Flexible HPC for Bio-informatics

Peter Clapham

Overview

- Overview of the Sanger Institute
- How our data flow works today
- New scientific demands
- Private cloud deployment
- Transitional and future challenges

Scientific Programmes

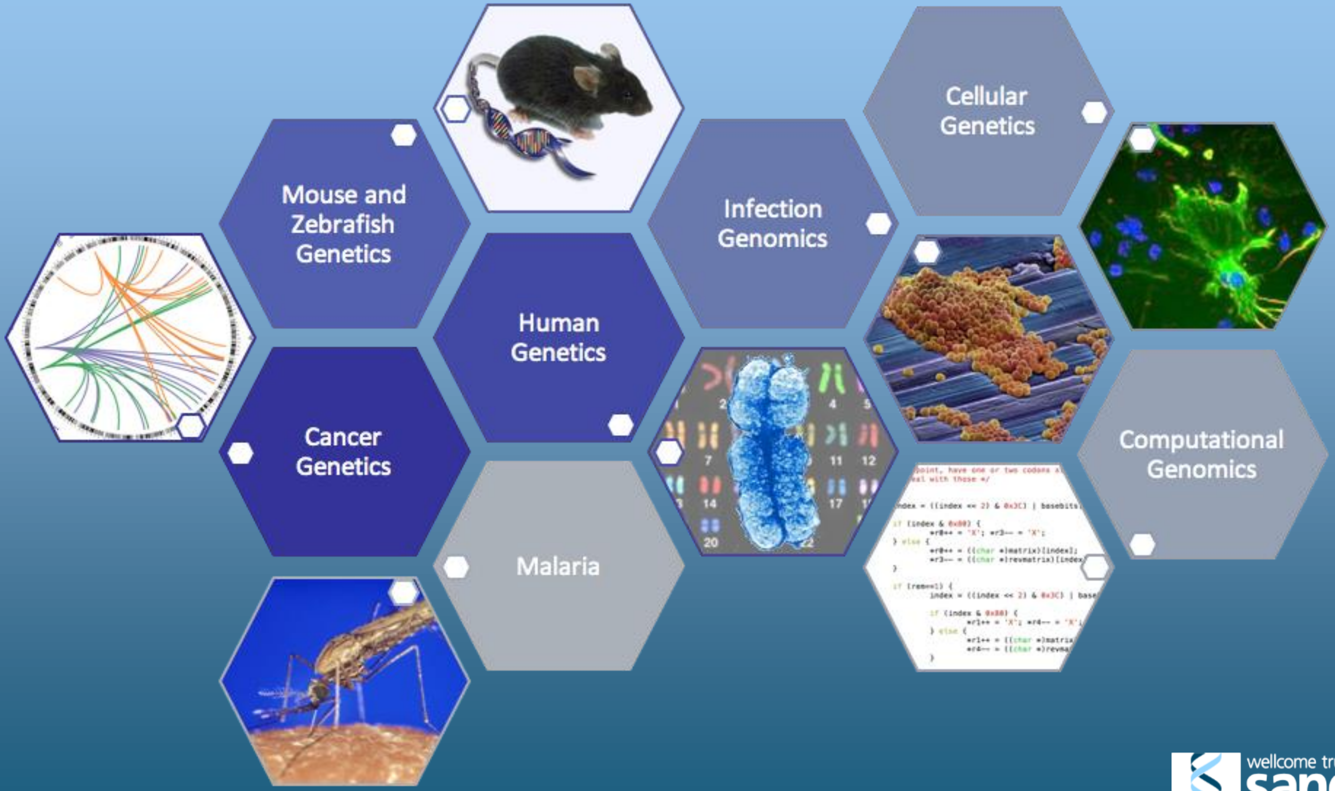

Core Facilities

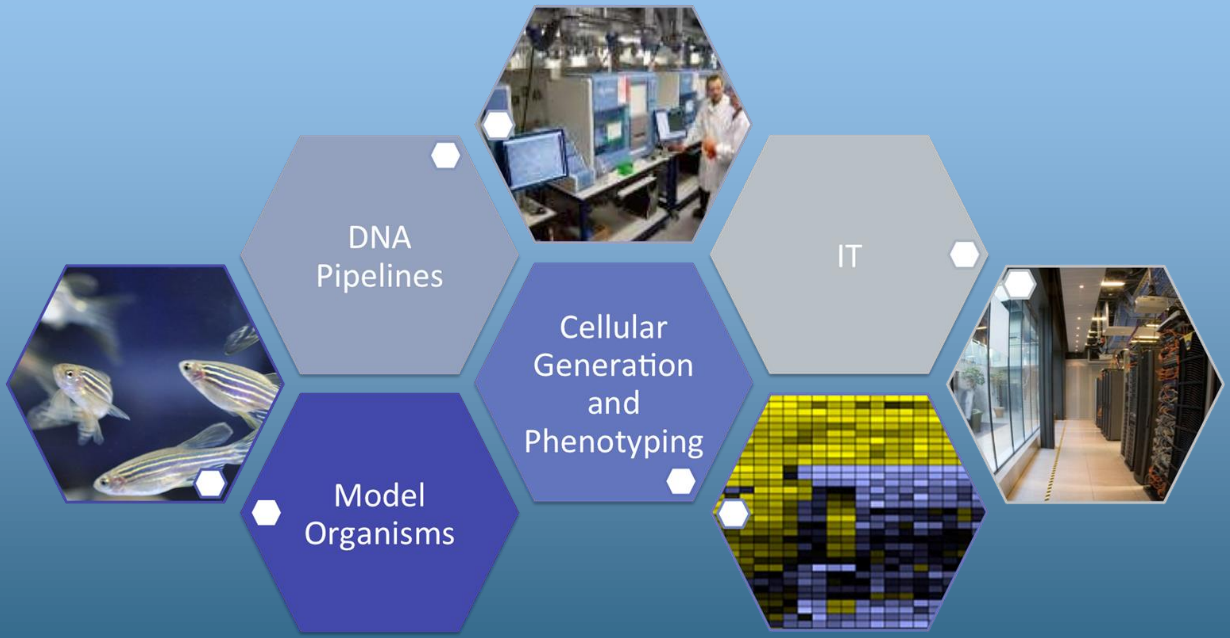

Sequencing Yields

Human Genome, 30-40x Coverage, 120GBases

	 MiniSeq System	 MiSeq Series	 NextSeq Series	 HiSeq Series	 HiSeq X Series*
Key Methods	Amplicon, targeted RNA, small RNA, and targeted gene panel sequencing.	Small genome, amplicon, and targeted gene panel sequencing.	Everyday exome, transcriptome, and targeted resequencing.	Production-scale genome, exome, transcriptome sequencing, and more.	Population- and production-scale whole-genome sequencing.
Maximum Output	7.5 Gb	15 Gb	120 Gb	1500 Gb	1800 Gb
Maximum Reads per Run	25 million	25 million†	400 million	5 billion	6 billion
Maximum Read Length	2 × 150 bp	2 × 300 bp	2 × 150 bp	2 × 150 bp	2 × 150 bp
Run Time	4–24 hours	4–55 hours	12–30 hours	<1–3.5 days (HiSeq 3000/HiSeq 4000) 7 hours–6 days (HiSeq 2500)	<3 days
Benchtop Sequencer	Yes	Yes	Yes	No	No

