Scientific OpenStack FY2021 Progress

Resource Management, CI/CD, Cloud Portal

August 2021

StackHPC

StackHPC Company Overview

StackHPC

- Formed 2016, based in Bristol, UK
  - Based in Bristol with presence in Cambridge, France and Poland
  - Currently 17 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

StackHPC

Open Infrastructure FOUNDATION

SILVER MEMBER

Overview

StackHPC

- What is Scientific OpenStack?
- Resource Management
- CI/CD Improvements
- Cloud Portal

What is Scientific OpenStack?

StackHPC

HPC Stack 1.0

StackHPC

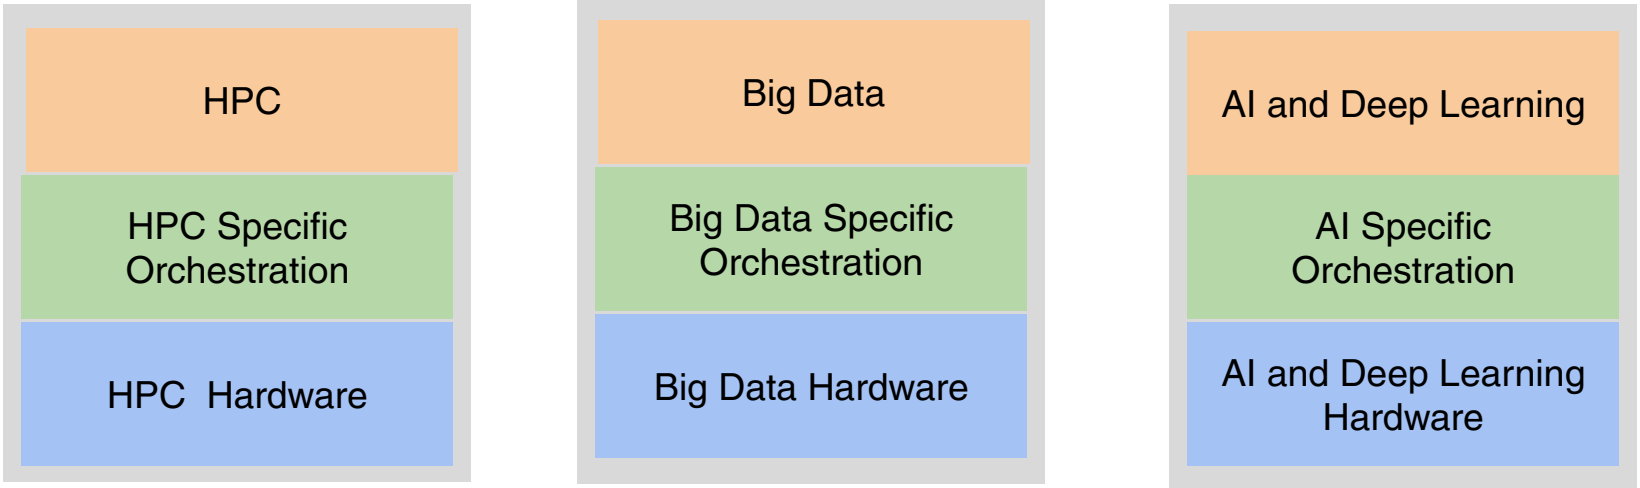

Why DevOps & ResOps in HPC?

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

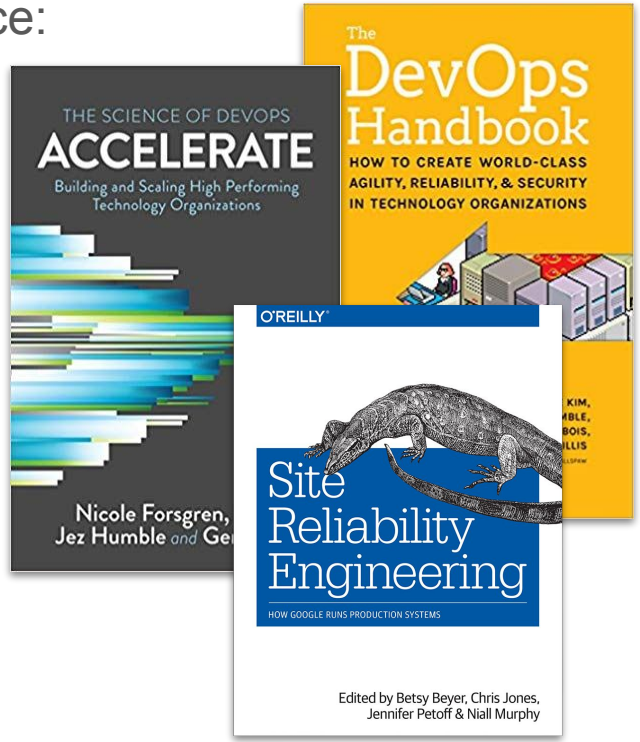

DevOps: “The Three Ways”

