iRODS and Data Management at..

wellcome trust
sanger
institute

Jon Nicholson
jjn@sanger.ac.uk

In this episode…

- **The story so far**
- **Overview of the Sanger Institute**
- **How our data flow works today**
- **How we use iRODS to manage petabytes of genomics data with diverse scientific needs**
- **Safeguarding our data**
- **Expanding our iRODS infrastructure**
- **Future works….**

wellcome trust sanger institute

Background to the Sanger Institute

wellcome trust sanger institute

Scientific Programmes

Mouse and Zebrafish Genetics

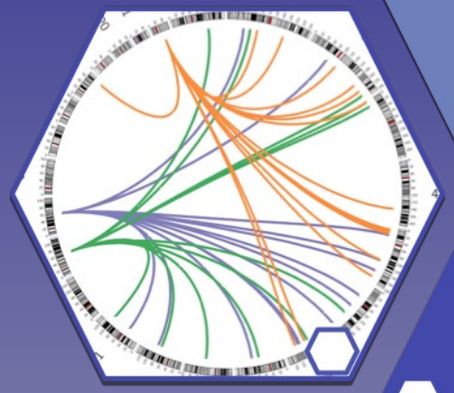

Human Genetics

Infection Genomics

Cellular Genetics

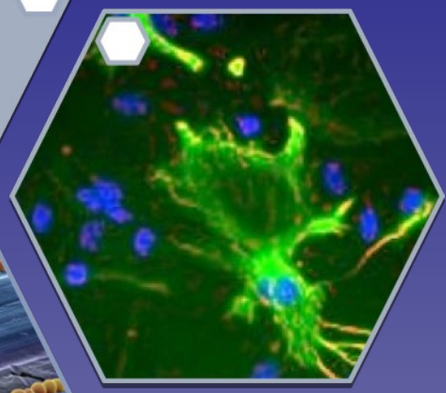

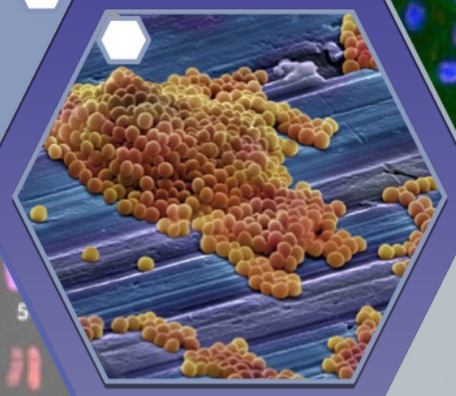

Cancer Genetics

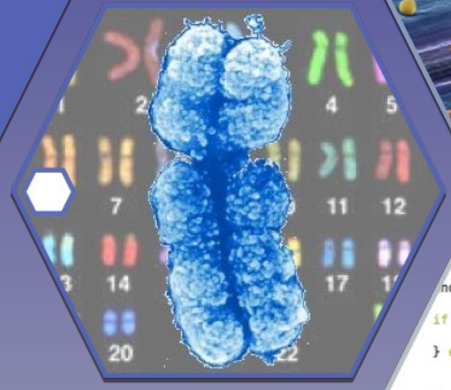

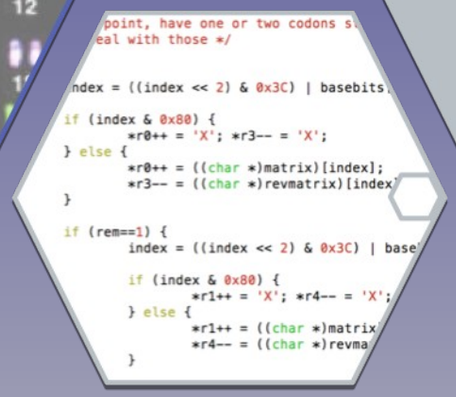

Computational Genomics

wellcome trust sanger institute

Core Facilities

DNA Pipelines

Cellular Generation and Phenotyping

IT

Model Organisms

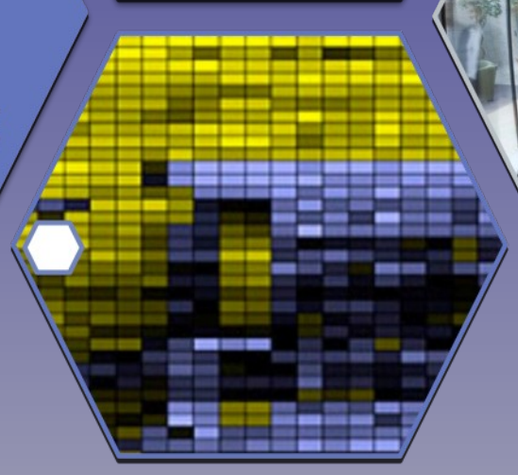

wellcome trust sanger institute

The Story so Far

- **At the iRODS user group meeting at SC2013 we left with the parting future work slide.**
- **What follows is a brief recap and where we are now...**

