



# HEP*i*X Spring 2019 Summary

<https://indico.cern.ch/e/hepix-spring2019>

Andrei Dumitru ▪ Quentin Barrand ▪ Piotr Mrowczynski

Hosted by University of California in San Diego at San Diego Supercomputing Center

# HEPiX

CERN IHEP DESY INFN PIC NIKHEP RAL FZU BNL ...

Future plans

Working groups

Challenges



AGLT2 KIT GSI TRIUMF FNAL KEK IN2P3 NERSC ...

Recent work

Status reports

Experiences

<https://www.hepixon.org>

**Andrei**

Spring 2019 Meeting and General HEPiX News  
Security and Networking  
Basic IT Services  
Site Reports

Computing and Batch Services  
Storage and Filesystems  
Miscellaneous

**Piotr**

**Quentin**

Grid, Cloud and Virtualisation  
IT Facilities and Business Continuity  
End-User IT Services and Operating Systems



# HEPiX 2019 Spring in numbers

88 registered participants

39 from Europe

31 from North America

4 from Asia

14 from companies/corporate sponsors





# HEPiX 2019 Spring in numbers

64 contributions

Site Reports (19)

Networking & Security (7)

End-User IT Services & OS (3)

Storage & Filesystems (10)

Computing & Batch Services (7)

IT Facilities & Business Continuity (7)

Basic IT Services (2) Misc. (1)

Grids, Clouds & Virtualisation (8)







35+ years of operational history

SDSC is home to:

- COMET supercomputer
- Open Science Grid
- Sherlock Cloud on AWS (burst to cloud)

Users:

- Main users: biology, neuroscience
- Small HEP community

# Computing and Batch Services


- Performance benchmarking and optimization for grid computing
- WLCG Site infrastructure developments
- Computing with Jupyter notebooks

# Computing and Batch Services

## Performance benchmarking and optimization for computing

### Presentation from HEPIX Benchmarking working group:

- Benchmarks to classify hardware (single hosts) running in WLCG
- HS06 benchmark obsolete, evaluating SC2017 (similar results). New benchmarks based on example jobs provided by experiments. Gitlab group hep-benchmarks for contributions



ALICE  
– A p-p simulation job using Geant3  
ATLAS  
– A ttbar simulation job using Geant4  
– A digitization + reconstruction job  
– A DxAOD derivation job  
CMS  
– Generation + simulation of ttbar events  
– Digitization and trigger simulation with premixed pile-up  
– Reconstruction job producing ADDSIM and MINIAODSIM  
LHCb  
– Generation + simulation using Geant4 of  $D^0 \rightarrow \pi(D^0 \rightarrow K \pi) \pi \pi$  events

### Analysis of fair-sharing in grid from IN2P3—proposals to address and fine-tune batch policy:

- Analysis and proposals on optimisation and configuration of queues to speed-up local jobs without harming grid jobs (UGE Univa Grid Engine used - looking also at HTCondor. Storage GPFS, HPSS, iRODS, dCache, XrootD)

#### Redefinition of the scheduling queues

Workload	Scenario	Average	Percentiles				Maximum
			50 <sup>th</sup>	75 <sup>th</sup>	90 <sup>th</sup>		
Grid	Baseline	1h 10m	0s	8m 18s	1h 18m	15d 21h 54m	
	Modified	1h 45m	0s	14m	2h 2m	14d 4h 33m	
Local	Baseline	2h 3m	4m 30s	1h 40m	6h 40m	11d 21h 41m	
	Modified	1h 58m	8s	1h 10m	6h 20m	4d 19h 6m	

#### Quota Relaxation

Workload	Scenario	Average	Percentiles				Maximum
			50 <sup>th</sup>	75 <sup>th</sup>	90 <sup>th</sup>		
Grid	Conservative	1h 53m	0s	16m	2h 21m	13d 15h 21m	
	Extreme	1h 57m	4s	17m 41s	2h 47m	14d 4h 41m	
Local	Conservative	1h 39m	2s	45m 40s	5h 8m	3d 16h 58m	
	Extreme	1h 14m	1s	21m 55s	2h 30m	3d 23h 11m	