StackHPC

The First Way:
Systems Thinking

The Second Way:
Amplify Feedback Loops

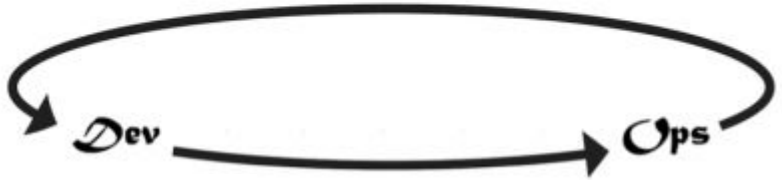

The Third Way:
Culture Of Continual Experimentation And Learning

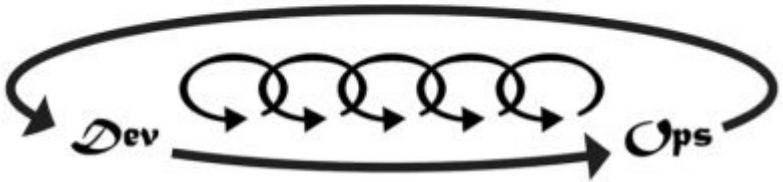

HPC Stack 1.0

StackHPC

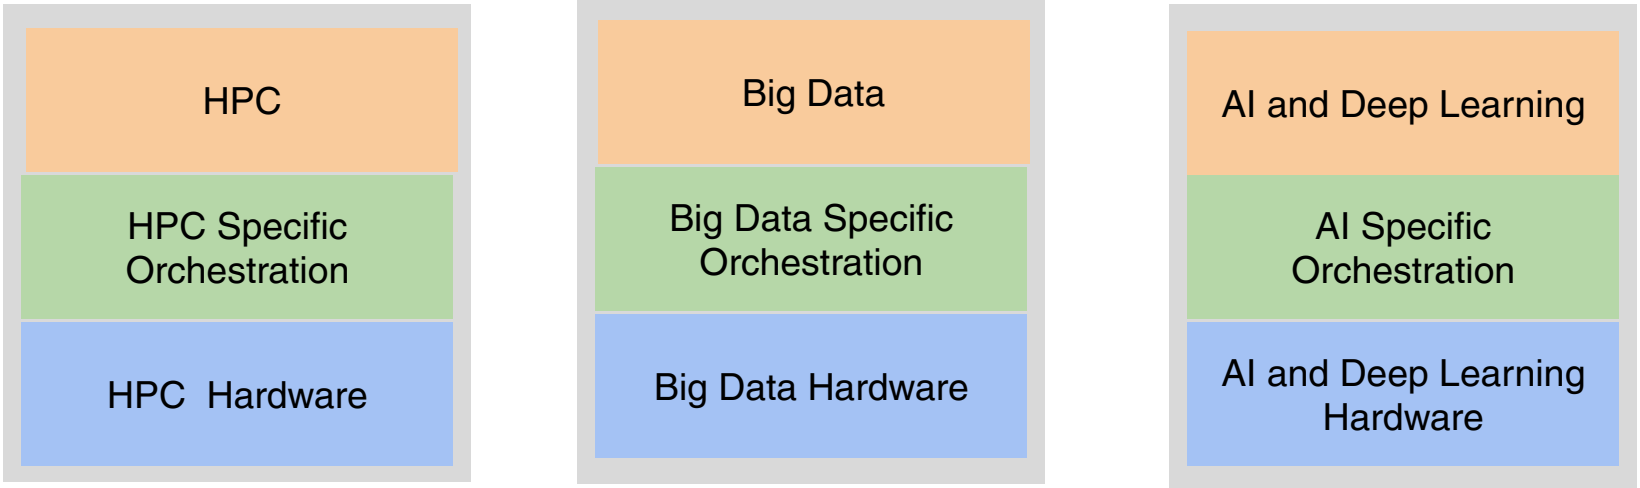

HPC Stack 2.0

StackHPC

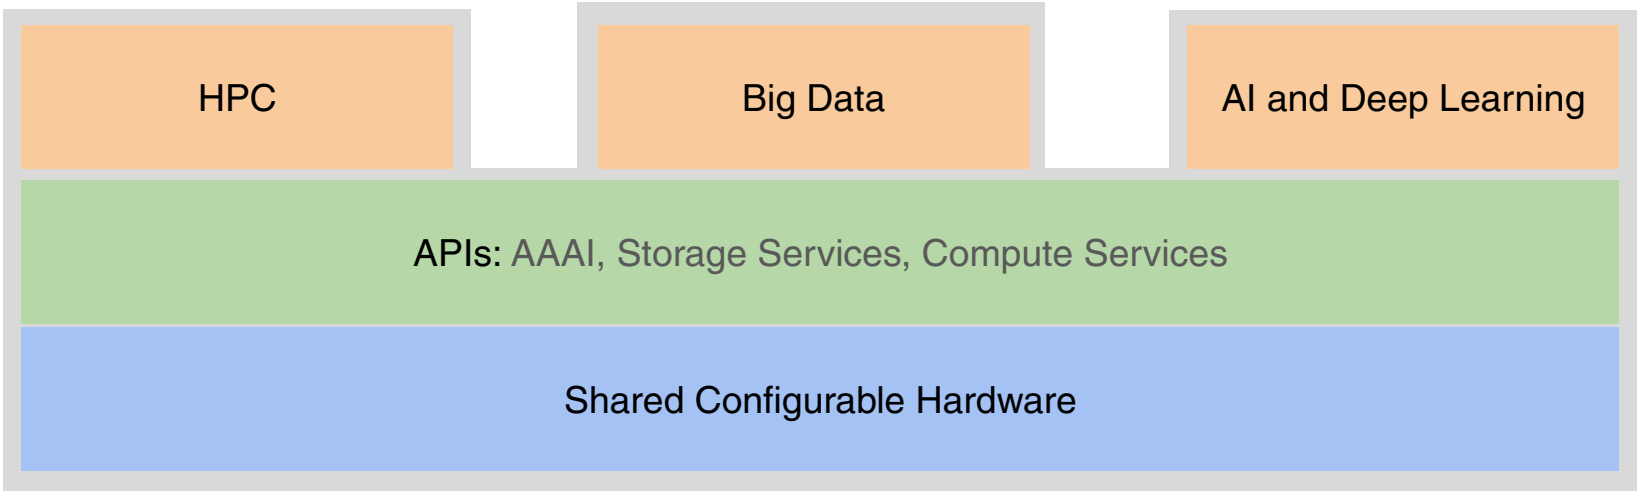

HPC Stack 2.0

StackHPC

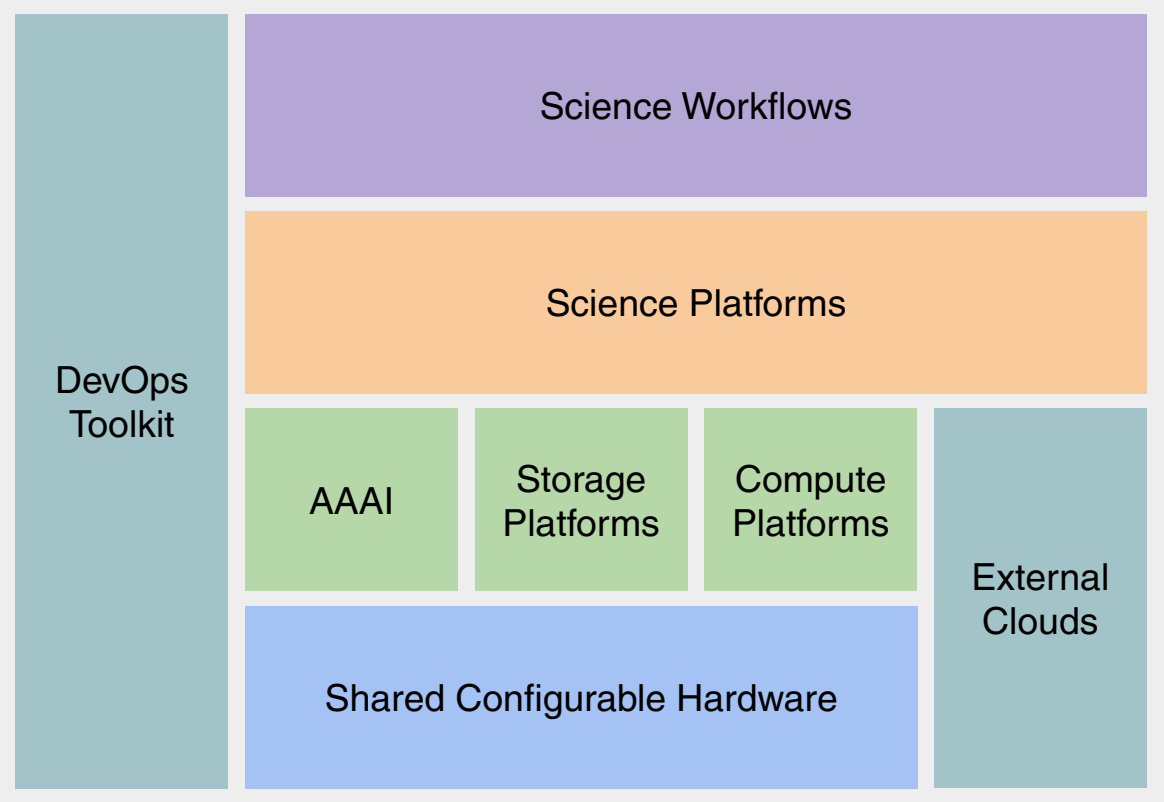

Scientific OpenStack

StackHPC

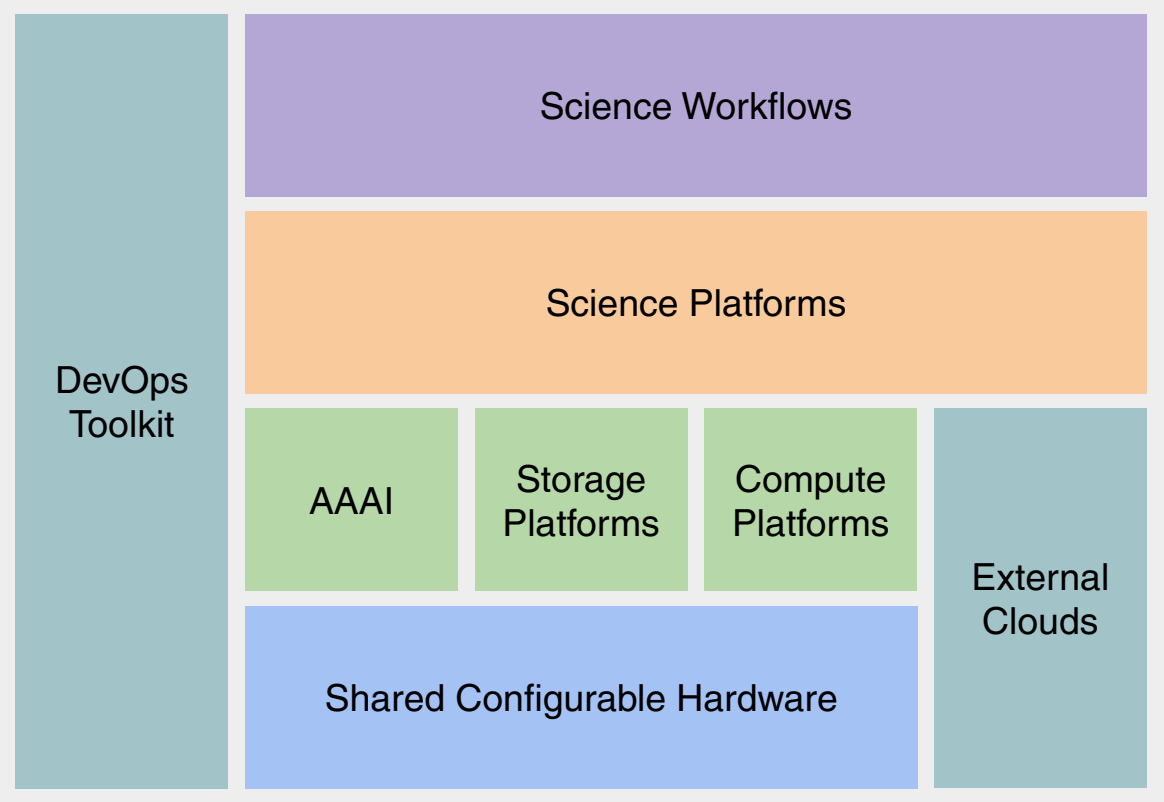

HPC Stack 2.0

StackHPC

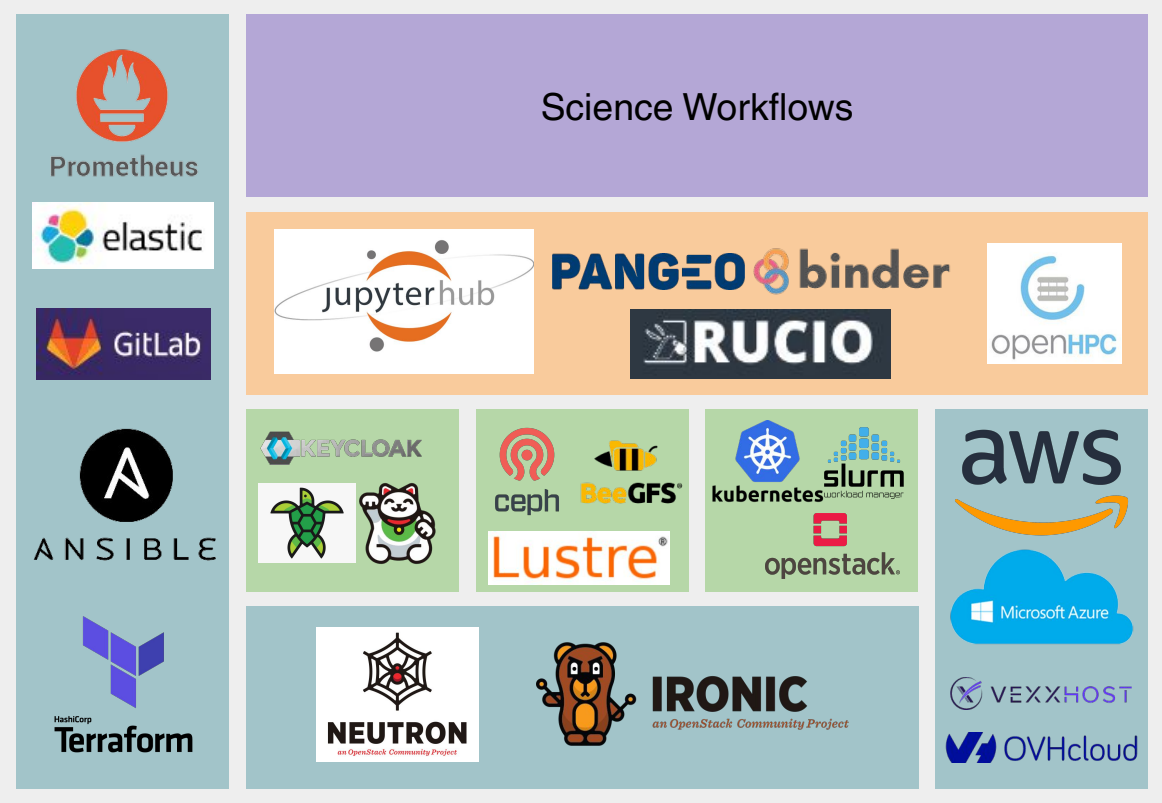

Resource Management

StackHPC


The Coral Reef Cloud

Reservations and Preemptibles

StackHPC

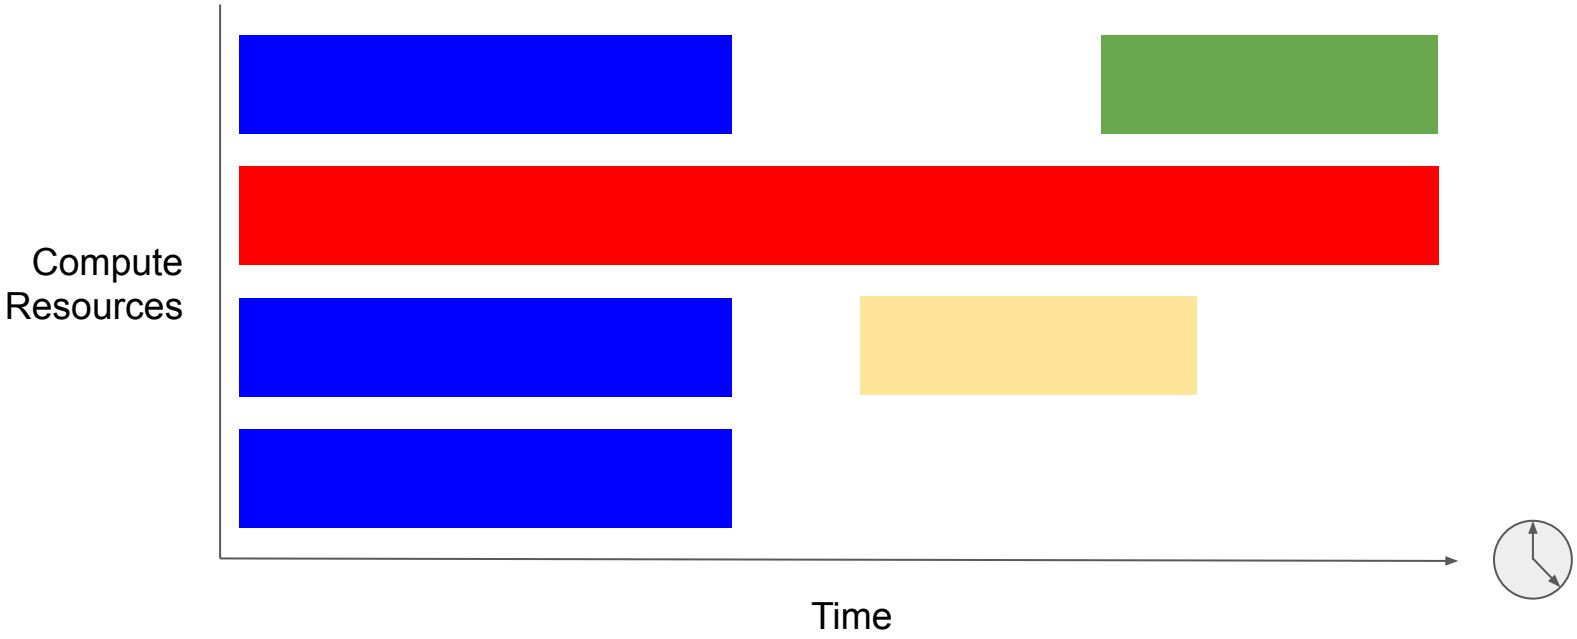

Reservations and Preemptibles

StackHPC

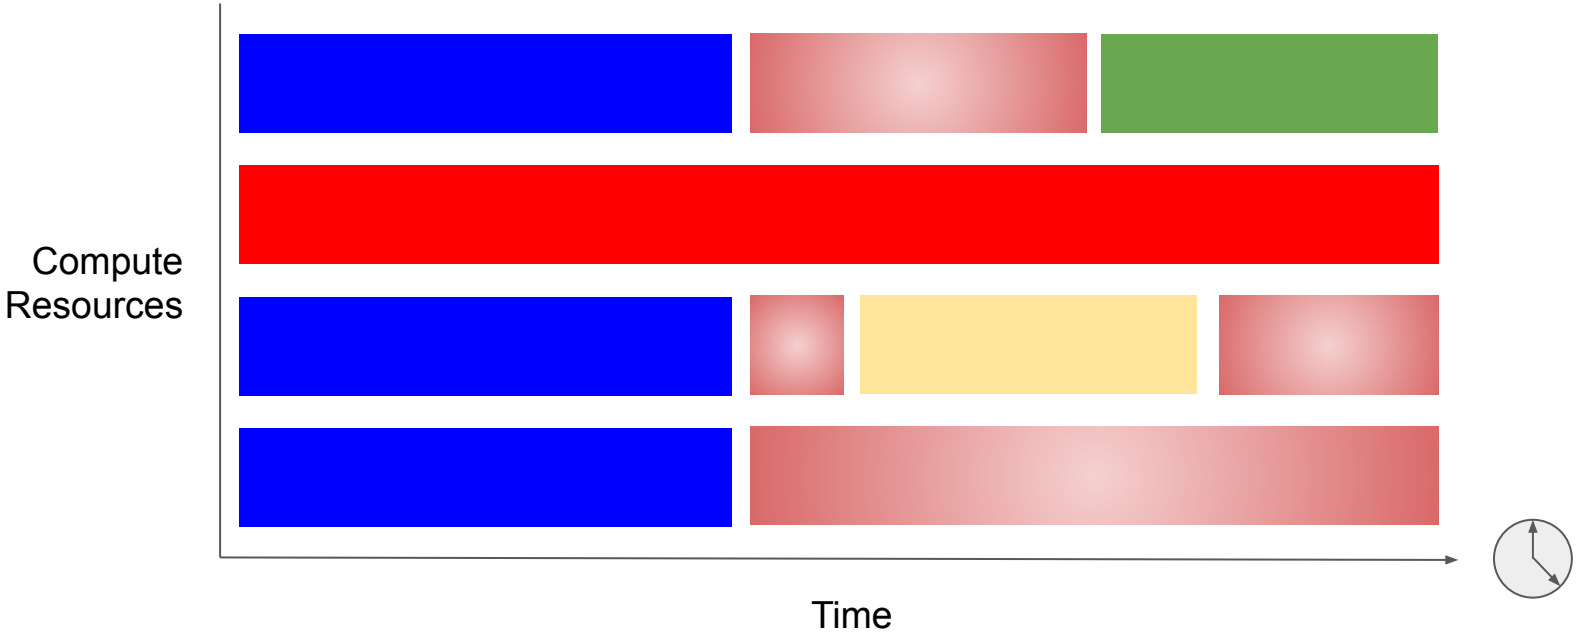

Why change?

StackHPC

- Quota system has some big limitations within fixed capacity clouds
  - People want to reserve a chunk of resource next month
- Need to better expose cloud capacity vs current utilization
- Need to expose project allocation vs utilization
- For each VM, expose efficiency of its resource usage
  - Did you create a VM that was too big?
  - Did you forget to kill your VM
  - Did you forget to stop your reservation when you finished early
- Pre-emptables to use the gaps between agreed reservations
  - Use of GridPP to fill the gaps

OpenStack Blazar

StackHPC

- Reserve resources when you need them
  - We made it easier to start a server inside a reservation
  - We limit what can be reserved based on cloud credits (i.e. IRIS allocation of CPU hours)
  - Currently you reserve an integer number of hypervisors (future plans: instance reservation)
- Allow preemptible instances in unreserved space
  - Minimum guaranteed lifetime, defaults to one hour
- Added utilization reporting
  - (WIP) Utilization metering using Cloud Kitty
- Allocation no longer has to be constant throughout the whole year
- Stop people “squatting” on resources so others don’t “steal” them
- Easier to understand how to balance / negotiate future resource allocation

