

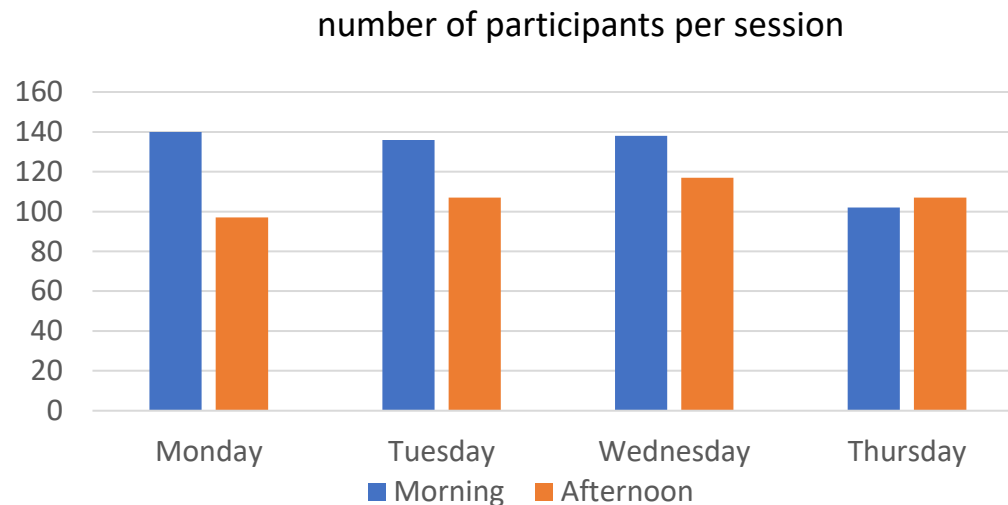
HEPIX Spring 2022 Wrap-up

Tony Wong

April 28, 2022

Statistics

- Over 340 registered participants
 - Participation varied by time and track
- 42 presentations covering all standard HEPIX tracks
 - 9 in Site Reports
 - 2 in End-User IT & Operating Systems
 - 5 in Networking & Security
 - 4 in Computing & Batch Services
 - 13 in Storage & Filesystems
 - 2 in IT Facilities & Business Continuity
 - 4 in Grid, Cloud & Virtualization
 - 3 in Basic IT Services
- Following slides present a brief overview of various presentations—full details can be found on <https://indico.cern.ch/e/hepix-spring2022>
- Entire workshop was recorded—recordings will be added to Indico agenda soon and a corresponding announcement will be made
- Shared HedgeDoc document with comments, questions, feedback, etc covering all presentations is available here: https://notes.desy.de/ou_F-tY6SDKkaE8kUoi6gA



Site Reports

- PIC—early results of the Tape Challenge'22 and leveraging Barcelona Supercomputing Center (BCS) resources for ATLAS and CMS
- CERN—news on 2FA for new SSO, conclusions of the MALT project, equivalency among RHEL, CentOS Stream, ELC (Rocky, Alma, etc)
- ASGC—exponential growth in gpu resource utilization, consolidation of TW HEP resources at ASGC, ML-enabled anomaly detection and system efficiency optimization
- KEKCC—current generation of resources to run (at least) until August 2024, Belle-II is dominant user, migration from SINET5 to SINET6
- IHEP—usage of ARM resources, anomaly detection in Lustre I/O behavior, R&D progress for HEPS
- RAL—plans for a new data center, focus on cyber-security concerns, new erasure-coding Ceph cluster
- FNAL—support for DUNE, HEPCloud, migration from Enstore to CTA
- BNL—migration to new data center, on-going tape challenge, R&D on future disk storage
- Diamond—usage of STFC cloud resources, MFA implementation, RHEL upgrade

Networking & Security (1)