# Computing and Batch Services

## Site infrastructure developments


**Presentation from DESY on infrastructure evolution for future online processing requirements:**

- With future evolution current workflow (Detector -> GPFS -> offl. Ana -> tape) will not be efficient.
- Evaluating online analysis with ASAP:0 system: Detector -> Reduction Techniques (no raw data stored) -> Storage -> Online Analysis on data with Jupyter Notebooks




## Presentation from Swiss HPC Tier-2 (CSCS):

- Investigating use of HPC (including GPUs) dominantly for HL-LHC with dCache and GPFS scratch space



## Presentation from IN2P3 on GPUs farm:

- Investigating use of GPUs with Tensorflow on grid GPU farm, with containers (Singularity)



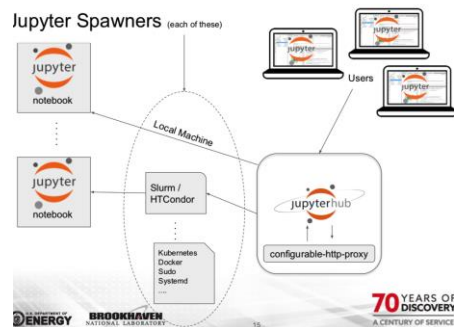
1. Install the tensorflow module you need depending on your python environment (python 2.7 or python 3.6) and the GPU type you want (K80 or V100)
2. Submit your code to the GPU farm, specifying which image you want to run it, and your python environment
3. The GPU farm computes your code through the specified environment

# Computing and Batch Services

## Computing with Jupyter notebooks

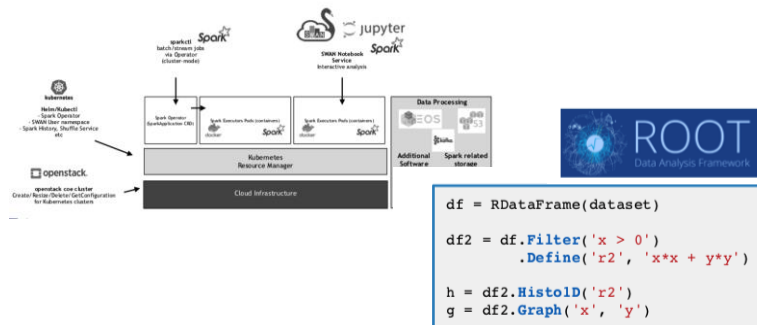
### Presentation from Brookhaven National Laboratory on Notebooks integrations:

- Integration of Jupyter Notebooks into their computing environment (Slurm/HTCondor - interactive submission of set of jobs or scheduling notebook as single job)
- Notebooks scheduled on HTC/HPC resources



### Presentation from CERN IT-DB-SAS on Notebooks, Spark and RDataFrame:

- Challenges and motivation for integration of Jupyter/Spark Notebooks with K8s Container Service, compared to existing on-premise Hadoop
- Interactive analysis for physics with Distributed RDataFrame on real user analysis example (Proton's Elastic Scattering)



# Storage and Filesystems

- Updates on Software
- Site infrastructure developments and updates



# Storage and Filesystems


## Updates on Software

### Presentation from OpenAFS:

- 1.8.X is new stable release of OpenAFS (platform support for new releases of linux, fixes)
- 1.6.X old stable will be still supported and no EOL so far (backports from 1.8.x)
- 1.8.3 will soon be released (platform support and maintenance fixes)

### Presentation from WEKA.IO:

- presented a GPU-optimized storage filesystem that outperforms other popular filesystems on data processing on GPUs (storage in S3-compatible systems)



### Presentation from FZU on Disk Pool Manager:

- DPM old code will be discontinued, moved to DOME (Disk Operations Management Engine). DPM Sites should move to DOME