OpenStack Utilization using Blazar

StackHPC

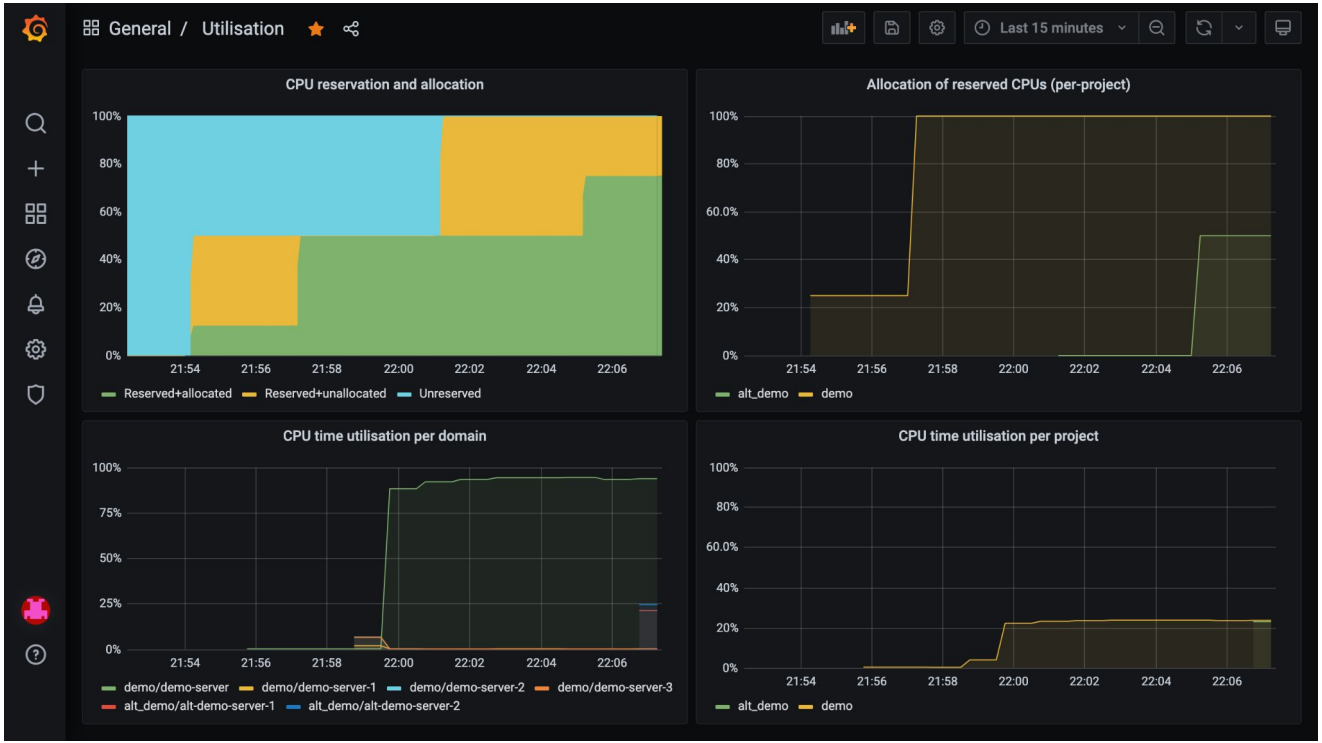

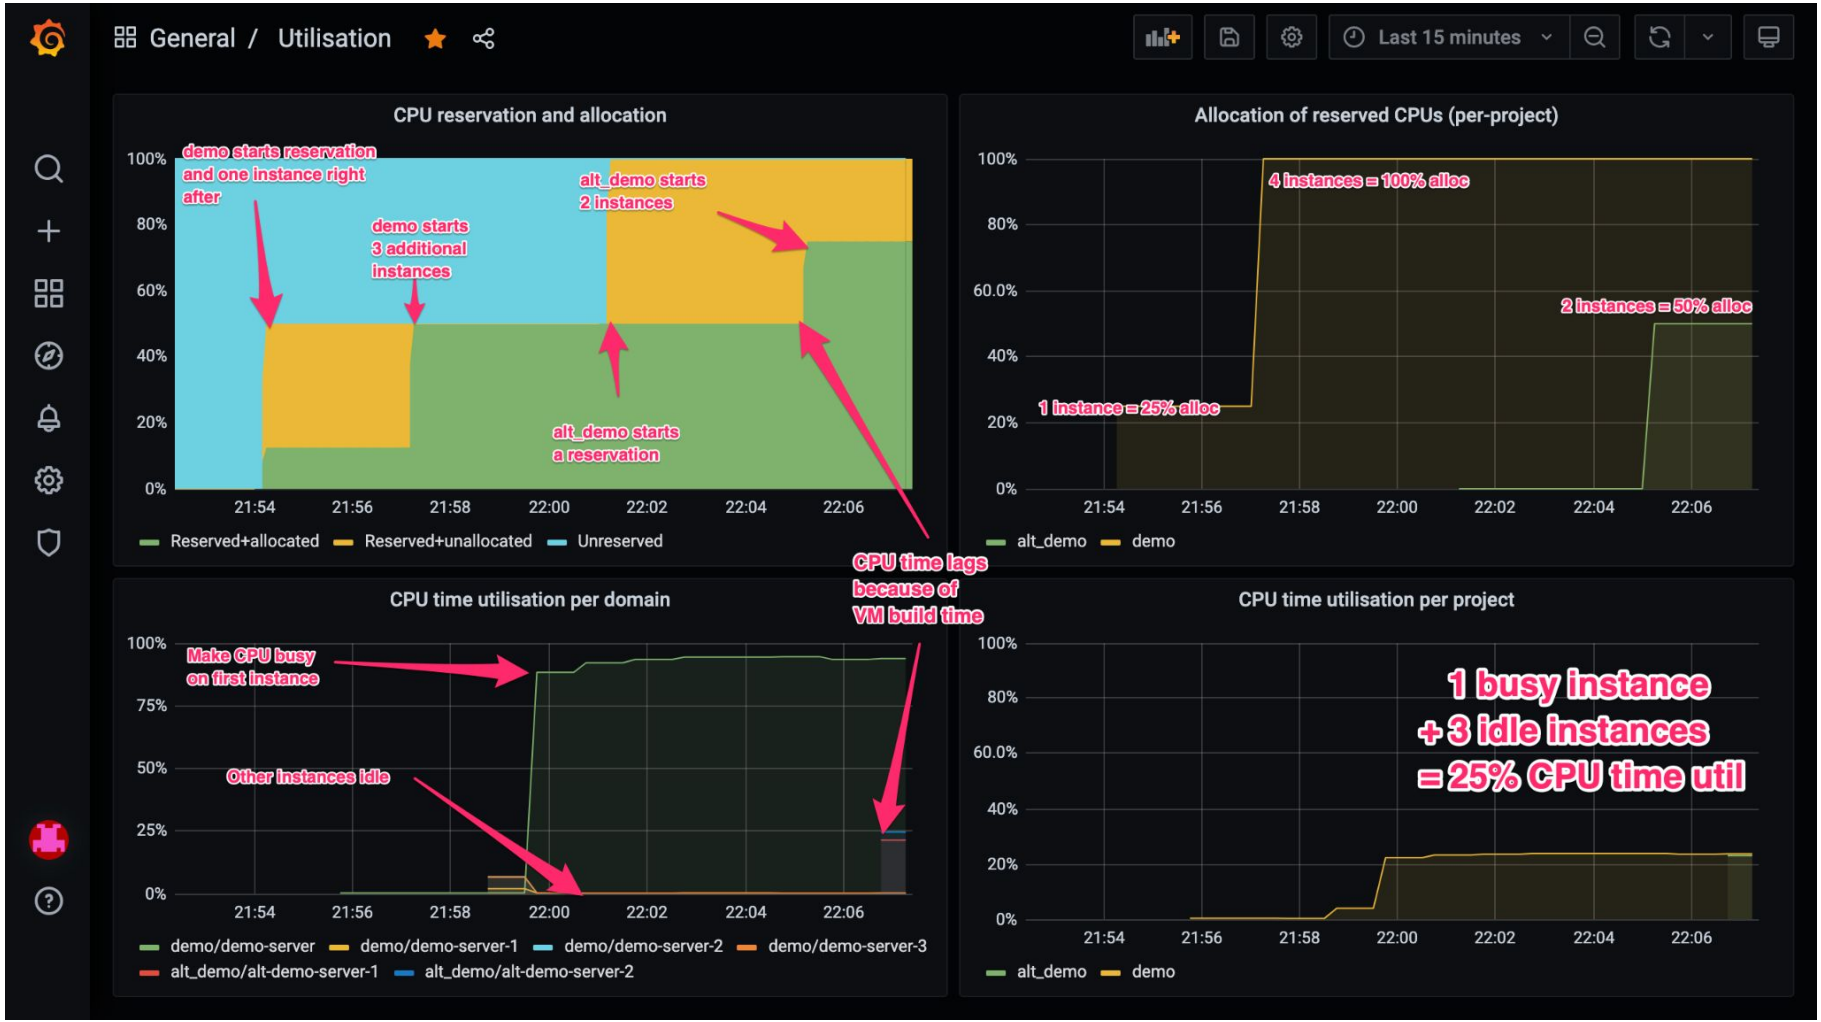
General / Utilisation
Last 15 minutes
CPU reservation and allocation
demo starts reservation and one instance right after
demo starts 3 additional instances
alt_demo starts 2 instances
alt_demo starts a reservation
100%
80%
60%
40%
20%
0%
21:54
21:56
21:58
22:00
22:02
22:04
22:06
Reserved+allocated
Reserved+unallocated
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 is next for Blazar?

StackHPC

- More testing and **more operator and user feedback required!**
- Chameleon Cloud's Blazar Calendar UI
- Preemptible instances in unused reserved space
- Usability refinements around requesting and modifying reservations

CI / CD Improvements

StackHPC

OpenStack Kayobe Deployment

StackHPC

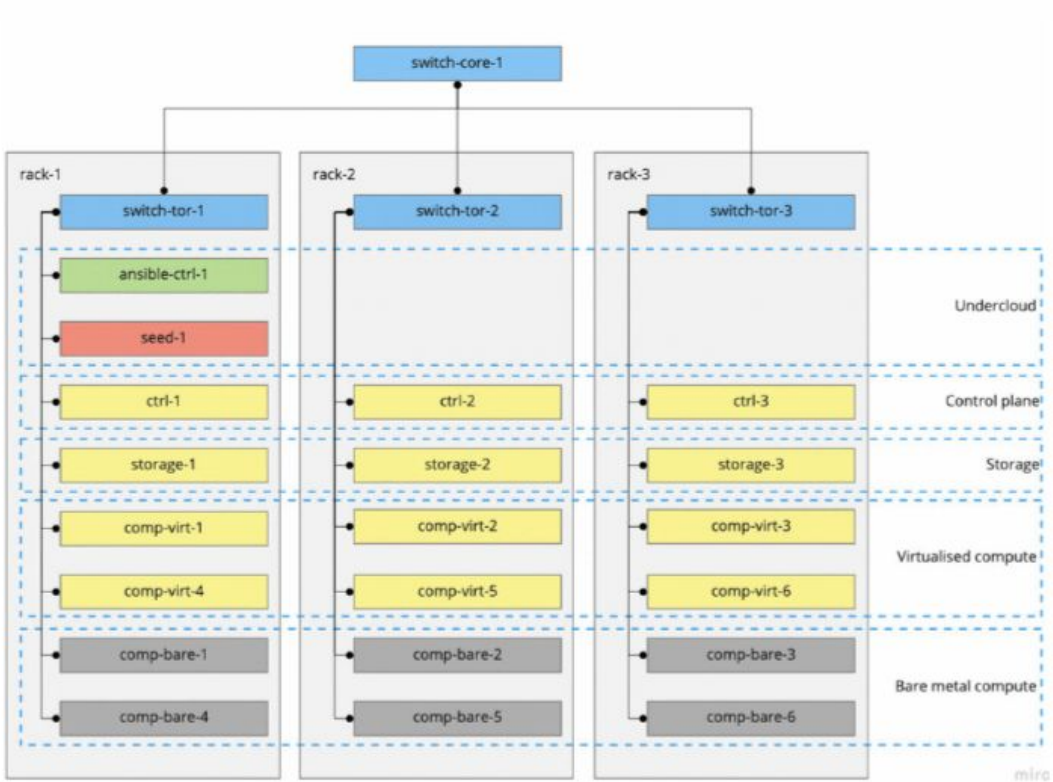

IRONIC
an OpenStack Community Project

OpenStack Kayobe Architecture

StackHPC

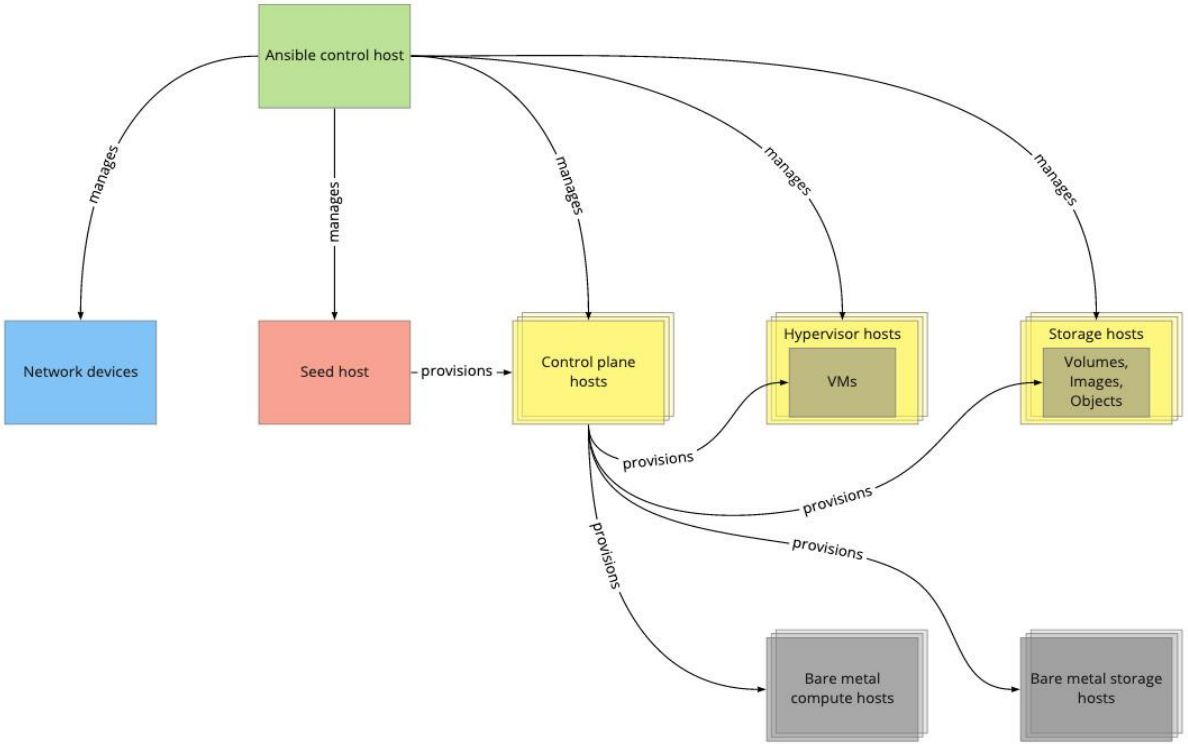

IRONIC
an OpenStack Community Project

OpenStack Continuous Integration

StackHPC

Better testing of merge requests:

- Sharing configuration between multiple environments
- Configuration Diff on merge requests
- Create and deploy all-in-one cloud, run OpenStack RefStack tests
- Building of binary artifacts: container images, IPA ramdisks, etc
- (WIP) Test new K8s cluster templates and run Sonoboy conformance tests
- (WIP) Using Pulp to snapshot OS repositories