wellcome trust sanger institute

Sanger NGS Data Flow

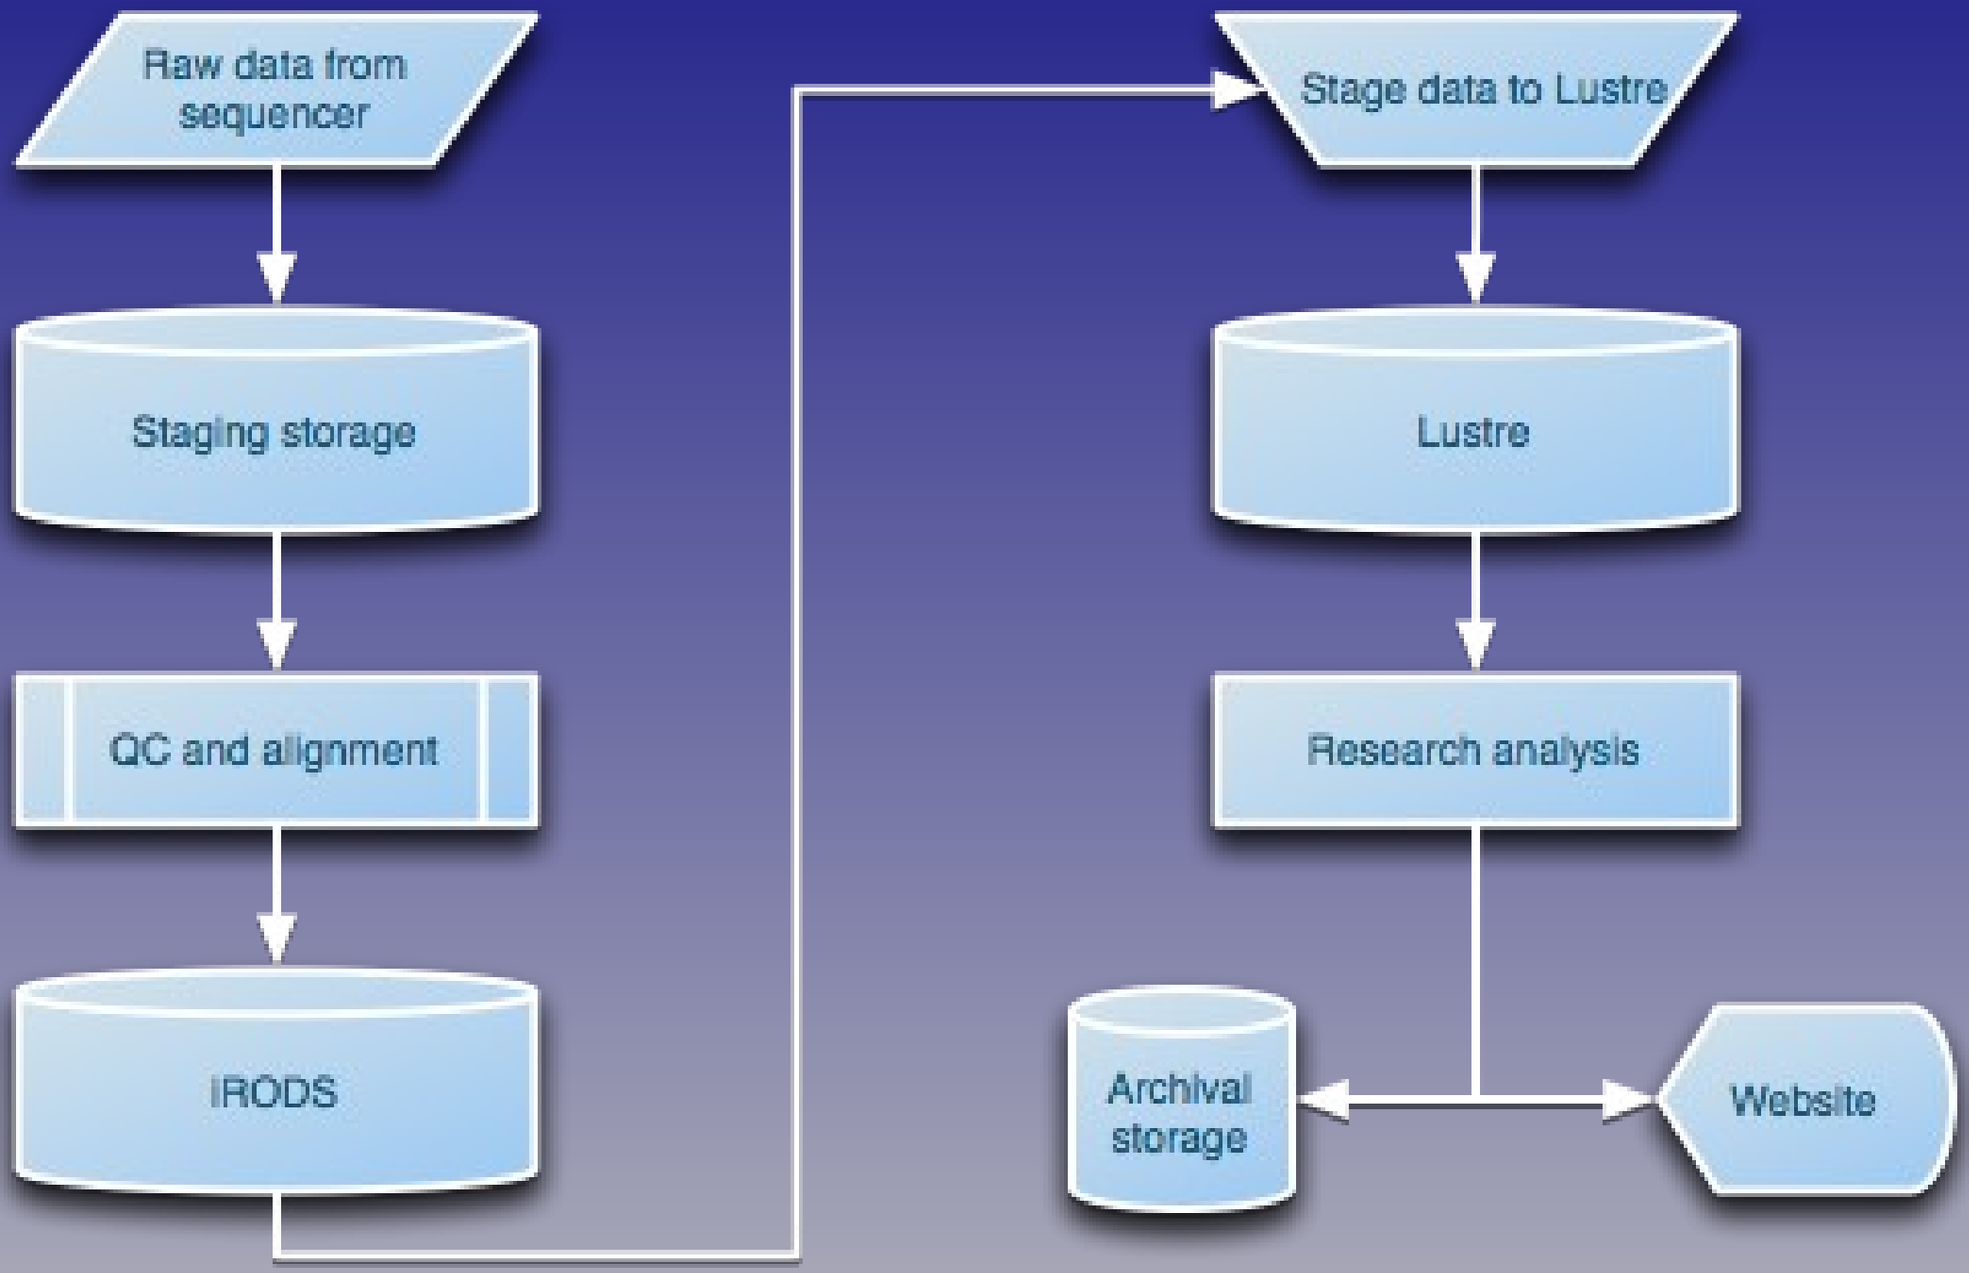

wellcome trust sanger institute

Staging Storage

- **Simple scale-out architecture**
  - **Server with ~50TB direct attached block storage**
  - **One per sequencer**
  - **Running SAMBA for upload from sequencer**
- **Sequencing data ~2 TB/day**
- **1000 core cluster**
  - **Read data from staging servers over NFS**
  - **Quality checks**
  - **Alignment to reference genome**
  - **Store aligned BAM and/or CRAM files in iRODS**

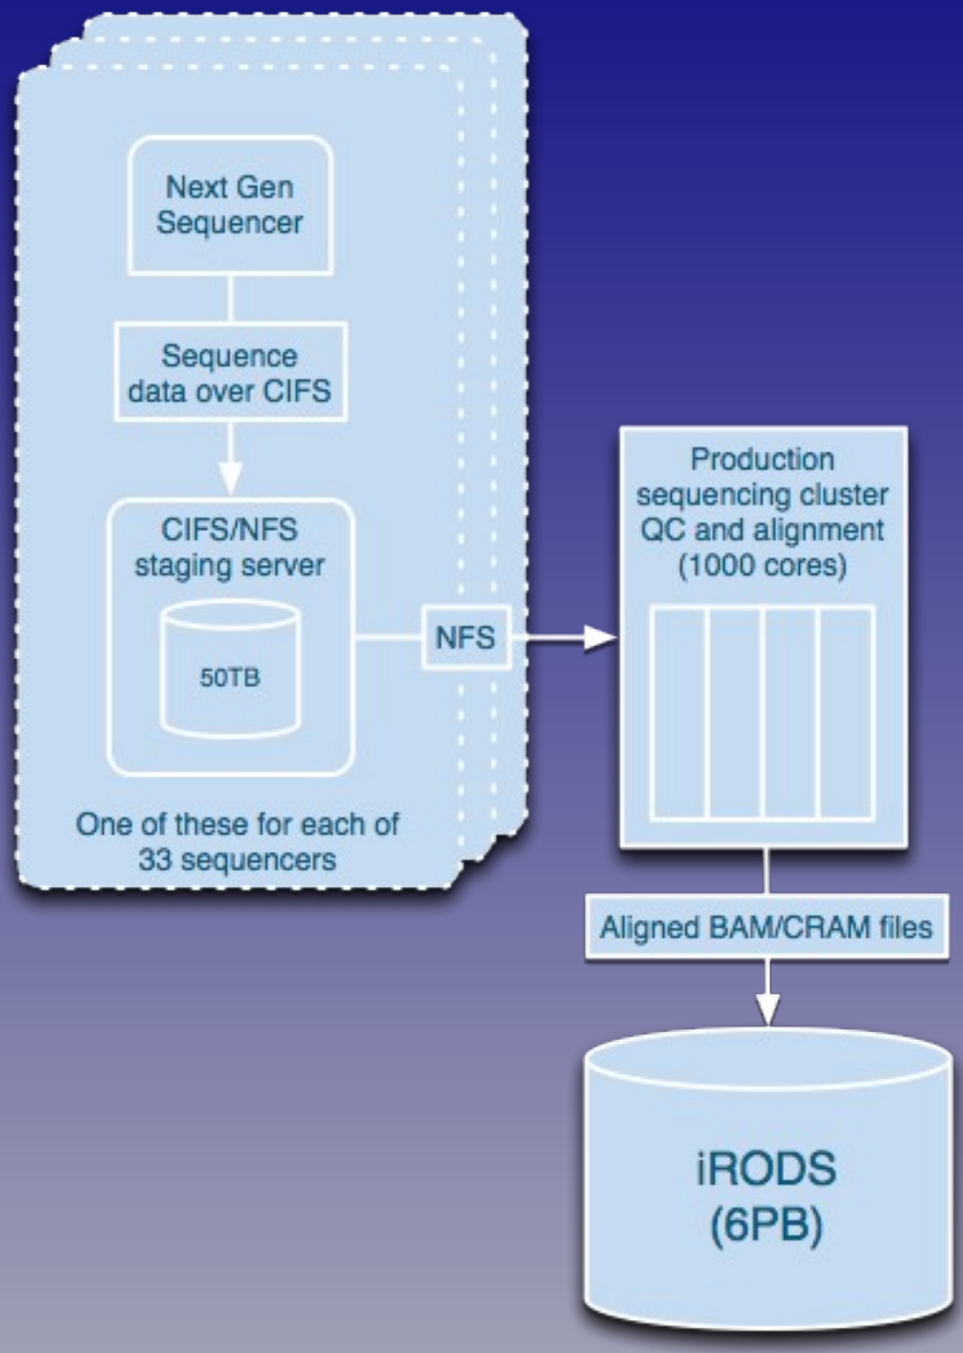

wellcome trust sanger institute

Next Steps: Simplification?