# Storage and Filesystems

## Site infrastructure developments and updates

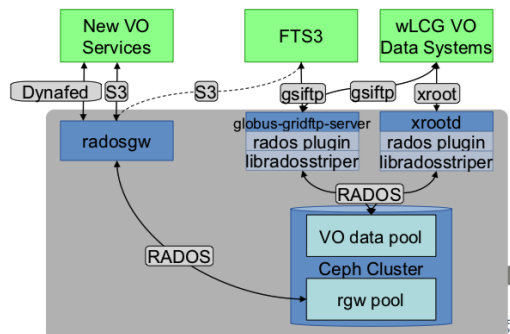
### Presentation from Vanderbilt Uni on LStore HDFS Integration:

- Presented cost-effective replacement for GPFS with LStore and HDFS:  
lstore:// connector via `org.apache.hadoop.fs.FileSystem.LStoreFileSystem`
- LStore is generic block level storage abstraction for HPC multi-user environment, integration of HDFS boosts read effectiveness



### Presentation from UK RAL (Tier 1) on ECHO:

- Presented ECHO disk storage - erasure-encoded high-throughput object store. CEPH based object storage, access via GridFTP/XRootD/S3. Goal of ECHO is maximising throughput. Random-access for their use-case no important. All data in erasure code, no replication used (even for hot data)
- Improvements for CASTOR (minimised configuration complexity)






# Storage and Filesystems

## Site infrastructure developments and updates

### Presentation from Karlsruhe Institute of Technology (Tier 1) on use of virtualized storage cluster

- Virtualization cluster using oVirt / RHV with KVM hypervisors. Hyperconverged setup with Gluster storage.
- Server VMs are “close to storage” in hyperconverged
- Tests with HPC access to storage cluster



### Presentation from Brookhaven Laboratory on their storage management practices

- Brookhaven presented how they manage their JBOD storage racks (they showcased scripting they use)

### Presentation from CERN on storage services:

- Update on status of storage services (EOS, Castor, AFS, Swan, CERNBox, Ceph, CVMFS)
- Plans to phase out AFS has been delayed but still on target for completion before run 3. Announced to sites using the AFS the disconnection test
- Presented the FuseX, new namespace architecture with QuarkDB, Migration to EOSHOME and CERNBox as App Hub with SWAN, Office Suites etc
- Presented evolution of CEPH (2nd region for s3, demand as backup-storage, CVMFS backend phase-out to s3)

# Grid, Cloud and Virtualization (1)

## Addressing the Challenges of Executing Massive Computational Clusters in the Cloud

Boyd Wilson (Omnibond)

- Omnibond Automaton / PAW tool to create, scale, use and destroy a cluster in commercial cloud environments
- CCQ: meta-scheduler over several clusters
- Two examples were presented:
  - Using CloudyCluster to spawn a cluster of 1.2 million cores on AWS in less than two hours; increased (~ x2) the public price of compute resources in AWS' US regions
  - Using traffic video data with AI algorithms for accident detection, congestion analytics and improve traffic flow

## CloudScheduler version 2

Rolf Seuster (University of Victoria)

- CloudScheduler starts VMs which run HTCondor when they boot
  - Currently compatible with OpenStack, AWS next
- Rewrite of CloudScheduler v1 released in 2009
  - Python 3 and MariaDB
  - The CLI and the GUI call the same REST endpoints
- In production in ATLAS since several months; in Belle II next

# Grid, Cloud and Virtualization (2)

## The GlideinWMS system: recent developments

Marco Mascheroni (UCSD)

- GlideinWMS looks at user jobs in the Condor schedulers
  - Matches them to Computing Elements
  - Determines the number of pilot jobs to submit via Glidein factories
- Supports EC2 and GCE
- New features have been introduced
  - Singularity support
  - Whole-node allocations

## The Experience and Challenge in Grid Computing at KEK

Go Iwai (KEK)

- KEKCC:
  - 10k CPU cores, 2TB memory, Linux cluster and storage system
  - Replaced every 4-5 years
  - Mostly running Belle II workloads now
- Data migration from Belle to Belle II
  - 1PB, 15M files
  - Estimated 2 months, took one year (busy system, software issues)
- Massive Data Transfer performance issues
  - Some workarounds found thanks to efforts from both KEKCC and Belle II
  - Highly anticipated KEKCC upgrade in 2020



# Grid, Cloud and Virtualization (3)

## Creating an opportunistic OSG site inside the PRP Kubernetes cluster

Igor Sfiligoi (UCSD)

