpu-h21a8-u26-svn4 mlx5_0 1 pps 0 pps
cpu-h21a8-u1-svn2 mlx5_0 1 pps 0 pps
cpu-h21a8-u12-svn1 mlx5_0 1 pps 0 pps
cpu-h21a8-u9-svn4 mlx5_0 1 pps 0 pps
cpu-h21a8-u30-svn1 mlx5_0 1 pps 0 pps
CPU usage (system)
50%
40%
30%
20%
10%
0%
percent
15:30
16:00
16:30
17:00
max
avg
node_cpu_system_seconds:record{instance="cpu-h21a8-u12-svn1"} 44.18% 7.21%
node_cpu_system_seconds:record{instance="cpu-h21a8-u11-svn4"} 44.01% 7.63%
150%
100%
50%
0%
15:30
node_cpu_user_seconds:rec
node_cpu_user_seconds:rec
Switch Traffic
20 Gbps
10 Gbps
0 bps
-10 Gbps
-20 Gbps
15:10
15:20
15:30
15:40
15:50
16:00
16:10
16:20
16:30
16:40
16:50
17:00
max
avg
current
ethsw-h22b5-u20 10.45.253.48:9100 recv bond01 15.00 Gbps 10.25 Gbps 13.36 Gbps
ethsw-h21c5-u20 10.45.253.36:9100 recv bond01 10.70 Gbps 6.43 Gbps 8.72 Gbps
ethsw-h21a8-u20 10.45.253.20:9100 trans bond01 10.53 Gbps 6.52 Gbps 10.17 Gbps
CPU Frequency Min
4.0 GHz
3.0 GHz
2.0 GHz
1.0 GHz
percent
15:30
16:00
16:30
17:00
min
max
current
Min cpu-h21a8-u32-svn1 1.300 GHz 3.000 GHz 1.560 GHz
3.5 GHz
3.0 GHz
2.5 GHz
2.0 GHz
1.5 GHz
percent
15:30
Max cpu-h21a8-u4-svn1
Switch Packet Rate
10 Mpps
5 Mpps
0 pps
-5 Mpps
-10 Mpps
15:30
16:00
16:30
17:00
max
avg
current
ethsw-h22b5-u20 10.45.253.48:9100 recv 4.135 Mpps 2.794 Mpps 4.089 Mpps
ethsw-h22b5-u20 10.45.253.48:9100 trans 3.999 Mpps 2.682 Mpps 3.700 Mpps
Switch Packet Drops
100 pps
75 pps
50 pps
25 pps
0 pps
-25 pps
15:30
16:00
16:30
17:00
max
avg
current
ethsw-h22a1-u29 10.45.253.4:9100 recv 69 pps 47 pps 44 pps
ethsw-h22a1-u33 10.45.253.1:9100 recv 58 pps 40 pps 38 pps


CPU usage (system)
25%
20%
15%
10%
5%
0%
percent
11:40
11:42
11:44
11:46
11:48
11:50
11:52
11:54
11:56
11:58
max
avg
node_cpu_system_seconds:record{instance="proto-hpc-2"} 20.6% 9.12%
node_cpu_system_seconds:record{instance="proto-hpc-1"} 20.6% 8.90%
CPU usage (user)
100%
75%
50%
25%
0%
max
avg
node_cpu_user_seconds:record{instance="proto-hpc-0"} 96.2% 69.7%
node_cpu_user_seconds:record{instance="proto-hpc-3"} 79.9% 37.1%
Network Traffic (transmit)
4 Gb/s
3 Gb/s
2 Gb/s
1 Gb/s
0 b/s
bits in (-) / bits out (+)
max
avg
proto-hpc-0 eth0_out 3.45 Gb/s 1.46 Gb/s
proto-hpc-1 eth0_out 3.13 Gb/s 1.40 Gb/s
proto-hpc-2 eth0_out 2.89 Gb/s 1.31 Gb/s
proto-hpc-3 eth0_out 2.89 Gb/s 1.31 Gb/s
Network Traffic (receive)
4 Gb/s
3 Gb/s
2 Gb/s
1 Gb/s
0 b/s
bits in
max
avg
proto-hpc-3 eth0_in 3.10 Gb/s 1.37 Gb/s
proto-hpc-2 eth0_in 3.09 Gb/s 1.37 Gb/s
proto-hpc-1 eth0_in 3.09 Gb/s 1.39 Gb/s
proto-hpc-0 eth0_in 2.94 Gb/s 1.35 Gb/s

Science Platforms:
On Demand

StackHPC

Supercomputing Cloud

StackHPC

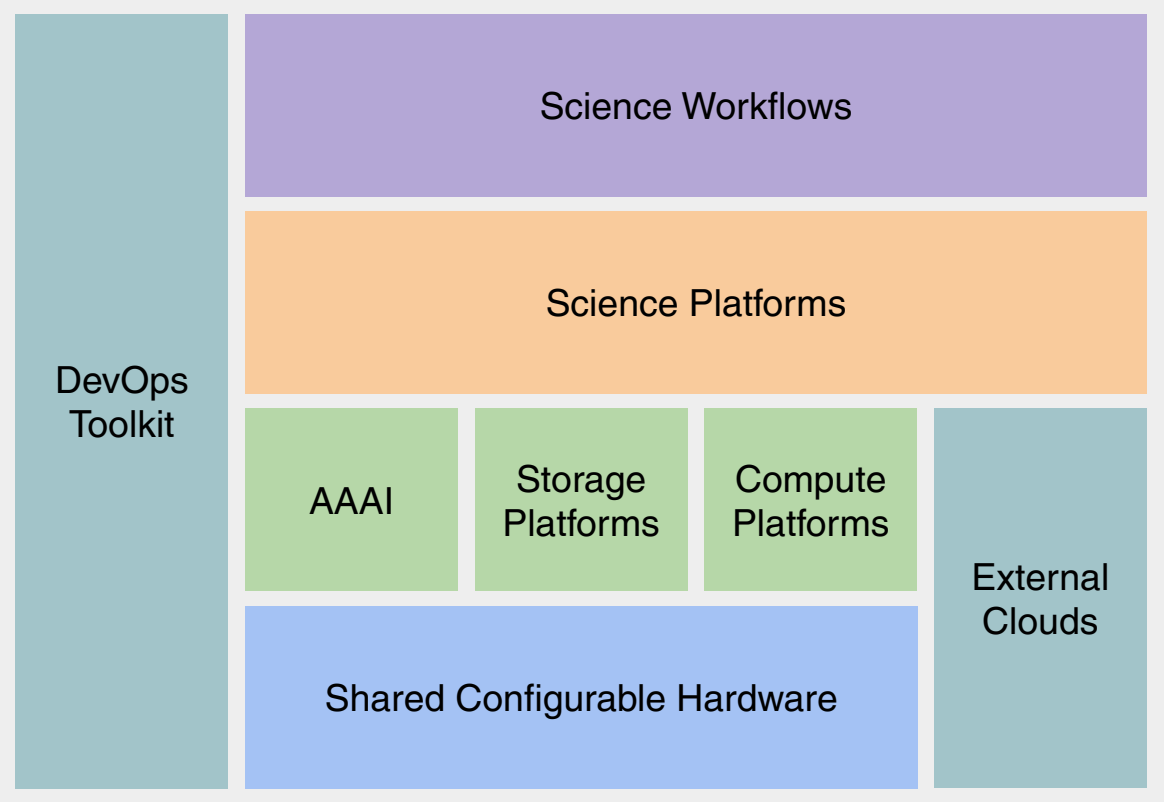

Azimuth Cloud Portal

StackHPC

- Reduce time to science & reduce operational effort of onboarding
- Easier to **find** and **reuse** common lego bricks, between science communities
- Builds on the STFC funded work done by JASMIN
- Target use cases
  - Works with OpenID Connect
  - Get me a:
    - Bigger laptop
    - Slurm Clusters, Kubernetes clusters
    - JupyterHub, Kubeflow, DaskHub, …

https://github.com/stackhpc/azimuth

Science Platforms:
Get me a bigger Laptop!

StackHPC

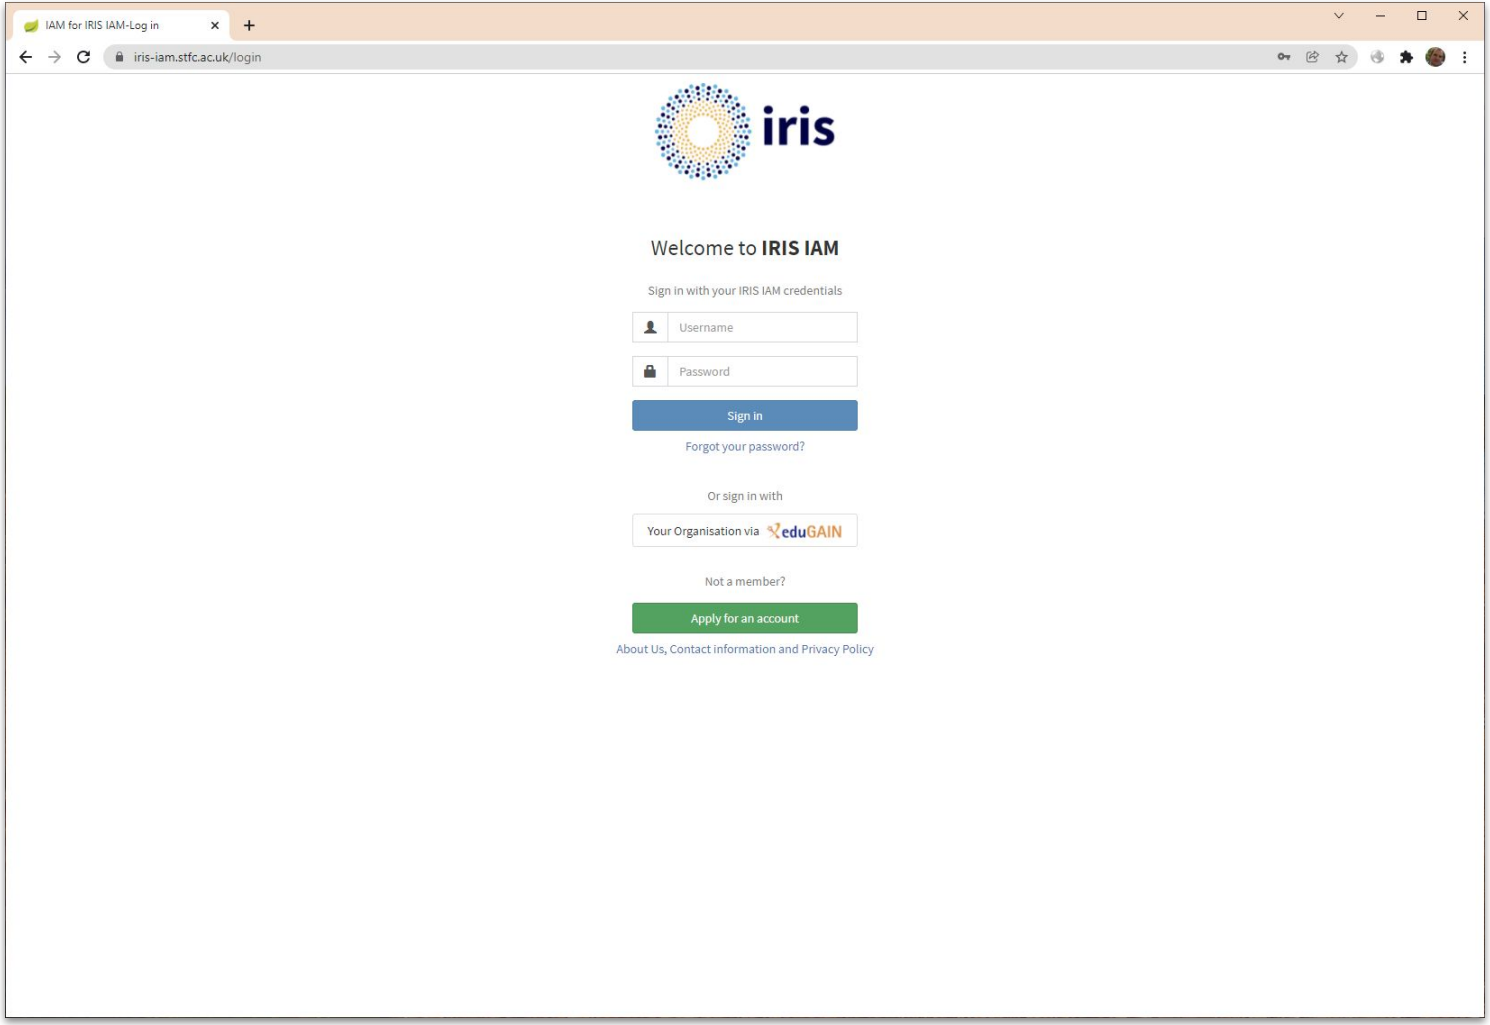
IAM for IRIS IAM-Log in
iris-iam.stfc.ac.uk/login
iris
Welcome to IRIS IAM
Sign in with your IRIS IAM credentials
Username
Password
Sign in
Forgot your password?
Or sign in with
Your Organisation via eduGAIN
Not a member?
Apply for an account
About Us, Contact information and Privacy Policy

