

Sharing Hardware between Big Data, AI, HTC and HPC

Lessons learnt expanding CSD3
using OpenStack Ironic

John Garbutt, StackHPC

StackHPC

StackHPC Company Overview

StackHPC

- Formed 2016, based in Bristol, UK
 - Based in Bristol with presence in Cambridge, France and Poland
 - Currently 16 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
- “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
- Develop skills and know-how in Hybrid Cloud




CSD3: Cambridge Service for Data-Driven Discovery

StackHPC

- Intel Xeon, Intel KNL, Nvidia P100
- #75 in the November 2017 Top500 list
- xCAT provisioned Slurm Cluster
 - Running node converted to image
- Resources only accessible via Slurm


<https://www.hpc.cam.ac.uk/>



Motivations Driving a Change

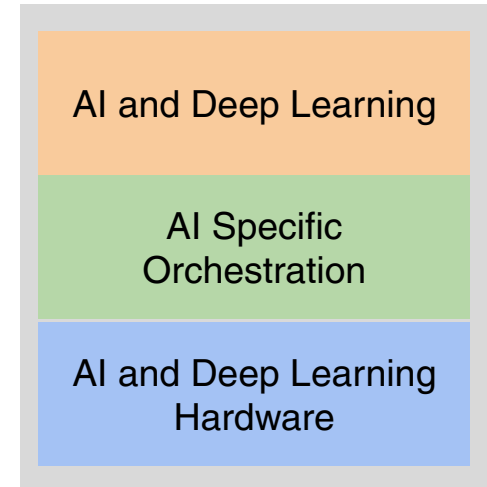


StackHPC

- Manage the increasing complexity
- Better knowledge sharing
- Move away from Resource Silos




