DELL EMC

# OpenStack for Research Computing - A University of Cambridge Perspective

Paul Calleja, Wojciech Turek and John Taylor

intel

UNIVERSITY OF CAMBRIDGE

# Contents



# Introduction

The High Performance Computing Service (HPCS) is part of the wider Research Computing Service (RCS) at the University of Cambridge and has, over the past 10 years, pioneered the exploitation of commodity hardware and open-source technologies in delivering industry-leading supercomputing solutions. This has helped to support world-leading research across a broad range of academic departments within the University as well as providing UK National e-Infrastructure services. Throughout this time the HPCS and latterly the RCS has had a long term partnership with numerous technology partners helping in the implementation and support of these groundbreaking research computing platforms.

Scientific, technical and medical research disciplines are currently undergoing a revolution in terms of a rapidly increasing and widening user base of computational techniques: this is also combined with a significant emergence of data intensive science. These developments promise significant advancement in terms of accelerating the discovery process, but in order to release this potential, traditional HPC solutions and service offerings need to change with the times: HPC systems need to be developed that continue to service traditional HPC use cases but which also service the new rapidly growing "long tail" user demographic [1,10], characterised by smaller usage patterns from non-computational experts bringing compute to the data.

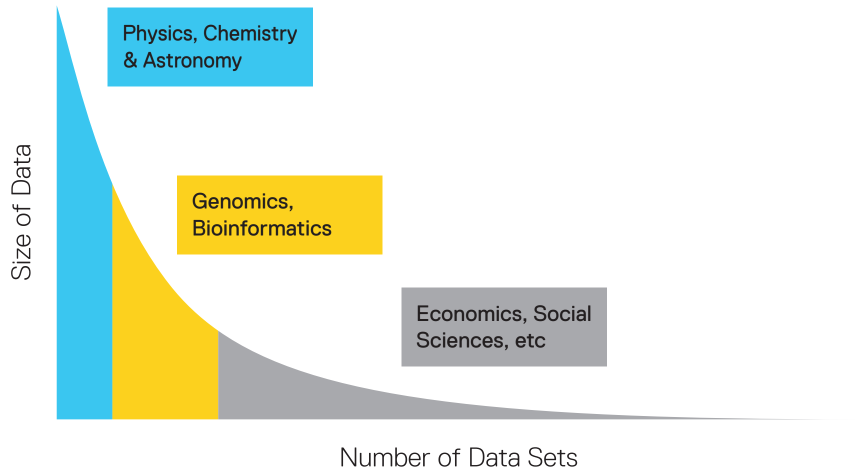


*Computational and Data Requirements for the 'Long Tail of Science'.*

Also HPC systems need to be designed for large scale data intensive workloads servicing the demands for what is being termed "High Performance Data Analytics" (HPDA). This evolving environment places new demands on how high performance can be delivered in a data-centric and data-intensive model supporting cross-discipline, and hence heterogeneous data-sets, analytics and federated data management all within a flexible and agile infrastructure that is cost-effective. Combining such modalities of operation calls for the convergence of HPC, HPDA and Cloud, which has been the precipitant of the new Research Computing Service at the University for which the RCS is using OpenStack as a means to enable the infrastructure.

This first paper will explore the use of OpenStack as a framework for delivering a range of Research Computing services notably in the bioinformatics arena. A separate paper will explore specifically the case for HPC and OpenStack in terms of the specialized performance and consequential high-performance technologies deployed in this market. Following papers will explore the widening of cloud computing techniques to embrace further scientific disciplines as additional infrastructure is enabled at the University.

# Motivation

This changing HPC landscape is seen across all disciplines, but is most dramatic within the life science and clinical medicine domains. Here we are witnessing an explosion in the uptake of computational techniques combined with a staggering increase in electronic sensor and instrument data: from bio-informatics [2] genomics [3] , medical imaging [4], microscopy [5] and medical informatics [6]. To meet this demand, the Cambridge University Information Service has worked extensively with its users and partners in research and industry to design and implement a new paradigm in HPC system design; mixing traditional HPC elements of performance and throughput with new OpenStack for ease of use and flexibility. This new HPC system is a Hybrid HPC/Cloud solution called the Cambridge 'Bio Cloud'.