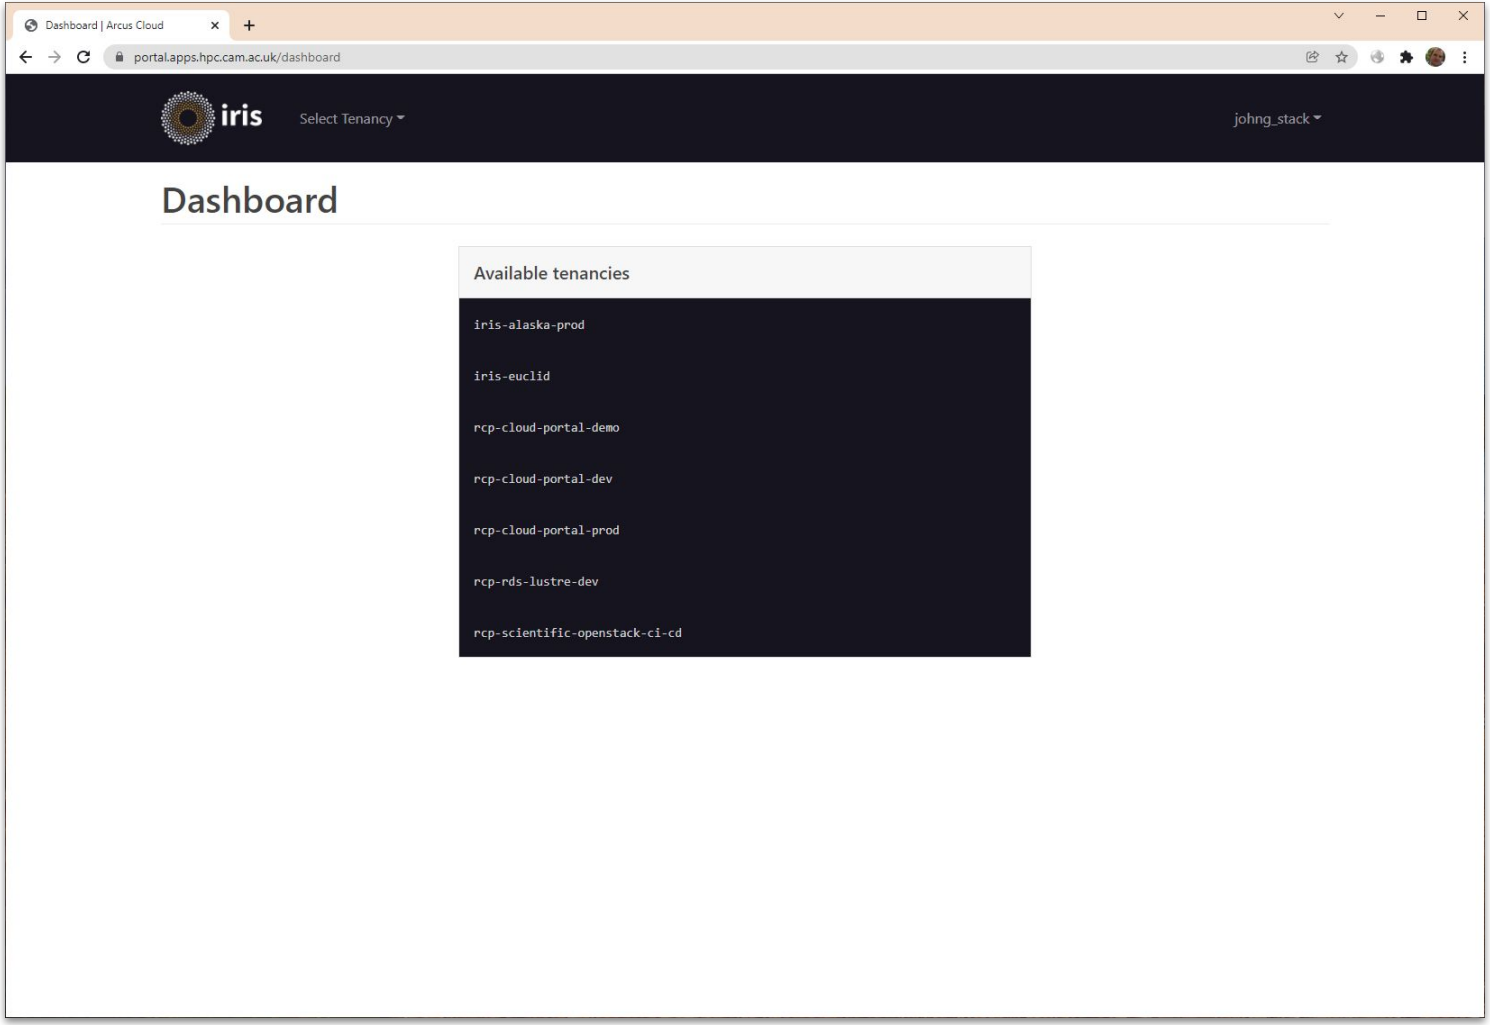
Dashboard | Arcus Cloud
portal.apps.hpc.cam.ac.uk/dashboard
iris
Select Tenancy
johng_stack
Dashboard
Available tenancies
iris-alaska-prod
iris-euclid
rcp-cloud-portal-demo
rcp-cloud-portal-dev
rcp-cloud-portal-prod
rcp-rds-lustre-dev
rcp-scientific-openstack-ci-cd

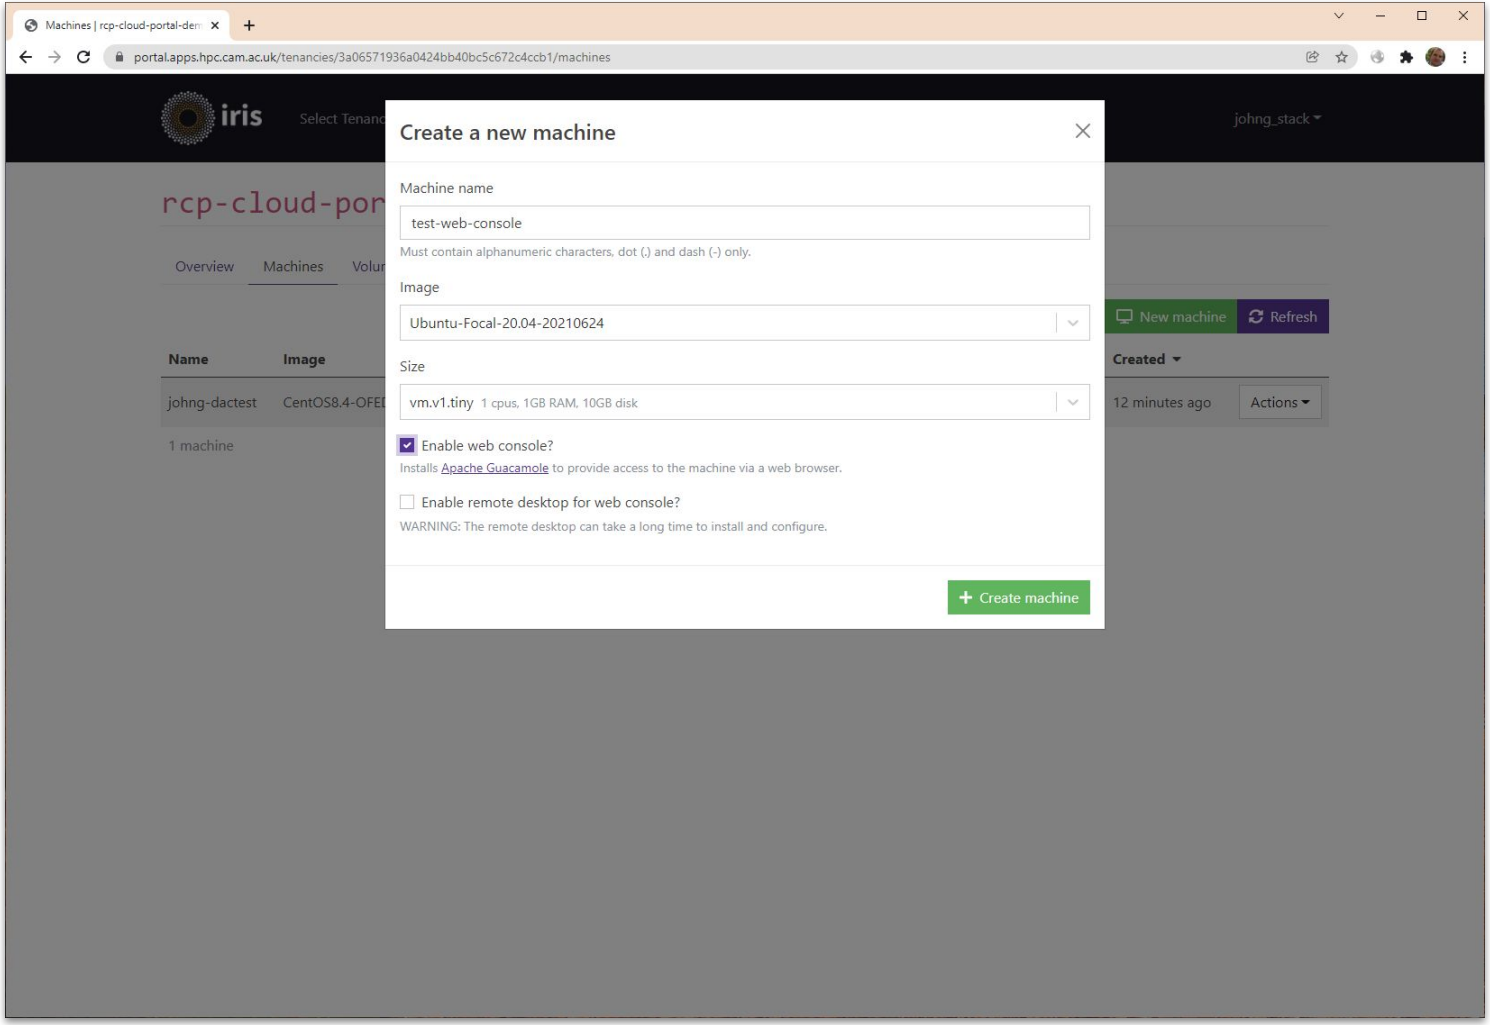
Machines | rcp-cloud-portal-dem
portal.apps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/machines
iris
Select Tenanc
johng_stack
rcp-cloud-por
Overview
Machines
Volur
New machine
Refresh
Name
Image
Created
johng-dactest
CentOS8.4-OFED
12 minutes ago
Actions
1 machine
Create a new machine
Machine name
test-web-console
Must contain alphanumeric characters, dot (.) and dash (-) only.
Image
Ubuntu-Focal-20.04-20210624
Size
vm.v1.tiny 1 cpus, 1GB RAM, 10GB disk
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