- **Simpler scale-out architecture**
  - **Storage servers with ~50TB internal storage**
  - **≥20 CPU cores per server**
  - **≥256 GB RAM**
  - **One per sequencer**
  - **Running SAMBA for upload from sequencer**
- **Run primary pipeline directly on the storage server**
- **Eliminate the separate cluster**
- **Eliminate NFS data access**
- **Smaller failure domain**

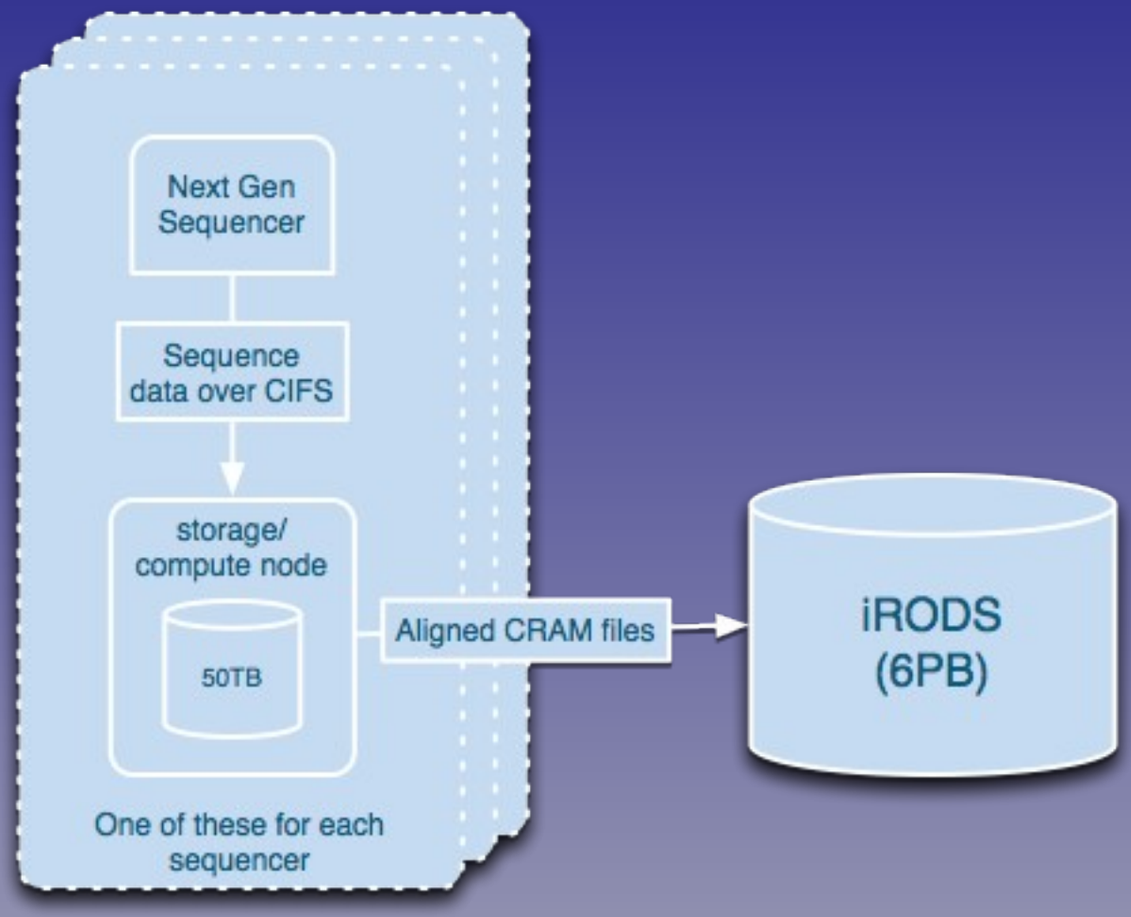

wellcome trust sanger institute

Logical Architecture

- Vendor-agnostic object store
  - arbitrary metadata
  - “microservices” operate on the data
  - Rules apply micro services to the data
- Multiple federated zones
  - Kerberos authentication from the “portal zone”
- Data replicated using a rule
- Each zone has different rules according to scientific need
  - Replication checksums
  - Data integrity checks
- Checksum by the application before sending to iRODS
- We keep rules light

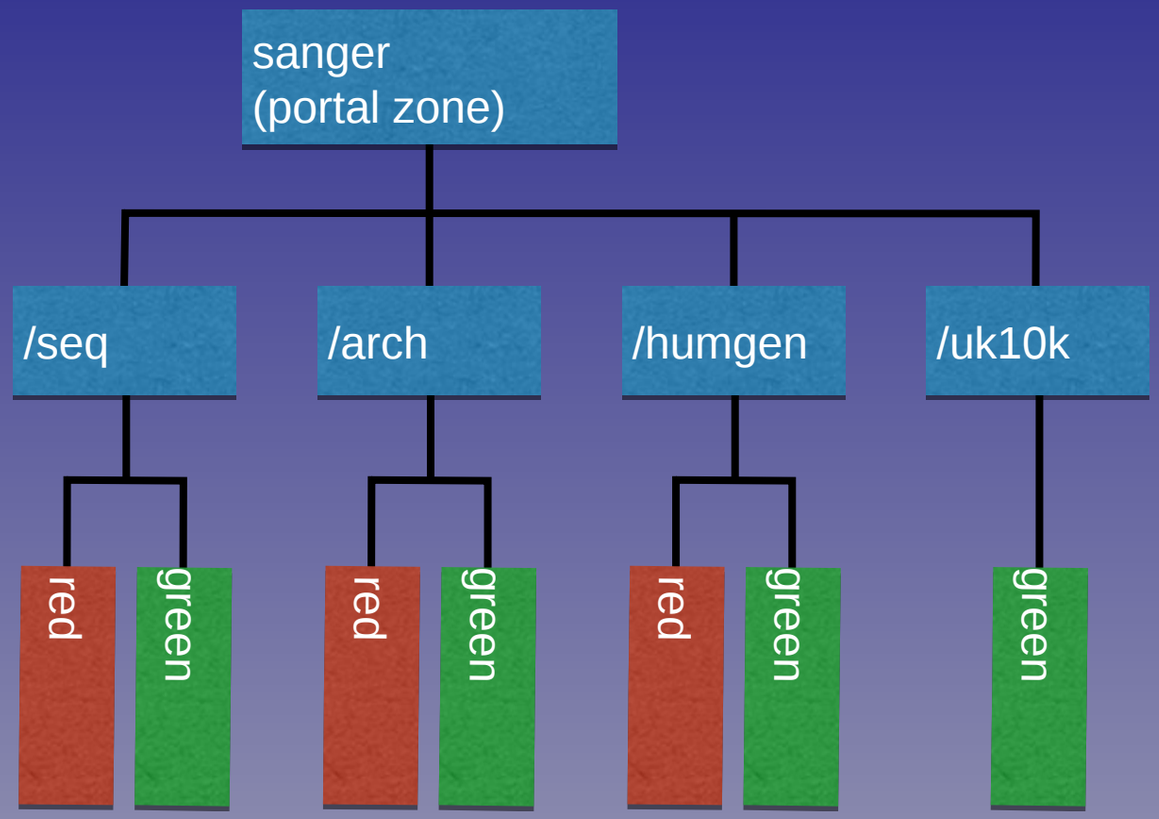

Physical Architecture

- **Hardware**
  - **DDN SFA 10K**
  - **Other vendors' storage**
- **Oracle RAC cluster holds metadata**
- **Two active-active resource servers in different rooms**
  - **8Gb FC to storage**
  - **10Gb IP to clients**
- **Series of 43 TB LVM volumes from 2x SFA 10K in each room**

