



# HEPiX Autumn 2018 Workshop Report

Helge Meinhard / CERN-IT  
WLCG Grid Deployment Board, 17 Oct 2018

Contributions by  
Manfred Aef, Sebastien Gadrat, Andreas Haupt, Michel Jouvin, David  
Kelsey, Shawn McKee, Michele Michelotto, Sandy Philpott, Ofer Rind,  
Wayne Salter, Peter van der Reest, Mattias Wadenstein, Tony Wong

# HEPiX Mission

- From our Web site <https://www.hepixon.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 2018 in Barcelona was the 58th workshop since 1991
- 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
  - Tracking of technology and market evolutions (tech watch)
  - SIG on AAI
- 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 28 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](mailto:hepixon-users@hepixon.org) (moderated)
  - To subscribe, send an e-mail to [listserv@listserv.hepixon.org](mailto: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 2018 Workshop

- Held 08 – 12 October in downtown Barcelona, hosted by Port d'Informació Científica (PIC)
- PIC: Tier-1 for WLCG (ATLAS, CMS, LHCb), support for particle physics, astrophysics, cosmology, earth sciences
- Sponsors: Fujifilm, Western Digital (!), Omega/Dell, Spectralogic, IFAE, CIEMAT, Nemix/Supermicro, Accelestor
- 137 registered attendees – record for HEPiX!
  - Many first-timers, many old friends
  - 105 from Europe (including 12 from PIC), 14 from North America, 9 from Asia, 9 from companies
  - 54 different affiliations: 37 from Europe, 7 from North America, 5 from Asia, 5 companies
- 69 contributions, not counting any extra session (see later summaries)
- Full details: <https://indico.cern.ch/e/hepix-autumn2018>

# Attendees



# Tracks and Trends (1)

- Miscellaneous: 4 contributions, 1h10'
  - CSBS is OUR journal. Let's use it and publish our work! Many interesting things this week – consider writing an article
- Basic IT services: 4 contributions, 1h40'
  - Using messaging services to improve scalability and reliability of systems management
  - New approach in the CERN Authentication and Authorisation management
- Grids, clouds, virtualisation: 3 contributions, 1h10'
  - Data lake proposal: To optimise costs, will need to switch to QoS (performance, reliability etc.) for storage – major change (technical, cultural, social)

# Tracks and Trends (2)

- Computing and batch services: 9 contributions, 3h35'
  - Benchmarks : 2 main activities
    - Fast benchmark for estimating the job slot CPU power : LHCb DB12 adopted by LHCb and ALICE but not appropriate for the procurements
    - Next generation benchmark for estimating the installed capacity and for procurements : work still in progress: SPECcpu 2017, set of HEP apps?
  - AMD EPYC architecture: promising evaluation by BNL, will restore some competition on the server CPU market
    - Performance and price competitive with Intel; need to see adoption by server vendors
    - Next generation next year will bring many more cores
  - Photon science support: several HEP sites involved with big projects for the coming decade
    - A 1st workshop at BNL discussing the specific issues, in particular those caused by the loose links between users and the photon facility
    - Idea of organising a workshop focused on computing for photon science co-located with HEPiX, like we did for LHCOPN/LHCONE
  - Improving OpenMP scaling using openssl:: Using set of "openssl speed" commands to optimize OpenMP performance
  - Commissioning CERN Tier-0 reconstruction workloads on Piz Daint at CSCS: Configuring and optimizing Piz Daint at CSCS for running ATLAS and CMS Tier-0 workloads
  - PDSF - Current Status and Migration to Cori
  - Spare compute resources of EOS storage nodes at CERN have been enabled to run user jobs in containers

# Tracks and Trends (3)

- End-user services and operating systems: 7 contributions, 2h55'
  - CERN service management: Focus on Service Catalogue, bringing services on board, the User experience, and Tool configuration
  - CERN Linux services: Update on CC7, SLC6, RHEL support distributions and services; software collections, virtualization, openstack SIGs; anaconda plugin, lockup; community work on alternative architectures; Koji and Gitlab; future support for lightweight containers, CC8, Freeipa, s3 for static content
  - Indico: Upgrade to 2.1! Roadmap for next releases: new room booking, internationalisation, paper reviewing, CalDAV support, ...
  - Reducing dependencies on commercial software (providers): investigating open-source products to provide services and avoid vendor lock-in
  - Container orchestration: Enable hosting large range of Web applications within CERN; opportunity to consolidate all web hosting on common infrastructure
  - Jupyter-based analysis portal at BNL: Supports containerized applications started via a seamless integration into local batch with authentication token delegation
  - Rust: Worth a serious look by Python developers, good compromise between C++ and Python

