

# HEPiX Fall 2016 Workshop



## Report of Contributions

Contribution ID: 1

Type: **not specified**

## CERN Computing Facilities Evolution

*Thursday, October 20, 2016 11:10 AM (25 minutes)*

This talk will give an overview of current activities to expand CERN's computing facilities infrastructure. This will include a description of the 2nd Network Hub currently being constructed as well as its purpose. It will also cover the initial plans for a possible second Data Centre on the CERN site.

**Primary author:** SALTER, Wayne (CERN)

**Presenter:** SALTER, Wayne (CERN)

**Session Classification:** IT Facilities and Business Continuity

**Track Classification:** IT Facilities & Business Continuity

Contribution ID: 2

Type: **not specified**

## Wi-Fi service enhancement at CERN

*Tuesday, October 18, 2016 3:15 PM (25 minutes)*

Over the last few years, the number of mobile devices connected to the CERN internal network has increased from a handful in 2006 to more than 10,000 in 2015. Wireless access is no longer a “nice to have” or just for conference and meeting rooms, now support for mobility is expected by most, if not all, of the CERN community. In this context, a full renewal of the CERN Wi-Fi network has been launched in order to provide a state-of-the-art Campus-wide Wi-Fi Infrastructure. Which technologies can provide an end-user experience comparable, for most applications, to a wired connection? Which solution can cover more than 200 office buildings, which represent a total surface of more than 400.000 m<sup>2</sup>, while keeping a single, simple, flexible and open management platform? The presentation will focus on the studies and tests performed at CERN to address these issues, as well as some feedback about the global project organisation.

**Primary author:** DUCRET, Vincent (CERN)

**Presenter:** DUCRET, Vincent (CERN)

**Session Classification:** Security & Networking

**Track Classification:** Security & Networking

Contribution ID: 3

Type: **not specified**

## Security Update

*Tuesday, October 18, 2016 2:25 PM (25 minutes)*

What's been happening in security for HEP? We will discuss the recent trends in the ever changing threat landscape, and the new initiatives being put in place to protect our people, data and services. One such initiative to highlight is our focus on bootstrapping international collaboration within research and academia, encouraging communities to participate in intelligence sharing and incident response. We will also discuss developments in the technologies being used to target us and the rest of the academic community.

**Primary author:** SHORT, Hannah (CERN)

**Presenter:** SHORT, Hannah (CERN)

**Session Classification:** Security & Networking

**Track Classification:** Security & Networking

Contribution ID: 4

Type: **not specified**

## Can we trust eduGAIN?

*Tuesday, October 18, 2016 4:35 PM (25 minutes)*

EduGAIN, the international identity federation, allows users from all over the world to access a globally distributed suite of academic resources. You are most likely already able to use your primary account, from CERN or your home organisation, to tap in to these services! Federated Identity Management, the technology underpinning eduGAIN, brings many benefits for users and organisations alike but... how can we trust these users with our HEP services? This is one of the questions that the AARC project (<https://aarc-project.eu>), in which CERN is a partner, is seeking to answer. We will discuss the measures being put in place to allow WLCG to reap the rewards of eduGAIN without exposing itself to increased risk.

**Primary authors:** SHORT, Hannah (CERN); WARTEL, Romain (CERN)

**Presenter:** SHORT, Hannah (CERN)

**Session Classification:** Security & Networking

**Track Classification:** Grid, Cloud & Virtualisation

Contribution ID: 5

Type: **not specified**

## Effective Data Retrieval from Massive Amounts of Tape-Resident Data

*Wednesday, October 19, 2016 4:10 PM (25 minutes)*

Randomly restoring files from tapes degrades the read performance primarily due to frequent tape mounts. The high latency and time-consuming tape mount and dismount is a major issue when accessing massive amounts of data from tape storage. BNL's mass storage system currently holds more than 80 PB of data on tapes, managed by HPSS. To restore files from HPSS, we make use of a scheduler software, called ERADAT. This scheduler system was originally based on code from Oak Ridge National Lab, developed in the early 2000s. After some major modifications and enhancements, ERADAT now provides advanced HPSS resource management, priority queuing, resource sharing, web-browser visibility of real-time staging activities and advanced real-time statistics and graphs. ERADAT is also integrated with ACSLS and HPSS for near real-time mount statistics and resource control in HPSS. ERADAT is also the interface between HPSS and other applications such as the locally developed Data Carousel providing fair resource-sharing policies and related capabilities.

ERADAT has demonstrated great performance at BNL and other scientific organizations.

**Primary author:** YU, David (Brookhaven National Laboratory (US))

**Presenter:** YU, David (Brookhaven National Laboratory (US))

**Session Classification:** Storage and Filesystems

**Track Classification:** Storage & Filesystems

Contribution ID: 6

Type: **not specified**

## INFN-T1 Status report

*Tuesday, October 18, 2016 10:00 AM (15 minutes)*

A short update on what's going on at the Italian T1 center.

**Primary author:** CHIERICI, Andrea (INFN-CNAF)

**Presenter:** CHIERICI, Andrea (INFN-CNAF)

**Session Classification:** Site Report

**Track Classification:** Site Reports

Contribution ID: 7

Type: **not specified**

## Deploying Open Compute hardware at CERN

*Thursday, October 20, 2016 10:45 AM (25 minutes)*

The Open Compute Project, OCP, was launched by Facebook in 2011 with the objective of building efficient computing infrastructures at lowest possible cost. Specifications and design documents for Open Compute systems are released under open licenses following the model traditionally associated with open source software projects. In 2014 we presented our plans for a public procurement activity for a small-size Open Compute hardware installation aimed at assessing the maturity of OCP market and whether it could be identified as a possible competitor of “traditional” hardware (Open Compute at CERN, HEPiX Spring 2014). We have finally deployed in September 2015 six Open Compute racks populated with CPU servers and storage enclosures in CERN’s Meyrin data-centre. We were presented with interesting challenges during all phases of the project and at all levels of the stack, from the power distribution to hardware monitoring. I will outline some of the hurdles we had to overcome and the lessons we have learnt along the way, together with the results obtained during the evaluation of the systems.