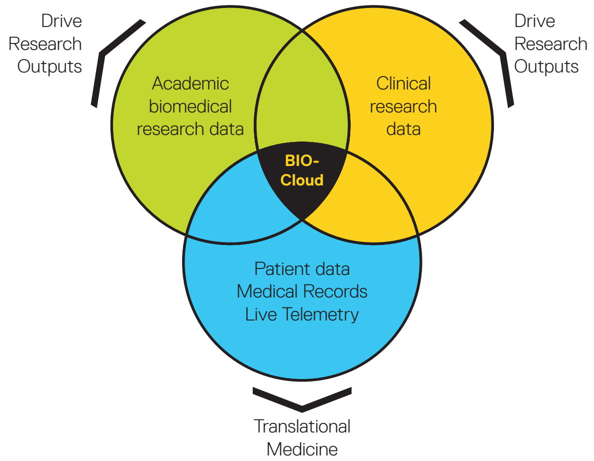


*The Cambridge Bio-Cloud brings together Research Fields and Medical Practitioners to Drive Translational Medicine*

Cloud computing is recognised as an enabling technology for the research community. Superficially it is seen as a method for provisioning resources on demand and in an elastic manner, however more importantly from a collaborative research aspect, it heralds a key solution of bringing the user and their compute to the data thereby allowing secure, multi-tenant, permanent and ephemeral virtualised environments to be created and appropriate service provision enabled. The adoption of a cloud computing model at the University was therefore a key strategy in bringing together multiple disciplines from bio-medical researchers to medical practitioners providing a substrate for new application development, new methods and ultimately new science, easily extensible to both on- as well as off-premise infrastructure.

In addition to the scientific benefits, the adoption of OpenStack is seen as a key to widening the skills base for the delivery of HPC, moving away from the bespoke methods employed today in HPC for provisioning, management and control to a consistent and widely used method [9].

Taking these factors into consideration, a key motivation therefore for the UIS was to provision a cloud that was not only functional in respect of agility and flexibility but that was also oriented for performance and scalability to promote new data analytics methods. Given the experience and skill-set of the RCS in HPC this led to a process of exploring HPC technologies within OpenStack, and in particular RDMA-enabled storage and networking, to provide the cloud.

### Hybrid Cloud - Balancing Capital and Operational Expenditure

Cloud models provide new resourcing options for research compute. A traditional HPC infrastructure is procured, installed and operated in a budgetary model that is heavy on capital expenditure. After a system is commissioned, it becomes inflexible or expensive to adapt it to changing requirements.

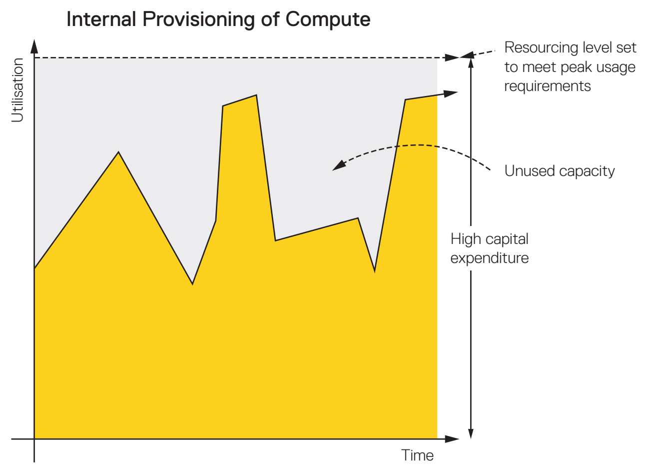


Conversely, migrating all resources to public cloud raises the operational expenditure for the same resource. Public cloud pricing typically is markedly higher than the operating costs for compute resources.