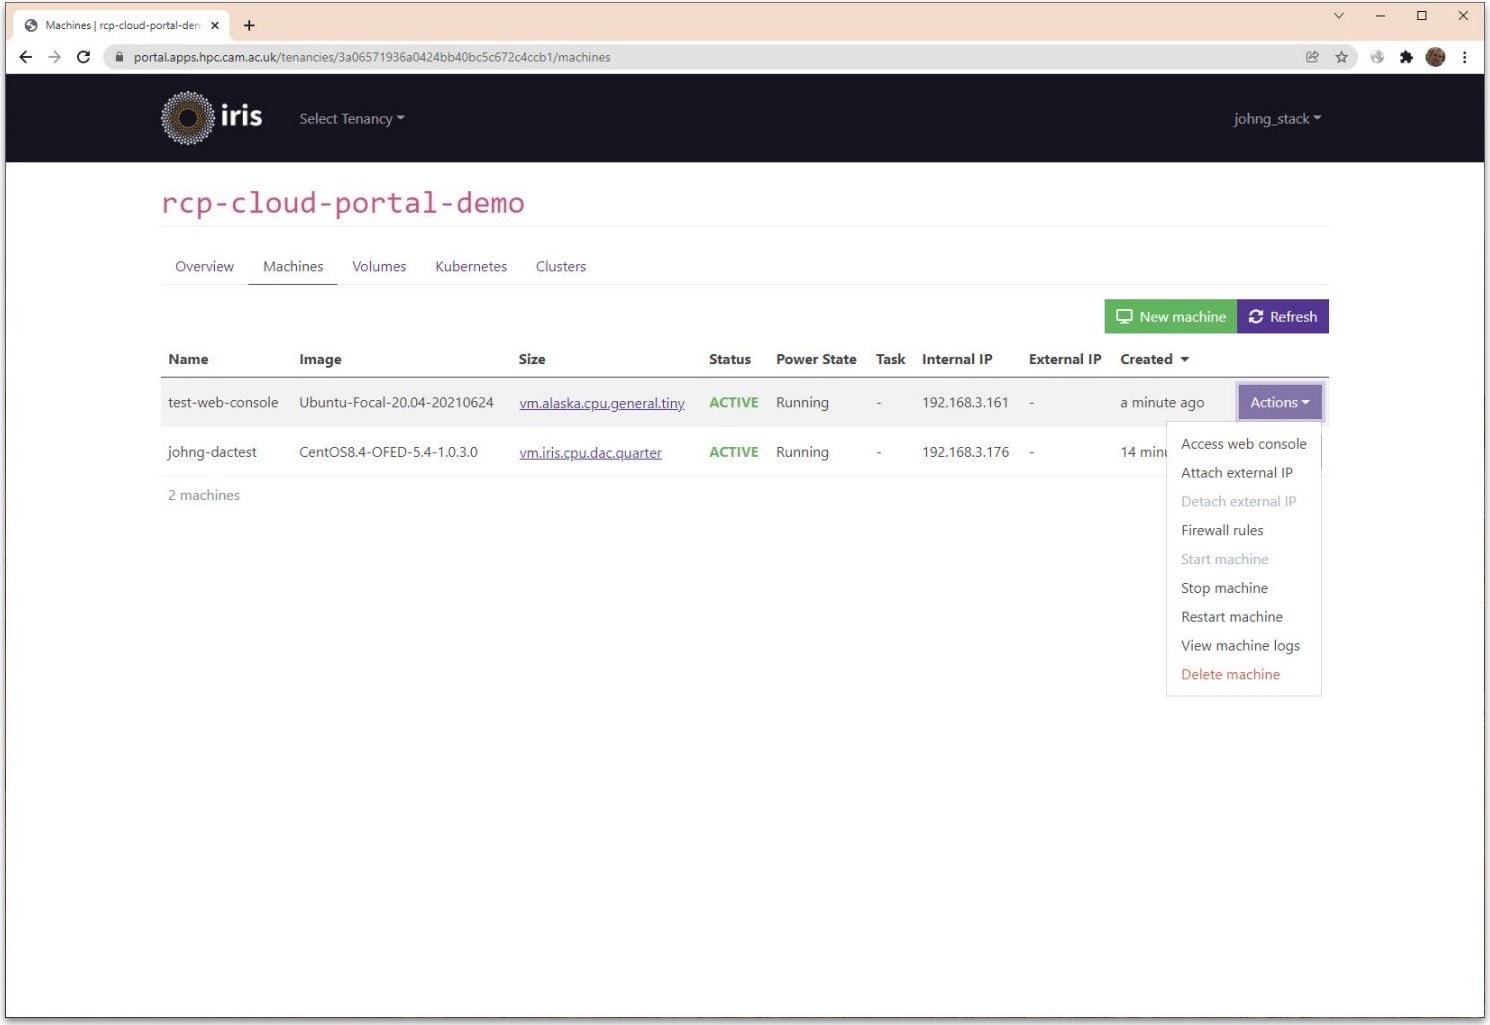
Machines | rcp-cloud-portal-dem
portal.apps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/machines
iris
Select Tenancy
johng_stack
rcp-cloud-portal-demo
Overview
Machines
Volumes
Kubernetes
Clusters
New machine
Refresh
Name
Image
Size
Status
Power State
Task
Internal IP
External IP
Created
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
2 machines
Access web console
Attach external IP
Detach external IP
Firewall rules
Start machine
Stop machine
Restart machine
View machine logs
Delete machine

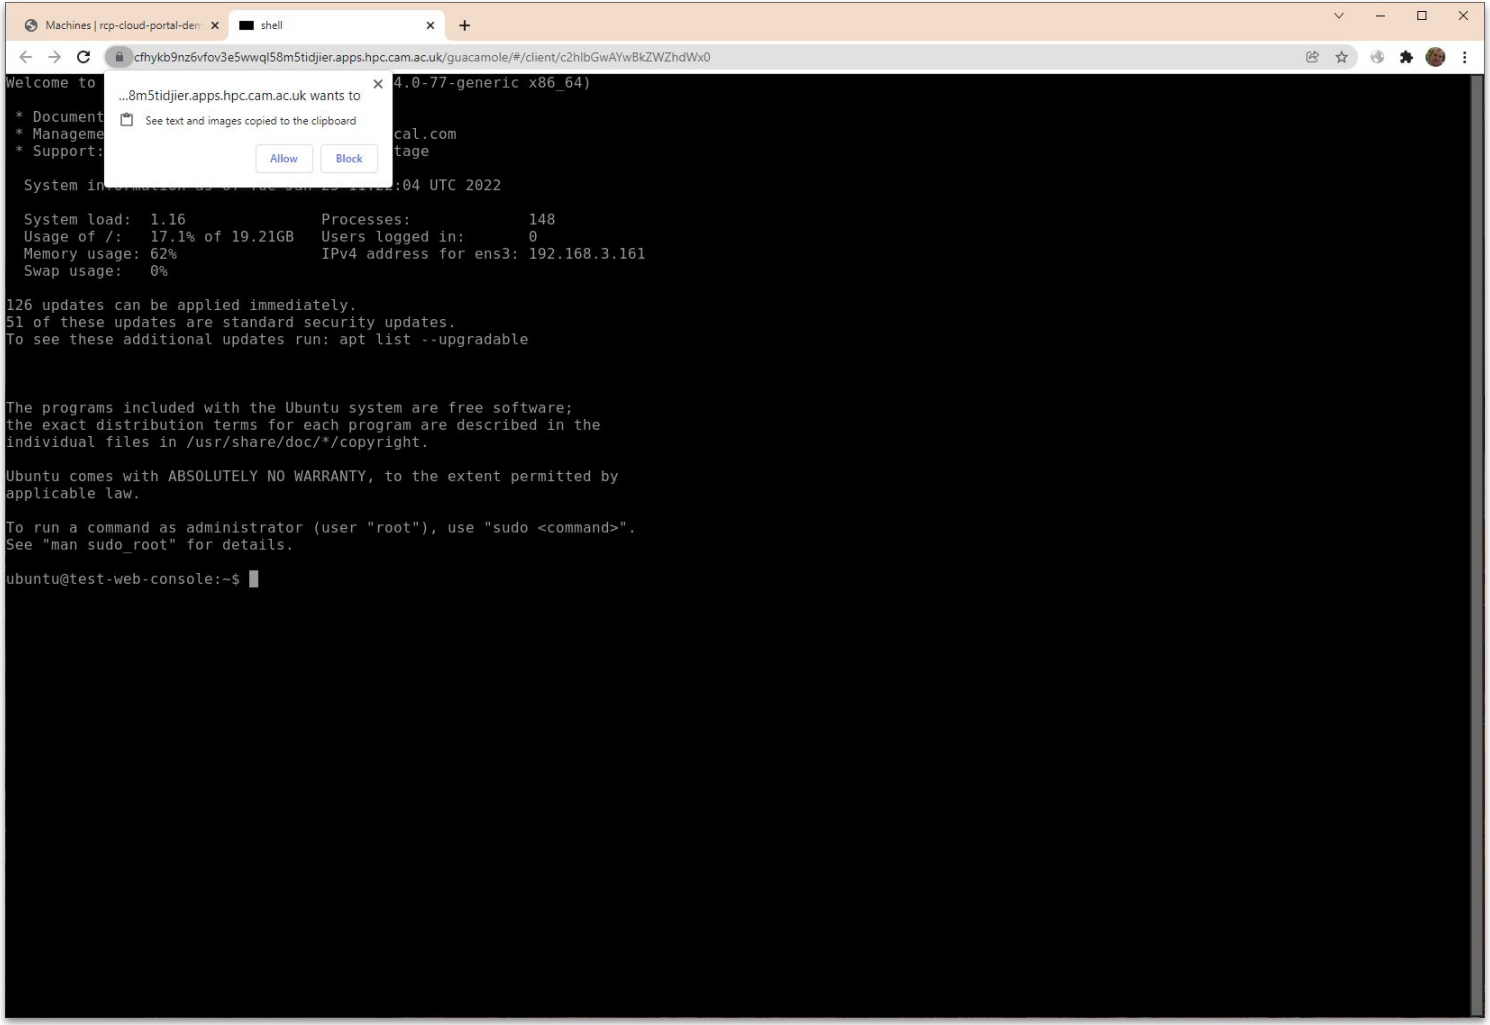
Machines | rcp-cloud-portal-dem
shell
cfhykb9nz6vfov3e5wwql58m5tidjier.apps.hpc.cam.ac.uk/guacamole/#/client/c2hlbGwAYwBkZWZhdWx0
...8m5tidjier.apps.hpc.cam.ac.uk wants to
See text and images copied to the clipboard
Allow
Block
Welcome to
4.0-77-generic x86_64)
* Document
* Manageme
cal.com
* Support:
tage
System in
:04 UTC 2022
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:~$

How did the VM access work?

StackHPC

- There is nothing hiding the OpenStack API here
  - Similar to Exposhere
- Login to OIDC integrated OpenStack Keystone
  - AARC Blueprint, IRIS IAM
- Create OpenStack server via API
  - Cloud-init configures guacamole and proxy
- Zenith Proxy
  - Authenticate users via OpenStack
  - Authorize based on OpenStack project membership
  - Don’t always burn a public IP, via ingress controller style proxy

Science Platforms:
Get me a Slurm Cluster!

StackHPC

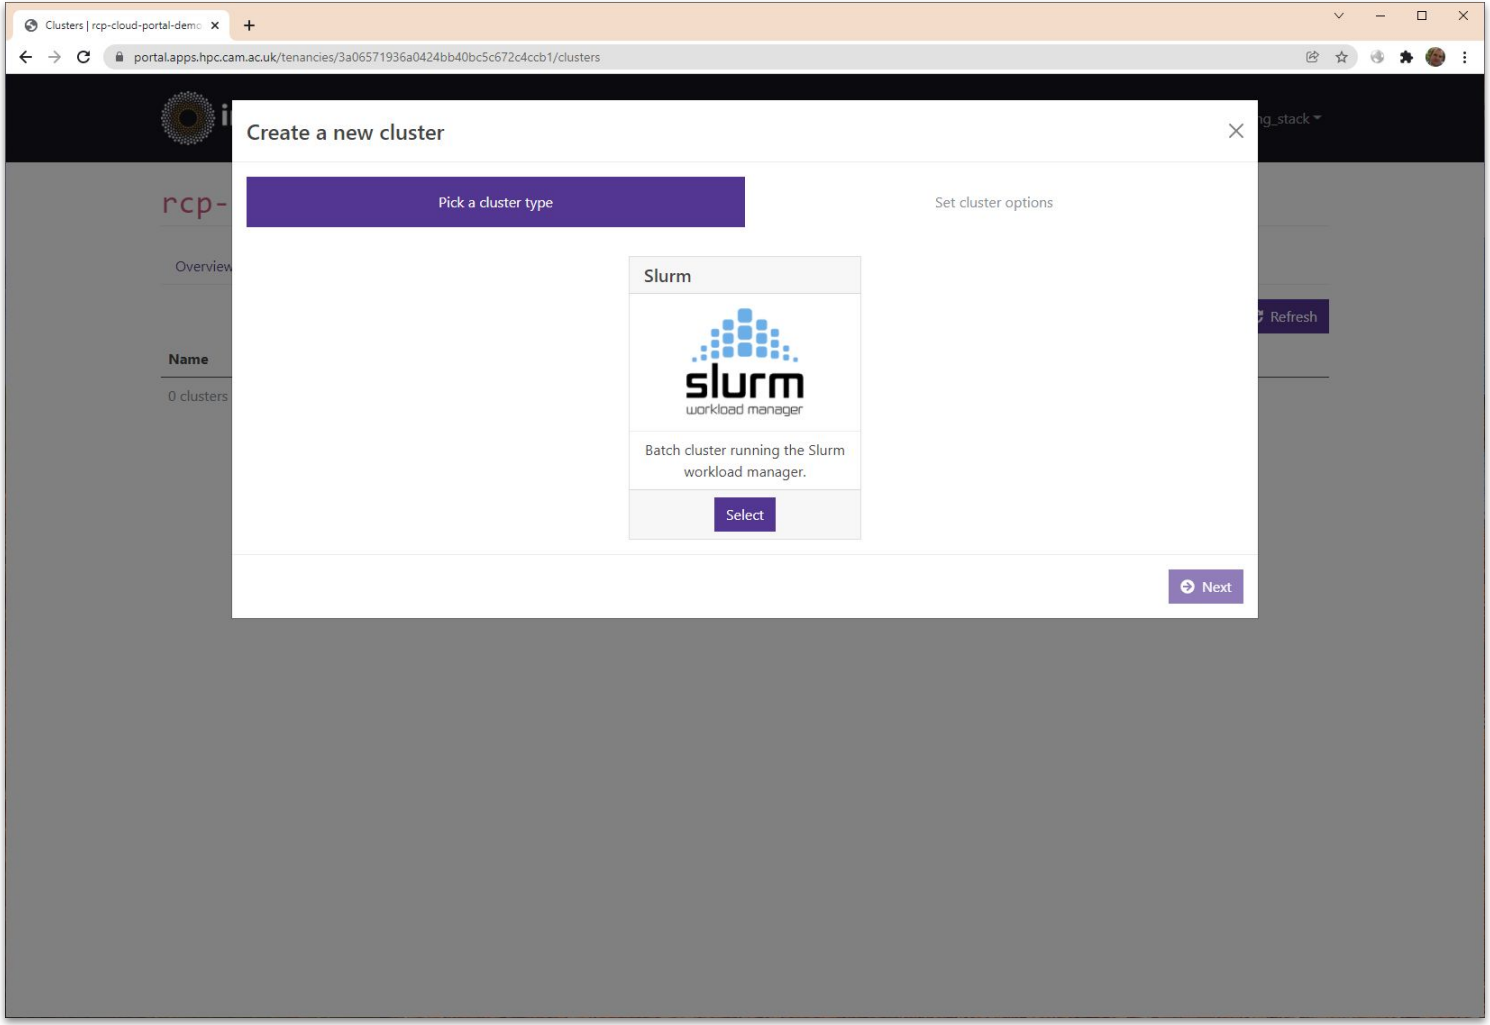
Clusters | rcp-cloud-portal-demo
portal.apps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/clusters
Create a new cluster
Pick a cluster type
Set cluster options
Slurm
slurm
workload manager
Batch cluster running the Slurm workload manager.
Select
Next
rcp-
Overview
Refresh
Name
0 clusters