- Zookeeper Policy Audit Tool Project @ CERN
 - (Meta)data security concerns—restrictions to prevent unauthorized access
 - Goal is to create set of tools for auditing, monitoring and enforcement of security policies
- Security Update (CERN)
 - Geopolitically-inspired phishing and electronic fraud activities
 - Supply chain security procedures (monitor security advisories, use trusted sources, enhance preventive measures, etc)
- Collaborative Threat Intelligence and Security Operation Centers
 - More coordinated, targeted attacks of academic research resources
 - Formation of a WLCG Working Group (WG) and description of recent activities at Nikhef and STFC SOC's
 - Streamline deployment and minimize effort for further SOC deployments at other WLCG sites

Networking & Security (2)

- Report by IPv6 WG
 - Continued progress on dual-stacking, an intermediate step towards migration to IPv6-only at a future date
 - U.S. gov't policy encouraging 80% of services to be IPv6-only in a few years
- Report by Research Network Technical WG
 - Background context of WLCG network requirements for Run3 and beyond
 - Focus on network visibility, usage optimization and management

Storage & Filesystems (1)

- Antares @ RAL
 - Migration to CTA (1M directories, 57M files, 70+ PB)
 - Participation in LHC tape challenges (Oct. 2021, March 2022)
- Ceph @ RAL
 - Cephadm—automate deployment and management of clusters
 - Demonstrated usage on new cluster (Arided) as a blueprint for future production deployment
- EOS Report
 - Review of recent workshops and activities
 - On-going migration to EOS5
- IO Shaping in EOS
 - Tools to manage data and meta-data rates to preserve service guarantees to core activities
 - Very useful to manage resource sharing, specially during peak usage of storage resources
- Third-Party Copy transfer service @ JUNO
 - ~3 PB of data transfer to remote (CNAF, IN2P3, JINR) sites
 - Recent transfer data challenge to guide future developments

Storage & Filesystems (2)

- Bulkrequests
 - Tool to manage dCache file location and QoS in bulk (multiple files simultaneously), seems to scale well
 - Continued development ahead with further integration with dCache
- CERN Tape Infrastructure
 - Goals for Run3 → 10 GB/s for each LHC experiment, with ~720 PB accumulated over 4 years
 - New hardware (servers, drives, media) ahead of Run3
 - LTO-9 performance data presented
- CTA Tier-0
 - Data challenges (June 2021 and March 2022) to measure throughput to (disk) buffer and tape
 - Ready for Run3

Storage & FileSystems (3)

- dCache integration with CTA @ DESY
 - All access to data on tape goes exclusively through dCache
 - Promising results from recent (2022) pilot deployment, currently in production use for EuXFEL
- EOS and Xcache performance for data analysis
 - Study data access in EOS to identify bottlenecks for Run3 (and beyond) to determine if more resources are needed
 - Ran tests (synthetic and realistic workloads), existing hardware adequate for current activities
- Open-Source Erasure Coding (EC) Technologies
 - Introduction to EC technical details and technologies
 - Comparison of open-source EC storage systems (MinIO, Ceph, EOS, XRootD)
- Native EC-based file storage and S3 Proxy
 - Alice O2 use case of XRootD EC-based storage
 - Set-up of S3 gateway to interface with commercial cloud provider resources
- PostgreSQL Table Partitioning in dCache
 - Responsiveness concerns with large dCache PostgreSQL instances at DESY
 - Observed improvements with table partitioning schemes

End-user IT and Operating Systems

- CERN's new Web Service Portal
 - Modern user-friendly entry-point to replace old, Windows-centric portal
 - Access to centrally-provided web services
 - Full management of (user-created) webEOS and Drupal sites
 - Gradual roll-out with new features (GitLab, web analytics, etc) coming soon
- Kernel modules in CentOS Stream
 - How to get non-standard kernel modules working on CentOS Stream?
 - Presentation describes the challenges due to available kernel interfaces and dynamic nature of CentOS kernels