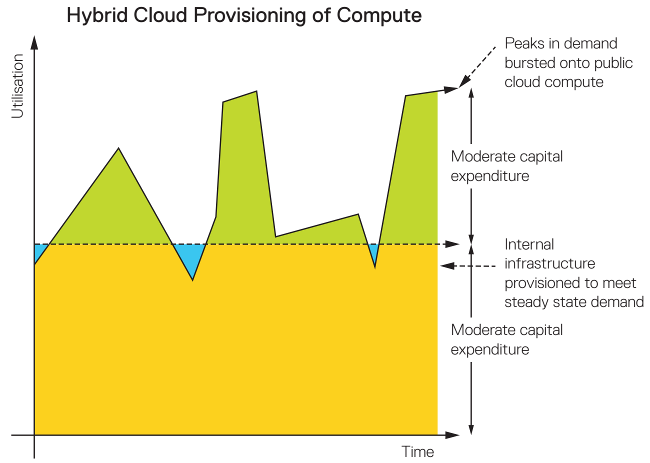


A hybrid cloud solution couples on-site compute infrastructure with public cloud. The on-site infrastructure is deployed as a private cloud, offering similar APIs and capabilities to public cloud. The process of seamlessly transferring new workloads to public cloud compute when the internal resource is fully utilised is known as cloud bursting.

The hybrid cloud approach has the potential to offer increased flexibility and reduced total cost of ownership over other approaches. Over time this can inform the business of capacity planning leading to potentially even more cost reduction.

# What is OpenStack

To realise the private cloud within Cambridge, the RCS has adopted OpenStack. OpenStack is an orchestrator. The OpenStack platform connects and coordinates compute, storage and networking systems to create a coherent environment, managed according to site policies but offering new levels of flexibility to the users it supports. This provides an ability to offer cloud services to a range of workloads/applications within a single framework.

