- How a large site will support Atlas and other VOs and how other internal pressures interact with this.
- Such as using off site computer rooms, the interaction with local (tier-3) clusters, and funding sources and pressures

The World According to the QMUL Grid cluster

Daniel Traynor, GridPP38

QMUL setting

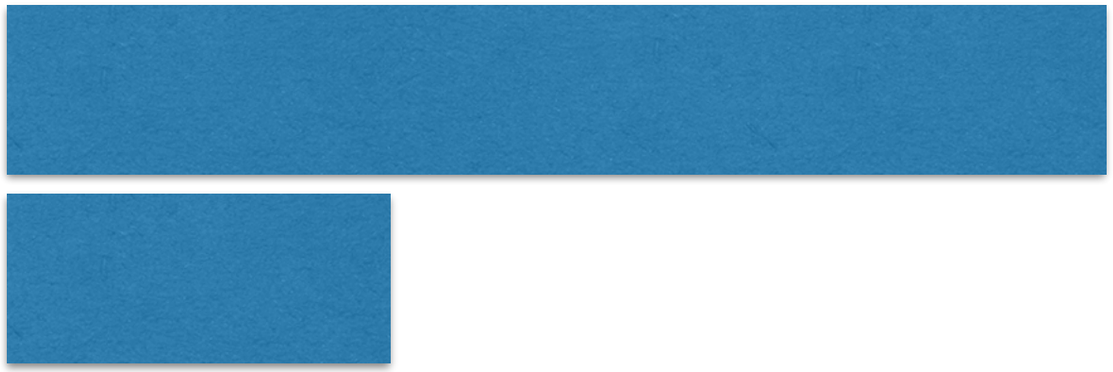

Central IT
4000 job slots,
1PB DDN GPFS

pay for support for:
GPFS, GE,…
performance - infiniband, DDN
big deals site is offline, data backed up
moved to slough, remote hands

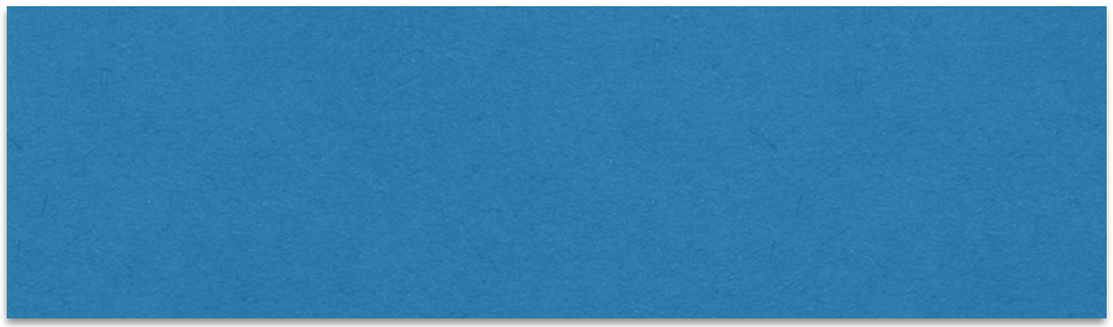

GridPP
4000 job slots
3.3PB Lustre

Open source self support:
Lustre , Slurm
performance - good enough
Down time needs to be limited
dedicated on campus DC

Physics
500 job slots
500TB NFS

Open source self support, NFS, GE
performance - good enough
users are in next door office
sited in Physics