**Primary author:** GUERRI, Marco (CERN)

**Presenter:** GUERRI, Marco (CERN)

**Session Classification:** IT Facilities and Business Continuity

**Track Classification:** IT Facilities & Business Continuity



Contribution ID: 8

Type: **not specified**

## A Race for the Data Center: POWER8 and AArch64

*Wednesday, October 19, 2016 10:45 AM (25 minutes)*

x86 processors have been the long-time leaders of the server market and x86\_64 the uncontested target architecture for the development of High Energy Physics applications. Up until few years ago, interests in alternative architectures targeting server environments that could compete in terms of performance, power efficiency and total cost of ownership with x86 could not find any concrete response. However, the past few years have seen the introduction of new processor architectures and initiatives aimed at challenging the leading position of x86. With the introduction in 2011 of the ARMv8 Instruction Set Architecture supporting 64-bit, ARM set the first milestone for the expansion into the server landscape. The OpenPOWER Foundation founded in 2013 set as its main goal the development of the POWER ecosystem in the server market, initially embracing under this initiative the POWER8 processor family. In 2015 we presented performance and power consumption benchmarks of uni-socket platforms that proved the existence of a significant gap between x86 and other competitors (A look beyond x86: OpenPOWER8 & AArch64, HEPiX Spring 2015). The ecosystem has grown both in terms of availability of hardware platforms and software support. I will present new performance and power consumption results covering recent dual-socket ARMv8 and POWER8 platforms.

**Primary authors:** BROSA IARTZA, Aritz (Universidad de Oviedo (ES)); GUERRI, Marco (CERN)

**Presenter:** GUERRI, Marco (CERN)

**Session Classification:** Computing and Batch Services

**Track Classification:** Computing & Batch Services

Contribution ID: 9

Type: **not specified**

## Nikhef Site Report

*Tuesday, October 18, 2016 9:45 AM (15 minutes)*

Update from Nikhef

**Primary author:** KUIPERS, Paul (Nikhef)

**Presenter:** KUIPERS, Paul (Nikhef)

**Session Classification:** Site Report

**Track Classification:** Site Reports

Contribution ID: **10**Type: **not specified**

## Tokyo Tier-2 Site Report

*Monday, October 17, 2016 2:00 PM (15 minutes)*

The Tokyo Tier-2 site, which is located in International Center for Elementary Particle Physics (ICEPP) at the University of Tokyo, is providing resources for the ATLAS experiment in WLCG. In December 2015, almost all hardware devices were replaced as the 4th system. Operation experiences with the new system and a migration plan from CREAM-CE + Troque/Maui to ARC-CE + HTCondor will be reported.

**Primary author:** KISHIMOTO, Tomoe (University of Tokyo (JP))

**Presenter:** KISHIMOTO, Tomoe (University of Tokyo (JP))

**Session Classification:** Site Report

**Track Classification:** Site Reports

Contribution ID: 11

Type: **not specified**

## Deep dive into Spectrum Scale (fomerly known as GPFS)

*Tuesday, October 18, 2016 5:00 PM (25 minutes)*

Intent of this presentation is to give current (or potential) users of Spectrum Scale a deep dive into various key components and functions of the Product and its usage in High Performance Computing. i will share Performance data for problematic filesystem workloads like shared directory or file access as well as demonstrate some new capabilities that have been added into the 4.2.1 release. i will further explain some i/o optimization technologies like LROC and HAWC that allow the use of FLASH technologies of various sorts to accelerate workloads. if time permits i can show some of the advanced performance and problem determination capabilities that were recently added to the product as well, including a live realtime performance demo.

**Primary author:** OEHME, Sven

**Presenter:** OEHME, Sven

**Session Classification:** Storage and Filesystems

**Track Classification:** Storage & Filesystems

Contribution ID: 12

Type: **not specified**

## University of Wisconsin-Madison CMS T2 site report

*Monday, October 17, 2016 11:40 AM (15 minutes)*

As a major WLCG/OSG T2 site, the University of Wisconsin-Madison CMS T2 has consistently been delivering highly reliable and productive services towards large scale CMS MC production/processing, data storage, and physics analysis for last 10 years. The site utilises high throughput computing (HTCondor), highly available storage system (Hadoop), scalable distributed software systems (CVMFS), and provides efficient data access using xrootd/AAA. The site fully supports IPv6 networking and is a member of the LHCCONE community with 100Gb WAN connectivity. An update on the activities and developments at the T2 facility over the last year (since the BNL meeting) will be presented.

**Primary author:** MOHAPATRA, Ajit (University of Wisconsin-Madison (US))

**Presenter:** MOHAPATRA, Ajit (University of Wisconsin-Madison (US))

**Session Classification:** Site Report

**Track Classification:** Site Reports

Contribution ID: 13

Type: **not specified**

## Upgrade of network connection between KEK and SINET

*Monday, October 17, 2016 4:15 PM (25 minutes)*

Since last Apr.1, SINET that is NREN for universities in Japan has started the operation of 5th generation infrastructure, SINET5. It accepts 100Gbps connection to the backbone from each institutes, and newly provides the direct path from Japan to Europe. KEK is connected to SINET by 120Gbps bandwidth in total and mostly the bandwidth will be used by the mass data transmission via LHCONE. We will report how we upgrade and change the monitoring scheme to keep the security level.

**Primary author:** SUZUKI, Soh

**Presenter:** SUZUKI, Soh

**Session Classification:** Security & Networking

**Track Classification:** Security & Networking

Contribution ID: 14

Type: **not specified**

## Experience of Development and Deployment of a Large-Scale Ceph-Based Data Storage System at RAL

*Wednesday, October 19, 2016 2:25 PM (25 minutes)*

A new data storage system, Echo, has been developed as a replacement for CASTOR disk-only storage of LHC data at the RAL Tier-1 for the past two years. This presentation will share the RAL experience of developing and deploying a new, ceph-based storage service at the 13 PB scale to the standard required for production use.

This is the first new service that we have developed at this scale for some time and ceph is a very different technology from our existing storage solution. This presentation will explore the changes required to accommodate such a service: from the location of servers in the data centre; development of the network topology and the effect this has on data placement; the design and construction of a system that is more manageable, maintainable and upgradable by a system administrator; the adaptation of existing software in order to support LHC VO workflows and the implementation of new software to support industry standard protocols for both LHC VOs and other user communities. I will also discuss the changes brought by the deployment of a new OS major version and the change from sysVinit to systemd for process management, the changes to monitoring and alerting required to support the continuous operation of the service and the risks and impacts of transitioning to this technology.