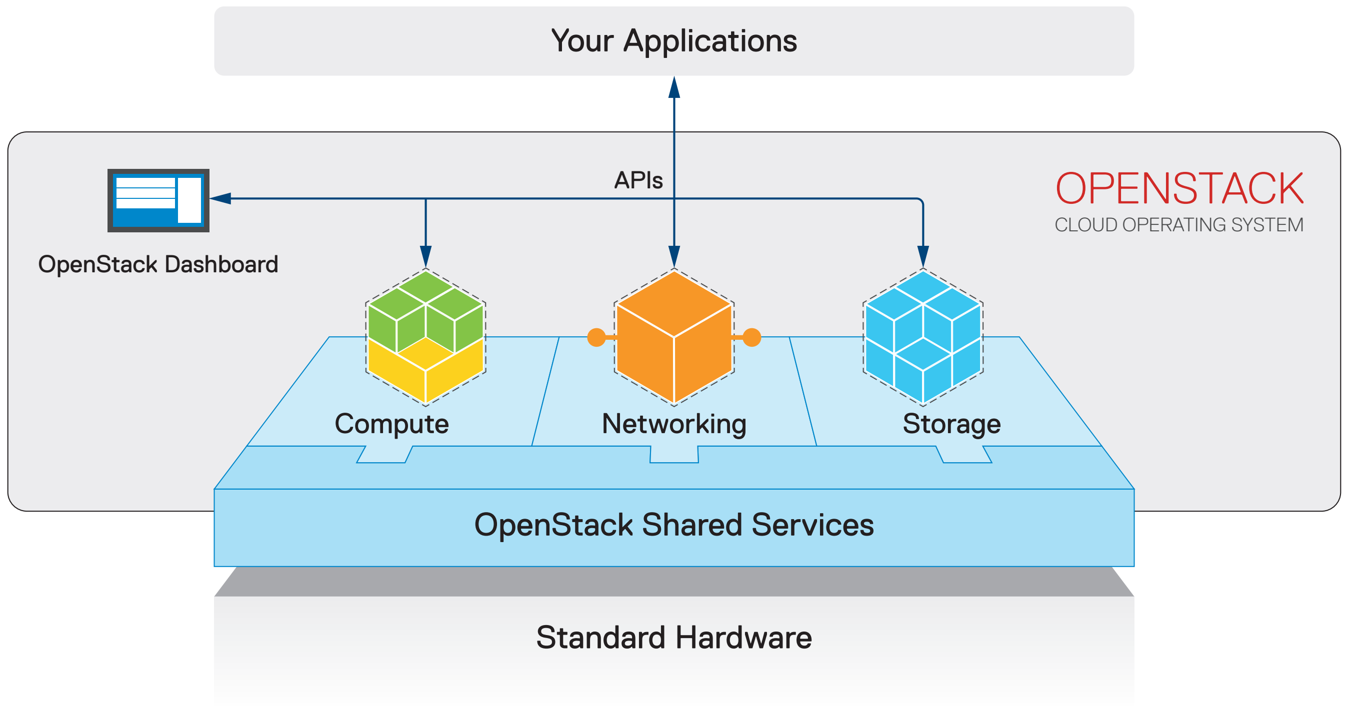


*OpenStack Cloud Environment.*

Cloud computing is built on the concept of software-defined infrastructure. It goes beyond grid computing to add new capabilities and new levels of flexibility. With every new release, OpenStack develops new capabilities and asOpenStack releases are made twice a year the pace of development is very high. With every new release the arguments in favour of using OpenStack to deliver the needs of Research Computing become increasingly attractive. Research Computing can exploit the strengths of software-defined infrastructure to redefine the way in which compute resources are procured, managed and used. The use of OpenStack for creating private cloud infrastructure is driving this paradigm shift.

OpenStack officially became an independent non-profit organisation in September 2012. The OpenStack community, which is overseen by a board of directors, is comprised of many direct and indirect competitors see www.openstack.org

# Community Engagement

A key tactic in the OpenStack strategy at Cambridge was recognition from many discussions with other institutions that we were not alone in this view of the future and therefore it was critical that the research community had the means to articulate its requirements and define use-cases as a single voice, albeit with individual undertones. This, if enabled, would share risk and provide consistency of approach, avoiding the multiplicity of good-enough approaches, and promote development from this consistency. The means to achieving this was thought about for many, many months and eventually catalysed, more by serendipity, than forethought by helping to construct a sushi-inspired (in Tokyo) Scientific Working Group (Scientific-WG) within the OpenStack Foundation, thereby benefiting from the network of OpenStack and providing a means to drive the Science agenda. From small beginnings the Scientific-WG s has now developed into a bona fide stream at community-led events. The progress of the Scientific-WG can be followed here [7].

# OpenStack at University of Cambridge

## Analysis of Solutions

The University Information Services spent some time in consultation with its user base to determine a suitable cloud environment. The major considerations from this requirements capture process was the need to provide flexibility in storage, ease of data-sharing, security, elastic provisioning and ease-of-use. These requirements have been constructed into a potential series of service offerings. These are expounded below:

| Service | Workload and Motivation |
|---|---|
| Research computing and storage (IaaS). | Provides central IT infrastructure, for departmental local IT staff to construct persistent IT services. Local IT staff build and support platforms. Enfranchises local IT staff, drives efficiencies, increases agility |
| Research computing and storage platform as a service (PaaS) | Provides central IT platforms, for research group or departmental local IT staff to build persistent IT services. Local IT staff build and support platforms. Enfranchises research staff and local IT staff, drives efficiencies and increases agility. |
| Application development cloud (IaaS, PaaS) | Bare metal or virtualised provision of dynamic resources for platform or application development and testing. Uptake by local IT staff and research staff Drives innovation, in terms of application development, cloud computing utilisation, infrastructure. |
| Research computing-as-a-service (PaaS, SaaS) | Allowing easy access to virtualised or containerised compute infrastructure, long tail science, new user communities, science portals, can include application software. Easy access, collaboration, sharing methods, workflows data greatly democratising research computing and driving discovery. |
| HPC-as-a-service (SaaS) | HPC cluster, with workload scheduler, MPI workload, parallel efficiency, either shared tenant pool or individual, Private Virtual Cluster, user defined HPC environment, run anywhere Hybrid HPC, customisable, flexible repeatable HPC. |
| HPDA-as-a-service (SaaS) | System architecture optimised for large storage and Hadoop Big Data Frameworks, allow virtual I/O to integrate large data sets, virtualised Spark/Hadoop, data federation. Bring users to the data, can use own environment, portable, flexible. Hadoop is customisable, grow on demand, tenant separation. |

# Openstack Realisation

The RCS has over the past year implemented a proof-of-concept cloud infrastructure based on OpenStack and is currently migrating this to a production-level service. The infrastructure consists of Dell commodity compute and storage servers (comprising Intel Enterprise Lustre, Nexenta and CEPH Software-Defined Storage for persistence, Intel SSD technologies and tape for archive) coupled with high performance Ethernet networking from Mellanox. These component technologies provide a major Medical and Bioinformatics infrastructure that couples agile cloud infrastructure with traditional High Performance Computing technologies at the storage level.

