



GridPP

UK Computing for Particle Physics



UNIVERSITY
of
GLASGOW

UK HTC: GridPP

12th September 2018

DiRAC Days, Swansea

David Britton



First circle of hell: Limbo

STFC has seven Directorates

Finance

Corp
Services

HR

Strategy,
Policy,
Comms

Business &
Innovation

National
Laboratories
Directorate

Programmes
Directorate


External
Networks

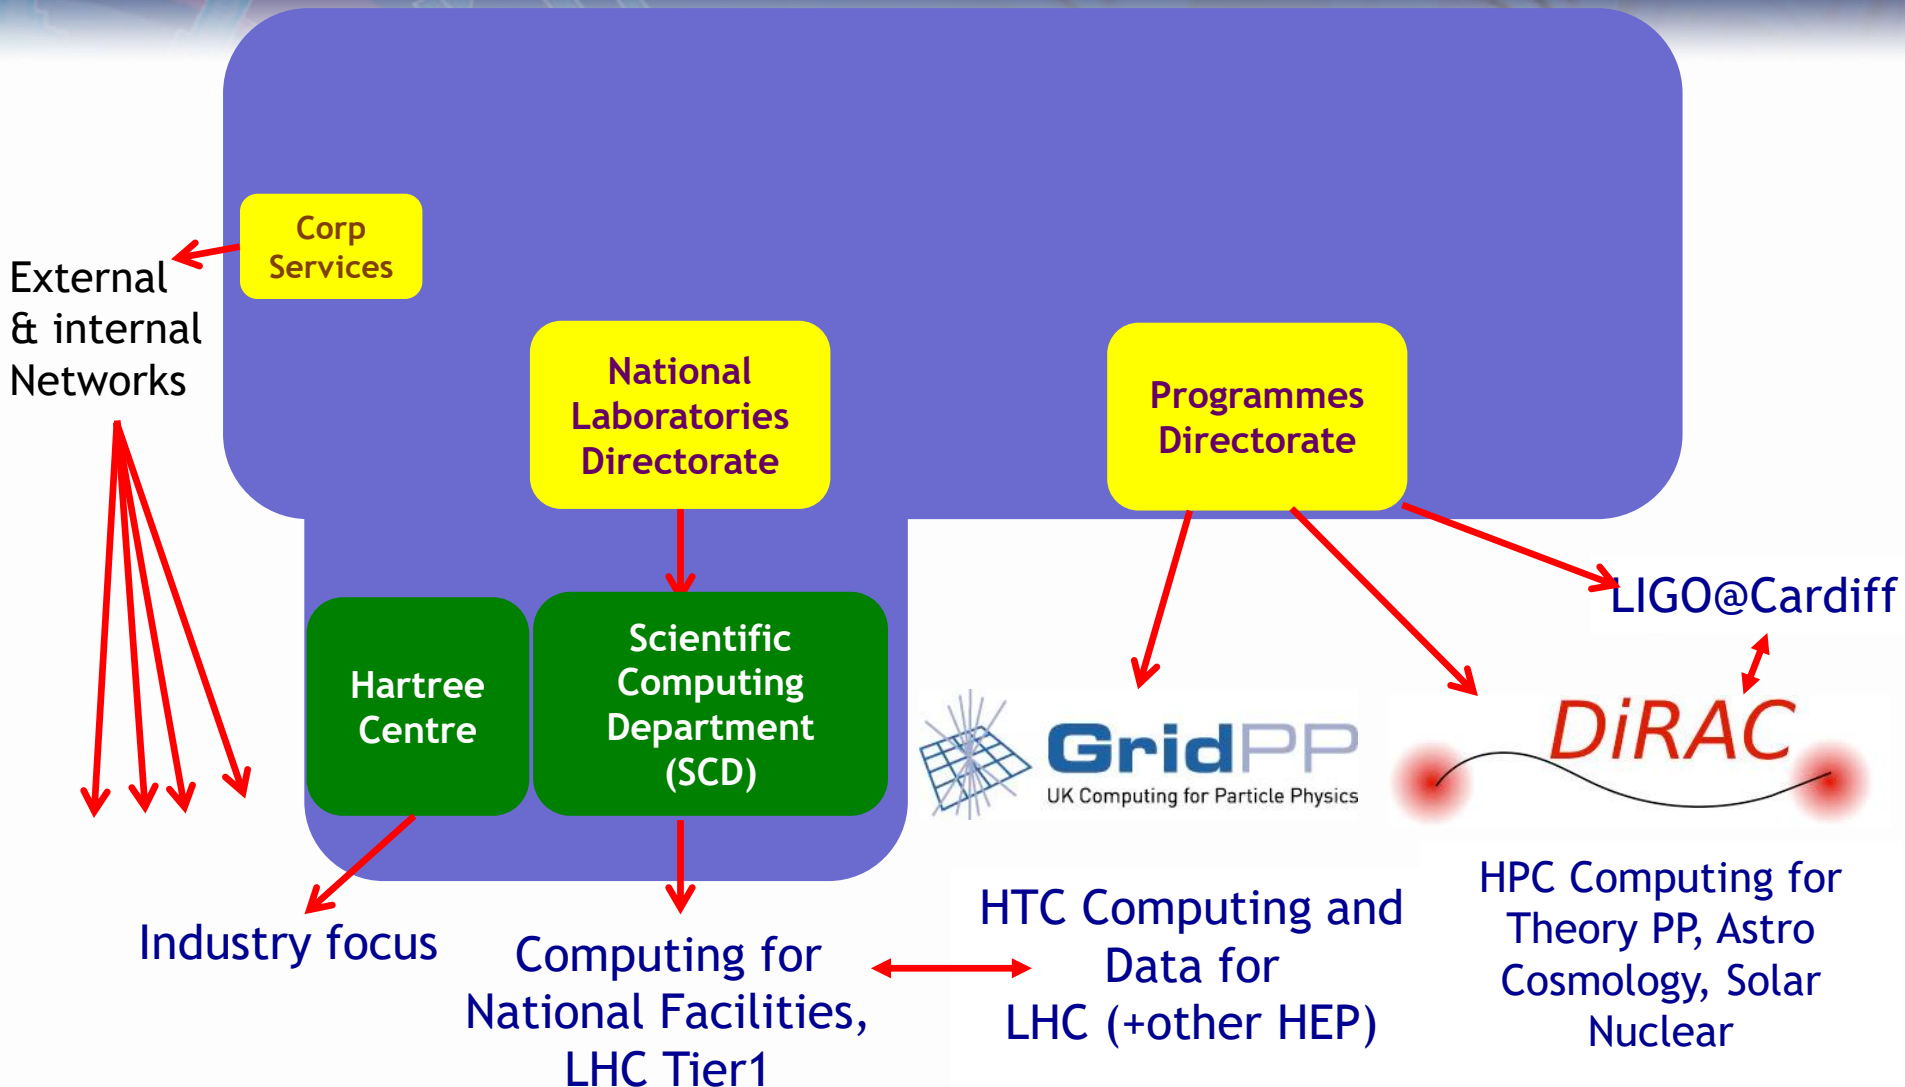
Runs National Facilities:

Diamond Light Source
ISIS (neutrons)
Central Laser Facility
RAL Space

Supports programmes in:

HEP,
Astronomy,
Astro-Particle, Nuclear







Sisters but not Twins

GridPP and DiRAC are complementary:

- GridPP provides HTC and DiRAC HPC resources.
- GridPP is a collaborative part of a bigger international endeavor (WLCG).
- DiRAC is a flagship UK project that competes on the international stage.
- GridPP provides distributed resources with a common (completely transparent and global) interface (software/middleware/access-layer).
- DiRAC provides focused resources with capabilities specialised to solve specific classes of problems with simpler AAI and site specific protocols.
- GridPP is funded as part of the programme with both capital and resources; DiRAC is funded, predominantly, by making capital cases to BEIS (or the equivalent) and a relatively small amount of resource funding from the programme.

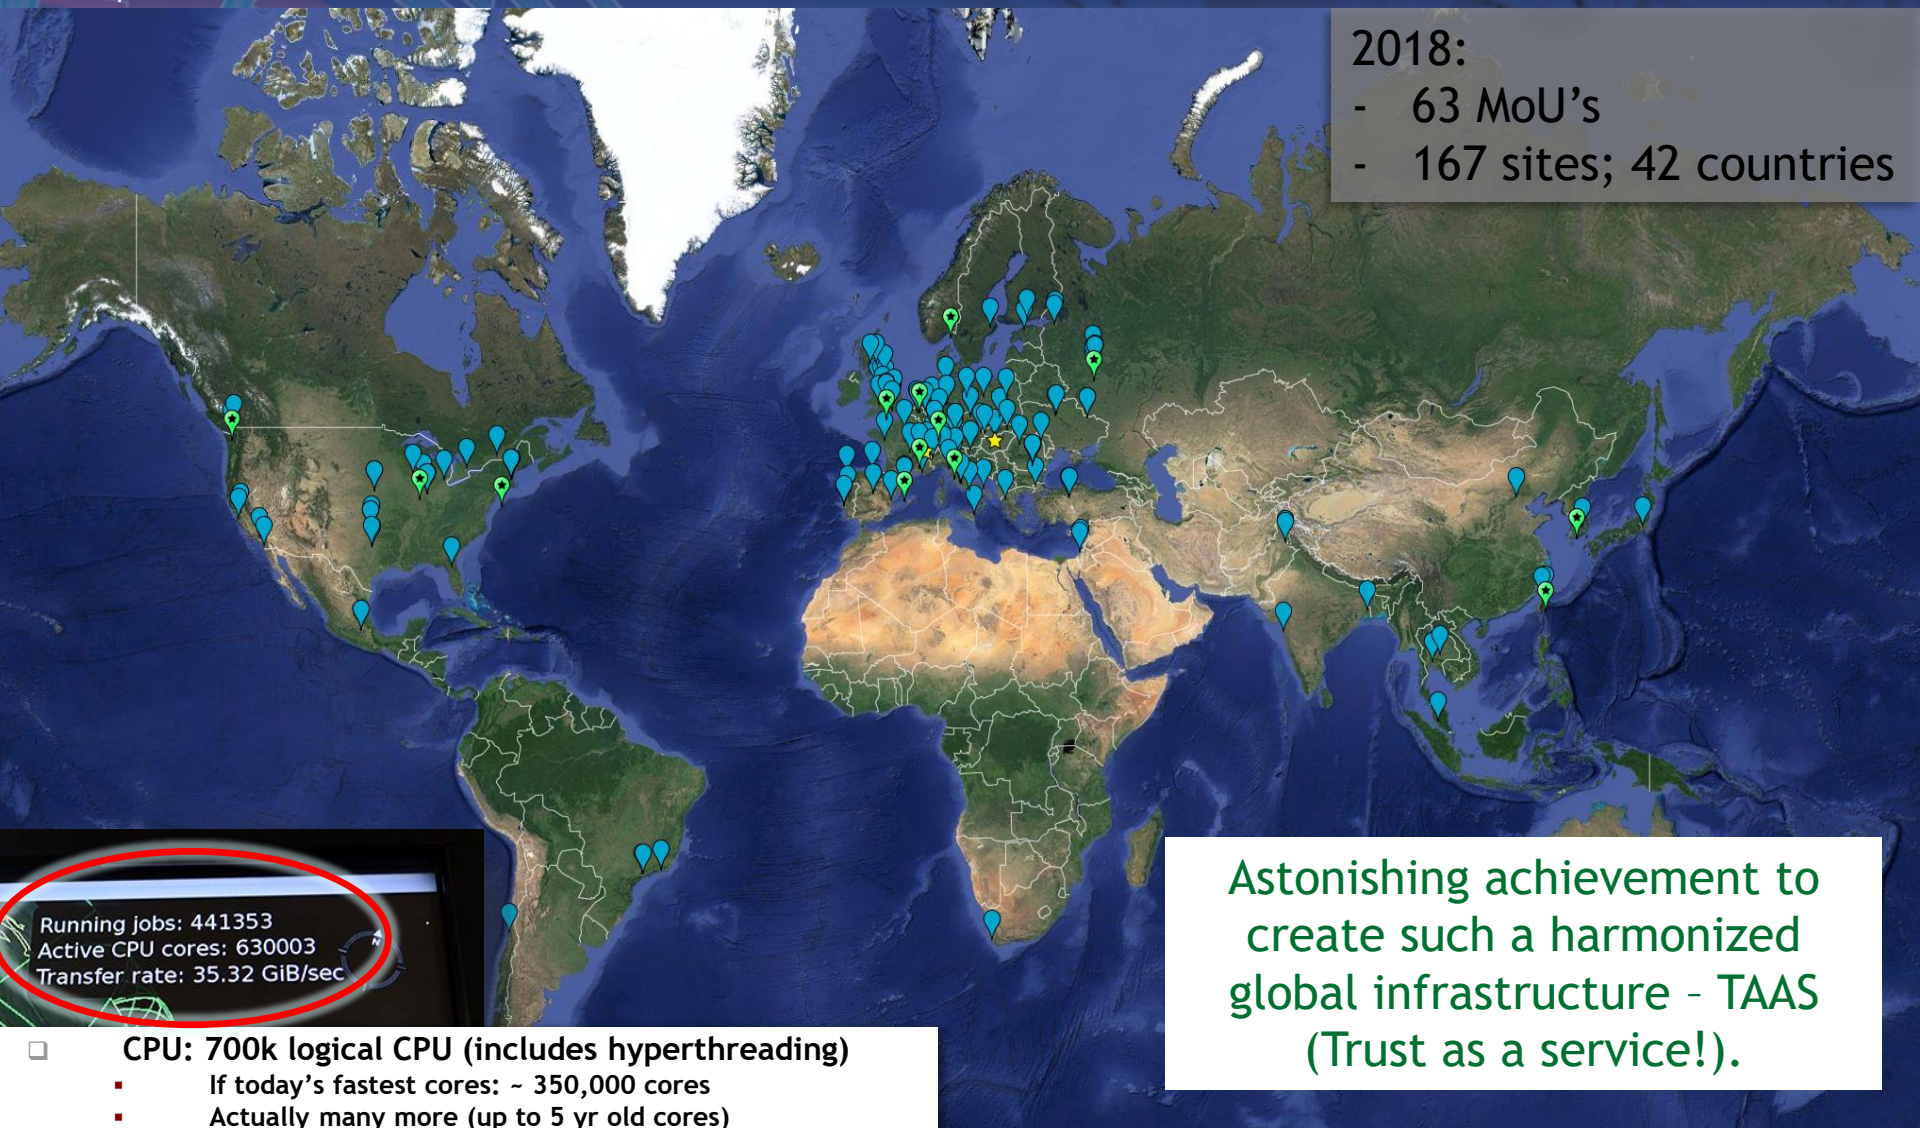


GridPP
UK Computing for Particle Physics

WLCG Collaboration

2018:

- 63 MoU's
- 167 sites; 42 countries



Running jobs: 441353
Active CPU cores: 630003
Transfer rate: 35.32 GiB/sec

Astonishing achievement to
create such a harmonized
global infrastructure - TAAS
(Trust as a service!).

- CPU: 700k logical CPU (includes hyperthreading)
 - If today's fastest cores: ~ 350,000 cores
 - Actually many more (up to 5 yr old cores)
- Disk 370 PB
- Tape 390 PB



A bit of Networking...