http://www.illumina.com/systems/sequencing.html

Genome-wide association studies

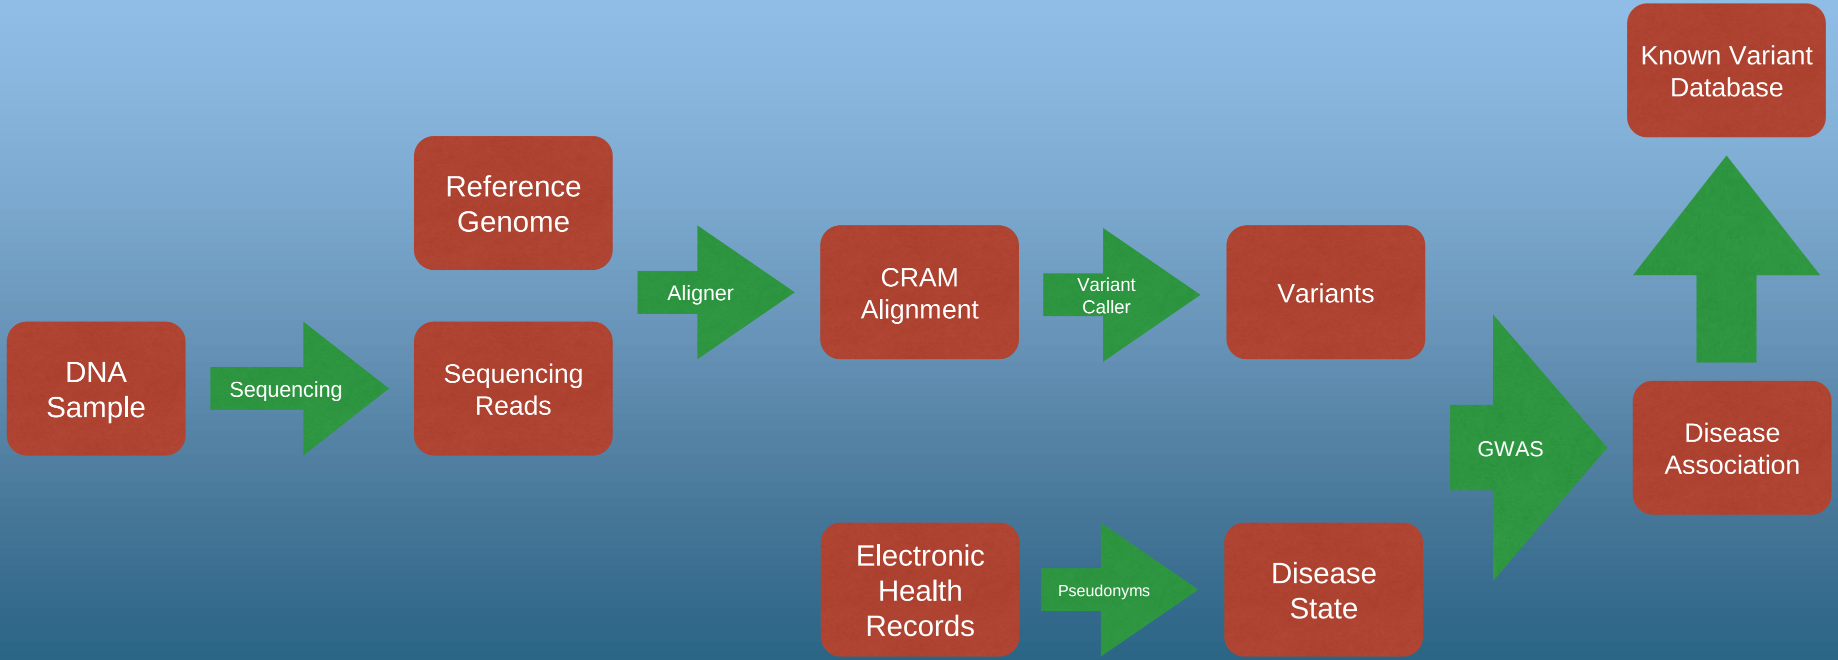

Sanger NGS Data Flow: IT View

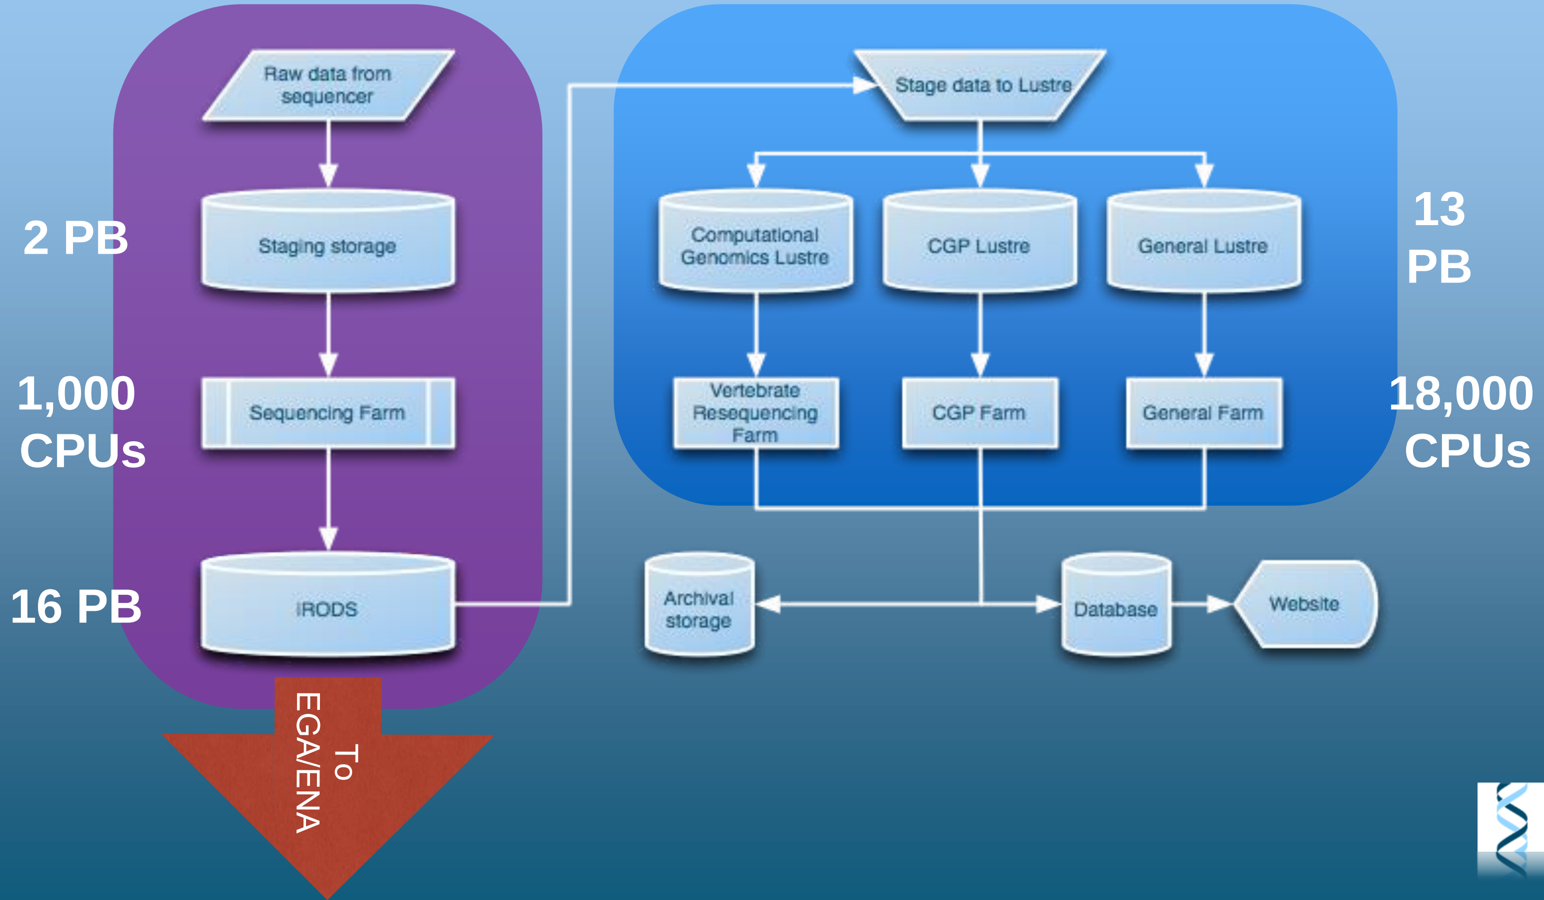

Increasing Collaboration

- Academic collaborations
  - Ensembl (EMBL-EBI)
  - Pan-cancer (OICR, Broad, DKFZ etc)
  - Global Alliance for Genomics & Health
  - eMedLab (Crick, UCL, EMBL-EBI *et al*)
  - FARR Institute
- Industrial collaborations
  - OpenTargets (GSK, EMBL-EBI, Biogen, Sanger)
  - AstraZeneca
- Spin-out companies
  - Congenica
  - Mouse Colony Management System
  - … up to 30 others
- Clinical resources
  - DECIPHER
  - COSMIC
  - Pathogen Surveillance

GA4GH

- Federated data analysis
- Federated AAAI
- Keep data secure
- Open Standard data formats
- Open Standard APIs
- Freedom within the VM