**Primary author:** CANNING, Bruno (RAL)

**Presenter:** CANNING, Bruno (RAL)

**Session Classification:** Storage and Filesystems

**Track Classification:** Storage & Filesystems

Contribution ID: 15

Type: **not specified**

## Big Data: Genomics vs. Physics

*Wednesday, October 19, 2016 9:25 AM (25 minutes)*

Big data is typically characterized by only a few features, such as Volume, Velocity and Variety. This is a simplification that overlooks many factors that affect the way data is used and managed, factors that can have a profound effect on the computing systems needed to serve different communities.

I compare the computing and data-management needs of the genomics domain with those of big physics experiments, highlight the differences between them and discuss the implications of those differences.

**Primary author:** WILDISH, Tony (Lawrence Berkeley National Laboratory)

**Presenter:** WILDISH, Tony (Lawrence Berkeley National Laboratory)

**Session Classification:** Computing and Batch Services

**Track Classification:** Miscellaneous



Contribution ID: 16

Type: **not specified**

## ITER siter eport

*Tuesday, October 18, 2016 11:30 AM (15 minutes)*

Critical to the success of ITER reaching its scientific goal ( $Q \geq 10$ ) is a data system that supports the broad range of diagnostics, data analysis, and computational simulations required for this scientific mission. Such a data system, termed ITERDB in this document, will be the centralized data access point and data archival mechanism for all of ITER's scientific data. ITERDB will provide a unified interface for accessing all types of ITER scientific data regardless of the consumer (e.g., scientist, engineer, plant operations) including interfaces for data management, archiving system administration, and health monitoring capabilities.

Due to the INB nature of ITER, there are two parts – one located in POZ (Plant Operation Zone) to collect experimental data and another one located in XPOZ (outside Plant Operation Zone) to allow offline analysis execution and storage. In this paper, we will focus on ITERDB-POZ part, the other part being still under-designed.

ITER is the international project consisting of seven Das (Domestic Agencies). Its procurement makes it quite challenging. To smooth integration, we developed the CODAC Core system which is a mini-platform based on RHEL and EPICS which simulates the functional CODAC behaviour. Since its first version (2010), it has been increased with new features and new APIs. ITER consists of roughly 200 systems (roughly millions of variables). In this paper, we will focus on the Data Acquisition Network (DAN). Many systems will stream data over DAN at various rates from a few hundred kB/sec to 50GB/sec). We describe in this document the various components involved in the data acquisition and a data storage chain.

**Primary author:** ABADIE, lana (ITER)

**Presenter:** ABADIE, lana (ITER)

**Session Classification:** Site Report

**Track Classification:** Site Reports

Contribution ID: 17

Type: **not specified**

## ZFS on Linux

*Wednesday, October 19, 2016 5:00 PM (25 minutes)*

ZFS is a combination of file system, logical volume manager, and software raid system developed by SUN Microsystems for the Solaris OS. ZFS simplifies the administration of disk storage and on Solaris it has been well regarded for its high performance, reliability, and stability for many years. It is used successfully for enterprise storage administration around the globe, but so far on such systems ZFS was mainly used to provide storage, like for users home directories, through NFS and similar network related protocols. <br>

Within GridPP, ZFS was also used before for the management of user home directories through NFS. These systems were based on Solaris or similar systems like the ones provided by Nexenta. However, most of the Grid Middleware run on Linux systems and not on Solaris and therefore ZFS wasn't used so far for Grid storage management or in general for Grid middleware servers.<br>

Since ZFS is available in a stable version on Linux now, here I will present our experience made with ZFS on Linux since we started to updated all GridPP storage (about 1PB) at our site at the end of last year to be managed by ZFS using the current Linux version of it. Since with larger growing disk capacity raid6 rebuild times get soon too large to be feasible, ZFS built in raid functionality was tested as an alternative to hardware raid systems and the results will be presented. I'll also report on other ZFS specific properties like compression,nfs sharing, and snapshots and how it is working in the Linux port.

ZFS on Linux could be an efficient and cost effective alternative to hardware raid and Solaris based systems, which has characteristics no other file system can provide and which can provide real data safety and reliability.

**Primary author:** EBERT, Marcus (University of Edinburgh (GB))

**Presenter:** EBERT, Marcus (University of Edinburgh (GB))

**Session Classification:** Storage and Filesystems

**Track Classification:** Storage & Filesystems

Contribution ID: 18

Type: **not specified**

## The role of dedicated computing centers in the age of cloud computing

*Thursday, October 20, 2016 11:35 AM (25 minutes)*

BNL anticipates significant growth in scientific programs with large computing and data storage needs in the near future and has recently re-organized support for scientific computing to meet these needs. A key component is the enhanced role of the RHIC-ATLAS Computing Facility (RACF) in support of HTC and HPC at BNL.

This presentation discusses the evolving role of the RACF at BNL, in light of its growing portfolio of responsibilities and its increasing integration with cloud (academic and for-profit) computing activities. We also discuss BNL's plan to build a new computing center to support the new responsibilities of the RACF and present a summary of the cost benefit analysis done, including the types of computing activities that benefit most from a local data center vs. cloud computing. This analysis is partly based on an updated cost comparison of Amazon EC2 computing services and the RACF, which was originally conducted in 2012.

**Primary author:** WONG, Tony (Brookhaven National Laboratory)

**Co-author:** Dr LANCON, Eric (Brookhaven National Laboratory)

**Presenter:** WONG, Tony (Brookhaven National Laboratory)

**Session Classification:** IT Facilities and Business Continuity

**Track Classification:** IT Facilities & Business Continuity

Contribution ID: **19**

Type: **not specified**

## **Welcome To NERSC/LBNL**

*Monday, October 17, 2016 9:10 AM (15 minutes)*

Sudip Dosanjh, NERSC

**Session Classification:** Miscellaneous

Contribution ID: 20

Type: **not specified**

## **Logistics & Safety Announcement**

*Monday, October 17, 2016 9:00 AM (10 minutes)*

**Presenters:** MEINHARD, Helge (CERN); WONG, Tony (Brookhaven National Laboratory)

**Session Classification:** Miscellaneous

Contribution ID: 21

Type: **not specified**

## **Plans to Support Data-Intensive Computing on the NERSC 8 System**

*Monday, October 17, 2016 9:25 AM (30 minutes)*

**Session Classification:** Miscellaneous

Contribution ID: 22

Type: **not specified**

## Monitoring HTCondor with Clustered Graphite and Grafana

*Thursday, October 20, 2016 2:00 PM (25 minutes)*

Grafana is a popular tool for data analytics, and HTCondor generates large amounts of time-series data appropriate for the kinds of analysis Grafana provides. We use a Graphite cluster, which will be described in some detail, as a back-end for metric storage, and adapted some scripts from Fermilab for metric gathering. This work is in the context of the batch-monitoring working group.