The environment will provide a transformational medical informatics environment to accelerate clinical outcomes using Hadoop and Spark Big Data environments, Machine Learning, traditional HPC simulation activities as well as persistent Web Services for data access and analysis. A reference architecture study is in the process of being written, although a schematic of the implementation is shown below.

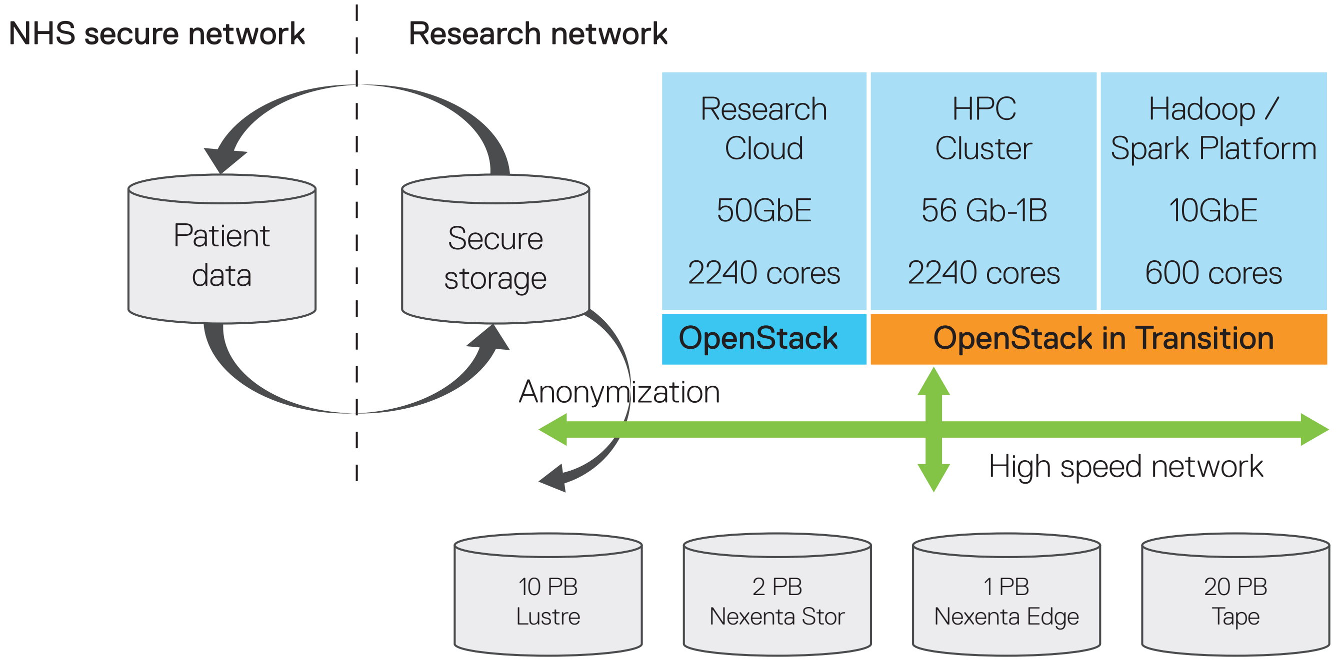


*Overview of the Cambridge Bio-Cloud showing the current OpenStack Cloud and the transition of OpenStack into the HPC and Big Data workloads*

Over time we are developing methods to ensure that the HPC and Big Data platforms will converge into an OpenStack mediated environment. This is being done by engagement with the HPC User Community, the OpenStack Foundation and vendors, together with Cambridge-specific resource. These challenges are discussed here.

