

# HEPIX Spring 2019 Summary

Tony Wong

BNL

# Meeting Statistics

- 88 participants
  - 31 from North America
  - 39 from Europe
  - 4 from Asia and elsewhere
  - 14 from corporate sponsors
- 64 accepted abstracts
  - Site Reports (19)
  - Storage & Filesystems (10)
  - Grids, Clouds & Virtualization (8)
  - Networking & Security (7)
  - Computing & Batch Systems (7)
  - IT Facilities & Business Continuity (7)
  - End User IT Services & OS (3)
  - Basic IT Services (2)
  - Other (1)
- Webcasting most presentations—recordings to be made available

# SDSC/UCSD

- Brief introduction by Mike Norman on SDSC history, usage and current developments
  - 35+ years of operational history
  - 99% of jobs use single rack and ~10% use single node (16 cores)
  - Support for HTC and OSG computing
  - Support for bursting out to cloud computing

# Site Reports

- NERSC turning off PDSF in April and Edison in May 2019, migrating workload to Cori
- Migration to FreeIPA to BNL—federation and SSO on-going
- Canada's Tier1 services migrating from TRIUMF to SFU—support continuity provided by TRIUMPF staff
- Optimization of computing resources at AGLT2 with HTCondor to increase cpu utilization rate
- UW-Tier2 – opportunistic usage of CHTC (average of ~1.5k cores for CMS production)
- Nebraska Tier2 reported extending useful lifetime of equipment to stretch resource availability –general push to use Leadership Class Facilities HPC resources
- KEK reported on ramp-up in resource utilization due to Belle-II and J-PARC physics data
- Tokyo Tier2 completed migration to new cpu and disk hardware in January (total downtime was 34 hours)

# Site Reports

- UW's CEMPA description of bringing up a re-purposed NERSC cluster for usage
- Jlab reported on a new GPU cluster for LQCD and providing support for cloud computing
- IHEP presented remotely—expand services in HPC activities and plans to leverage remote, heterogeneous clusters to augment resources
- CERN's decommissioning of the Wigner remote site—equipment to go into containers while 2<sup>nd</sup> data center plans are being discussed
- Infrastructure work still on-going at INFN as part of plan to address previous data center flooding
- RAL's SCD Cloud delivers dynamic compute resources across STFC and externally (including LSST and Diamond)
- DESY's developments regarding future AFS and Windows upgrades

# Site Reports

- Increasing deployment of EPYC-based servers (Nikhef, FZU, GSI)—reported variable gains in price/performance based on historical data
- Diamond leverages services at STFC RAL and described existing resources available to Photon Science community
- FNAL updates on current and future HEP experimental activities, storage updates, HEPCloud, InCommon, etc

# End-User IT Services & OS

- SDSC presentation on ML algorithm to classify user tickets and possibly automate response to common questions
- CERN presentation on user end-point management—survey available on Indico page for the San Diego meeting
- NERSC's project to consolidate cluster management & provisioning—simplify staff effort and present consistent experience to users

# Networking & Security

- CERN has plans to modernize support code for DNS and DHCP services
- Security presentation by Stefan Lueders—freedom, security, convenience, choose two
- BNL presentation on migration to FreeIPA to support SSO and federated access
- Nikhef presentation on performance and scaling tests on selected network equipment
- Shawn gave an update on OSG/WLCG Networking activities
- Andrea presented the status of IPv6 readiness across sites (T1, T2, etc)—progress, but still ways to go
- Shawn reviewed status of activities in the Network Functions Virtualization WG—white paper coming soon



# Storage & Filesystems

- OpenAFS description of latest developments (releases, bug fixes, security updates, etc) and contributions by user community
- VU presentation on a HDFS-LS integration project to optimize processing performance of large volume of small data files
- RAL presented status of migration to ECHO (Ceph-based disk storage) and tape services (currently based on CASTOR but in discussions about the future)
- Update on Dynafed by Marcus—redirector for a dynamic data federation
- Presentation on OSIRIS and mechanism for federated access with CoManage and usual IDP's (InCommon, EduGain, etc)
- KIT's presentation on oVirt/RHV-based virtualized online storage cluster activities
- FZU described activities related to upgrade of legacy grid storage system to DPM Dome

# Storage & Filesystems

- CERN update on status of storage services (EOS, Castor, AFS, Swan, CERNBox, Ceph, CVMFS, etc)
- plans to phase out AFS has been delayed but still on target for completion before run 3—disconnection test for one week beginning April 3<sup>rd</sup> dddd
- BNL activities on large-scale storage hardware management
- WekaIO presented a GPU-optimized storage filesystem that significantly outperforms other popular filesystems
- BoF session on storage with contribution from AuriStor