**Primary author:** STRECKER-KELLOGG, William (Brookhaven National Lab)

**Presenter:** STRECKER-KELLOGG, William (Brookhaven National Lab)

**Session Classification:** Basic IT Services

**Track Classification:** Basic IT Services

Contribution ID: 23

Type: **not specified**

## Ceph Based Storage Systems at the RACF

*Wednesday, October 19, 2016 2:50 PM (25 minutes)*

We give a report on the status of Ceph based storage systems deployed at the RHIC & ATLAS Computing Facility (RACF) that are currently providing 1 PB of data storage capacity for the object store (with Amazon S3 compliant Rados Gateway front end), block storage (RBD), and shared file system (CephFS with dCache/GridFTP front-ends) layers of Ceph storage system. The hardware and software upgrades performed over the duration of the last year are reported, including the results of performance tuning for the Rados Gateway subsystem of the cluster in order to support the high concurrency (up to 24k simultaneous connections), high granularity (about 1-10 MB payloads per client session), and high bandwidth (up to 1 GB/s of aggregate bandwidth on the WAN) data transfers via Amazon S3 compatible API in order to match the growing requirements of the ATLAS Event Service. The results of boosting the performance of our Ceph clusters using the low latency PCIe NVMe SSD storage devices and the future plans for our Ceph based storage systems are also discussed.

**Primary author:** ZAYTSEV, Alexandr (Brookhaven National Laboratory (US))

**Presenter:** ZAYTSEV, Alexandr (Brookhaven National Laboratory (US))

**Session Classification:** Storage and Filesystems

**Track Classification:** Storage & Filesystems



Contribution ID: 24

Type: **not specified**

## NDGF Site Report

*Tuesday, October 18, 2016 10:45 AM (15 minutes)*

News and interesting events from NDGF and NeIC.

**Primary author:** WADENSTEIN, Erik Mattias (University of Umeå (SE))

**Presenter:** WADENSTEIN, Erik Mattias (University of Umeå (SE))

**Session Classification:** Site Report

**Track Classification:** Site Reports

Contribution ID: 25

Type: **not specified**

## TRIUMF Site Report

*Monday, October 17, 2016 10:25 AM (15 minutes)*

Updates on the status of the Canadian Tier-1 and other TRIUMF computing news will be presented.

**Primary author:** DEATRICH, Denice

**Presenter:** DEATRICH, Denice

**Session Classification:** Site Report

**Track Classification:** Site Reports

Contribution ID: 26

Type: **not specified**

## KEK Site Report

*Monday, October 17, 2016 12:10 PM (15 minutes)*

The new KEK Central Computer system started the service on September 1st, 2016 after the renewal of all hardware. In this talk, we would like to introduce the performance of the new system and improvement of network connectivity with LHCONE.

**Primary author:** NAKAMURA, Tomoaki (KEK)

**Co-author:** SUZUKI, Soh (KEK)

**Presenter:** NAKAMURA, Tomoaki (KEK)

**Session Classification:** Site Report

**Track Classification:** Site Reports

Contribution ID: 27

Type: **not specified**

## **RAL Site Report**

*Tuesday, October 18, 2016 9:30 AM (15 minutes)*

Latest news of activities at the RAL Tier1.

**Primary author:** BLY, Martin (STFC-RAL)

**Presenter:** BLY, Martin (STFC-RAL)

**Session Classification:** Site Report

**Track Classification:** Site Reports

Contribution ID: 28

Type: **not specified**

## EOS, DPM and FTS developments and plans

*Wednesday, October 19, 2016 4:35 PM (25 minutes)*

The CERN IT-ST Analytics and Development section is responsible for the development of Data Management solution for Disk Storage and Data Transfer, namely EOS, DPM and FTS.

The talk will describe some recent developments in those 3 software solutions

### EOS

The integration and evaluation of various technologies to do the transition from a single active in-memory namespace to a scale-out implementation distributed over many meta-data servers. The new architecture aims to separate the data from the application logic and user interface code, thus providing flexibility and scalability to the namespace component.

### DPM

The implementation of a new core daemon (DOME) based on the fast-CGI and RESTful technologies. This brings the opportunity of working in a totally SRM-free mode, the implementation of quotas, free/used space on directories, and the implementation of volatile pools that can pull files from external sources, which can be used to deploy data caches.

### FTS

The extension to better support data transfer workflows between Grid, Cloud and HPC systems. This includes FTS3 implementing protocol translations and performing efficient 3rd party transfers over HTTP. One of the core component ( Optimizer ) has been also rewritten to allow ranges of active transfers and better exploitation of the network resources.

**Primary author:** MANZI, Andrea (CERN)

**Presenter:** MANZI, Andrea (CERN)

**Session Classification:** Storage and Filesystems

**Track Classification:** Storage & Filesystems

Contribution ID: 29

Type: **not specified**

## Plans to support IPv6-only CPU on WLCG - an update from the HEPiX IPv6 Working Group

*Tuesday, October 18, 2016 2:00 PM (25 minutes)*

This report from the HEPiX IPv6 Working Group will present activities during the last 6-12 months. With IPv4 addresses running out and with some sites and Cloud providers now wishing to offer IPv6-only CPU, together with the fact that several WLCG sites are already successfully running production dual-stack storage services, we have a plan to support IPv6 CPU from April 2017 onwards. This plan will be presented.