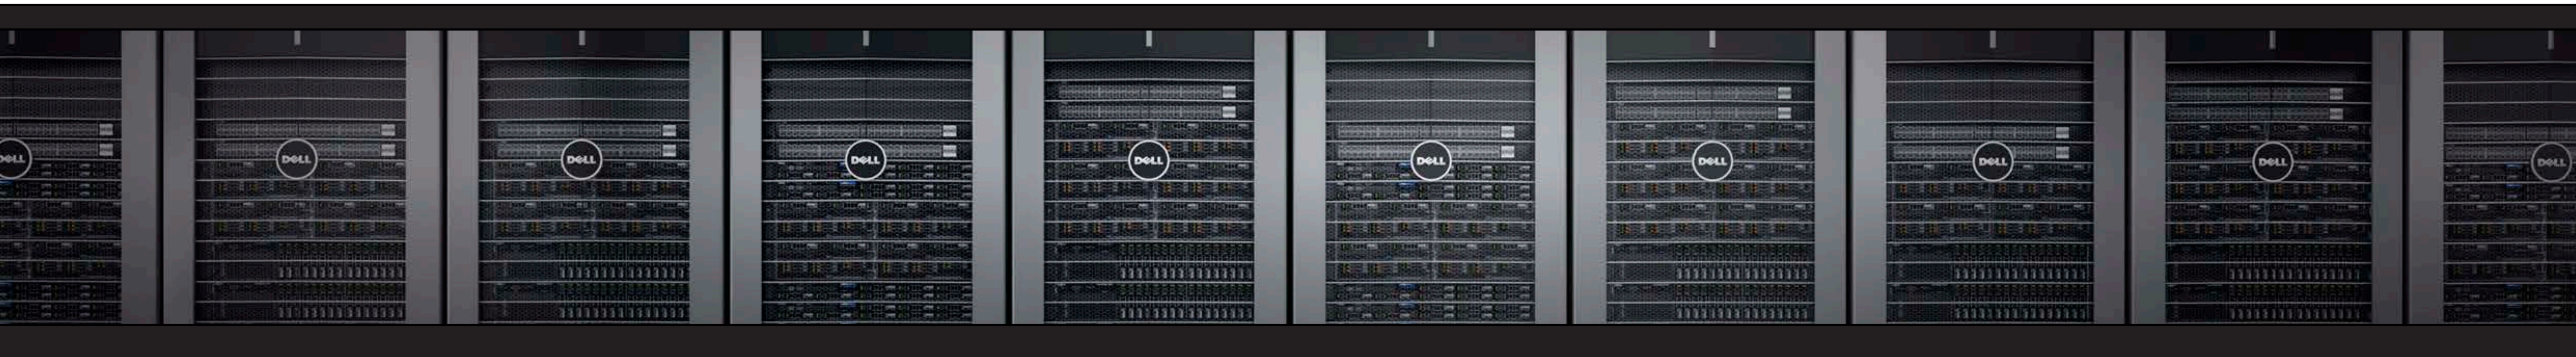

# Hardware Inventory

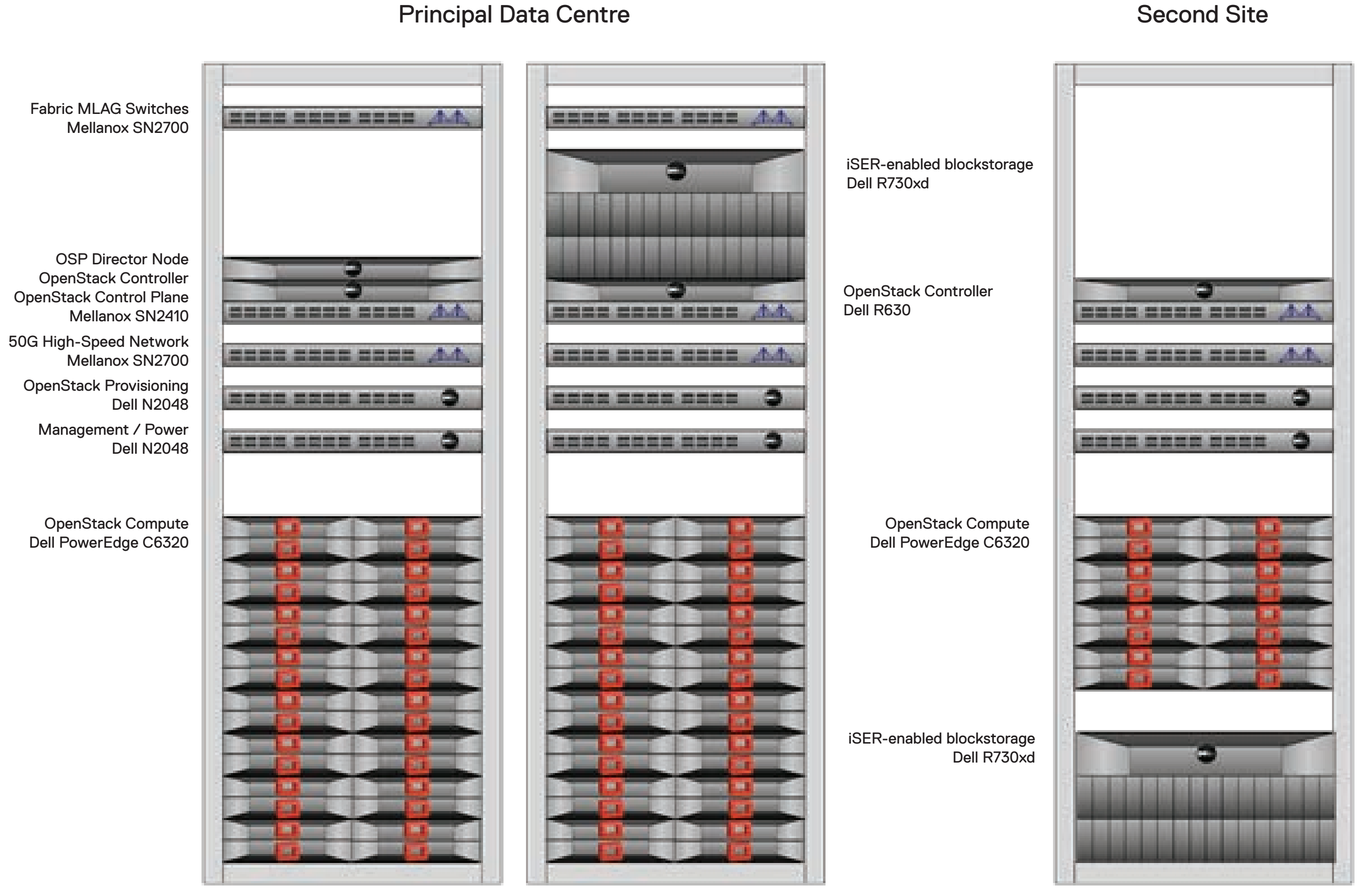


*Physical compute hardware configuration. Nexenta and Ceph storage hardware is not depicted in this layout.*

# OpenStack Challenges in HPC and HPDA

## HPC-as-a-Service (SaaS)

High Performance Computing represents a special case for Research Computing within OpenStack as it currently relies on a number of technologies that lie outside the immediate realm of the OpenStack framework, although they are being rapidly being adopted as future releases emerge. However, the motivation at Cambridge remains consistent. To engage the community to gather use-cases and requirements; to engage with partners and identify amenable technologies that support the approach and help develop configuration blueprints that can be readily supported at the appropriate OpenStack distribution level; whether community or vendor supported. Currently identified threads of development are highlighted here (and are further elucidated here [8]). The currently identified topics are OpenStack and:

- HPC and the Overheads of Virtualization
  - Processor technology developments that assist with the efficiency of virtualization, bare-metal provisioning (including microkernels, LWKs), SR-IOV, trends in hypervisor optimisation, accelerator enablement being Xeon-Phi, GPU, FPGA, etc.
- HPC and Network Fabrics
  - Benefit of RDMA-enabled clouds, configuration of RDMA, VXLAN encapsulation, Infiniband, Intel Omni-Path architechture and RoCE.
- HPC Workload Management
  - SLURM integration, pre-emptible instances.
- HPC and Infrastructure Management
  - Challenges with deploying and managing large-scale but homogenous infrastructure, advantages over traditional HPC system management, support for accelerators.
- HPC and Parallel File Systems coupled with Non-Volatile Memory Technologies
  - Block-based data services iSER, Object Stores, NvMe over Fabrics, Lustre, OrangeFS, GlusterFS, Spectrum Scale, etc.
- Intel Scalable Systems Framework
  - PaaS/SaaS integration of SSF and OpenHPC.

## HPDA-as-a-service (SaaS)

High Performance Data Analytics, in a similar vein to HPC, can span new technologies that do not necessarily fit within OpenStack, today. Research Computing centred around novel techniques in Machine Learning and Big Data platforms seek to extract new insights from complex data involving multiple-disciplinarian teams. Centralising, and minimising data movement is a key enabler in bringing compute to the data where high performance and scalability are paramount for groundbreaking research.

