Why smart organisations no longer procure storage

Kurt Kiefer - Red Hat Storage
Simon Thompson - Swansea University

Storage in the past

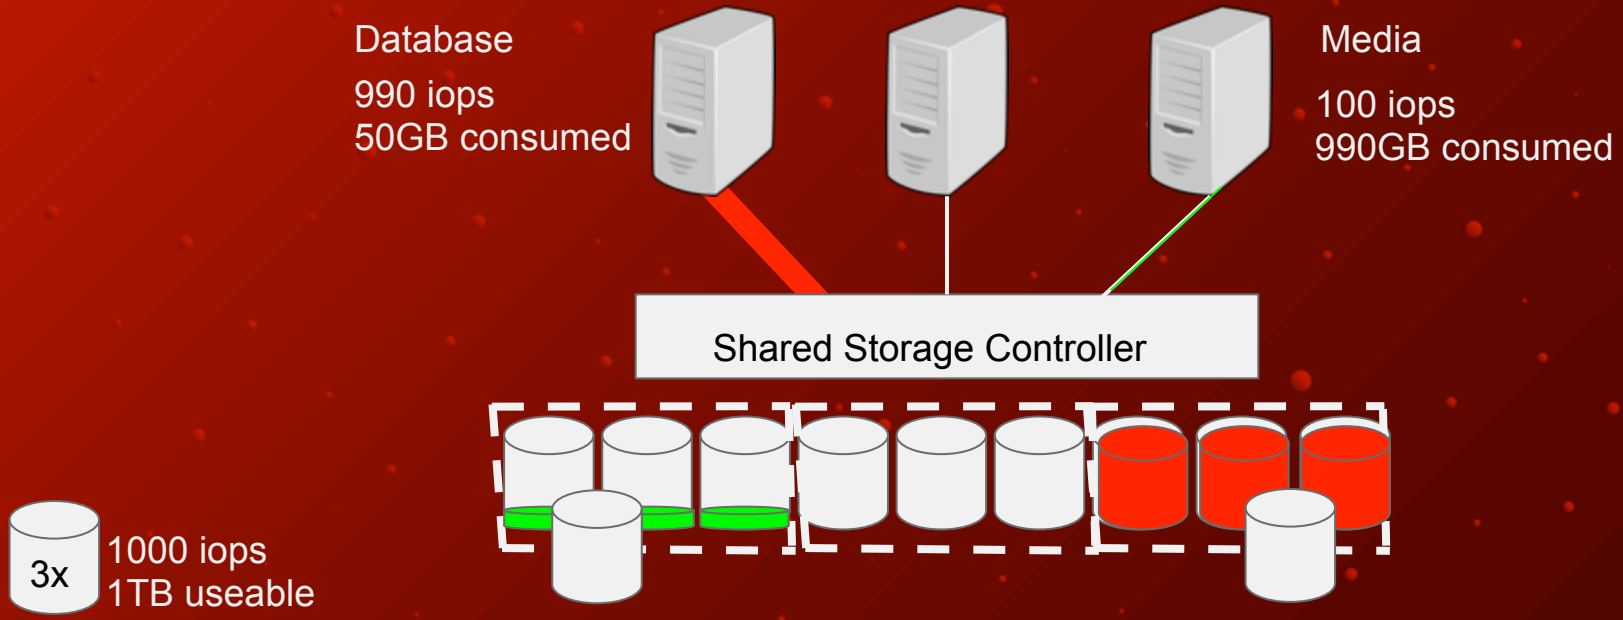

Storage in the present

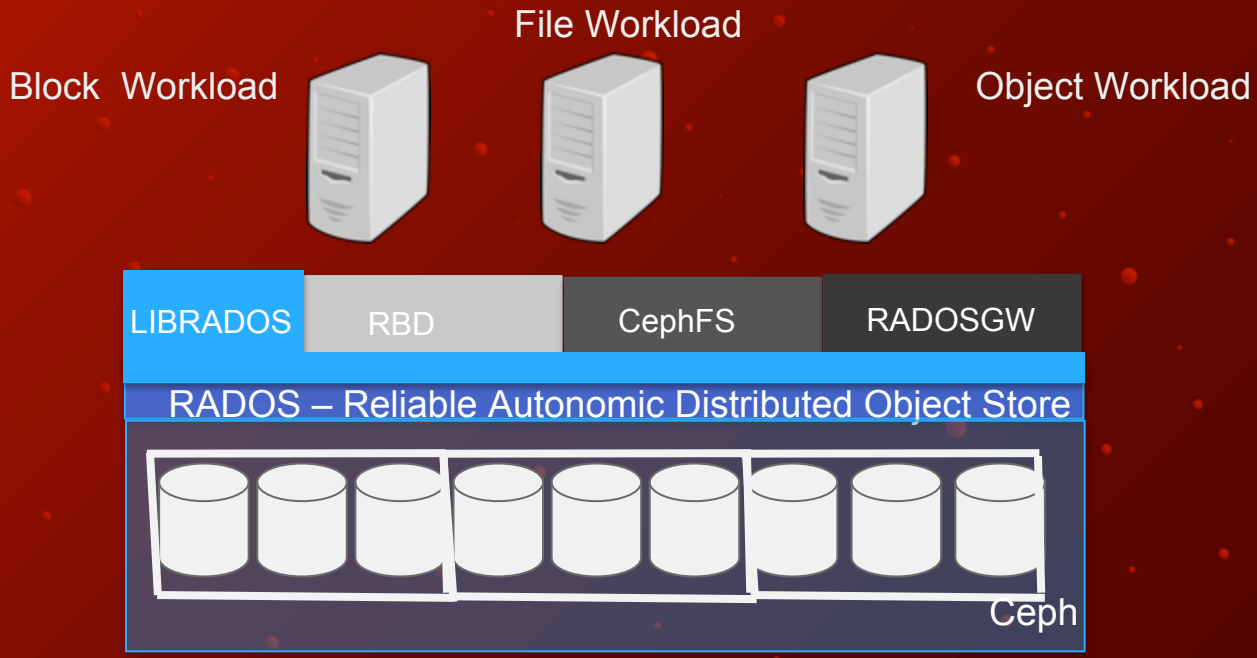

Simon Ellwood-Thompson

Head of Research IT

Swansea University Medical School

The SAIL Databank (SAIL) +

Cloud Infrastructure Microbial Bioinformatics (CLIMB)

+

The UK Secure e-Research Platform (UKSeRP)

~~***The What, The How, The Why***~~

What we have done/doing with RedHat

Background: Coming together of several major initiatives

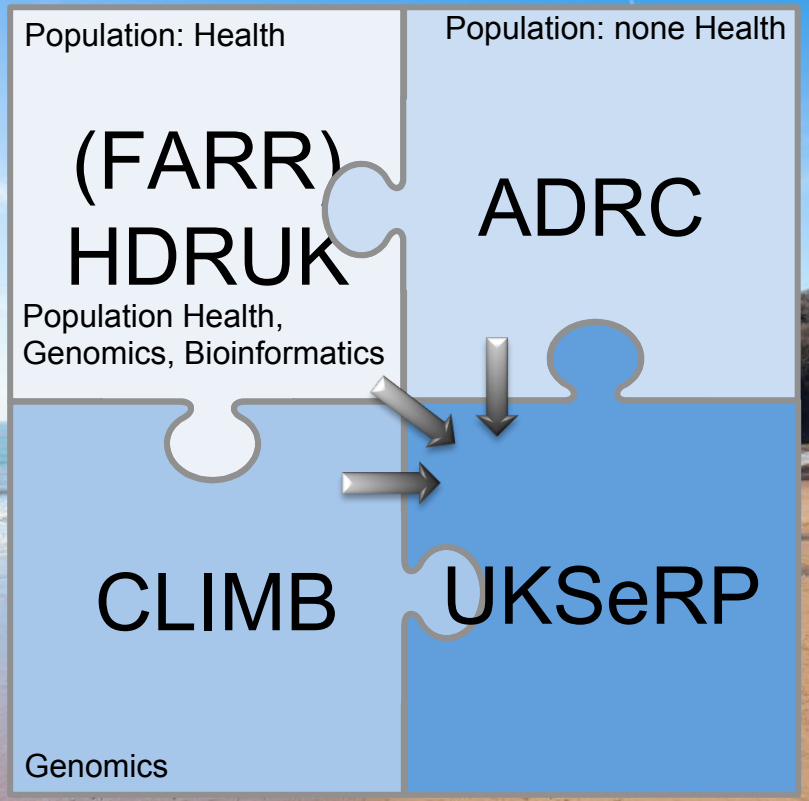

Example: SAIL Databank

- Governance Model and Privacy Protection
  - Research to data not data to researcher
  - Rich collaborative virtual space
- Large Data Collection
  - Lots of health data but other types too
  - No exclusively Welsh data but has all Wales datasets and holdings
- UKSeRP as infrastructure
  - performance & secure remote access
  - Multi Modality + Omics.*, Imaging, NLP, GIS

- Over 26 billion records for >5 million people
- Much data goes back 10-20 years
- All pre-linked data
- >300 users,
- >£200m projects from UKRCs (enable by Swansea)
- 350+ approved SAIL projects, with 152 active today
- 120 staff in Swansea working on Health Informatics related projects

UKSeRP today..