On campus DC

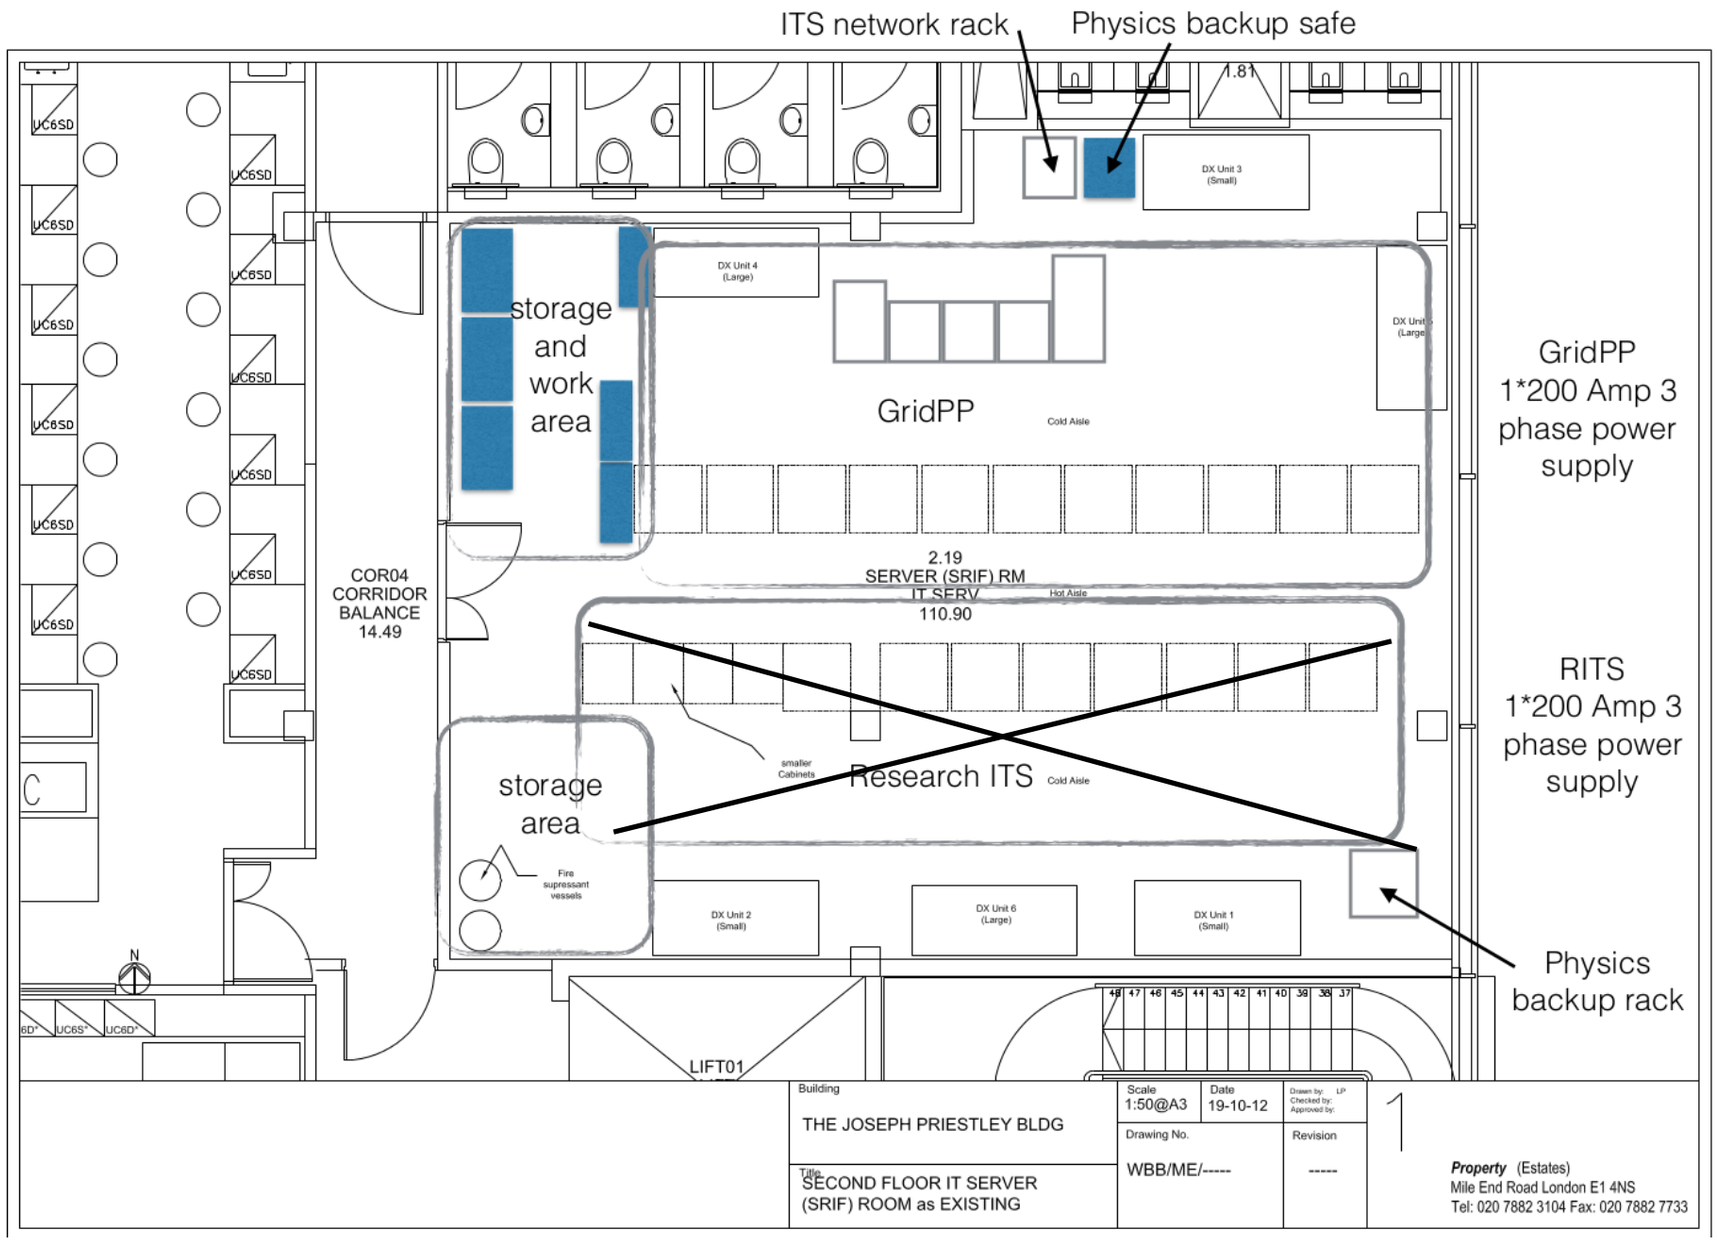

SRIF room layout, GridPP occupies about half the room. In addition to storage and working areas backup services for physics are also located in this room.

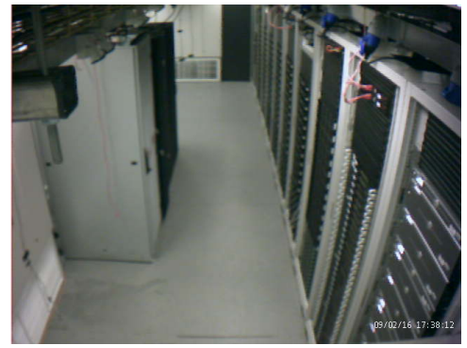

SRIF room cold aisle row 1/2

SRIF room hot aisle

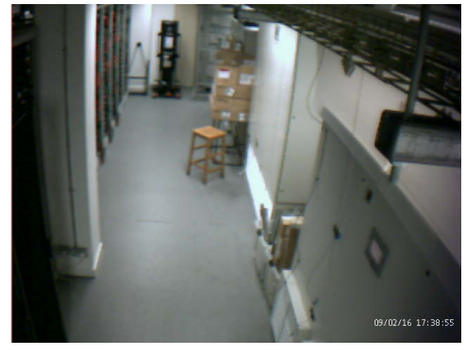

SRIF room cold aisle row 3/4

SRIF room working area

Future Stress

- Two power cuts building in summer due to lack of building power. building also contains MRI, wet labs,
- Require refresh of room infrastructure. Cooling units spare parts have to come from Canada now.
- ITS research (1/2 room) have moved to Slough Janet DC.
- Major project to upgrade building infrastructure, more power, more cooling. Could benefit us.
- Room would be a good lab space

Atlas support

ATLAS say Jump

We say How High

Other VO

If Atlas works most other VOs support should be noise

General Aims

- Make site as HTC/HPC standard as possible:
  - SLURM, Lustre, Private network
- Make sure software has longterm support and development:
  - LBNL, LLNL, Intel
- Keep the hardware simple vs resilience
  - 1U pizza boxes vs “Blades” vs Big Disk arrays
- Diversification of software compared to other grid sites.
- Unique resources: GPUs, High mem queue (4G/job), (Xeon Phi?)