**Primary author:** KELSEY, Dave (STFC - Rutherford Appleton Lab. (GB))

**Presenter:** KELSEY, Dave (STFC - Rutherford Appleton Lab. (GB))

**Session Classification:** Security & Networking

**Track Classification:** Security & Networking

Contribution ID: 30

Type: **not specified**

## User/group based access control for ElasticSearch + Kibana

*Thursday, October 20, 2016 3:45 PM (25 minutes)*

Kibana and ElasticSearch are used for monitoring in many places. However, by default they do not support authentication and authorization features. In the case of single Kibana and ElasticSearch services shared among many users, any user that can access Kibana can retrieve any information from ElasticSearch.

In this talk, we will report on our latest R&D experience in securing the Kibana and ElasticSearch services. We will describe a Kibana plugin that allows Kibana dashboards to be separated based on user/group. We will also describe the effect on performance from using SearchGuard, which is an ElasticSearch plugin enables user/group based access control.

**Primary author:** TAKASE, Wataru (High Energy Accelerator Research Organization (JP))

**Presenter:** TAKASE, Wataru (High Energy Accelerator Research Organization (JP))

**Session Classification:** Basic IT Services

**Track Classification:** Basic IT Services

Contribution ID: **31**

Type: **not specified**

## **KIT Site Report**

*Tuesday, October 18, 2016 11:00 AM (15 minutes)*

News about GridKa Tier-1 and other KIT IT projects and infrastructure.

**Primary author:** PETZOLD, Andreas (KIT - Karlsruhe Institute of Technology (DE))

**Presenter:** PETZOLD, Andreas (KIT - Karlsruhe Institute of Technology (DE))

**Session Classification:** Site Report

**Track Classification:** Site Reports



Contribution ID: 32

Type: **not specified**

## On-demand provisioning of HEP compute resources on cloud sites and shared HPC centers

*Friday, October 21, 2016 9:00 AM (25 minutes)*

This contribution reports on solutions, experiences and recent developments with the dynamic, on-demand provisioning of remote computing resources for analysis and simulation workflows. Local resources of a physics institute are extended by private and commercial cloud sites, ranging from the inclusion of desktop clusters over institute clusters to HPC centers.

We report on recent experience from incorporating a remote HPC center (NEMO Cluster, Freiburg University) and resources dynamically requested from a commercial provider (1&1 Internet SE), which have been seamlessly tied together with the ROCED scheduler 1 such that, from the user perspective, local and remote resources form a uniform, virtual computing cluster with a single point-of-entry. On a local test system, the usage of Docker containers has been explored and shown to be a viable and light-weight alternative to full virtualization solutions in trusted environments.

1 O. Oberst et al. Dynamic Extension of a Virtualized Cluster by using Cloud Resources, J. Phys.: Conference Ser. 396(3)032081, 2012

**Primary author:** GIFFELS, Manuel (KIT - Karlsruhe Institute of Technology (DE))

**Co-author:** PETZOLD, Andreas (KIT - Karlsruhe Institute of Technology (DE))

**Presenter:** PETZOLD, Andreas (KIT - Karlsruhe Institute of Technology (DE))

**Session Classification:** Grid, Cloud and Virtualisation

**Track Classification:** Grid, Cloud & Virtualisation

Contribution ID: **33**

Type: **not specified**

## **CERN Site Report**

*Tuesday, October 18, 2016 9:15 AM (15 minutes)*

News from CERN since the DESY workshop.

**Primary author:** BELLEMAN, Jerome (CERN)

**Presenter:** BELLEMAN, Jerome (CERN)

**Session Classification:** Site Report

**Track Classification:** Site Reports

Contribution ID: 34

Type: **not specified**

## Update on HNSciCloud project

*Friday, October 21, 2016 9:25 AM (25 minutes)*

Overview of what has happened in HNSciCloud over the last five months

**Primary author:** MEINHARD, Helge (CERN)

**Presenter:** MEINHARD, Helge (CERN)

**Session Classification:** Grid, Cloud and Virtualisation

**Track Classification:** Grid, Cloud & Virtualisation

Contribution ID: 35

Type: **not specified**

## GSI Site Report

*Tuesday, October 18, 2016 11:15 AM (15 minutes)*

During the last few months, HPC @ GSI has moved servers and services to the new data center *Green IT Cube*. This included moving the users from the old compute cluster to the new one with a *new* scheduler, and moving several Petabytes of data from the old to the new Lustre cluster.

**Primary author:** Dr ROTH, Thomas (GSI Darmstadt)

**Presenter:** Dr ROTH, Thomas (GSI Darmstadt)

**Session Classification:** Site Report

**Track Classification:** Site Reports

Contribution ID: 36

Type: **not specified**

## Chameleon: A Computer Science Testbed as Application of Cloud Computing

*Thursday, October 20, 2016 4:35 PM (25 minutes)*

Did you ever need hundreds of state-of-the-art nodes that you could use to scalably test new ideas on? Run experiments that are not disrupted by what other users are doing? A platform that allows you to reinstall the operating system, recompile the kernel, and gives you access to the console so that you can debug the system? A place where your research team can easily reproduce experiments carried out weeks ago? A lab where your students can work with different hardware configurations, from Infiniband to GPUs, either as part of a class or homework?

This talk will introduce Chameleon, a large-scale, deeply reconfigurable NSF-funded testbed for Computer Science research and education ([www.chameleoncloud.org](http://www.chameleoncloud.org)). The testbed consists of ~600 nodes (~14,000 cores) and a total of 5PB disk space hosted at the University of Chicago and TACC, and leverages 100 Gbps connection between the sites. The hardware consists primarily of homogenous nodes to support large-scale experiments – but subgroups of those nodes are equipped with additional capabilities including Infiniband networking, high-bandwidth I/O storages nodes, GPUs, and storage hierarchies with a mix of HDDs, SSDs, NVRAM, and high memory. To support Computer Science experiments, ranging from operating system and virtualization to security research, Chameleon provides a configuration system giving users exclusive access to bare metal nodes on an “as if it were in your lab basis”, i.e., full control of the software stack including root privileges, kernel customization, and console access. In addition, to facilitate educational and application exploratory projects Chameleon also provides a KVM cloud.

I will describe user facing Chameleon capabilities, describe some of the project that the testbed supported in the past, and explain how the testbed was built and will continue to develop.

