HEPiX Fall 2024 Workshop Recap

Ofer Rind
SDCC-NPPS Joint Meeting
14 Nov, 2024

With contributions from:
David Britton, Dino Conciatore, David Crooks, Josep Flix, Sebastien Gadrat, Mary Hester, Shawn McKee, Andreas Petzold, Matthias Schnepf, Horst Severini, Elvin Sindrilaru

What is HEPiX?

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. Participating sites include ASGC, BNL, CERN, DESY, FNAL, IHEP, IN2P3, INFN, JLAB, KEK, KIT, Nikhef, PIC, RAL, SLAC, TRIUMF and many others. The HEPiX organization was formed in 1991, and its semi-annual meetings are an excellent source of information and sharing for IT experts in scientific computing. https://www.hepix.org/

Attendees and Program

- 25 Institutions
  - 13 from North America, 10 from Europe, 2 from Asia
- 44 registered participants
  - Many were first time attendees
  - 23 from North America, 18 from Europe, 3 from Asia
    - Several registrants were unable to participate in the end due to travel restrictions
    - Timing played a role in limiting attendance - sandwiched between CHEP and SuperComputing, registration opened with insufficient lead time, etc.
- 42 submitted abstracts for 48 presentations
  - The lighter numbers allowed us to give speakers a little bit more time than usual (20 min for Site Reports, 30 min for other talks)
  - Two special topical sessions on Security Operations Centers and Sustainability

Timetable


HEPiX
HEPiX Fall 2024
University of Oklahoma, USA
4-8 November 2024
The UNIVERSITY of OKLAHOMA
Photo credit: Alexandr Mikula

Topical Session - Security Operations Center

- Most of Wednesday was devoted to improving cybersecurity at our sites through sharing of information (threat intelligence)
  - Morning session introduced the tools and the drivers - ZEEK, MISP, pDNSSOC
  - Afternoon devoted to demos and detailed discussion

SOC Model

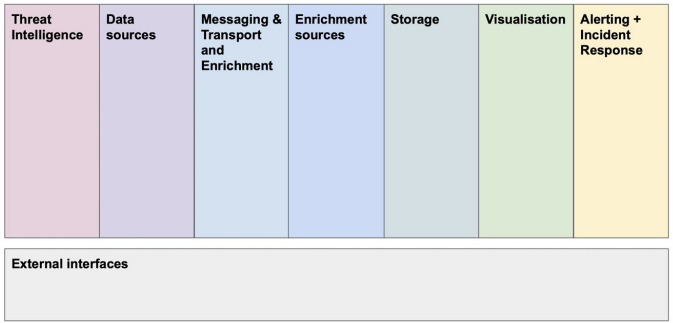

Collaborative Operational Security: The future of Cybersecurity for Research and Education

Topical Session - Carbon and Sustainability

- Environmental Sustainability continues to be a hot topic in the community and beyond
- 1.5 hrs devoted to this topic with three presentations on recent, increasingly in-depth analysis of carbon (not just power and not just CPU!) consumption
- Demos of two new tools: a utility to account for carbon in procurement, and one to optimize power reduction on compute farms
- Discussion to continue in December at the WLCG Sustainability Workshop
  [call for abstracts closes soon: 11th Nov.]

