


# UK Status and Plans

Catalin Condurache - STFC RAL  
ALICE Tier-1/Tier-2 Workshop  
University of Torino, 23-25 February 2015



Science & Technology Facilities Council  
Rutherford Appleton Laboratory



**GridPP**  
UK Computing for Particle Physics



# Content

- UK GridPP Collaboration
- Tier2s Status and Plans
  - Birmingham
  - Oxford
- RAL Tier-1 Centre
  - Last 12 months
  - Status
  - Plans and future
- Updates on IPv6





# UK Grid Collaboration

- GridPP is a collaboration of 19 UK universities + RAL + CERN with the primary goal of providing computing resources to LHC particle physics experiments



# GridPP resources

| Site                       | CPU kSI2k (vs 03/2014)  | Storage TB (vs 03/2014)                  |
|----------------------------|-------------------------|--|
| <b>RAL</b>                 | <b>35,596 (24,107)</b>  | <b>10,389 + 13,612 (10,517 + 11,390)</b> |
| Edinburgh                  | 21,624 (9,015)          | 1,186 (355)                              |
| Glasgow                    | 10,992 (8,472)          | 1,402                                    |
| Queen Mary, London         | 7,946 (7,946)           | 1,675                                    |
| Imperial College, London   | 26,872 (6,833)          | 2,872                                    |
| Lancaster                  | 7,820 (5,210)           | 1,729 (1,135)                            |
| Manchester                 | 8,920 (5,798)           | 1,430 (951)                              |
| Sheffield                  | 3,264 (2,540)           | 521 (363)                                |
| <b>RALPP</b>               | <b>8,940 (6,200)</b>    | <b>2,686 (1,607)</b>                     |
| Royal Holloway, London     | 6,946 (3,704)           | 1,768 (728)                              |
| Brunel, London             | 5,412 (4,027)           | 634                                      |
| <b>Oxford</b>              | <b>4,197 (3,107)</b>    | <b>1,219 (709)</b>                       |
| Liverpool                  | 4,337 (2,840)           | 784 (544)                                |
| <b>Birmingham</b>          | <b>2,130 (2,130)</b>    | <b>600 (389)</b>                         |
| Cambridge                  | 932 (707)               | 350 (278)                                |
| Sussex                     | 502 (502)               | 82 (77)                                  |
| Bristol                    | 2,610 (1,467)           | 209 (120)                                |
| University College, London | 502 (502)               | 215                                      |
| EFDA Jet                   | 332 (332)               | 2  |
| Durham                     | 2,491 (2,419)           | 53                                       |
| <b>TOTAL</b>               | <b>162,365 (91,566)</b> | <b>29,733+13,612 (23,897 + 11,390)</b>   |



# Tier2s Status and Plans

- **Birmingham**

- in the last 12 months
  - from 816 cores to 1216 cores (8121 HS06 to 12489 HS06)
  - from 389 TB to 600 TB
- ~69% of UK T2 ALICE CPU allocation
- currently CREAM, ready for ARC if ALICE happy
- storage for ALICE
  - from 120 TB to 280 TB
  - native XRootD


# Tier2s Status and Plans

- Birmingham
  - 58% overall CPU fairshare for ALICE
    - ATLAS 28%, LHCb 9%, others 5%

Developed by CESGA "EGT View": / normcpu / 2014:1-2014:12 / SITE-VO / all (x) / GRBAR-LIN / 1

2015-02-02 18:38

UKI-SOUTHGRID-BHAM-HEP Normalised CPU time (kSI2K) per VO






# Tier2s Status and Plans

- Birmingham

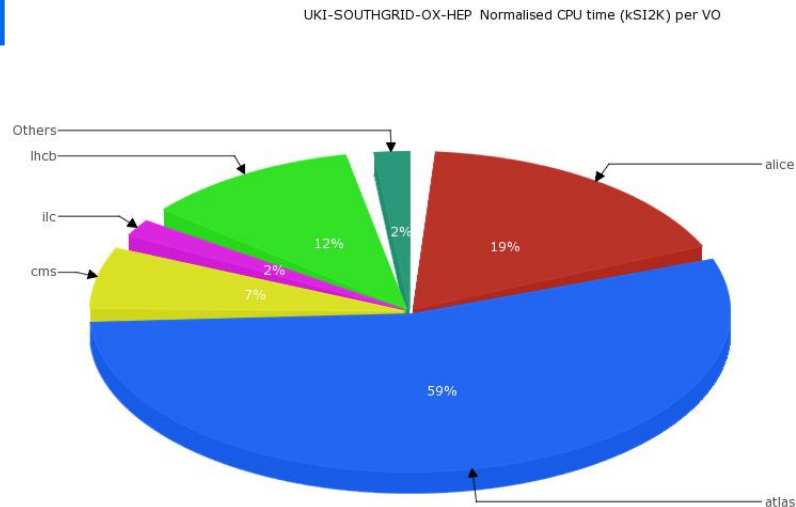
- running and queued jobs since October 2014



