

HEPIX – Spring 2018 Wrap-up

Madison, Wisconsin

Introduction

- Welcome speech provided by Prof. Sridhara Dasu (Physics Dept. Chair)
- 94 registered participants
 - 37 from the Americas
 - 40 from Europe
 - 10 from Asia
 - 7 from non-academic institutions
- 66 contributions
 - Site Reports (13)
 - Networking & Security (10)
 - End-User IT Services & OS (4)
 - Storage & Filesystems (9)
 - Computing & Batch Services (11)
 - IT Facilities & Business Continuity (6)
 - Basic IT Services (6)
 - Grids, Clouds & Virtualization (7)

Site Reports

- Purdue Site report
 - Dedicated + opportunistic access to campus-wide resources
 - Plans to explore cloud + HPC resources
- BNL Site report
 - New support for Belle-II and HPC activities
- AGLT2 Site report
 - Virtual hosts via VMWare at both UM and MSU
 - dCache transition to IPv6 (dual-stack) completed
- Nebraska Tier 2 Site report
 - Access to other local resources (HPC, excessed equipment, etc)
 - IPv6 support with select tools/services (Ganglia, HDFS, Docker, etc)
- NERSC Site report
 - Slurm @ PDSF for serial jobs (STAR Tier-1, ALICE Tier-2, ATLAS Tier-3, etc)
 - Migration of PDSF workloads to Cori by April 2019
- IHEP Site report
 - Support for HTCondor (HTC) and Slurm (HPC) and resource sharing for all experiments
 - EOS and IHEPBox deployment

Site Reports

- DESY Site report
 - Photon Science activities ramping up
- CERN Site report
 - Increased usage of various services (InfluxDB, CVMFS, EOS, SWAN, Digital Library, etc)
- INFN Site report
 - Facility back online after flood restoration, some data loss and equipment damage
- PIC Site report
 - Under-performance issue with oil-immersed cpu's
- NIKHEF Site report
 - Issues with auto-installation of Active Directory preview patches
- RAL Site report
 - Large HPC cluster with low-latency interconnect
 - Problem with UPS/generator back-up power
- FZU Site report
 - Comparison of Huawei and Intel servers using ATLAS workloads
 - ATLAS access to HPC resources (~20% of ATLAS capacity provided)

End-User IT Services & OS

- Troy's presentation on Modularity
 - Run bleeding-edge software vs. keep a stable software environment – hybrid OS
 - Demo available on slides
- Scientific Linux update from FNAL
 - RH6 support EOL is Nov 2020, RH7 supported till 2024, no news on RH8 yet
 - Brief description of new features and firmware updates
- IN2P3 report on the User Portal
 - Instrumented website as entry point for users
 - Initial deployment April 2018—user feedback for further development
- Trident presentation from CERN
 - Tool to collect and analyze workload performance
 - Sample analysis of HEP workloads—focus on select hardware metrics

Networking & Security

- WLCG/OSG networking update
 - OSG services undergoing transition due to funding changes
 - Increased monitoring to improve diagnostics and troubleshooting
- Update on IPv6 working group
 - Increasing level of dual-stack readiness
 - FTS performance difference between IPv6 vs. IPv4
- IPv6 experience at KIT
 - Issues with FTS transfers to IPv4-only sites – migration to dual-stack not so straightforward
- Security update from CERN
 - Bitcoin mining and mitigation procedures
 - Lively discussion about the continuing tension between security, freedom and convenience
- Federated Identity Management news
 - AAI improvements since 2012 white paper—effort to update white paper
 - IGTF effort to enable global trust relations among infrastructure providers and identity providers
 - Activities by WLCG Authorization WG on AAI providing membership management & token translation services

Networking & Security

- Network Functions Virtualization WG
 - Formed to evaluate if SDN/NFV should be deployed in HEP—initial report in early 2019
- KEK network updates
 - Renewal of campus LAN this Fall—current equipment no longer meets demand
- Cyber infrastructure at CSTC (CAS)
 - Description of research activities & IT capabilities
- Network infrastructure status at IHEP
 - Web interface for network traffic monitoring system
 - Joined LHCONe in March 2018
- Computer Network Information Center plans
 - Provides flexible network service to research institutes across China
 - Created own custom white box switch

Storage & Filesystems

- OpenAFS
 - 1.6.x stable release, 1.8 released April 2018
 - Presentation by OpenAFS Foundation
 - Strengthen, modernize and advance OpenAFS
- WLCG Archival Storage group
 - Report on progress in establishing a knowledge-sharing community and outline future activities
- KIT Online Storage System
 - New, scalable GPFS-based storage with IB interconnect fabric
- New Storage Services @ CERN
 - Description of various services (EOS, CERNBox, SWAN, CVMFS, etc) deployed
- AuriStorFS
 - Latest developments and updates on Linux AFS, AuriStorFS, etc
- Storage performance @ NIKHEF
 - Balancing I/O throughput with increasing storage capacity

Storage & Filesystems

- xCache/xRootd Cache for Southern California
 - Merge of Caltech and UCSD namespace
 - Scaling test results are promising—think about scaling and further expansion (Northern California?)
- AFS and Linux containers
 - AuriSTorFS integration with container and orchestration tools