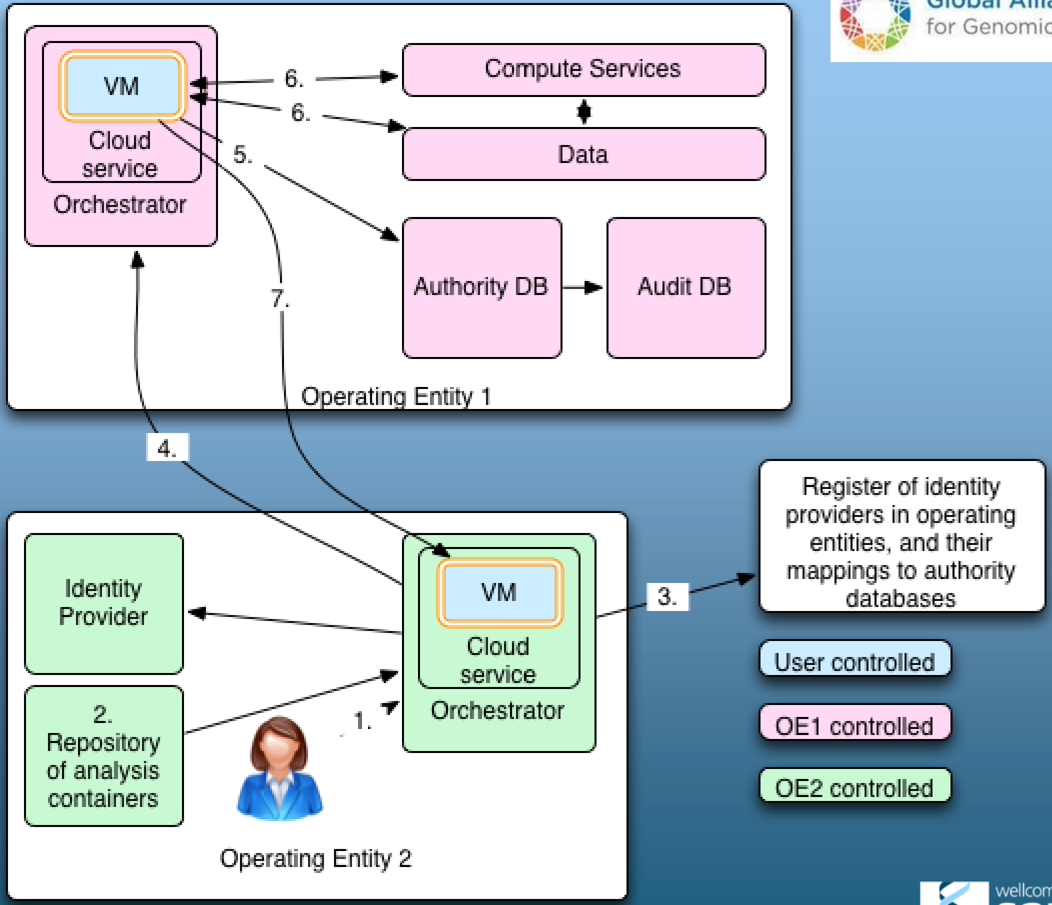

Wellcome Genome Campus Vision

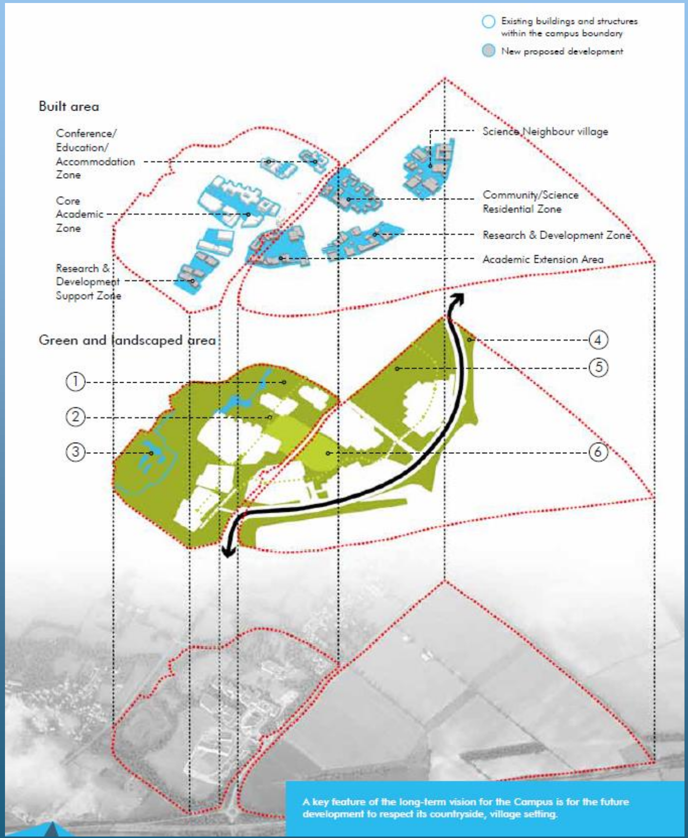

Our vision is to achieve extraordinary improvements in human and animal health.

Our mission is to support the brightest minds in biomedical research and the medical humanities.

Our funding focuses on:
1. Supporting outstanding researchers
2. Accelerating the application of research
3. Exploring medicine in historical and cultural contexts

— Wellcome Trust Strategic Plan 2010-2020

wellcome trust sanger institute

Biodata Innovation Centre

Ready in July 2016

wellcome trust sanger institute

Pressures in delivery

- Staff time
- Flexibility
  - Time to deployment
  - No self-service
  - Admin rights and agile development
- Security
- Lustre
  - Labour intensive
  - Not multi-tenant
  - Not well suited to start-up scale
- Limited monitoring of resource utilisation