- Pacific Research Platform
  - Kubernetes cluster of 2.5k CPU cores, 250 GPUs
  - Nautilus web interface for self-service
  - Mostly ML/AI use cases
  - Little contention: happy users, but GPUs underutilized
- Opportunistic jobs for OSG users
  - Mostly from the IceCube community
  - One worker node per available GPU
  - Pod priority very low to allow preemption; HTCondor supports it
- Container-Storage-Interface (CSI) to mount CVMFS in the containers
  - Work In Progress in liaison with CERN
- Native OSG support for Kubernetes is desirable to reduce the amount of glue code

# Grid, Cloud and Virtualization (4)

## Public cloud for high throughput computing

Dr Gregory Parker (Entonos)

- Compute-intensive simulations of modern hard drive techniques (heat assisted magnetic recording - HAMR)
- Leveraged AWS Spot (preemptible) instances and S3 for storage
- Fleets of 100k concurrent simulations achieved with 99,99% of the cost just for compute
  - Eventually scaled out to a million without modifying the abstraction layer
- Cloud surprises: variations in compute performance, S3 latency at scale

## 6 years of CERN Cloud - From 0 to 300k cores

Belmiro Moreira (CERN)

- LxCloud, CVI, OpenStack
  - From 2 to 70+ cells; upgraded to CellsV2
  - Storage: Cinder, Manila, S3
  - Deprecation of Ceilometer, as it was hard to maintain
- In 2019
  - 2 cloud regions: Batch vs. Personal / Services
  - Preemptible instances
  - OpenStack control plane in container clusters
  - nova-network → Neutron migration, SDN
  - Enrolling of hypervisors in Ironic

# Grid, Cloud and Virtualization (5)

## Developing for a Services Layer At The Edge (SLATE)

Ben Kulbertis (University of Utah)

- SLATE: Kubernetes-based distributed service orchestration platform
- Developers deploy services based on Application Catalogs (Helm charts)
  - Can be contributed to a Git repository via a merge request
- Developer environment:
  - Minikube considered, but it has limitations and overhead
  - MiniSLATE: local, dind one-node SLATE setup

## Changes to OSG and how they affect US WLCG sites

Jeffrey Michael Dost (UCSD)

- Spring 2018: OSG transitioned to Indiana University
- OSG CA discontinued; sites should use Let's Encrypt or IGTF InCommon
- OIM transitioned to OSG Topology (catalog of all OSG sites stored on GitHub)
- Hosted CE: OSG-managed compute elements that just require remote SSH login
- OASIS (CVMFS repository), StashCache (XRootD federation for VOs), OSG Glidein factory...

# IT Facilities and Business Continuity (1)

## **Cost and system performance modelling in WLCG and HSF: an update**

Jose Felix Molina (CIEMAT), Dr Andrea Sciabà (CERN)

- Gap between available resources and expected needs for HL-LHC
  - 10x increase in trigger rates, 5x increase in pile-up
  - CPU and disk short by a factor of 2
- The WG identified some reference workloads for each experiment
  - Having a common framework for modelling the computing requirements turns out to be unrealistic
  - A spreadsheet implements a model for calculating TCOs and budget; discussions ongoing to reach an approach usable by all sites

## **RACF/SDCC Datacenter Transformation**

Alexandr Zaytsev (BNL)

- Moving from their 1.5MW air-cooled datacenter to a new building (former Light Source)
  - Fully power redundant
  - Water-cooled racks
- Construction will start in April 2019; move expected to be completed by FY23
- 1.6Tbps between the two data centers
- No service interruption foreseen



# IT Facilities and Business Continuity (2)

## **Omni, N9 and the Superfacility**

Cary Whitney (LBNL)

- Perlmutter (2020): new Cray system providing 3-4x the performance of the existing Cori installation
- Leverages Cray's Slingshot networking: 3 hops max between any node, ethernet compatible
- All-flash filesystem: 30PB, 4Tbps bandwidth, 7 000 000 IOPS
- Support for containers and Python-based analytics with GPUs
- 12MW upgrade to the facility required

## **HEPiX Technology Watch working group**

Helge Meinhard (CERN)