At the University, such new Data Analytical techniques, involving cross-departmental research is a key motivator in exploring Big Data frameworks within OpenStack. Many of the topics are mentioned specifically for HPC transfer to HPDA, with much more focus on underlying storage and memory hierarchies coupled to Big Data environments such as Spark.

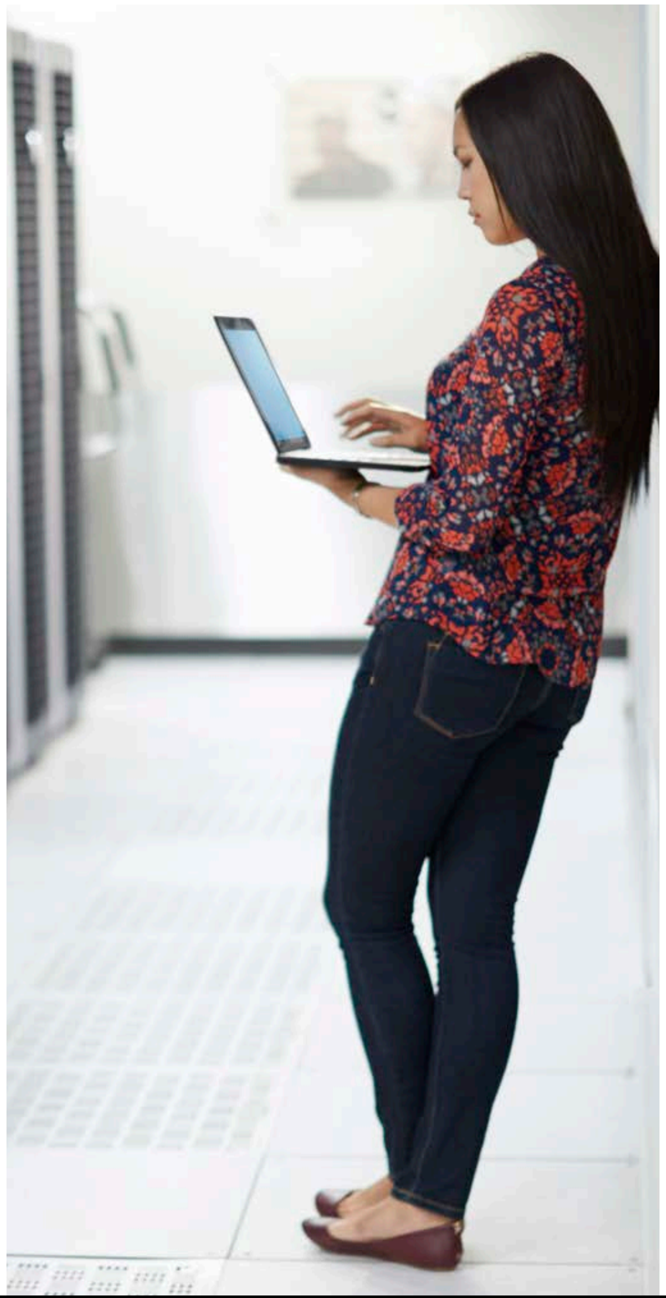

# References


[1] http://jetstream-cloud.org/technology.php

[2] http://www.climb.ac.uk/

[3] http://indico.cern.ch/event/246453/contribution/1566598/attachments/423144/587517/life-sciences-overview.ppt

[4] http://www.wbic.cam.ac.uk/facilities/high-performance-hub-for-clinical-informatics

[5] http://i.dell.com/sites/doccontent/corporate/case-studies/en/Documents/2016-monash-univ-10022600-openstack-cloud-hpc-sds.pdf

[6] https://www.hindawi.com/journals/bmri/2013/658925/

[7] https://wiki.openstack.org/wiki/Scientific_working_group

[8] https://www.amazon.co.uk/Crossroads-Cloud-HPC-Scientific-scientific/dp/1539673987

[9] http://ska-sdp.org/publications/sdp-performance-prototype-platform

[10] https://support.ehelp.edu.au/support/home


# Cambridge University Partners

- Ceph
- Intel
- Lustre
- Mellanox
- Nexenta
- OpenFabrics
- OpenStack
- Red Hat
- StackHPC

Dell Corporation Limited. Registered in England.
Reg. No. 02081369 Dell House, The Boulevard,
Cain Road, Bracknell, Berkshire, RG12 1LF, UK.