



GridPP

UK Computing for Particle Physics

GridPP

Access for non-LHC activities

PPAP Meeting

Imperial, 24/25th Sep 2015

Pete Clarke
University of Edinburgh


GridPP Status

(see talk from Dave Britton)


- GridPP5 was recently renewed in the PPGP round.
- Resources were awarded at ~ 90% of flat cash.
- Features:
 - Tier-1 site at RAL remains.
 - Tier-2 sites will be consolidated into ~ 5 largish ones
 - Other Tier-2 sites retained at minimal staff support level
- GridPP strongly wishes to continue to support non-LHC activities

Non-LHC usage of GridPP today

T2K




Pheno




SNO+




ILC



Biomed



Hone




NA62



Fusion




CERN@School



See also: https://indico.cern.ch/event/299622/session/1/contribution/7/attachments/564613/777890/twhyntie_gridpp32_otherVOs_v1-0.pdf

slide from D.Britton's talk yesterday



9% of Tier-2 CPU
4% of Tier-1 CPU

used by

32 non-LHC VOs

between Jan 2012 and Dec 2014



The changing landscape

- Data rates are increasing very significantly across the science domains
 - No longer just LHC - SKA will be a major data source, others as well (DLS, Telescopes..)
 - It is a challenge to work out how STFC can support all of these !
- Funding realities
 - Flat cash or less ?
 - All countries are facing this
- EU-T0
 - European fundi
 - They all want to

**→ Do more - do it for less
- be more joined up**

a consortium.
support
- UK-T0
 - Initiative to join up STFC computing across science and facilities (SLIDE AT END)
- H2020, CSR
 - If funds are going to be accessible for computing, then this will only be for a more joined up approach.



Non-LHC activities : Future

- All of the foregoing leads to an increased mandate for GridPP to support non-LHC activities.
 - Part of GridPP5 brief from Swindon
 - This is great - it has always been the spirit of GridPP anyway.
- Formal position:
 - GridPP welcomes non-LHC activities to discuss sharing the resources
 - You are welcome to raise this through your local GridPP contacts if you have them
 - You can contact myself (peter.clarke@ed.ac.uk) or Jeremy Coles (jeremy.coles@cern.ch)
 - It is helpful if you could provide a ~few page document describing
 - your computing requirement
 - your resource requirement profile
 - Technical recipe already available on GridPP website
 - GridPP staff will then liaise with you to discuss timescales, get you going.
 - We will assemble a description of all of this for PIs on the web site
- Resources
 - In order to get going resources are provided within the ~ 10% allocation for non-LHC work
 - In you have a particularly large CPU and Storage resource requirement then in due course you will need to seek funding for the marginal cost of this - SEE LATER SLIDE

Non-LHC support : some of the common services

APEL (accounting/usage). VO Nagios (monitoring)

FTS
(bulk file
transfers)

VOMS
(authorisation)

CA
(authentication)

GridPP DIRAC (job submission
framework)
+
Ganga (for bulk operations)

resources
(hardware at
incremental cost)

CVMFS (software repository)

GGUS (support – help desk)/Documentation/Examples/User interface

+ access to GridPP expertise and experience

Ease of access to new communities


- Under the wider “UK-T0” banner it is obvious that to enable new/smaller communities in the future will also require development
 - A “single sign on” type AAA system (using University credential)
 - A “cloud” deployment (facility for you to deploy your virtual environment)
 - Easy to use services for managing and moving even larger data volumes
- There are no resources awarded under GridPP5 to develop all of this, but - at the margins we are trying
 - Some marginal RAL SCD effort as SCD have responsibilities for all of STFC science
 - H2020 projects such as AARC (authentication), DataCloud (cloud/virtualisation)
 - EGI funded staff work on community services
 - Shared GridPP-SKA and GridPP-LSST posts already in place.




GridPP
UK Computing for Particle Physics

Non-LHC activities ramping up

DIRAC




LIGO



LOFAR




LSST



LZ (Data Centre at IC)




QCD




PRaVDA (Proton Radiotherapy)



GalDyn



GHOST



- Pre-LSST

- Pilot activity using DES shear analysis at Manchester
- Joe Zunst (LSST) and Alessandra Forti (GridPP)



Galax



Job submission



- Fit a model to 10^{10} g
 - Maybe $\mathcal{O}(100)$ image
 - Time taken up to 1s
 - => 100s of millions of
 - Will need to speed up
 - Many many painful issues
- So far
 - Ganga Direct Submission
 - ~ 5500 with North
 - ~ 7000 with LSST
 - Brokering two choices
 - Dirac
 - Instance at Imperial
 - Bigpanda
 - In contact with developers

Using Ganga

- Submitting & managing jobs with Ganga
- Pros
 - Good job organisation
 - Many submission backends
 - Very scriptable
- Cons:
 - Could do with more documentation
 - Very CERN-focused
 - Sometimes loses track of jobs






GridPP


UK Computing for Particle Physics

Non-LHC support : DiRAC

- DiRAC Storage
 - Use of STFC RAL tape store for the DiRAC HPC
 - Lydia Heck (Durham) + GridPP staff enabled this
 - Excellent co-operation between GridPP and DiRAC



Blue Gene
Edinburgh




Cosmos
Cambridge




Complexity
Leicester



Data
Centric
Durham




Data
Analytic
Cambridge



Geant Human Oncology Simulation Tool I

One of our most recent use-cases has come from the STFC funded GHOST project for evaluating Late Toxicity Risk for RT Patients through the use of Geant 4 Simulation of X-Ray Dose Deposition. (see [this talk](#) from GridPP35)

The approach:





UK-T0 meeting

- UK-T0 is an initiative to bring STFC science communities together to address future computing and data centre needs
- First meeting arranged for non-pure-PP communities on Oct 21/22 at RAL.
(pure PP communities are already part of GridPP (T2K, NA62, ILC..))
- To discuss:
 - Sharing of the infrastructure and services where this makes sense.
 - How to ease access to smaller communities.
 - How to go for funding opportunities in both UK and EU
- Contacted so far
 - LOFAR, LSST, EUCLID, Advanced-LIGO, SKA, DiRAC, Fusion (Culham), LZ, CTA, Facilities computing.
- If there are other experiments/projects/activities interested - please contact me at the end of the meeting.

Practicalities and caveats

- There is no magic wand
- GridPP5 has been at flat cash for 8 years → 19% reduction in resources.
- Non-LHC activities are typically not awarded computing capital resources by PPRP, and in some cases asked to talk to GridPP
- The incremental capital cost of CPU and Storage for these activities falls between the cracks
 - If 10 non-LHC activities require 10% of GridPP → would double the resource requirement !
 - Mitigated by leverage at Tier-2 sites.
 - This is as yet an unsolved situation, but we have ideas.
- Some key software services which would have helped other smaller communities have had their support cut (e.g. Ganga)
- GridPP is seeking capital resources from outside the science line aggressively
 - Lobbying for CSR capital injection
 - Working hard to be involved in H2020 bids



GridPP
UK Computing for Particle Physics

Questions ?