OpenStack GitLab CI/CD

StackHPC

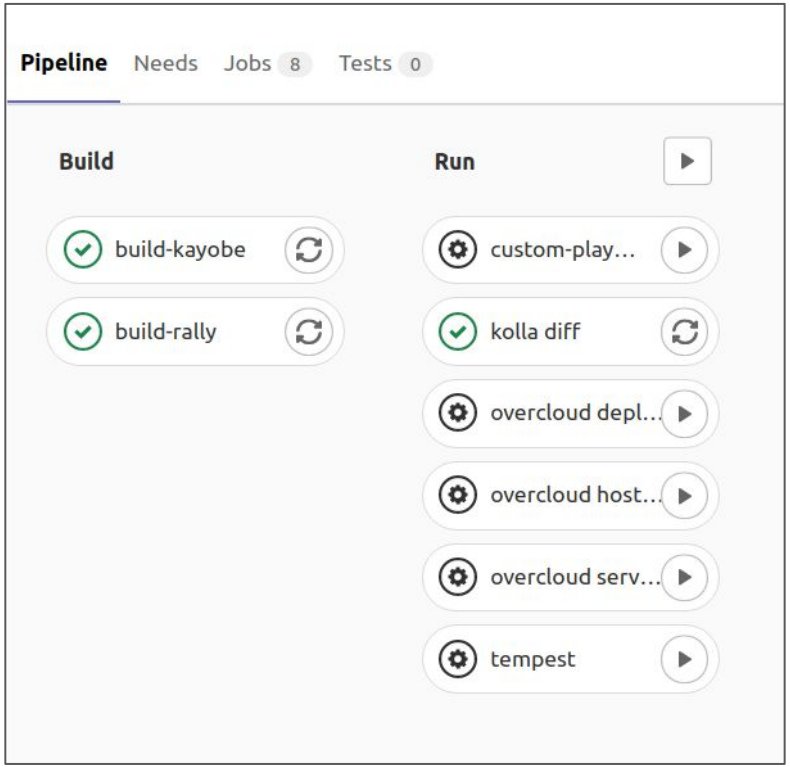

GitLab CI/CD

StackHPC

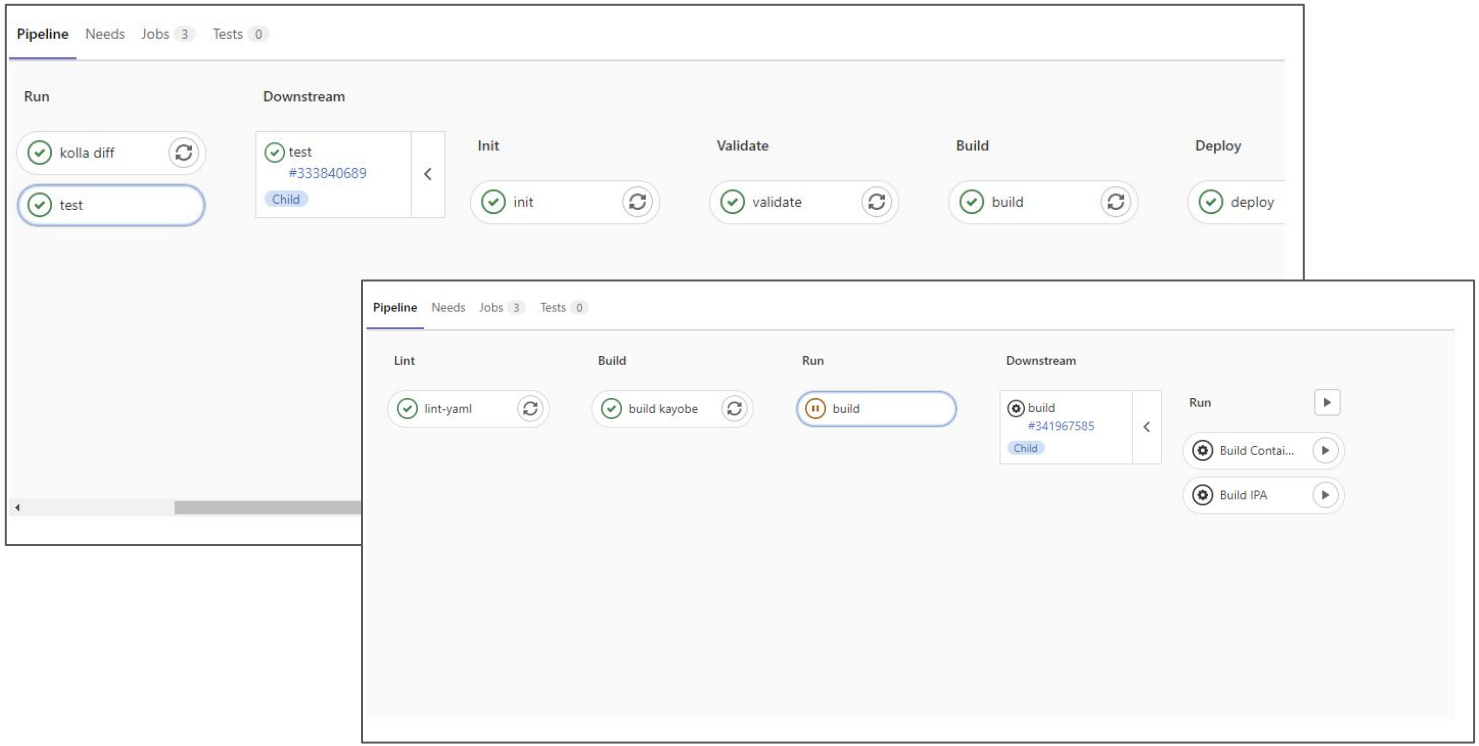

GitLab CI/CD

StackHPC

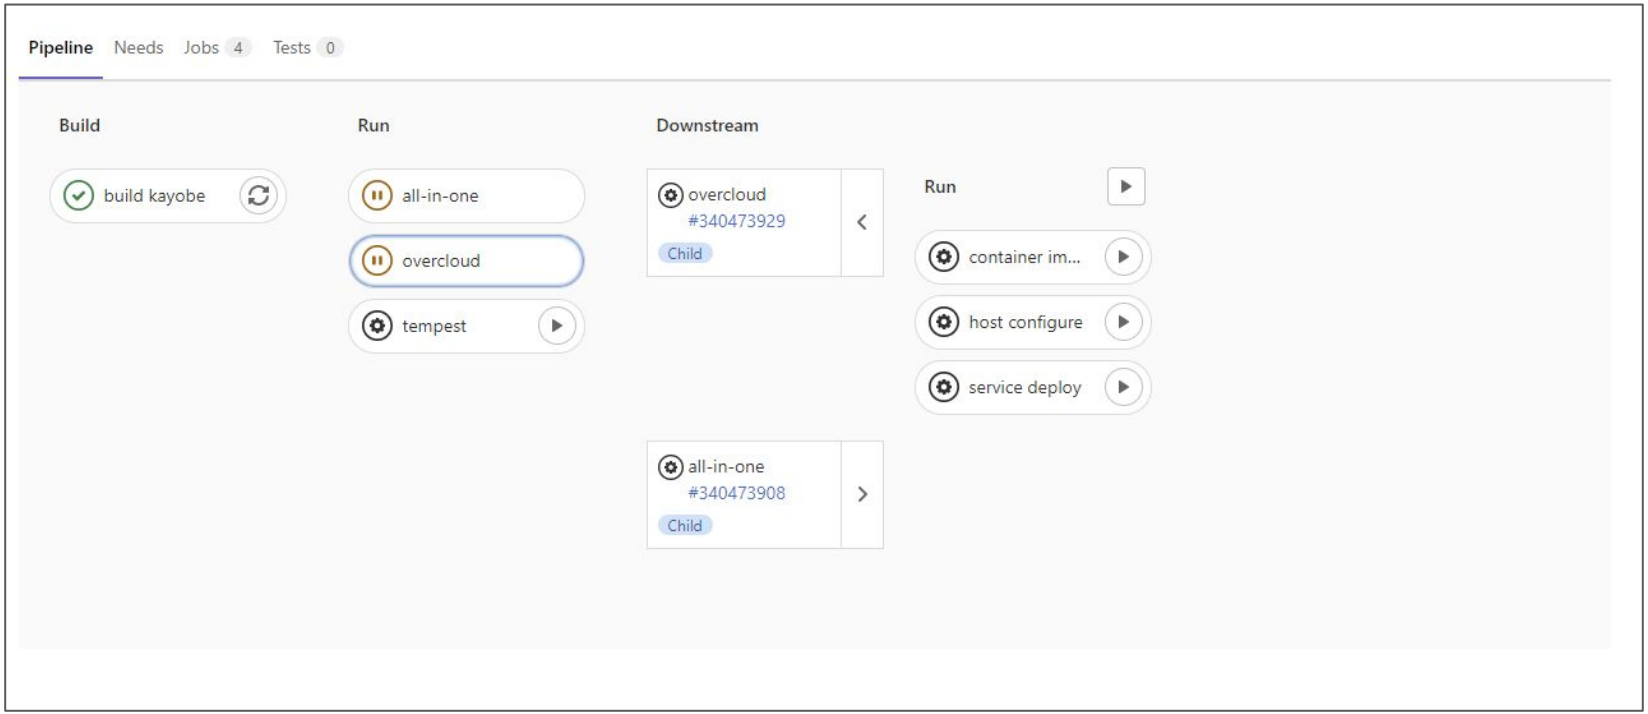

Improved Testing

StackHPC

- OpenStack
  - Rally uses tempest to run the RefStack conformance tests
  - Compliments the improved monitoring (not covered here)
- Slurm
  - Many tools packaged as an ansible collection
  - Multiple environments in a single branch
  - … some of this tooling has now helped to build several Top500 supercomputers
- Kubernetes (WIP)
  - Investigate kubeapps to ease deployment of helm charts
  - Investigate Harbor as a Helm repository that packages site specific enhancements
  - Document using helm chart versions to manage upgrades
  - Looking at Flux v2.0 vs Argo for platform development workflows

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

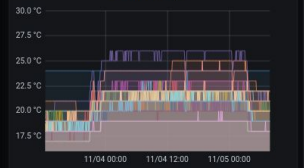

	max	avg	current
max inlet h22b8	26.00 °C	24.20 °C	21.00 °C
max inlet h24c5	25.00 °C	22.59 °C	21.00 °C
max inlet h21c5	25.00 °C	22.74 °C	22.00 °C
max inlet h22d8	24.00 °C	24.00 °C	24.00 °C
max inlet h24b8	23.00 °C	21.36 °C	22.00 °C

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

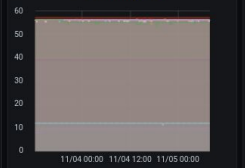

	max	current
Powered up h22d8	57.00	57.00
Powered up h22d5	57.00	57.00
Powered up h24d5	56.00	56.00
Powered up h24c5	56.00	56.00
Powered up h24b8	56.00	56.00

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

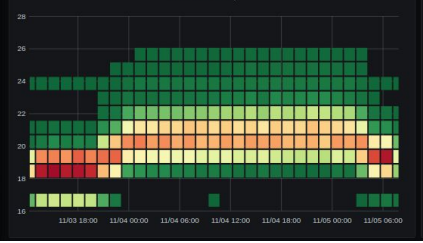

Max server fan speed

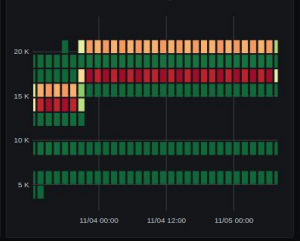

CPU1 Temp

CPU2 Temp

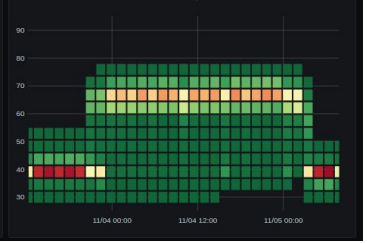


CPU usage (system)
25%
20%
15%
10%
5%
0%
percent
11:40
11:42
11:44
11:46
11:48
11:50
11:52
11:54
11:56
11:58
max
avg
node_cpu_system_seconds:record{instance="proto-hpc-2"} 20.6% 9.12%
node_cpu_system_seconds:record{instance="proto-hpc-1"} 20.6% 8.90%
CPU usage (user)
100%
75%
50%
25%
0%
max
avg
node_cpu_user_seconds:record{instance="proto-hpc-0"} 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

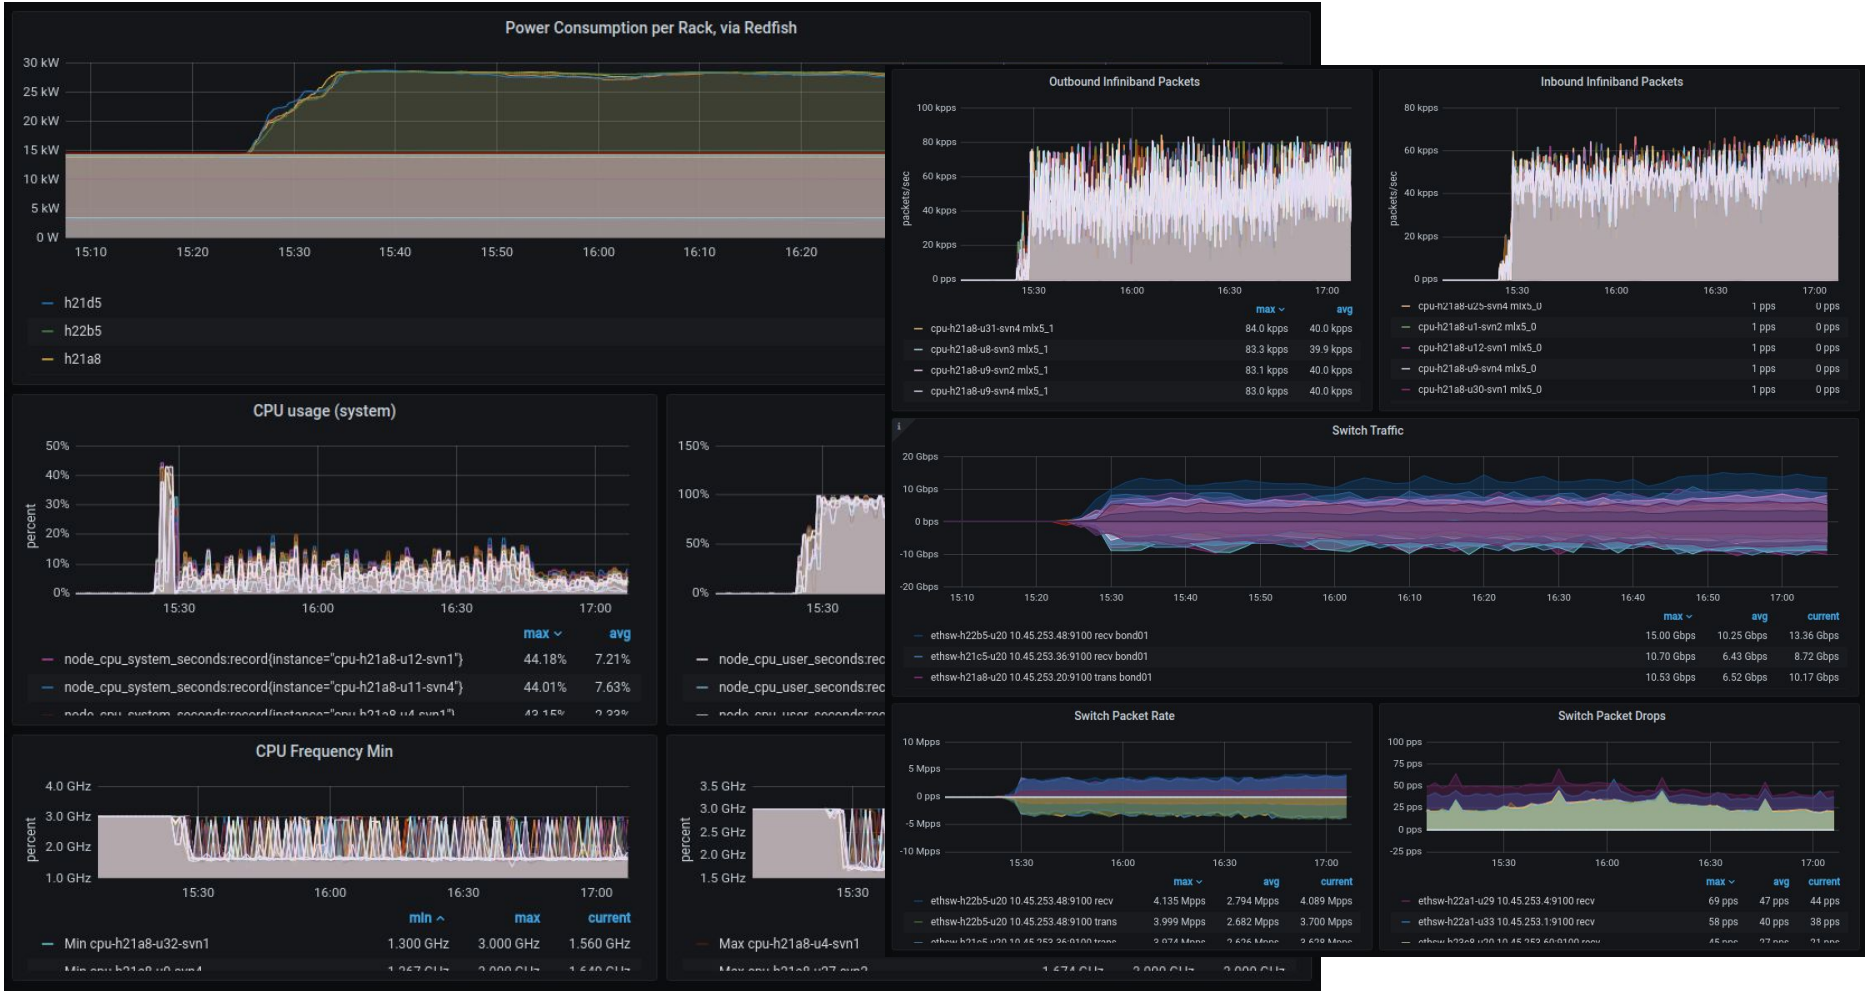
Power Consumption per Rack, via Redfish
30 kW
25 kW
20 kW
15 kW
10 kW
5 kW
0 W
15:10
15:20
15:30
15:40
15:50
16:00
16:10
16:20
h21d5
h22b5
h21a8
Outbound Infiniband Packets
100 kpps
80 kpps
60 kpps
40 kpps
20 kpps
0 pps
packets/sec
15:30
16:00
16:30
17:00
max
avg
cpu-h21a8-u31-svn4 mlx5_1 84.0 kpps 40.0 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

Cloud Portal

StackHPC

HPC Stack 2.0

StackHPC

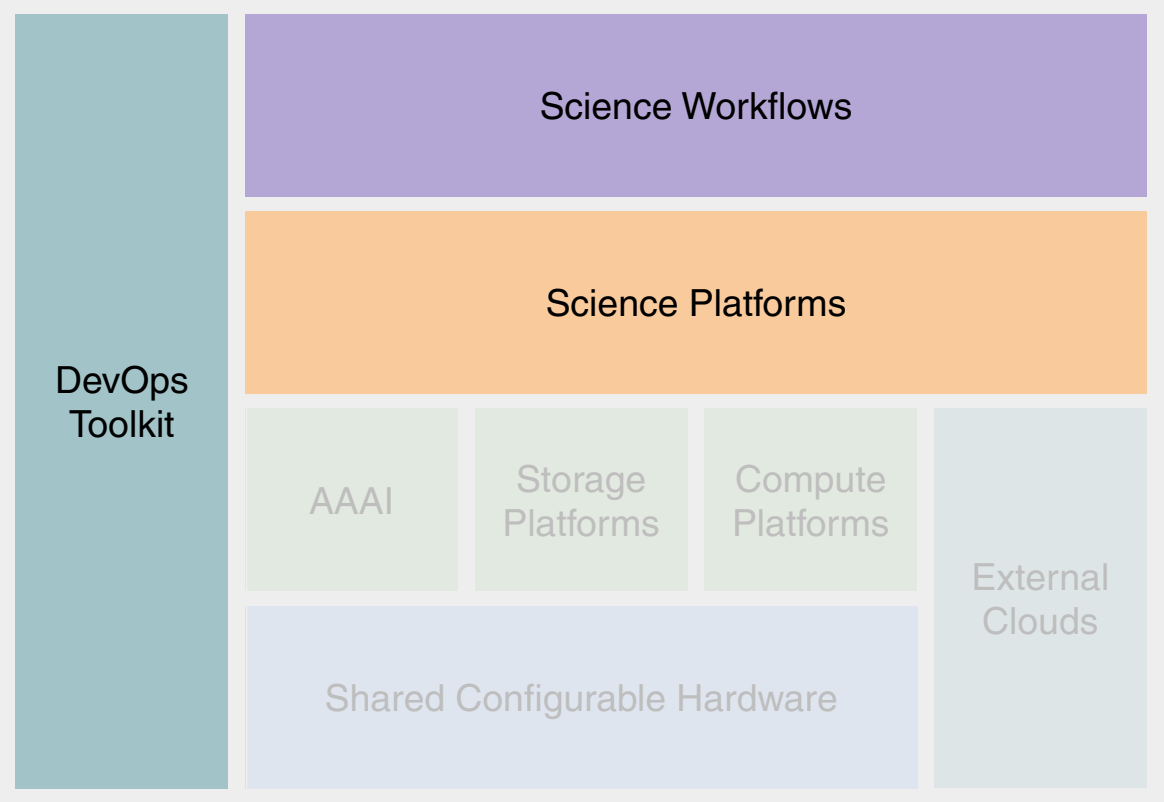

Cloud Portal

StackHPC

- Builds on the STFC funded work done by JASMIN
- Reduce time to science & reduce operational effort of onboarding
- Easier to **find** and **reuse** common lego bricks, between science communities
- Target use cases
  - Login with IRIS IAM to the projects associated with my groups
  - Get me a bigger laptop, or a bigger k8s clusters than my minikube
  - (WIP) Get me a Slurm Cluster with Open OnDemand, Get me a JupyterHub, ...
  - (WIP) Get me access to the monitoring dashboards
  - (WIP) Help me build my group its own cluster using the lego bricks
- Future ideas include:
  - Simpler access data sets your groups can access
  - Proxy security hardening
  - Mediate access to public cloud resources