Computing & Batch Services

- Rebalancing HTCondor fairshare
 - Heterogeneous (different cpu generations) and unpredictable single and multi-core jobs not well-addressed by default fairshare tools (defrag and static provisioning)
 - Implemented new strategy refined job scheduling and improved resource utilization
- Report by HEP-SCORE TaskForce
 - Organize effort for new benchmark to replace obsolete HEP-SPEC06
 - Migration scenario to be presented in the Fall of 2022
- Report by Benchmarking WG
 - Suite of HEP workloads (LHC and others) as part of a new HEP-SCORE benchmark
 - Studying correlations among workloads and refine workload composition for (final?) benchmark
- FNAL's HEPCloud
 - Built on top of GlideinWMS and HTCondor
 - Decision Engine regulates job routing
 - Single point of access to various (local, clouds, HPC) resources

IT Facilities & Business Continuity

- Vendor support (Nikhef)
 - Responsiveness by hardware warranty support suppliers is worsening amid recent events (pandemic, geopolitical upheavals, etc)
 - Share experiences with support issues
- New Data Center @ BNL
 - Low PUE, water-cooled racks
 - Available for operations in late 2021, expect to vacate most of old data center by Sep. 2023
 - Significant increase in capacity (260 to 478 racks) and power (1.5 MW to 9.6 MW) with dedicated space for network, tape silo and HPC systems

Grid, Cloud and Virtualization

- CERN Cloud Infrastructure
 - Hardware upgrades
 - Enabling live migration of VM's (minimize disruptions, improve service level)
- Anomaly Detection System
 - Monitoring for CERN Cloud Infrastructure
 - Comparative study between Grafana and ML-based alarms
- Getting FTS ready for Run3 @ CERN
 - Completed migration to Python 3 and MySQL 8.0
 - Data Challenges confirmed readiness for production-level service
- Integration of JLAB resources with OSG cyber-infrastructure
 - Usage of OSG tools and resources and providing access points to experiments
 - Description of plans to migrate to token authentication/authorization

Basic IT Services

- Transcoding as a Service
 - OPenCast – open-source (web)lecture recording system
 - Integration with existing publishing services at CERN (Indico, CDS)
- Databases @ DESY
 - Large footprint of DB services for ~30 years—compliance with new requirements
 - End reliance on Oracle, provide MariaDB and PostgreSQL-based services
- Moving from Elasticsearch to OpenSearch @ CERN
 - Elasticsearch is centrally-managed with user-friendly access, but it has maintainability issues (hard to keep in sync with external plugins)
 - Migration to OpenSearch on-going to address EOL for various EL versions deployed @ CERN

Some Broad Trends (my own view, not necessarily the presenters or the institutes represented at HEPIX)

- LHC run 3 (and beyond) driving several upgrades and activities
 - Network R&D and improvements
 - Tape libraries and systems
 - Storage systems (Erasure Coding, etc)
 - Leveraging of non-dedicated (clouds, HPC, etc) resources
- Adoption of modern tools for data analysis and resource management
 - Jupyter
 - Kubernetes, OKD
 - ML-enabled algorithms for anomaly detection and resource monitoring
 - Increasing security awareness

Fall 2022 workshop

- On-going discussions about in-person vs. remote meeting
 - Favor in-person format, if conditions allow
 - Support for remote participation, with some constraints (TBD)
- Targeting Early Fall—announcement in due time
- New co-chair (to serve alongside Peter van der Reest) to be elected at the board meeting in the Fall workshop
 - Search & Election committee formed
 - David Kelsey (chair), Gang Chen, Jose Flix Molina and Ofer Rind
 - Announcement coming soon

Note to Presenters

- If you have uploaded CodiMD, Google, etc formatted files to Indico, please consider re-uploading with PDF files. This will help us preserve the presentations and maintain uniformity with all other slides on the Indico timetable.

Manfred Alef

- Manfred is retiring from KIT very soon
- On behalf of the HEPIX board and the extended community, thank you, Manfred, for your extensive contributions to HEPIX and to our field. We will miss your presence at HEPIX, and we wish you the best in the next stage of life!

Finally...

- Thanks to all for another interesting and productive workshop
 - Speakers – for your interesting contributions
 - Participants – for engaging in active discussions
- Thanks as well to the conveners, session chairs for helping with the organizational details (including editing of recordings!)
- Wish you all a healthy and safe year and looking forward to meeting you again soon!

HEPIX Spring 2022

