

# HEPiX Fall 2013 Highlights

Michel Jouvin  
LAL, Orsay  
[jouvin@lal.in2p3.fr](mailto:jouvin@lal.in2p3.fr)

LCG France, Lyon, November 2013





# HEPiX Facts

- 22 years old, self-organized forum of people involved in “fabric management” in HEP
  - More sysadmins than IT managers
  - Open to anybody interested
  - <http://www.hepixon.org> : archive of all meetings
- Mix of large and small sites, many involved in WLCG
  - Better understanding of each others
  - Each one benefit of the others
- 1-week meeting twice a year
  - Last one at Ann Arbor (UMich) end of October
- Increasing number of WG on specific issues to reach consensus or tackle specific issues when needed



# Ann Arbor Meeting at a Glance

- 115 participants: the largest US meeting
  - > 20+ countries including Australia, 42 affiliations
  - > Strong US participation: many new comers (WLCG T2/T3)
- Dense agenda
  - > No dominating topic, all usual threads
  - > <http://indico.cern.ch/conferenceTimeTable.py?confId=247864#all.detailed>
- Difficulty to establish a permanent track around green computing and business continuity
  - > No talk submission for the Ann Arbor meeting... doesn't mean it is dead!

# IPv6

- Lots of progresses: many sites have joined the WG testbed, TF created in WLCG Ops Coord
  - Main activity is testing dual stack services
  - Ensure that IPv6 doesn't break IPv4, assess IPv6 service readiness
  - <http://hepixonweb.cern.ch/wlcg-applications>
- IPv6 in the world: Google reported 2.3% IPv6 access
  - Switzerland: most IPv6 friendly country (11%)!
- CERN is "IPv6 ready": first IPv6-only WNs in 2014
  - Non IPv6 services (e.g. AFS) accessed through private IPv4
- IPv6 GridFTP: 2 PB in 6 months, 87% success rate
  - Very good result, best-effort infrastructure, misconfiguration...
- IPv6-only DPM at 2 sites in UK used by dual-stacked FTS3
- IPv6 workshop planned at CERN Spring/Summer 2014



# Security & Networking

- Security: growing activity around identity federation and federated identity management
  - German project demonstrating feasibility even for non Web apps
- 100 Gb is now available for WAN in several places, in particular in the US
  - Changing dramatically the traditional division between LAN and WAN
  - Several demonstration of efficient usage
- BNL mentioning looking at IPoIB as an alternative to 10 GbE
  - Potential cost advantage
  - Preliminary test demonstrated 40 Gb over 56 Gb QDR
  - Lower latency (10x) than 10 GbE: open the path for new use cases

# Storage

- ◉ WG mandate ended
  - Due to personal constraints of A. Maslennikov, no final report given this time
- ◉ Very interesting talk by Western Digital on drive reliability
  - Look at slides: highly technical!
  - New features to improve future failure prediction
- ◉ openAFS refuses to die!
  - No clear alternative yet, some sites shutting it down (e.g. RAL)
  - No plan for IPv6 support
- ◉ Hype peak for CEPH: several successful pre-prod deployments
  - CERN, RAL
  - Currently, mainly distributed object storage (block devices)



# Benchmarking

- Lots of HS06 result updates on HEPiX wiki
  - <http://w3.hepik.org/afs/hepik.org/project/benchmarks/doku.php>
  - Several SL6 results but HS06 still defined by SL5/32-bit for comparison with previous results (pledges)
- SL6: +5% on similar HW compared to SL5 for recent procs
  - Benefit varying with processor generations
- Last Intel processors (Ivy Bridge): 25% HS06/watt improvement
  - Good scalability of consumption and performances with the number of cores
- New SPEC benchmark announced for Fall 2014
  - Restart beginning of next year work with experiment to assess scalability of HSxx with HEP apps
    - Has been very time consuming to prepare HS06
    - Take advantage of the LS1



# Batch Systems

- ◉ Regularly discussed since a few years
  - A “WG” started by CCIN2P3 a few years ago around GE
- ◉ Several (large) sites moved to GE
  - All moved at the end to UNIVA GE
    - Oracle transferred GE ownership to UNIVA
  - Open-source versions seem almost dead... or dying
- ◉ SLURM hype is over
  - Several disappointing scalability tests: good scaling with high number of nodes but not with high number of jobs (100+ Kjobs)
- ◉ Several sites looking at HTCondor
  - Scalability seems impressive, dynamic membership
  - Successfully used at OSG sites for many years
  - RAL moved its production CE, CERN planning to move from LSF
    - RAL chose ARC CE but CREAM CE working too





# Configuration Management

- More and more sites using Puppet
  - Cfengine sites or sites without a configuration tool
  - Every site has deployed some sort of modules... not always very generic
  - Many sites ready to share their work... but how to decide between conflicting pieces?
- HEPiX WG about Puppet created at the previous meeting
  - Slow start... first real meeting after the Ann Arbor meeting
  - Urgency if the group wants to come up with a set of generic modules that can be reused without development by other sites
    - Quattor has been successful on this but required a lot of effort and sites with their own configuration have never migrated...



# Upcoming Meetings

- Spring 2014: Annecy (LAPP)
  - > 19-23/5/2014
  - > <http://indico.cern.ch/conferenceDisplay.py?confId=274555>
- Fall 2014: Univ. of Nebraska (CMS T2)
  - > Dates to be set soon



# Conclusions

- HEPiX remains healthy: not a surprise for regular attendants!
- HEPiX is open to every site interested
  - > No formal membership
  - > Register to the mailing list: [hepixonhepixon.org](mailto:hepixonhepixon.org)
  - > Participate to meetings
- Hope to see **more French sites joining at the LAPP meeting**