- Created in Spring 2018; 58 subscribers to the ML
- Organized in six sub-groups:
  - General market trends
  - CPUs and accelerators
  - Servers
  - Memory
  - Storage
  - Network
- Server costs: CPU and storage price/performance thought to be improving by 15% and 20% per year; ultimately found to be lower than that. WG participants should find a way to gather some data to better understand the trends

# IT Facilities and Business Continuity (3)

## Network TechWatch report

Rolf Seuster (University of Victoria)

- Internet traffic is expected to increase by a factor of 3 between 2017 and 2022
- Ethernet evolutions
  - P802.3cn: 50/200/400Gbps over single-mode fiber
  - P802.3cp: bidi 10/25/50 Gbps for the Access
  - P802.3ct: 100/400Gbps over DWDM
- PAM4 and FEC to drive higher speeds and port density
- 100 Gbps server interfaces introduced at SC18

## Server TechWatch report

Michele Michelotto (Università e INFN, Padova), Martin Gasthuber (DESY)

- Data centers use ~3% of the world's electricity out, could reach 20% in 2025
- 42U rack can draw around 40kW; watercooling becomes more and more popular, but there are concerns about the power consumptions
- Open19: new open-source platform for servers
  - Rack, cage or network switch
  - 4 form factors for servers
  - Allegedly more economical compared to Facebook's OCP
- Open Computing Platform: new Open Rack v2, Server or Drive Tray

# IT Facilities and Business Continuity (4)

## Storage TechWatch report

Andrea Chierici (INFN-CNAF)

- Hard Drive Disks
  - IOPS per TB continue to fall: dual actuator coming in 2019
  - Reaching the density limit for perpendicular magnetic recording: MAMR & HAMR should allow 40TB drives by 2025
  - Concerns about the shrinking market; narrowing price with SSDs
- Solid State Drives: expanding beyond SATA
  - NVMe eliminates some software layers in the OS stack; it's the center of industry attention
  - NVDIMM: persistent memory, DDR4 / DDR5 form factors (CPU support required)
- Tape: still key technology for large-scale archival
  - LTO: Gen-8 released in Q4 2017-360MB/s, 12TB
  - Concerns about the shrinking competition on both drives and media

# End-user IT Services and Operating Systems (1)

- **Text Classification via Supervised Machine Learning for an Issue Tracking System**

Martin Kandes (UCSD)

- HPC User Services Group: handles end-user support tickets
- Goals:
  - Classify incoming user tickets
  - Provide generic but useful automated replies as a first level of support
- Several models experimented: naive / scikit Bayes classifier, TensorFlow...
  - Scikit MLP showing the best accuracy
- Results promising
  - 0.5 to 0.95 accuracy, depending on the topic
  - Plans to explore more completely some ML approaches, such as multilayer perception (MLP) and Support Vector Machine (SVM) more completely
  - OAuth2 support with Outlook Mail API in the pipeline

# End-user IT Services and Operating Systems (2)

## Endpoint user device and app provisioning models

Michal Kwiatek (CERN)

- Goal: to understand endpoint user device landscape at CERN
- Data collection: anonymized SSO / email logs, user survey
  - 57% of users who access CERN e-mail use CMF
  - 26% of users who access CERN e-mail use also Self-Service
  - 5000 Android users, 4000 iOS users
- Strategy: focus on application management rather than device
  - Offer apps repositories for all platforms
  - Provision licences

## How to make your Cluster look like a Supercomputer (for Fun and Profit)

Georg Rath (LBNL)

- SMWFlow: Git-based management tool to manage system states, developed in partnership with Cray
  - Code reviews, testing, reproducibility
- Added a SMW layer on top of all systems for them to be manageable by SMWFlow
- Custom management stack built on top of xCAT and cfengine



# Site Reports

## NERSC/PDSF

- Migration to Cori (Cray XC40, 12.076 nodes, 700k cores) for workloads complete
- PDSF computing cluster turned off