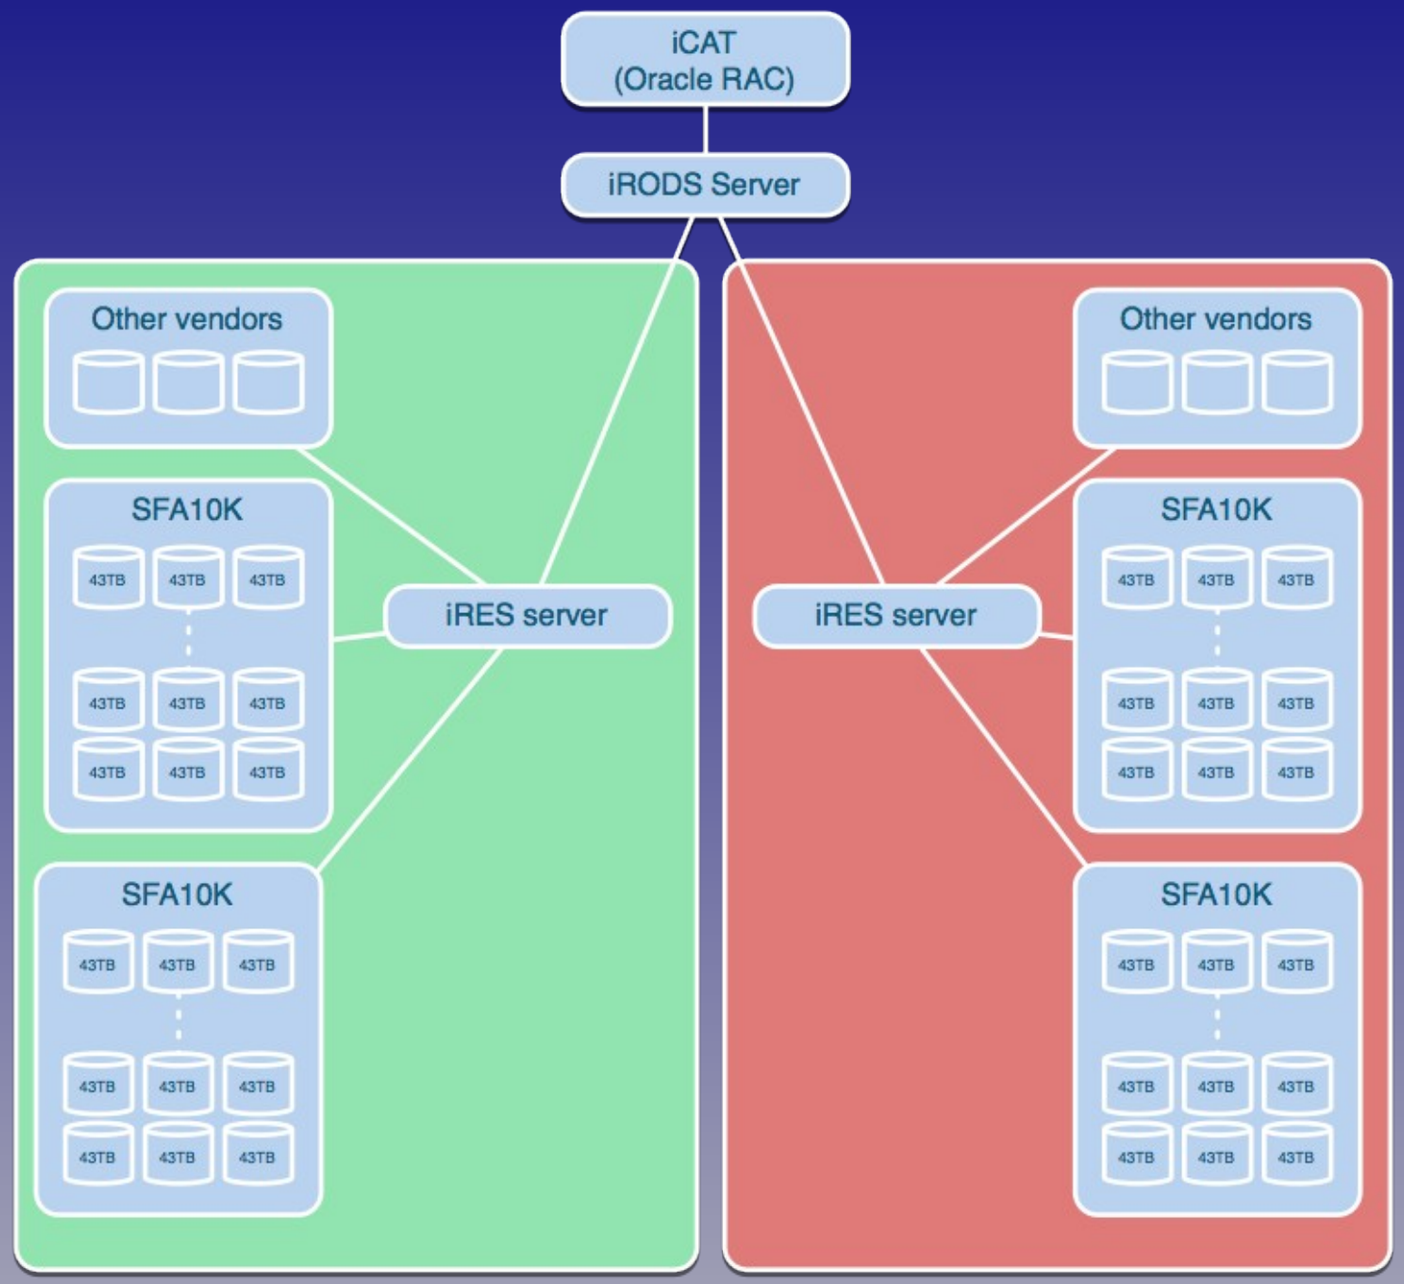

Current iRODS architecture concerns

- Only two iRES servers per zone
  - Large failure domain
  - Limited bandwidth
- Large RAID arrays are expensive
- Large numbers of sizeable POSIX filesystems
- Both copies still on site

wellcome trust sanger institute

Safeguarding Our Data

wellcome trust sanger institute

Sometimes bad things happen

BBC NEWS

LIVE BBC NEWS CHANNEL

News Front Page
World
UK
England
Northern Ireland
Scotland
Wales
Business
Politics
Health
Education
Science & Environment
Technology
Entertainment
Also in the news

Last Updated: Sunday, 13 July, 2003, 09:22 GMT 10:22 UK

E-mail this to a friend Printable version

Air show deaths inquiry starts

The Ministry of Defence has started an investigation after two men were killed when an historic aeroplane crashed at an air show.

Lieutenant Commander Bill Murton, 45, and Neil Rix, 29, were crewing the Fairey Firefly vintage
came d
airfield
the eas

BBC NEWS

LIVE LI BBC NEWS CHANNEL

News Front Page
World
UK
England
Northern Ireland
Scotland
Wales
Business
Politics
Health
Education
Science & Environment
Technology
Entertainment
Also in the news

Last Updated: Tuesday, 19 August, 2003, 09:

E-mail this to a friend Prin

Airfield to examine safety

Officials at the Imperial War Museum airfield are to examine airshow safety after the crash of a World War II bomber on Monday - the third crash there in two months.

The Blenheim crash was the third at Duxford in two months

Two men had a lucky escape when the Bristol Blenheim bomber crash "belly-flopped" on the grass near the busy M11 motorway.

NEWS ENGLAND

Home World UK England N. Ireland Scotland Wales Business Politics Health Education Sci/Env

Fighter planes collide at Duxford air show

11 July 2011 Last updated at 16:25

wellcome trust sanger institute

So following on from our last User Group presentation....

- **We should probably keep our second copy offsite...**
- **Should be easy... We have two copies of the data, just ship one elsewhere.**
- **A second DataCentre needed:**
  - **JANET Shared Datacentre Project**
  - **Infinity SDC – Slough (Just outside London)**
- **Moving Pb of data is not easy...**