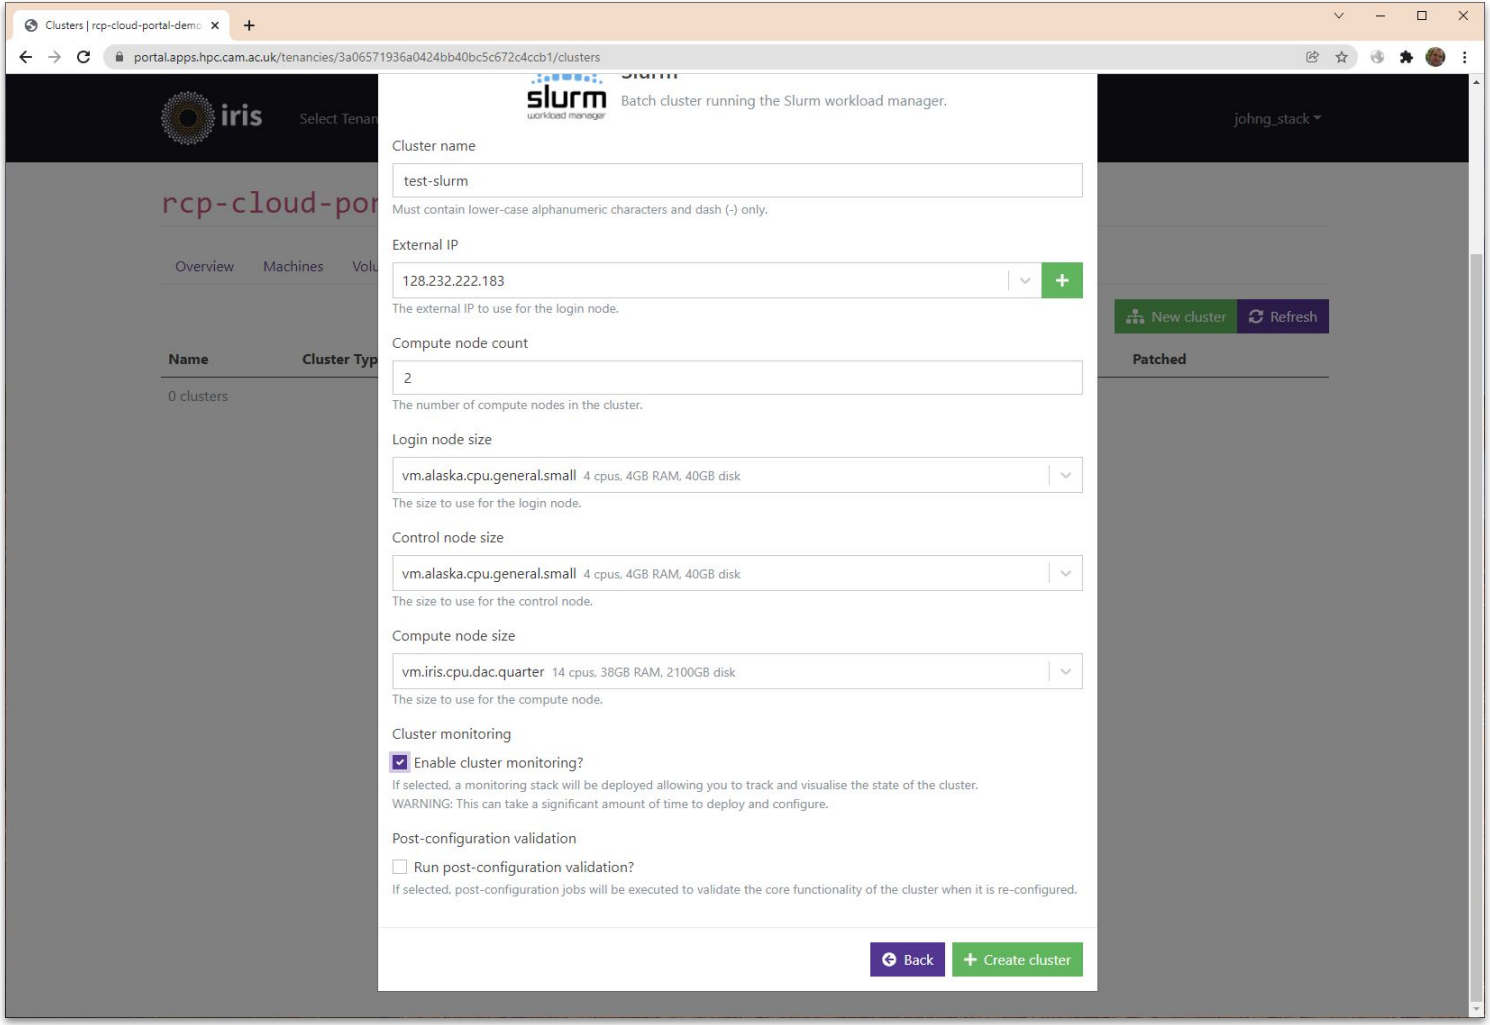
Clusters | rcp-cloud-portal-demo
portal.apps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/clusters
iris
Select Tenant
johng_stack
rcp-cloud-por
Overview
Machines
Volu
New cluster
Refresh
Name
Cluster Typ
Patched
0 clusters
slurm
workload manager
Batch cluster running the Slurm workload manager.
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

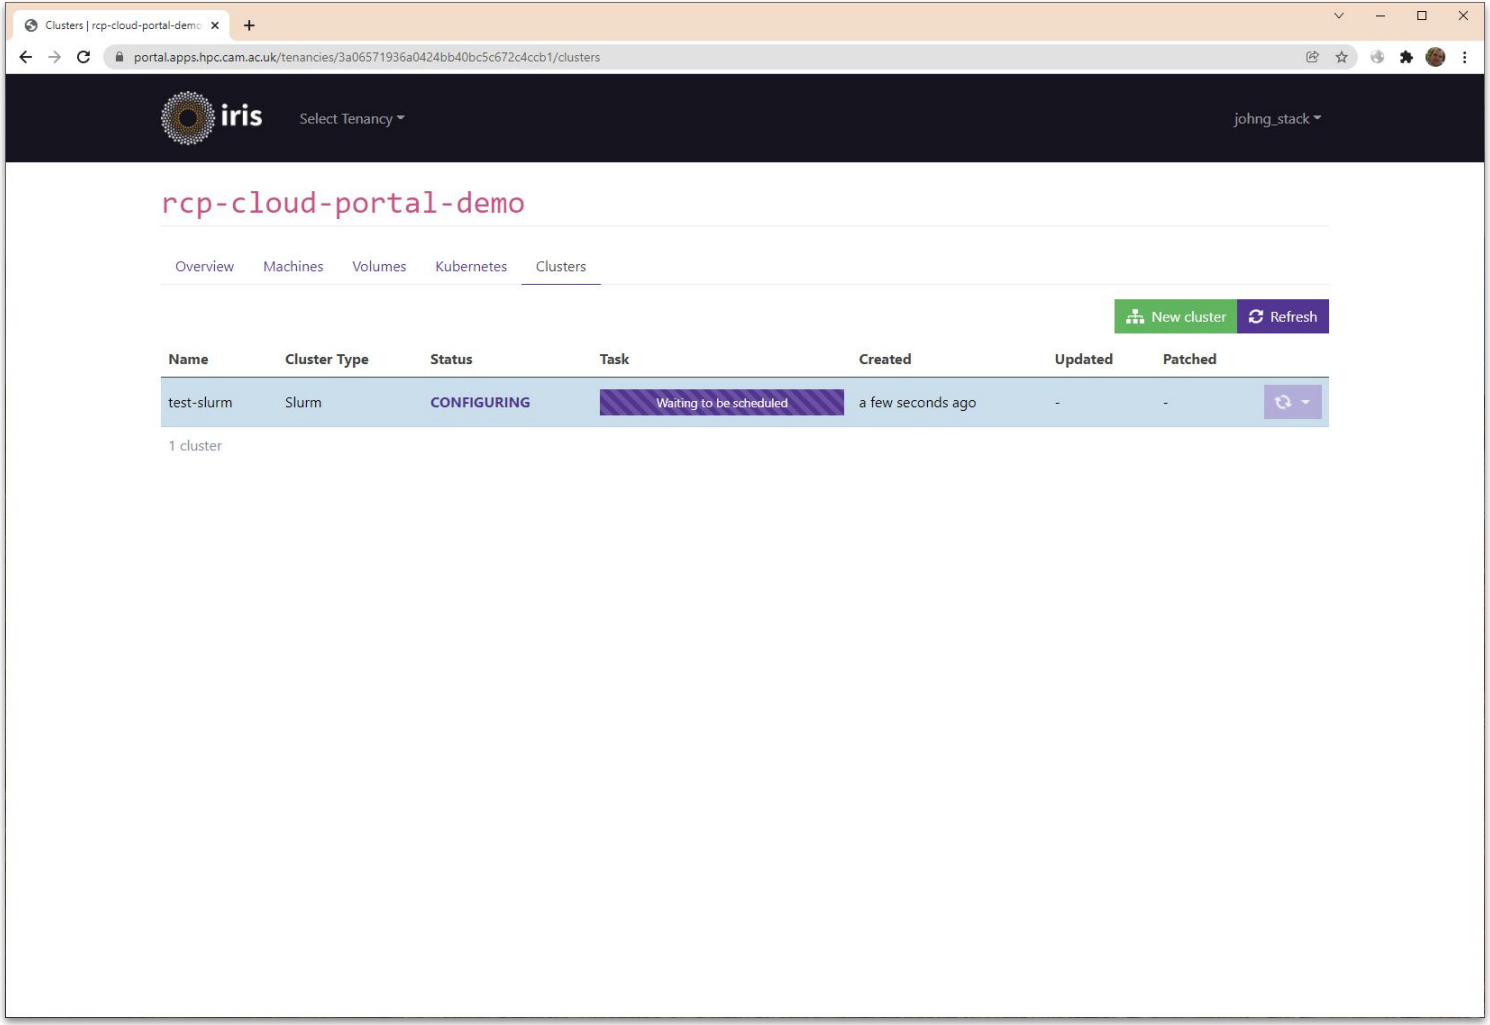
Clusters | rcp-cloud-portal-demo
portal.apps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/clusters
iris
Select Tenancy
johng_stack
rcp-cloud-portal-demo
Overview
Machines
Volumes
Kubernetes
Clusters
New cluster
Refresh
Name
Cluster Type
Status
Task
Created
Updated
Patched
test-slurm
Slurm
CONFIGURING
Waiting to be scheduled
a few seconds ago
-
-
1 cluster