Transformation

- [x] NAT gateway -> Failover Router (supported)
- [x] Physically separate Out Of Band network for PDUs, switches and IPMI.
- [x] IPv6 on all front end nodes (CEs,SEs,Squid,…)
- [ ] Home brew VMhosts -> Proxmox (open source, can buy support)
- [ ] Home brew deployment -> Ansible+Cobbler (open source, can buy support)
- [ ] IPv6 on the worker node
- [ ] SL/Centos 7 + containers (singularity)
- [ ] NIS -> LDAP+Kerberos

Resiliency

- [x] Failover Lustre MDS servers.
- [x] Failover router.
- [x] distributed core network.
- [ ] High availability VMhosts.
- [ ] Redundant power/ UPSs on fronted rack with core services.
- [ ] 90 Storage servers -> several with OSS redundancy (failover Lustre OSS setup) and power and network redundancy.

Also Warranty NBD 5 years for storage + critical servers. Every thing with access via Out Of Band management network. 1G backup network.

Site Funding

- University facility maintenance and power (85KW)
- GridPP 2 FTE + Hardware
- ~ £5K a year from University for…
  - Development and experimentation e.g. GPUs, Knights landing box, ARM server, SSDs testing, etc…
  - Spares, new tools, gadgets, bits & bobs (labels, velcro ties, storage boxes etc..), new backup solution.

Gadgets


DYMO
Clear
Settings
Format
Cancel
OK
Insert


StarTech.com
2017/04/04 15:35
8.Secure Erase
OK
ESC


ASUS
IN SEARCH OF INCREDIBLE
BLU-RAY
WRITER
16X
BW-16D1

Tools and storage

Power Measurement

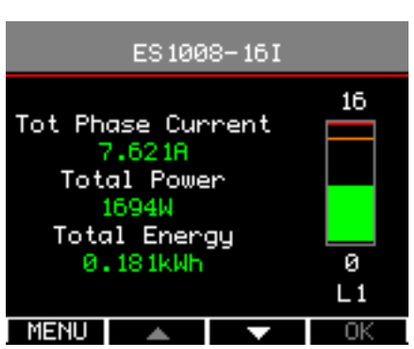

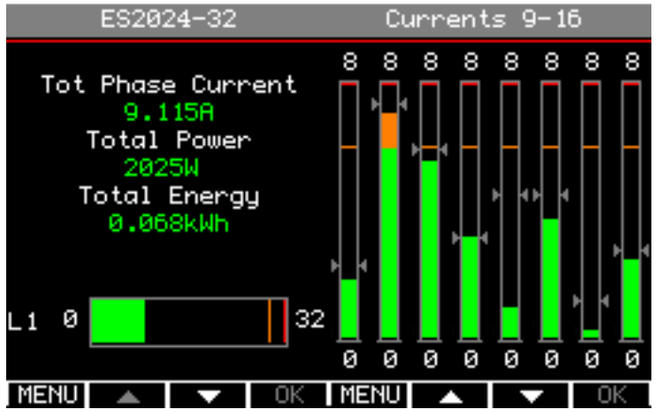

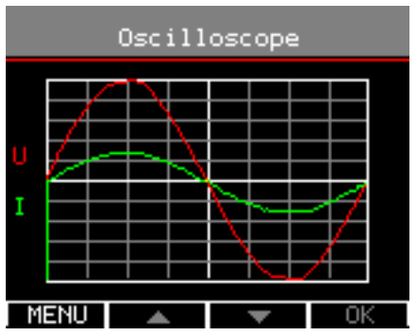

PDU capable of measurement power per port,
Real-Time, True RMS measurements.
Aim to optimise HepSpec/power consumption
monitor long term usage