# Tier2s Status and Plans

- Oxford

- need to supplement the support given by B'gham
- ~31% of UK T2 ALICE CPU allocation
- no storage provided







# Tier2s Status and Plans

- Oxford
  - soft limit - 150 jobs
    - no upper hard limit (if cluster is free)
    - very efficient to fill the empty job slots
  - gradual move from CREAM+torque to ARC+condor
    - two thirds of the Grid Cluster now running HT Condor behind ARC CE
    - remaining third running legacy torque/maui driven by CREAM CE
  - will continue to provide only CPU



# RAL Tier-1 Centre

- ...in the last 12 months - Hardware
  - CPU: ~127k HS06 (~13k cores)
  - storage: ~13PB disk
  - FY14/15 procurement
    - expected 6PB and 42k HS06
    - storage capable of both Castor and CEPH
      - Extra SSDs for CEPH journals
    - 10GbE for WNs
  - tape: 10k slot SL8500
    - migration to T10KD tapes (finished - Jan 2014)
  - RAL Tier-1 network upgrade - new pair of routers



## RAL Tier-1 Centre

- ...also in the last 12 months - Middleware
  - migration to EMI3 - early 2014
  - CREAM-CE still kept for ALICE, small VOs
    - only ALICE - since January 2015
  - FTS2 stopped - September 2014
  - various developments on HT Condor, ARC
  - ALICE jobs increased - August 2014
    - 3500 - since October 2014




# RAL Tier-1 Centre

- ...more in the last 12 months - Storage
  - Castor upgrade to new major version (2.1.14) - Jun 2014
    - various improvements (disk rebalancing, xroot internal protocol)
    - upgrade complete
  - Castor logging system with Elasticsearch
  - Draining disk servers still slow - major production problem
  - CEPH
    - Evaluations continue on small test cluster
    - SSDs for journals installed on cluster nodes
      - mixed performance results - needs more study



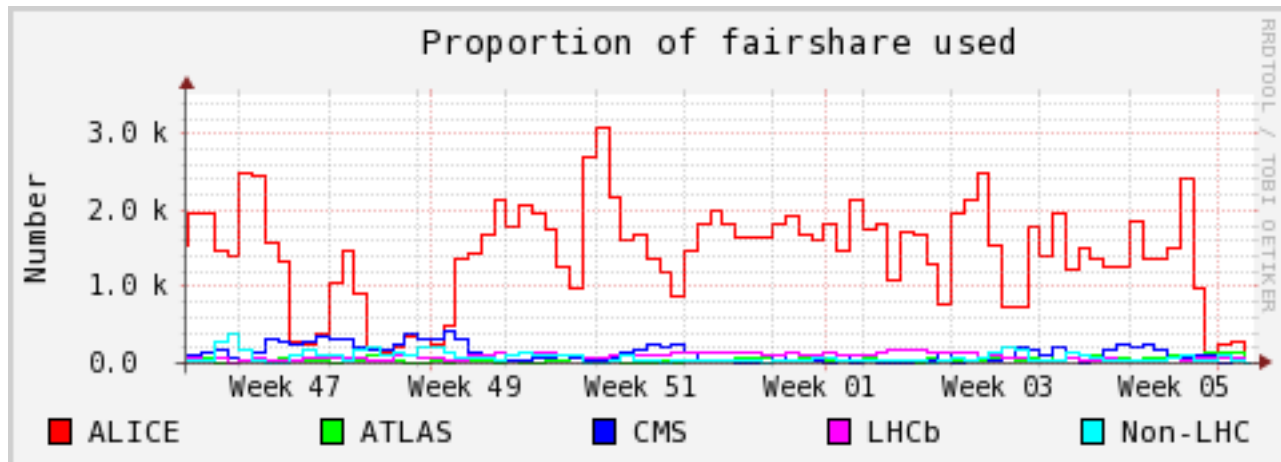
# RAL Tier-1 Centre - More on ALICE

- CPU fairshare - 1.9% (2200 HS06)
- spare cycles up to 3500 jobs
- CPU efficiencies
  - 90 - 91% since August 2014
  - 2<sup>nd</sup> among LHC VOs
- 200 TB disk



## RAL Tier-1 Centre - More on ALICE

- Good use of a significant amount of opportunistic CPU
- In the graph below, the expectation is 100 for everyone
- Average for ALICE - at least 10 times the fairshare!!