**Primary author:** KEAHEY, Kate (Argonne National Laboratory)

**Presenter:** KEAHEY, Kate (Argonne National Laboratory)

**Session Classification:** Grid, Cloud and Virtualisation

**Track Classification:** Grid, Cloud & Virtualisation

Contribution ID: 37

Type: **not specified**

## CEPHFS: a new generation storage platform for Australian high energy physics

*Wednesday, October 19, 2016 2:00 PM (25 minutes)*

In this paper we present a CEPHFS use case implementation at the Center of Excellence for Particle Physics at the TeraScale (CoEPP). CoEPP operates the Australia Tier-2 for ATLAS and joins experimental and theoretical researchers from the Universities of Adelaide, Melbourne, Sydney and Monash. CEPHFS is used to provide a unique object storage system, deployed on commodity hardware and without single points of failure, used by Australian HEP researchers in the different CoEPP locations to store, process and share data, independent of their geographical location. CEPHFS is also working in combination with a SRM and XROOTD implementation, integrated in ATLAS Data Management operations, and used by HEP researchers for XROOTD or/and POSIX-like access to ATLAS Tier-2 user areas. We will provide details on the architecture, its implementation and tuning, and report performance I/O metrics as experienced by different clients deployed over WAN. We will also explain our plan to collaborate with Red Hat Inc. on extending our current model so that the metadata cluster distribution becomes multi-site aware, such that regions of the namespace can be tied or migrated to metadata servers in different data centers.

**Primary author:** BORGES, Goncalo (University of Sydney (AU))

**Presenter:** BORGES, Goncalo (University of Sydney (AU))

**Session Classification:** Storage and Filesystems

**Track Classification:** Storage & Filesystems

Contribution ID: 38

Type: **not specified**

## Pre-Studies for Wi-Fi service enhancement at CERN

*Tuesday, October 18, 2016 2:50 PM (25 minutes)*

Over the last few years, the number of mobile devices connected to the CERN internal network has increased from a handful in 2006 to more than 10,000 in 2015. Wireless access is no longer a “nice to have” or just for conference and meeting rooms, now support for mobility is expected by most, if not all, of the CERN community. In this context, a full renewal of the CERN Wi-Fi network has been launched in order to provide a state-of-the-art Campus-wide Wi-Fi Infrastructure. Which technologies can provide an end-user experience comparable, for most applications, to a wired connection? Which solution can cover more than 200 office buildings, which represent a total surface of more than 400.000 m<sup>2</sup>, while keeping a single, simple, flexible and open management platform? The presentation will focus on the pre-studies which were done at CERN to review the full Wi-Fi infrastructure across the Campus. Moreover modern demands for Wi-Fi connectivity, as well as designing process of new CERN Wi-Fi network (RF planning, simulation, site survey) will be presented.

**Primary author:** SOSNOWSKI, Adam Wojciech (AGH University of Science and Technology (PL))

**Co-author:** DUCRET, Vincent (CERN)

**Presenter:** SOSNOWSKI, Adam Wojciech (AGH University of Science and Technology (PL))

**Session Classification:** Security & Networking

**Track Classification:** Security & Networking

Contribution ID: 39

Type: **not specified**

## BNL Site Report

*Monday, October 17, 2016 10:10 AM (15 minutes)*

The site report contains the latest news and updates on computing at BNL.

**Primary author:** STRECKER-KELLOGG, William (Brookhaven National Lab)

**Presenter:** STRECKER-KELLOGG, William (Brookhaven National Lab)

**Session Classification:** Site Report

**Track Classification:** Site Reports



Contribution ID: 40

Type: **not specified**

## Dynamical Provisioning of Cloud Computing Resources for Batch Processing

*Wednesday, October 19, 2016 11:10 AM (25 minutes)*

We aim to build a software service for provisioning cloud-based computing resources that can be used to augment users' existing, fixed resources and meet their batch job demands. This service must be designed to automate the delivery of compute resources (HTCondor execute nodes) to match user job demand in such a way that cloud-based resource utilization is high and, thus, cost per cpu-hour is low. In addition, since this provisioning service will acquire resources on behalf of its users, acting as a third-party buyer for them, it is also our fiduciary responsibility to ensure the system is stable or, at least, that stability can be maintained. In order to assess if stable resource utilization is possible, a dynamical systems approach is developed to provide a framework for understanding how the provisioning service will respond to user job demand. We will present our latest results on the project and give an overview of the development plan moving forward.

**Primary author:** Dr KANDES, Martin (Univ. of California San Diego (US))

**Presenter:** Dr KANDES, Martin (Univ. of California San Diego (US))

**Session Classification:** Computing and Batch Services

**Track Classification:** Computing & Batch Services

Contribution ID: 41

Type: **not specified**

## HEPiX Benchmarking Working Group - Status Report HEPiX Fall 2016

*Wednesday, October 19, 2016 9:00 AM (25 minutes)*

The HEPiX Benchmarking Working Group has been relaunched in spring 2016. First tasks are:

- Development and proposal of a fast benchmark to estimate the performance of the provided job slot (in traditional batch farms) or VM instance (in cloud environments)
- Preliminary work for a successor of the HS06 benchmark

This talk provides a status report of the work done so far.

**Primary author:** ALEF, Manfred (Karlsruhe Institute of Technology (KIT))

**Presenter:** ALEF, Manfred (Karlsruhe Institute of Technology (KIT))

**Session Classification:** Computing and Batch Services

**Track Classification:** Computing & Batch Services

Contribution ID: 42

Type: **not specified**

## T2\_FI\_HIP Site Report

*Tuesday, October 18, 2016 11:45 AM (15 minutes)*

- hardware renewal
- dCache and OS upgrade
- ansible

**Primary author:** GULDMYR, Johan Henrik (Helsinki Institute of Physics (FI))

**Presenter:** GULDMYR, Johan Henrik (Helsinki Institute of Physics (FI))

**Session Classification:** Site Report

**Track Classification:** Site Reports

Contribution ID: 43

Type: **not specified**

## **JLab Scientific and High Performance Computing**

*Monday, October 17, 2016 9:55 AM (15 minutes)*

JLab high performance and experimental physics computing environment updates since the spring 2016 meeting, including recent hardware installs of KNL and Broadwell compute clusters, Super-micro storage; our Lustre Intel upgrade status; 12GeV computing updates; and Data Center modernization progress.