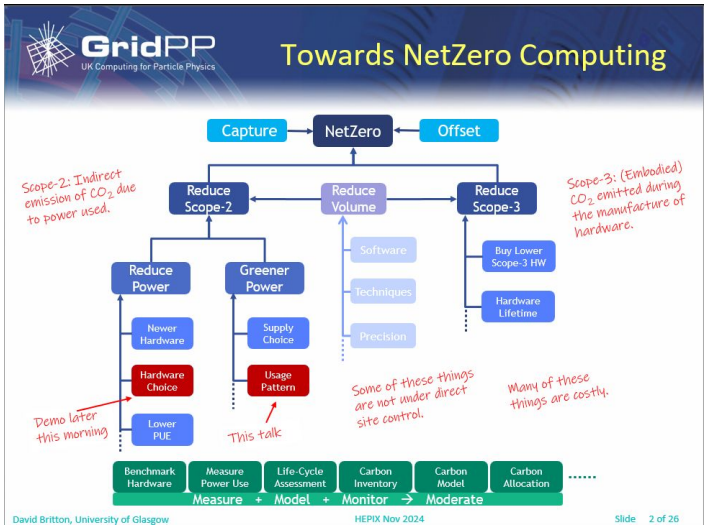

Tracks and Trends - Site Reports I

- **9 Site Reports** (4 from US, 3 from EU and 2 from Asia)
  - SWT2, AGLT2, S3DF (SLAC), JLab, PIC, RAL, CERN, KEK and IHEP.
- **Computing**
  - batch : Slurm & HTCondor
  - Still mainly CPU, but some ARM and more GPU in production
  - GPU mainly for AI oriented projects (on notebooks based service)
  - HTC and Google computing resources also integrated and used.
- **Storage**
  - Lustre, EOS, XRootD, dCache (as usual)
- **Operating Systems**
  - Migration from CentOS7 to RHEL8/9 and related flavours (Alma or Rocky) on-going, almost completed at most of the sites.
- **Config Management and Deployment**
  - **challenging** in an environment where **bare metal and cloud native infrastructures** need to be managed. Many sites (not all) converging on Foreman/RH Satellite for bare metal and VMs, together with Puppet + Ansible or Chef.

Tracks and Trends - Site Reports II

- **Tape**
  - Mostly Spectra Logic TFinity with LTO8 and LTO9 cartridges
- **(Still Tape)** Humidity sensitivity of LTO8 & LTO9 cartridges (PIC Site report)
  - positioning errors and drives degradation
  - New room required, will be constructed in Q1 2025

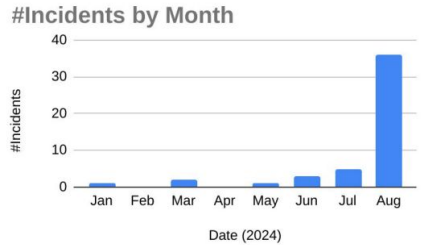

- Services more and more **containerized** (Docker), k8s / Openshift based infrastructure.
- Sites often provide "science platform" based on notebooks (**Jupyter**) and related tools.
  - Used for end-user analysis, especially for IA related projects.

Tracks and Trends - Computing and Batch Services

- Youngest track by speakers
- 3 contributions
  - High throughput setup for HL-LHC analysis facility and fast turnaround
    - XCache setup provides 200Gbps
    - interactive analysis environment via jupyter and Dask
  - Accounting Tool AUDITOR
    - EGI and in-house accounting
    - subsite accounting
  - Benchmark WG
    - improvements and new features for workloads and benchmark tools
    - various analyses within the WG
    - GPU benchmark and score are currently on hold

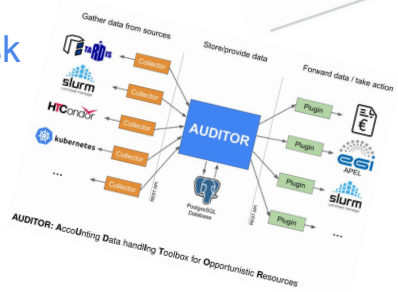

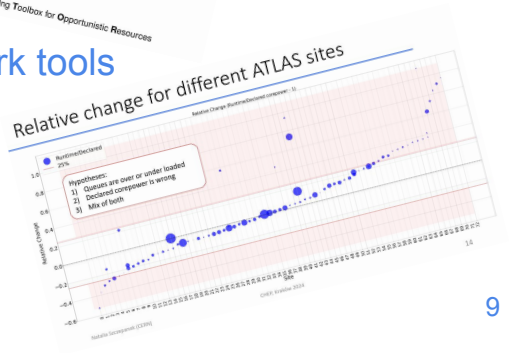

Tracks and Trends - OS, Clouds, Virtualization, Grids I

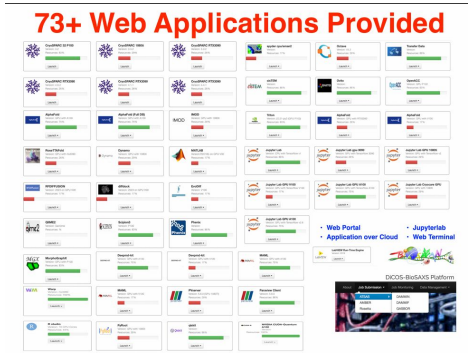

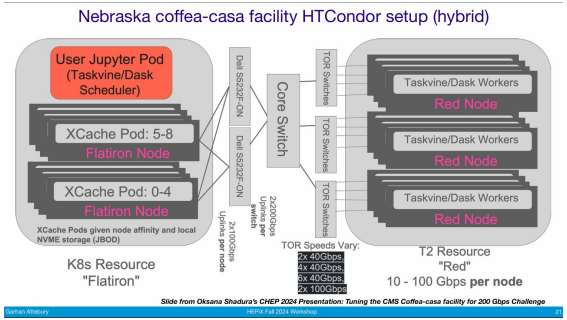

Science Cloud at Academia Sinica (Taiwan): Taiwan's cloud infrastructure leverages WLCG tech, with PanDA and RUCIO integration, to support scientific communities. The focus is on data efficiency, bandwidth, and security, with ASGC expanding capacity. Upcoming updates include 2FA in 2024 and CTA migration in 2025.