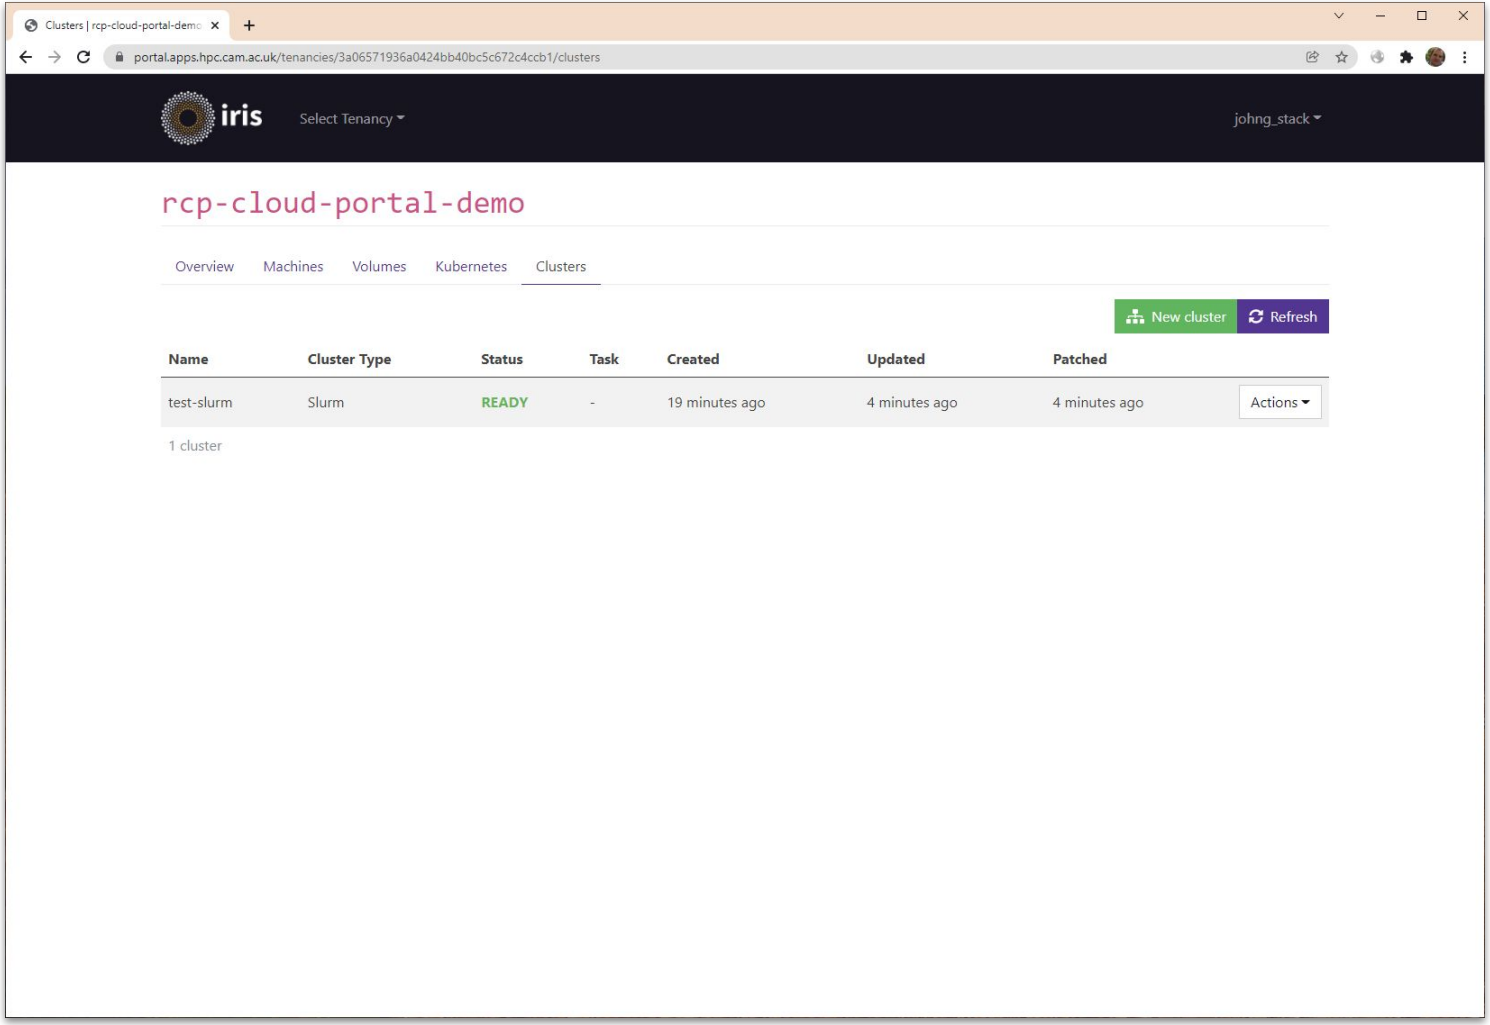
Clusters | rcp-cloud-portal-demo
portal.apps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/clusters
iris
Select Tenancy
johng_stack
rcp-cloud-portal-demo
Overview
Machines
Volumes
Kubernetes
Clusters
New cluster
Refresh
Name
Cluster Type
Status
Task
Created
Updated
Patched
test-slurm
Slurm
READY
-
19 minutes ago
4 minutes ago
4 minutes ago
Actions
1 cluster

How did you create the cluster?

StackHPC

- Generic, not specific to Slurm
- AWX runs Ansible playbooks
  - https://github.com/stackhpc/caas-slurm-appliance
- Terraform creates infrastructure
  - State stored in Consul
  - Inventory updated from Terraform output
- SSH in to configure hosts
- Optionally using image build via Packer pipeline
  - Referencing Pulp repo mirrors
- Auto-detect if project has RDMA provider network

Science Platforms:
Get me a Kubernetes Cluster

StackHPC

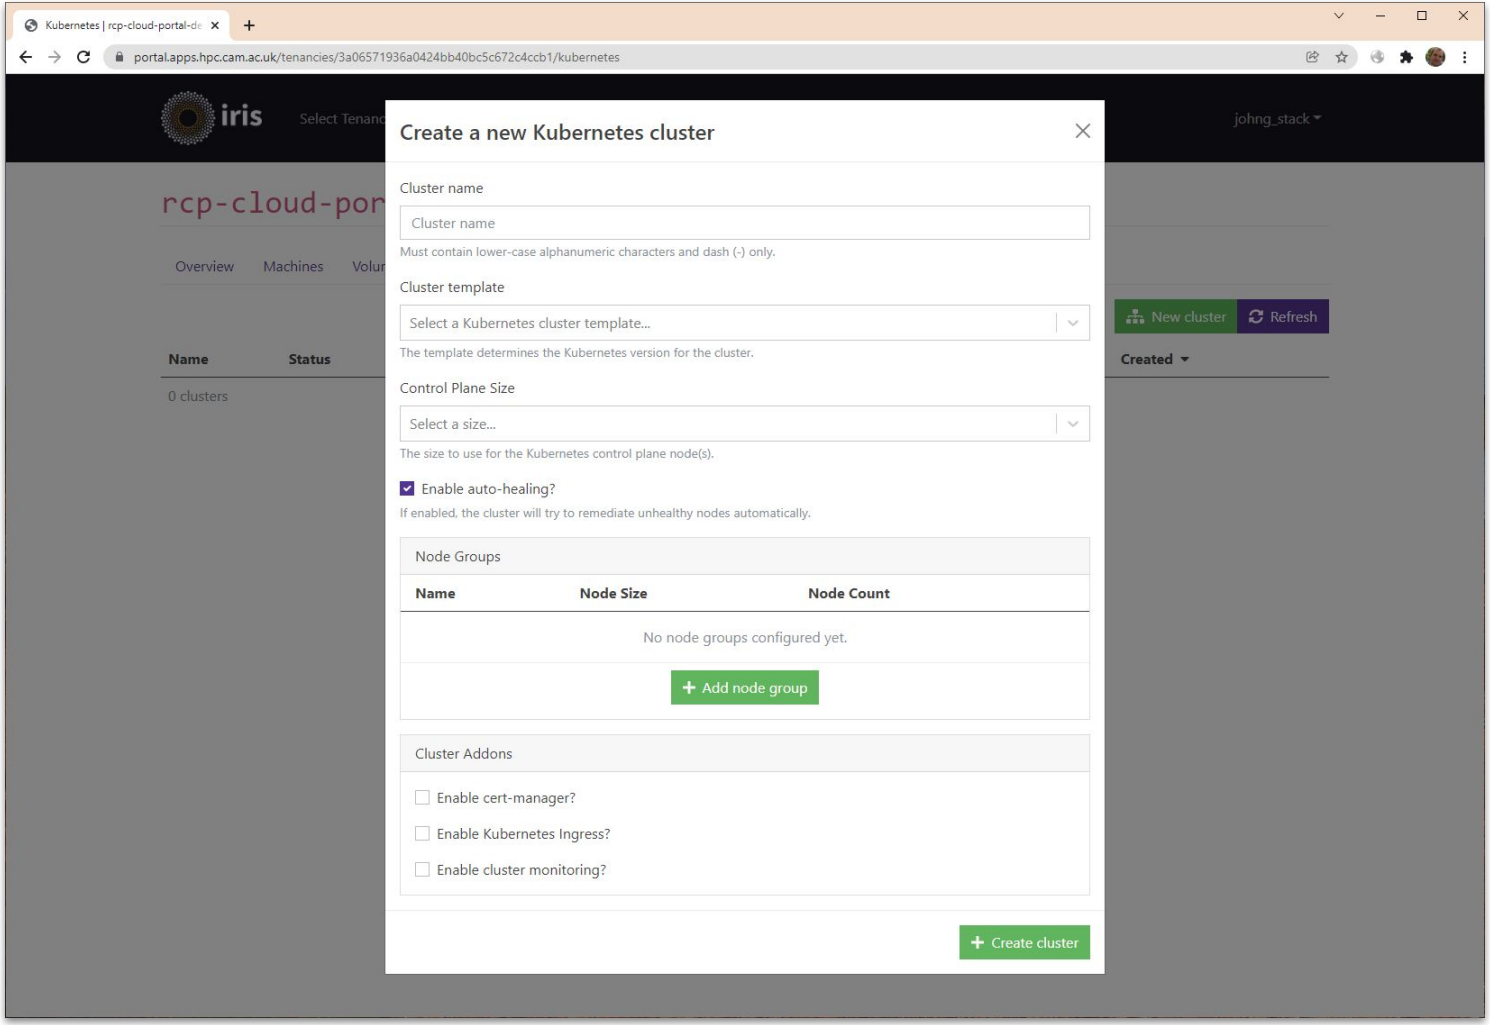
Kubernetes | rcp-cloud-portal-de
portal.apps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/kubernetes
iris
Select Tenanc
johng_stack
rcp-cloud-por
Overview
Machines
Volur
New cluster
Refresh
Name
Status
Created
0 clusters
Create a new Kubernetes cluster
Cluster name
Cluster name
Must contain lower-case alphanumeric characters and dash (-) only.
Cluster template
Select a Kubernetes cluster template...
The template determines the Kubernetes version for the cluster.
Control Plane Size
Select a 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 Ingress?
Enable cluster monitoring?
Create cluster

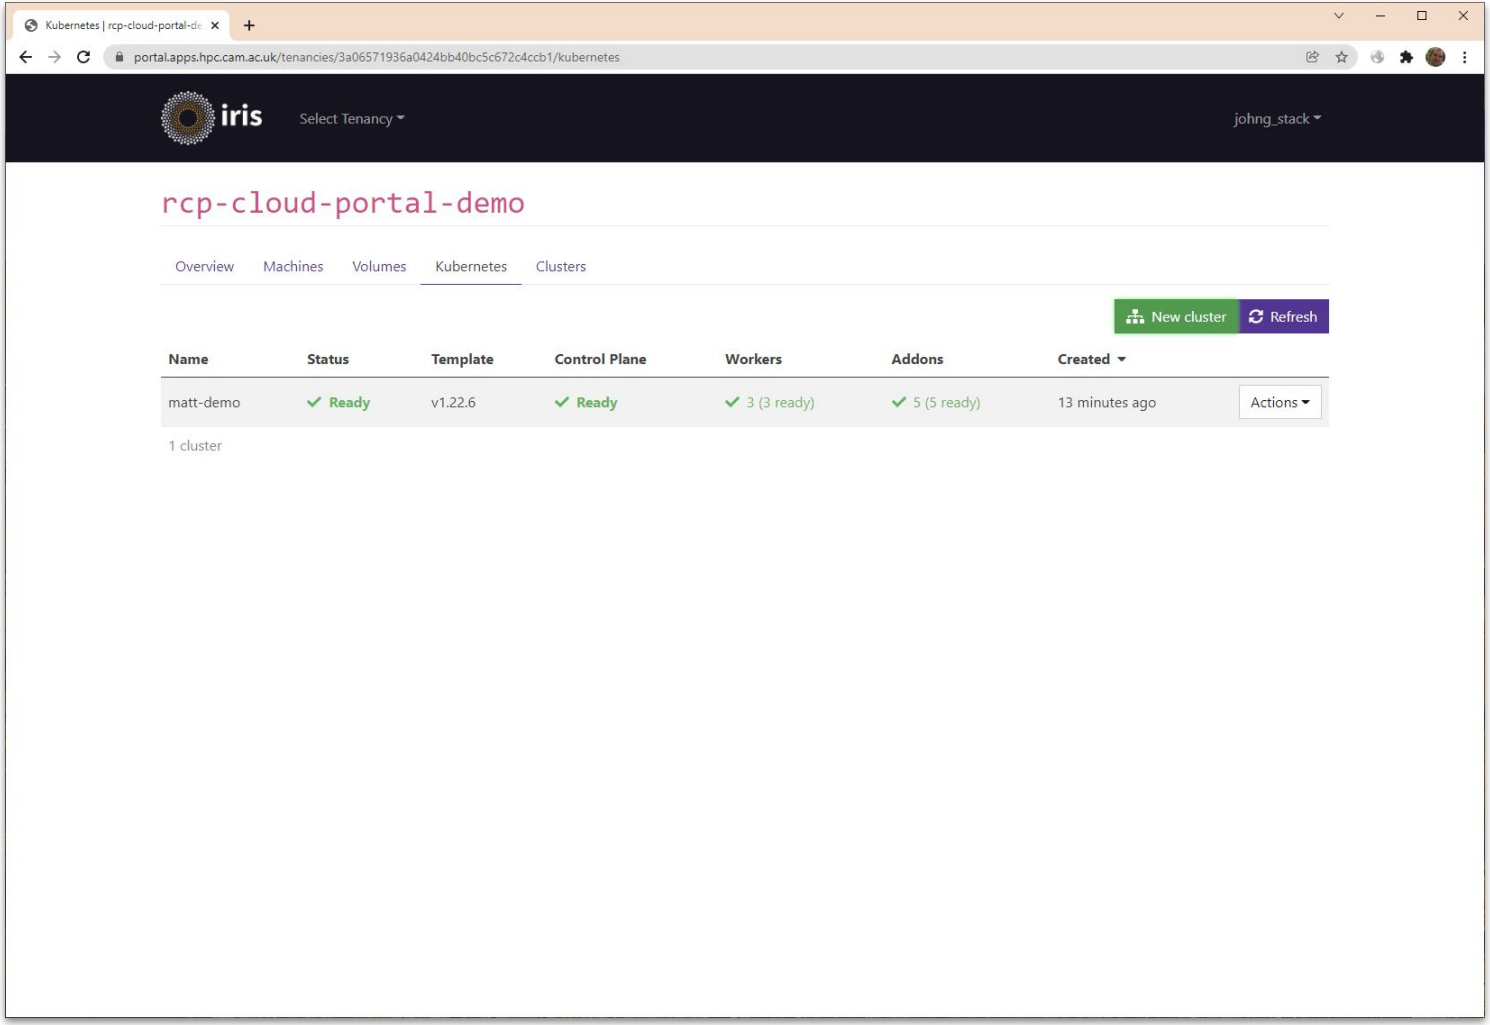
Kubernetes | rcp-cloud-portal-de
portal.apps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/kubernetes
iris
Select Tenancy
johng_stack
rcp-cloud-portal-demo
Overview
Machines
Volumes
Kubernetes
Clusters
New cluster
Refresh
Name
Status
Template
Control Plane
Workers
Addons
Created
matt-demo
Ready
v1.22.6
Ready
3 (3 ready)
5 (5 ready)
13 minutes ago
Actions
1 cluster