**Primary author:** PHILPOTT, Sandy

**Presenter:** PHILPOTT, Sandy

**Session Classification:** Site Report

**Track Classification:** Site Reports

Contribution ID: 45

Type: **not specified**

## Introduction of load balancers at a Tier-1 site

*Thursday, October 20, 2016 2:25 PM (25 minutes)*

Historically at the RAL Tier-1 we have always directly exposed public-facing services to the internet via static DNS entries. This is far from ideal as it means that users will experience connection failures during server maintenance (both planned and unplanned) and any changes to the servers behind a particular service require DNS changes. Since April we have been using in production HAProxy and Keepalived to facilitate a highly-available load balancer in front of FTS3 in order to avoid the issues resulting from the use of DNS aliases. We are also making extensive use of HAProxy and Keepalived for our OpenStack cloud which is under development. Here we will describe our setup, experience with load balancers for FTS3 and OpenStack as well as our progress and plans for other services.

**Primary authors:** LAHIFF, Andrew David (STFC - Rutherford Appleton Lab. (GB)); COLLIER, Ian (STFC - Rutherford Appleton Lab. (GB))

**Presenter:** COLLIER, Ian (STFC - Rutherford Appleton Lab. (GB))

**Session Classification:** Basic IT Services

**Track Classification:** Basic IT Services

Contribution ID: 46

Type: **not specified**

## Status of IHEP Site

*Monday, October 17, 2016 11:55 AM (15 minutes)*

This talk will give a brief introduction to the status of computing center IHEP, CAS, including local cluster, Grid Tier2 site for Atlas and CMS, file and storage system, cloud infrastructure, planned HPC system, Internet and domestic network.

**Primary author:** CHENG, Yaodong (IHEP)

**Presenter:** CHENG, Yaodong (IHEP)

**Session Classification:** Site Report

**Track Classification:** Site Reports

Contribution ID: 47

Type: **not specified**

## An e-mail quarantine with open source software

*Monday, October 17, 2016 2:55 PM (25 minutes)*

Filtering e-mails for security reasons is a common procedure. At DESY e-mails with suspicious content are quarantined, users are notified and may request delivery of those e-mails. DESY is in the process of shifting from a commercial product to a quarantine solution made of open source and self-made software. This solution will be presented in context with DESY's e-mail infrastructure.

**Primary author:** Mr JAHNKE-ZUMBUSCH, Dirk (DESY)

**Presenter:** Mr JAHNKE-ZUMBUSCH, Dirk (DESY)

**Session Classification:** End-User IT Services & Operating Systems

**Track Classification:** End-User IT Services & Operating Systems

Contribution ID: 48

Type: **not specified**

## AGLT2 Site Update

*Monday, October 17, 2016 11:10 AM (15 minutes)*

We will present an update on our site since the Spring 2016 report, covering our changes in software, tools and operations.

We will also report on our recent significant hardware purchases during summer 2016 and the impact it is having on our site.

We conclude with a summary of what has worked and what problems we encountered and indicate directions for future work.

**Primary author:** MC KEE, Shawn (University of Michigan (US))

**Presenter:** MC KEE, Shawn (University of Michigan (US))

**Session Classification:** Site Report

**Track Classification:** Site Reports



Contribution ID: 49

Type: **not specified**

## University of Nebraska CMS T2 Site Report

*Monday, October 17, 2016 11:25 AM (15 minutes)*

Updates from T2\_US\_Nebraska covering our experiences operating CentOS 7 + Docker/SL6 worker nodes, banishing SRM in favor of LVS balanced GridFTP, and some attempts at smashing Open-Flow + GridFTP + ONOS together to live the SDN dream.

**Primary author:** ATTEBURY, Garhan (University of Nebraska-Lincoln (US))

**Presenter:** ATTEBURY, Garhan (University of Nebraska-Lincoln (US))

**Session Classification:** Site Report

**Track Classification:** Site Reports

Contribution ID: 50

Type: **not specified**

## **JLab's SciPhi-XVI Knights Landing Cluster**

*Wednesday, October 19, 2016 9:50 AM (25 minutes)*

Jefferson Lab recently installed a 200 node Knights Landing cluster, becoming an Intel® Parallel Computing Center. This talk will give an overview of the cluster installation and configuration, including its Omni-Path fabric, benchmarking, and integration with Lustre and NFS over Infiniband.

**Primary author:** PHILPOTT, Sandy

**Presenter:** PHILPOTT, Sandy

**Session Classification:** Computing and Batch Services

**Track Classification:** Computing & Batch Services

Contribution ID: 51

Type: **not specified**

## Resilient dCache and other news

*Wednesday, October 19, 2016 3:15 PM (25 minutes)*

New developments in dCache, in particular resilient features of redundant headnode services where we can now do automatic failover and rolling upgrades with low to none service impact.

Some other news too, on recent development in other areas like ceph support.

**Primary author:** WADENSTEIN, Erik Mattias (University of Umeå (SE))

**Co-author:** MILLAR, Paul (DESY)

**Presenter:** WADENSTEIN, Erik Mattias (University of Umeå (SE))

**Session Classification:** Storage and Filesystems

**Track Classification:** Storage & Filesystems

Contribution ID: 52

Type: **not specified**

## Platform Providing Network Awareness to ATLAS and Beyond

*Monday, October 17, 2016 3:50 PM (25 minutes)*

With the change of the ATLAS computing model from hierarchical to dynamic, processing tasks are dispatched to sites based not only on availability of resources but also network conditions along the path between compute and storage, which may be topologically and/or geographically distant. We describe a system developed to collect, store, analyze and provide timely access to the network conditions for ATLAS sites, which is also generalized for broader use. We describe the data we collect from four different sources giving orthogonal views of network performance and utilization. The pre-existing ATLAS Distributed Computing Analytics platform is used for data transport and storage. The platform provides interactive monitoring dashboards, and serves as a backend to an alarm and alert system which we have developed for site operators. A co-located Jupyter service is used to perform in-depth interactive data analysis, train different Machine Learning algorithms and test models on historical data. We discuss how the derived knowledge gets used by ATLAS for network anomaly detection, job scheduling and data brokering.