Cloud-Native Control Plane at CSCS: A Control Plane for infrastructure-as-code, enabling reusable composite resources and single-manifest deployment. Demo showcased Slurm cluster creation on VMs with reusable code.

Coffea-Casa Analysis Facility: Provides high-throughput, interactive data analysis with ML tools and reproducibility. Currently handling up to 1000 TB and exploring HL-LHC data speeds (200 Gbps Challenge). Future plans involve Tier2 integration, improved VPN, and user support enhancements.

Tracks and Trends - OS, Clouds, Virtualization, Grids II

dCache on Kubernetes at CSCS: Custom dCache on Kubernetes with Helm for deployment, HA PostgreSQL, and ArgoCD for GitOps-driven continuous delivery. Ongoing improvements aim to boost robustness and monitoring.

Xrootd & FTS Workshop: CERN FTS team reports 1 billion transfers and introduces FTS v3.13. ATLAS and CMS scalability issues highlight token coordination needs. Xrootd v6.0 adds RDMA support and error handling. The Pelican platform advances a CDN for scientific data sharing.

Purdue Analysis Facility (AF) provides an interactive and comprehensive toolset deployed on a Kubernetes cluster at Purdue University. Connected to Purdue's CMS Tier-2 storage and utilizes Slurm for job scheduling.

Transitioning from RHEV to Openshift: Transitioning from RHEV to OpenShift includes configuring NetApp storage with CSI drivers, setting up a bare-metal cluster with custom network configurations, using MTV and NMstate operators for VM migration and networking, and planning to complete migration by February 2025.

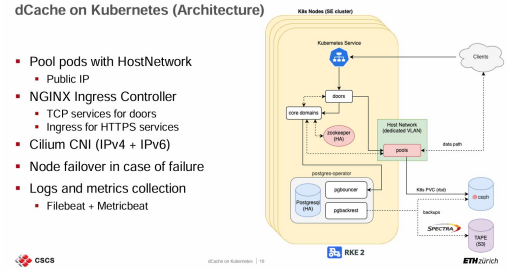

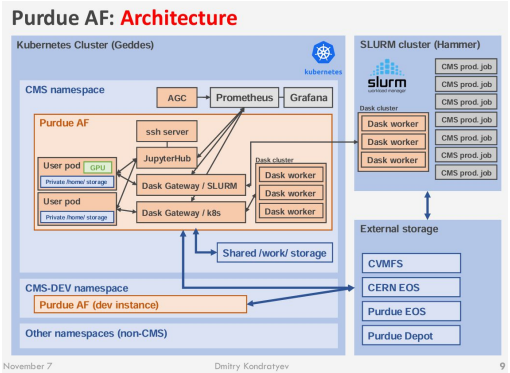

Tracks and Trends - Storage and Filesystems

- Two main areas of storage and file systems were covered in 4 talks:
  - **Tape technologies**
    - Nice roadmap about the future of tape from Spectra Logic
    - Focus on new media technology and ways of using tape libraries for increasing data reliability
    - Lessons learned and experience of moving from TSM to HPSS
  - **Data placement and access analysis using Big Data and AI**
    - Focusing on real data analysis from two of the main LHC experiments
    - Showcasing valuable insights and trends that can be extracted from logs
    - Uncovering emerging patterns in both data management and application frameworks that are otherwise impossible to detect

Tracks and Trends - Networking and Security

- We had 7 talks (of 8 scheduled), 1 on Wednesday and 6 on Thursday with 4 on security and 3 on networking
- Security
  - Firewall under attack: operational security rollercoaster (IPv6 to the rescue :))
    - Patrick Storm and Romain Wartel
  - One year into the CERN Cyber-Security Audit, Stefan Lueders
  - Computer Security Update, Stefan Lueders
  - HELP! I have DataCenter Nightmares Stefan Lueders
- Networking
  - Networking Topics for WLCG, Shawn McKee
  - Network tests at CZ Tier-2, Jiri Chudoba
  - Getting closer to an IPv6-only WLCG – update from the HEPiX IPv6 Working Group
    - Martin Bly and Dave Kelsey