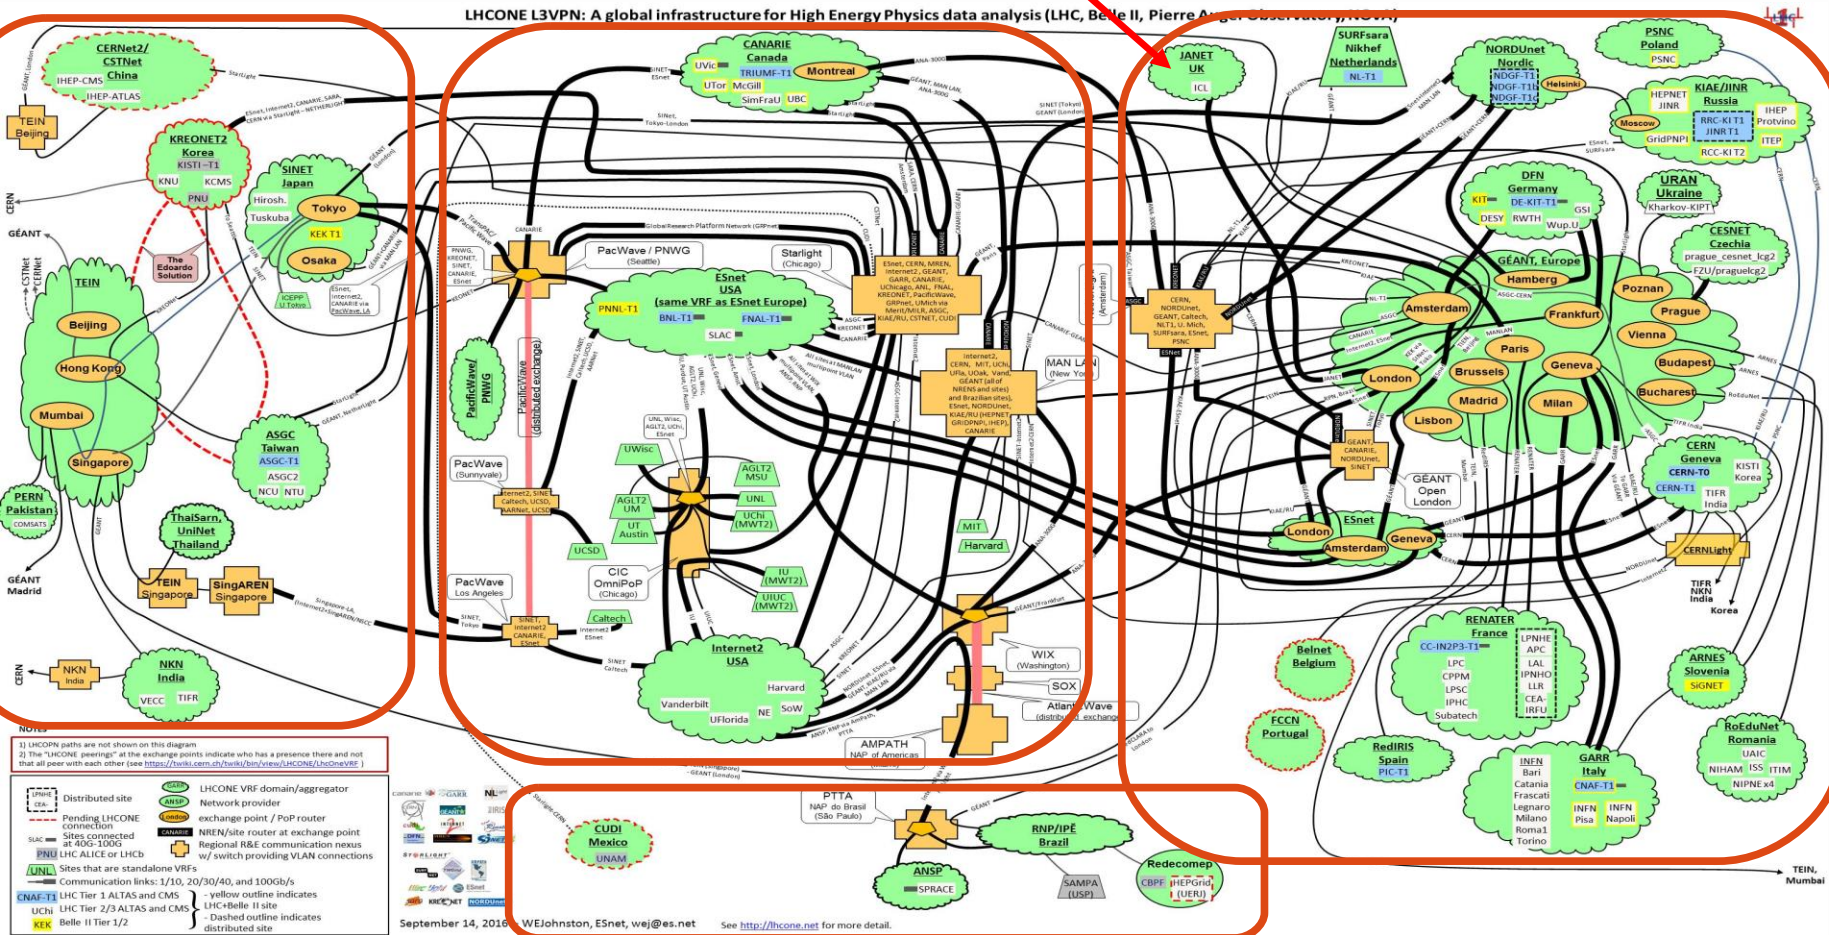
“Far out in the uncharted backwaters of the unfashionable end of the Western spiral arm of the galaxy...”