Development + testing

SSD - performance impact and wear

1 year of testing
250GB SSDs

new 500GB NVMe

40 Gb/s

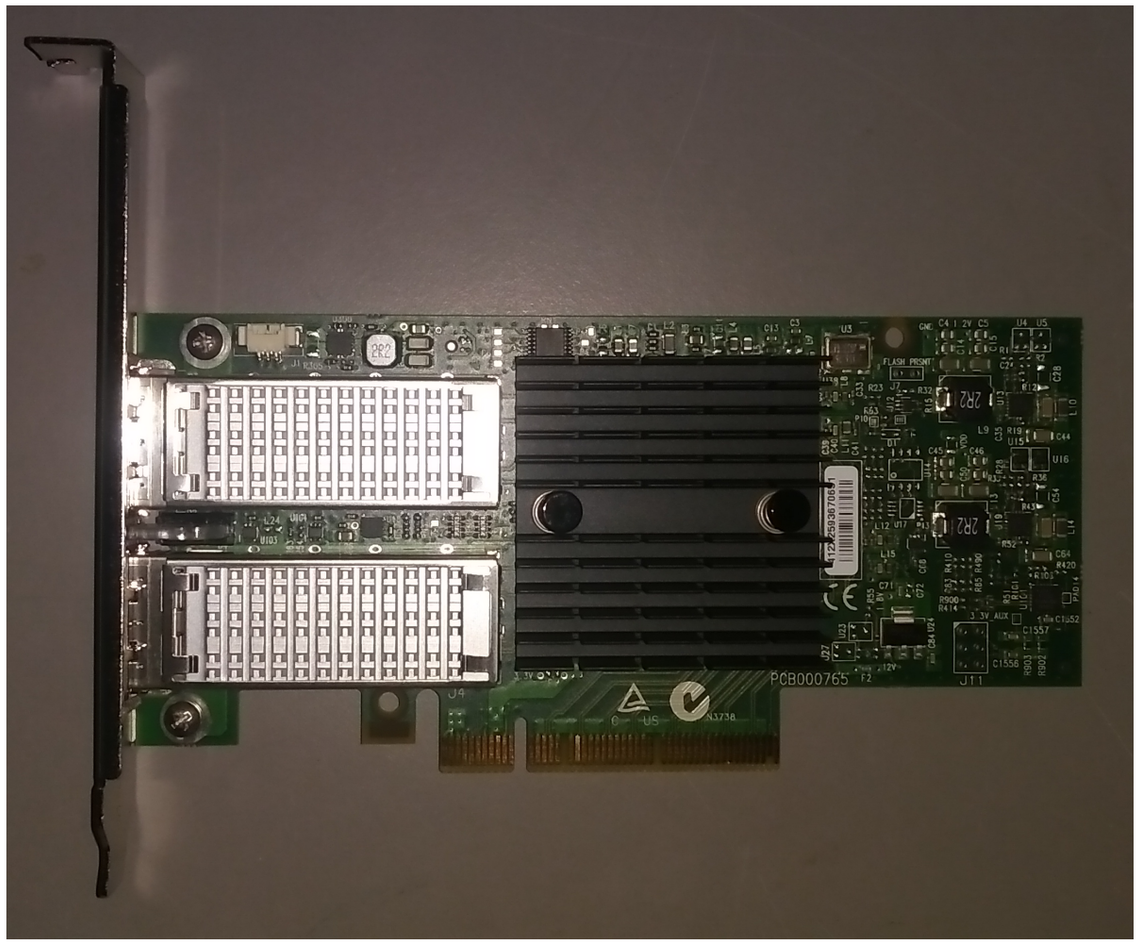

Need to make sure we can run at higher data rates.
Test SEs, Lustre OSSs.
Optimise and specify future hardware requirements

GPUs


PCI Module Removal
PCI Module Replacement
PCI Module Removal
PCI Module Replacement

ZFS

ZFS fully supported in Lustre.
Benefits from raid rebuilds, data consistency, cost.
Build on existing work.
make sure next storage buy is correctly specified.

Machine learning

- Lots of work on this in HEP but more so in wider world.
  - TMVA / Tensor Flow / Caffe etc..

- What impact for GridPP resources? GPUs, Vectorisation (e.g. intel MKL libs).

Big Data

- magpie: https://github.com/LLNL/magpie
  - Magpie contains a number of scripts for running Big Data software in HPC environments. Thus far, Hadoop, Spark, Hbase, Storm, Pig, Mahout, Phoenix, Kafka, Tachyon, Zeppelin, and Zookeeper are supported. It currently supports running over the parallel file system Lustre and running over any generic network filesytem. There is scheduler/resource manager support for Slurm, Moab, Torque, and LSF.

~100 Gb/s wan link

- Have already stared interaction with centralIT/Networks to build business case for ~100Gb/s connection for Run3.
- Science DMZ/Data Transfer Zone
- Some Resilience or redundancy (Don’t need full 100Gb failover)
- IPv6 DNS, NTP, etc…

Summary