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

StackHPC

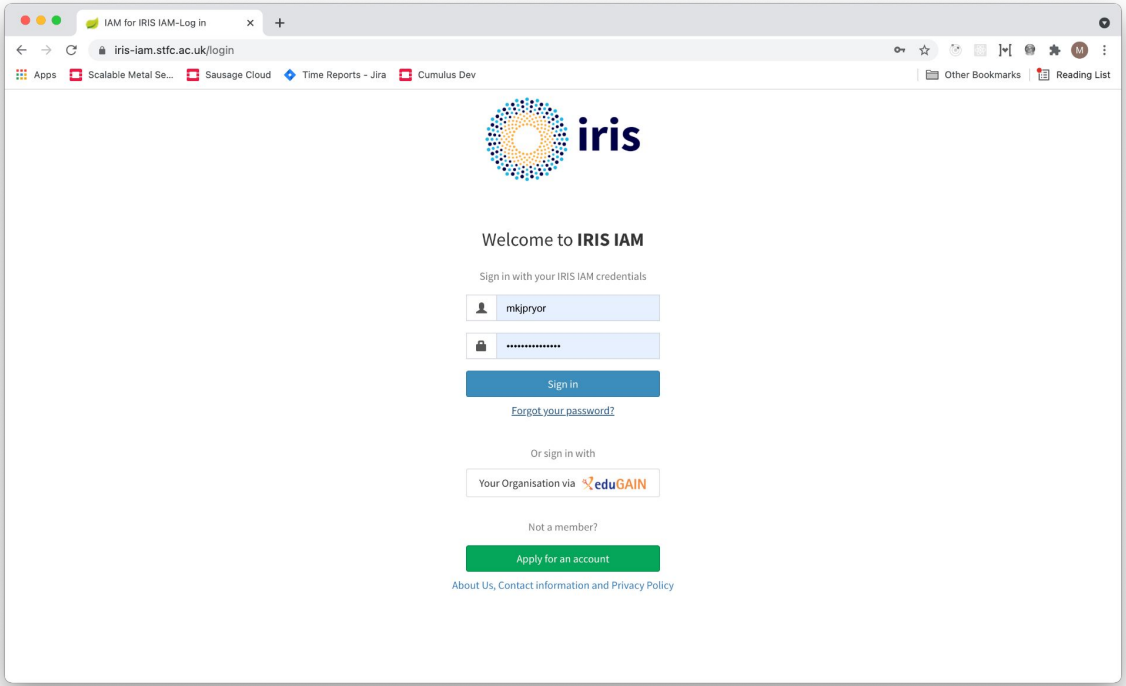
IAM for IRIS IAM-Log in
iris-iam.stfc.ac.uk/login
Apps
Scalable Metal Se...
Sausage Cloud
Time Reports - Jira
Cumulus Dev
Other Bookmarks
Reading List
iris
Welcome to IRIS IAM
Sign in with your IRIS IAM credentials
mkjpryor
Sign in
Forgot your password?
Or sign in with
Your Organisation via eduGAIN
Not a member?
Apply for an account
About Us, Contact information and Privacy Policy

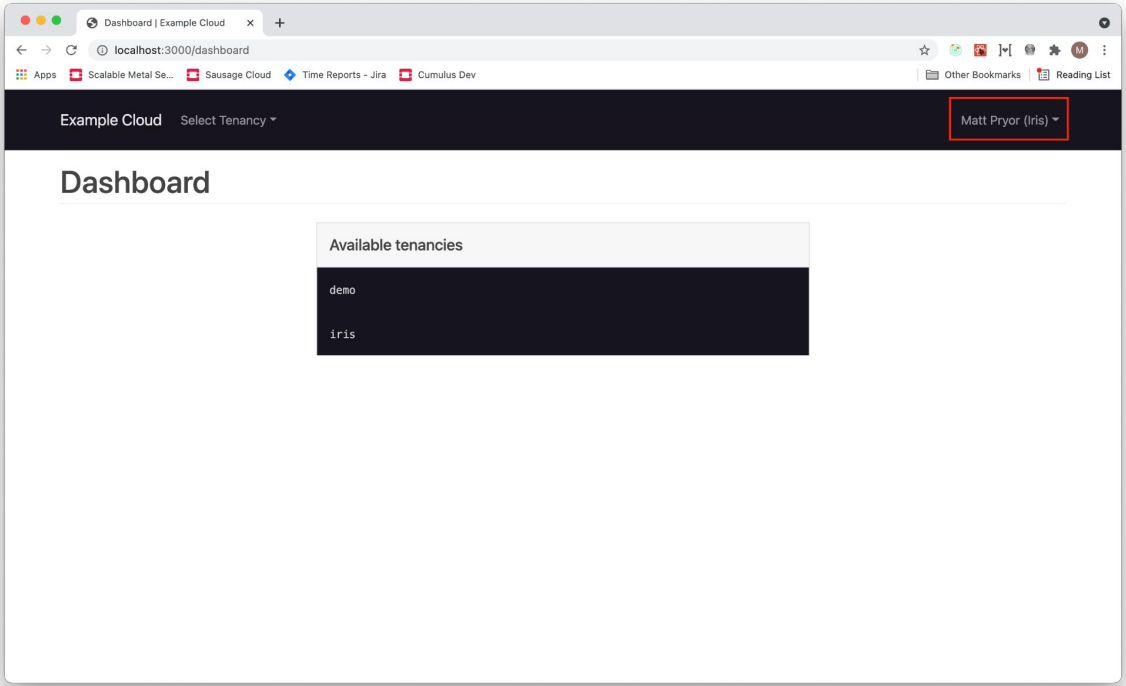
Dashboard | Example Cloud
localhost:3000/dashboard
Apps
Scalable Metal Se...
Sausage Cloud
Time Reports - Jira
Cumulus Dev
Other Bookmarks
Reading List
Example Cloud
Select Tenancy
Matt Pryor (Iris)
Dashboard
Available tenancies
demo
iris

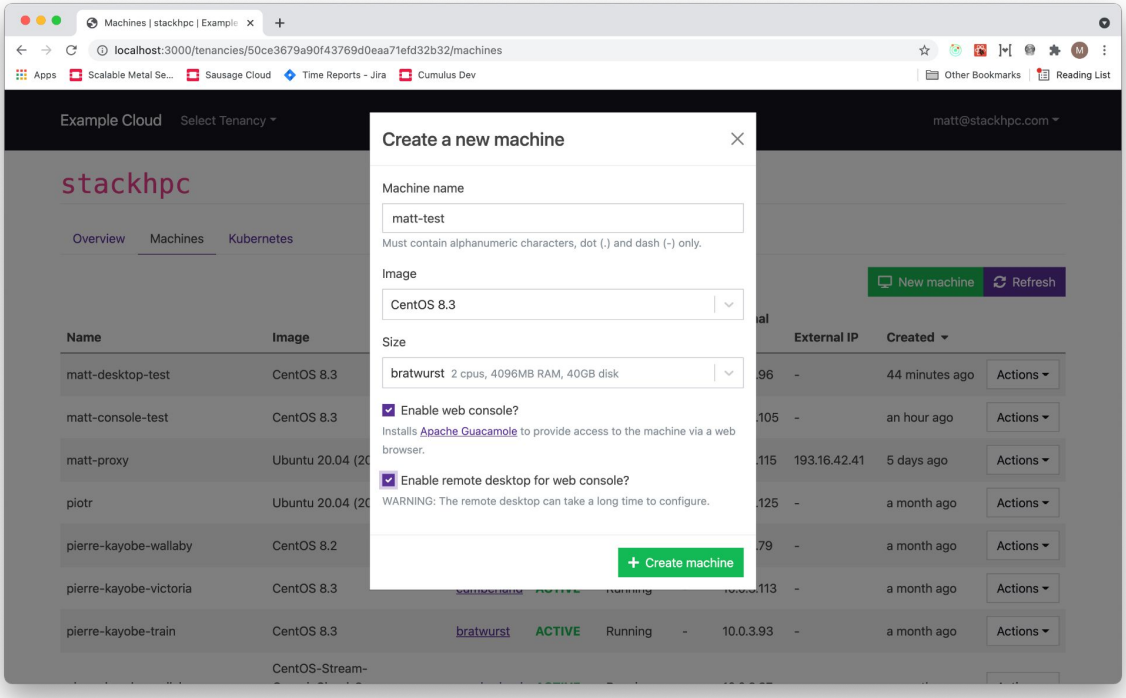
Machines | stackhpc | Example
localhost:3000/tenancies/50ce3679a90f43769d0eaa71efd32b32/machines
Apps
Scalable Metal Se...
Sausage Cloud
Time Reports - Jira
Cumulus Dev
Other Bookmarks
Reading List
Example Cloud
Select Tenancy
matt@stackhpc.com
stackhpc
Overview
Machines
Kubernetes
New machine
Refresh
Name
Image
External IP
Created
matt-desktop-test
CentOS 8.3
44 minutes ago
Actions
matt-console-test
CentOS 8.3
an hour ago
Actions
matt-proxy
Ubuntu 20.04 (20
193.16.42.41
5 days ago
Actions
piotr
Ubuntu 20.04 (20
a month ago
Actions
pierre-kayobe-wallaby
CentOS 8.2
a month ago
Actions
pierre-kayobe-victoria
CentOS 8.3
a month ago
Actions
pierre-kayobe-train
CentOS 8.3
bratwurst
ACTIVE
Running
10.0.3.93
a month ago
Actions
CentOS-Stream-
Create a new machine
Machine name
matt-test
Must contain alphanumeric characters, dot (.) and dash (-) only.
Image
CentOS 8.3
Size
bratwurst 2 cpus, 4096MB RAM, 40GB 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 configure.
Create machine

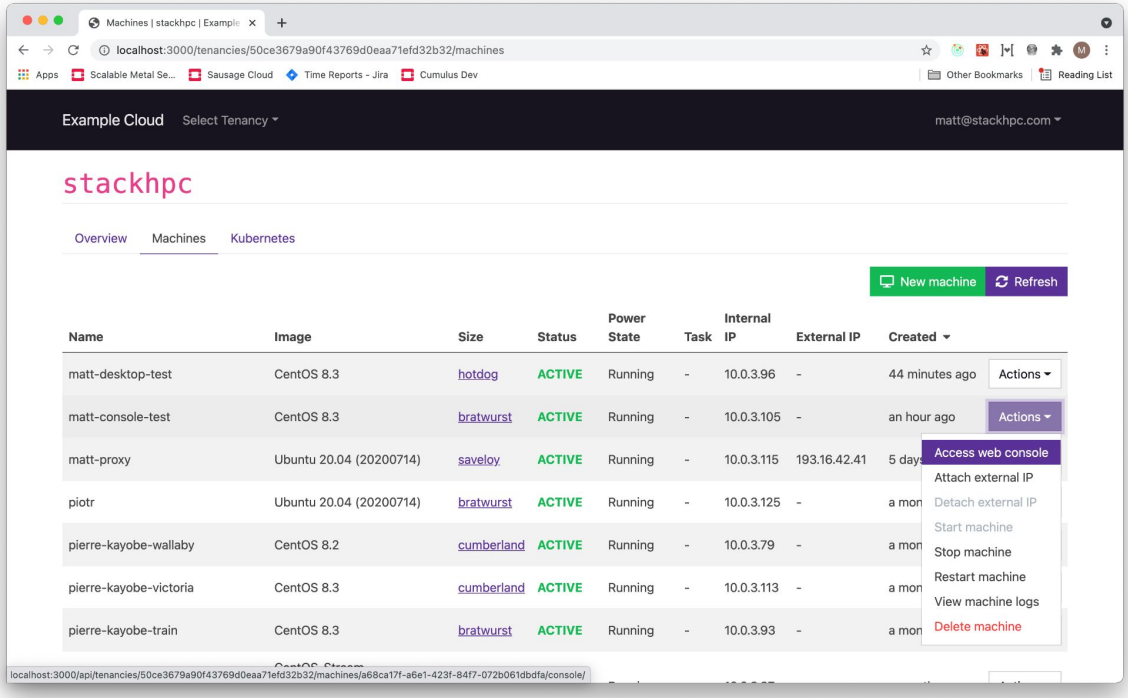
Machines | stackhpc | Example
localhost:3000/tenancies/50ce3679a90f43769d0eaa71efd32b32/machines
Apps
Scalable Metal Se...
Sausage Cloud
Time Reports - Jira
Cumulus Dev
Other Bookmarks
Reading List
Example Cloud
Select Tenancy
matt@stackhpc.com
stackhpc
Overview
Machines
Kubernetes
New machine
Refresh
Name
Image
Size
Status
Power State
Task
Internal IP
External IP
Created
matt-desktop-test
CentOS 8.3
hotdog
ACTIVE
Running
-
10.0.3.96
-
44 minutes ago
Actions
matt-console-test
CentOS 8.3
bratwurst
ACTIVE
Running
-
10.0.3.105
-
an hour ago
Actions
matt-proxy
Ubuntu 20.04 (20200714)
saveloy
ACTIVE
Running
-
10.0.3.115
193.16.42.41
5 days
piotr
Ubuntu 20.04 (20200714)
bratwurst
ACTIVE
Running
-
10.0.3.125
-
a mon
pierre-kayobe-wallaby
CentOS 8.2
cumberland
ACTIVE
Running
-
10.0.3.79
-
a mon
pierre-kayobe-victoria
CentOS 8.3
cumberland
ACTIVE
Running
-
10.0.3.113
-
a mon
pierre-kayobe-train
CentOS 8.3
bratwurst
ACTIVE
Running
-
10.0.3.93
-
a mon
Access web console
Attach external IP
Detach external IP
Start machine
Stop machine
Restart machine
View machine logs
Delete machine
localhost:3000/api/tenancies/50ce3679a90f43769d0eaa71efd32b32/machines/a68ca17f-a6e1-423f-84f7-072b061dbdfa/console/

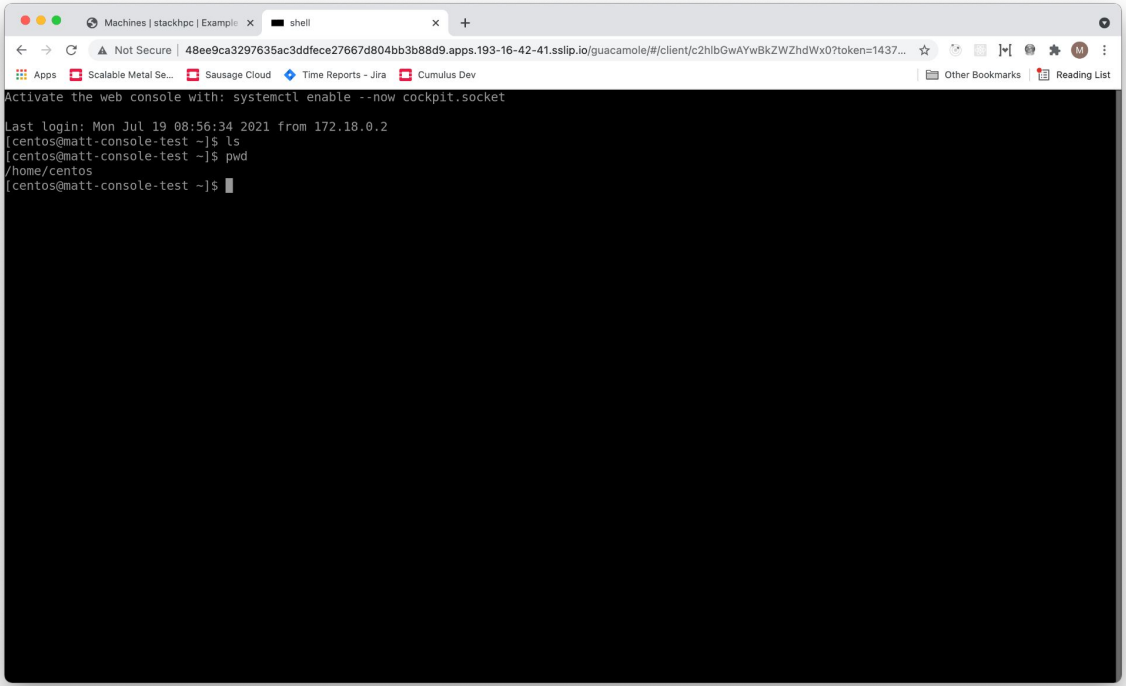
Machines | stackhpc | Example
shell
Not Secure | 48ee9ca3297635ac3ddfece27667d804bb3b88d9.apps.193-16-42-41.sslip.io/guacamole/#/client/c2hlbGwAYwBkZWZhdWx0?token=1437...
Apps
Scalable Metal Se...
Sausage Cloud
Time Reports - Jira
Cumulus Dev
Other Bookmarks
Reading List
Activate the web console with: systemctl enable --now cockpit.socket
Last login: Mon Jul 19 08:56:34 2021 from 172.18.0.2
[centos@matt-console-test ~]$ ls
[centos@matt-console-test ~]$ pwd
/home/centos
[centos@matt-console-test ~]$

How did you get access to that VM?