HPC Stack 1.0

StackHPC



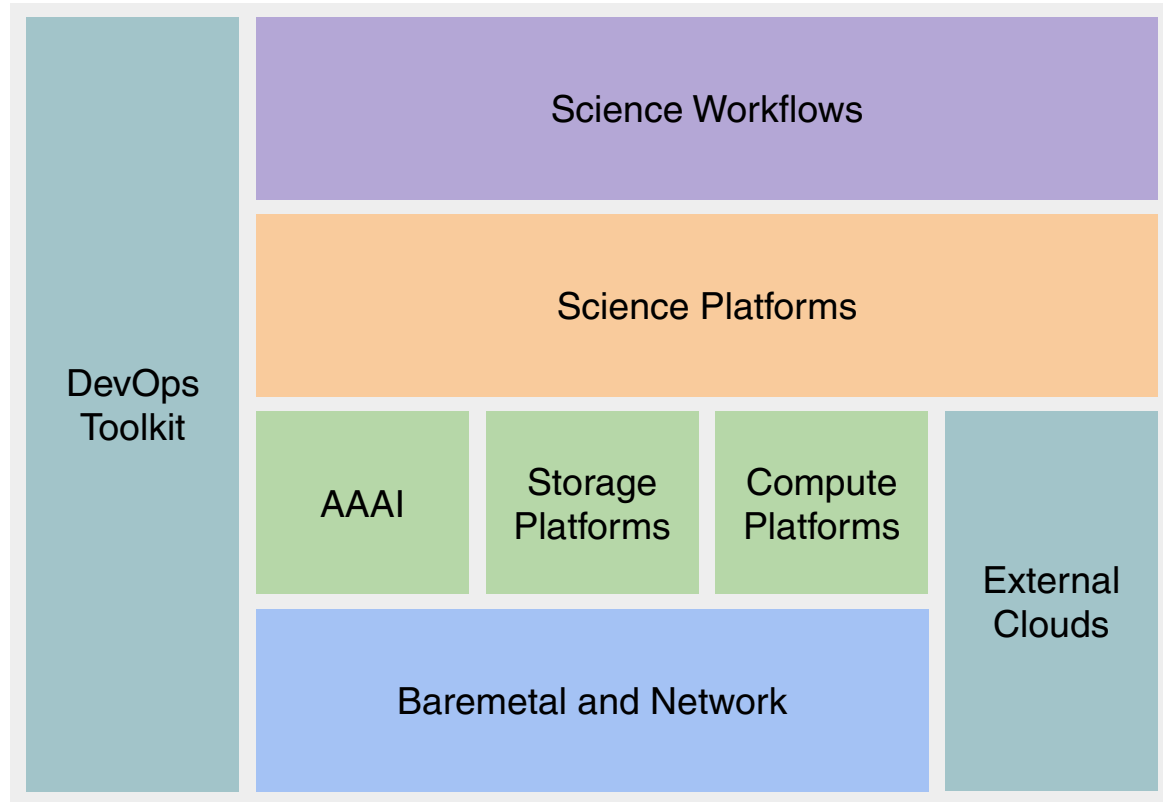
HPC Stack 2.0

StackHPC



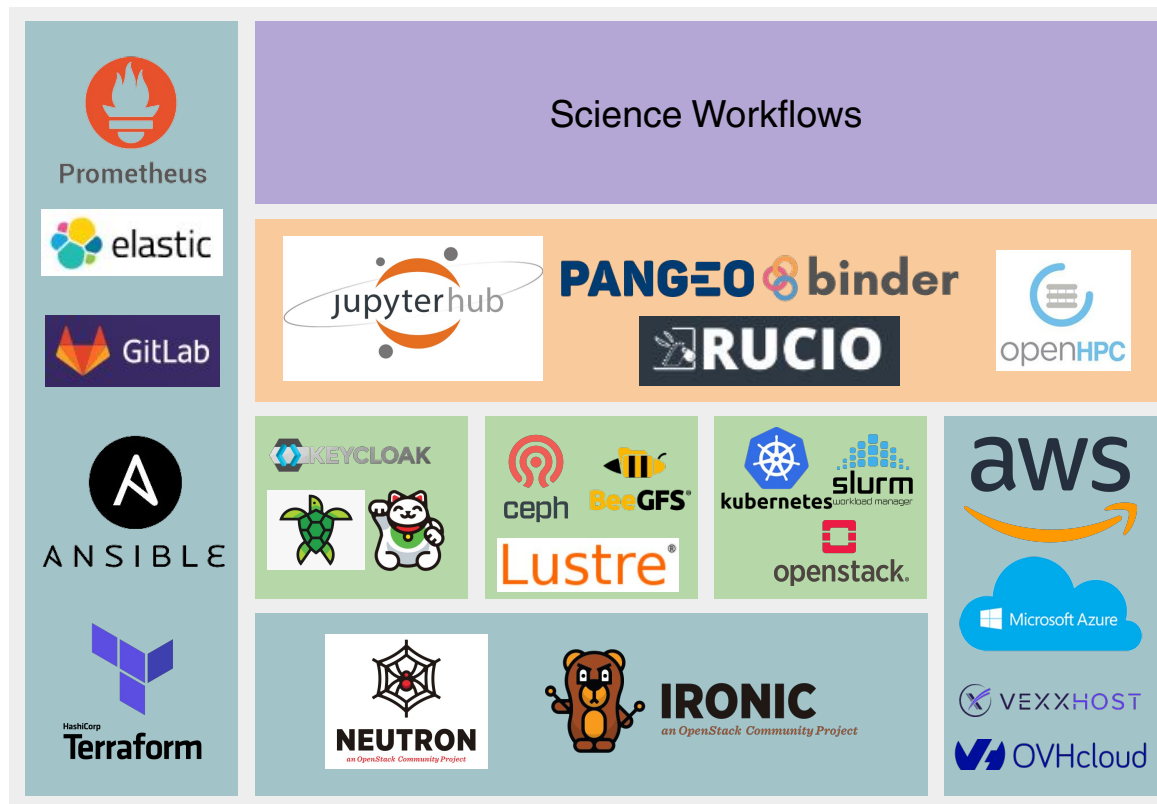
HPC Stack 2.0

StackHPC



HPC Stack 2.0

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)

