

HEPiX Spring Highlights

Michel Jouvin

LAL, Orsay

jouvin@lal.in2p3.fr

<http://grif.fr>

June 11, 2008

GDB, CERN



- 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 and benchmarking with

- Focus on LHC preparation and computing
 - Held 2 weeks after WLCG Collaboration workshop
- Main topics (between ½ and 1 day per topic)
 - Site reports : very valuable part, update on changes, give a « picture » of what happening around the world
 - Scientific Linux status and future directions
 - Discussion reviewing main priorities
 - Data Centers : cooling, power consumption, ...
 - Storage : convened by File System WG
 - Preliminary “final” report
 - Presentation from Google on their data management infrastructure
 - Unbearable slowness of tape by C. Curran
 - Benchmarking : what is the right benchmark ?
 - First conclusions of benchmarking effort
 - Networking : LHCOPN, Atlas online farm network
 - Security : several talks by Roman on new threats...

- Connected to WLCG MoU based on SI2K
 - Si2K no longer available and maintained
 - Find a benchmark accepted by HEP and others as many sites served different communities
- Initial phase : review of existing benchmarking practices (CERN, FZK, INFN, ...)
- Last 6 months : setup of benchmarking testbed with dedicated HW at CERN and others
 - Covering of wide range of processor with typical HEP configuration (2GB/core)
 - Run of different standard (SPEC) benchmarks with agreed flags, mainly SL4 32-bit/gcc 3.4
 - Look at SL5, 64-bit, gcc4
 - Run of benchmarks from LHC experiments and comparison of scaling with SPEC
- No major scaling problem with either SI2K or SI06
 - Should allow a smooth transition

- 3 phases:
 1. Questionnaire about storage at T1s
 2. Technology review and selection
 - Posix FS (TFA) : LUSTRE, GPFS, AFS
 - SRM : CASTOR, dCache, DPM
 - Xrootd
 3. Performance comparison between selected technologies
- Testbed setup at CERN with 10 servers and 60 8-core clients with 1 Gb/s connection, 4-5 6TB
 - 480 simultaneous client tasks
 - 3 tests : writing, sequential read, pseudo-random read
 - Most implementations able to sustain wire-speed in writes and sequential reads
 - Significant performance advantage for LUSTRE in pseudo-random reads but must clarify test conditions
 - Use case may be an advantage for LUSTRE client-side caching

- Both WGs willing to continue
 - Good relationship established between participants
 - Still many things to evaluate or review
 - Gained experience running an appropriate testbed
- Initial mandated ended, applying for a new mandate from IHEPCCC...
 - Mandate from an official body important for participant to justify time spent and to obtain loan of HW resources from sites
 - IHEPCCC seen as a broader spectrum than WLCG : important to take into account non-pure HEP context at most sites
- But IHEPCCC future seems uncertain
 - Reason unclear... but not under our control!
 - WGs are probably one positive work from IHEPCCC as an attempt to identify new issues or elaborate "consensus"



- Debate started during last meeting in St Louis
 - First meeting in a non HEP place
 - Several genomics centers attended
 - In fact, already 1 from CEA Tera10 (19 + 26 in Top500 C)
- General agreement opening HEPiX to non-HEP is a good move
 - Lots of common problems, different experiences
 - E.g. global file systems (LUSTRE, GPFS...)
 - Many HEPiX sites collaborating with other communities
- Question is not participation to meetings but involvement in HEPiX Board
 - No real power...
 - In charge of defining HEPiX orientations
 - Support received from IHEPCCC
- A few afraid of not being HEP-focus anymore..

- HEPIX is 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 at Taipeh, October 20-24th
 - First HEPIX meeting in Asia, looking for involvement of new people/sites
 - Look for announcement soon at <http://www.hepixon.org>
 - Ask to be registered on HEPIX mailing list (low volume...)
- Spring 2009 meeting should be in Nordic countries
 - Another extension to a region not very much involved so far
- Agreement to open HEPIX to non HEP participation
 - Offering them to participate to HEPIX board if they are regular attendees with active participation and interest