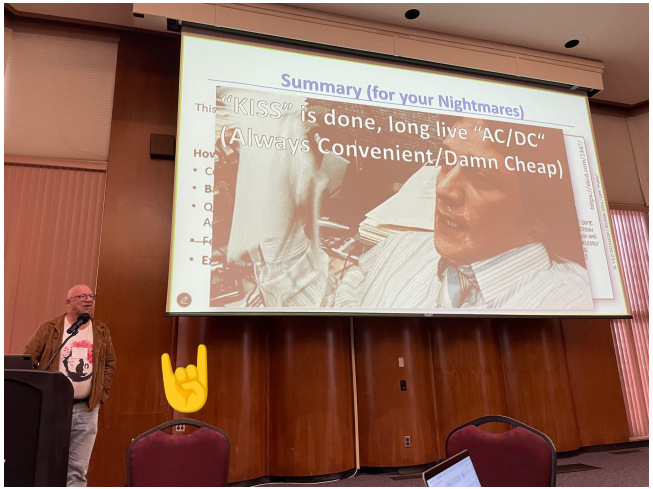

Tracks and Trends - Basic and End-User IT Services

3 talks submitted for the Basic & End-User IT services track

1. **CI4FPGA: Continuous Integration for FPGA/SoC Projects:** CERN's Engineering Software Services team centralized infrastructure to use FPGAs and Gitlab CI/CD with pre-configured Docker images.
2. **CRISP: Collaborative Tools for the ePIC Experiment**: BNL reported on ePIC experiment which modularly manages collaboration information. One point of interest was the incorporation of the InvenioRDM software.
3. **pkcli: A Framework for Scripts to Manage Applications: (***Unable to present due to travel issues)* RadiaSoft developed an open source Python framework, pykern.pkcli, that simplifies the creation of these application-specific CLIs. They provided an overview of our framework and share examples of how we've used it to administer our systems [in their slides.]

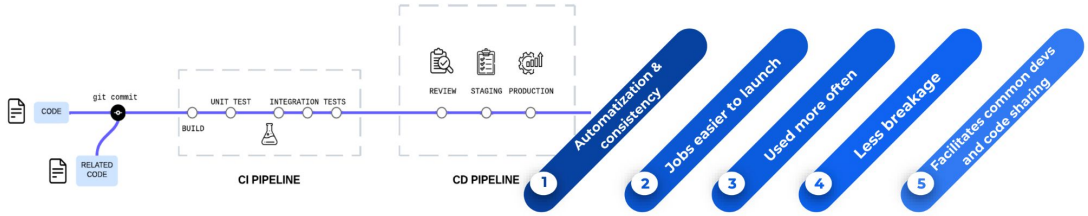

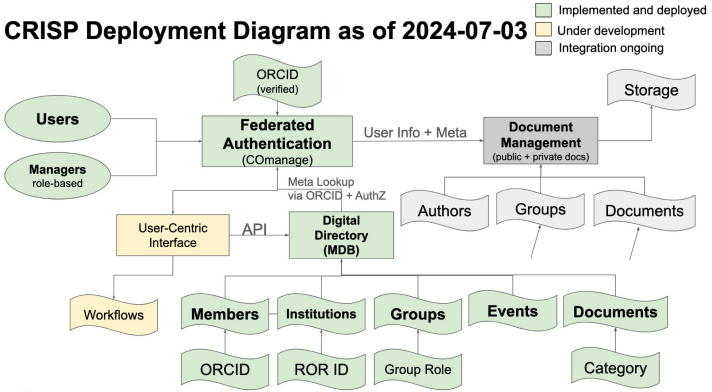

Tracks and Trends - Miscellaneous

- Henry Neeman started us off with an excellent view into the local OSCER Supercomputing facility
- Andrew Fagg from OU gave a really nice overview of the real-life application of using AI/ML techniques for Atmospheric Visibility Estimation
- Shigeki Misawa, speaking for the HEPiX Techwatch Working Group, provided insight into what advances in tech hardware really mean for existing practices.
  - Strategic concerns raised regarding performance and cost for storage. Looking to have a dedicated OTF session on this.