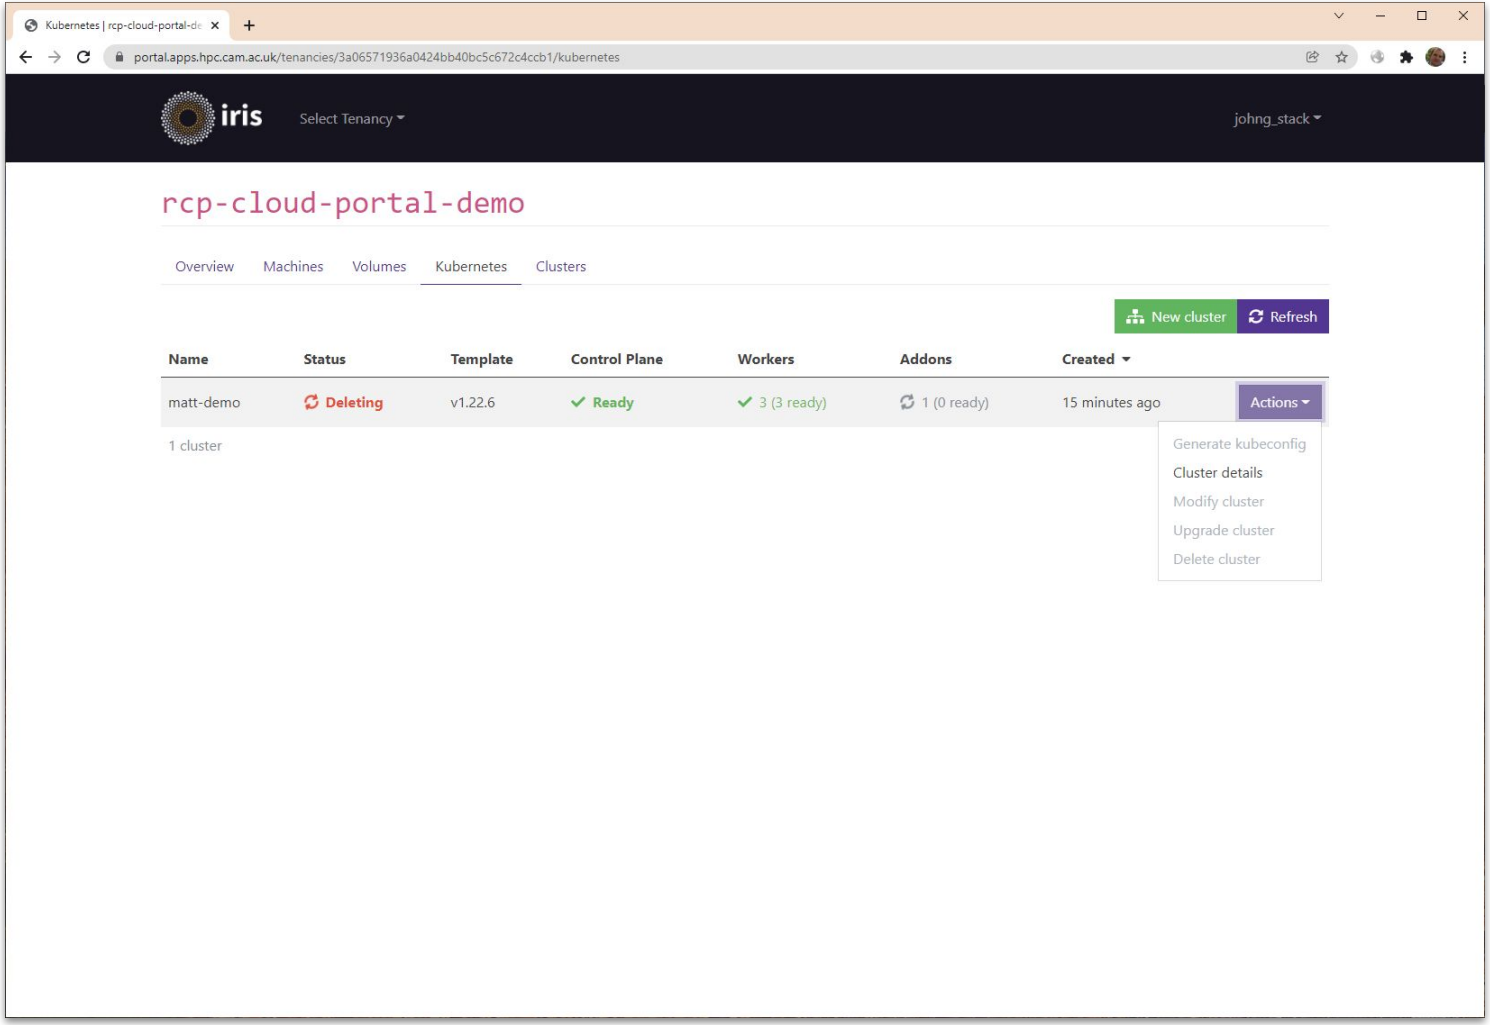
Kubernetes | rcp-cloud-portal-de
portal.apps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/kubernetes
iris
Select Tenancy
johng_stack
rcp-cloud-portal-demo
Overview
Machines
Volumes
Kubernetes
Clusters
New cluster
Refresh
Name
Status
Template
Control Plane
Workers
Addons
Created
matt-demo
Deleting
v1.22.6
Ready
3 (3 ready)
1 (0 ready)
15 minutes ago
Actions
Generate kubeconfig
Cluster details
Modify cluster
Upgrade cluster
Delete cluster
1 cluster

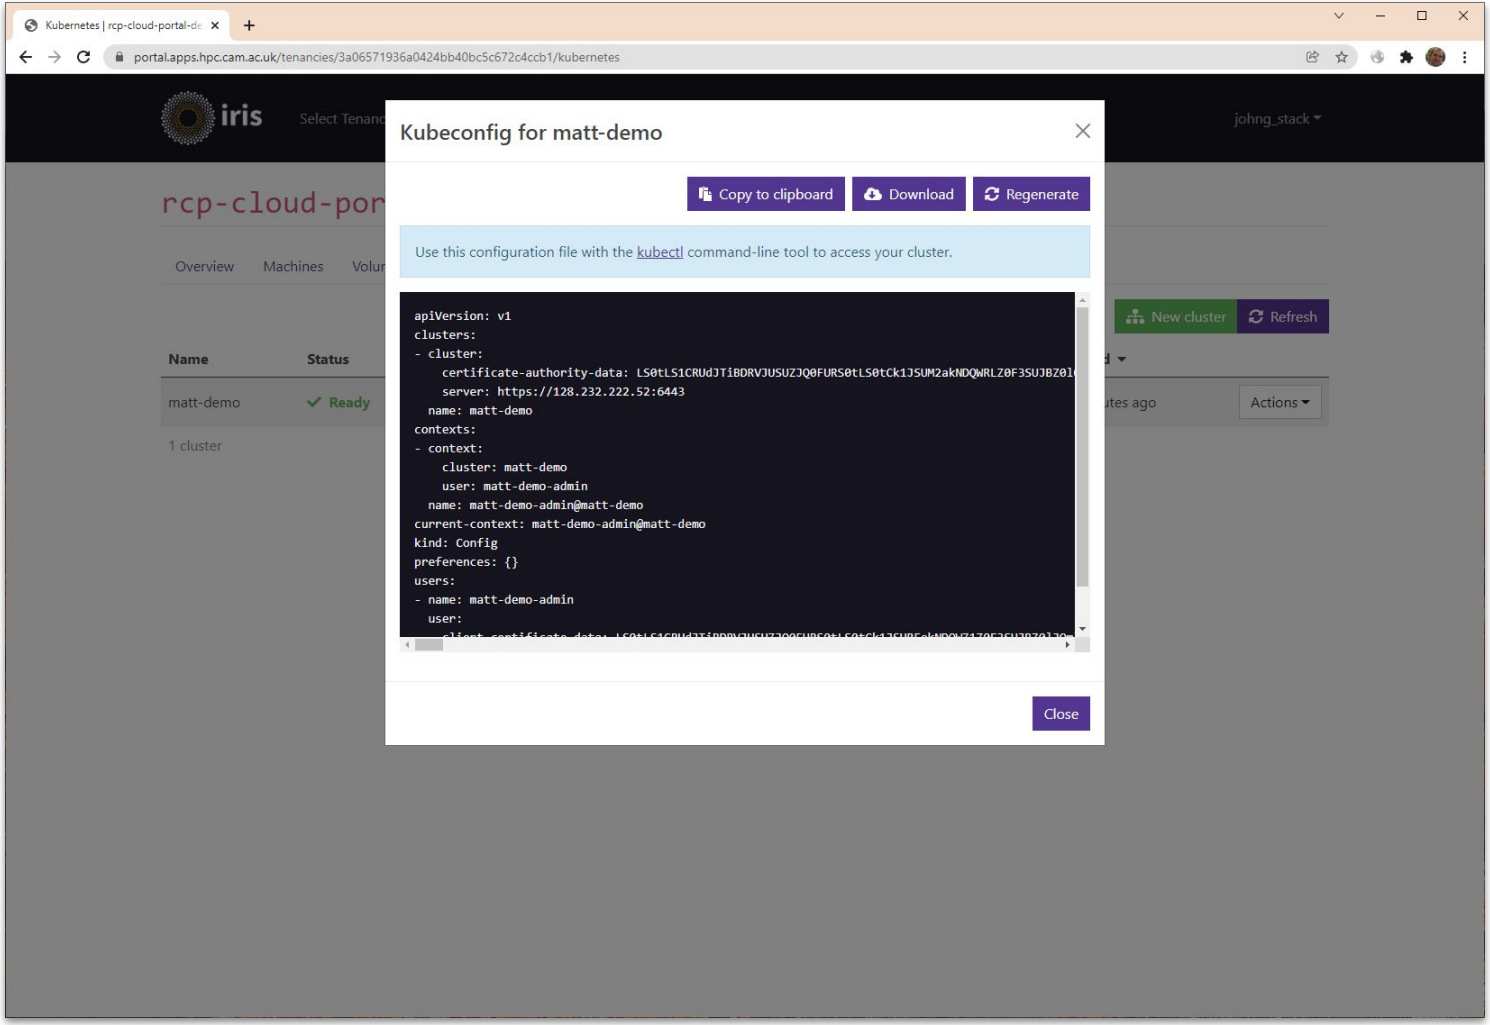
Kubernetes | rcp-cloud-portal-de
portal.apps.hpc.cam.ac.uk/tenancies/3a06571936a0424bb40bc5c672c4ccb1/kubernetes
iris
Select Tenanc
johng_stack
rcp-cloud-por
Overview
Machines
Volur
New cluster
Refresh
Name
Status
matt-demo
Ready
utes ago
Actions
1 cluster
Kubeconfig for matt-demo
Copy to clipboard
Download
Regenerate
Use this configuration file with the kubectl command-line tool to access your cluster.
apiVersion: v1
clusters:
- cluster:
certificate-authority-data: LS0tLS1CRUdJTiBDRVJUSUZJQ0FURS0tLS0tCk1JSUM2akNDQWRLZ0F3SUJBZ0l
server: https://128.232.222.52:6443
name: matt-demo
contexts:
- context:
cluster: matt-demo
user: matt-demo-admin
name: matt-demo-admin@matt-demo
current-context: matt-demo-admin@matt-demo
kind: Config
preferences: {}
users:
- name: matt-demo-admin
user:
Close

How did you create K8s Clusters?

