

HEPiX Spring 2009 Highlights

Michel Jouvin

LAL, Orsay

jouvin@lal.in2p3.fr

<http://grif.fr>

June 10, 2009

GDB, CERN



HEPiX at a Glance

- 15+ years old informal forum of people involved in “fabric management” in HEP
 - No mandate from any other body
 - More sysadmins than IT managers
 - Open to anybody interested
 - <http://www.hepix.org> : archive of all meetings
- Main activity is a 1-week meeting twice a year
 - 60-100 attendees per meeting, “stable” core
 - Focus : exchange of experience and technology review
- Mix of large and small sites
 - Better understanding of each others
 - Each one benefit of the others
 - Most of the sites involved in grid computing
- “On-demand” working groups
 - Currently distributed file systems

HEPiX Extending

- Last 2 meetings in « new » places:
 - Fall 2008: ASCG
 - Spring 2009: Umea (Sweden)
- In both cases, successful to bring new sites in
 - Asia : lot of participation from all around Asia
 - Including Australia and US!
 - Several sites from Nordic countries with many interesting contributions
- Several other European sites joined
 - PIC, Lisbon (LIP T2), CNAF...
 - Next Spring meeting in Lisbon!
- Still attracting new sites out of HEP
 - Univ. of Gand (Belgium)
 - Generally because of people who spent some time in a HEP lab

Main Topics

- Main topics are always the same but main focus changes at each meeting
 - HEPIX value is the periodic report on work in progress
 - Umea: focus on virtualization
- Main tracks
 - Site reports : very valuable part, update on changes, give a « picture » of what's happening around the world
 - Scientific Linux status and future directions
 - Data Centers : cooling, power consumption, ...
 - Less active, projects in building phase
 - Storage : convened by File System WG
 - LUSTRE growing
 - Evolving as a storage-focused forum to share experience
 - Virtualization
 - Benchmarking : still active, new activities related to virtualization
 - Security and networking

Virtualization

- A track at each meeting for at least 2 years...
 - Initial focus was mainly on service consolidation
- Coverage extended to virtualized environments for applications
 - Virtualized WN
 - Integration with batch schedulers
 - Image management and resource allocation (openNebula)
 - Grid and clouds (eg. StratusLab)
 - CERNVM: very minimal and generic VM approach
- Position of sites about VO-maintained images
 - "First" discussion rather decision : sites not comfortable but no "No"
 - Feasibility of a limited number of images per VO ?
 - Part of SW area ?
 - GDB is the place for further discussions
- Benchmarking: CPU, IO, Xen vs. KVM...

File Systems WG

- Set up 2 years ago with a mandate to review distributed file systems technologies
 - Initial mandate from former IHEPCCC
- Continued since on a voluntary basis with 2 objectives
 - Benchmarking activities with more realistic and diverse use cases: LUSTRE outperforming (2x) all other solutions in every case so far.
 - Share experience and expertise with new technologies
 - Mainly LUSTRE in fact as it is used in several places now
- New members joined: CERN (LUSTRE evaluation), FNAL, Lisbon
- Producing a report at each HEPiX meeting
 - This meeting new topic: potential AFS + LUSTRE combination

Miscellaneous

- iSCSI evaluation at CERN
 - Alternative to FC ?
- SLURM, an alternative to Torque (MAUI?)
 - Command-compatible with Torque
- First benchmarks of Nehalem-based machines
 - 50% improvement in power efficiency compared to Hapertown
- CERN R&D on network anomaly (CINBAD)

Conclusions

- HEPIX is a very « useful » forum opened to any site interested
 - Complementary to GDB, focused on fabric management rather than grid services
 - No formal membership : just register to the meeting..
- Next meeting in Berkeley, October 26-30th
 - Look for announcement soon at <http://www.hepix.org>
 - Ask to be registered on HEPIX mailing list (low volume...)
 - New track on monitoring tools and practices ?
- Material produced is mainly presentations during the workshops
 - Look at agendas if interested by a presentation. Start at <http://www.hepix.org>