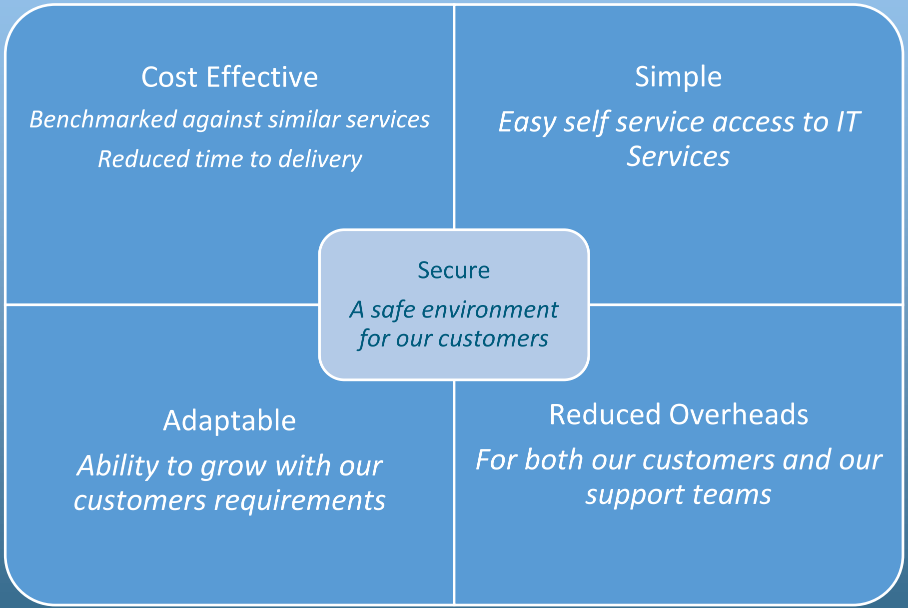

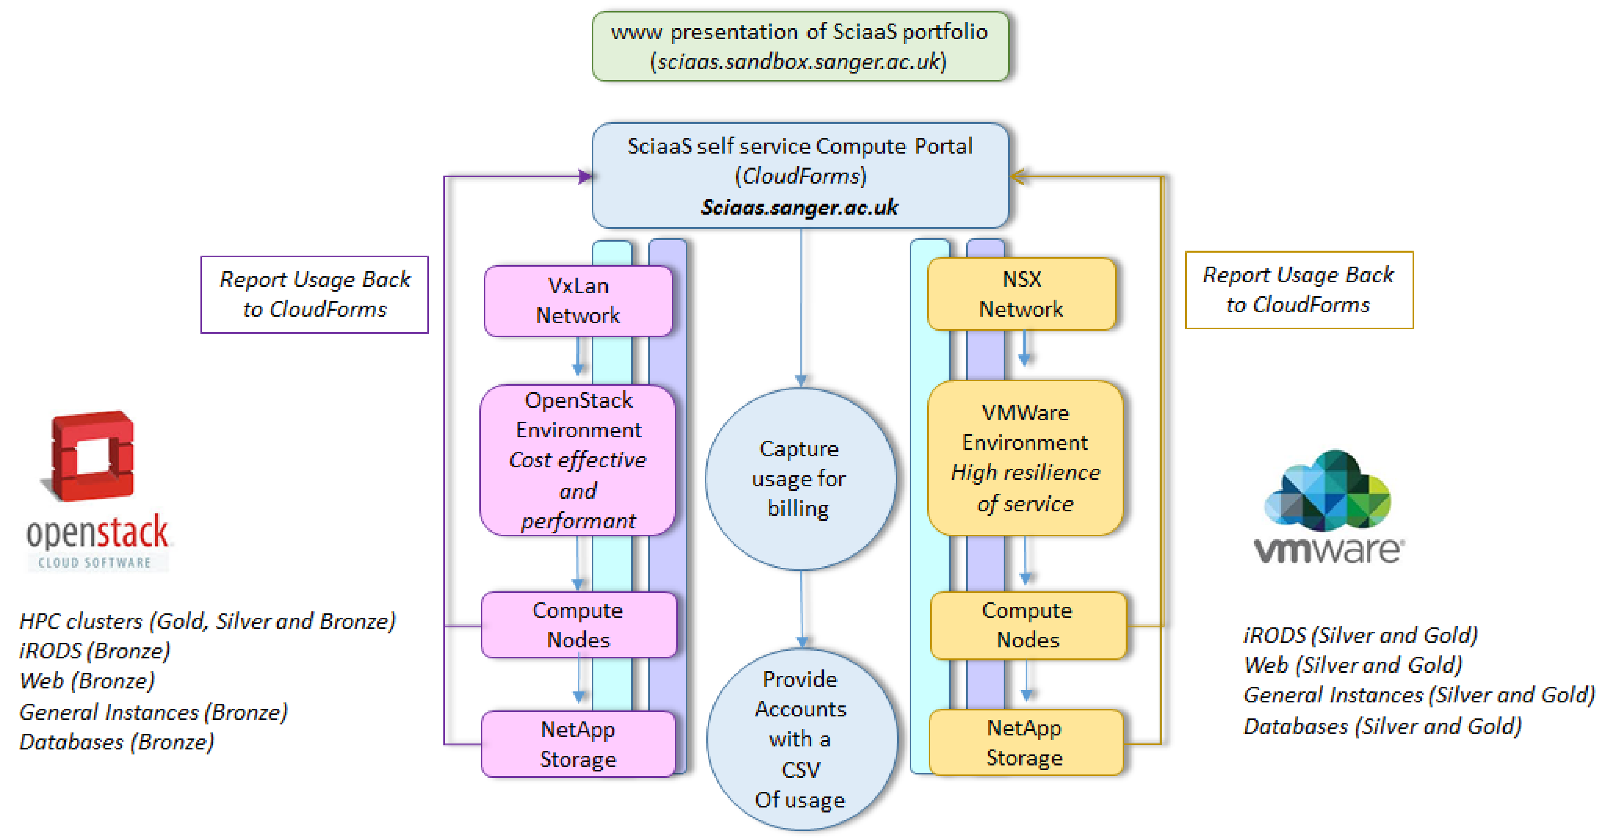
www presentation of SciaaS portfolio
(sciaas.sandbox.sanger.ac.uk)
SciaaS self service Compute Portal
(CloudForms)
Sciaas.sanger.ac.uk
Report Usage Back to CloudForms
VxLan Network
OpenStack Environment
Cost effective and performant
Compute Nodes
NetApp Storage
Capture usage for billing
Provide Accounts with a CSV Of usage
NSX Network
VMWare Environment
High resilience of service
Compute Nodes
NetApp Storage
Report Usage Back to CloudForms
openstack
CLOUD SOFTWARE
HPC clusters (Gold, Silver and Bronze)
iRODS (Bronze)
Web (Bronze)
General Instances (Bronze)
Databases (Bronze)
vmware
iRODS (Silver and Gold)
Web (Silver and Gold)
General Instances (Silver and Gold)
Databases (Silver and Gold)