StackHPC

- Kubernetes Cluster API
  - OpenStack Cluster API Provider
  - Cinder CSI and Cloud Provider OpenStack
  - General good stuff: Rolling upgrades, Self-healing, Auto-scaling
  - Add additional RDMA network when available
- Opinions wrapped with Helm charts
  - Add and remove “add-ons” (e.g. OFED and CUDA)
  - Wrapped up options, driven by an operator
  - https://github.com/stackhpc/capi-helm-charts
  - https://github.com/stackhpc/azimuth-capi-operator
- Previously used OpenStack Magnum
  - … looking at adding a Cluster API driver into Magnum

Science Platforms:
Get me a JupyterHub

StackHPC

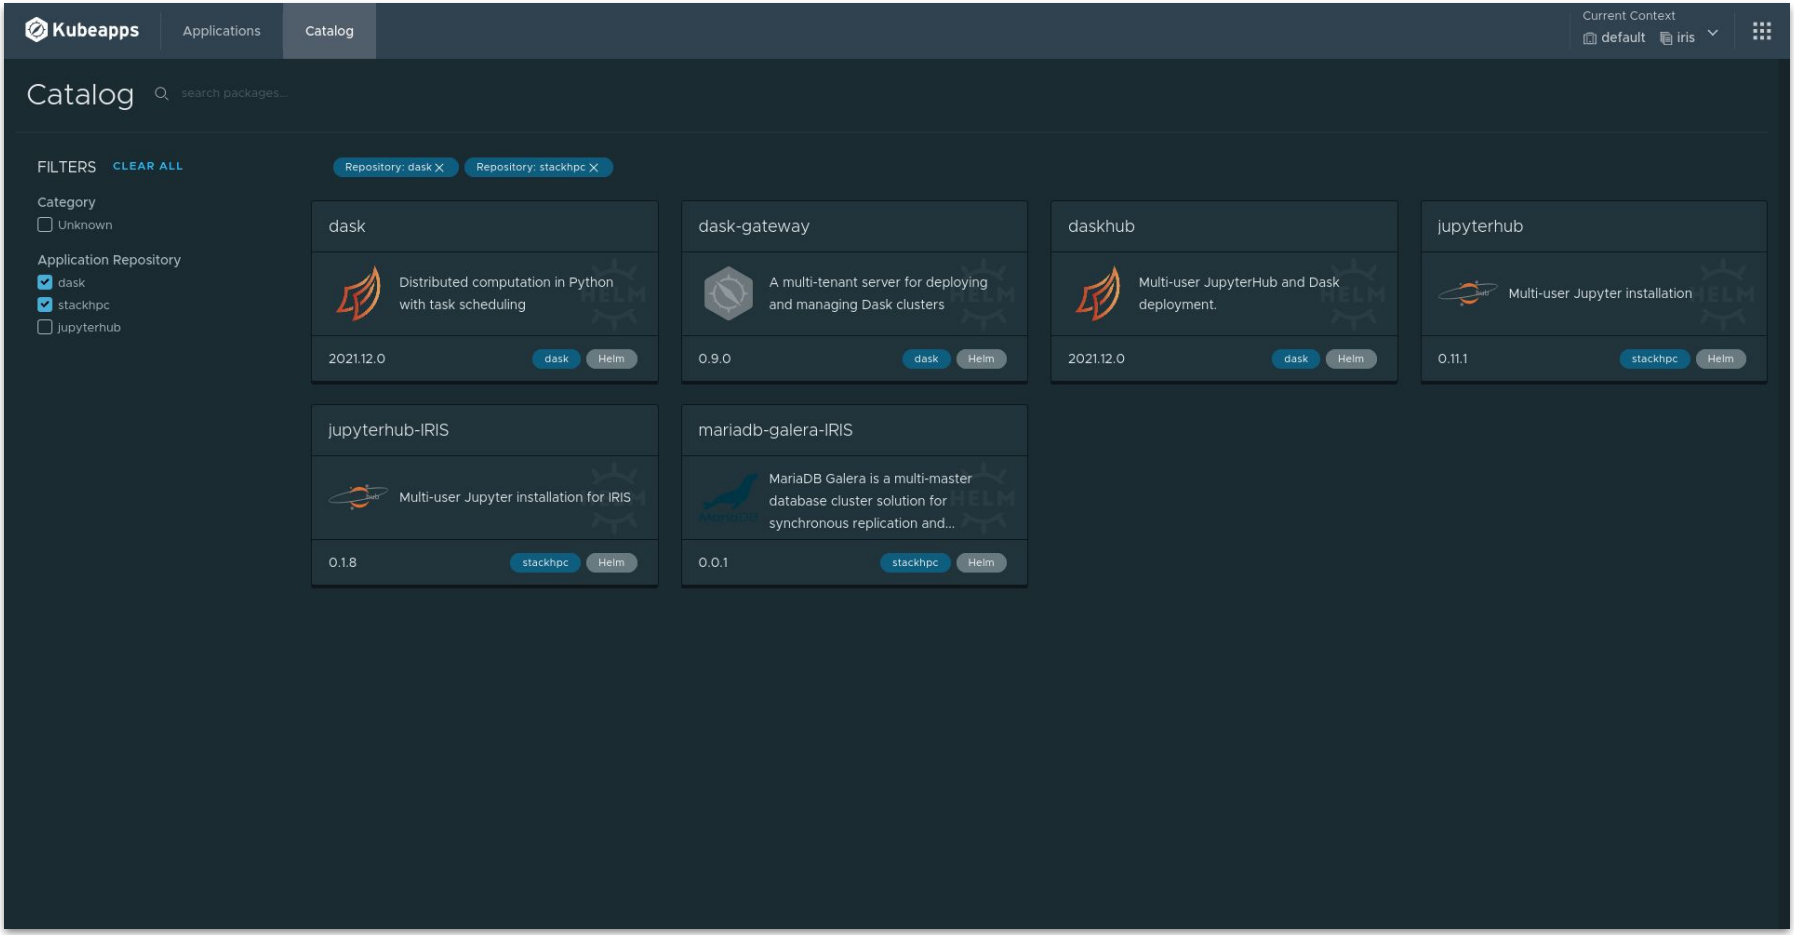
Kubeapps
Applications
Catalog
Current Context
default
iris
Catalog
search packages...
FILTERS
CLEAR ALL
Category
Unknown
Application Repository
dask
stackhpc
jupyterhub
Repository: dask
Repository: stackhpc
dask
Distributed computation in Python with task scheduling
2021.12.0
dask
Helm
dask-gateway
A multi-tenant server for deploying and managing Dask clusters
0.9.0
dask
Helm
daskhub
Multi-user JupyterHub and Dask deployment.
2021.12.0
dask
Helm
jupyterhub
Multi-user Jupyter installation
0.11.1
stackhpc
Helm
jupyterhub-IRIS
Multi-user Jupyter installation for IRIS
0.1.8
stackhpc
Helm
mariadb-galera-IRIS
MariaDB Galera is a multi-master database cluster solution for synchronous replication and...
0.0.1
stackhpc
Helm

What are you planning?

StackHPC

- Pre-install KubeApps
  - With a curated Helm repository
- Expose from the Portal
  - Kubernetes Dashboard
  - Grafana
  - KubeApps
  - … and JupyterHub created by KubeApps
- Zenith Proxy for Kubernetes Service
  - Why: Share auth with the Portal
  - How: Zenith Client as a sidecar

Science Platforms:
Can I get some resources?

StackHPC

OpenStack Quota

StackHPC

- Resource Allocation converted into OpenStack Quota
  - Defines the high watermark of your usage
  - … very hard to “catch up” if you fall behind in your usage
- Dedicated Quota
  - Underutilization a problem
  - Sometimes implemented using Tenant isolation filters
- Overcommit Quota
  - Might not be able to use your Quota

… but what if you want need half your allocation in June, the rest in September?