- 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

StackHPC



openstack®

Expanding CSD3 using OpenStack

Arcus Architecture Overview

StackHPC


- 672 Dell PowerEdge C6420
 - 2 x Intel Xeon 8276, local SSD
 - 1GbE for iDRAC, optionally visible in OS
 - Mellanox ConnectX-6
- Networks
 - HDR100 Infiniband, Mellanox MQM8790
 - 50GbE Ethernet, Mellanox MSN3700
 - 1GbE Ethernet, Dell S3048
- Currently building a similar system that is twice as big



Bootstrapping Compute Nodes

StackHPC

- Scan nodes
 - Re-using existing scanning procedure
 - Encouraged by OpenDev
- 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



<https://github.com/stackhpc/arcus-terraform-idrac>

Building Slurm on OpenStack

StackHPC

- Create cluster
 - Terraform: AZ mapping used to target specific baremetal node
 - Image based deploy, auto joins existing CSD3 Slurm cluster
 - NFS used to provide configuration
 - Ansible for image build and ad-hoc changes
- Rebuild via Slurm
 - Slurm can schedule a reboot after current job completes
 - Can use accounting to limit number of concurrent reboots
 - Slurm reboot takes a reason, include target image
 - Custom reboot script talks to OpenStack



<https://github.com/stackhpc/slurm-openstack-tools>

<https://github.com/stackhpc/ansible-role-openhpc>

Tuning Ironic for Scale

- Focus on on 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
 - Turned off power sync
 - Server Delete: increased rpc_timeout
- connect_timeout in HAproxy and MariaDB



Building Storage on OpenStack

StackHPC

Ceph

- Object Storage
- OpenStack Cinder (RBD)
- OpenStack Manila (CephFS)
- Kubernetes Cloud Provider OpenStack



Storage Managed Services (Software Defined)

- Local or virtual storage, Dedicated disaggregated or Hyperconverged
- Lustre: Spinning Disk, NVMe Data Accelerator, Hadoop, BeegFS...
- ... but data preservation inherently limits flexibility

Monitoring

StackHPC

- kolla-ansible built in
 - Prometheus and Grafana
 - Fluentd, Elasticsearch and Kibana
- Useful extras
 - Node Exporter on Cumulus
 - Redfish exporter
 - GMR: Guru Meditation Reports



Prometheus

What is next?

StackHPC

Baremetal Node <-> Hypervisor

- Works today to deliver STFC IRIS cloud on CSD3
- Kayboe and Kolla-Ansible, OpenStack Ussuri
 - Kayobe uses ansible to configure the network switches
- Moving to use overcloud Ironic to deploy hypervisors
 - Usually Bifrost used to deploy hypervisors and controllers
- Networking
 - Currently no trunk support in ngs cumulus driver
 - Work around using flat network driver

StackHPC



KOLLA

an OpenStack Community Project

Future Consolidation

StackHPC

“One OpenStack to rule them all”

- Migrate previous hardware generations from xCAT to OpenStack
 - More flexibility around accessing GPU resources
- SR-IOV: get RDMA inside VMs with both IB and RoCE
- More diverse Kubernetes clusters
 - Already using OpenStack Magnum
 - ... add in SR-IOV and GPU templates, for things like Horovod
- Improve the “Day 2” Operational Tooling
 - Moving to Prometheus metric based Unified Alerts

<https://github.com/RSE-Cambridge/iris-magnum>



kubernetes

How can I get involved?

OpenStack Special Interest Groups


StackHPC

Scientific OpenStack SIG

<https://www.openstack.org/use-cases/science/>

Bare Metal SIG

<https://www.openstack.org/use-cases/bare-metal/>



Thank You!

StackHPC

- OpenStack Community
- Cambridge University HPCS
- EPSRC, STFC, MRC, HDR UK, UKAEA
- Dell, Intel, NVIDIA, Mellanox Technologies

<https://epsrc.ukri.org/newsevents/news/supercomputing-services-to-drive-scientific-breakthroughs/>