- UK Secure e-Research Platform (UKSeRP) x 12
  - SAIL - Health related person data
  - ADRC - None Health person level data
  - DPUK - 35 dementia cohorts + Imaging + Genomics
  - ALSPAC - From birth cohort, deep phenotype
  - UK Biobank (outcomes) – sub project of Biobank
  - UK MS Register - UK register of people with MS – EHR & PROMS
  - MRC Pathfinder - Mental health platform(s)
  - CLIMB - Microbial Genomics
  - UKCRIS - Mental health unstructured data
  - ELGH - East London Genomics and Health Programme
  - DSB - Collection of smaller projects
  - GOV - Welsh government use

Swansea University Prifysgol Abertawe
Medical School Ysgol Feddygaeth

MRC Medical Research Council

CRIS NETWORK

MS REGISTER

SAIL DATABANK

Llywodraeth Cymru Welsh Government

Dementias Platform UK Medical Research Council

Administrative Data Research Network

biobank UK Imaging study

MRC Cloud Infrastructure for Microbial Bioinformatics

East London Genes & Health

CHILDREN OF THE 90S

CLIMB – plug into the side of UKSeRP https://www.climb.ac.uk

CLIMB is a collaboration between Birmingham, Cardiff, **Swansea** and Warwick Universities, to develop and deploy a world leading cyber infrastructure for microbial bioinformatics.

Provision of high memory or high CPU count VM “servers” for research groups

At Swansea:-
2880 vCPU
50 TB Memory
400 TB fast SAN
1.4 PB data/object storage

Cluster: 4 Identical Sites

Target operating model : 1000 VM’s

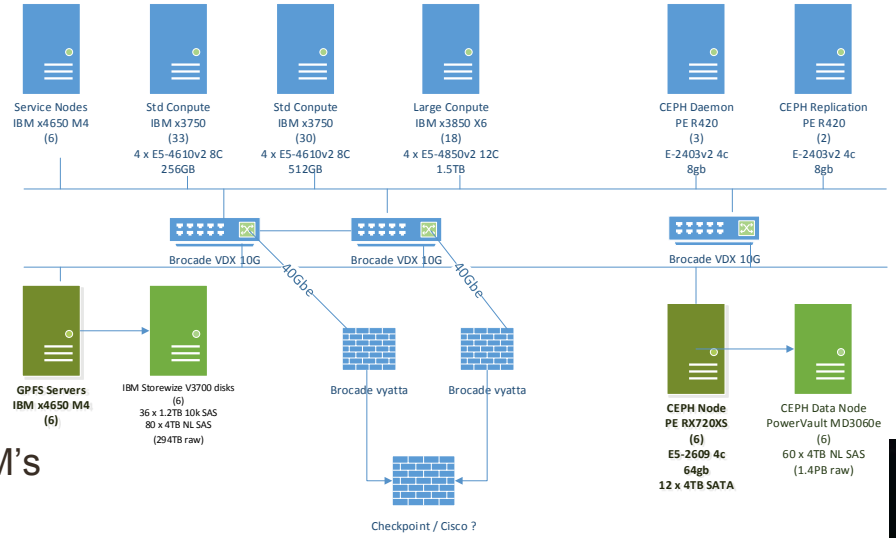

Not traditional HPC - more large resource compute / modest dynamically created clusters

Swansea plan: CLIMB v1.5 for security partitioned extension of CLIMB and expand CEPH by end of 2018