## University of Nebraska CMS Tier2

- Extending useful lifetime of equipment to stretch resource availability

## BNL

- New data center - beginning of construction in April 2019
- Migrated from OpenLDAP to RedHat IPA
- Transitioning from Shibboleth to Keycloak for service auth
- Jupyter service has been re-architected with a unified environment on Red Hat Enterprise Virtualization cluster

# Site Reports

## Canadian T1 / T2

- all T1 services previously at TRIUMF moved to SFU (~30km or 1ms RTT away)
- new T2s in Canada at UVic, SFU, Waterloo, ETS (T3 facility, might be upgraded to T2)

## AGLT2 – ATLAS Great Lakes Tier 2

- Significant work to improve the utilization of the cluster (increase CPU utilization rate)
- Reconfigure low memory nodes, making all the cores being claimable by HTCondor

## University of Wisconsin-Madison CMS T2

- Using opportunistic resources at CHTC for CMS production jobs (average of ~1.5k cores used)
- Deployed XrootD over HTTPS and upgraded cluster from Scientific Linux 6 to CentOS 7.6

# Site Reports

## KEK

- Ramp-up in resource utilization due to Belle-II and J-PARC physics data
- The renewal of KEK campus network was completed in September 2018

## Tokyo Tier-2 - International Center for Elementary Particle Physics

- Hardware in the computing center supplied by a lease, replaced every 3 years
- Completed migration to new hardware as a new system was delivered in Jan 2019
- all Grid services and worker nodes are running on CentOS 7

## FNAL

- Moving service certs to InCommon CA after OSG CA shut down last year
- HEPCloud now live (scientific gateway to resources beyond local worker nodes and grids, expanding HPC centers and the cloud)

# Site Reports

## UW CENPA

- Repurposed previous generation of PDSF hardware donated by NERSC to upgrade their existing cluster

## JLab

- New GPU cluster installation is underway (32 servers, each with 8 RTX 2080 GPUs)
- Increase offsite resource usage at NERSC, OSG, and also in the cloud

## IHEP / BEIJING-LCG2

- Became LHCbTier 2 site
- Significant improvements to their monitoring and fault analysis platform

# Site Reports

## INFN

- Infrastructure work still on-going as part of plan to address previous data center flooding
- Next-Generation Firewall installed(Palo Alto PA-5250)
- 2 racks now implement an ISMS (CNAF got ISO27001 certification in 2017)
- Several scientific collaborations interested in hosting medical data there

## NIKHEF

- More systems now configured with Salt (LDAP, DNS, dCache)
- Migrated from Nagios to Icinga2
- Working on setting up a high-throughput private compute cloud with Openstack



# Site Reports

## RAL

- SCD cloud delivers dynamic compute resources to scientists across STFC and externally (IaaS platform built on OpenStack)
- New tape cache pool running CASTOR version 2.1.17-35
- Decommissioned “disk-only” CASTOR pools in favour of Ceph-based storage

## DESY

- Jupyter Notebooks very well received (especially by Photon Science users)
- Migration to Windows 10 - no DESY accounts or information in Microsoft Clouds
- Large fully featured MS Windows 10 roll-out (now 40:60% ratio for Win10 vs Win7 usage)
- Storage options considered for the future
  - replace OpenAFS with AuristorFS ? (tests underway)
  - going Sync&Share only? (based on NextCloud and dCache)

# Site Reports

## FZU / Prague

- New cluster based on HPE Apollo 35 (AMD EPYC 7301 CPU)
- Network upgrade - Catalyst 6509 and Nexus 3172 switches replaced with Catalyst 9500

## GSI

- GDPR: several inquiries requesting access to personal data (challenging for IT)
- Windows 7 to 10 migration in progress

## Diamond Light Source

- New scientific data processing cluster (2.560 Xeon Gold 6148 cores, 8PB GPFS 5)
- Making use of on-premise or STFC OpenStack cloud resources

# Basic IT Services

## Configuration Management and Deployment Setup at KIT

- many different approaches including Ansible, KITE (homegrown system) and Katello/Satellite/Foreman + Puppet
- Foreman/Katello manage all aspects of host life cycle
- Next steps: migration to Puppet 6