StackHPC

- 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
- Proxy security needs improving in future work packages
  - But this demonstrates the utility of the proxy
  - And work has been done to design with those improvements from the beginning
  - HA possible via use of Nginx, consul, consul-template, ssh
  - https://github.com/stackhpc/tunnel-server

Cloud Portal:
Get me a K8s Cluster, via IRIS IAM

StackHPC

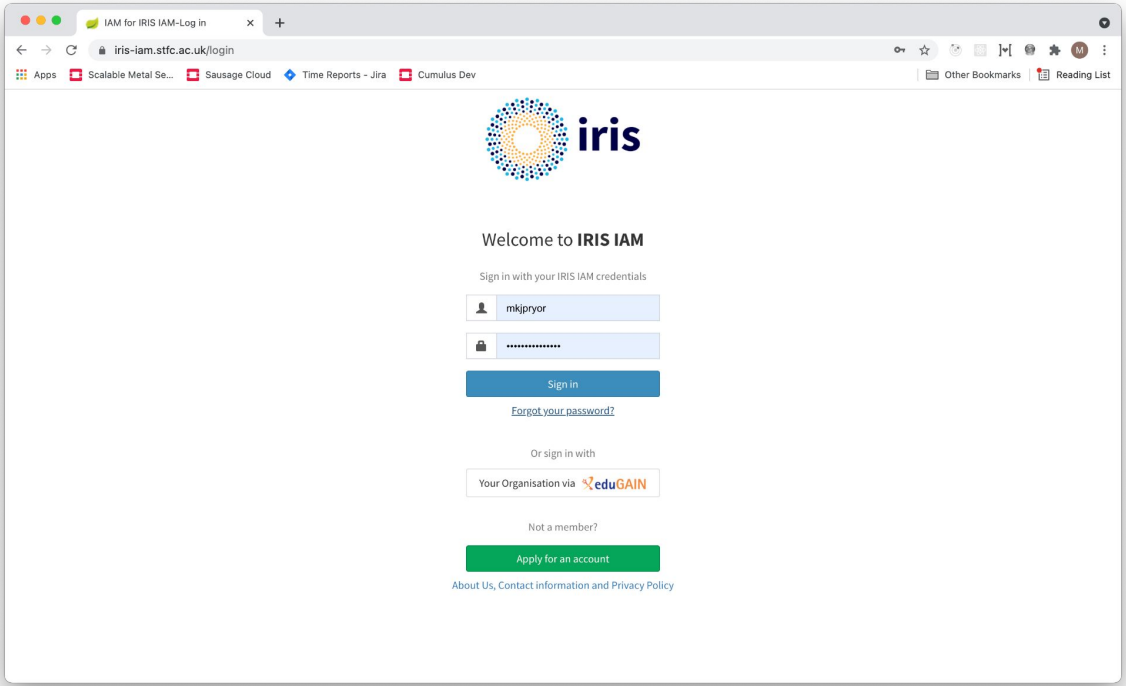
IAM for IRIS IAM-Log in
iris-iam.stfc.ac.uk/login
Apps
Scalable Metal Se...
Sausage Cloud
Time Reports - Jira
Cumulus Dev
Other Bookmarks
Reading List
iris
Welcome to IRIS IAM
Sign in with your IRIS IAM credentials
mkjpryor
Sign in
Forgot your password?
Or sign in with
Your Organisation via eduGAIN
Not a member?
Apply for an account
About Us, Contact information and Privacy Policy

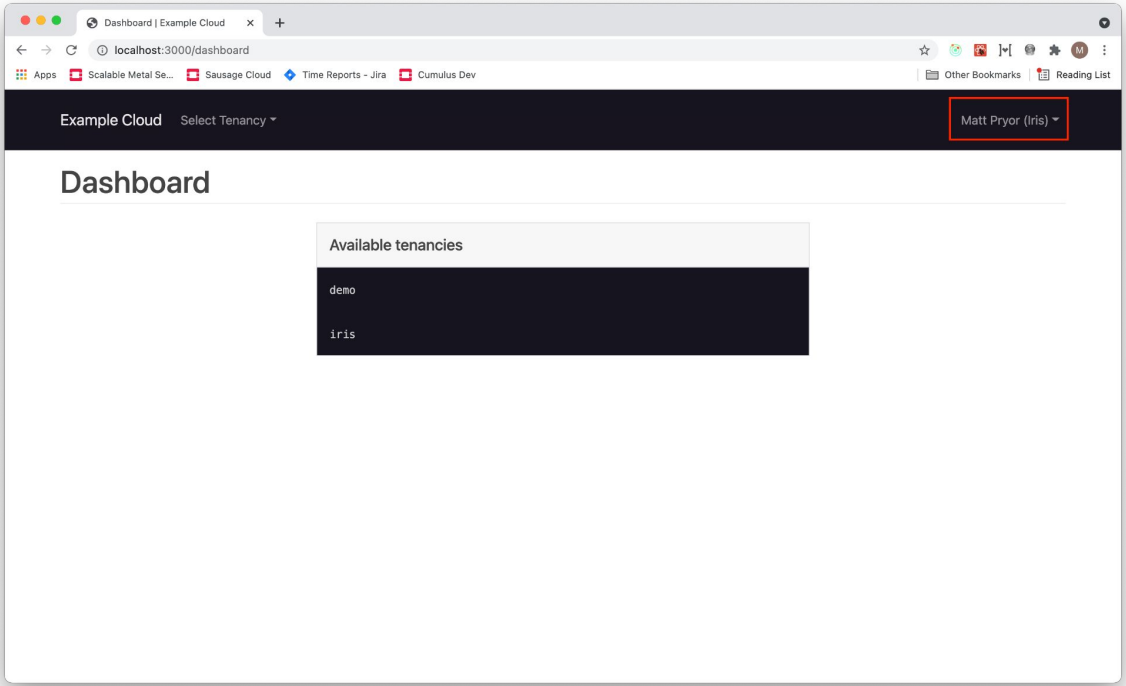
Dashboard | Example Cloud
localhost:3000/dashboard
Apps
Scalable Metal Se...
Sausage Cloud
Time Reports - Jira
Cumulus Dev
Other Bookmarks
Reading List
Example Cloud
Select Tenancy
Matt Pryor (Iris)
Dashboard
Available tenancies
demo
iris

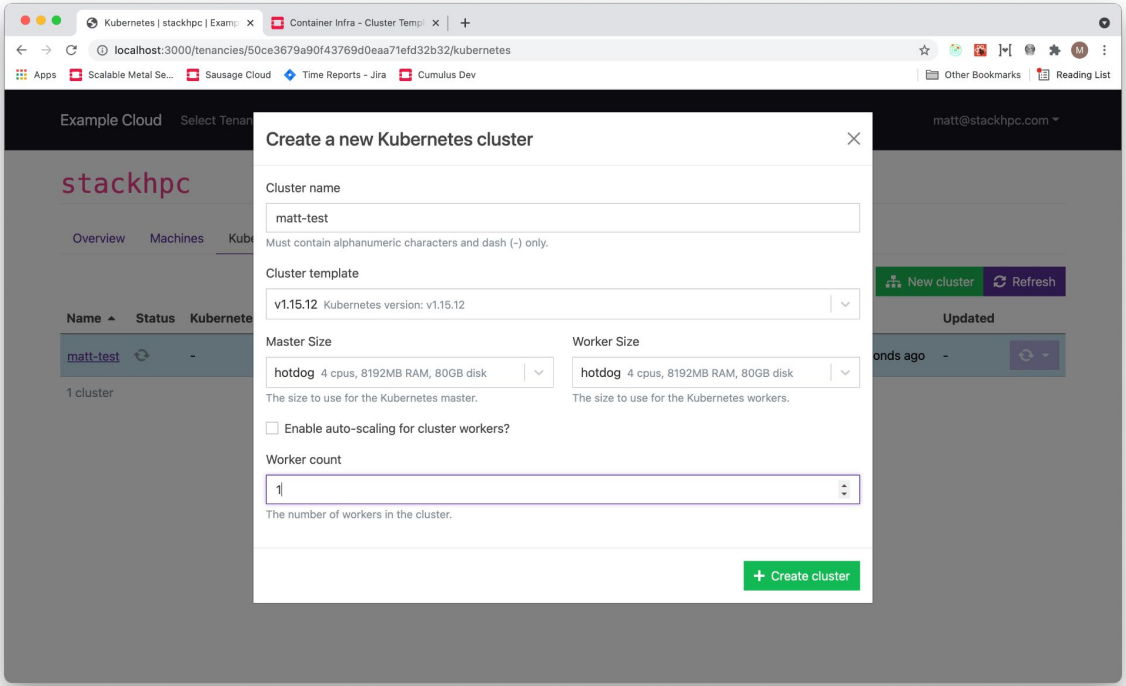
Kubernetes | stackhpc | Exam
Container Infra - Cluster Temp
localhost:3000/tenancies/50ce3679a90f43769d0eaa71efd32b32/kubernetes
Apps
Scalable Metal Se...
Sausage Cloud
Time Reports - Jira
Cumulus Dev
Other Bookmarks
Reading List
Example Cloud
matt@stackhpc.com
stackhpc
Overview
Machines
New cluster
Refresh
Name
Status
Updated
matt-test
1 cluster
Create a new Kubernetes cluster
Cluster name
matt-test
Must contain alphanumeric characters and dash (-) only.
Cluster template
v1.15.12 Kubernetes version: v1.15.12
Master Size
hotdog 4 cpus, 8192MB RAM, 80GB disk
The size to use for the Kubernetes master.
Worker Size
hotdog 4 cpus, 8192MB RAM, 80GB disk
The size to use for the Kubernetes workers.
Enable auto-scaling for cluster workers?
Worker count
1
The number of workers in the cluster.
Create cluster

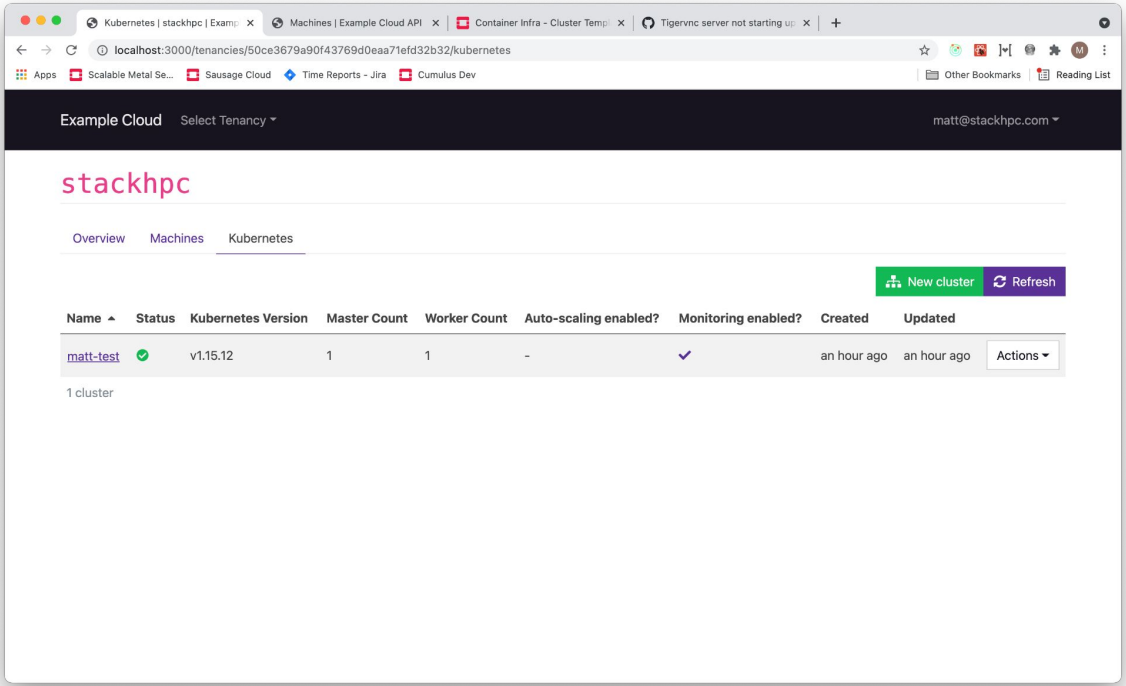
Kubernetes | stackhpc | Examp
Machines | Example Cloud API
Container Infra - Cluster Temp
Tigervnc server not starting up
localhost:3000/tenancies/50ce3679a90f43769d0eaa71efd32b32/kubernetes
Apps
Scalable Metal Se...
Sausage Cloud
Time Reports - Jira
Cumulus Dev
Other Bookmarks
Reading List
Example Cloud
Select Tenancy
matt@stackhpc.com
stackhpc
Overview
Machines
Kubernetes
New cluster
Refresh
Name
Status
Kubernetes Version
Master Count
Worker Count
Auto-scaling enabled?
Monitoring enabled?
Created
Updated
matt-test
v1.15.12
1
1
-
an hour ago
an hour ago
Actions
1 cluster

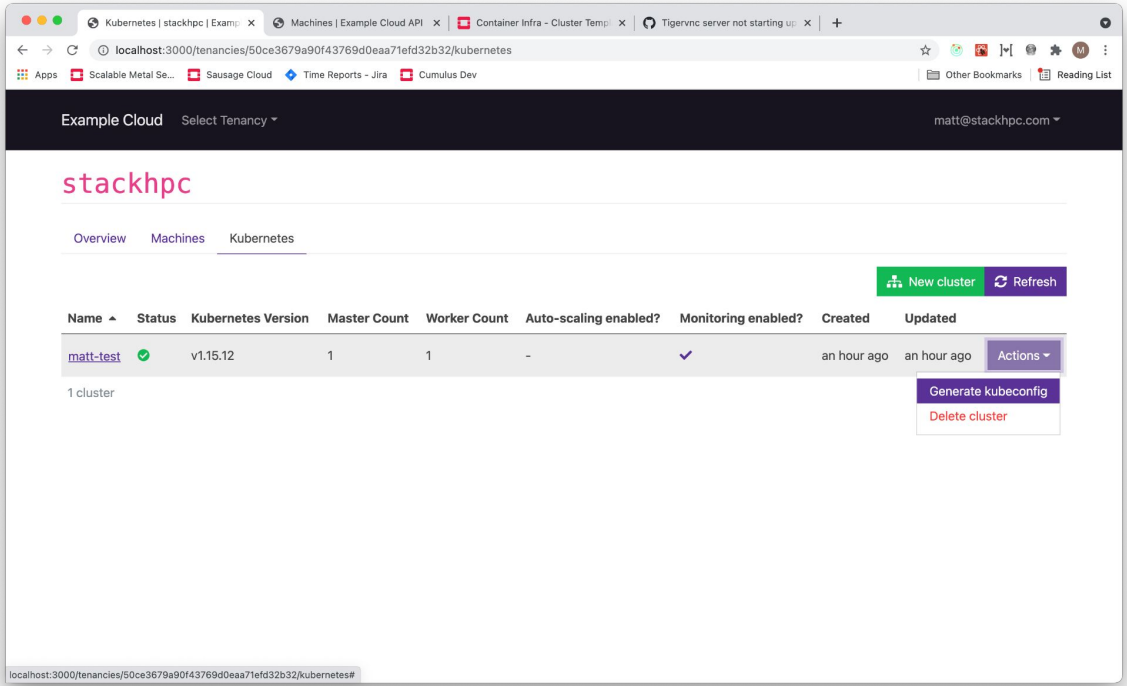
Kubernetes | stackhpc | Examp
Machines | Example Cloud API
Container Infra - Cluster Temp
Tigervnc server not starting up
localhost:3000/tenancies/50ce3679a90f43769d0eaa71efd32b32/kubernetes
Apps
Scalable Metal Se...
Sausage Cloud
Time Reports - Jira
Cumulus Dev
Other Bookmarks
Reading List
Example Cloud
Select Tenancy
matt@stackhpc.com
stackhpc
Overview
Machines
Kubernetes
New cluster
Refresh
Name
Status
Kubernetes Version
Master Count
Worker Count
Auto-scaling enabled?
Monitoring enabled?
Created
Updated
matt-test
v1.15.12
1
1
-
an hour ago
an hour ago
Actions
1 cluster
Generate kubeconfig
Delete cluster
localhost:3000/tenancies/50ce3679a90f43769d0eaa71efd32b32/kubernetes#

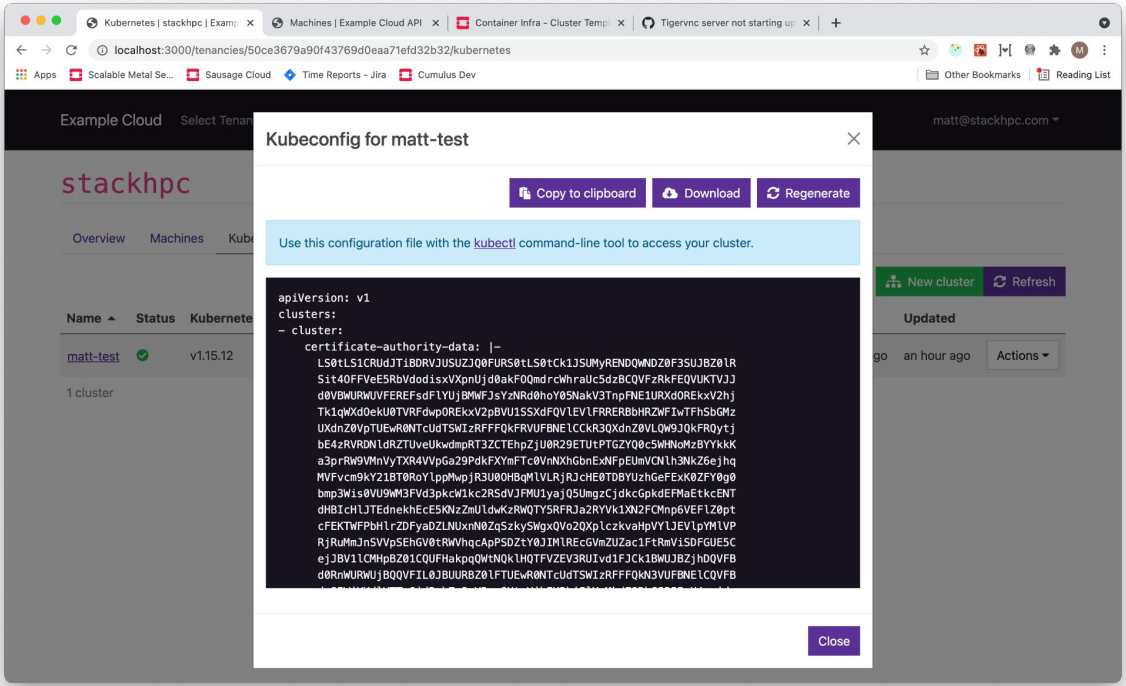
Kubernetes | stackhpc | Exam
Machines | Example Cloud API
Container Infra - Cluster Temp
Tigervnc server not starting up
localhost:3000/tenancies/50ce3679a90f43769d0eaa71efd32b32/kubernetes
Apps
Scalable Metal Se...
Sausage Cloud
Time Reports - Jira
Cumulus Dev
Other Bookmarks
Reading List
Example Cloud
matt@stackhpc.com
stackhpc
Overview
Machines
New cluster
Refresh
Name
Status
Updated
matt-test
v1.15.12
an hour ago
Actions
1 cluster
Kubeconfig for matt-test
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: |-
LS0tLS1CRUdJTiBDRVJUSUZJQ0FURS0tLS0tCk1JSUMyRENDQWNDZ0F3SUJBZ0lR
Sit4OFFVeE5RbVdodisxVXpnUjd0akFOQmdrcWhraUc5dzBCQVFzRkFEQVUKTVJJ
d0VBWURWUVFEREFsdFlYUjBMWFJsYzNRd0hoY05NakV3TnpFNE1URXdOREkxV2hj
Tk1qWXdOekU0TVRFdwpOREkxV2pBVU1SSXdFQVlEVlFRERBbHRZWFIwTFhSbGMz
UXdnZ0VpTUEwR0NTcUdTSWIzRFFFQkFRVUFBNElCCkR3QXdnZ0VLQW9JQkFRQytj
bE4zRVRDNldRZTUveUkwdmpRT3ZCTEhpZjU0R29ETUtPTGZYQ0c5WHNoMzBYYkkK
a3prRW9VMnVyTXR4VVpGa29PdkFXYmFTc0VnNXhGbnExNFpEUmVCNlh3NkZ6ejhq
MVFvcm9kY21BT0RoYlppMwpjR3U0OHBqMlVLRjRJcHE0TDBYUzhGeFExK0ZFY0g0
bmp3Wis0VU9WM3FVd3pkcW1kc2RSdVJFMU1yajQ5UmgzCjdkcGpkdEFMaEtkcENT
dHBIcHlJTEdnekhEcE5KNzZmUldwKzRWQTY5RFRJa2RYVk1XN2FCMnp6VEFlZ0pt
cFEKTWFPbHlrZDFyaDZLNUxnN0ZqSzkySWgxQVo2QXplczkvaHpVYlJEVlpYMlVP
RjRuMmJnSVVpSEhGV0tRWVhqcApPSDZtY0JIMlREcGVmZUZac1FtRmViSDFGUE5C
ejJBV1lCMHpBZ01CQUFHakpqQWtNQklHQTFVZEV3RUIvd1FJCk1BWUJBZjhDQVFB
d0RnWURWUjBQQVFIL0JBUURBZ0lFTUEwR0NTcUdTSWIzRFFFQkN3VUFBNElCQVFB
Close

How did you do create the K8s cluster?

StackHPC

- Similar to the VM, except create k8s cluster with Magnum
  - Magnum APIs to generate the kubeconfig
- Future work will look at exposing services running in K8s via the Proxy:
  - Grafana, Kubernetes Dashboard, and kubeapps are the first targets

Cloud Portal:
Get me a Slurm Cluster, via IRIS IAM (WIP)

StackHPC

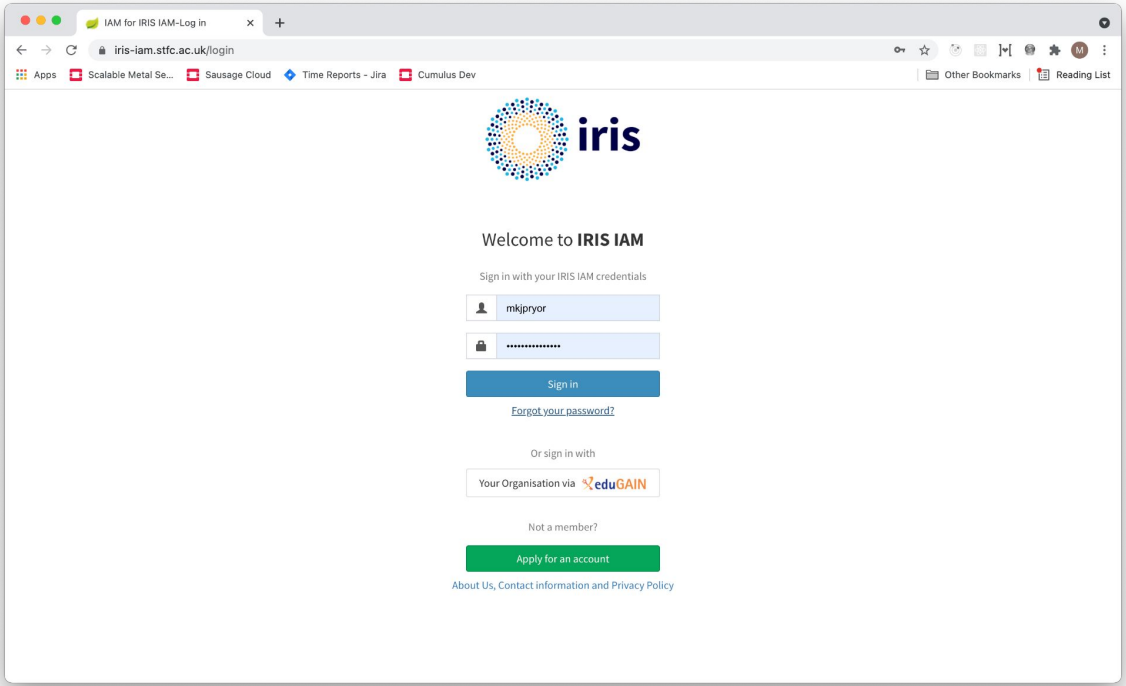
IAM for IRIS IAM-Log in
iris-iam.stfc.ac.uk/login
Apps
Scalable Metal Se...
Sausage Cloud
Time Reports - Jira
Cumulus Dev
Other Bookmarks
Reading List
iris
Welcome to IRIS IAM
Sign in with your IRIS IAM credentials
mkjpryor
Sign in
Forgot your password?
Or sign in with
Your Organisation via eduGAIN
Not a member?
Apply for an account
About Us, Contact information and Privacy Policy

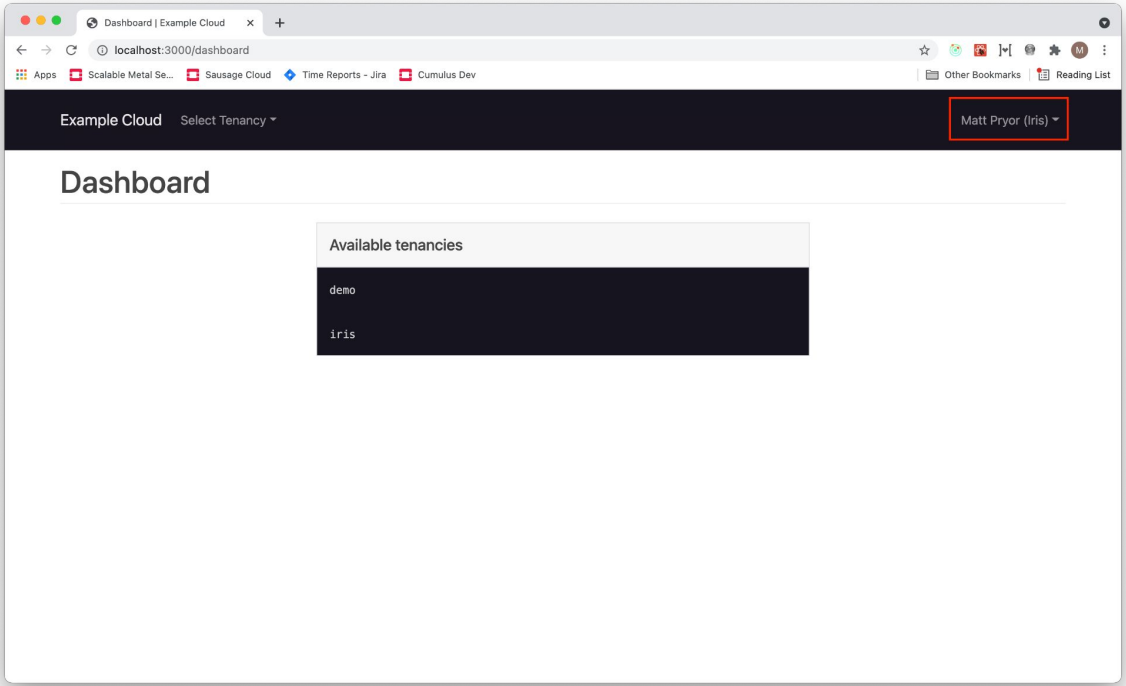
Dashboard | Example Cloud
localhost:3000/dashboard
Apps
Scalable Metal Se...
Sausage Cloud
Time Reports - Jira
Cumulus Dev
Other Bookmarks
Reading List
Example Cloud
Select Tenancy
Matt Pryor (Iris)
Dashboard
Available tenancies
demo
iris

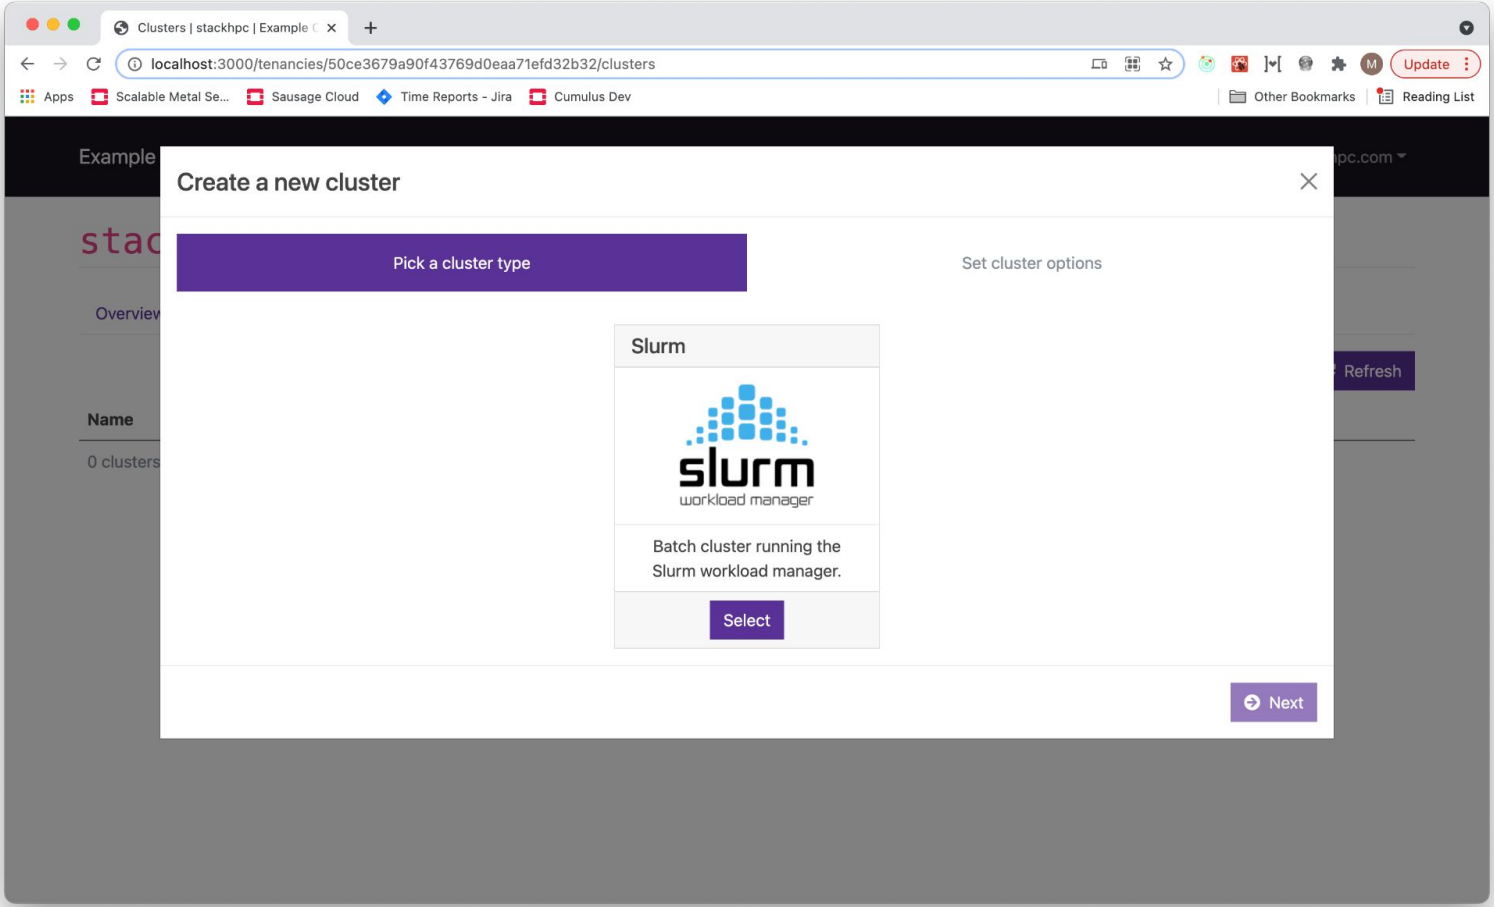
Clusters | stackhpc | Example
localhost:3000/tenancies/50ce3679a90f43769d0eaa71efd32b32/clusters
Update
Apps
Scalable Metal Se...
Sausage Cloud
Time Reports - Jira
Cumulus Dev
Other Bookmarks
Reading List
Example
Create a new cluster
Pick a cluster type
Set cluster options
Slurm
slurm
workload manager
Batch cluster running the Slurm workload manager.
Select
Next
Overview
Refresh
Name
0 clusters