wellcome trust sanger institute

Capacity Expansion and Offsite Migration

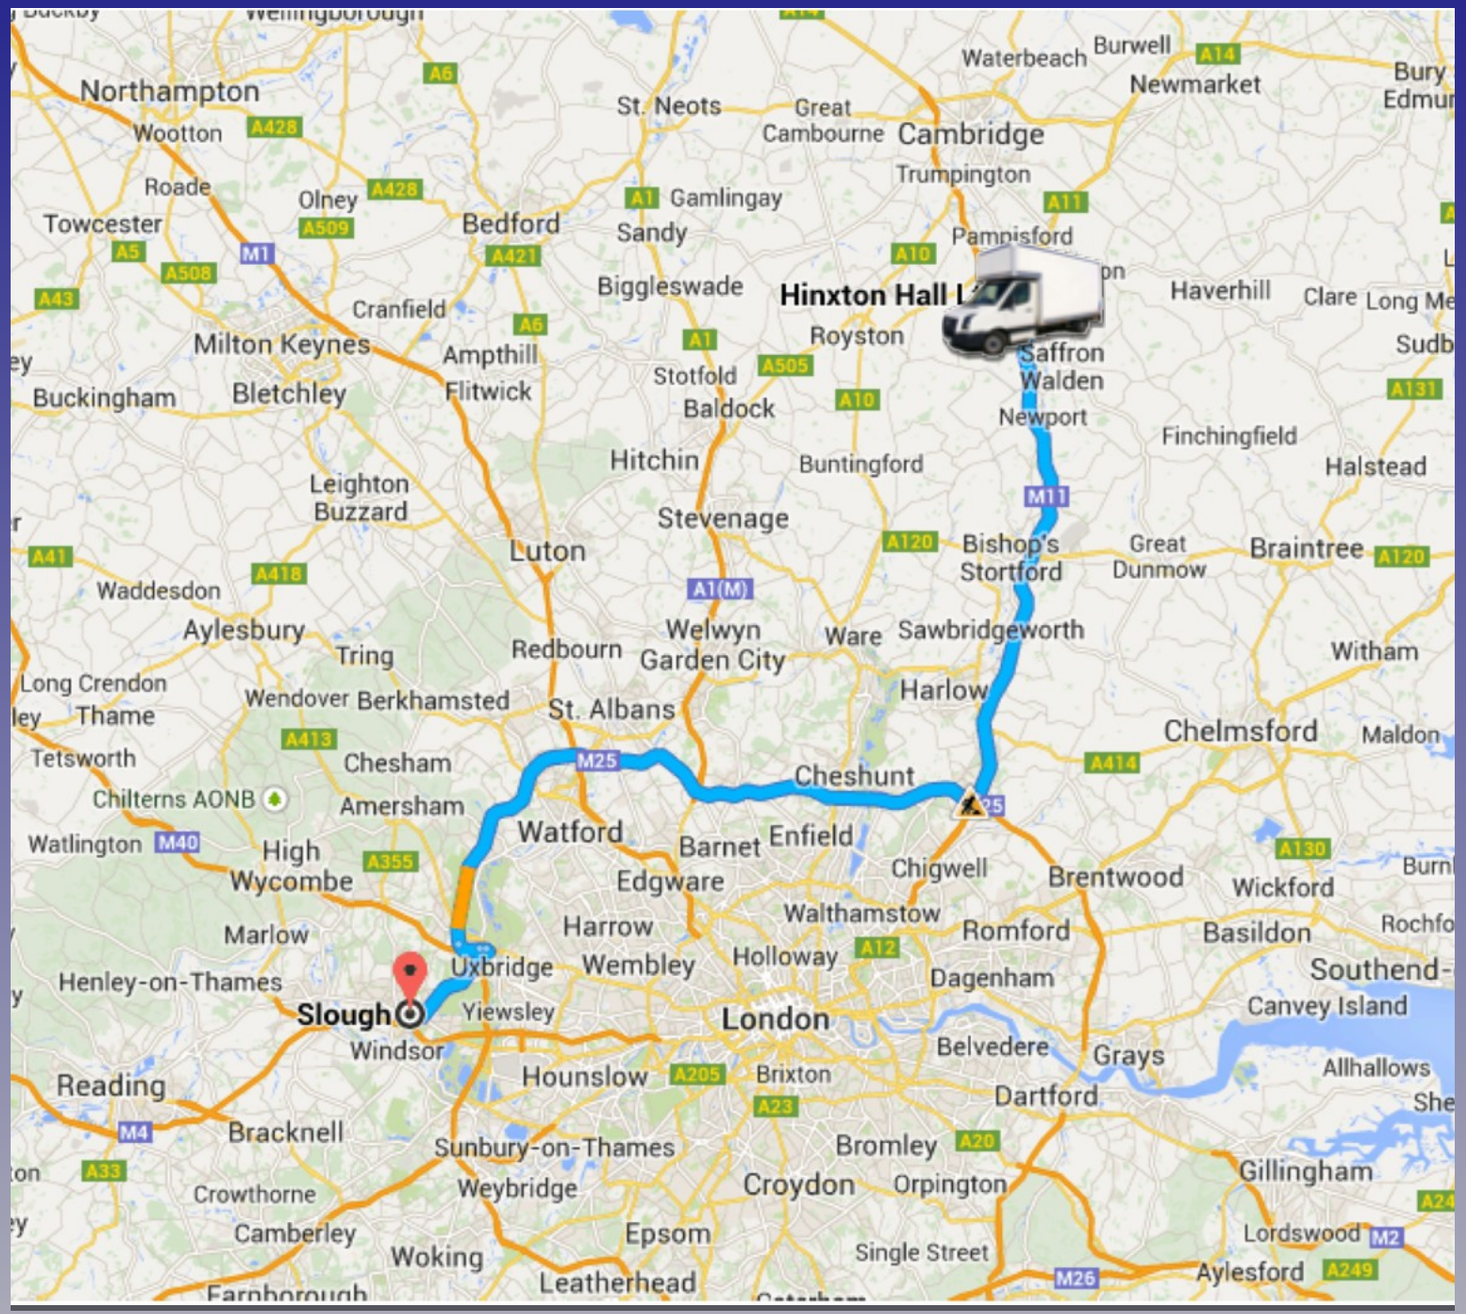

Capacity Expansion and Offsite Migration

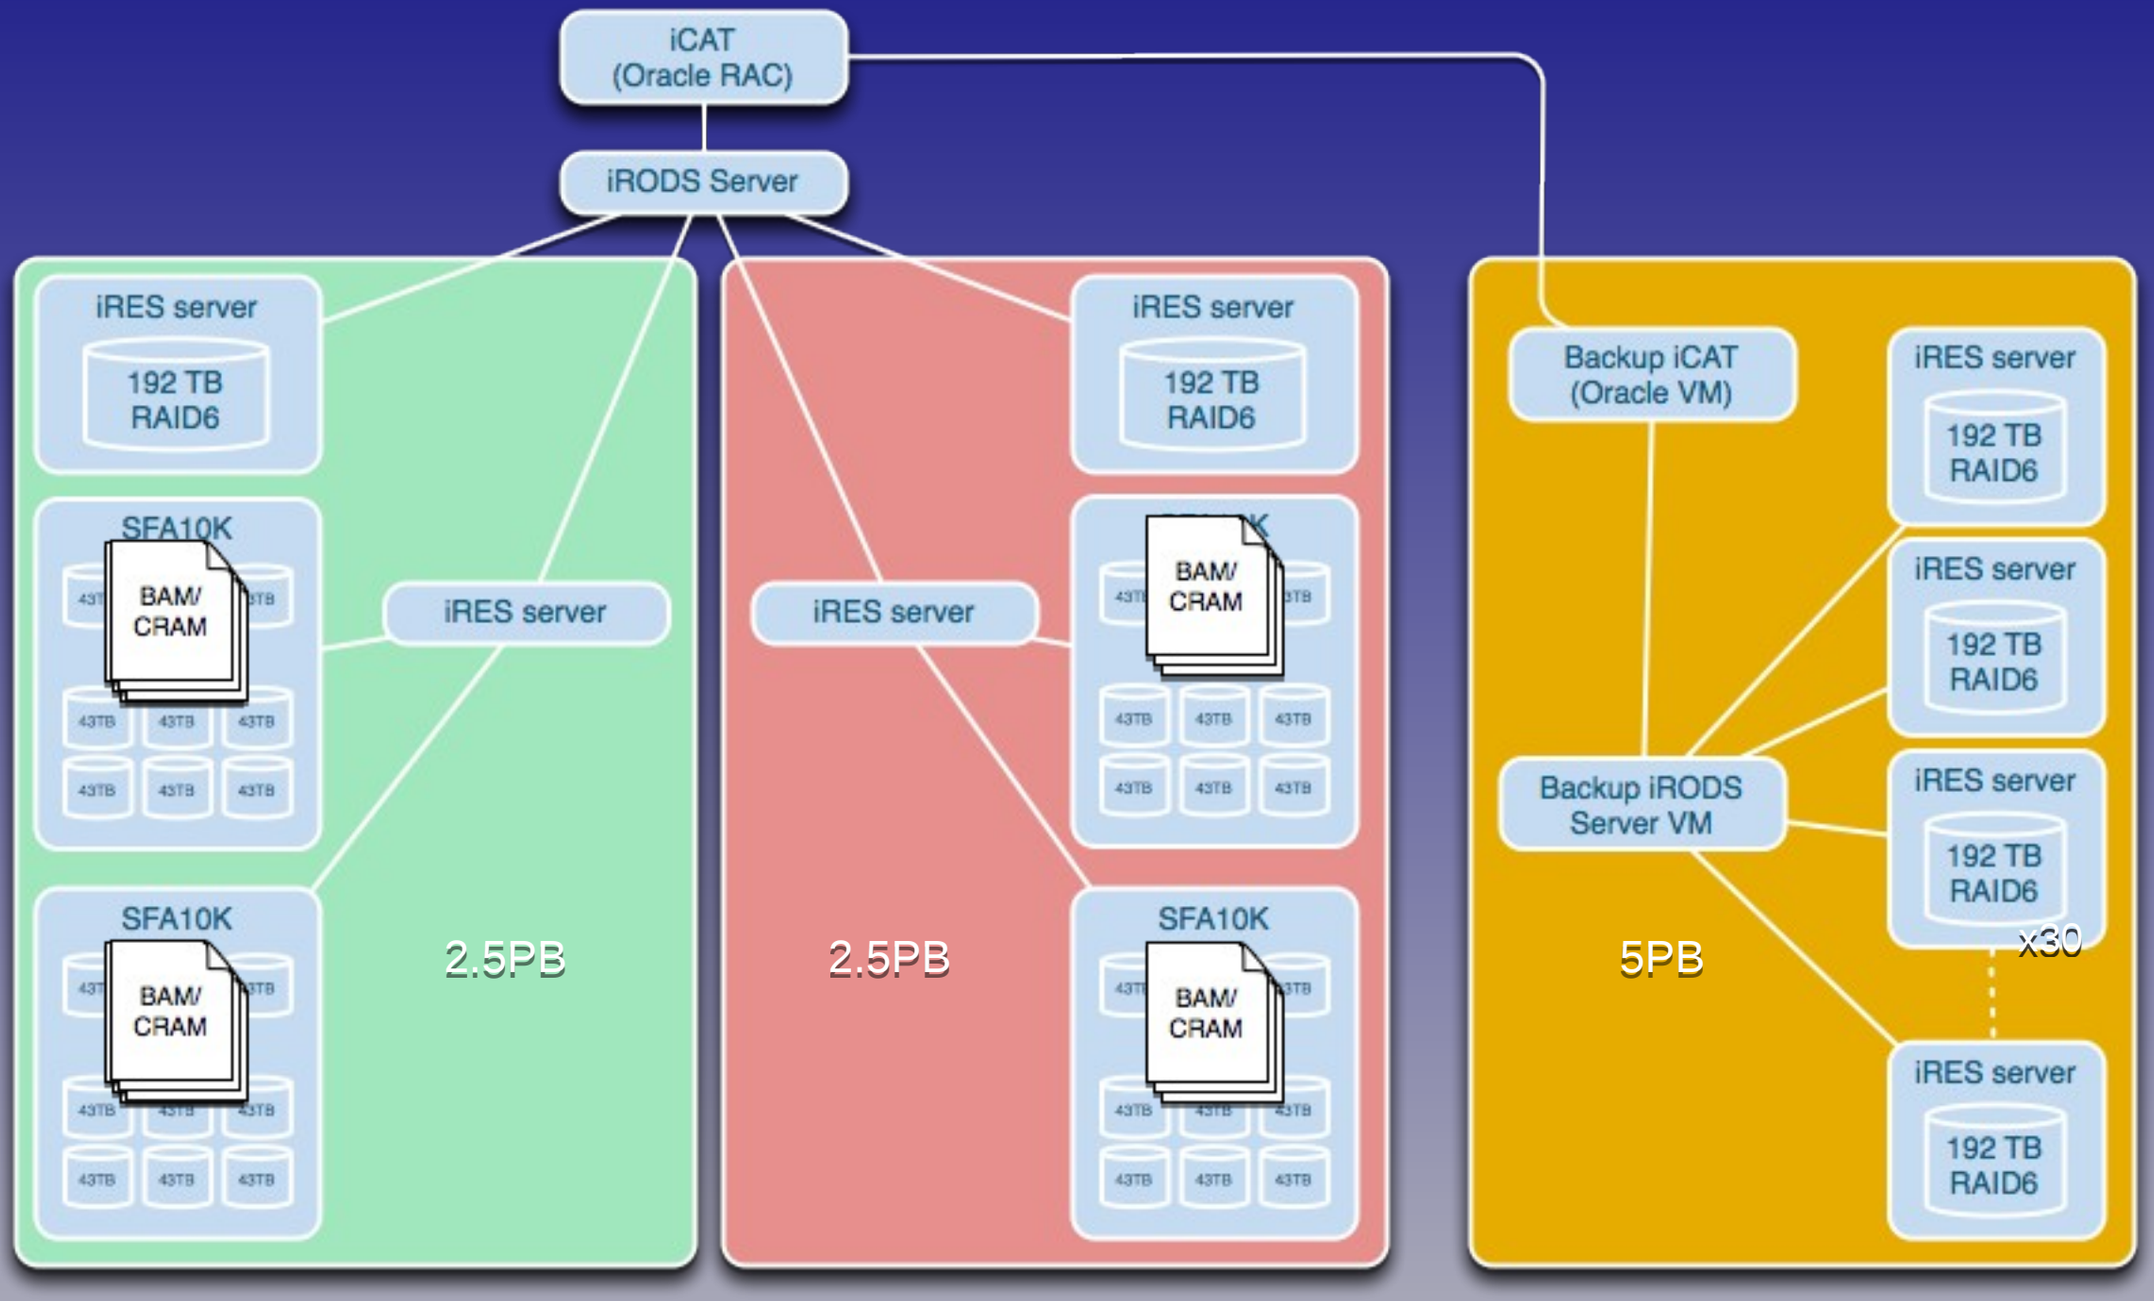

wellcome trust sanger institute

iRODS Expansion Key Points

- We had two copies available at all times
  - Built the new third copy on site
  - Test
  - Shipped offsite
  - Tested again
  - Now freeing one local copy for re-use as new capacity
- Self-contained storage servers for small iRODS zones
- Manage the storage nodes in the same way as our compute nodes
  - cfengine, automated installation, automated firmware updates
- Doesn’t solve the POSIX filesystem issue
- Reduces failure domain
- Higher bandwidth
- Cost effective, but will it be reliable enough?

wellcome trust sanger institute

And the results

So far so good.

- **Only one errant file permission so far observed.**
- **We did find some files that managed to sneak under the hood. This required run through the iRODS iquest and checks on the filesystem**
- **The last mile was, as always the hardest.**
- **By including a VM infrastructure and a second readonly copy of the ICAT db server we now have a viable disaster recovery option.**

wellcome trust sanger institute

The future…

- As always hard to predict.

  But we know we need:
  - Metric & Monitoring
  - Security
    - Federated authentication
    - Secure data
      - In Flight (4.x)
      - ? At Rest
  - Improved iCAT and meta data access

- (it's not awesome... yet)

wellcome trust sanger institute

Metrics are really very useful.

