



# HEPiX Autumn 2019 Workshop Report

Helge Meinhard / CERN-IT  
WLCG Grid Deployment Board  
20 November 2019

Contributions by session chairs: Andrea Chierici, Sebastien Gadrat, Martin Gasthuber, Pete Gronbech, Andreas Haupt, Chris Hollowell, Shawn McKee, Michele Michelotto, Tomoaki Nakamura, Georg Rath, Wayne Salter, Jingyan Shi, Mattias Wadenstein

# 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, nuclear fusion, ...
  - Just attend the workshops (see below)

# HEPiX Activities: Workshops

- Twice per year, one week each, in Europe, North America and Asia
  - Typically 80...140 participants
  - Autumn 2019 in Amsterdam was the 60th workshop since 1991
- Workshop style
  - All plenary, no parallel sessions
  - Abstracts are usually accepted - no formal Programme Committee reviewing
  - No proceedings
  - 60...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 27 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 2019 Workshop

- Held 14 – 18 October at Nikhef, Amsterdam, The Netherlands
- Nikhef: Dutch national institute for subatomic physics; major Tier-1 centre for WLCG
- Sponsors: Fujifilm; Dell / AMD; Western Digital; Fujitsu, Mellanox, Promise, Tyan
- 113 registered attendees
  - Many first-timers, many old friends
  - Academia: 81 from Europe, 13 from North America, 8 from Asia; 11 from companies
- 67 contributions (see later summaries), all recorded
  - All slides and videos available from...
- Full details: <https://indico.cern.ch/e/hepix-autumn2019>
- Followed by ARC-CE workshop (18 Oct pm) and SOC workshop (21-23 Oct) at Nikhef

# Attendees



# Tracks and Trends (1)

- Basic IT services: 6 contributions, 2h30'
  - Several talks on logins, credential managements, AAI
  - Monitoring: Prometheus+ELK at DESY; Infiniband networks
- Grids, clouds, virtualisation: 5 contributions, 2h05'
  - Openstack at CERN: Nice success story – for CERN and Openstack
  - Distributed computing for LHAASO using cloud technology successful
  - Data lake R&D in Russia
  - Life Sciences (JGI) has very similar requirements to HEP on distributed computing – started successfully, extension to public clouds planned
  - SLATE: Service delivery via container orchestration, a step towards no-ops operations
    - Challenges on trust and security



# Tracks and Trends (2)

- Computing and batch services: 8 contributions, 3h20'
  - HTCondor: Workshop report; coming features; bulk deployment in Bonn (Singularity) and at DESY; performance of file transfer mechanism
    - No other batch systems mentioned at all
  - Benchmarking: Impressive progress on framework as well as workload from LHC (SPEC-CPU 2017 not suitable). Other communities? How to move to HS06 successor?

# Tracks and Trends (3)

- End-user services and operating systems: 6 contributions, 2h30'
  - CERN MALT: Telephony, mail, home directories moving away from commercial solutions
  - W7 to W10: upgrade rather than re-install

# Tracks and Trends (4)

- IT facilities: 2 contributions, 0h50' (***we really can do with more!***)
  - CERN repatriates equipment from Wigner in containers at LHCb (for Run 3); moving ahead with new data centre (early days, budget cuts)
  - Open Compute Project: Industry support, more variety, more market penetration
    - Not quite competitive yet, but getting closer



# Tracks and Trends (6)

- Site reports: 17 contributions, 4h15'
  - HTCondor everywhere
  - JupyterHub spreading, spawner scripts allow to submit jobs on the local batch system, with sometimes switches to select the right resources to submit the job to (CPU or GPU for instance).
  - Kubernetes trend re-enforced (used for managing JupyterHub instances or containers)
  - Licensing problems everywhere: Many sites moving to Open Source or other less traditional software
  - Increasing usage of GPUs
    - Sometimes only few nodes to start a transition, sometimes with hundreds of GPU board
  - Increasing use of intermediate Ethernet speed (25 - 50 Gbps when 10 Gbps non sufficient)
  - Many sites study how to exploit HPC machines and/or commercial clouds
    - Peaks of demand; special grants; opportunistic usage
  - Presentations by photon science sites
    - Willingness to collaborate inter alia and with HEP despite cultural differences of community with respect to HEP

# Tracks and Trends (7)

- Storage and file systems: 9 contributions, 3h55'
  - Tape: Moving out of Oracle drives and libraries; moving out of Castor; LTO increasingly deployed
    - Data Carousel proposal: new challenges for archive/cold storage systems in terms of guaranteed 'bring online' latency
  - CERN: MALT-related changes/challenges (e.g. SAMBA on EOS); CERNBox as focal point provided sync&share plus app integration (EU project); DboD: growing usage, HA features added to PostgreSQL and MySQL
  - Ceph and CephFS becoming commodity; interesting backup options
  - Storage usage statistics (2 presentations)

# Next Meetings

- Spring 2020: Academia Sinica Grid Computing (ASGC), Taipei, Taiwan
  - 02 – 06 March
  - Same tracks as this time
  - ISGC is the week after (at ASGC), followed by the OSG AHM (at Oklahoma University)
  - First announcement is out – mark your calendars, register once open!
- HTCondor workshop in Europe: LAL, Orsay (near Paris), France
  - 22 – 25 September
- Autumn 2020: IPHC, Strasbourg, France
  - (very probably) 12 – 16 October
- Quite firm ideas about 2021 and 2022, some expressions of interest even beyond 2022 from all regions

Expressions of interest and proposals still very welcome



See you at the next meeting in  
Academia Sinica – ASGC  
Taipei, Taiwan  
March 2 – 6, 2020

# Questions, Comments?