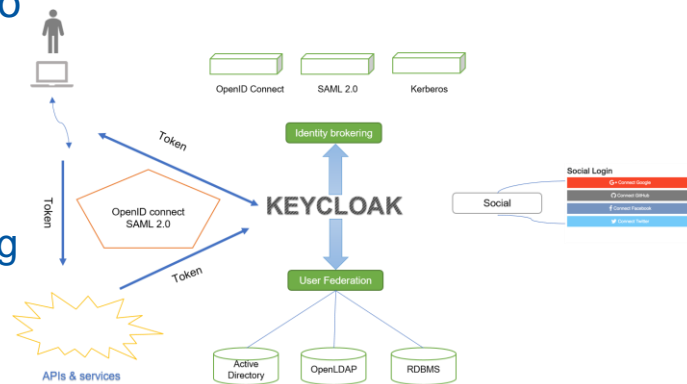
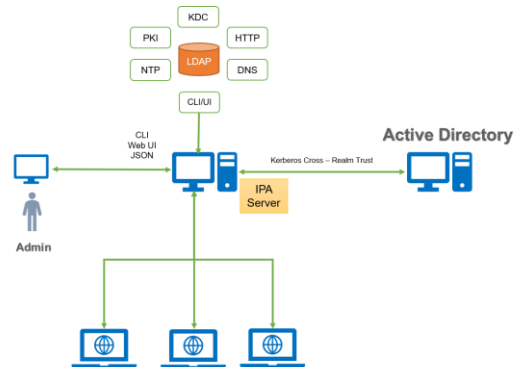
## Token Renewal Service at SLAC

- Need to upgrade its token renewal service – used since the late 1990s
- Many options available: kcron, k5cron, etc.
- Work in progress

# Network & Security

## BNL activities on federated access and Single Sign-On

- Why: several Kerberos realms, Shibboleth config/mgmt is complex, SSO and Federated access wanted
- Based on RedHat Identity Management (easy install and setup, redundancy with multi-master replication and read-only replicas)
- Migrated users from OpenLDAP to IPA
- Keycloak deals with authentication (2F auth easy to setup)
- Status:
  - Invenio test prototype now has federated access
  - Jupyter test instance has MFA access enabled using KeycloakQR codes



# Network & Security

## Situational Awareness: Computer Security

- Data leaks, CPU vulnerabilities, crypto mining
- Credential exposed, software licences, etc.

## DNS and DHCP service evolution plans @ CERN

- Rewrite DNS-related software in Go
- Use the ISC Kea DHCP server (database backend, working HA, REST API)

# Network & Security

## Differences in network equipment and how to break them

- Same chips in products from different vendors but differences in software stack, CLI, support, etc.
- Performance and scaling tests on selected network equipment

## IPv6 & WLCG – updates from the HEPiX IPv6 WG

- All T1 sites connected to LHCONE except RRC-KI
- TRIUMF, CERN, KISTI and RRC-KI still have some problem with IPv6 on perfSonar
- Tier-1s should already have production storage accessible over IPv6 (more than 85% of data is available via IPv6)
- 37% of FTS transfers today over IPv6



# Network & Security

## WLCG/OSG Network Activities, Status and Plans

- perfSONAR – 261 active instances
  - new plugins: network traffic capture and TWAMP (two-way active measurement protocol)
  - new central web interface and toolkit configuration mechanism
- ‘New’ Collaborating Projects: IRIS-HEP and SAND (coordinates network monitoring and diagnostics)

## HEPiX Network Functions Virtualisation WG Update

- Role: evaluate whether and how SDN and NFV should be deployed in HEP
- Wide range of SDN solutions for data centres exists
- Focus on Open vSwitch/Open Virtual Networking and Tungsten Fabric (both open source)
- SmartNIC now offered from multiple vendors
- SENSE and BigData Express leading projects in programmable networks
- Initial report/white paper - in progress

HEPIX

**Fall / Autumn 2019**  
14 - 18 October, Amsterdam

Nikhef





# Next meetings

NIKHEF

Amsterdam, The Netherlands

14 - 18 October 2019

ASGC

Taipei, Taiwan

Spring 2020



HEPIX

Fall / Autumn 2019  
14 - 18 October, Amsterdam

Nikhef