Asia

North America

Europe

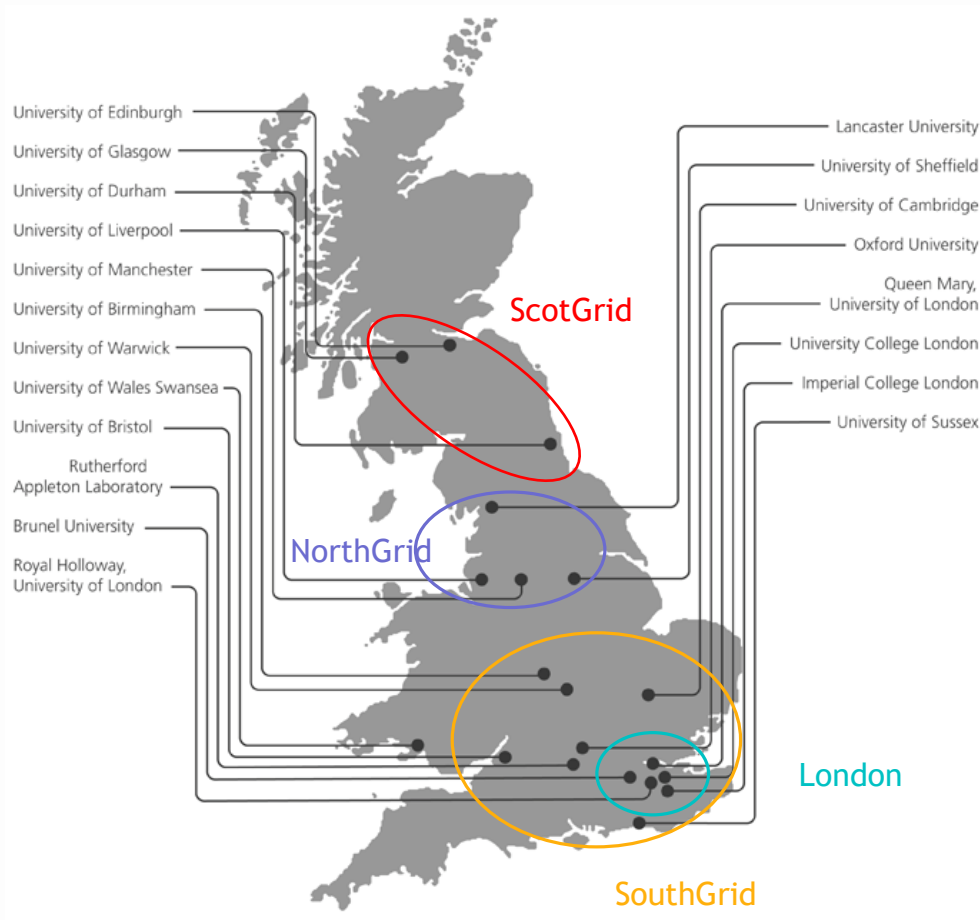
LHCONE L3VPN: A global infrastructure for High Energy Physics data analysis (LHC, Belle II, Pierre Auger Observatory, etc.)



South America



UK part of WLCG: GridPP




GridPP:

~10% of WLCG; 18 sites hosting hardware.