Swansea University Medical School - Journey

Phase I – The Past
Phase II – The Present
Phase III – The near future

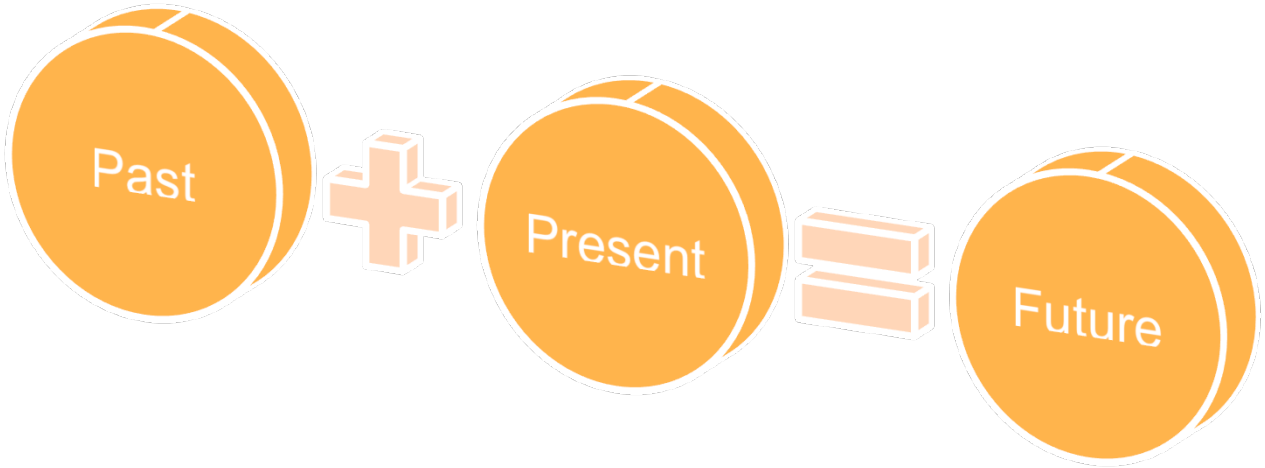

Phase I - CEPH as a storage subsystem

MRC | Cloud Infrastructure for Microbial Bioinformatics

CLIMB Project – Genomic VM's

- Chose to base storage on CEPH (with Redhat support)
- Looking for vender neutral storage
- Horizontally scalable 1.4PB per site
  (4 sites – Swansea, Cardiff, Birmingham, Warwick)
- Supported by Openstack (Cinder)
- Support inter-site replication
- Scripted deployment for identical site configuration(SALT) – Swansea switched to Ansible
- Willing to take risks (sort of – procured some GPFS local SAN just incase)

Per Site
5 x Monitor Nodes
2 x Gateway servers
28 x OSD/Storage

In the end CEPH out performed GPFS SAN's (repurposed)
Very little problems and has recovered from occasional hardware failure very quickly

Phase Ib – Openstack for imaging

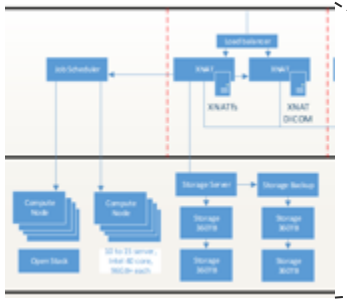

UKSeRP Extension

SWANSEA

Central Hub

Farr

UK Biobank

biobank

Oxford

UCL

Cambs

Edinburgh

Imperial

Newcastle

Cohorts

M'chester

Kings

Cardiff

MRI, PET & MEG

Powered by XNAT

Openstack Cluster

- 10 compute hosts for image analysis pipelines
- Two Dell SAN's – each with 750TB on 180 disks
- SAN based replication between two sites

Phase 2 - CEPH as core storage for main workloads under pinning SAIL Databank, HDRUK, ADRC and UKSeRP

- Use CEPH to end the SAN nightmare and vender lock in.
    - 2 x large IBM, 3 x Fujitsu, 4 x Dell, 1 x HP
- 24 Vmware ESXi hosts (hardware replacement/upgrade)
- Connected by 25GBe Min to storage network
- 2PB over two local data centres on campus, 20% SSD, 80% HHD. A proportion used for workloads needing HA/DR
    - iSCSI block storage as "normal"
    - CEPH object access for VM's supporting this
    - Look at S3 interface for system compatibility
        - such as Research cloud software / FTPS+SFTP using S3 as storage

Phase 2 - CEPH as core storage for main workloads under pinning SAIL Databank, HDRUK, ADRC and UKSeRP

- New GPU servers
    - Small local SSD caching storage
    - Connect to CEPH for main storage subsystem
- New HADOOP cluster
    - Looking at using CEPH as a replacement for HDFS - need to explore more
- New Kubernetes cluster
    - Connect to CEPH for main storage subsystem
- Migrate 2nd Openstack (imaging) cluster to CEPH and retire SAN's
- 10 Host Test and Dev cluster – connect to CEPH for storage

Storage Network – 2PB over two sites

New VMware ESXi infrastructure

Migrate 2nd Openstack cluster from SAN to CEPH

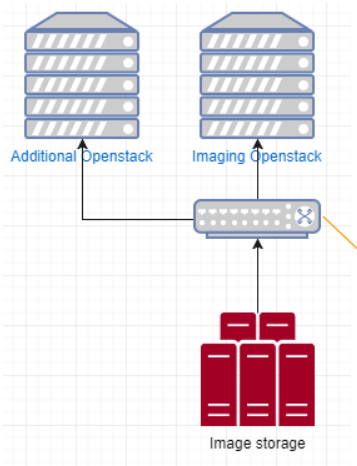

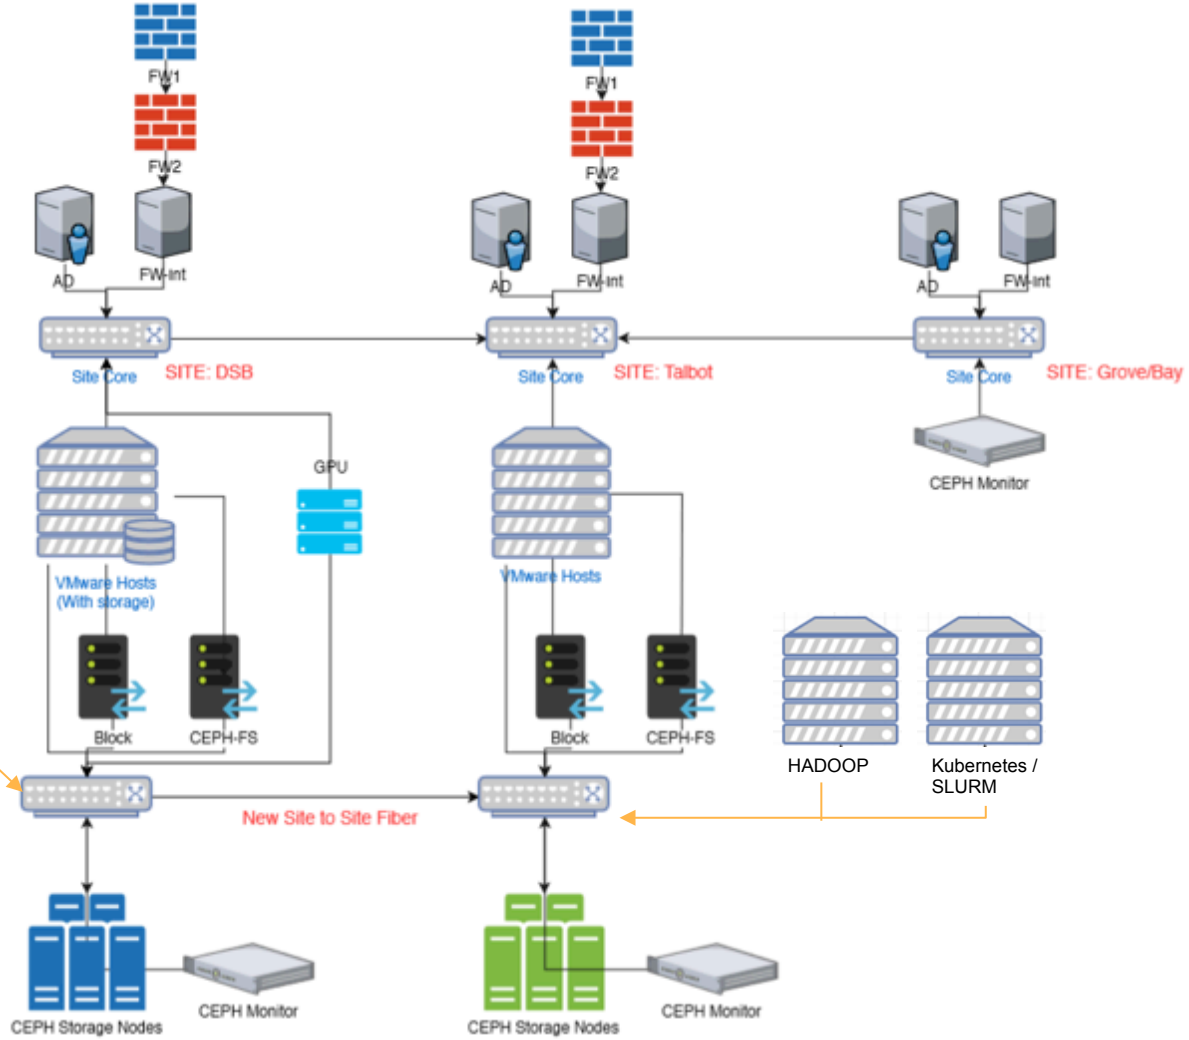

HA/DR vs Capacity

Some storage pools replicating

Others site specific

Phase III - capitalise on other institutions investment in CEPH

- Looking to create secure inter site replication storage pools, based on CEPH, across multiple universities existing systems to support a new genomic project. Allowing each university to bring its unique infrastructure, such as HPC Pipelines to process the data as part of

- Cloud Ready – move to architecture where object storage is primary storage strategy. Although on prem – AWS, Azure, Google – all have object stores

- Ready for significant increase in genomic data storage

It's a Promotion or P45 moment ☺

Data Analytics Evolution

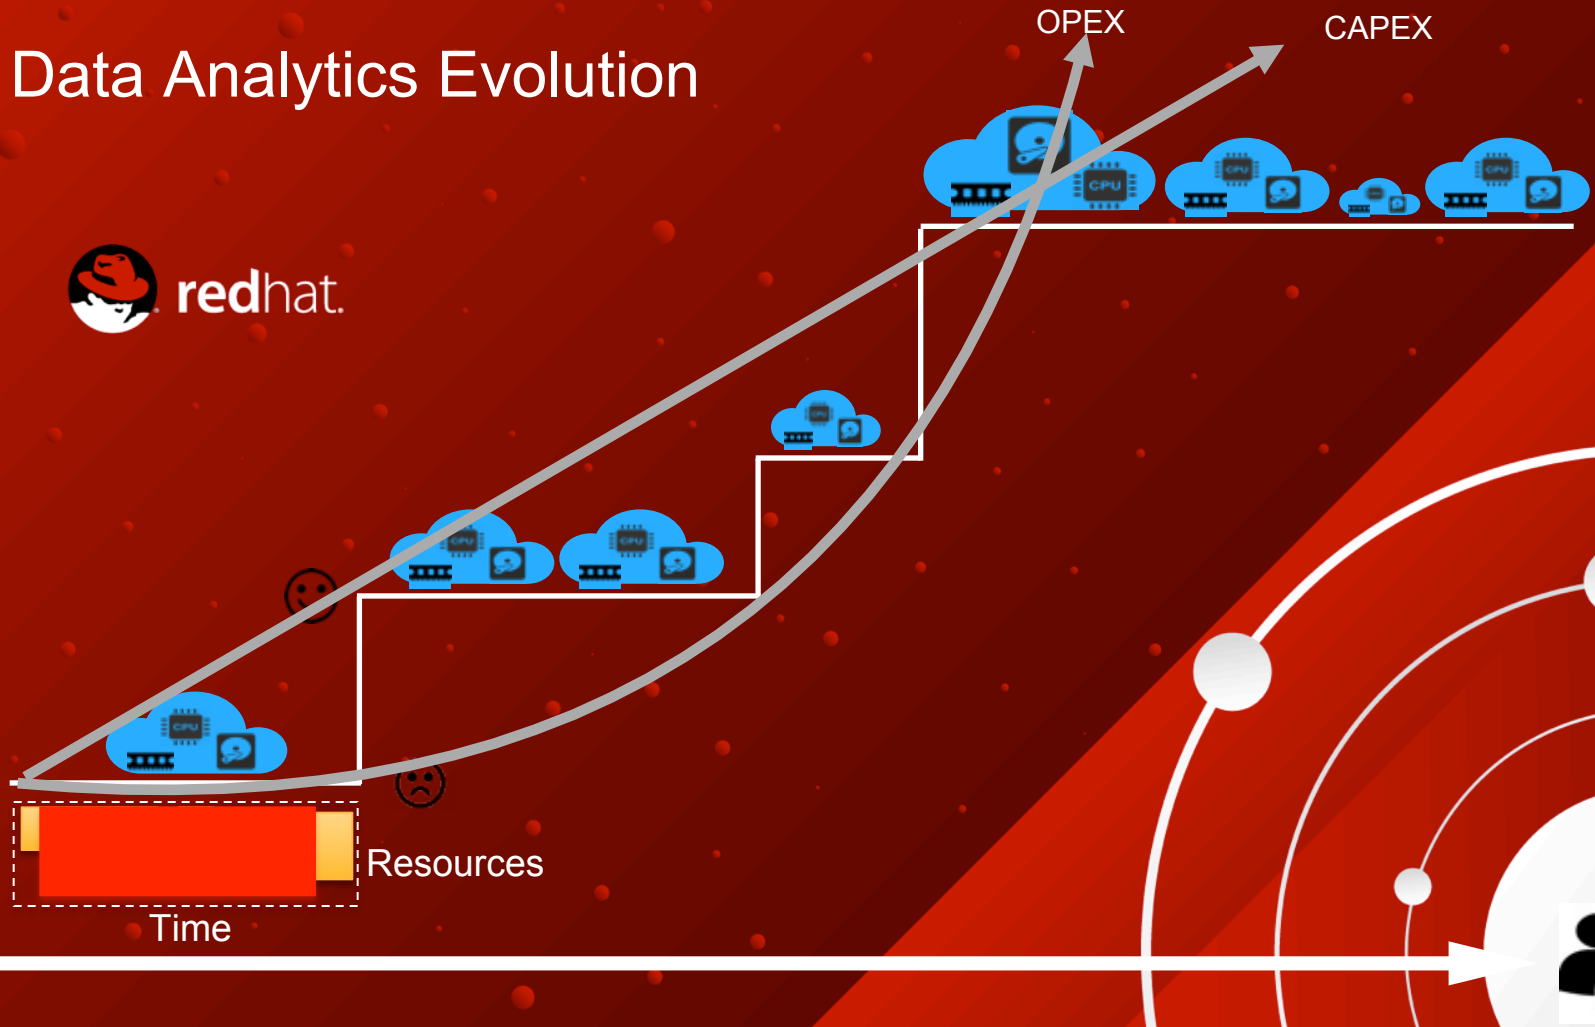

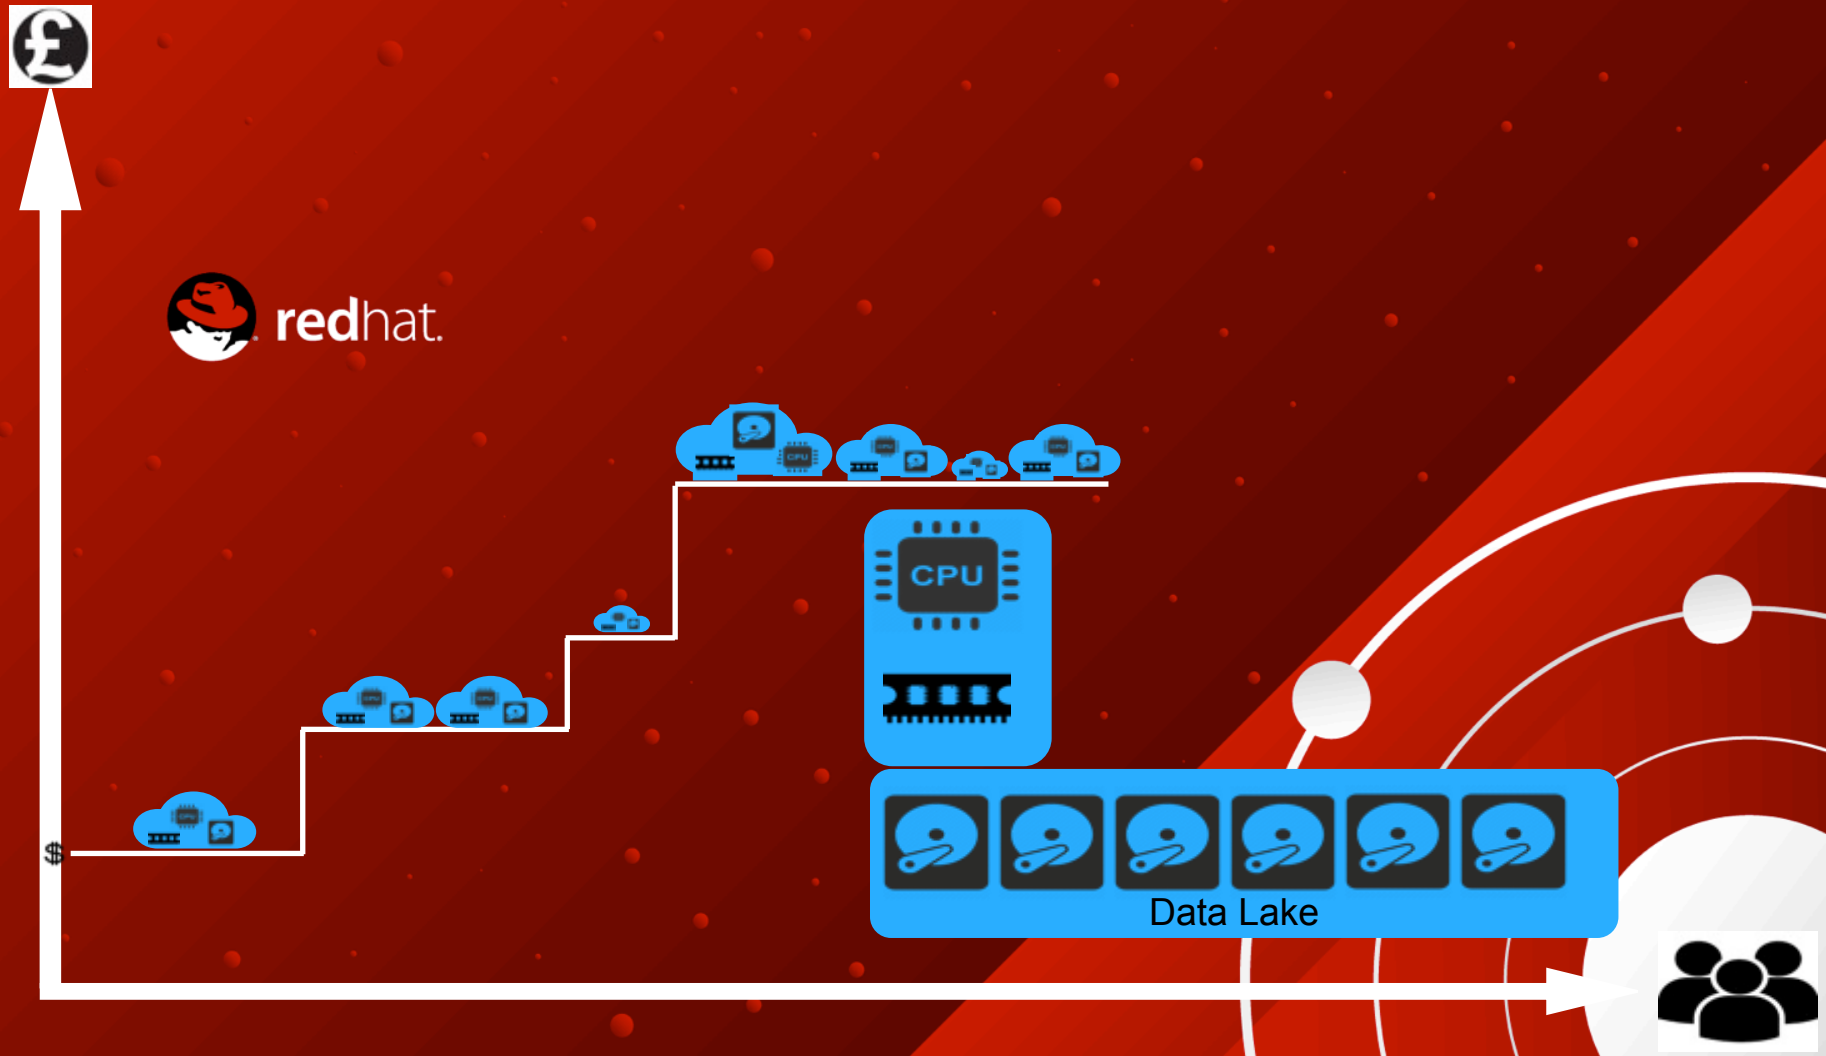
redhat.
CPU
Data Lake

redhat.

OpenStack

CEPH

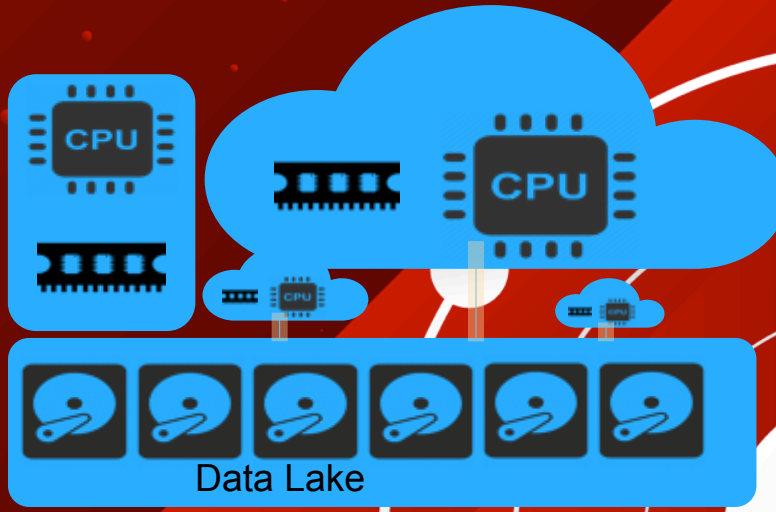

Questions?