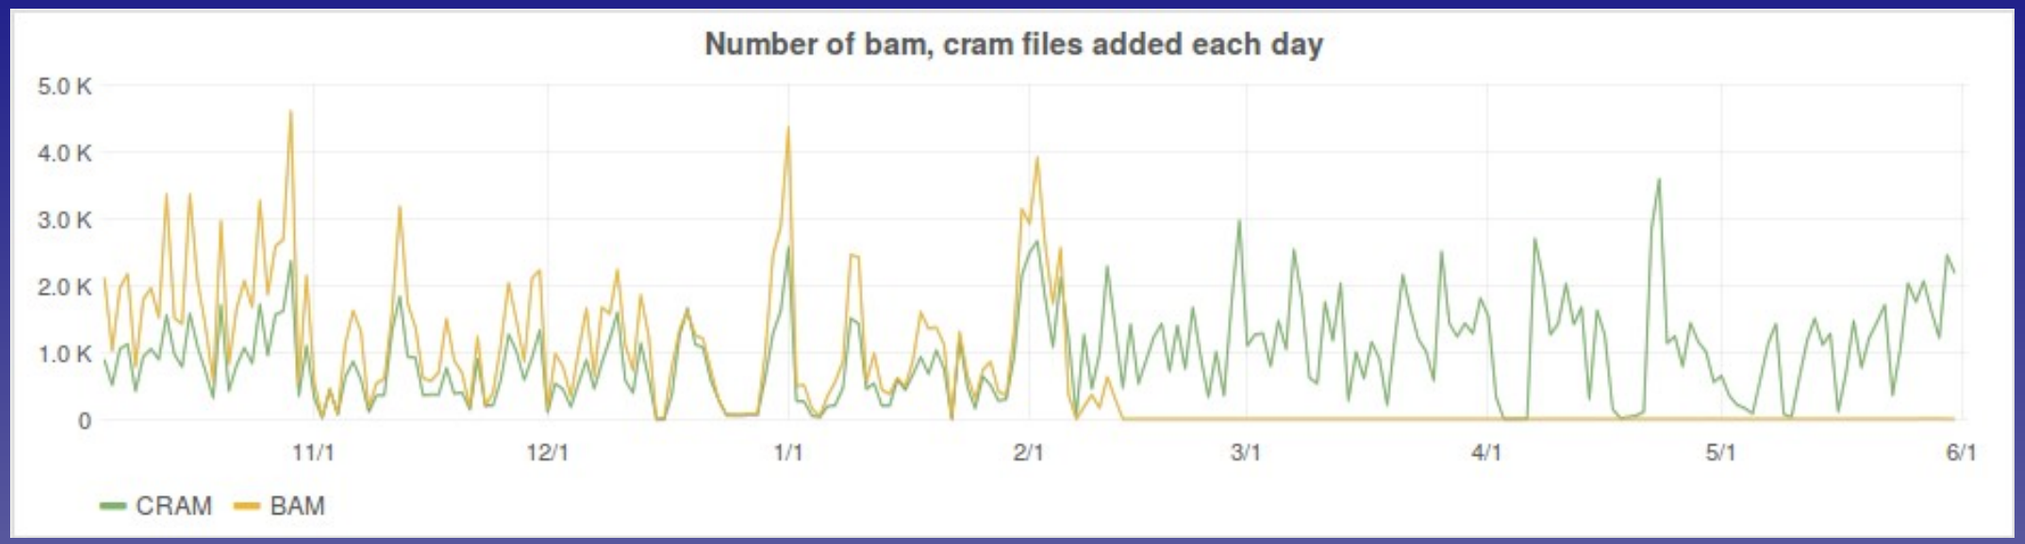

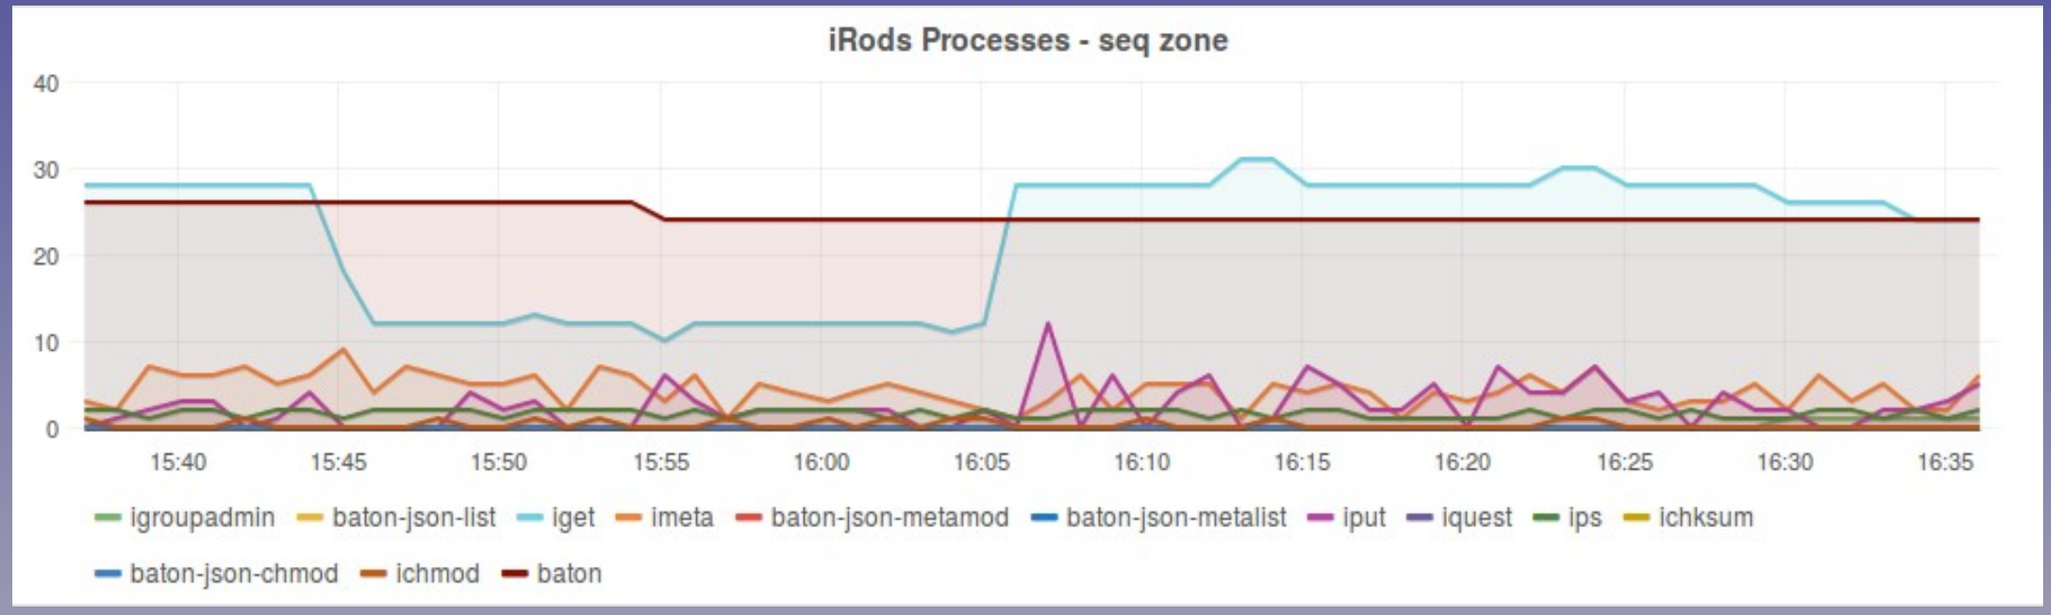

wellcome trust sanger institute

Very useful indeed

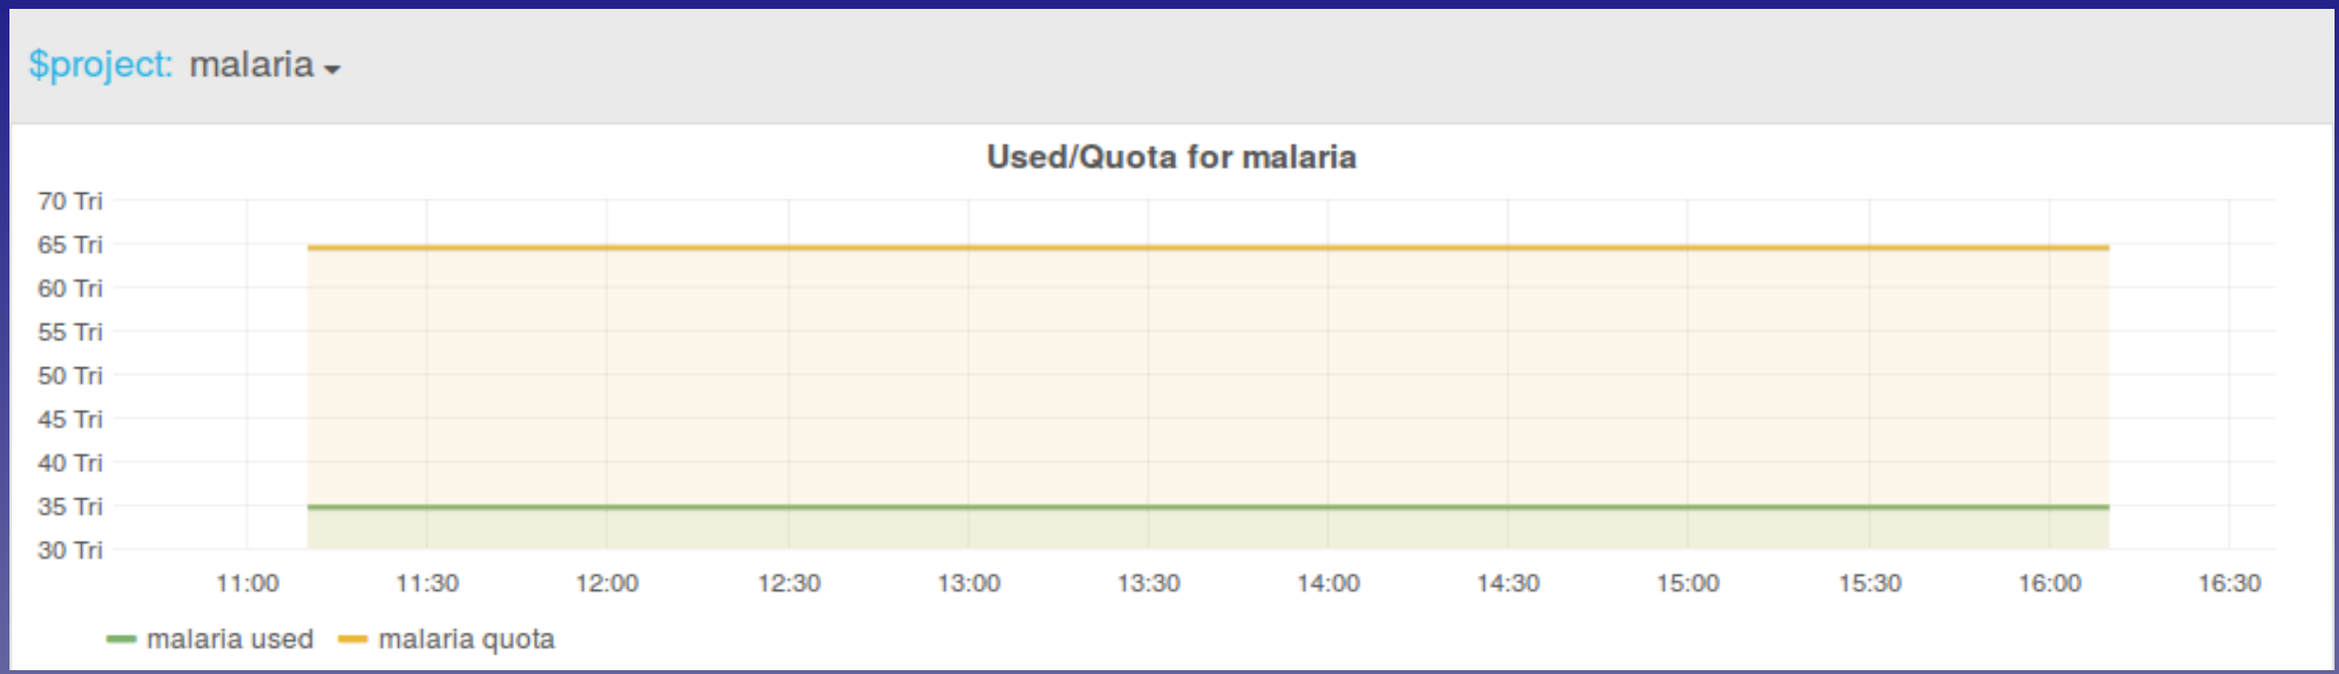

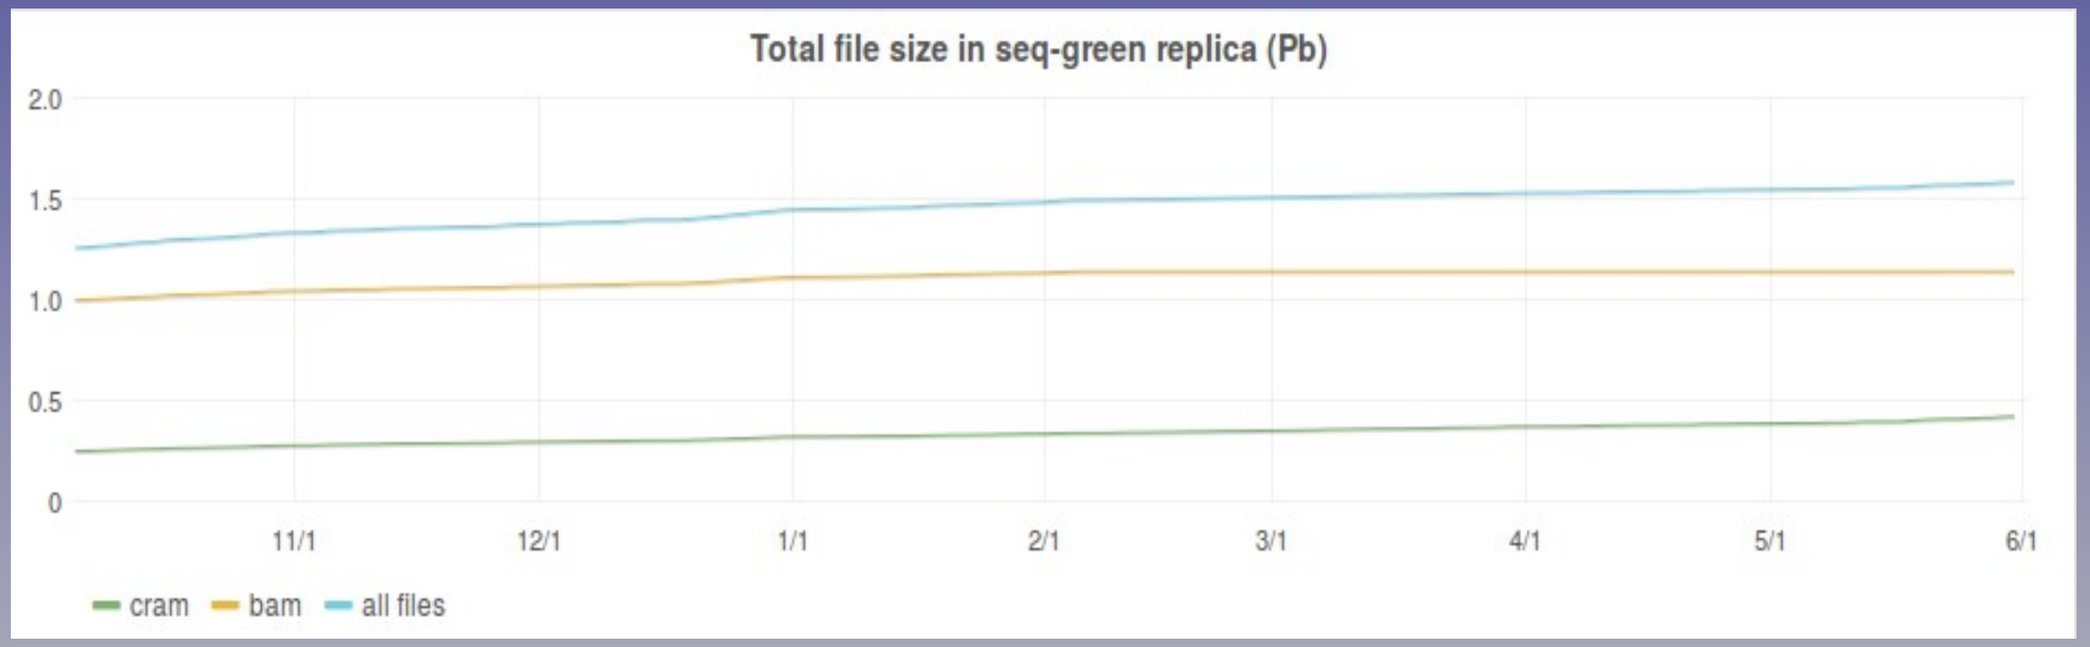

wellcome trust sanger institute

Metrics

- **Graphana with Influxdb, have proven very powerful, flexible and allows for user reports as well:**
  - **Scalability prediction**
  - **Capacity management**
  - **Problem management**
  - **User advice**
  - **Is the system working as intended ?**