**Primary author:** VUKOTIC, Ilija (University of Chicago (US))

**Presenter:** VUKOTIC, Ilija (University of Chicago (US))

**Session Classification:** Security & Networking

**Track Classification:** Security & Networking

Contribution ID: 53

Type: **not specified**

## Cloud Services – Network realities

*Tuesday, October 18, 2016 4:10 PM (25 minutes)*

HEP use of cloud services has brought to light various network issues that hamper the full integration of such services with WLCG resources. In this presentation we comment on the issues that have been encountered and present the ongoing actions of the international network community to facilitate the integration of cloud services into the research computing environment.

**Primary author:** CASS, Tony (CERN)**Presenter:** CASS, Tony (CERN)**Session Classification:** Security & Networking**Track Classification:** Security & Networking

Contribution ID: 54

Type: **not specified**

## Extending the farm to external sites: the INFN Tier-1 experience

*Thursday, October 20, 2016 5:00 PM (25 minutes)*

The Tier-1 at CNAF is the main INFN computing facility offering computing and storage resources to more than 30 different scientific collaborations including the 4 experiments at the LHC. A huge increase in computing needs is foreseen in the next years mainly driven by the experiments at the LHC (especially starting with the run 3 from 2021) but also by other upcoming experiments such as CTA.

While we are considering the upgrade of the infrastructure of our data center, we are also evaluating the possibility of using CPU resources available in other data centers or even leased from commercial cloud providers.

Hence, at INFN Tier-1 we have pledged a small amount of computing resources (~2000 cores located at the Bari ReCaS) for the WLCG experiments for 2016 and we are testing the use of resources provided by a commercial cloud provider. While the Bari ReCaS data center is directly connected to the GARR network with the obvious advantage of a low latency and high bandwidth connection, in the case of the commercial provider we rely only on the General Purpose Network.

In this presentation we describe the setup phase and the first results of these installations, started in the last quarter of 2015, focusing on the issues that we had to deal with and discussing the measured results in terms of efficiency.

**Primary author:** DELL'AGNELLO, Luca (INFN-CNAF)

**Co-author:** CHIERICI, Andrea (INFN-CNAF)

**Presenter:** CHIERICI, Andrea (INFN-CNAF)

**Session Classification:** Grid, Cloud and Virtualisation

**Track Classification:** Grid, Cloud & Virtualisation

Contribution ID: 55

Type: **not specified**

## Australia-ATLAS Site report

*Monday, October 17, 2016 2:15 PM (15 minutes)*

Will provide updates on technical and managerial changes to Australia's only HEP grid computing site.

**Primary author:** BOLAND, Lucien Philip (University of Melbourne (AU))

**Co-author:** BORGES, Goncalo (University of Sydney (AU))

**Presenter:** BOLAND, Lucien Philip (University of Melbourne (AU))

**Session Classification:** Site Report

**Track Classification:** Site Reports

Contribution ID: 56

Type: **not specified**

## Renewal of Puppet for Australia-ATLAS

*Thursday, October 20, 2016 2:50 PM (25 minutes)*

Australia-ATLAS has been running Puppet for all infrastructure and Grid nodes since 2012. With the release of Puppet 4, and the move to Centos 7, we decided to rejig our Puppet configuration using what we've learnt in 4 years, and best practice methodologies. This talk will describe the problems we had with the old Puppet config, the decisions we made constructing the new system, and how the new system makes configuration management much easier.

**Primary author:** Mr CROSBY, Sean (University of Melbourne (AU))

**Co-authors:** BORGES, Goncalo (University of Sydney (AU)); BOLAND, Lucien Philip (University of Melbourne (AU))

**Presenter:** Mr CROSBY, Sean (University of Melbourne (AU))

**Session Classification:** Basic IT Services

**Track Classification:** Basic IT Services



Contribution ID: 57

Type: **not specified**

## SDN-enabled Intrusion Detection System

*Monday, October 17, 2016 4:40 PM (25 minutes)*

CERN networks are dealing with an ever-increasing volume of network traffic. The traffic leaving and entering CERN has to be precisely monitored and analysed in order to properly protect the networks from potential security breaches. To provide the required monitoring capabilities, the Computer Security team and the Networking team at CERN have joined efforts in designing and deploying a scalable Intrusion Detection System (IDS) setup. The setup features symmetrical load-balancing of monitored traffic across a pool of IDS servers with optional OpenFlow-based traffic shunting (offloading) and selective packet capturing capabilities. Having an experimental instance deployed, the solution is currently under testing with a promising perspective of putting it in production in the near future.

**Primary author:** KRAJEWSKI, Adam Lukasz (CERN)

**Presenter:** KRAJEWSKI, Adam Lukasz (CERN)

**Session Classification:** Security & Networking

**Track Classification:** Security & Networking

Contribution ID: 58

Type: **not specified**

## What's new in HTCondor? What is upcoming?

*Wednesday, October 19, 2016 11:35 AM (25 minutes)*

The goal of the HTCondor team is to to develop, implement, deploy, and evaluate mechanisms and policies that support High Throughput Computing (HTC) on large collections of distributively owned computing resources. Increasingly, the work performed by the HTCondor developers is being driven by its partnership with the High Energy Physics (HEP) community.

This talk will present recent changes and enhancements to HTCondor, including details on some of the enhancements created for the imminent HTCondor v8.6.0 release, changes created on behalf of the HEP community, and advancements on interactions with Docker and public cloud services. It will also discuss the upcoming HTCondor development roadmap, and seek to solicit feedback on the roadmap from HEPiX attendees.

**Primary author:** TANNENBAUM, Todd

**Presenter:** TANNENBAUM, Todd

**Session Classification:** Computing and Batch Services

**Track Classification:** Computing & Batch Services

Contribution ID: 59

Type: **not specified**

## Fermilab Site Report

*Monday, October 17, 2016 12:25 PM (15 minutes)*

News and updates from Fermilab.

**Primary author:** SCOTT, Rennie (Fermilab)

**Presenter:** SCOTT, Rennie (Fermilab)

**Session Classification:** Site Report

**Track Classification:** Site Reports

Contribution ID: **60**

Type: **not specified**

## Scientific Linux Status Update

*Monday, October 17, 2016 2:30 PM (25 minutes)*

Scientific Linux status and news.

**Primary author:** SCOTT, Rennie (Fermilab)

**Co-author:** RIEHECKY, Pat (Fermilab)

**Presenter:** SCOTT, Rennie (Fermilab)

**Session Classification:** End-User IT Services & