OpenStack at the Sanger Institute - the first 18 months in production

Dave Holland

What I’ll talk about

- The Sanger Institute
- Motivations for using OpenStack
- Our journey
- Some decisions we made (and why)
- Some problems we encountered (and how we addressed them)
- Projects that are using OpenStack so far
- Next steps

Sanger Science

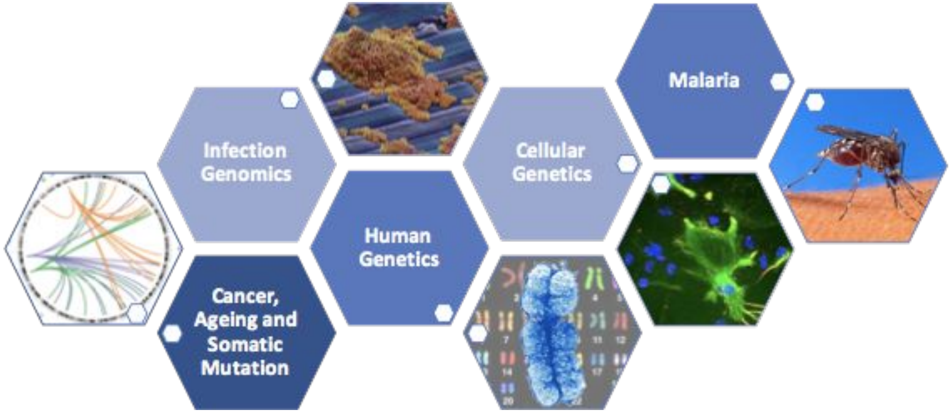

Infection Genomics

Cellular Genetics

Malaria

Human Genetics

Cancer, Ageing and Somatic Mutation

Scientific Research Programmes

Core Facilities

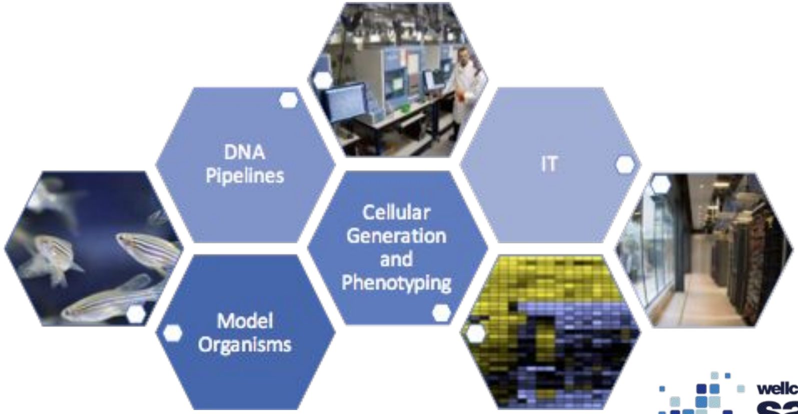

DNA Pipelines

IT

Cellular Generation and Phenotyping

Model Organisms

The Sanger Institute
Traditional HPC Environment

LSF 9

~10,000 cores in main compute farm

~10,000 cores across smaller project-specific farms

~15PB Lustre high-performance storage

Limited security - “isolation” is based on POSIX file permissions

Limited flexibility - no root access, incompatible software dependencies etc

Pipelines and stacks are complex, and scientific reproducibility is hard

HPC and cloud computing are complementary

Traditional HPC

- Highest possible performance
- A mature and centrally managed compute platform
- High-performance Lustre filesystems for data intensive analysis

Cloud compute

- Full segregation of projects ensures data security
- Developers no longer tied to a single stack
- Reproducibility through containers / images and infrastructure-as-code

But there’s a catch or two…

- Large number of traditional/legacy pipelines
    - They require a performant shared POSIX filesystem, while cloud workloads support object stores
- We do not always have the source code or expertise to migrate
- We need multi-gigabyte per second performance
- The tenant will have root
    - and could impersonate any user, but Lustre trusts the client’s identity assertions, just like NFSv3
- The solution must be simple for the tenant and administrator

The learning curve

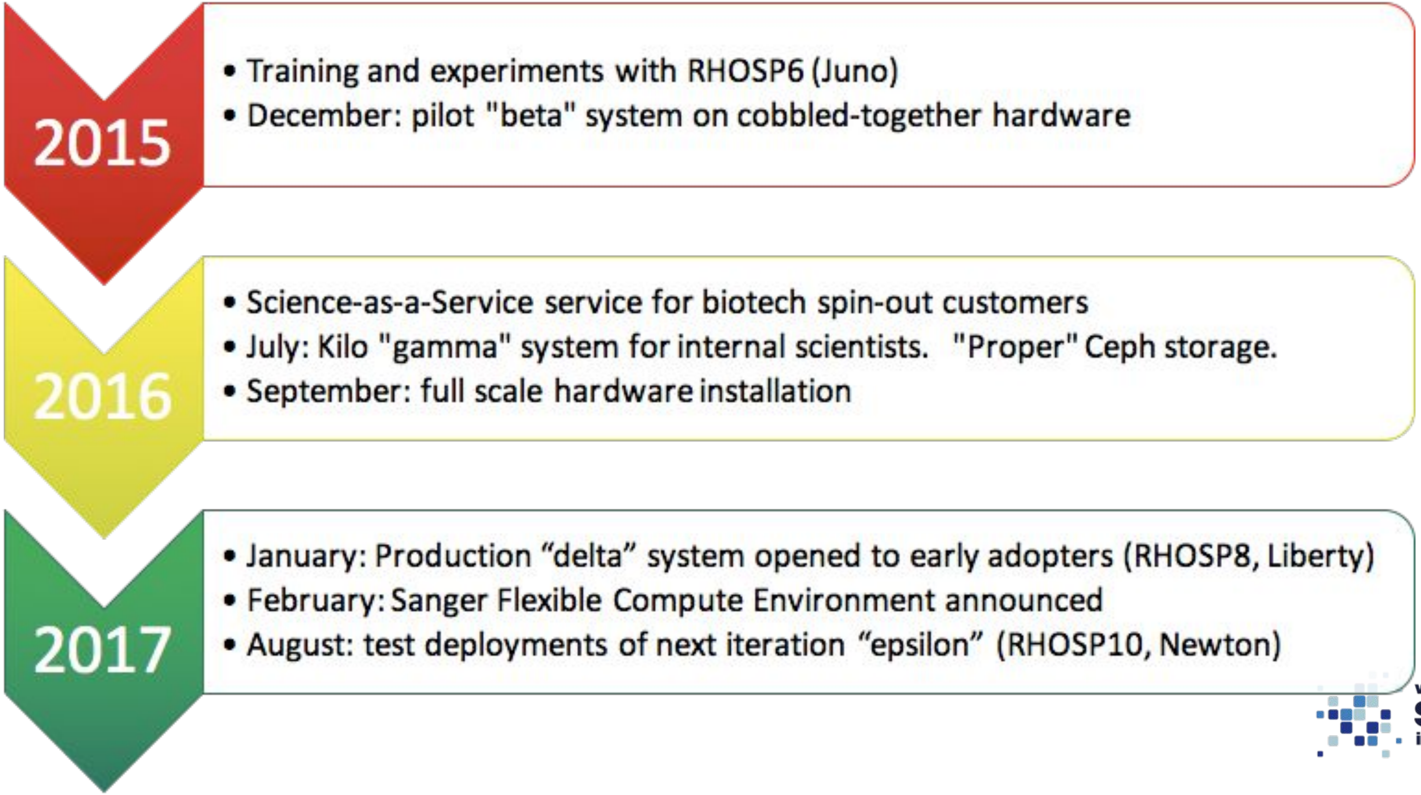

2015

- Training and experiments with RHOSP6 (Juno)
- December: pilot "beta" system on cobbled-together hardware

2016

- Science-as-a-Service service for biotech spin-out customers
- July: Kilo "gamma" system for internal scientists. "Proper" Ceph storage.
- September: full scale hardware installation

2017

- January: Production “delta” system opened to early adopters (RHOSP8, Liberty)
- February: Sanger Flexible Compute Environment announced
- August: test deployments of next iteration “epsilon” (RHOSP10, Newton)

Production hardware

- 107 Compute nodes (Supermicro) each with:
  - 512GB of RAM, 2x 25Gbit/s network interfaces,
  - 1x 960GB local SSD, 2x Intel E52690v4 (14 cores @ 2.6GHz)
- 6 Control nodes (Supermicro) allows 2 versions side by side
  - 256 GB RAM, 2x 100 Gbit/s network interfaces,
  - 1x 120 GB local SSD, 1x Intel P3600 NVMe (/var)
  - 2x Intel E52690v4 (14 cores @ 2.6GHz)
- Total of 53 TB of RAM, 2996 cores, 5992 with hyperthreading
- Red Hat OSP8 (“Liberty”) deployed with Triple-O

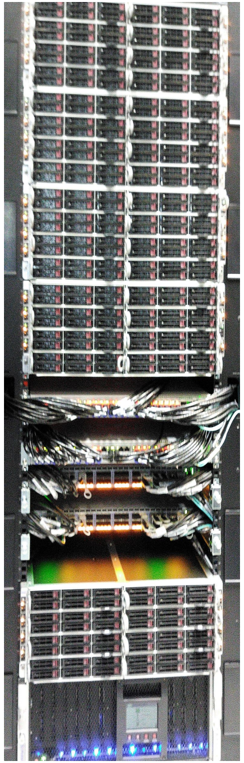

Ceph

- 9 storage nodes (Supermicro) each with:
  - 512GB of RAM, 2x Intel E52690v4 (14 cores @ 2.6GHz)
  - 2x 100Gbit/s network interfaces,
  - 60x 6TB SAS discs, 2 system SSD, 4TB of Intel P3600 NVMe used for journal.
- Ubuntu Xenial, Ceph “Jewel”
- 3PB of disc space, 1PB usable.
- Single node (1.3 GBytes/sec write, 200 MBytes/sec read)
- Ceph benchmarks imply 7 GBytes/second.
- Rebuild traffic of 20 GBytes/second.

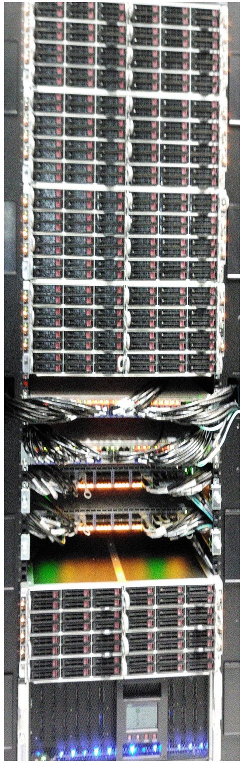

Networking

- 10 Arista 7060CX-32S switches.
  - 1U, 32 * 100Gb/s -> 128 * 25Gb/s
  - Hardware VXLAN support integrated with OpenStack*
  - Layer two traffic limited to rack, VXLAN used inter-rack.
  - Layer three between racks and interconnect to legacy systems.
  - All network switch software can be upgraded without disruption.
  - 400 Gb/s from racks to spine, 160 Gb/s from spine to legacy systems.

(* VXLAN in ml2 plugin not used in first iteration because of software issues)

- 3 racks of equipment, 24 kW load per rack.

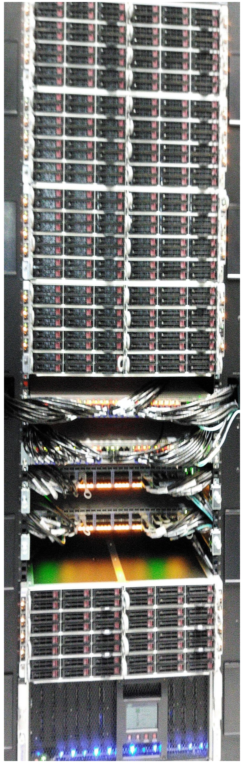

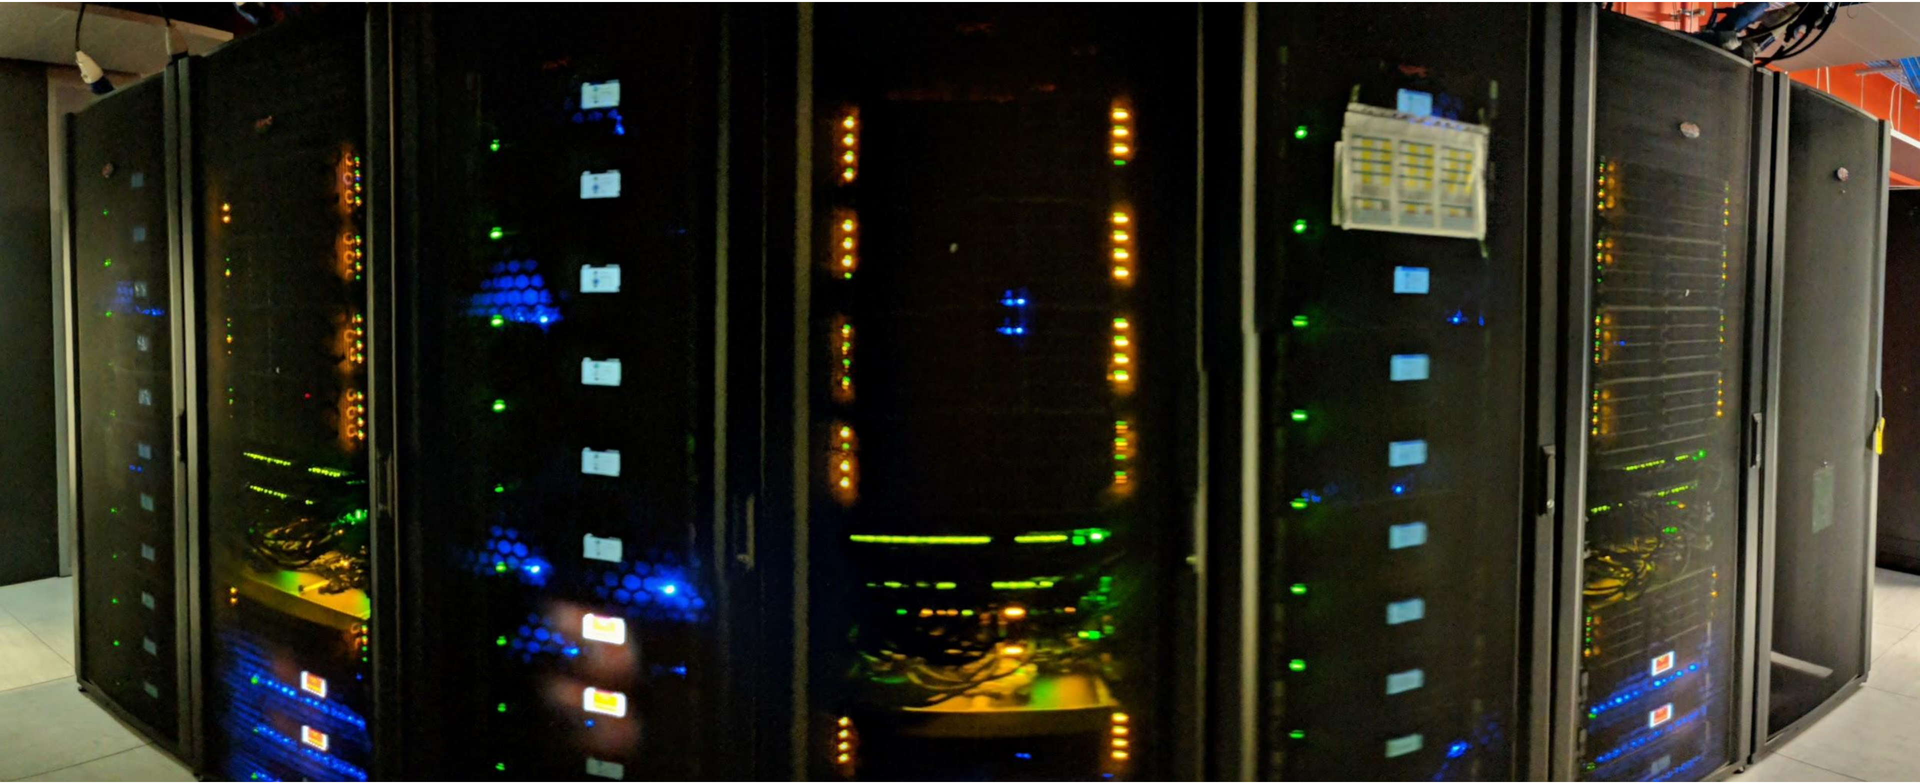

Technologies

Liberty (RHOSP8) - “Delta”

- default file descriptor limits for Glance, Cinder, RabbitMQ etc way too low
- NIC hardware acceleration works fine with the correct kernel/driver (thanks to Mellanox)
- races when scheduling many instances; mitigate by directing nova_osapi to a “primary” server
- faulty fibres caused NIC flaps and RabbitMQ problems

RabbitMQ

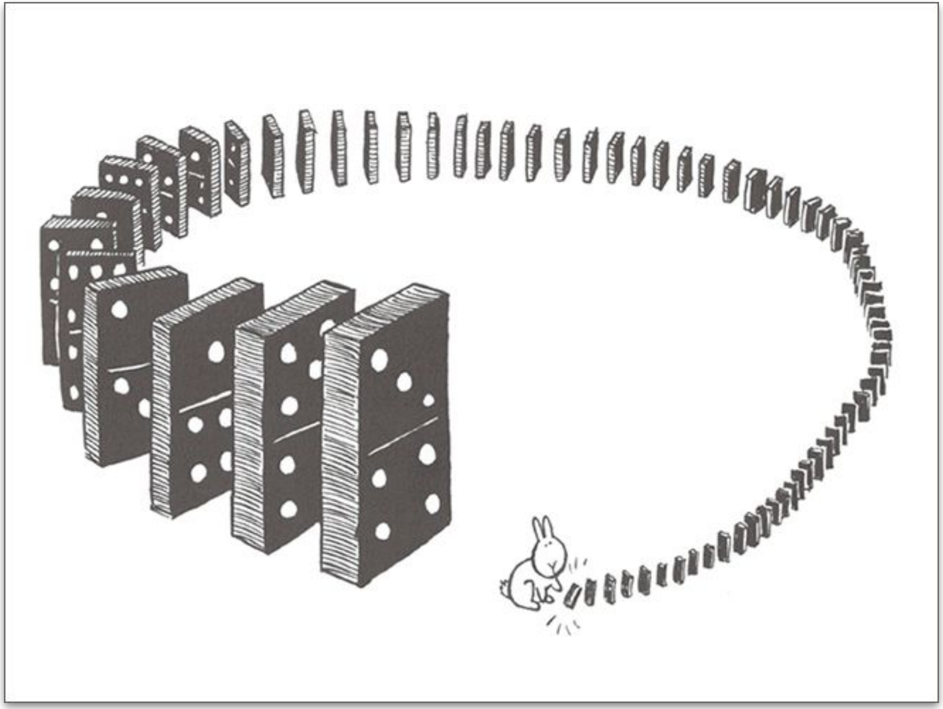

Andy Riley's "Bunny Suicides"

Liberty (RHOSP8) - “Delta”

- not enough hypervisor memory reserved, neutron-openvswitch-agent fails to allocate memory
- various package upgrades (Glance, Neutron, Dnsmasq etc) for particular bugs encountered or security fixes
- work with Arista on portchannel/LACP instability

Ceph

- standalone Ceph to support multiple OpenStack deployments
- generally robust in the face of spontaneous machine hangs and rack failures
- OSD start-up race due to single shared lock
- niggles with ceph-ansible

Ansible for customisation

- scheduler tweaks (stack not spread, CPU/RAM overcommit)
- hypervisor tweaks (instance root disk on Ceph or hypervisor)
- enable SSL for Horizon and API
- change syslog destination
- add “MOTD” to Horizon login page
- change session timeouts
- register systems with RedHat
- and more…

- but deployer’s Puppet overwrites some of these

Monitoring

- evolving from “bare minimum” to customer and engineer views, scorecard/availability report
- custom Nagios scripts, active and passive checks
- Grafana for metrics
- rsyslog and ELK for logging

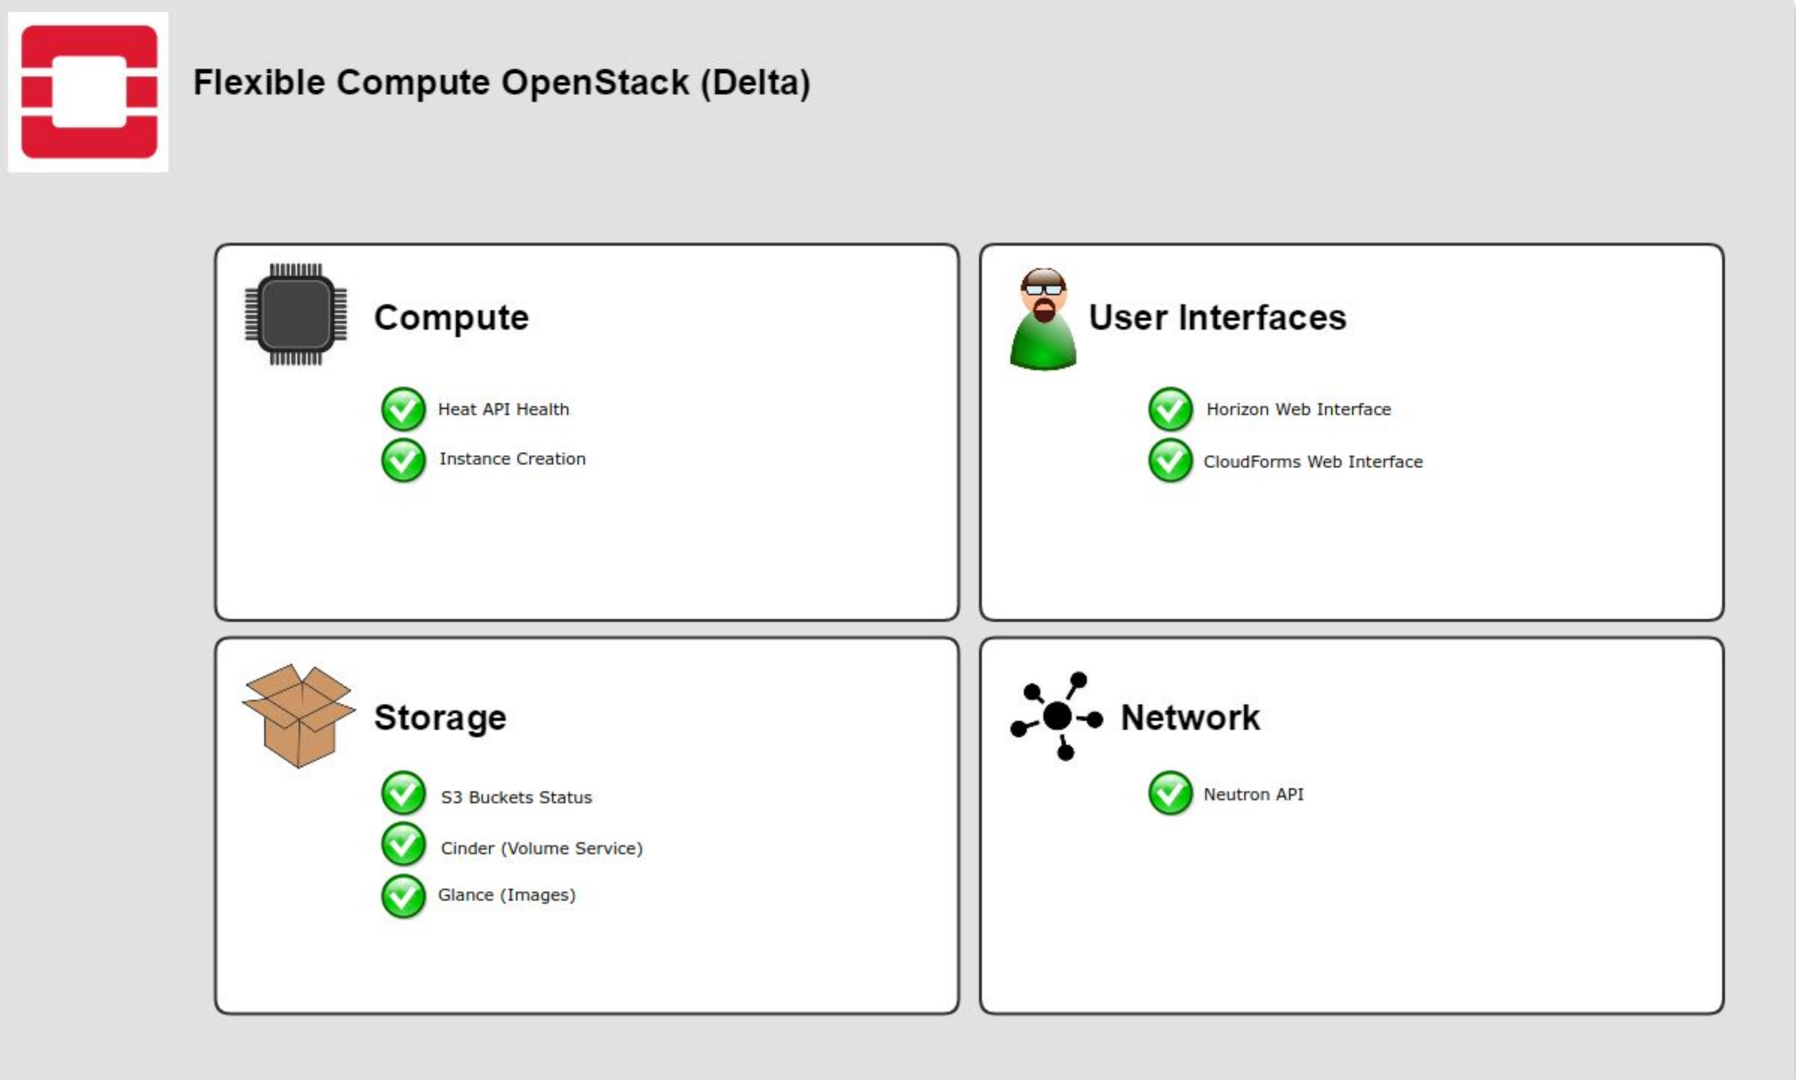

Flexible Compute OpenStack (Delta)

Compute

- ✓ Heat API Health
- ✓ Instance Creation

User Interfaces

- ✓ Horizon Web Interface
- ✓ CloudForms Web Interface

Storage

- ✓ S3 Buckets Status
- ✓ Cinder (Volume Service)
- ✓ Glance (Images)

Network

- ✓ Neutron API

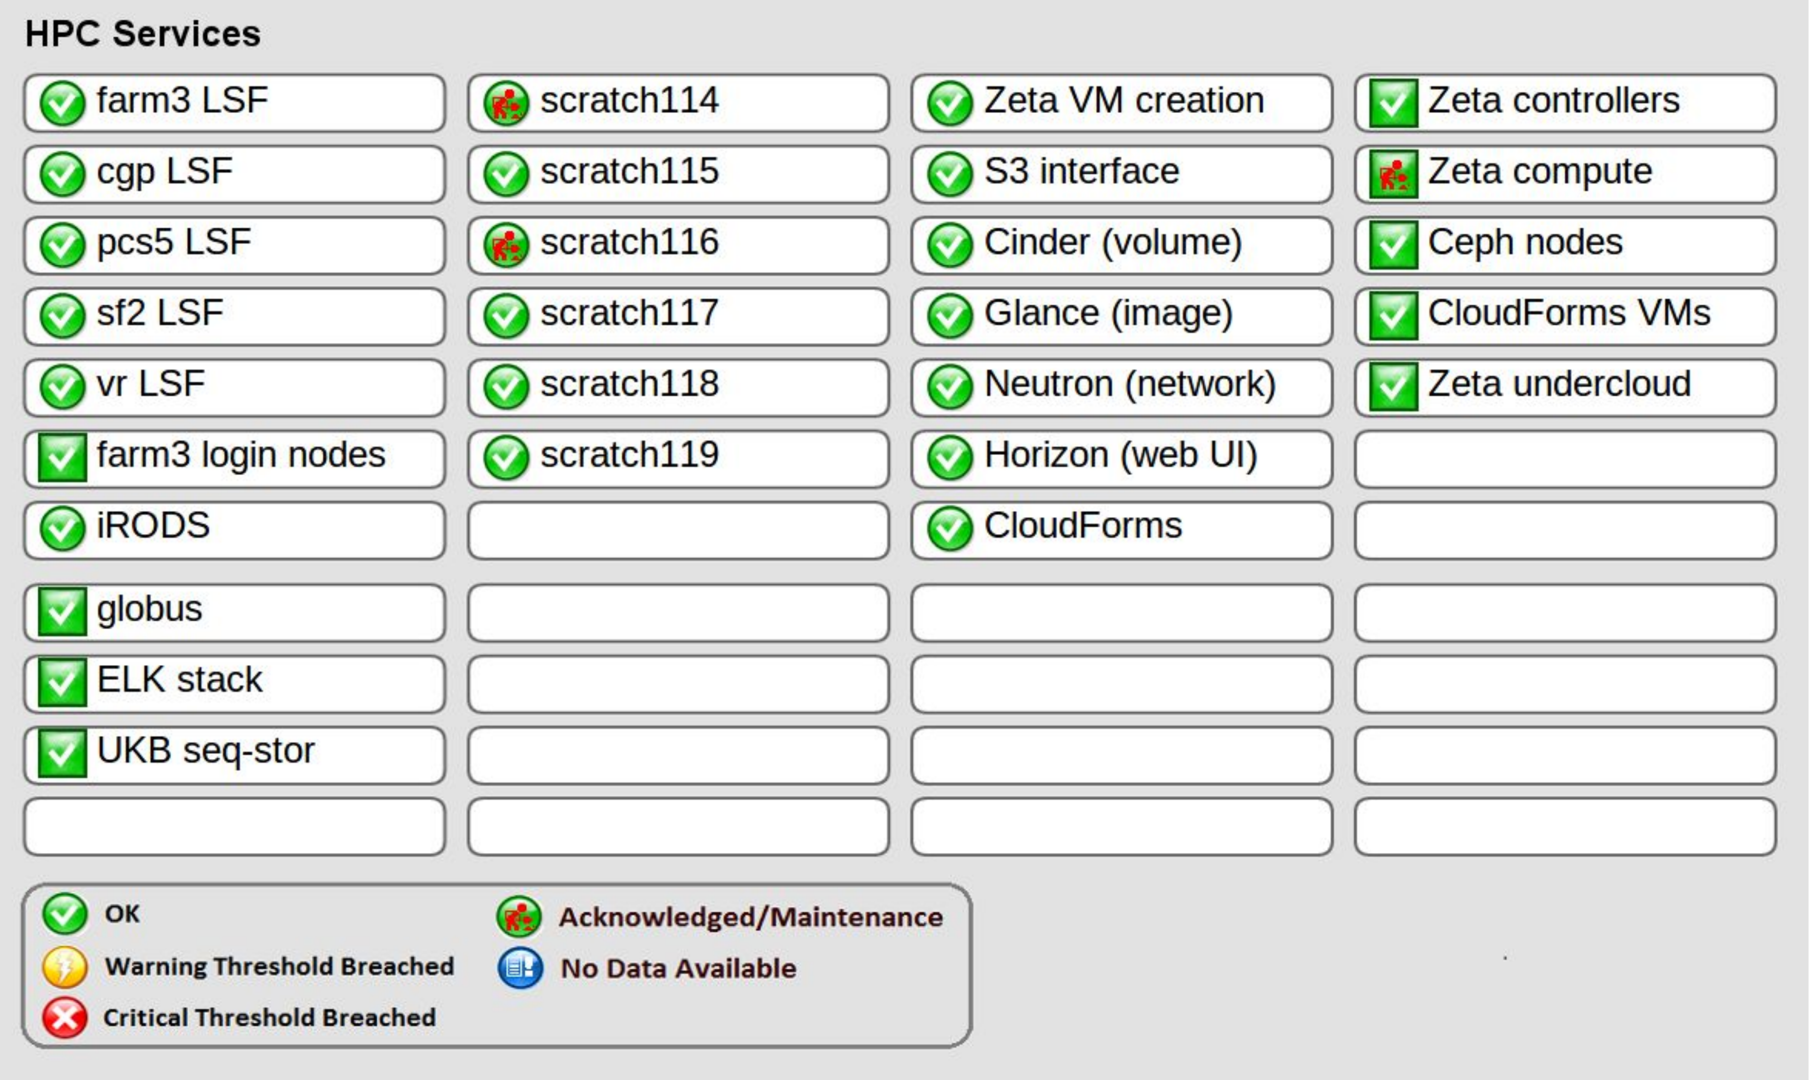

HPC Services

✓ farm3 LSF	🔧 scratch114	✓ Zeta VM creation	☑ Zeta controllers
✓ cgp LSF	✓ scratch115	✓ S3 interface	🔧 Zeta compute
✓ pcs5 LSF	🔧 scratch116	✓ Cinder (volume)	☑ Ceph nodes
✓ sf2 LSF	✓ scratch117	✓ Glance (image)	☑ CloudForms VMs
✓ vr LSF	✓ scratch118	✓ Neutron (network)	☑ Zeta undercloud
☑ farm3 login nodes	✓ scratch119	✓ Horizon (web UI)	
✓ iRODS		✓ CloudForms	
☑ globus			
☑ ELK stack			
☑ UKB seq-stor			

✓ OK	🔧 Acknowledged/Maintenance
⚡ Warning Threshold Breached	ℹ No Data Available
✖ Critical Threshold Breached	

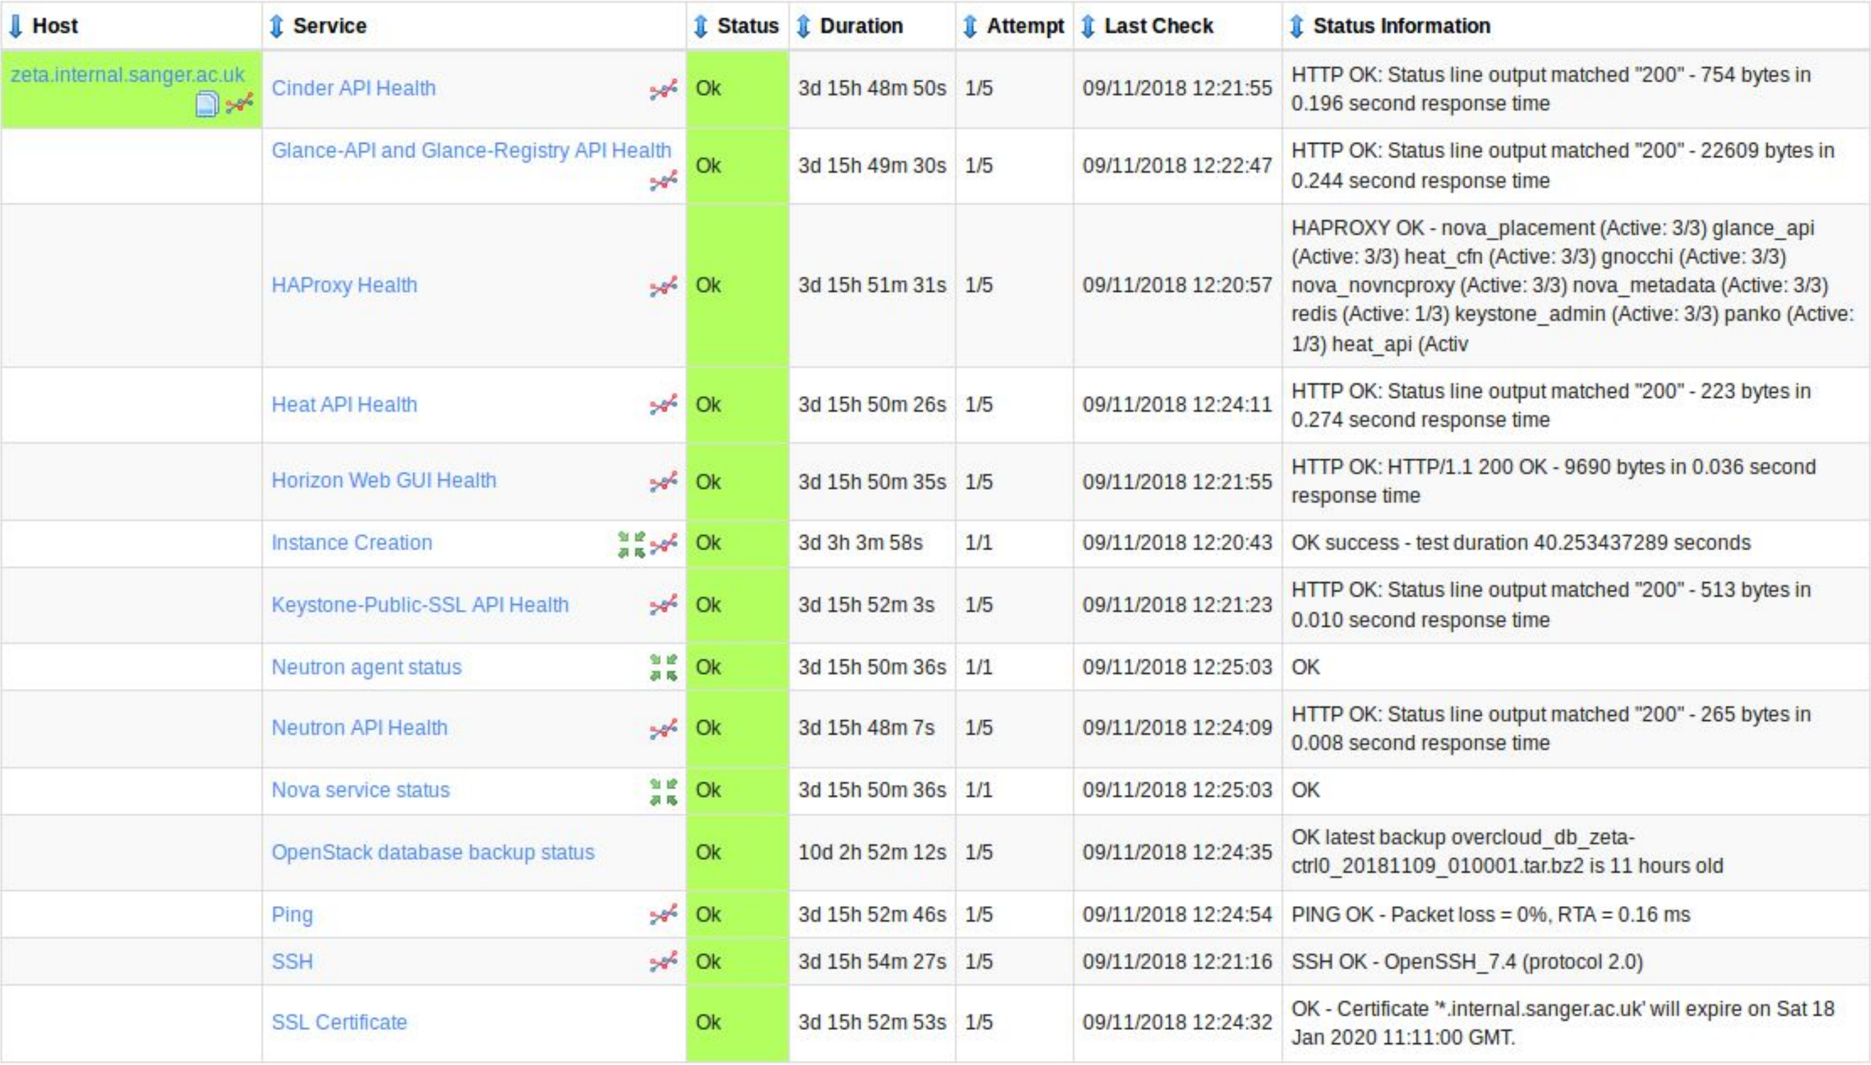

Host	Service	Status	Duration	Attempt	Last Check	Status Information
zeta.internal.sanger.ac.uk	Cinder API Health	Ok	3d 15h 48m 50s	1/5	09/11/2018 12:21:55	HTTP OK: Status line output matched "200" - 754 bytes in 0.196 second response time
	Glance-API and Glance-Registry API Health	Ok	3d 15h 49m 30s	1/5	09/11/2018 12:22:47	HTTP OK: Status line output matched "200" - 22609 bytes in 0.244 second response time
	HAProxy Health	Ok	3d 15h 51m 31s	1/5	09/11/2018 12:20:57	HAPROXY OK - nova_placement (Active: 3/3) glance_api (Active: 3/3) heat_cfn (Active: 3/3) gnocchi (Active: 3/3) nova_novncproxy (Active: 3/3) nova_metadata (Active: 3/3) redis (Active: 1/3) keystone_admin (Active: 3/3) panko (Active: 1/3) heat_api (Activ
	Heat API Health	Ok	3d 15h 50m 26s	1/5	09/11/2018 12:24:11	HTTP OK: Status line output matched "200" - 223 bytes in 0.274 second response time
	Horizon Web GUI Health	Ok	3d 15h 50m 35s	1/5	09/11/2018 12:21:55	HTTP OK: HTTP/1.1 200 OK - 9690 bytes in 0.036 second response time
	Instance Creation	Ok	3d 3h 3m 58s	1/1	09/11/2018 12:20:43	OK success - test duration 40.253437289 seconds
	Keystone-Public-SSL API Health	Ok	3d 15h 52m 3s	1/5	09/11/2018 12:21:23	HTTP OK: Status line output matched "200" - 513 bytes in 0.010 second response time
	Neutron agent status	Ok	3d 15h 50m 36s	1/1	09/11/2018 12:25:03	OK
	Neutron API Health	Ok	3d 15h 48m 7s	1/5	09/11/2018 12:24:09	HTTP OK: Status line output matched "200" - 265 bytes in 0.008 second response time
	Nova service status	Ok	3d 15h 50m 36s	1/1	09/11/2018 12:25:03	OK
	OpenStack database backup status	Ok	10d 2h 52m 12s	1/5	09/11/2018 12:24:35	OK latest backup overcloud_db_zeta-ctrl0_20181109_010001.tar.bz2 is 11 hours old
	Ping	Ok	3d 15h 52m 46s	1/5	09/11/2018 12:24:54	PING OK - Packet loss = 0%, RTA = 0.16 ms
	SSH	Ok	3d 15h 54m 27s	1/5	09/11/2018 12:21:16	SSH OK - OpenSSH_7.4 (protocol 2.0)
	SSL Certificate	Ok	3d 15h 52m 53s	1/5	09/11/2018 12:24:32	OK - Certificate '*.internal.sanger.ac.uk' will expire on Sat 18 Jan 2020 11:11:00 GMT.

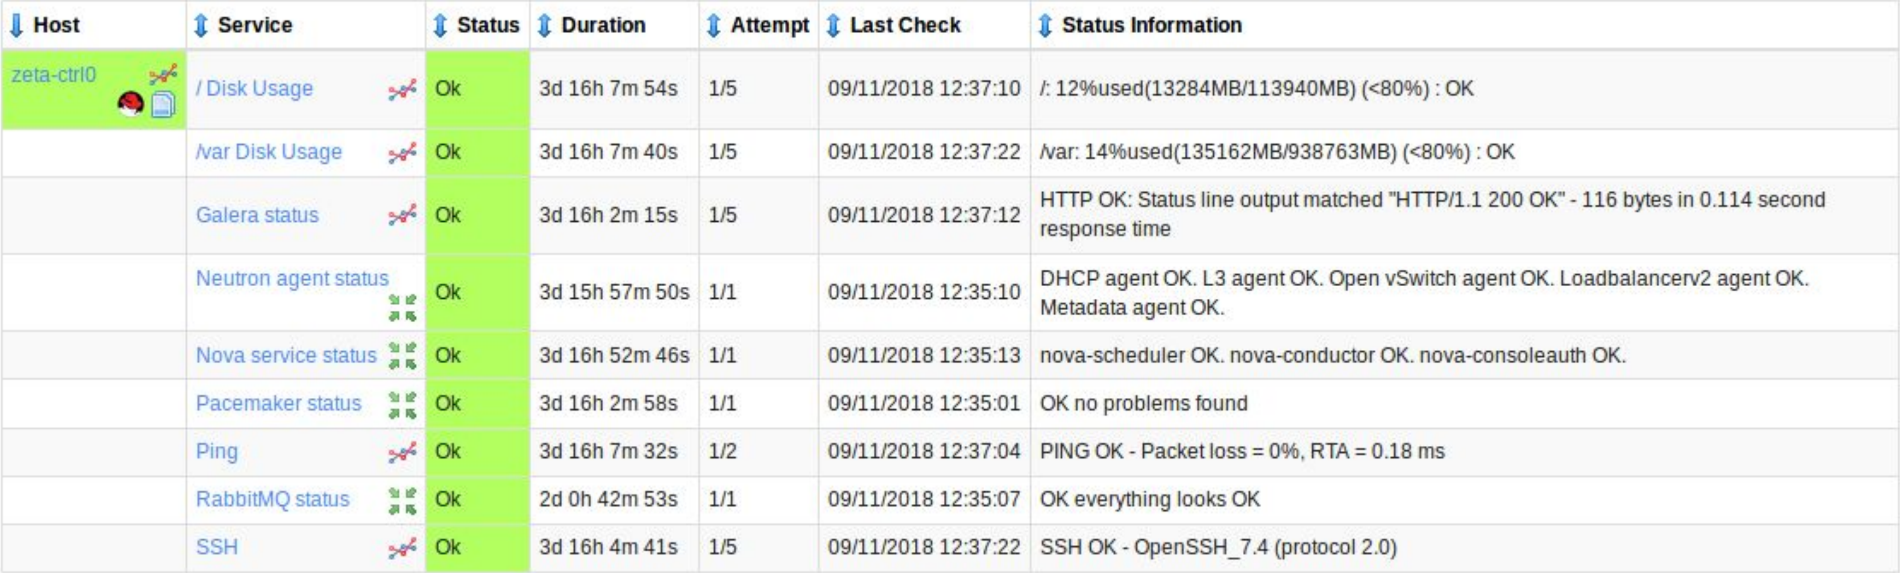

Host	Service	Status	Duration	Attempt	Last Check	Status Information
zeta-ctrl0	/ Disk Usage	Ok	3d 16h 7m 54s	1/5	09/11/2018 12:37:10	/: 12%used(13284MB/113940MB) (<80%) : OK
	/var Disk Usage	Ok	3d 16h 7m 40s	1/5	09/11/2018 12:37:22	/var: 14%used(135162MB/938763MB) (<80%) : OK
	Galera status	Ok	3d 16h 2m 15s	1/5	09/11/2018 12:37:12	HTTP OK: Status line output matched "HTTP/1.1 200 OK" - 116 bytes in 0.114 second response time
	Neutron agent status	Ok	3d 15h 57m 50s	1/1	09/11/2018 12:35:10	DHCP agent OK. L3 agent OK. Open vSwitch agent OK. Loadbalancerv2 agent OK. Metadata agent OK.
	Nova service status	Ok	3d 16h 52m 46s	1/1	09/11/2018 12:35:13	nova-scheduler OK. nova-conductor OK. nova-consoleauth OK.
	Pacemaker status	Ok	3d 16h 2m 58s	1/1	09/11/2018 12:35:01	OK no problems found
	Ping	Ok	3d 16h 7m 32s	1/2	09/11/2018 12:37:04	PING OK - Packet loss = 0%, RTA = 0.18 ms
	RabbitMQ status	Ok	2d 0h 42m 53s	1/1	09/11/2018 12:35:07	OK everything looks OK
	SSH	Ok	3d 16h 4m 41s	1/5	09/11/2018 12:37:22	SSH OK - OpenSSH_7.4 (protocol 2.0)

Metrics

- collectd + libvirtd + Python graphitesend
- custom scripts to aggregate CPU use for efficiency “score”
- need to scale up number of carbon-cache processes
- Grafana dashboards

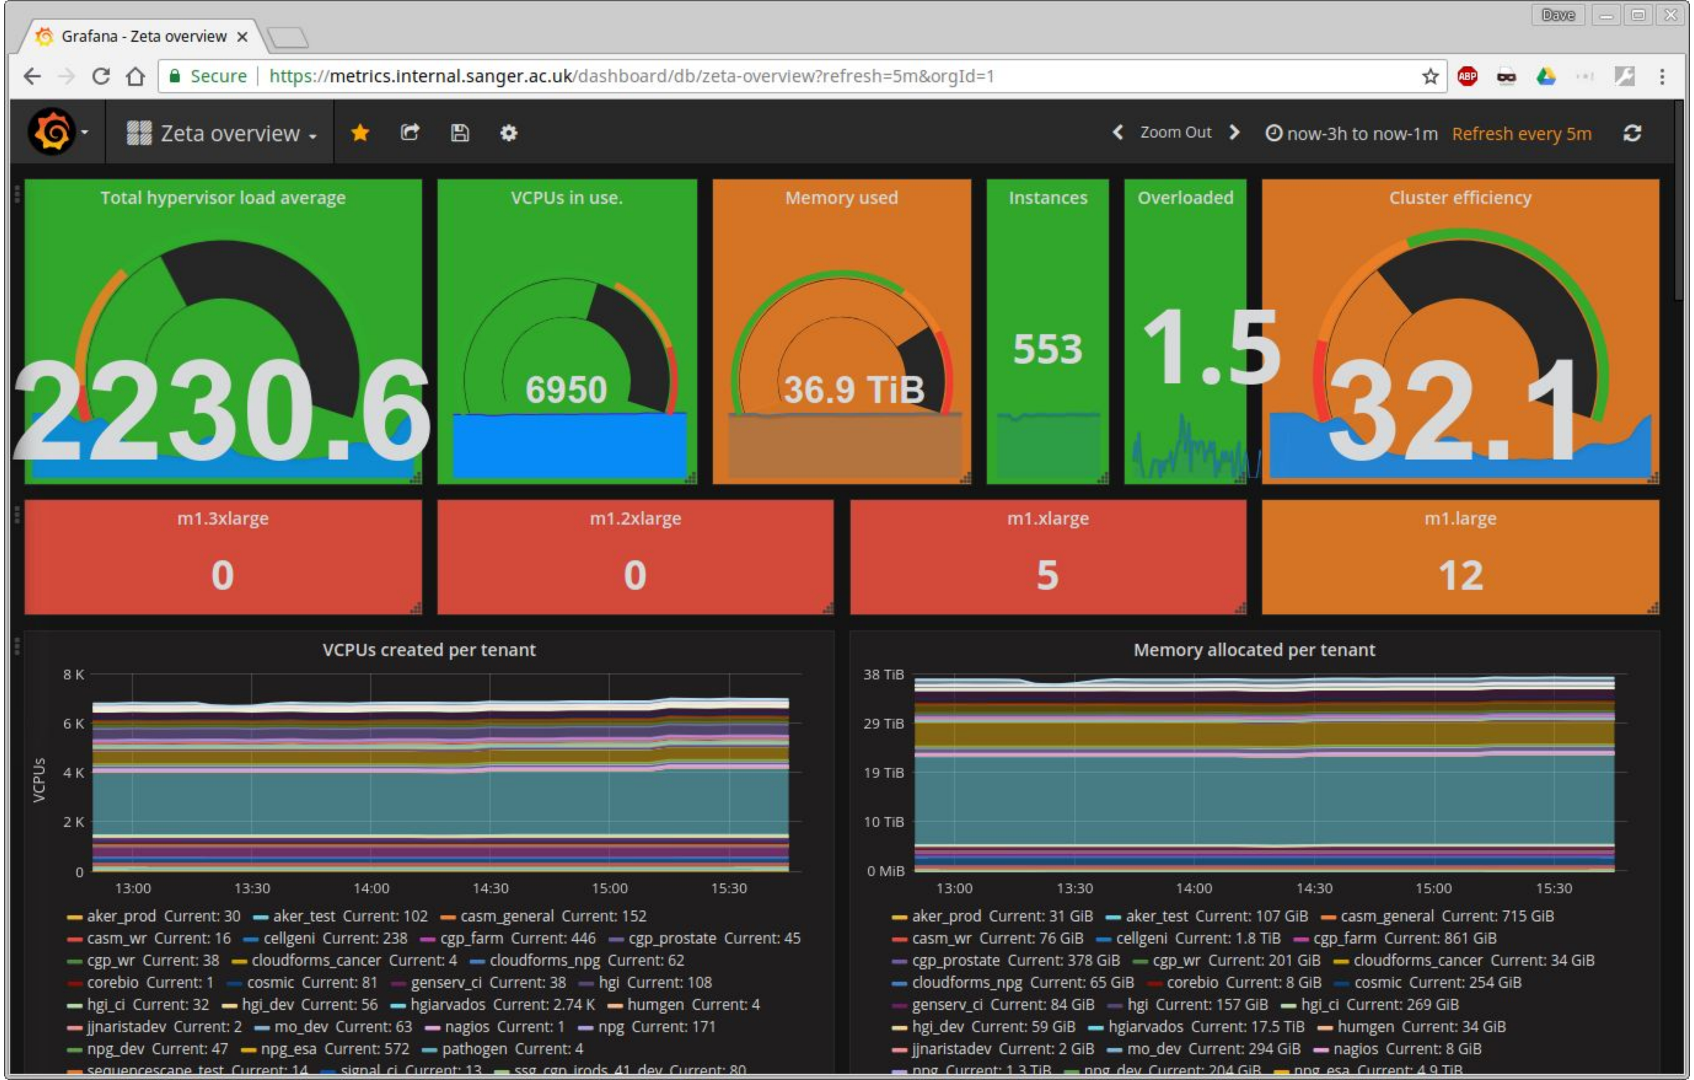

Total hypervisor load average 2230.6

VCPUs in use. 6950

Memory used 36.9 TiB

Instances 553

Overloaded 1.5

Cluster efficiency 32.1

m1.3xlarge 0

m1.2xlarge 0

m1.xlarge 5

m1.large 12

VCPUs created per tenant

aker_prod Current: 30 aker_test Current: 102 casm_general Current: 152
casm_wr Current: 16 cellgeni Current: 238 cgp_farm Current: 446 cgp_prostate Current: 45
cgp_wr Current: 38 cloudforms_cancer Current: 4 cloudforms_npg Current: 62
corebio Current: 1 cosmic Current: 81 genserv_ci Current: 38 hgi Current: 108
hgi_ci Current: 32 hgi_dev Current: 56 hgiarvados Current: 2.74 K humgen Current: 4
jjnaristadev Current: 2 mo_dev Current: 63 nagios Current: 1 npg Current: 171
npg_dev Current: 47 npg_esa Current: 572 pathogen Current: 4
sequencescape_test Current: 14 signal_ci Current: 13 ssg_cgp_irods_41_dev Current: 80

Memory allocated per tenant

aker_prod Current: 31 GiB aker_test Current: 107 GiB casm_general Current: 715 GiB
casm_wr Current: 76 GiB cellgeni Current: 1.8 TiB cgp_farm Current: 861 GiB
cgp_prostate Current: 378 GiB cgp_wr Current: 201 GiB cloudforms_cancer Current: 34 GiB
cloudforms_npg Current: 65 GiB corebio Current: 8 GiB cosmic Current: 254 GiB
genserv_ci Current: 84 GiB hgi Current: 157 GiB hgi_ci Current: 269 GiB
hgi_dev Current: 59 GiB hgiarvados Current: 17.5 TiB humgen Current: 34 GiB
jjnaristadev Current: 2 GiB mo_dev Current: 294 GiB nagios Current: 8 GiB
npg Current: 1.3 TiB npg_dev Current: 204 GiB npg_esa Current: 4.9 TiB

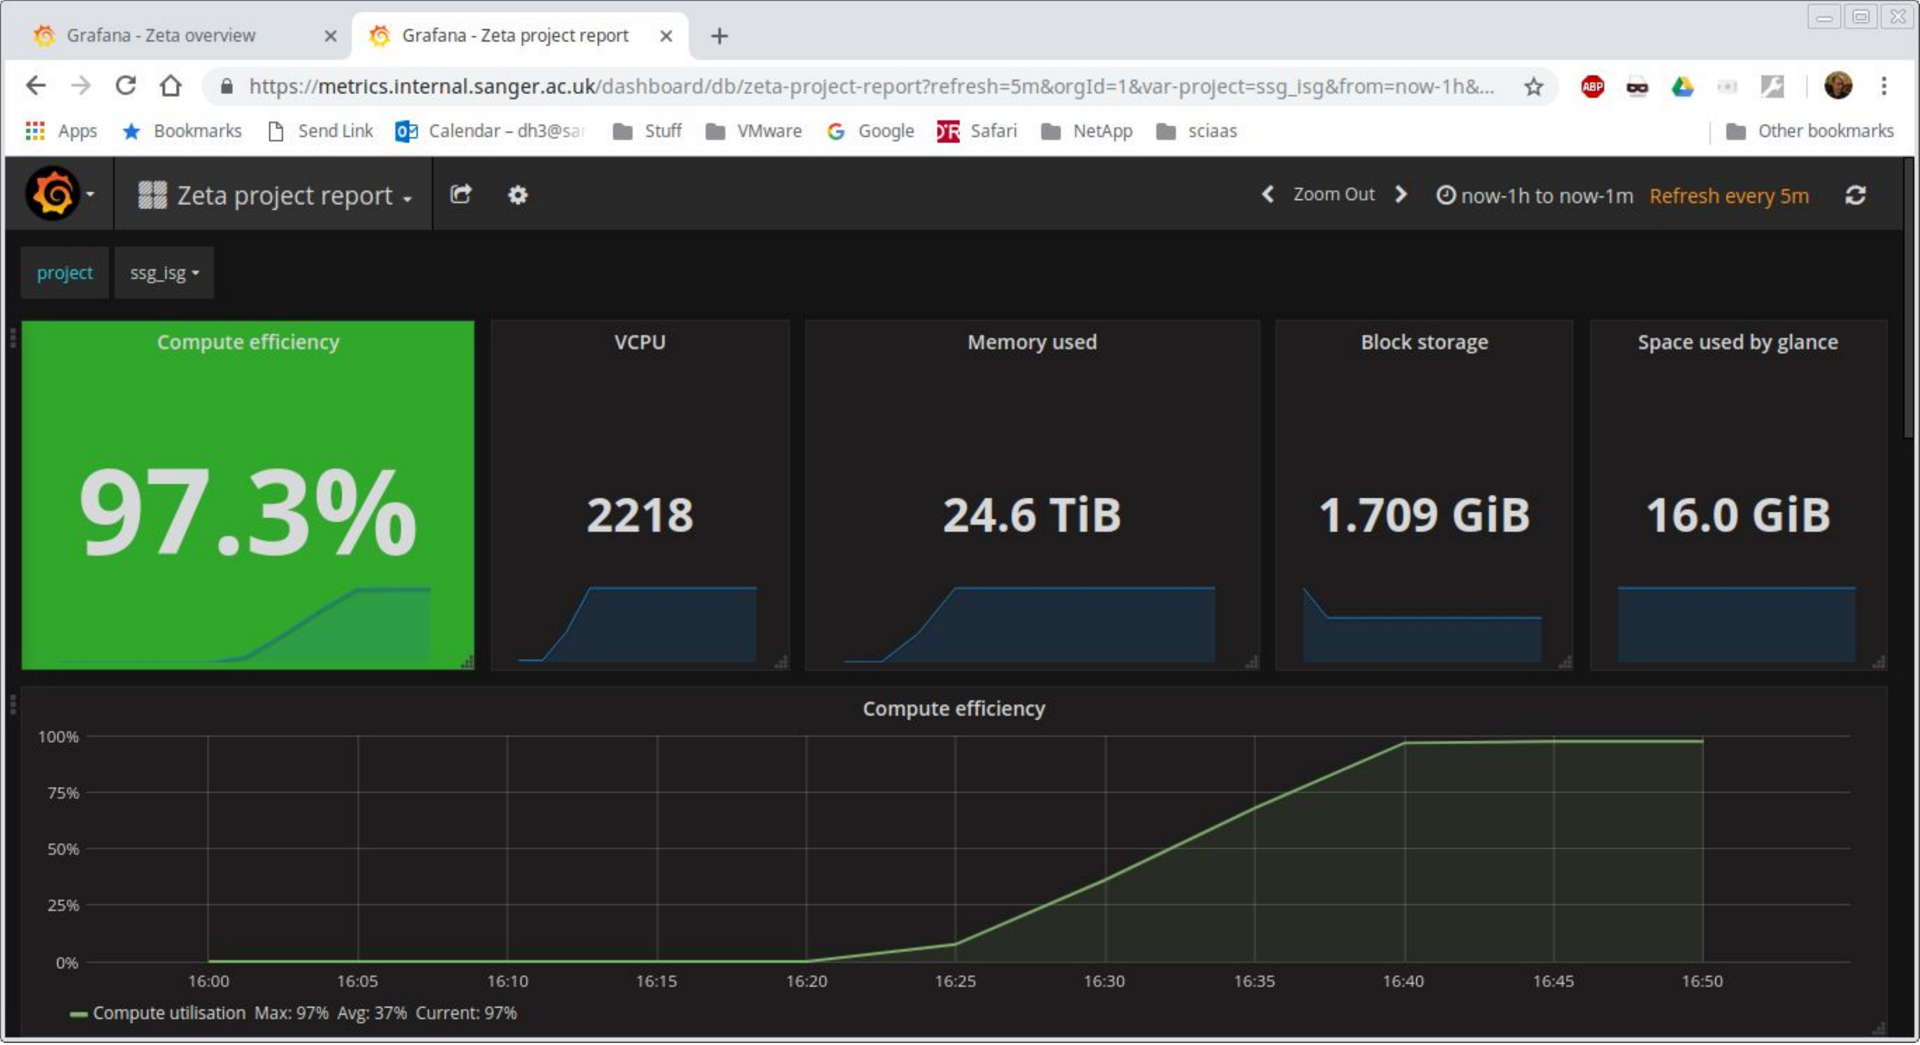

https://metrics.internal.sanger.ac.uk/dashboard/db/zeta-project-report?refresh=5m&orgId=1&var-project=ssg_isg&from=now-1h&...

project ssg_isg

now-1h to now-1m Refresh every 5m

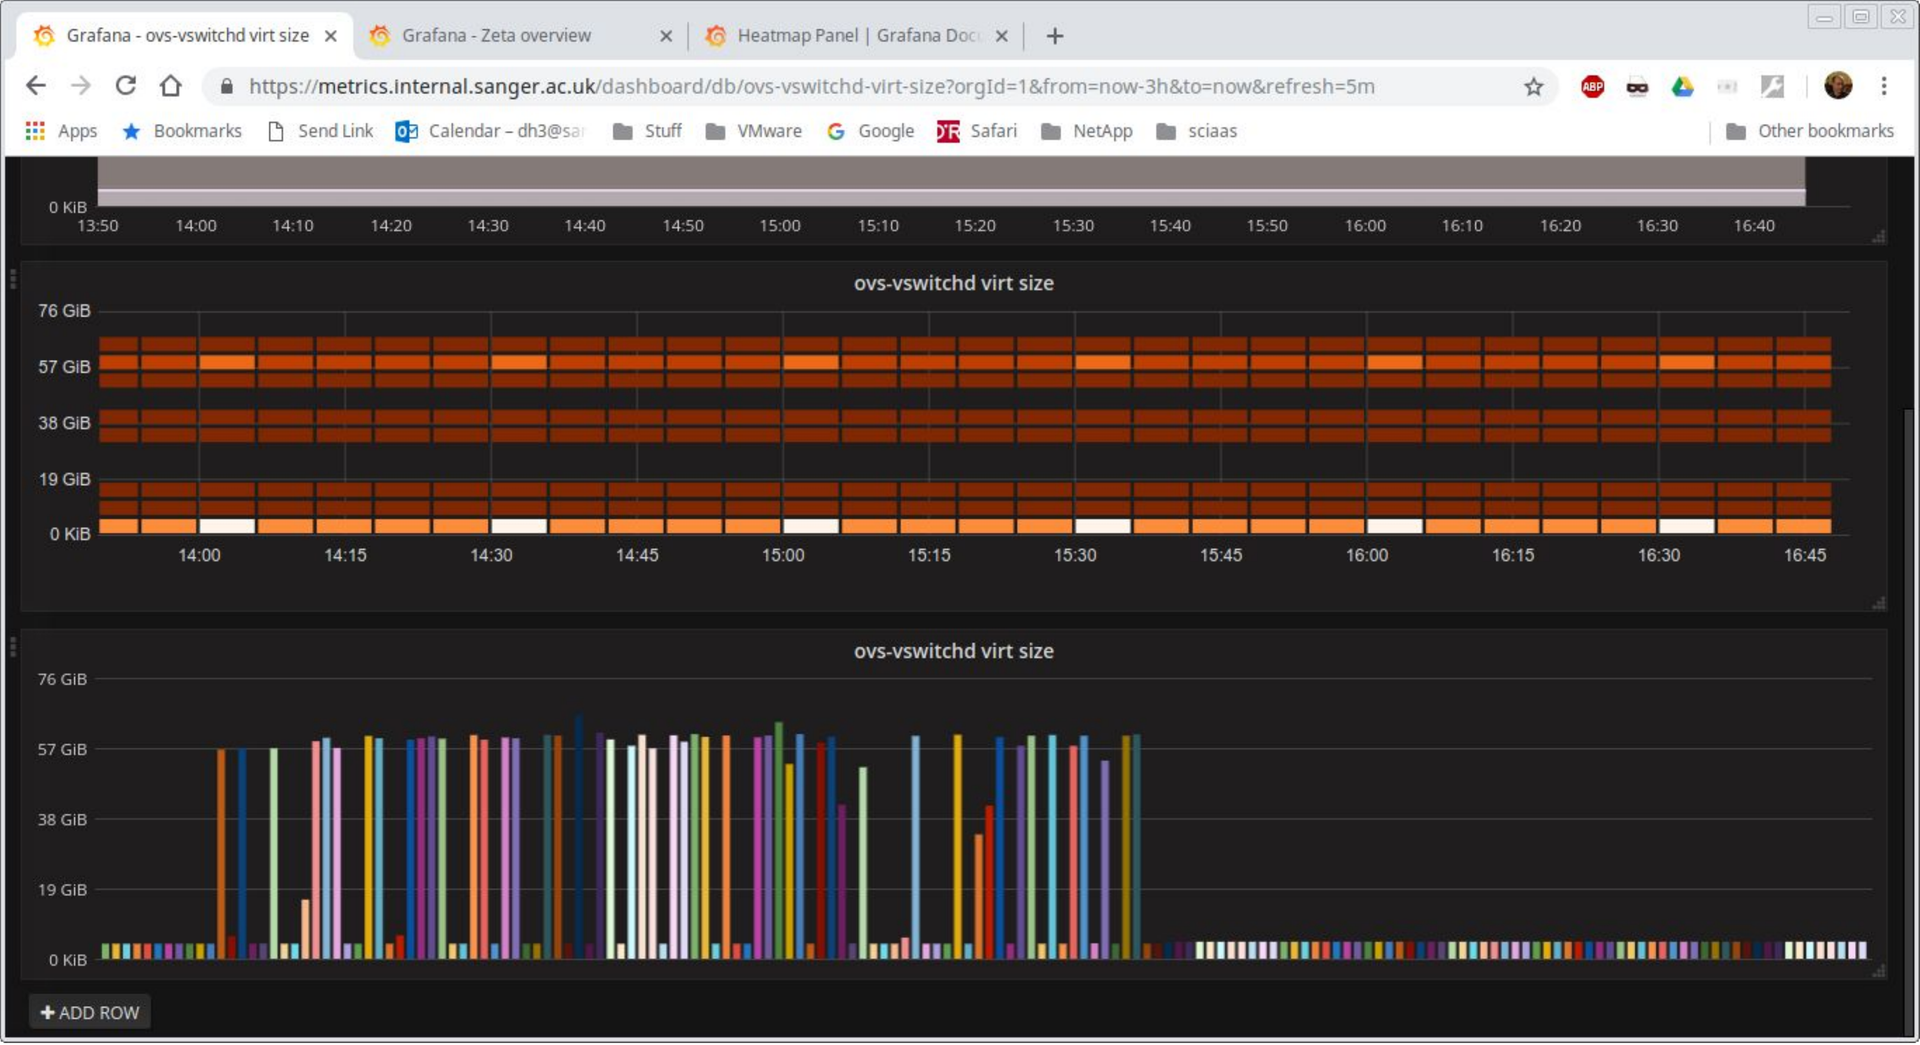
Grafana - ovs-vswitchd virt size
Grafana - Zeta overview
Heatmap Panel | Grafana Doc
https://metrics.internal.sanger.ac.uk/dashboard/db/ovs-vswitchd-virt-size?orgId=1&from=now-3h&to=now&refresh=5m
Apps
Bookmarks
Send Link
Calendar – dh3@sa
Stuff
VMware
Google
Safari
NetApp
sciaas
Other bookmarks
0 KiB
13:50
14:00
14:10
14:20
14:30
14:40
14:50
15:00
15:10
15:20
15:30
15:40
15:50
16:00
16:10
16:20
16:30
16:40
ovs-vswitchd virt size
76 GiB
57 GiB
38 GiB
19 GiB
0 KiB
14:00
14:15
14:30
14:45
15:00
15:15
15:30
15:45
16:00
16:15
16:30
16:45
ovs-vswitchd virt size
76 GiB
57 GiB
38 GiB
19 GiB
0 KiB
+ ADD ROW

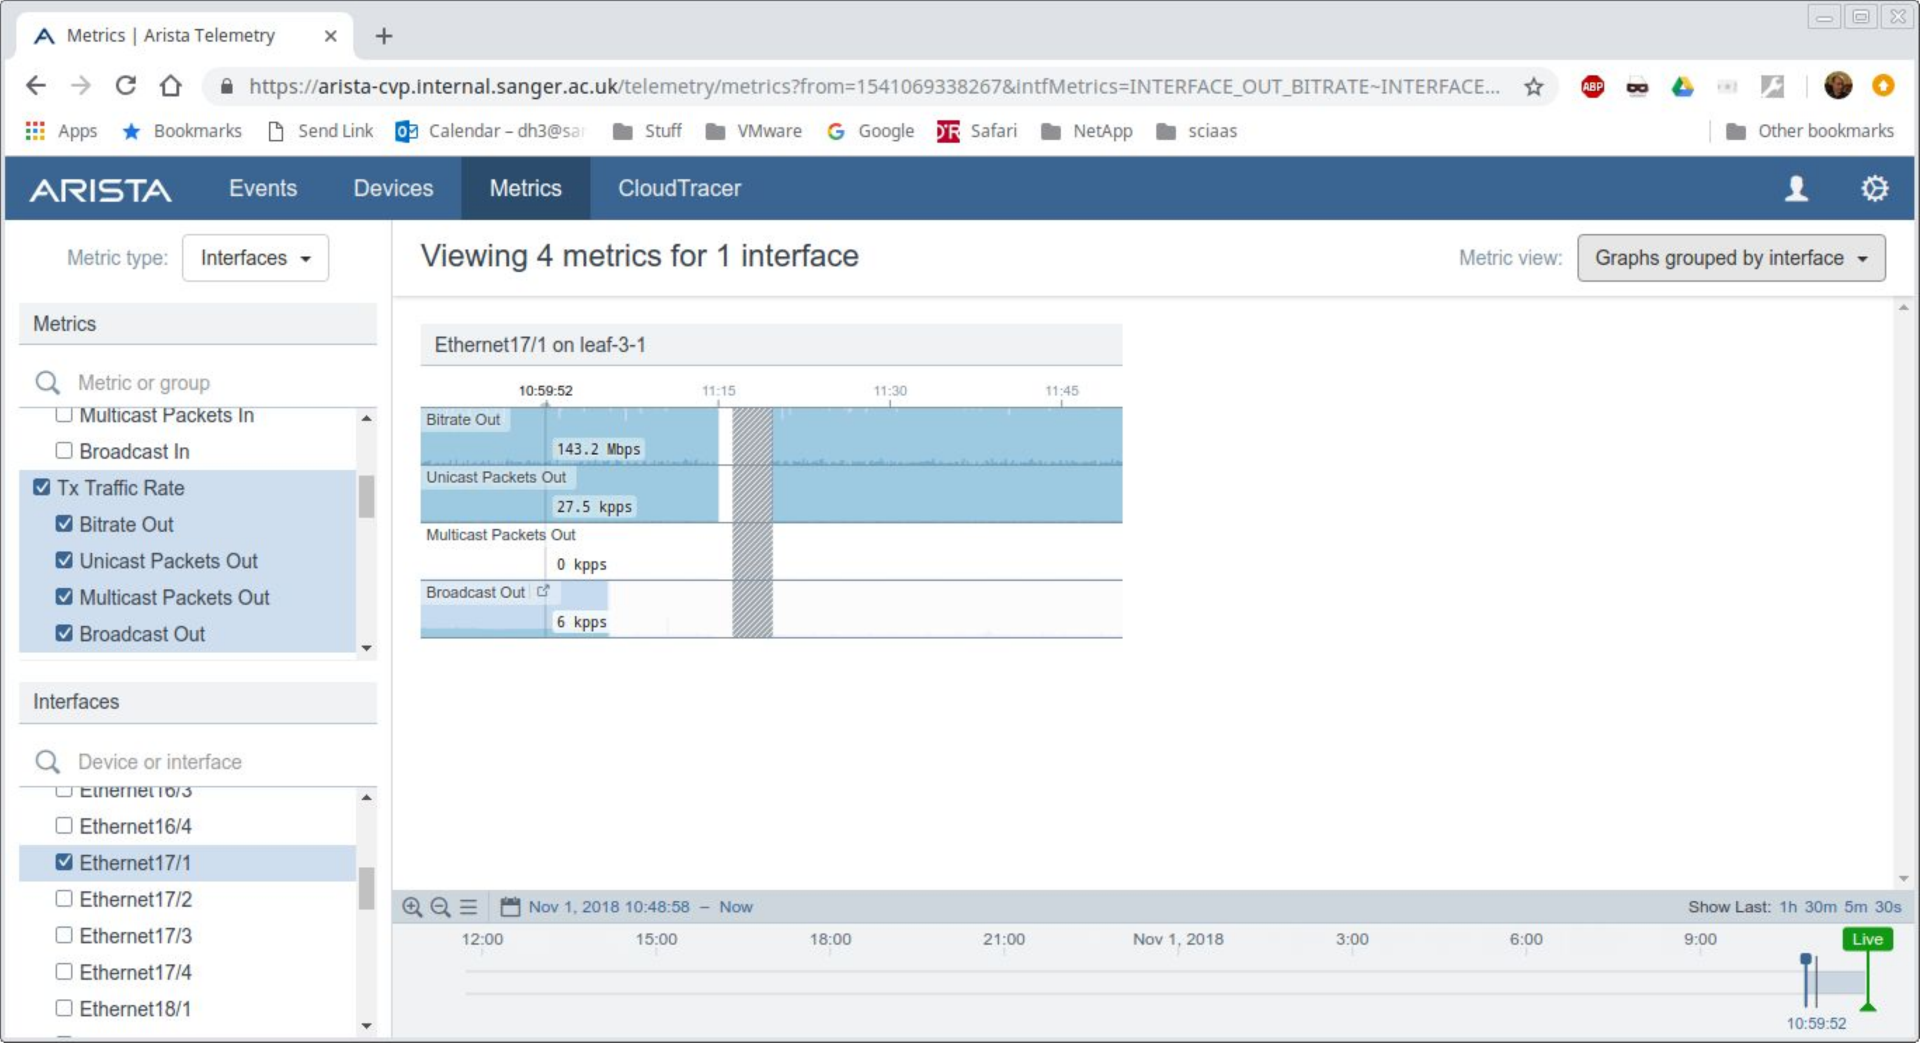

Metrics | Arista Telemetry
https://arista-cvp.internal.sanger.ac.uk/telemetry/metrics?from=1541069338267&intfMetrics=INTERFACE_OUT_BITRATE~INTERFACE...
Apps
Bookmarks
Send Link
Calendar – dh3@sa
Stuff
VMware
Google
Safari
NetApp
sciaas
Other bookmarks
ARISTA
Events
Devices
Metrics
CloudTracer

Metric type: Interfaces

Viewing 4 metrics for 1 interface

Metric view: Graphs grouped by interface

Metrics

Metric or group

- ☐ Multicast Packets In
- ☐ Broadcast In
- ☑ Tx Traffic Rate
  - ☑ Bitrate Out
  - ☑ Unicast Packets Out
  - ☑ Multicast Packets Out
  - ☑ Broadcast Out

Interfaces

Device or interface

- ☐ Ethernet16/3
- ☐ Ethernet16/4
- ☑ Ethernet17/1
- ☐ Ethernet17/2
- ☐ Ethernet17/3
- ☐ Ethernet17/4
- ☐ Ethernet18/1

Ethernet17/1 on leaf-3-1

Nov 1, 2018 10:48:58 – Now

Show Last: 1h 30m 5m 30s

Upgrades and expansions

Newton (RHOSP10) - “Epsilon”

- experiment for experience of “sidegrade”
- not released beyond guinea pigs

Pike (RHOSP12) - “Zeta”

- current production system
- many of our original customisations and fixes no longer necessary
- side-by-side upgrade was time-consuming (live migration not working in Liberty)

Pike (RHOSP12) - “Zeta”

- containerised overcloud services mean different ways of managing/customising services
- tenant networks can be VLAN via Arista ml2 plugin (but still VXLAN by default)
- enabled CPU overcommit, 8x, “o1.*” flavours
- enabled jumbo frames
- instance live migration works out of the box

Ceph growth

- expansion: 1PB → 4.5PB usable

  9 servers → 51 servers

  540 OSDs → 3060 OSDs
- current use: 1.3PB as S3 objects, 800TB as Cinder volumes

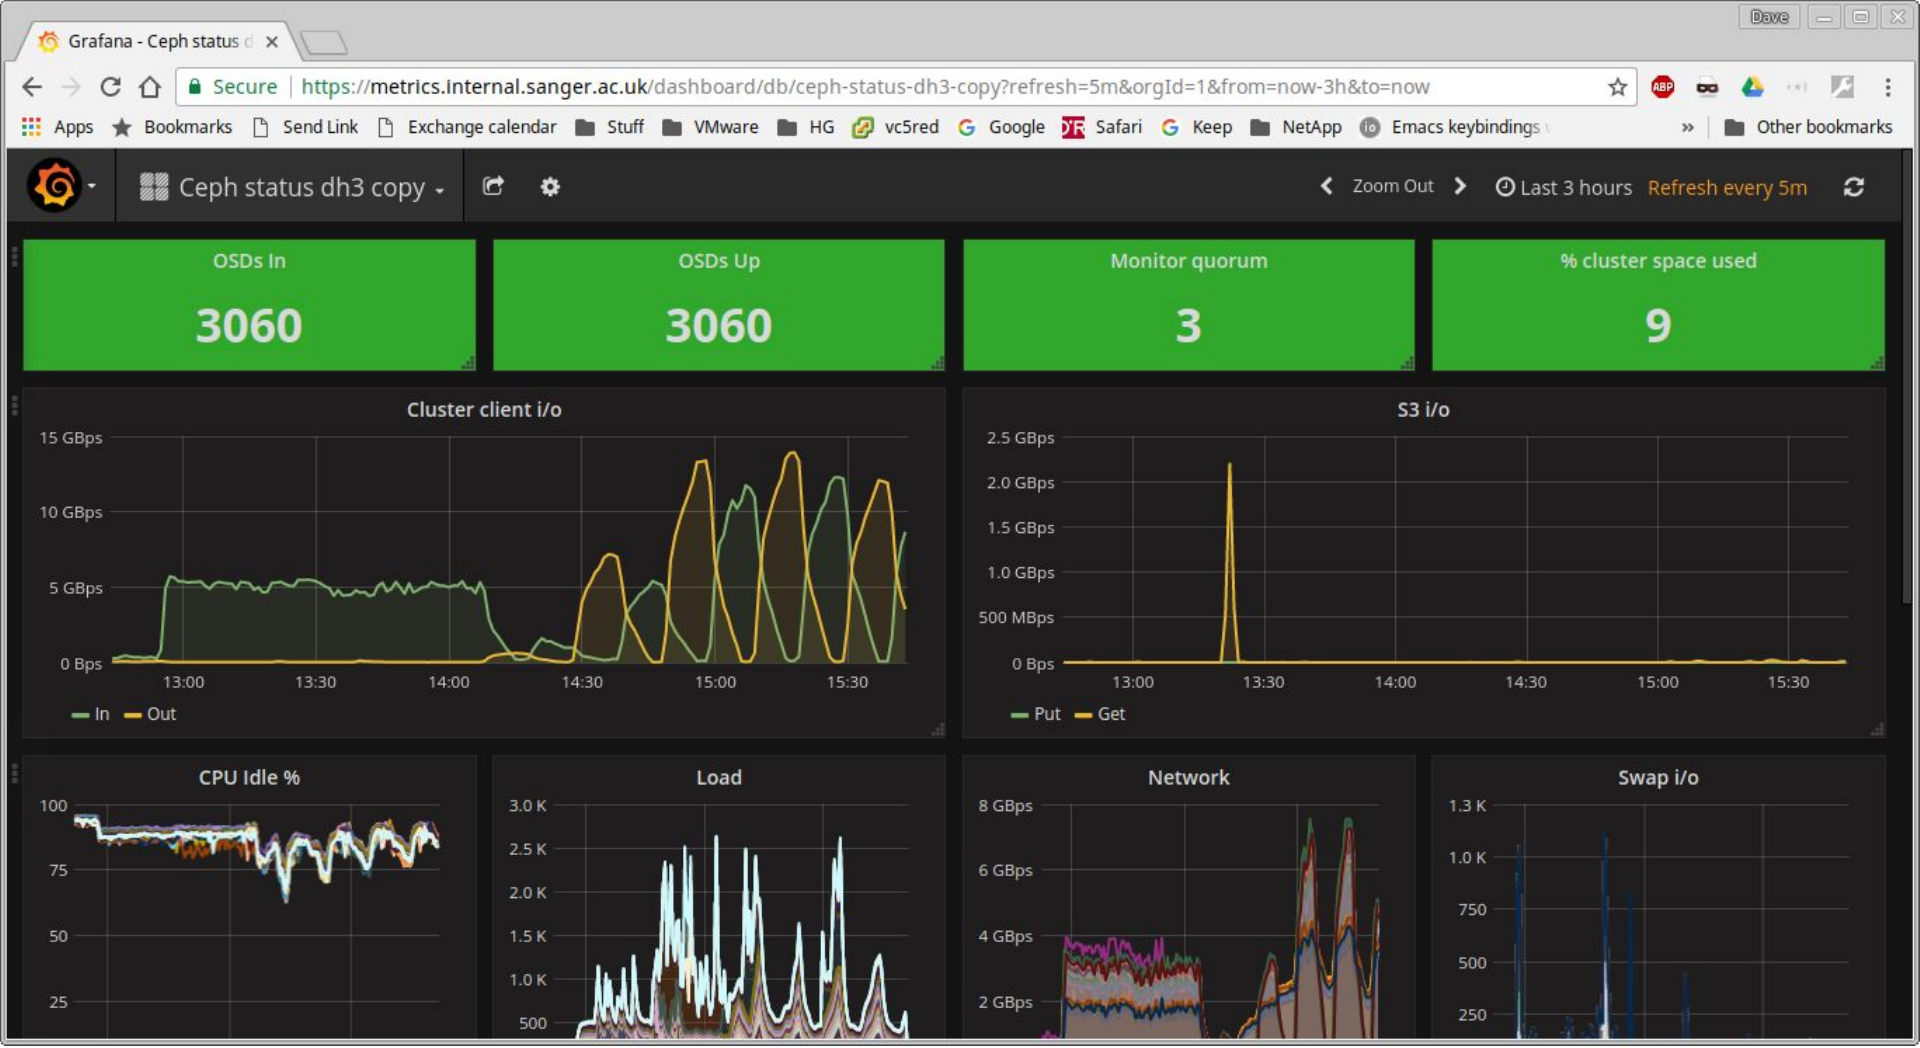

OSDs In: 3060

OSDs Up: 3060

Monitor quorum: 3

% cluster space used: 9

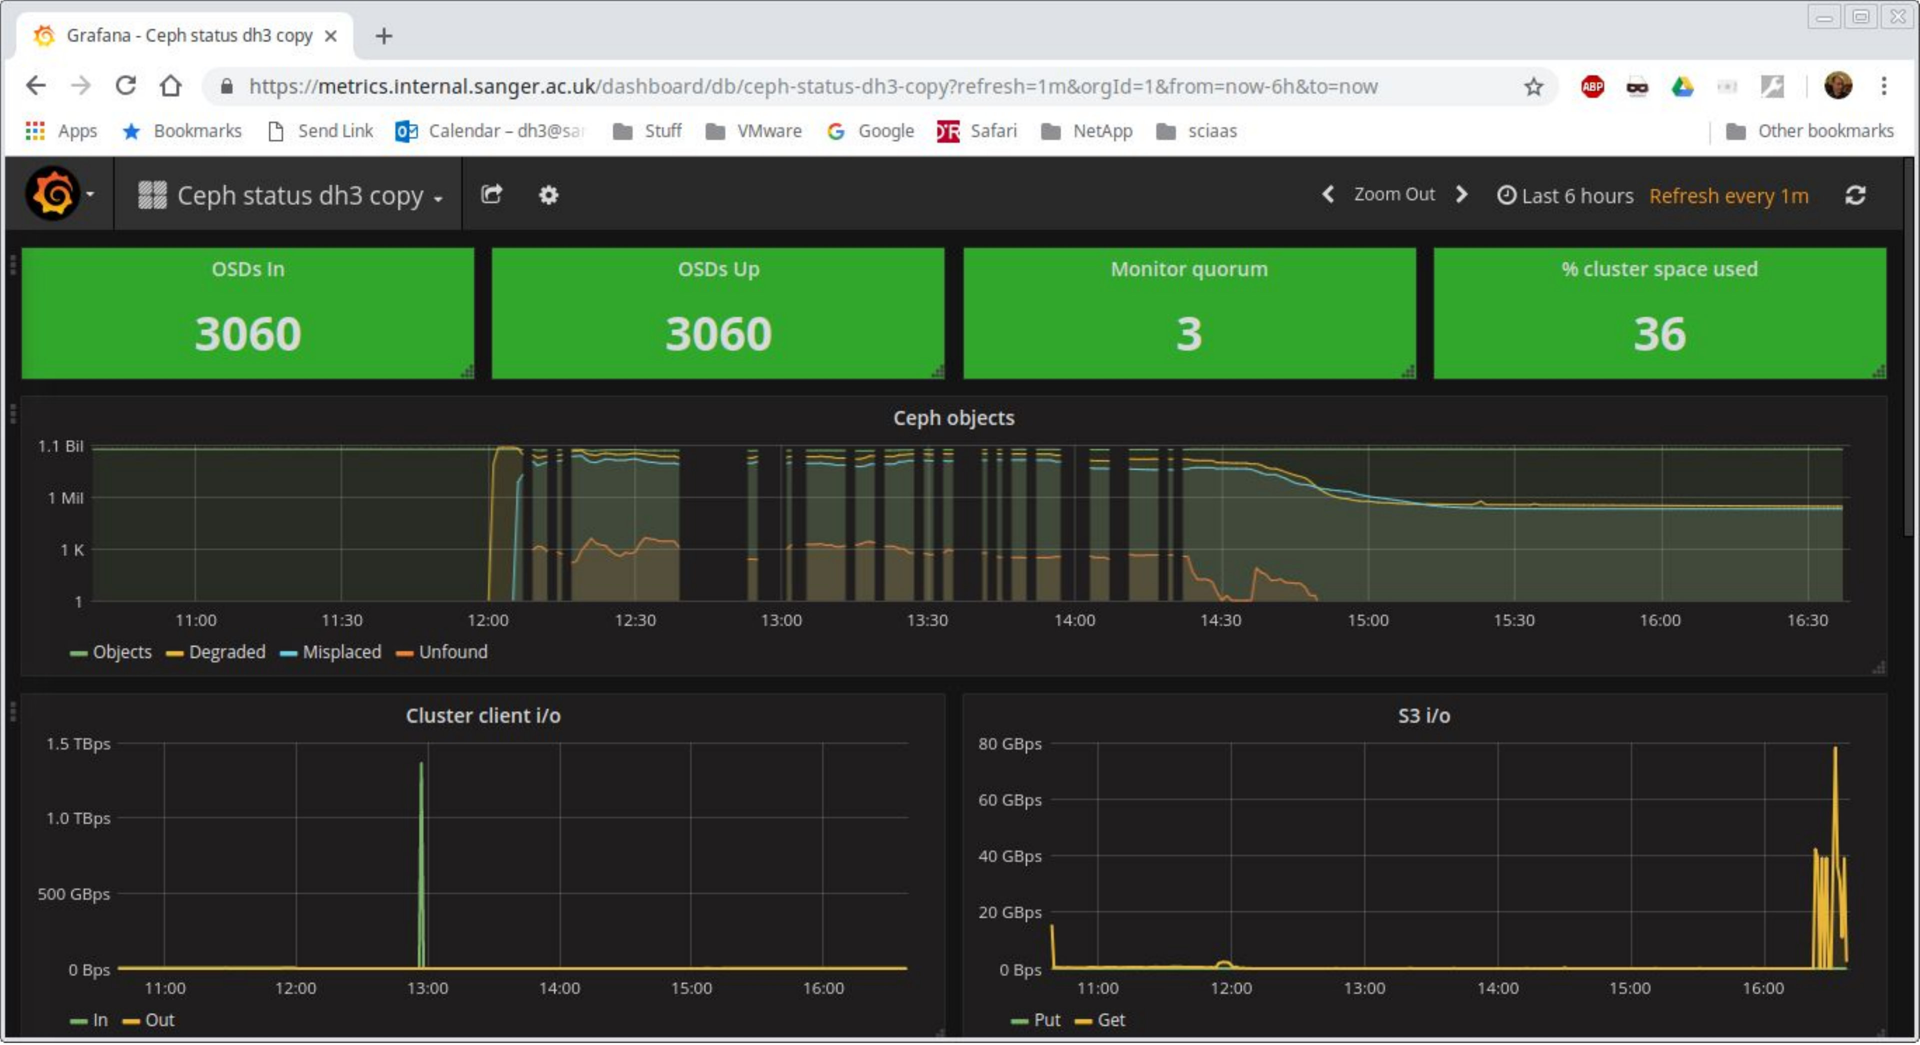

OSDs In 3060

OSDs Up 3060

Monitor quorum 3

% cluster space used 36

Ceph objects

Cluster client i/o

S3 i/o

Expansion and growth

- entire system physically relocated in July 2018 due to power requirements
- 67 compute nodes added during September and October 2018, another 32 to add shortly

ColdLogik
RDC-07

Enhancements

Ceph

- radosgw S3 pre-signed URLs as a stand-in for bucket policies
- FineUploader JavaScript framework
- radosgw bug (presence of Content-Type in upload policy) through ceph-users, ceph-devel, Red Hat support, to release

Ceph/S3 - secure file upload

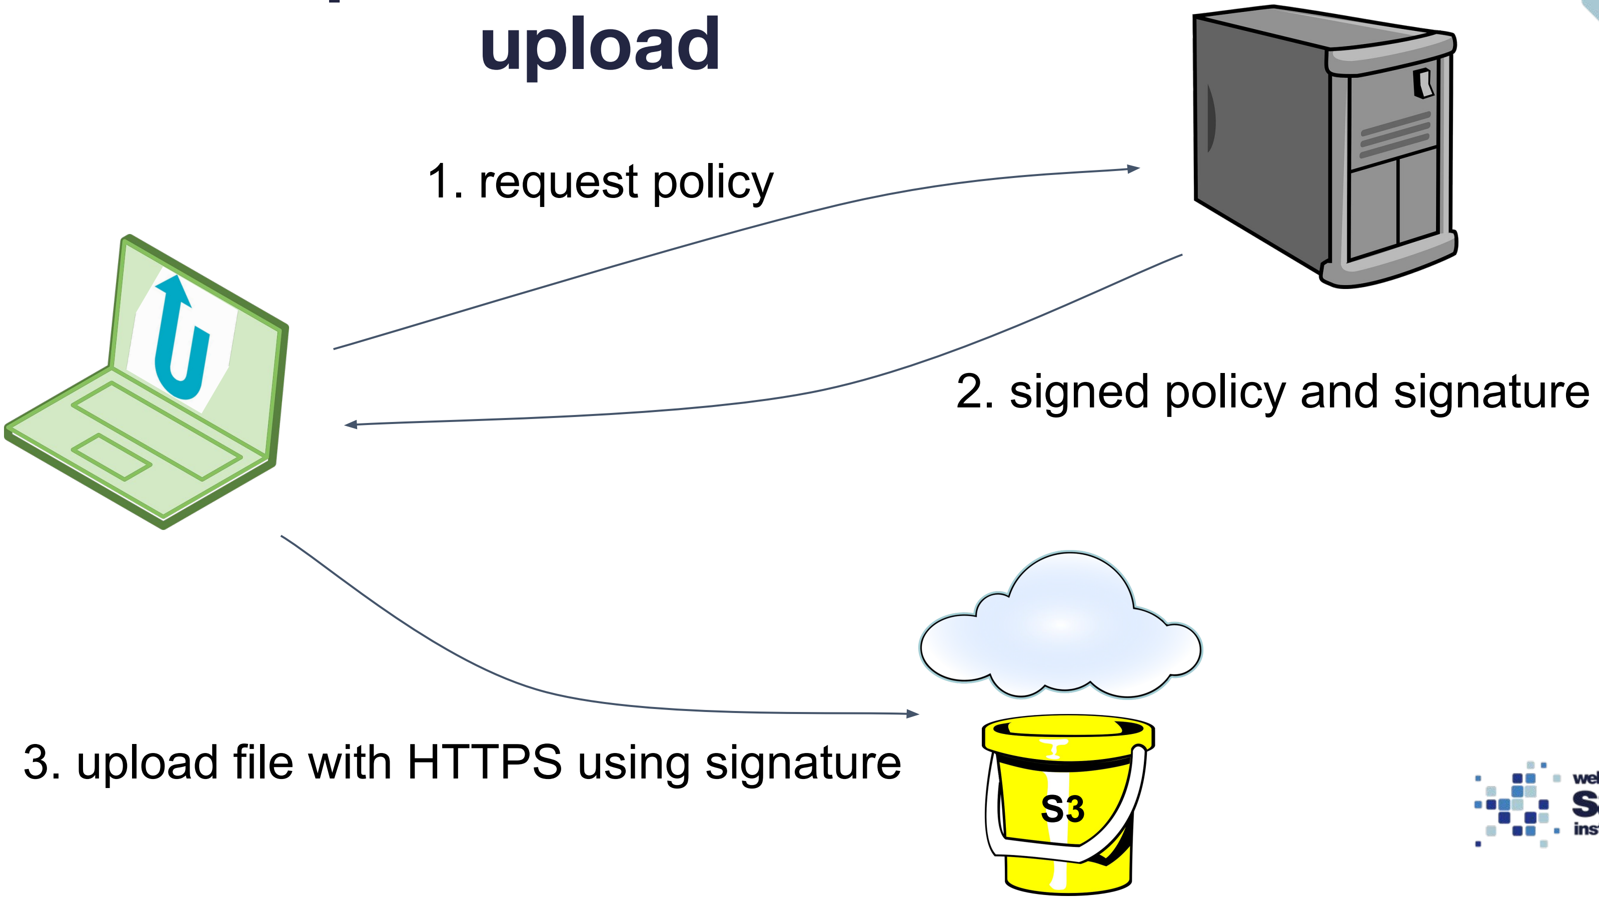

1. request policy

2. signed policy and signature

3. upload file with HTTPS using signature

Ceph

- automated radosgw/S3 public bucket scan

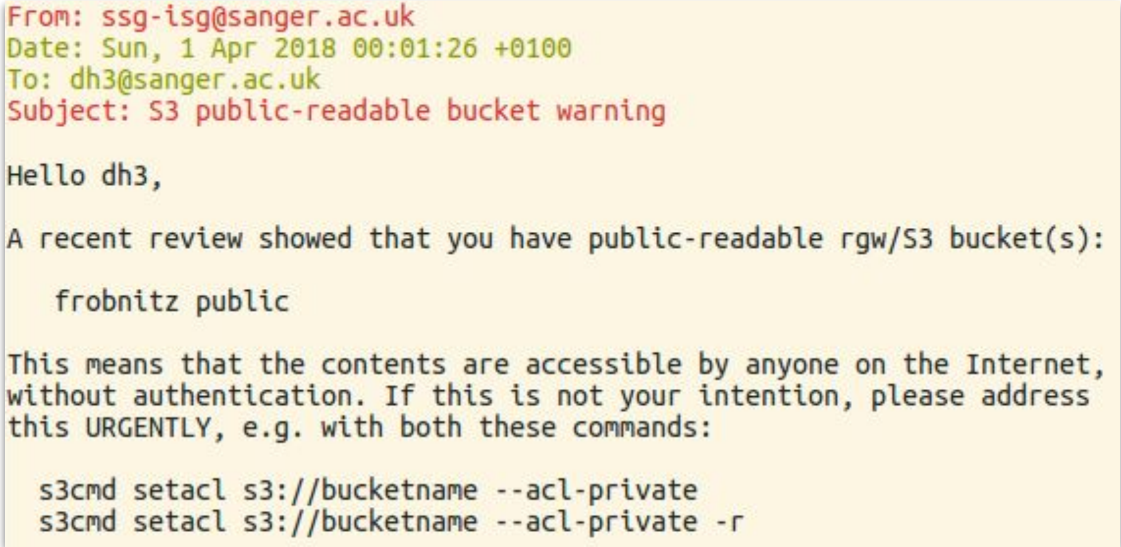

```
From: ssg-isg@sanger.ac.uk
Date: Sun, 1 Apr 2018 00:01:26 +0100
To: dh3@sanger.ac.uk
Subject: S3 public-readable bucket warning

Hello dh3,

A recent review showed that you have public-readable rgw/S3 bucket(s):

   frobnitz public

This means that the contents are accessible by anyone on the Internet,
without authentication. If this is not your intention, please address
this URGENTLY, e.g. with both these commands:

  s3cmd setacl s3://bucketname --acl-private
  s3cmd setacl s3://bucketname --acl-private -r
```

wellcome sanger institute

Provider networks

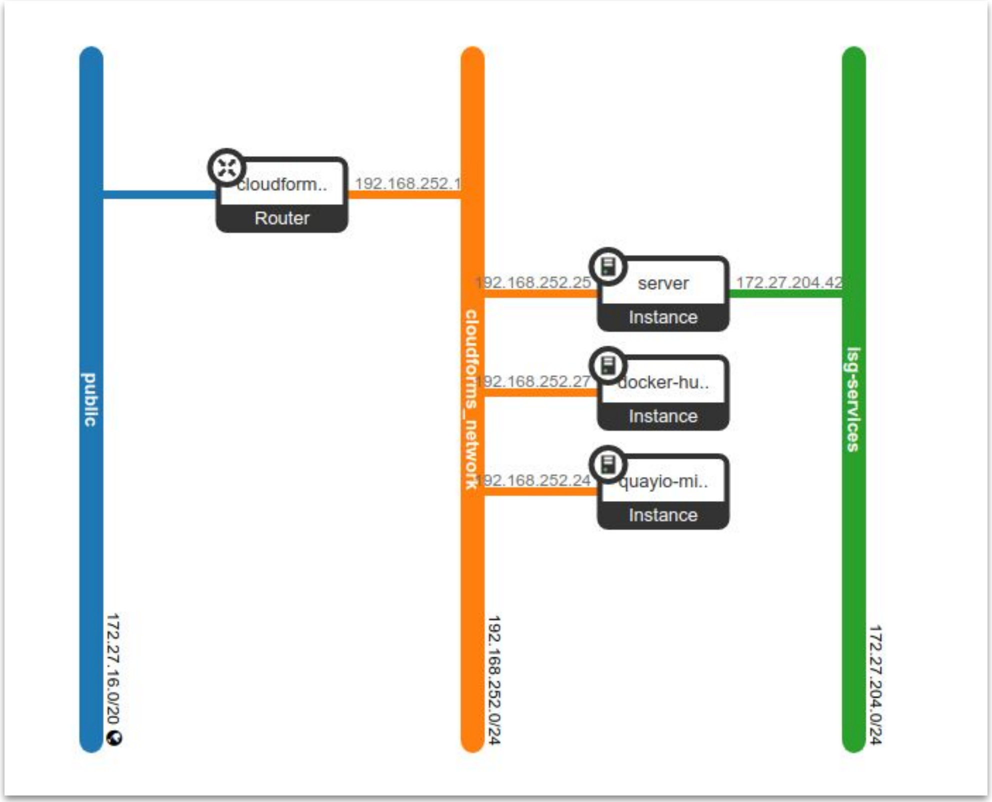

Provider networks

- “Secure Lustre” - multi-tenant
- scientific instruments
- farm4
- niggles with security groups
- “a hammer for all screws”? e.g. faster S3 access

“Secure Lustre”

- strong tenant isolation for a POSIX filesystem
- presented at LAD’17, SC’17, London OpenStack Day
- published in “The Crossroads of Cloud and HPC”

“Secure Lustre”

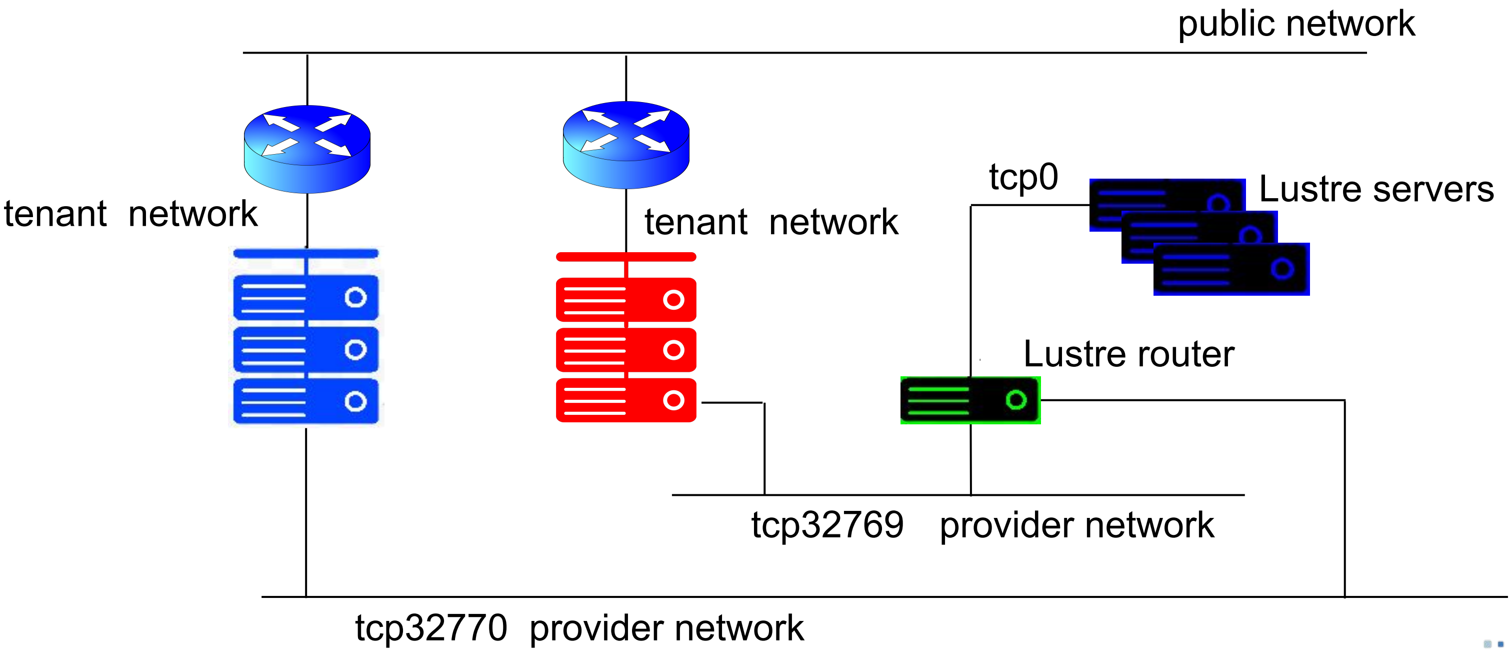

“But what’s it all for?”

Services and applications implemented on OpenStack

- Mutational Signatures
- HGI Arvados
- Gitlab CI runners
- farm4 (extending LSF into OpenStack)
- Mattermost
- CellphoneDB
- CloudForms
- ...

wellcome sanger institute

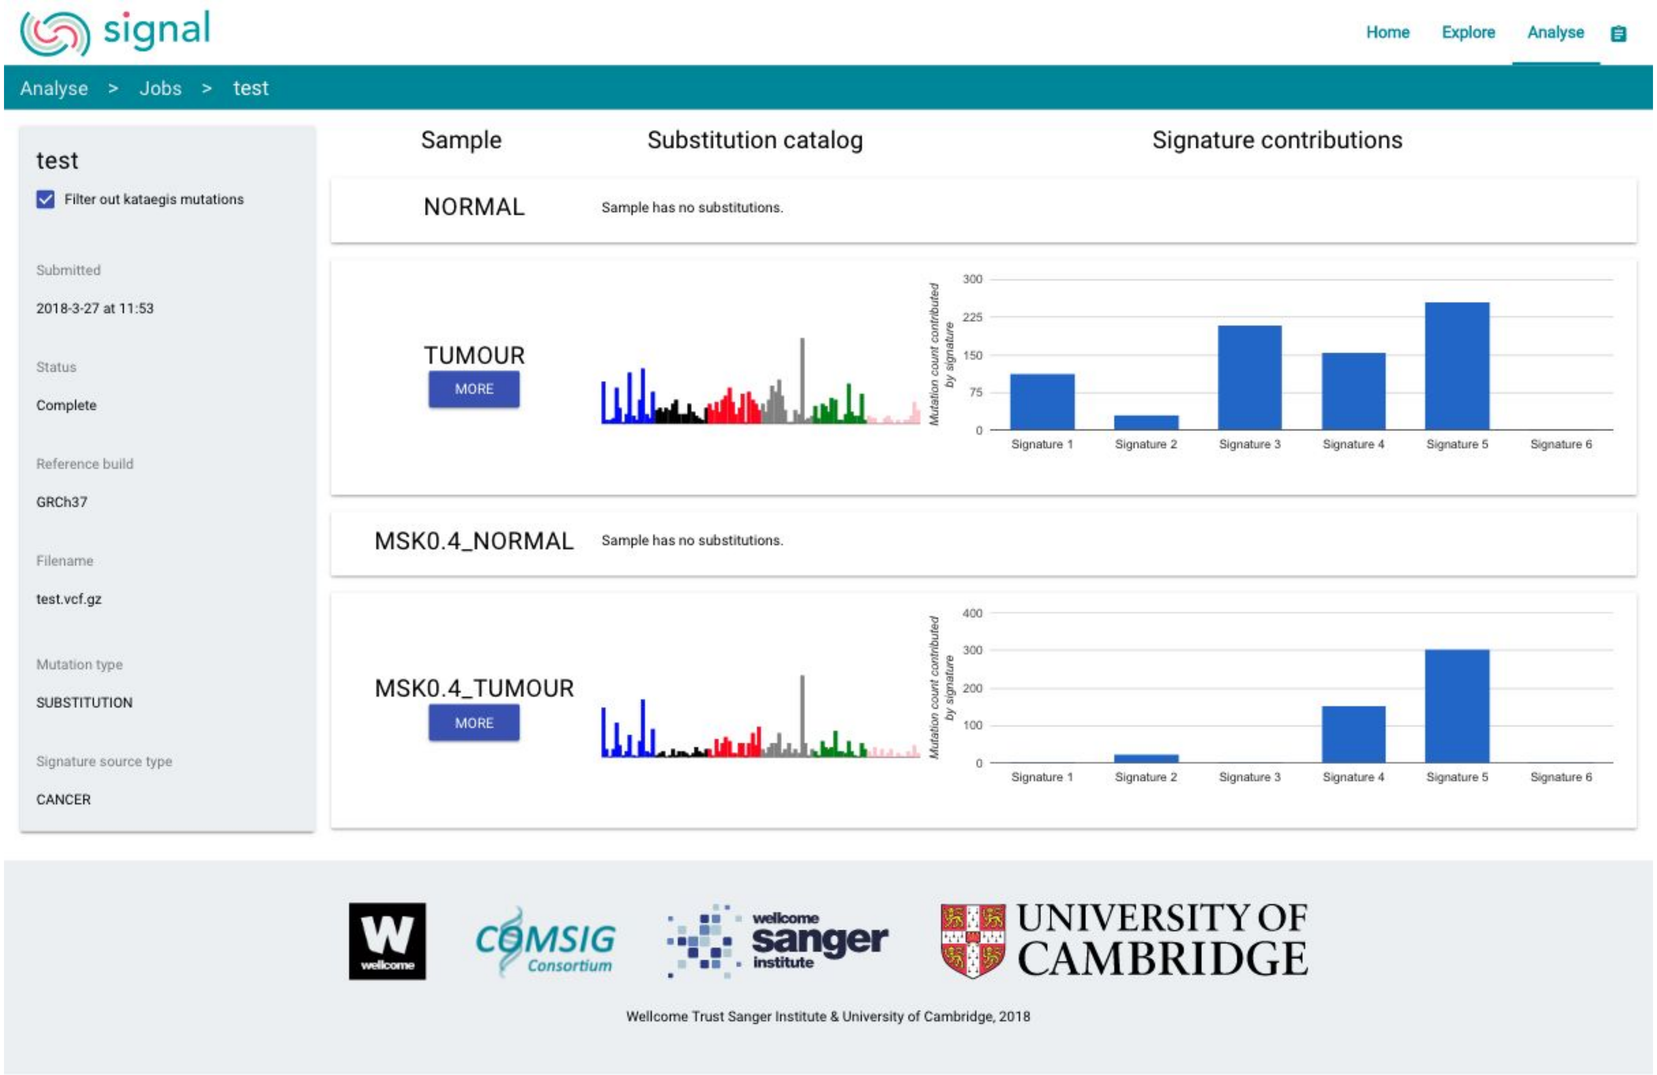

test

☑ Filter out kataegis mutations

Submitted

2018-3-27 at 11:53

Status

Complete

Reference build

GRCh37

Filename

test.vcf.gz

Mutation type

SUBSTITUTION

Signature source type

CANCER

NORMAL

Sample has no substitutions.

TUMOUR

MSK0.4_NORMAL

Sample has no substitutions.

MSK0.4_TUMOUR

Wellcome Trust Sanger Institute & University of Cambridge, 2018

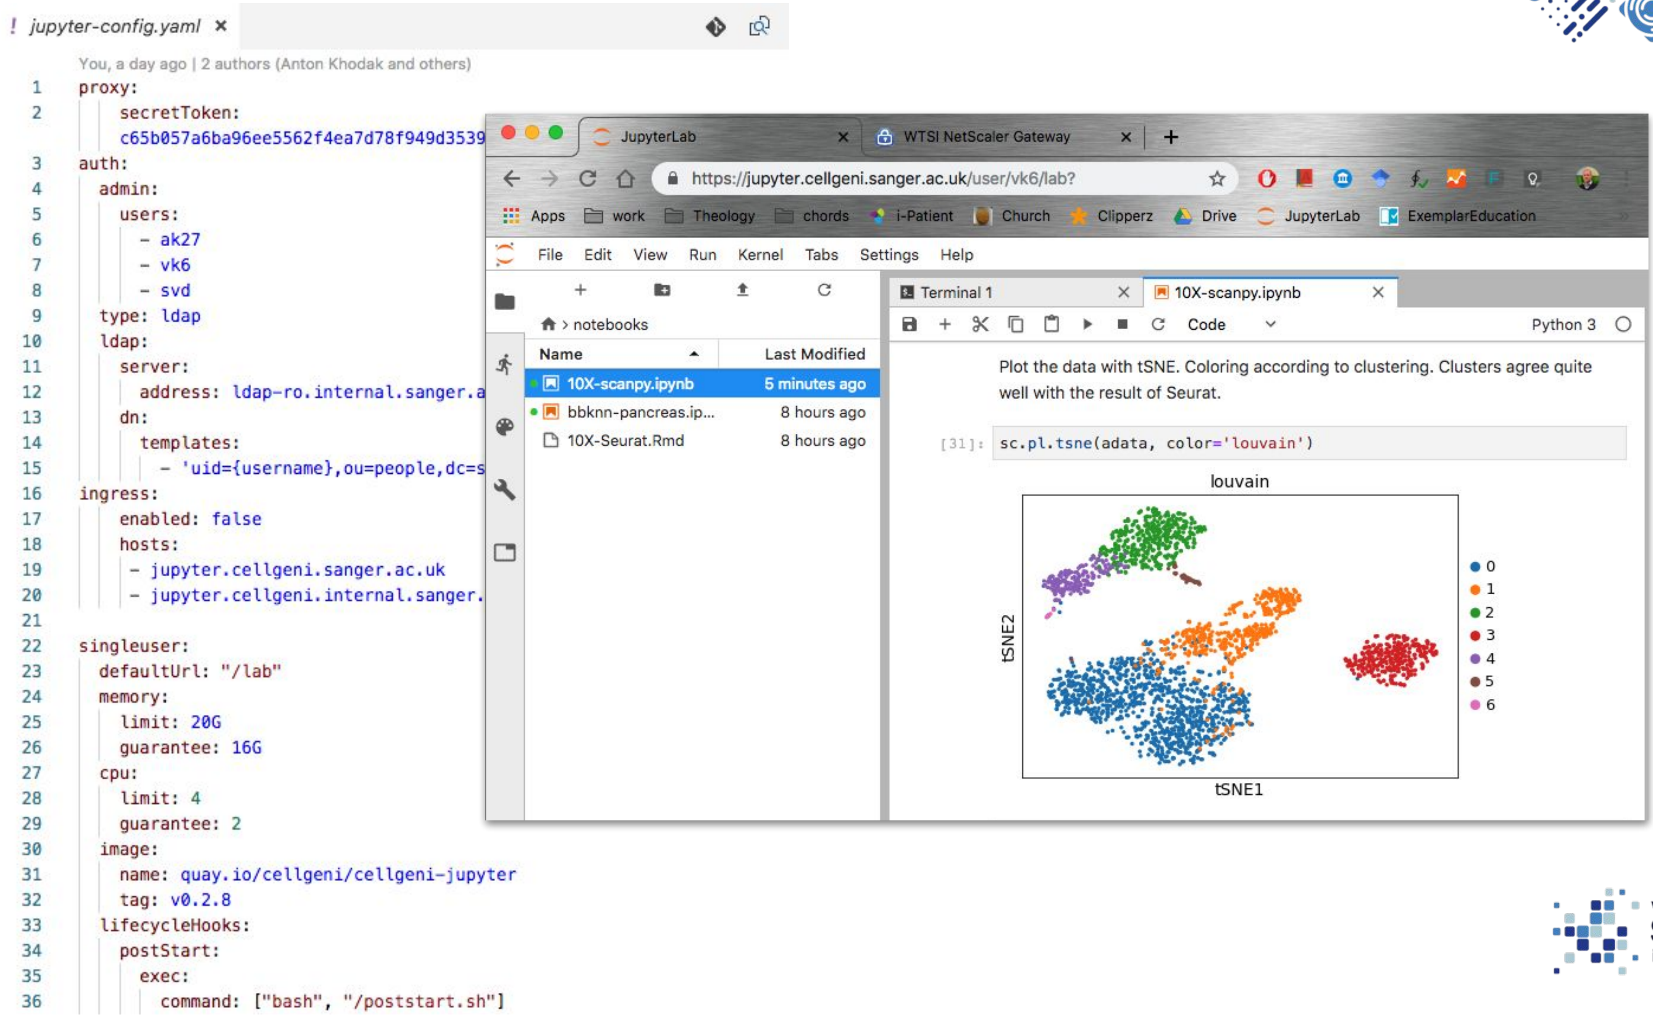

```yaml
proxy:
    secretToken:
    c65b057a6ba96ee5562f4ea7d78f949d3539
auth:
  admin:
    users:
      - ak27
      - vk6
      - svd
  type: ldap
  ldap:
    server:
      address: ldap-ro.internal.sanger.a
    dn:
      templates:
        - 'uid={username},ou=people,dc=s
ingress:
    enabled: false
    hosts:
    - jupyter.cellgeni.sanger.ac.uk
    - jupyter.cellgeni.internal.sanger.

singleuser:
  defaultUrl: "/lab"
  memory:
    limit: 20G
    guarantee: 16G
  cpu:
    limit: 4
    guarantee: 2
  image:
    name: quay.io/cellgeni/cellgeni-jupyter
    tag: v0.2.8
  lifecycleHooks:
    postStart:
      exec:
        command: ["bash", "/poststart.sh"]
```

Name	Last Modified
10X-scanpy.ipynb	5 minutes ago
bbknn-pancreas.ip...	8 hours ago
10X-Seurat.Rmd	8 hours ago

Plot the data with tSNE. Coloring according to clustering. Clusters agree quite well with the result of Seurat.

```
[31]: sc.pl.tsne(adata, color='louvain')
```

Science-as-a-service

- on-premise cloud for less technical users
- OpenStack and VMware with CloudForms orchestration
- pilot implementation well received

UK Biobank Vanguard

- 50,000 whole genome sequences over 18 months
- NovaSeq 6000: 6Tb/sequencer/2 days

HPCwire, 12 November 2018

HPC wire
HPC wire 2018 Readers' Choice Awards
CELEBRATING 30 YEARS
Best Use of HPC in the Cloud
The Wellcome Sanger Institute using a private OpenStack cloud enhances IT environment necessary to sequence and assemble 100 complete human genomes per day.

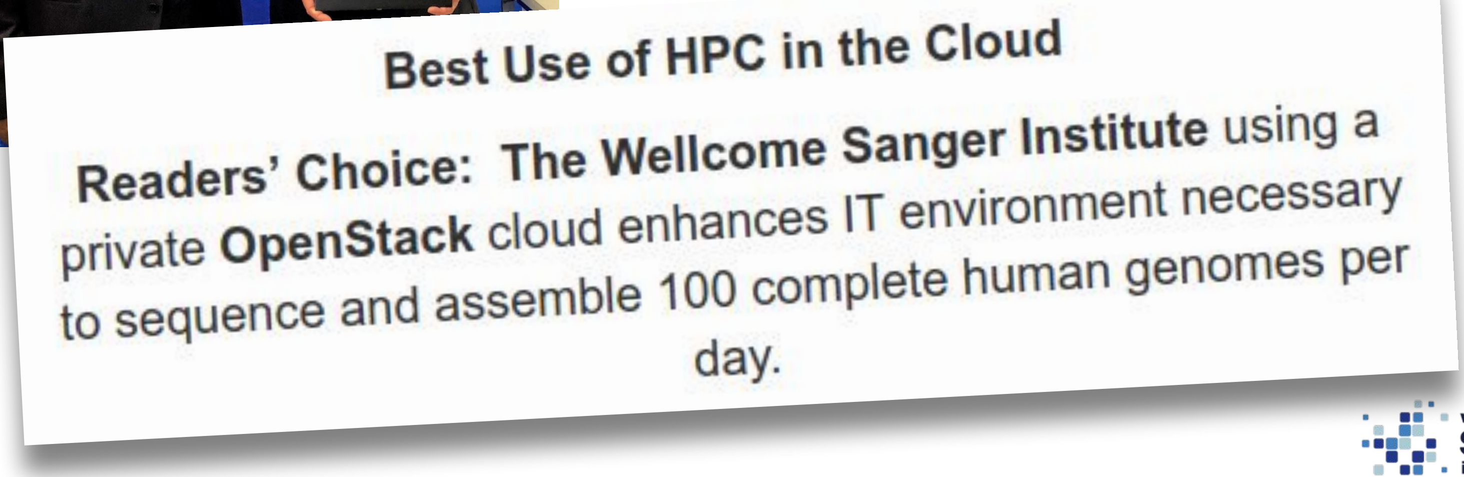

Best Use of HPC in the Cloud

Readers' Choice: The Wellcome Sanger Institute using a private **OpenStack** cloud enhances IT environment necessary to sequence and assemble 100 complete human genomes per day.

Compare and contrast

First draft of the human genome (Nature, 15 February 2001)

- took 10 years
- cost $2.7billion

Education/training

- Hashicorp products
- Ceph (for sysadmins)
- bespoke end-user training (OCF)

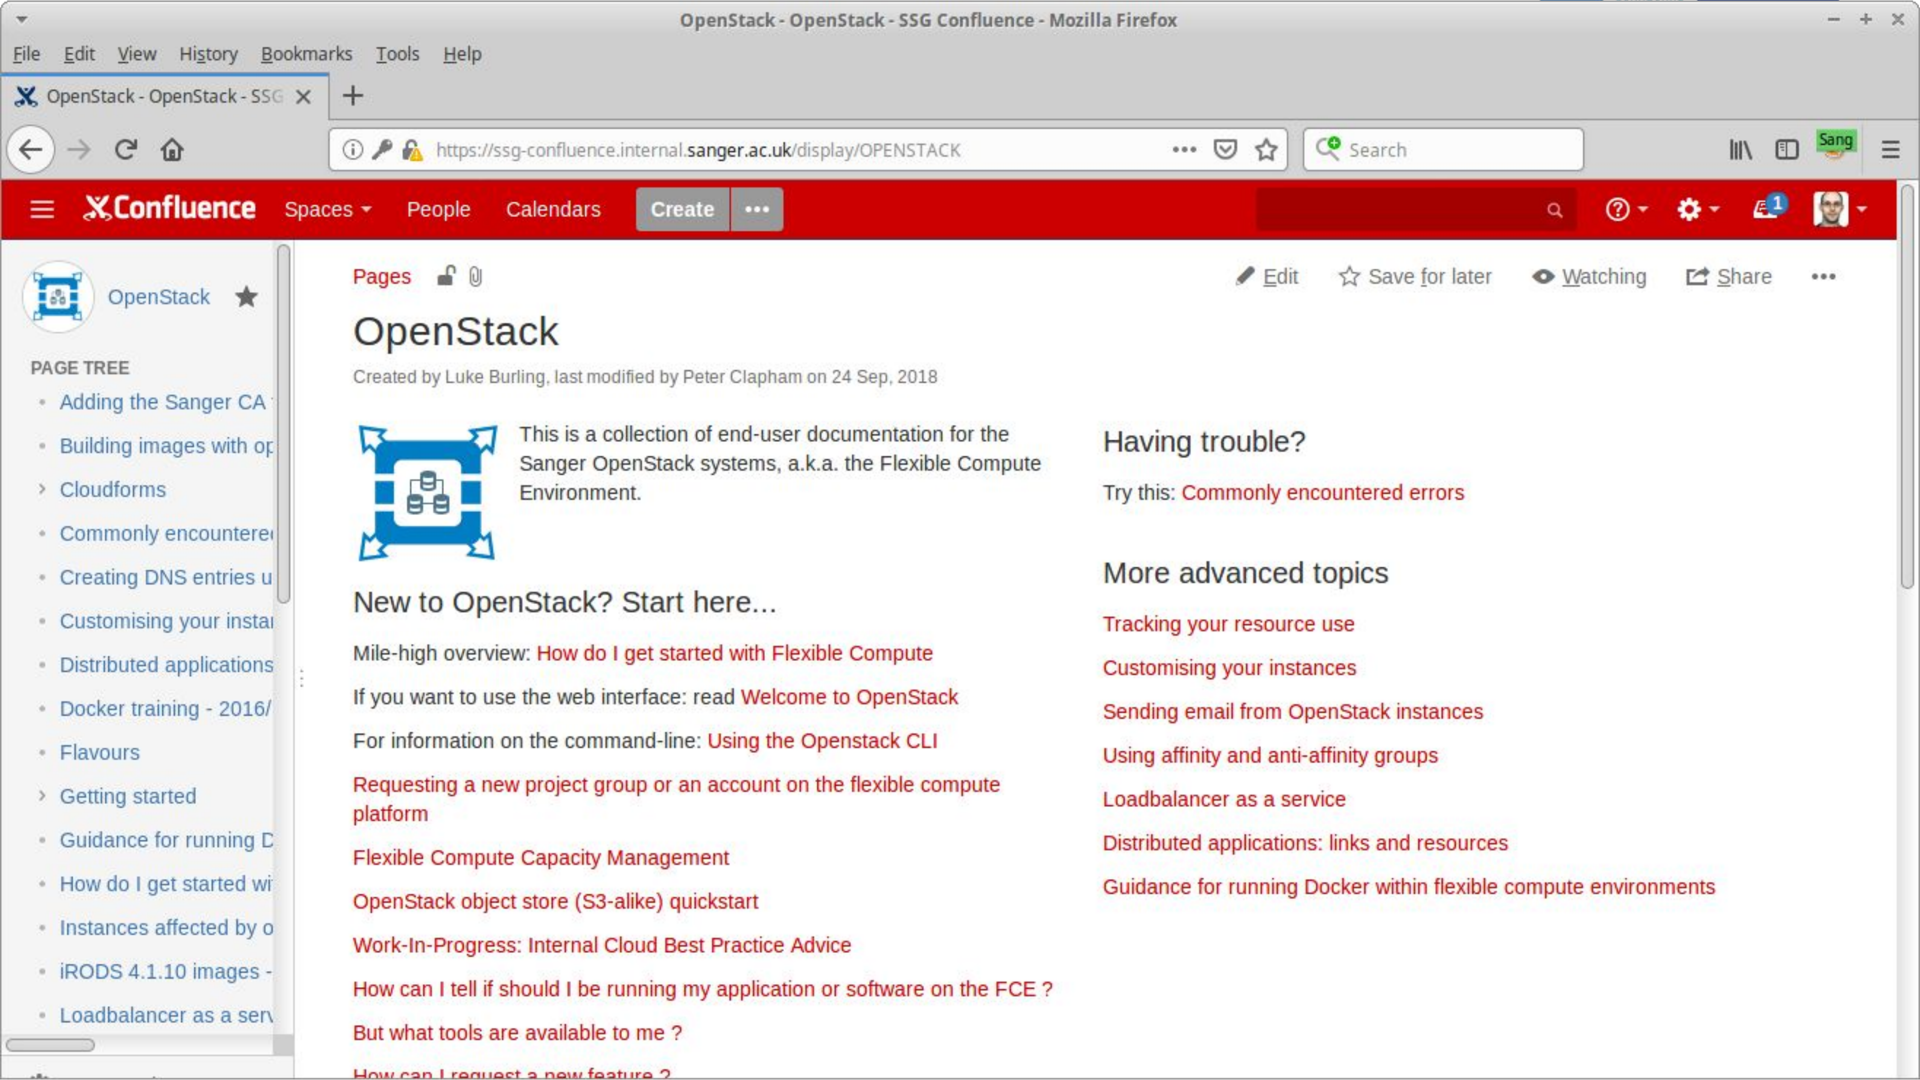

OpenStack

Created by Luke Burling, last modified by Peter Clapham on 24 Sep, 2018

This is a collection of end-user documentation for the Sanger OpenStack systems, a.k.a. the Flexible Compute Environment.

New to OpenStack? Start here...

Mile-high overview: How do I get started with Flexible Compute

If you want to use the web interface: read Welcome to OpenStack

For information on the command-line: Using the Openstack CLI

Requesting a new project group or an account on the flexible compute platform

Flexible Compute Capacity Management

OpenStack object store (S3-alike) quickstart

Work-In-Progress: Internal Cloud Best Practice Advice

How can I tell if should I be running my application or software on the FCE ?

But what tools are available to me ?

How can I request a new feature ?

Having trouble?

Try this: Commonly encountered errors

More advanced topics

Tracking your resource use

Customising your instances

Sending email from OpenStack instances

Using affinity and anti-affinity groups

Loadbalancer as a service

Distributed applications: links and resources

Guidance for running Docker within flexible compute environments

The Sanger community

- user engagement
- Slack
- coffee mornings

The wider community

- OpenStack summits
- Scientific SIG
- openstack-operators mailing list (soon to be openstack-discuss)

The future

Queens (RHOSP13)

- want to follow Red Hat LTS releases
- test deployments are underway
- will sideways-upgrade again
- dedicated networking nodes - composable roles
- move customizations from Ansible into deployment templates where possible

New features

Driven by user demand

- Barbican - enabling encrypted volumes
- ~~Manila~~
- Octavia
- Sahara?

Ceph upgrades

- version 10 (Jewel) to 12 (Luminous)
    - for bug fixes and features
- backend storage format: FileStore to BlueStore
    - for performance
- in-place upgrade is planned; testing is underway

Offsite DR/BCP

- small OpenStack and Ceph clusters at JSDC
- currently investigating global load-balancing options
- open questions about user federation/identity management
- need to investigate Ceph replication

Data flow model

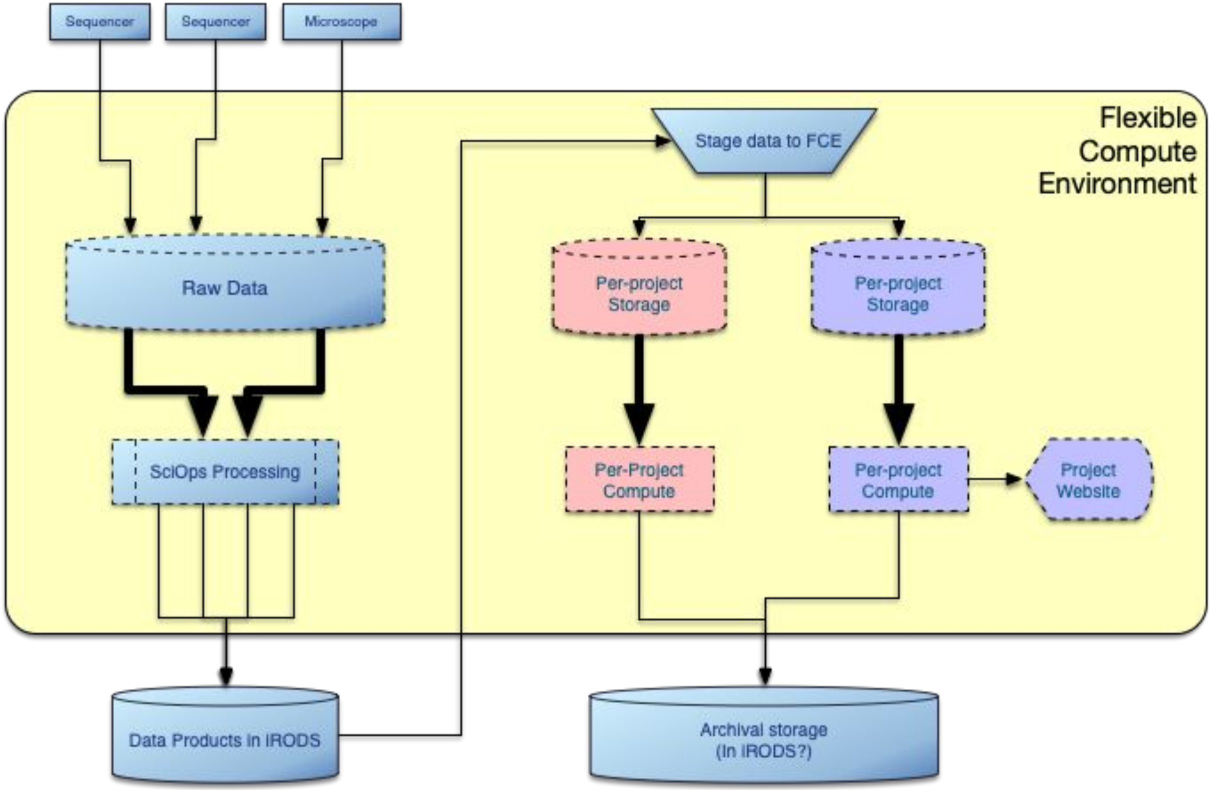

Evolution and federation

Public clouds

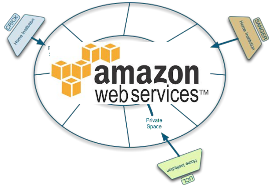

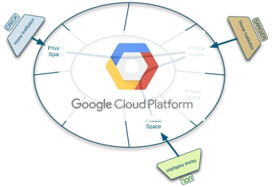

Global Alliance
GA4GH APIs

Private clouds

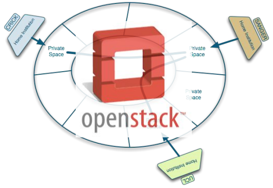

EBI
Embassy
Cloud

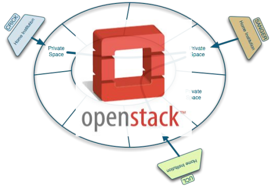

Sanger
Flexible
Compute

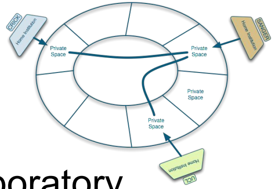

OICR
Collaboratory

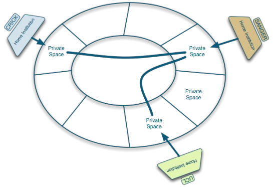

GDC Cloud
(Genomic Data Commons)

Cancer analysis pipelines

2013

2014

2015

2016

2017

2018

ICGC/TCGA Project

- Analyse 2,800 WGS Cancers
- 3 Analysis Pipelines
- (Broad, Sanger & DKFZ)
- Need to use multiple Data Centres

Problem: analysis pipelines NOT portable

Software Go To Data:
Analysis Pipeline available in multiple Data Centres

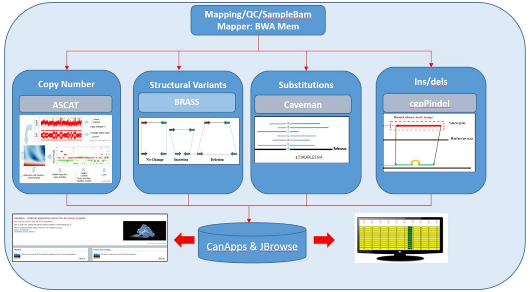

Science Advancements

Large Collaborative Projects

Arctic Ocean
Arctic Ocean
Greenland
Canada
United States
Mexico
North Pacific Ocean
North Atlantic Ocean
Russia
China
Indian Ocean
South Pacific Ocean
South Atlantic Ocean
Australia

IT Advancements

PAN PROSTATE CANCER GROUP

=

redhat

openstack

+

+

Tools Workflows Documentation Sign Up Login

Create, Share, Use

Search Docker Tools and Workflows for the Sciences:

Enter Keyword...

+

happy scientists

Summary

wellcome sanger institute

Dave Holland
dh3@sanger.ac.uk
https://hpc-news.sanger.ac.uk/