



HEPiX Fall 2017 Workshop

KEK, Tsukuba, Japan

Helge Meinhard, CERN-IT
WLCG GDB
08 November 2017

Contributions by Manfred Aef, Tony Cass, Ian Collier, Sebastien Gadrat, Pete Gronbeck, Michel Jouvin, David Kelsey, Edoardo Martelli, Shawn McKee, Michele Michelotto, Ofer Rind, Jingyan Shi, Mattias Wadenstein, Tony Wong

HEPiX Mission

- From our Web site <https://www.hepix.org>:
“The HEPiX forum brings together worldwide Information Technology staff, including system administrators, system engineers, and managers from the High Energy Physics and Nuclear Physics laboratories and institutes, to foster a learning and sharing experience between sites facing scientific computing and data challenges.”
- Emphasis is on site services as opposed to experiment software and middleware
- Originating in HEP (particle physics), but open to other sciences
 - Recent participation from life sciences, photon/material sciences, ...
 - Just attend the workshops (see below)

HEPiX Activities: Workshops

- Twice per year, one week each, in Europe, North America and Asia
 - Typically 100...140 participants
 - Autumn 2017 at KEK was the 56th workshop since 1991 - third Asian workshop after ASGC (Taipei) 2008 and IHEP (Beijing) 2012
- Workshop style
 - All plenary, no parallel sessions
 - Abstracts are usually accepted - no formal Programme Committee reviewing
 - No proceedings
 - 70...75 contributions per workshop
- Honest exchanges about experience, status and plans
 - People report successes as well as failures

HEPiX Activities: Working Groups

- Active
 - Migration to IPv6
 - CPU benchmarking
 - Batch system monitoring
 - Network function virtualisation
- Completed (selection)
 - Data preservation
 - Configuration management
 - Storage
 - Virtualisation
 - Unix environment scripts
 - AFS
 - Security

HEPiX Governance (*too strong a word*), Contacts

- HEPiX board: Representatives of major labs as well as organisers of past successful workshops
 - Currently 26 sites represented
- Two co-chairs: Helge Meinhard (CERN), Tony Wong (Brookhaven National Laboratory)

- Web site: <https://www.hepixon.org>
- Mailing list: hepixon-users@hepixon.org (moderated)
 - To subscribe, send an e-mail to listserv@listserv.hepixon.org with an empty subject and the following line in the e-mail body:
`subscribe hepixon-users <your first name> <your family name>`
- Co-chairs: Helge.Meinhard at CERN.ch, tony at BNL.gov

Autumn/Fall 2017 Workshop

- Held 16 – 20 October at KEK, Tsukuba, Japan
- KEK: Home to BELLE II, national lab, many other diverse research activities
- Co-hosted with LHCOPN/LHCONE (see next presentation) and HPSS User Forum
 - Common session of all three Mon morning, common sessions with LHCOPN/LHCONE Mon and Tue afternoon
 - Little commonality with HUF, but common interests with LHCOPN/LHCONE – repeat?
- Sponsors: Fujifilm, Lenovo, Fujitsu, IBM, Intel, Spectralogic, DDN, Sony Everspan, Tsukuba
- 149 registered participants for HEPiX and LHCOPN/LHCONE – very good attendance
 - Many first timers, many old friends
 - 41 from Asia/Pacific (including 19 from KEK), 72 from Europe, 21 from North America, 15 from companies
- 56 (!) different affiliations, including 9 from Asia
- 83 (!) contributions (see later summaries)
- Full details: <https://indico.cern.ch/e/hepixon-fall2017>
- Extensive workshop notes in preparation (yes, really – kindly contributed by Jérôme Belleman / CERN)

Attendees



Tracks and Trends... (2)

- Storage and file systems: 3 sessions, 7 contributions
 - *World's largest erasure-coded CEPH deployment in production and effectively serving multiple VO's through purpose-developed xrootd and gridftp gateways*
 - *Scaling up CEPH support for HPC; expanding CERNbox service roles within and outside of HEP*
 - *DynaFed deployments on the horizon*
 - *Strong community interest in (and concern about) archival storage roadmap; launch working group involving both facilities and experiments*
- Clouds, virtualisation, grids: 2 sessions, 4 contributions
 - *Containers (both Singularity and Docker), application to batch systems and hybrid clouds*
 - *Dynamic federation of data stores; distributed clouds successfully running substantial work for Belle II and Atlas with cloud scheduler*
 - *Large number of references to cloud and virtualisation in other tracks – few dedicated talks, but these tools are now embedded in much of our work (“commodity”)*

Tracks and Trends... (3)