Tier-1 at RAL.

4 distributed Tier-2 centres.





GridPP Infrastructure

Future Trend



Layer	Experiment – 1	Experiment – 2	Experiment - 3	Community
A	Experiment specific computing support to maintain and develop reconstruction and analysis software	Experiment specific computing support to maintain and develop reconstruction and analysis software	Experiment specific computing support to maintain and develop reconstruction and analysis software	Community common software development projects
B	Experiment specific effort for production and data management	Experiment specific effort for production and data management	Experiment specific effort for production and data management	
C	Common distributed computing software infrastructure deployment, operation, support and development. Global services such as security response, monitoring and accounting			



GridPP





Interface: Users

- For our main (LHC) clients, our interface is really to expert teams within the experiment collaborations who, in turn, interface with their quite-expert users.
- For the smaller VOs (HEP and non-HEP) we have a well established set of standard tools that has been crystalized out of the software/middleware stack used by the big LHC experiments e.g:
 - **CVMFS: Global software distribution [NA62, EUCLID, LSST, LZ, CTA....]**
 - This is effectively a global read-only file system that solves the problem of VM and software distribution across a global infrastructure - worthwhile considering if you need this.
 - **DIRAC: Workload management - submit jobs to sites [NA62, LZ, Pheno, SNO+....]**
- Infrastructure provides lots of other needed stuff: Certificates (for authentication); VOMS (for authorisation); APEL (for accounting); GOCDB (for resources identification); operational security; etc.



Case Study: NA62

- NA62 is a “small” HEP rare-decay experiment at CERN.
- Only (ever been) a single part-time Grid person in collaboration.
- Built custom interface on top of standard GridPP tools (DIRAC, CVMFS...)



VO	Percent
atlas	43.6%
lhcb	27.0%
cms	13.4%
pheno	6.7%
alice	5.0%
lz	1.1%
na62.vo.gridpp.ac.uk	1.0%
ilc	0.7%
icecube	0.3%
cepc	0.3%
biomed	0.2%
snoplus snolah ca	0.2%

Nine UK sites (GridPP) + UCL (Belgium) + CNAF (Italy) + CERN.

~4000 job slots

~1PB of storage currently in use

~Million (6-8hr) jobs so far in 2018

One Grid admin... but we benefit from a wealth of expertise from GridPP and local sites

Elephants and rooms...

	Percent	Non LHC %	Driven by...
RAL-LCG2	31.5%	5%	
UKI-LT2-Brunel	3.5%	1%	
UKI-LT2-IC-HEP	8.0%	19%	LZ
UKI-LT2-QMUL	6.0%	10%	
UKI-LT2-RHUL	5.2%	6%	
UKI-NORTHGRID-LANCS-HEP	4.7%	0%	
UKI-NORTHGRID-LIV-HEP	3.0%	7%	
UKI-NORTHGRID-MAN-HEP	9.1%	5%	
UKI-NORTHGRID-SHEF-HEP	1.5%	22%	LZ
UKI-SCOTGRID-DURHAM	7.1%	62%	Pheno
UKI-SCOTGRID-ECDF	1.0%	0%	
UKI-SCOTGRID-GLASGOW	6.6%	7%	
UKI-SOUTHGRID-BHAM-HEP	1.2%	2%	
UKI-SOUTHGRID-BRIS-HEP	0.9%	9%	
UKI-SOUTHGRID-CAM-HEP	1.3%	2%	
UKI-SOUTHGRID-OX-HEP	4.2%	16%	NA62
UKI-SOUTHGRID-RALPP	4.6%	4%	
UKI-SOUTHGRID-SUSX	0.6%	1%	
Total	100.0%	11%	

- GridPP has had great success in bringing in new communities.
- But we now face contention for resources.
- And we need to address the issues of “on-boarding” and “ongoing-support”.
- GridPP5’s strategy was to attract other users.... who would bring other funding... we have half succeeded!
- IRIS funding (£4m x 4-years) is both an opportunity (hardware) and a problem (manpower).
- IRIS is NOT GridPP... together with DiRAC we are all part of the e-Infrastructure... which we need to optimise