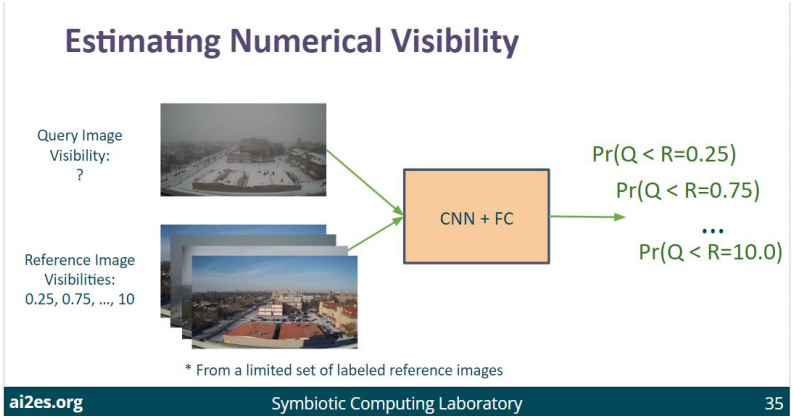

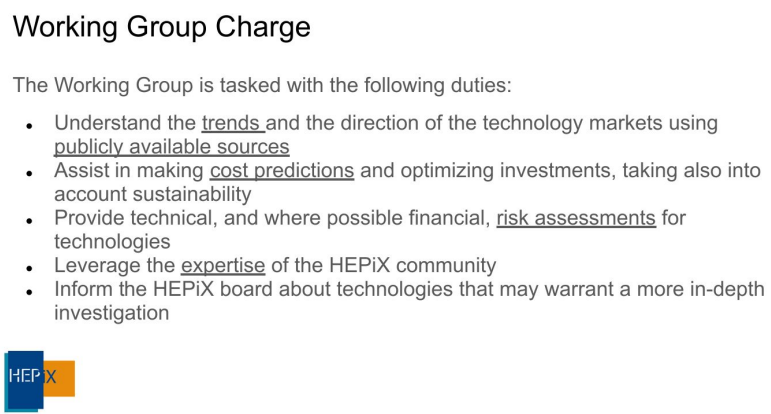

Tracks and Trends - Show Us Your Toolbox

- 2 really nice demos (TMUX, K9s) and 1 really engaging presentation from Dino, Elia and Garhan
  - Always one of the most lively and spontaneous sessions
  - Thank you to Mary Hester for stepping in to help organize it!

We do hope you get your dream datacenter some day, Garhan!

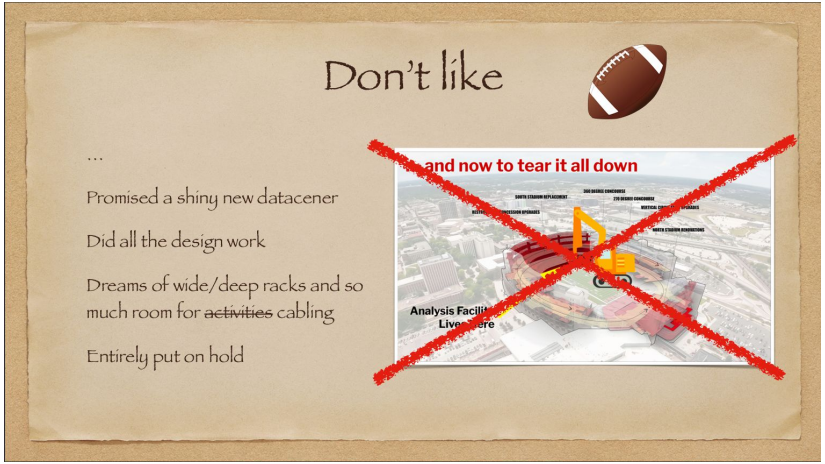

Board Meeting

- Lengthy review of the current meeting
  - Board was pleased with the way the meeting was running, the success of the SOC session, and the content in general
  - To review current tracks, as usual
  - Earlier announcements (including save-the-date) and clearer deadlines for future meetings
  - To compile a “guide” for hosting HEPiX (including indico tips!) and develop a “code of conduct”
- Report on progress for the Spring meeting, to be held 31 March – 04 April, 2025 at CSCS Lugano, Switzerland
  - Arrangements are in great shape with announcement forthcoming shortly
- Presentation of Lanzhou University proposal to host HEPiX, potentially in Fall 2025
  - Many details still to be understood, but the board maintains a strong commitment to holding workshops in the Asia region
- Reports from all three working groups
- **Thoughts on HEPiX tracks, content, organization? Fresh eyes are welcome...please contact us!**

Next year in Lugano

31 Mar - 4 Apr, 2025