- Computing and batch systems: 3 sessions, 8 contributions
 - *Fast benchmark: good progress, DB12 most promising candidate, work to understand the last scalability problems observed (mainly ATLAS reconstruction)*
 - *Work in progress to better understand problems of HS06 scalability with applications: number of configured slots per core shown to have a strong influence*
 - *Next generation procurement/pledge benchmark: SPEC2017 out, work in progress to identify whether it can be used to replace HS06*
 - *Benchmarking WG welcomes new members – talk to Manfred Alef*
 - *HTCondor used at ever more sites (community being built), GridEngine almost dropped.*
 - *TechLab: a quite unique facility to test and evaluate new approaches for efficient use of CPUs, open to the community.*
 - *HPC-HTC integration at several sites: room for sharing ideas and solutions, e.g. on HTCondor/SLURM integration*
 - *Understanding “costs”: starting work in WLCG to enable an informed global optimisation of our technical choices*
- IT facilities and business continuity: 1 session, 1 contribution
 - *Collecting large amounts of facility data allows for targeted display of critical metrics and for cost-saving optimisations*
 - *Air quality measurements important to monitor particulates in view of tape*

Tracks and Trends... (4)

- Basic IT services: 2 sessions, 9 contributions
 - *Stream processing and time-series databases in focus, riemann and InfluxDB hot*
 - *BMC/IPMI security concerns and work on getting suppliers to fix them*
 - *High-performance distributed puppet setup*
 - *Thermal reaction system shutting down servers based on cron and webpages*
- End-user IT services and operating systems: 2 sessions, 5 contributions
 - *Document conversion service: Handling large variety of formats, offering OCR*
 - *User portal: single entry for user-specific monitoring, documentation, user alerts, news*
 - *CI for Linux images: automated thorough testing of images for Openstack with gitlab*
 - *Printing at GSI: Moving all printers to dedicated VLAN connected only via dual-homed hosts*
 - *SELinux found useful, powerful and reasonably easy to configure*

Tracks and Trends... (5)

- Site reports – 4 sessions, 20 contributions
 - *Network upgrades to increase the bandwidth and cope with the upcoming needs (LHC; BelleII, AMS, LIGO, multiple experiments in China); migrating main line(s) to 100Gbps (40 Gbps mentioned as well a few times)*
 - *Volunteer computing as additional source of CPU capacity*
 - *Multiple sites looking into SDN technology*
 - *HPC clusters, partly with GPUs, at many sites; job flocking from HTC to HPC*
 - *IPv6 moving slowly, but with ongoing progress*
 - *Concerns about significant price increases of commonly used software*

Back to the Roots

- **CHEP (Computing in High-Energy Physics) conference in Tsukuba (Japan), March 1991: Presentation by Les Cottrell (SLAC)**

- Independent activities are in progress at many HEP sites to understand and define how to support Unix.
- At this meeting, people are talking in corridors, sessions etc., and recognizing common challenges & solutions in supporting Unix in the HEP community.
- We are learning much from one another & feel there is much to gain from further sharing of experiences and from coordinating or aligning our Unix activities.
- Most of those of us who have discussed this at this meeting feel that HEP is at an appropriate place in the Unix learning curve that it would be extremely valuable to
 - **SHARE UNIX EXPERIENCE** between HEP sites active in Unix
 - **ENCOURAGE ALIGNMENT** in the Unix user environment in the HEP community
- There is a critical need for mechanisms to be set up to provide information exchange
 - Within HEP via **REAL FACE-TO-FACE MEETINGS**, Email, netnews, info-servers, remote conferencing documentation, newsletters etc.
 - With vendors - make sure they know about HEP needs
 - With standards bodies - encourage HEP representation on critical bodies
 - ...

PROPOSAL: Set up a HEP Unix Coordination Group:

- **Brief:**
 - **Provide information exchange**
 - **Promote alignment of the Unix environment in the HEP community**
- **Recommend Regular Meetings:**
 - **Half Yearly**
 - **Aligned with CHEP & other meetings of relevance**
 - **1-2 people from HEP sites working on providing a Unix environment to users**
- **Name of Group: HEPIX**
- **First Meeting: FNAL - 1st week October 1991**
- **Email: HEPNET::HEPIX_L (after April); HEPIX_L@HEPNET.FNAL.GOV**
- **More Information:**

| | |
|-----------------|--|
| Les Cottrell | COTTRELL@SLACVM.SLAC.STANFORD.EDU |
| Judith Nicholls | NICHOLLS@FNAL.GOV |
| Harry Renshall | HRRCR@CERNVM.CERN.CH |
| Alan Silverman | ALAN@VXCERN.CERN.CH |

Next Meetings

- Spring 2018: University of Wisconsin in Madison (WI), USA
 - 14 – 18 May 2018 (followed by HTCondor week starting 21 May 2018)
 - Same tracks as this time
 - First public announcement early in 2018
- HTCondor workshop: RAL, Didcot, United Kingdom
 - 04 – 07 September 2018
- Autumn 2018: PIC, Barcelona (Spain – or will it be another country?)
 - 08 – 12 October 2018
- 2019: tentatively at SDSC, San Diego (CA), USA and Nikhef, Amsterdam, The Netherlands
 - No dates defined yet
- Comfortable number of proposals received from Europe, North America and Asia/Pacific for subsequent meetings

Expressions of interest and proposals are nonetheless always welcome



HEPiX

[More Information](#)

See you at the next meeting in
University of Wisconsin-Madison

Madison

May 14-18, 2018

Questions, Comments?