Multi-tenant and clinical NGS on a common platform

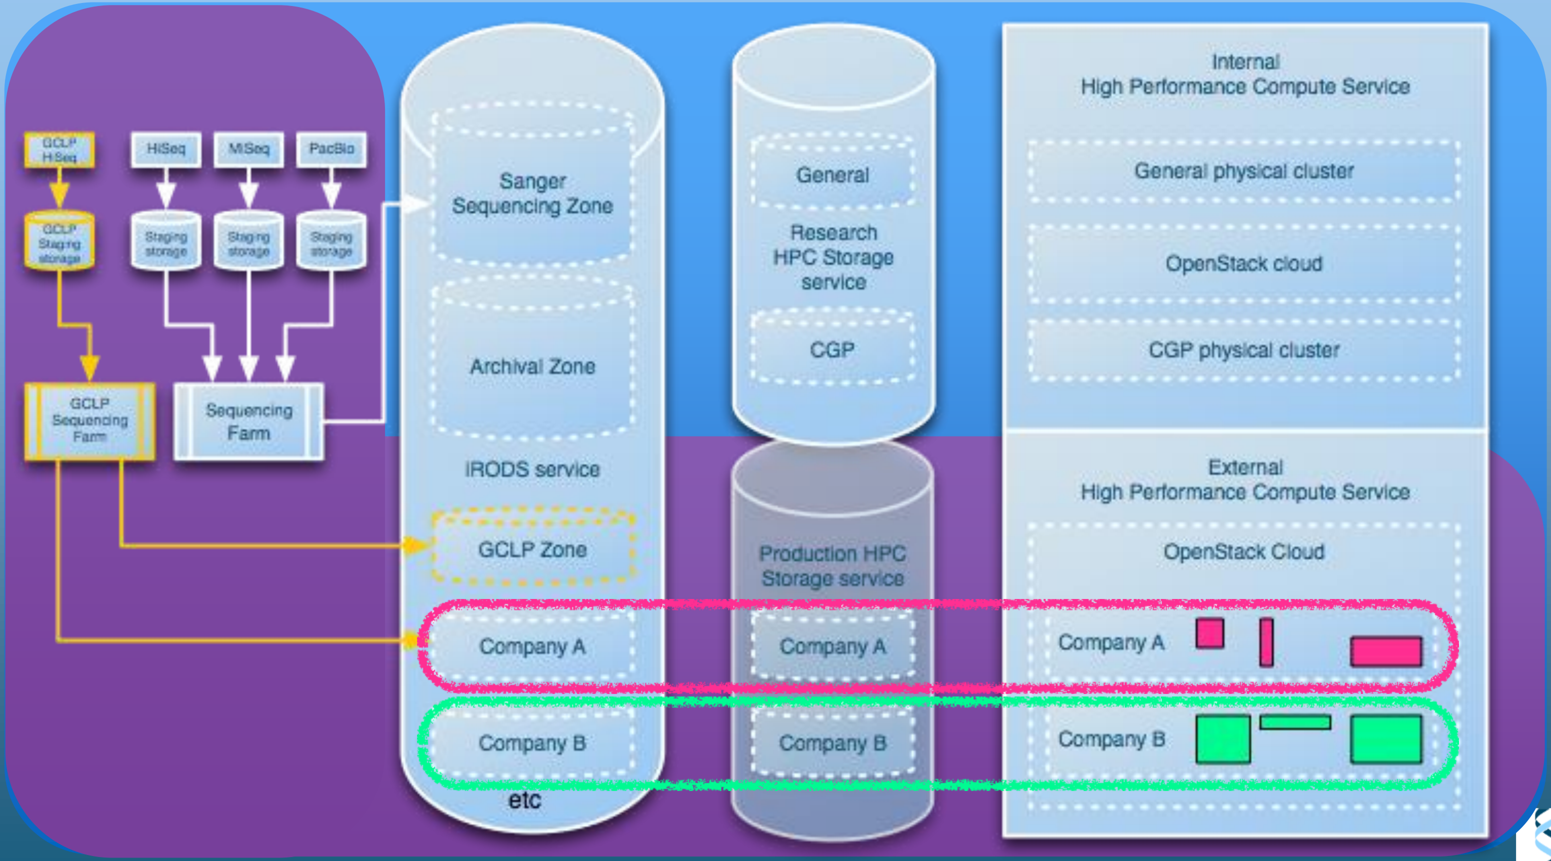

Starting with the basics

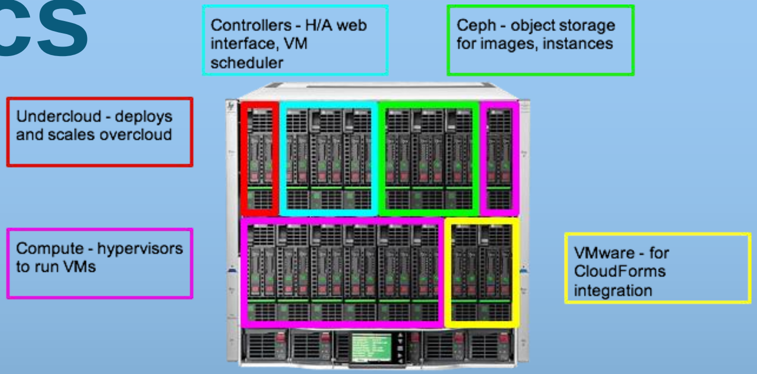

- **Red Hat OSP 7 (Kilo)**
- **Following Best Practices throughout**
- **Red Hat CloudForms for common orchestration of vSphere and OpenStack**
- **OpenIPA identity management with 2-factor authentication**
- **Build more complex services later**

Initial Service Catalogue

HPC
- HPC Cluster
- iRODS Data-management
- Standardised Linux Systems

DBA
- MySQL Environment
- PostgreSQL Environment

Web
- Drupal Environment
- WordPress Environment

Administrivia
- Recover Deleted Instances

Experience so far.

✔ Training first

Service Development
- ✔ Pilot up and running
- ✔ dev/test platforms
- ✔ Cluster in a box
- ✘ Licensing awkward
- ✔✘ OpenStack lifecycle
- ✘ Legacy migration
- ✔✘ CloudForms
  - Complex: training
  - VMware integration tricky
  - Not very multi-tenant (yet)

Security
- ✔ Secure tenant networks

✔✘ Documentation

✔ Support

Resilience
- ✔ Automated rebuilds

- Managing user expectations

wellcome trust sanger institute

Everything as code

Pros

- Consistency of approach
- Central management
- Audit trail
- Visibility of build status and testing

Cons (or at least, things to be aware of)