# Tracks and Trends (4)

- IT facilities: 4 contributions, 1h40'
  - Technology watch working group kicked off with 58 people subscribed. Number of subgroups covering individual domains, e.g. processors, memory, etc. Further volunteers to do the actual work are needed
  - Cost model WG report: better understanding of the workloads via defined metrics; define a common framework for estimating resources, to then look at scenarios to make improvements
  - Latest CERN procurements, changes in the team, impact of recent issues and technology changes
  - Superfacility at NERSC: introduce common workflow and API for access from multiple sciences and user to the facility (rather than dedicated access per science/user)



# Tracks and Trends (6)

- Site reports: 16 contributions, 4h00'
  - HPC sites / clusters more and more used by HEP, and discussed in HEPiX
  - Computing farms extended to clouds or HPC resources (unified pool).
    - HTCCondor + ArcCE a common choice
  - More improvements / upgrades on network, enabling IPv6 everywhere now
  - Sites need to cope with new and important resources needs of their users
  - SurfSara, the Dutch supercomputer center for the first time at HEPiX. Very active on GRID activities for non WLCG communities. Interesting among the several activities: WebDav security, and RcAuth, proxies without certificates
  - KiSTI notable Data Center Relocation, KISTI Grid CA system based on Hardware Security Module
  - LAL+GRIF: Expansion to new data center slower then expected, mainly bc of administrative problems
  - NERSC: move from PDSF to SLURM, move to CORI, CVMFS on the Cray

# Tracks and Trends (7)

- Storage and file systems: 11 contributions, 4h25', plus two BoF sessions
  - Tape Performance and accompanying infrastructure: centres have been busy w/ ATLAS Data Carousel; CERN field-testing CTA
  - Seven specialist presentations in Bird-of-a-Feather session leaving little time for discussion, but (hopefully) continued theme at HEPiX
  - Impact of vendor release changes on storage systems and experiments timelines: example of Oracle Databases at CERN
  - Large deployment of CEPH storage cluster (Osiris) around the Great Lakes: multi-site data placement challenging, caching very rewarding
  - Tools for data management being renewed and adapted (also touched upon in BoF)
    - Interesting R&D on tape drive characterisation (as essential information needed is not provided by vendors) to improve tape operations efficiency
  - Optimizing EOS based storage: splitting up storage metadata servers in multiple instances to remedy performance and operational issues
  - Introducing CEPHfs + Manila-based file storage as replacement of classic NFS file services
  - Overview of large backup infrastructures: 'new' themes both in infrastructure and in user requirements
  - Three industry talks focussing on tape media and tape drive tech developments, and on optimising flash storage

# Extra Sessions

- In addition to official workshop programme:
  - Ad-hoc BoF session on AAI at sites
    - Direct follow-up of workshop presentation and discussions during breaks
  - DOMA ACCESS meeting
  - BoF session on archival storage / tapes

# Ad-hoc BoF on AAI at Sites

- Well attended, 25...30 people in the room (if not more)
- Multiple sites planning to review (and partially redo) their stack
- Interest in digging into the issue further
- Paolo Tedesco (CERN) and Dave Kelsey (STFC-RAL) volunteered to take things in hand
- Interested? Subscribe to [hepux-aai@hepux.org](mailto:hepux-aai@hepux.org):
  - Send an e-mail to [listserv@listserv.hepux.org](mailto:listserv@listserv.hepux.org) with an empty subject and the following line in the e-mail body:  
`subscribe hepux-aai <your first name> <your family name>`

# Next Meetings

- Spring 2019: San Diego Supercomputer Center / University of California in San Diego, USA
  - 25 – 29 March (week following WLCG / HSF / OSG workshop at JLAB, week before ISGC at ASGC, Taipei, Taiwan)
  - Same tracks as this time
  - First announcement well before end of the year
- HTCondor workshop in Europe: Joint Research Centre of EC, Ispra (Lombardy), Italy
  - 24 – 27 September
- Autumn 2019: Nikhef, Amsterdam, The Netherlands
  - 14 – 18 October
- Quite firm ideas about 2020 and 2021, some expressions of interest even beyond 2021 from all regions

Expressions of interest and proposals still very welcome

See you at the next meeting in  
San Diego Supercomputer Center

San Diego, USA

March 25 - 29, 2019

# Questions, Comments?