The Coral Reef Cloud

OpenStack Blazar Reservations

StackHPC

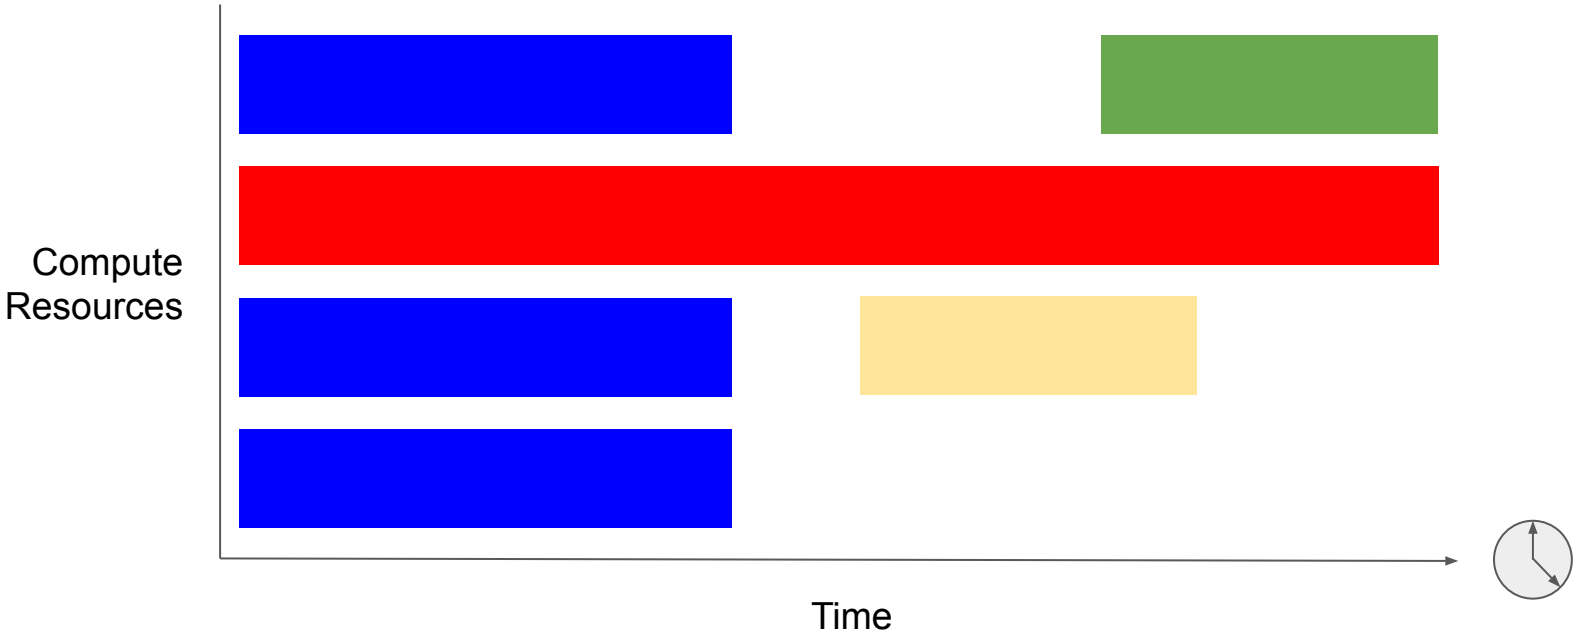

OpenStack Blazar Preemptibles

StackHPC

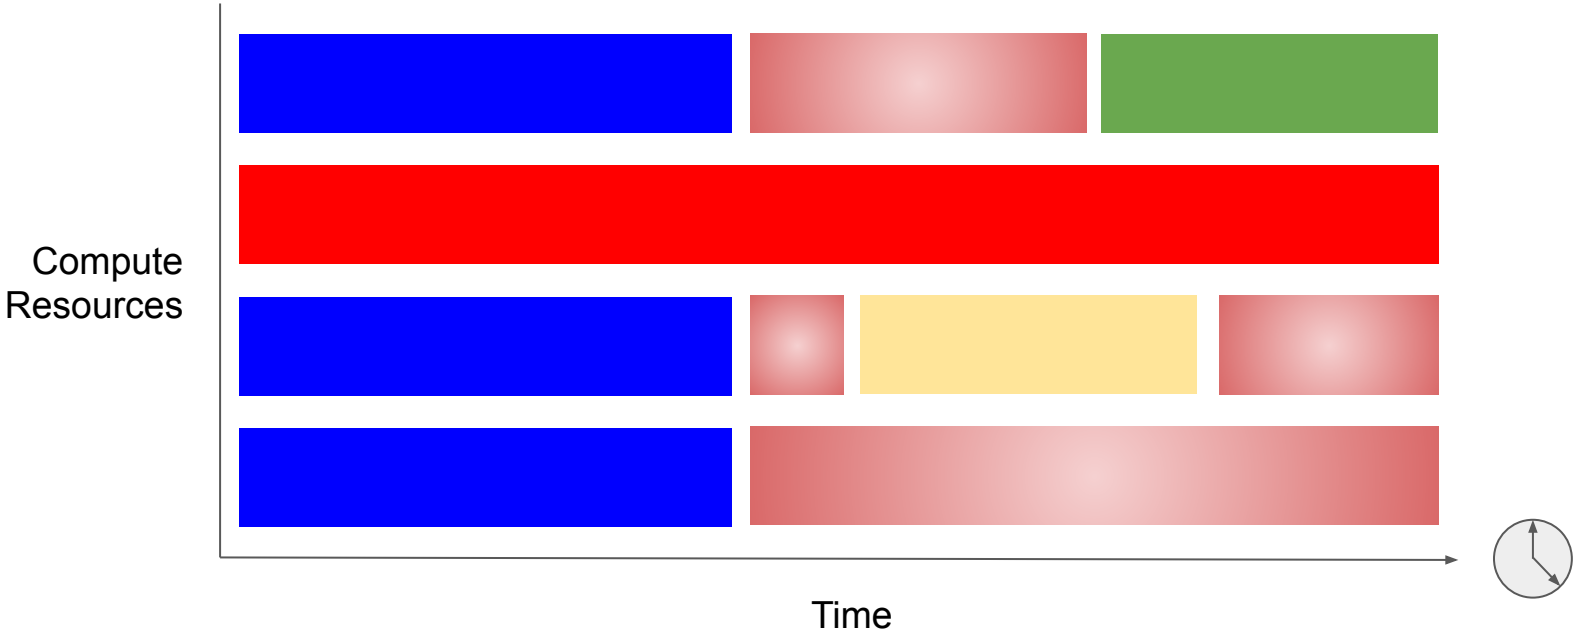

Tracking Efficiency / Utilization

StackHPC

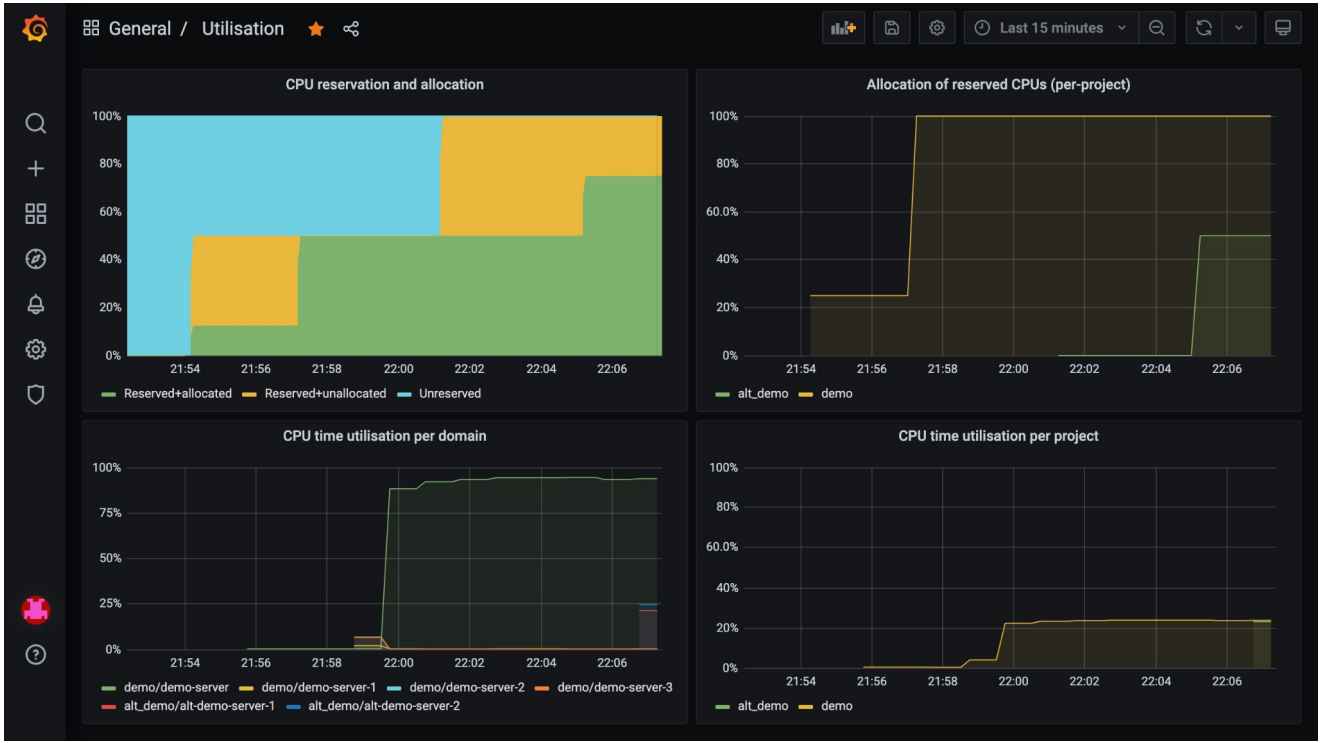

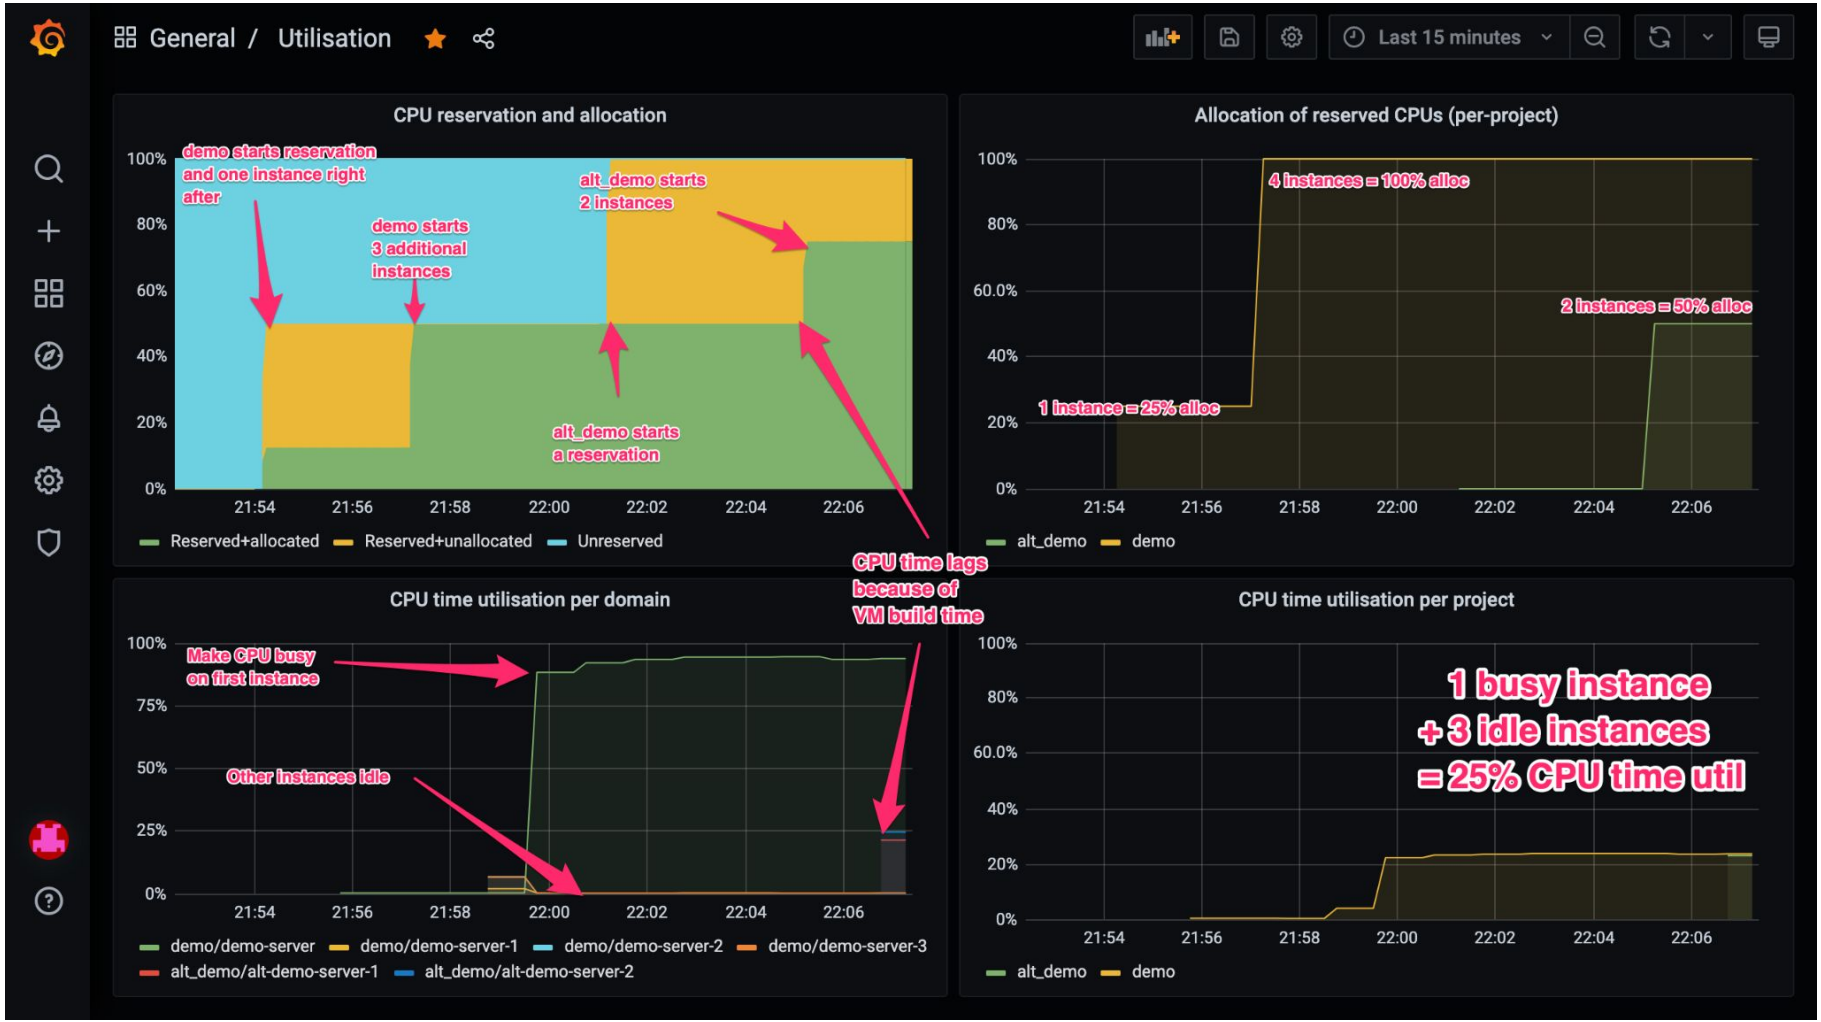
General / Utilisation
Last 15 minutes
CPU reservation and allocation
demo starts reservation and one instance right after
demo starts 3 additional instances
alt_demo starts 2 instances
alt_demo starts a reservation
100%
80%
60%
40%
20%
0%
21:54
21:56
21:58
22:00
22:02
22:04
22:06
Reserved+allocated
Reserved+unallocated
Unreserved
Allocation of reserved CPUs (per-project)
4 instances = 100% alloc
2 instances = 50% alloc
1 instance = 25% alloc
60.0%
alt_demo
demo
CPU time lags because of VM build time
CPU time utilisation per domain
Make CPU busy on first instance
Other instances idle
75%
50%
25%
demo/demo-server
demo/demo-server-1
demo/demo-server-2
demo/demo-server-3
alt_demo/alt-demo-server-1
alt_demo/alt-demo-server-2
CPU time utilisation per project
1 busy instance
+ 3 idle instances
= 25% CPU time util

What can I do?

StackHPC

OpenStack based HPC

StackHPC

- **OpenStack Ironic** to manage baremetal servers
- Slurm non-impacting updates by **reboot** triggering a **rebuild**
- **Kubernetes Cluster API** to create and upgrade K8s on OpenStack
- **Azimuth** Cloud Portal for on-demand Science Platforms
- **OpenStack Blazar** for resource reservations and preemption

Would you like to get involved?

StackHPC

- Scientific OpenStack SIG
  - https://www.openstack.org/use-cases/science/
- OpenStack Ironic
  - https://www.openstack.org/use-cases/bare-metal/
- OpenStack Kolla, Kolla-Ansible and Kayobe
  - https://docs.openstack.org/kayobe/latest/getting-started.html
- OpenHPC Slurm on OpenStack
  - https://github.com/stackhpc/ansible-slurm-appliance
- Azimuth Cloud Portal
  - https://github.com/stackhpc/azimuth

Questions?

StackHPC