## RAL Tier-1 Centre - More on ALICE

- Room to increase the share up to 4K+ jobs, but...
  - high ALICE zombie process numbers
  - risk to get alarms on a too high number of ‘zombies’ per WN made RAL Tier-1 to stay conservative
  - some feedback from ALICE, but no final resolution




# RAL Tier-1 Centre - More on ALICE

- Vulnerable Apache and OpenSSL releases within ALICE software
  - every now and then Networking@RAL query specific ports on RAL Tier-1 hosts
  - reports are issued for each host found vulnerable (i.e. ALICE VOBOX nodes) and actions are expected to be taken
  - some feedback from ALICE...



# RAL Tier-1 Centre - More on ALICE



After weeks and months of testing...


...first production jobs on ARC@RAL - Wed 8 October 2014

Smooth migration of production jobs from CREAMs to ARCs...

but...

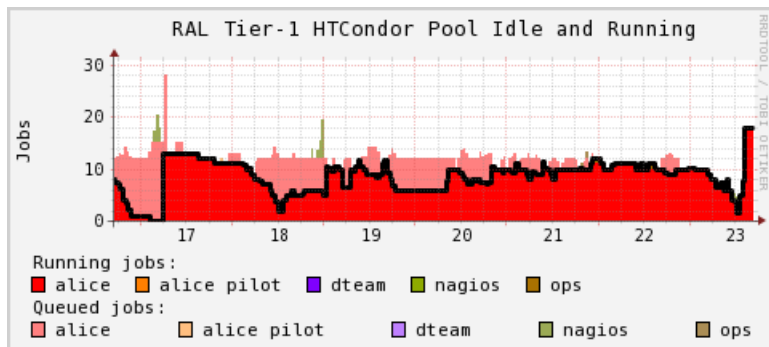


# RAL Tier-1 Centre - More on ALICE



...ALICE SAM tests still rely on CREAM (and WMS) at RAL

Not really worthy to keep 2 CREAMs and one CLUSTER node (+ another one in standby) in production (and part of the Call Out system) for ~30 jobs





## RAL Tier-1 Centre - More on ALICE

- Deadline to get rid of CREAM not yet met (mid ~~January~~, ~~February~~, March 201...)
- The best outcome of this workshop for RAL Tier-1
  - decision to terminate use of CREAM at RAL with immediate effect (or by 1<sup>st</sup> March)
  - any chance?



# RAL Tier-1 Centre - immediate future

- **Storage for ALICE**
  - 356TB (from 200TB) - by April 2015
  - the extra capacity will be in Castor
- **CPU for ALICE**
  - 10% more CPU
  - expect the amount of work to drop though...
    - ...as less spare CPU capacity available...
    - ...as other LHC experiments will increase their usage



# RAL Tier-1 Centre - More Plans

- Deployment of a new storage endpoint
  - CEPH as backend
  - additional storage capacity could be provided
    - but not all the traditional Grid functionality
    - only GridFTP and xrootd (ALICE specific xrootd security model?)
    - also S3 and Swift API access



# RAL Tier-1 Centre - Future

- GridPP project
  - currently GridPP4+ (1 year extension of GridPP4) - until March 2016
  - GridPP5 proposal underway
    - April 2016 - March 2020



# RAL Tier-1 Centre - Future

- ALICE @ (UK & RAL Tier-1)
  - by the time GridPP5 starts - ALICE supported for 3 years without any explicit funding to do so
  - however funding in place until March 2016
  - HW, other resources available during Run2 in 2016
  - the experiment features into the GridPP5 proposal
    - ~2% of the Tier-1 and Tier-2 global requirements
  - do not know yet the outcome



# IPv6 Updates

- **RAL**
  - IPv6 new testbed
    - separate from main production network
    - plan to install a dual stacked UI box and a perfsonar box
  - requirement on all Tier1s to provide IPv6 connectivity to the perfsonar machines by April 2015
    - will not be met at RAL
    - first need to move OPN links from UK Light router to the Tier-1 routers





# IPv6 Updates

- **Birmingham**
  - IPv6 addresses available at University level, but not yet for Tier-2
  - plans to get some services running in next 6 months
    - will greatly depend on central IT (routing etc)
- **Oxford**
  - IPv6 service is bandwidth limited
  - some test services available on either dual-stack or IPv6 only
  - plans for a production ready IPv6 system in ~1 year



**Any (other) questions?**

**Thank You!**