- Cultural change
- Time to set up needed infrastructure and processes

Continuous Integration and Image Management

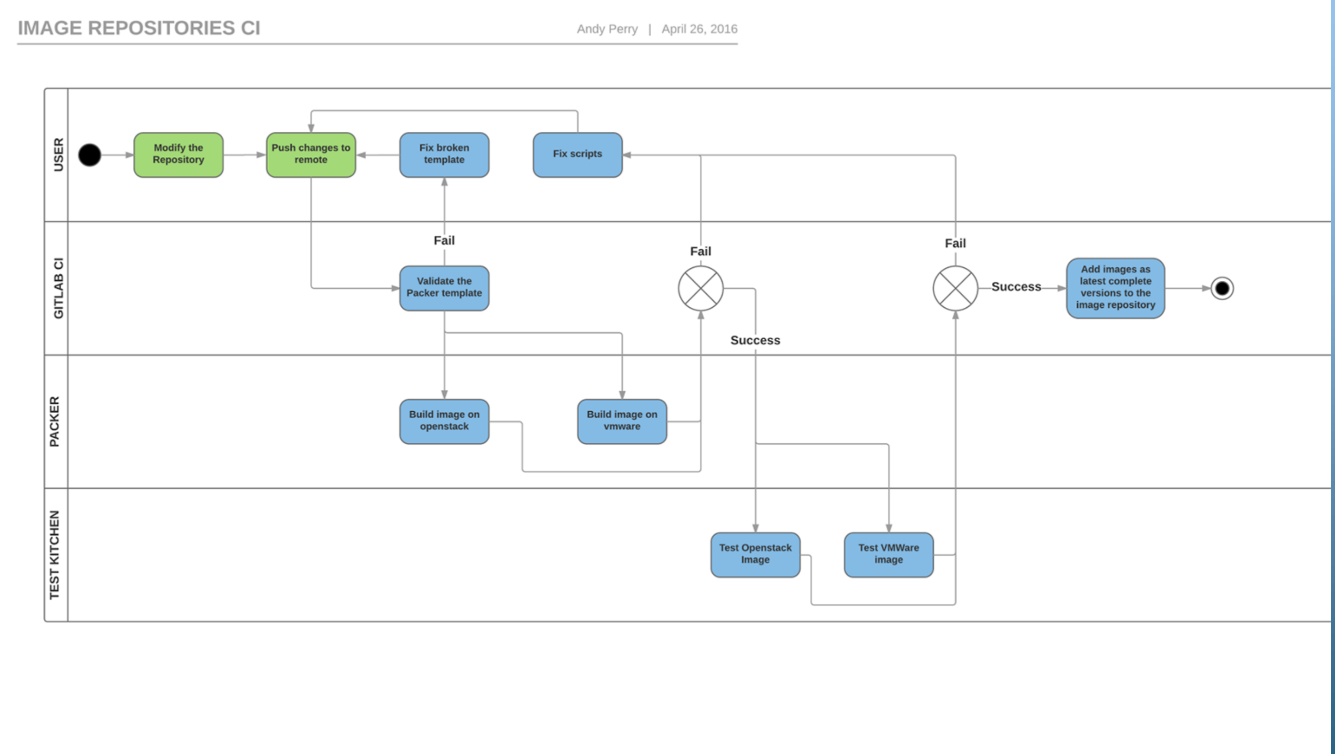

We now have

A fully codified method of creating and managing images for virtualisation platforms.

This provides:

- Providence of images provided
- Simple image management
- Ease of updates and maintenance over time
  - Particularly important for cattle over pets.
  - Reduced time for new image delivery
  - Consistency of builds and test systems as we look to develop our environment

Ongoing Challenges and transition

- Transition for services will take time.
  POSIX won’t go away overnight.
- As scale increases, failure becomes a certainty.
  services need to cope with component failures.
- HPC and security are strange bedfellows.
- Absolute freedom and absolute security even more so.

SDN is a key component of IaaS as we move forward, but what about storage and other supporting infrastructure components ?

Changing future toolsets

Area	Technology today	Limitations	The future
Compute	Linux clusters	Deployment weeks/months No multi-tenancy	OpenStack
Networking	Traditional, manually configured switches	Inflexible Time to delivery	Software Defined Networking: OpenStack, OpenFlow
Workload management	IBM Platform LSF	Focuses on batch work	???
Scientific pipelines	Embarrassingly parallel Single threaded	Fragility limits scale	Move to cloud-style?
Data formats	BAM → CRAM	Space hungry Analysis scaling problems	Research into better compression Graph assemblies, BWT, columnar formats
High performance storage	Lustre	No multi-tenancy Scale has to be large	???
General purpose NAS storage	Enterprise NAS	Expensive Limited multi-tenancy	Object storage? Cinder ?
Archival storage	POSIX filesystem on HP SL4540	Manageability	Object storage?
Data management	iRODS 3.3	No SSL, No atomic put	iRODS 4.1.x
Testing	Manual	Needs automation	Continuous integration
Configuration management	CFengine, git	Needs automation	Ansible, gitlab
Node deployment	Debian FAI	Deploys the OS only	OpenStack Heat/Ironic vagrant, packer

Storage, the missing magic

Providing access to secure shared filesystems
Needs to become:

- Lightweight
- Transparent for end users
- Low administration cost
- ACL’s via AAA and or strong OS control mechanisms
- Performant
- Cost effective
- Support POSIX (at least for now)

So fast, reliable, cheap and now secure by design.