- **It's hard to over state the importance of following these metrics. Next steps:**
- **Predictive management and machine learning ?!**

wellcome trust sanger institute

eMedLab

- **£9million MRC funded collaborative bio-research project to provide:**
- ***"The grant was given for Medical Bioinformatics: Data-Driven Discovery for Personalised Medicine. The funding will be used to create 'eMedLab', a shared offsite datacenter."***

  http://www.uclpartners.com/news/partnership-receives-9-million-from-mrc-towards-improving-data-research/

- **Initial partners include:**
- **University College London (UCL)**
- **The Francis Crick Institute**
- **Queen Mary University of London**
- **Sanger Institute**
- **European Bio-informatics Institute (EBI)**

wellcome trust sanger institute

Collaboration: eMedLab

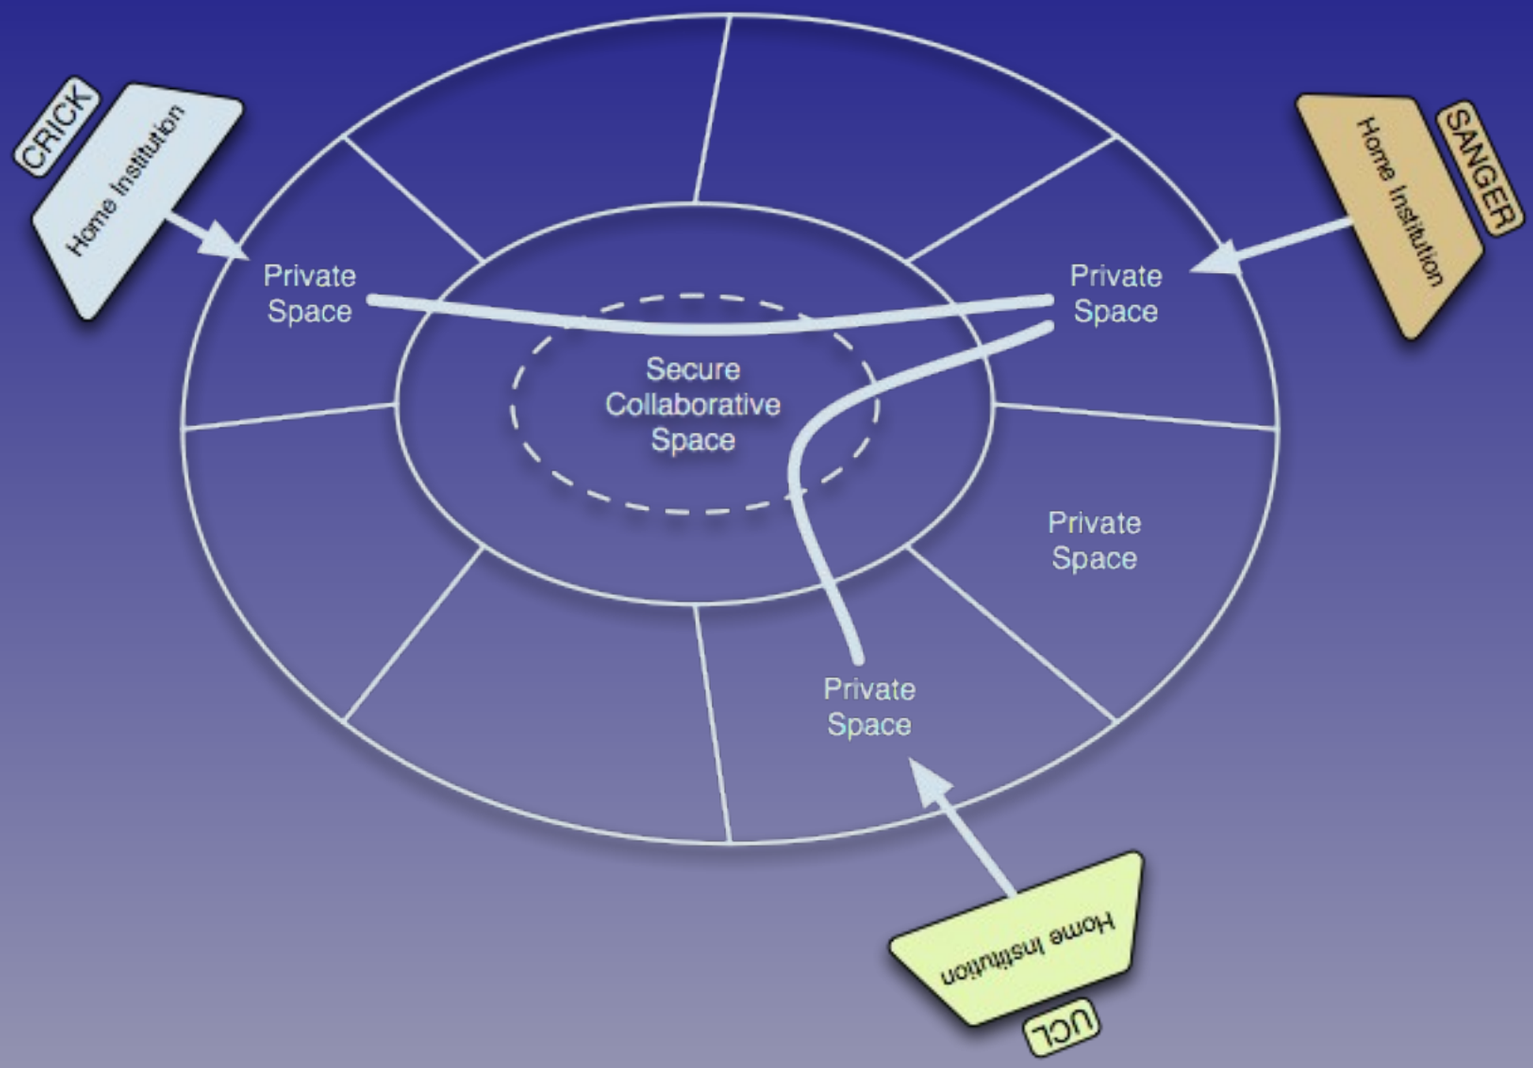

MRC Medical Research Council

THE FRANCIS CRICK INSTITUTE

UCL

EMBL-EBI

Queen Mary University of London

LONDON SCHOOL of HYGIENE &TROPICAL MEDICINE

wellcome trust sanger institute

Collaboration: eMedLab

- Hosted at the Janet Shared Datacentre in Slough
- Datacentre selected to meet IL3 and NHS security requirements
- Connectivity to Janet and NHS N3 networks
- 6,000 cores, 5PB of storage, multi-tenant virtualised HPC
- Due to begin operation in Q2 2015 (aggressive!)
- Pilot project in 2015 with the FARR Institute around federated AuthN/AuthZ to access clinical datasets using Janet Moonshot (ASSENT)

eMedLab

- **At it's core the facility aims to provide:**
  - **Common shared datasets**
    - ***IRODS* has been selected as a part of this solution**
- **Virualised compute infrastructure based around OpenStack**
- **Cross Institute federated access and multi-tenant environment.**

- **The first significant production use of JANET Assent Federated access.**
- **Opportunity to further develop multi-tenant iRODS federation with ASSENT integrated with PAM.**
- **New opportunities with 4.1 iRODS release for multi-tenant secure data management.**
- **All of the above are essential as we continue to see further medical developments developing from the world of bio-informatics !**

wellcome trust sanger institute

iRODS 4.1 ->

- **Currently still based on 3.3.1 in production**
- **Stagnating**
- **4.x stream has been a very rapidly developing feature stream, some very much in demand, however they include:**
  - **Many changes to the base**
  - **Some include significant changes at the config layer**
  - **Others include a complete re-structure to the resource and group layers.**

- **We have 7PB of storage behind the existing infrastructure..**
  - **This makes patching, upgrading and restructuring a twitchy option.**

wellcome trust sanger institute

iRODS 4.x ->

- **Re-evaluating 4.x in Q3 2015 and plan to upgrade shortly after if all goes well.**
- **Looking to impliment:**
  - **Hierarchal groups and replication.**
  - **PAM – SSL authentication**
  - **Full SSL data encryption of inflight data**
  - **Split out metadata to an external db and impliment solr or elastic search**
  - **iDROP 3.x for external data uploads from Sanger ?**
  - **JANET ASSENT federated data access for externally facing zones federated with other institutes.**

wellcome trust sanger institute

Thank you

- **The team**
  - **Paul Woobey, IT Director**
  - **Tim Cutts, Head of Scientific Computing**
  - **Peter Clapham, Acting Informatics Systems Group Team Leader**
  - **Guy Coates, former ISG TL**
  - **James Beal , John Constable Simon Fraser**
  - **Mark Jowett, Infrastructure Team Leader**
  - **Shanthi Sivadasan, DBA Team Leader**
  - **Numerous bioinformaticians**

wellcome trust sanger institute