# Computing & Batch Systems

- Michele provide a status report on the Benchmarking WG
  - HS06 obsolete, evaluating CPU Spec2017
  - Experiment-based benchmarks promising—use LHC workloads but other contributions welcome
- Fair(?)—share mechanism at IN2P3—proposals to address and fine-tune batch policy
- CERN’s presentation about on-demand provisioning and elastic scaling to large compute pools with Spark on Kubernetes and SWAN
- Will presented for BNL a MFA-enabled Jupyter cluster available on HTC and HPC resources
- DESY’s presentation on Photon Sciences computing practices—very different from HEPN community!
  - Small teams with little (no?) computing expertise
  - Challenges with detector and accelerator improvements
  - The necessity of data reduction techniques

# Computing & Batch Systems

- CSCS report on activities at the Swiss SuperComputing Center
  - ATLAS/CMS workloads on HPC environment
  - Potentially leveraging HPC technologies (ie, gpu's) for Tier2 facility usage
- IN2P3's presentation on deep learning in containers
  - Address needs (ie, simulations) in HEP
  - Potentially leverage GPU's for HL era

# IT Facilities & Business Continuity

- Andrea's presentation on WLCG/HSF on cost performance modeling
  - Gap between requirement and affordability is narrowing but still factor of  $\sim 2$
  - Future work includes refining requirement estimates, understand TCO of infrastructure and identify areas of savings
- BNL's description of new data center and migration plan without service interruption
- Superfacility project at NERSC to coordinate, manage and integrate various aspects needed for scientific computing workflow (computing, storage, network, software and infrastructure)

# Technology Watch WG summary

- Track evolution of IT technologies market trends
  - Current focus on HEP and LHC, but open to contributions from other fields
- Andrea presented a comprehensive summary of recent market developments for cpu, gpu, fpga, memory, etc
  - AMD return to the data center
- Martin highlighted a few trends in storage
  - Shrinking market for HDD, increasing usage of SSD
  - Concern with shrinking competition in tape storage market
- Michele showed some highlights in server trends
  - Rising market for x86 servers
  - Concern about power usage and transition to liquid-cooled racks
- Network highlights presented by Rolf
  - 5G will contribute to rising network traffic (3x higher in 2022 when compared to 2017)
  - 100 Gbps becoming commodity and 400 Gbps is available
  - Several network evolution areas discussed

# Basic IT Services

- centralized configuration and deployment at KIT for data-intensive computing
- SLAC presentation on the need to upgrade its token renewal service—many choices out there to cope with contemporary computing environments

# Grids, Clouds & Virtualization

- Boyd Wilson (Omnibond) described two projects
  - AWS cloud cluster (scale up 1.1M cores in 2 hours)
  - Use traffic video data with AI algorithms for congestion analytics and improve traffic flow
- UV's on-going work on CloudScheduler v2, an upgrade of v1 originally released in 2009
  - Used by ATLAS and to be extended to Belle-II
  - Report on scaling tests
  - Support for AWS soon
- GlideinWMS recent developments
  - Management system for dynamic HTCondor pools
  - Re-engineered front-end for FNAL's HEPCloud project
  - New features/improvements, such as whole-node allocation, singularity support
- KEK described activities and challenges to support grid computing needs , in particular in support of Belle-II requirements



# Grids, Clouds & Virtualization

- Igor described activities in provisioning an opportunistic OSG site within California's Pacific Regional Platform (PRP) cluster
- Entonos' experiences with HTC at public clouds—adjustments to computing models to maximize throughput/cost ratio
- A description of the evolution of cloud services at CERN over last 6 years—from 0 to 300k cores
- SLATE project (distributed service orchestration platform) for management of edge services remotely—stimulate creation of multi-institution research platforms
- Major OSG re-organization in Spring 2018—a summary of resulting changes and updates in services

# HEPIX Board meeting summary

- Edgar Fajardo Hernandez has accepted an invitation to join the HEPHX Board
- HEPHX Fall 2019 (Oct. 14-18) will be organized by Nikhef in Amsterdam
  - HEPHX last held in Amsterdam in 2003
  - Local organizational effort led by Dennis Van Dok
  - General announcement (and meeting webpage coming soon)
- On-going discussions on future meetings in 2020 and beyond—more details to be announced soon
- Presentations by working groups in San Diego
  - IPv6
  - Network Functions Virtualization
  - Benchmarking
  - Technology Watch
- A complete list of working groups can be found on [www.hepix.org](http://www.hepix.org). HEPHX encourages you to join and contribute. Please contact the WG coordinator (see website) for more information

# Acknowledgements

- Congratulations to Edgar, Frank, Cindy, Susan and the entire local organizing team for a well-run meeting and social events!
- Thanks to all sponsors (AuriStor, Weka.IO, Teradactyl, Spectra Logic, Entonos)
- Of course, thanks to all of you for making this an interesting meeting!



# See you in Amsterdam!

HEP*i*X

[More Information](#)

See you at the next meeting in

Nikhef - National Institute for Subatomic Physics

Amsterdam, The Netherlands

October 14 - 18, 2019