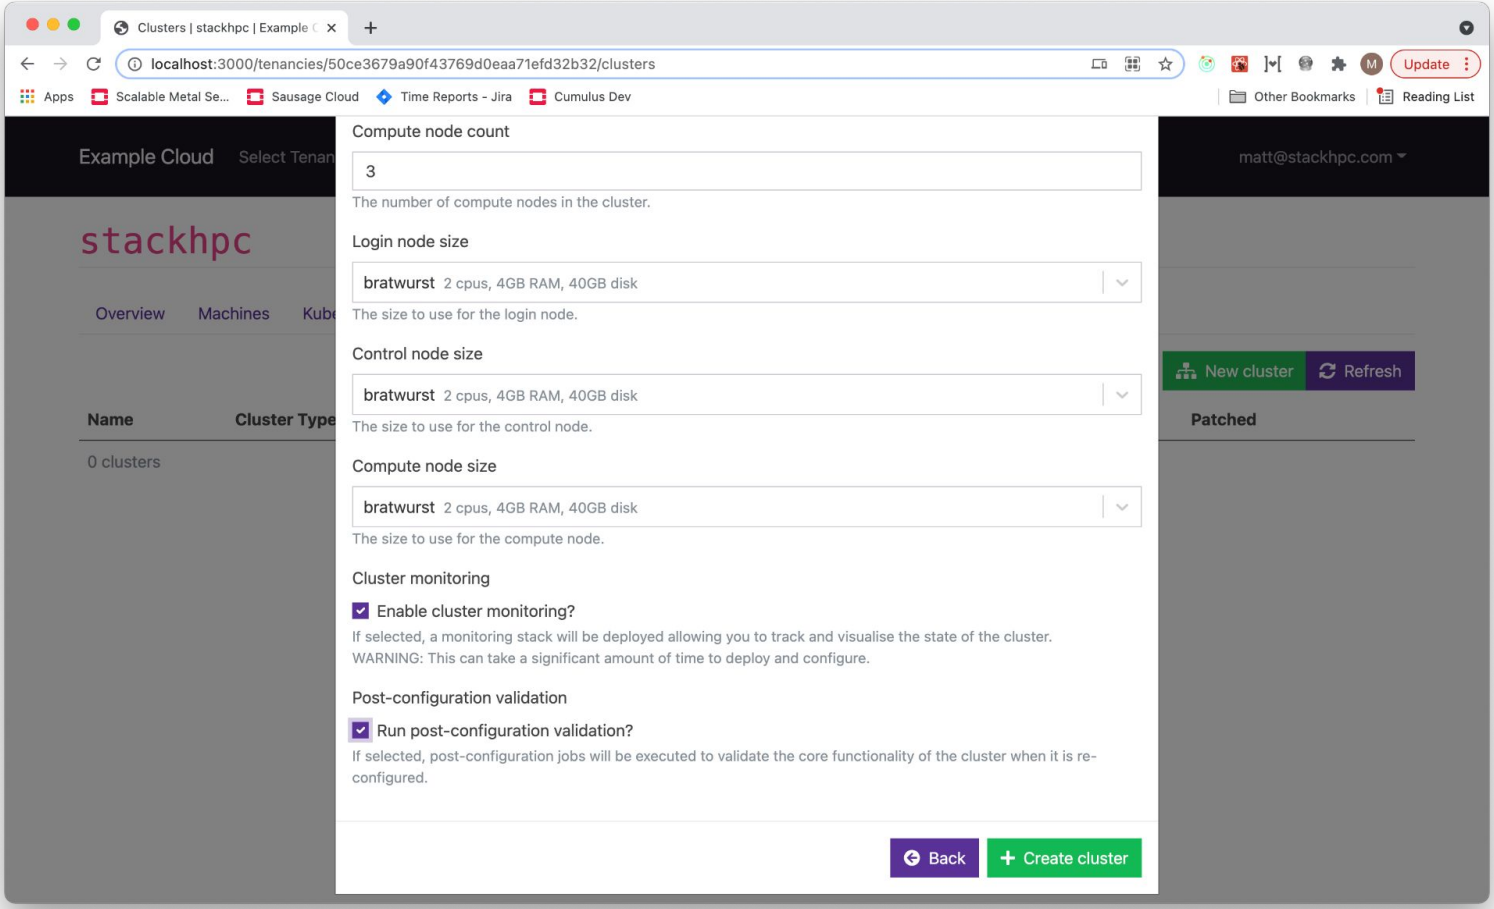
Clusters | stackhpc | Example
localhost:3000/tenancies/50ce3679a90f43769d0eaa71efd32b32/clusters
Update
Apps
Scalable Metal Se...
Sausage Cloud
Time Reports - Jira
Cumulus Dev
Other Bookmarks
Reading List
Example Cloud
Select Tenan
matt@stackhpc.com
stackhpc
Overview
Machines
Kube
New cluster
Refresh
Name
Cluster Type
Patched
0 clusters
Compute node count
3
The number of compute nodes in the cluster.
Login node size
bratwurst 2 cpus, 4GB RAM, 40GB disk
The size to use for the login node.
Control node size
bratwurst 2 cpus, 4GB RAM, 40GB disk
The size to use for the control node.
Compute node size
bratwurst 2 cpus, 4GB RAM, 40GB 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

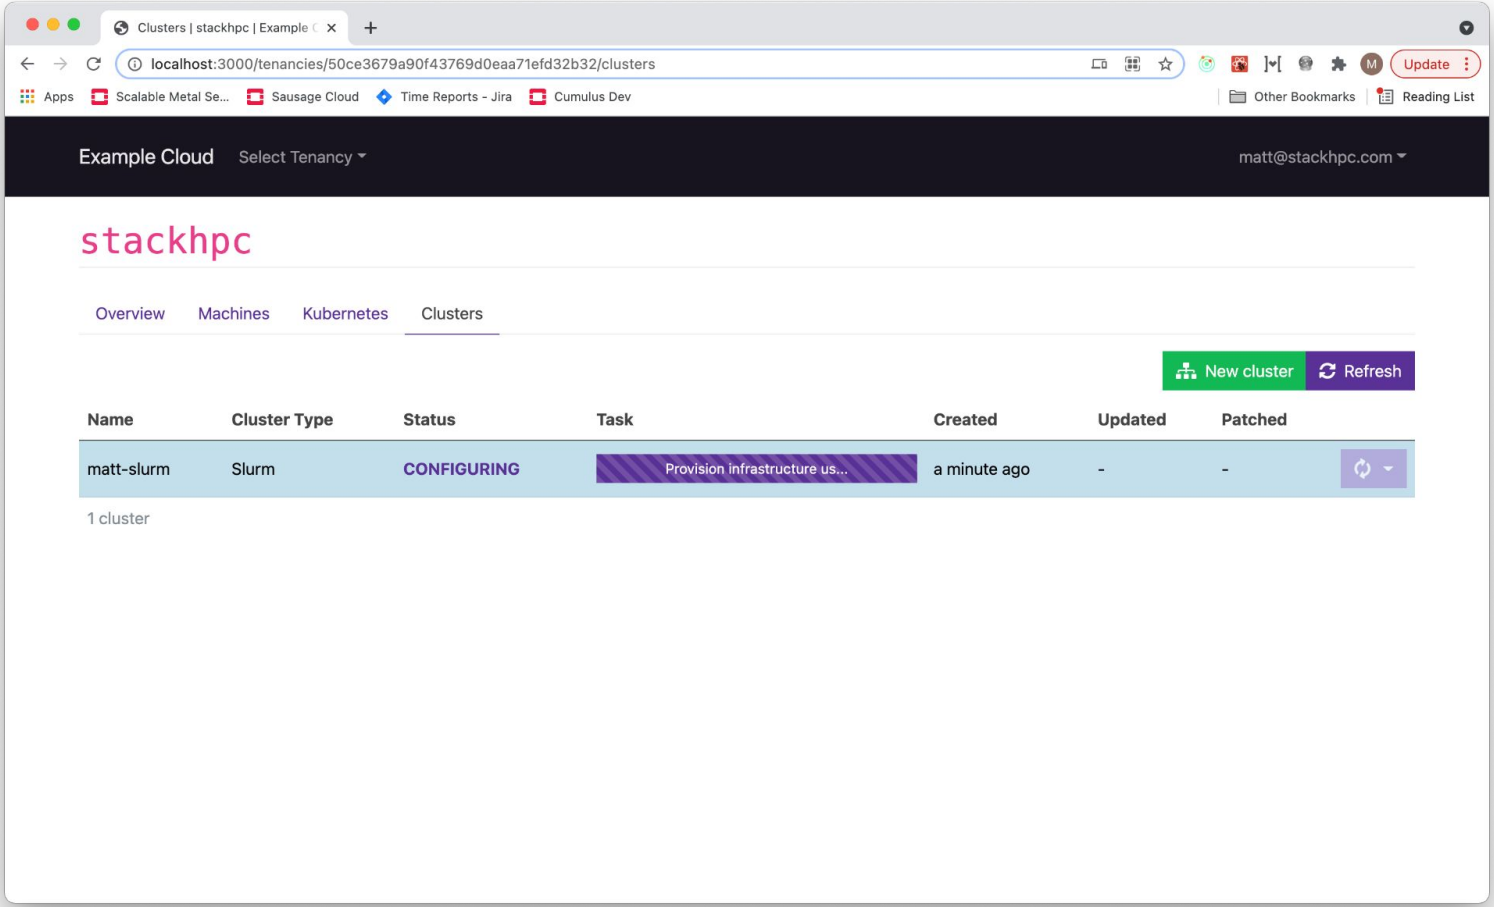
Clusters | stackhpc | Example
localhost:3000/tenancies/50ce3679a90f43769d0eaa71efd32b32/clusters
Update
Apps
Scalable Metal Se...
Sausage Cloud
Time Reports - Jira
Cumulus Dev
Other Bookmarks
Reading List
Example Cloud
Select Tenancy
matt@stackhpc.com
stackhpc
Overview
Machines
Kubernetes
Clusters
New cluster
Refresh
Name
Cluster Type
Status
Task
Created
Updated
Patched
matt-slurm
Slurm
CONFIGURING
Provision infrastructure us...
a minute ago
-
-
1 cluster

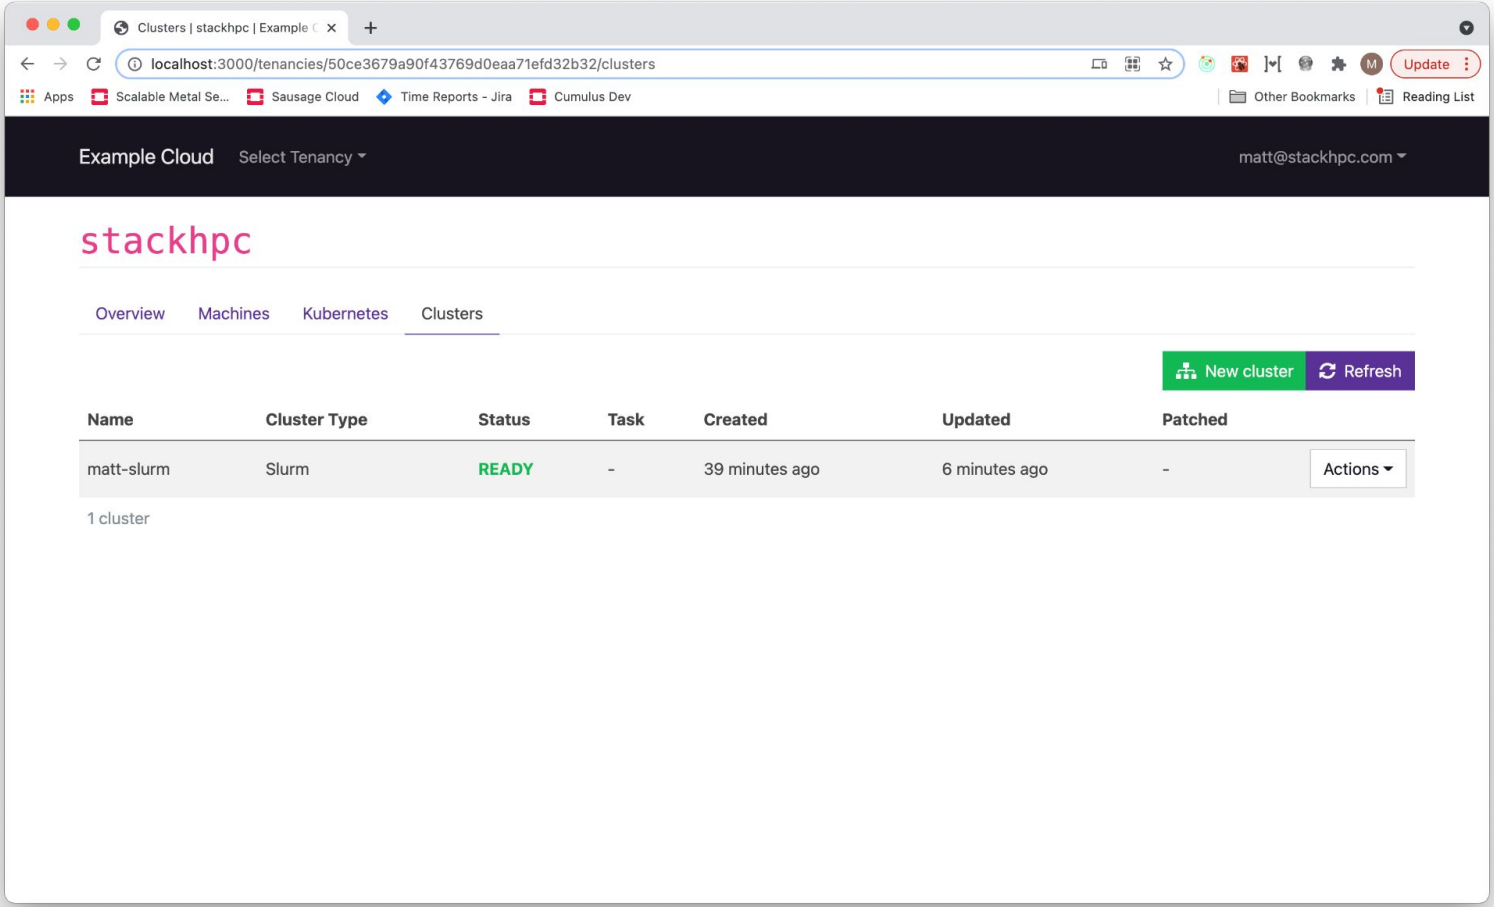
Clusters | stackhpc | Example
localhost:3000/tenancies/50ce3679a90f43769d0eaa71efd32b32/clusters
Update
Apps
Scalable Metal Se...
Sausage Cloud
Time Reports - Jira
Cumulus Dev
Other Bookmarks
Reading List
Example Cloud
Select Tenancy
matt@stackhpc.com
stackhpc
Overview
Machines
Kubernetes
Clusters
New cluster
Refresh
Name
Cluster Type
Status
Task
Created
Updated
Patched
matt-slurm
Slurm
READY
-
39 minutes ago
6 minutes ago
-
Actions
1 cluster

How did you create Slurm?

StackHPC

- Based on Slurm appliance from FY2020 IRIS Assets, rather than JASMIN
  - Moves to CentOS 8 and OpenHPC 2.0 (and associated slurm upgrades)
  - Includes monitoring that maps job id to VM metrics
- AWX runs ansible jobs on behalf of the Cloud Portal
- Terraform is used to create the infrastructure, storing state in Consul
- Ansible from the Slurm appliance configures the infrastructure
- (WIP) Guacamole and Proxy based access to the login node
- Lots of scope for exposing features of the appliance, such as the in place non-disruptive updates to compute nodes

What is next for Cloud Portal?

StackHPC

- More testing and **more operator and user feedback required!**
  - Work with other IRIS sites to trust a central cloud portal, similar to trusting their own horizon
- Making it easier to deploy and access platforms on Kubernetes
  - Test out new helm chart versions alongside existing versions
  - Try to integrate JupyterHub patterns from LSST digital asset
- More help with storage
  - sharing data sets (restricted and public)
  - into platforms, out from platforms
  - move between clouds
- More help with multiple clouds
  - Presents a common interface between multiple clouds
  - Non-OpenStack clouds could be supported using terraform approach
  - Bridge networks between sites using wireguard or openvpn, if needed

Questions?

StackHPC