Computing & Batch Services

- Benchmark WG report
 - Description of various activities, including searching for a replacement/alternatives to HS06
- Systems Performance and Cost Modeling WG report
 - Aim to characterize resource usage by workloads
- Shared HTCondor pool @ BNL
 - Proposed plan to convert dedicated clusters into a shared pool—need pre-emption feature in HTCondor
- DESY Batch news
 - Multiple batch systems to support heterogeneous user community (HTC and HPC)
 - Plan to keep AFS/Kerberos structure
- Batch On EOS Extra Resources (BEER)
 - Make better of CPU resources on storage servers
 - Configure and monitor HTCondor to minimize effect on I/O performance

Computing & Batch Services

- Techlab
 - CERN IT hardware lab available to HEP community
- HTCondor news
 - Description of several features in current release
 - New Stable series v 8.8 release this summer with various new features
- PDSF news
 - Retirement of legacy equipment and consolidation around Cori
- Swiss HPC Tier2 @ CSCS (Lugano)
 - Availability of HPC resources to LHC experiments (now and in the HL era)
- Benchmarking @ BNL
 - HPL and HPCG benchmark results on selected hardware platforms
- LNLS (Brazil) computing activities
 - Fast online reconstruction and analysis (common in Photon Sciences)

IT Facilities & Business Continuity

- ExDeMon
 - Scalable monitoring & alerting tool with machine-learning capabilities
- BNL data center news
 - Funding review this summer—occupancy in 2021
- Wigner datacenter network infrastructure
 - Redesign and upgrade capabilities
- INFN Flooding Incident update
 - Remediation and recovery of IT equipment and data underway
- Technology & Market Evolution
 - Market saturation, limitations of current technologies and pricing paint a challenging picture over the next few years
- Proposal for technology watch WG
 - Systematically follow trends, help predict costs for better investments
 - Present regularly at HEPIX, CHEP, WLCG and GCB meetings

Basic IT Services

- Jupyter notebook @ BNL
 - Initial deployment of new analysis platform
 - Integrate with other facility services and improve user experience
- HA logs/metrics collector @ CSCS
 - Facilitates post-incident analysis and response
- Monitoring Infrastructure @ CERN
 - Consolidate monitoring service to DC and IT services
- Salt & Reclass @ NIKHEF
 - New configuration management tools gradually being deployed
- Linux tools @ DESY
 - Enhancement or deployment of several services (2FA, Fail2ban, auto updates from repositories, kernel updates)
- Hadoop and Spark @ CERN
 - Growing demand for Big Data platforms and tools
 - Integration with other tools such as SWAN

Tape Storage BoF

- Market consolidation is decreasing choice of technologies, increasing costs and raising doubts about product availability
- Lively discussion on future deployment of tape technologies
 - Many plan to migrate to LTO-8
- Presentation by FNAL
 - Reliance on Oracle technology is a concern
 - Purchased two TS4500 (LTO8 drives) and retired two SL8500 libraries recently
- Presentation by BNL
 - Phased R&D plan to study tape storage performance with various ATLAS workflows – start at BNL and expand to other Tier 1 later

Grid, Cloud & Virtualization

- CERN Cloud update
 - Container clusters hosted at CERN and Wigner, increasing utilization
- HNSciCloud project
 - Pan-European collaboration aims to provide common cloud platform to its multi-disciplinary scientific community
- Ironic integration in CERN cloud
 - Baremetal provisioning to support more use cases
- STFC Cloud developments
 - Gradual migration from OpenNebula to OpenStack and increasing support to a diverse community
- STFC - DAaaS
 - Support for various facilities leverage existing tools and infrastructure

Grid, Cloud & Virtualization

- KEK integration with OpenStack and AWS
 - Performance comparison of on-site with AWS with various use cases
- LIGO activities
 - Leveraging existing tools (HTCondor, OSG, CVMFS, containers, etc) to support computing needs

HPC and AI Innovation Lab

- Dell-EMC initiative to provide a resource to scientific communities to design HPC systems, support R&D activities, develop best practices and evaluate new technologies (among other goals)

Major Themes & Trends

- HTCondor increasingly adopted at HEPIX sites
- AAI (Authorization & Authentication Infrastructure) was a frequent theme this week
 - Important feature for non-HEP domains as they join/leverage existing computing resources
- Increasing presence of HPC sites @ HEPIX
 - Potential for synergy and more collaborative effort
- Leveraging support (and consolidation of effort) for other (non-HEP) communities at various HEPIX sites—challenges with collaborative effort
- Tape Storage BoF highlighted narrowing options in technology choices

HEPIX Board Meeting Highlights

- Ajit Mohapatra has accepted an invitation to join the HEPHX Board
- Changes to Working Groups
 - Suspend Monitoring tools and Bit Preservation WG's (inactive)
 - Will Strecker-Kellogg has agreed to lead Batch Monitoring WG (announcement about first meeting soon)—initial focus on HTCondor
- HEPHX presence on LinkedIn coming soon

HEPIX Fall 2018

- Next meeting hosted by PIC in Barcelona (Oct. 8-12)
 - First HEPHX meeting in Spain
 - Local organizational effort (led by J. Flix) making good progress
 - General announcement (and website) coming very soon



Future Meetings

- San Diego in Spring 2019 (hosted by UCSD/SDSC)
- Amsterdam in Fall 2019 (hosted by NIKHEF)

Acknowledgements

- Congratulations to Ajit, Dan and the local organizing team for a well-run meeting and social events!
- Thanks to sponsors (Dell, KingStar Computer)
- Of course, thanks to all of you for making this an interesting meeting!



See you in Barcelona!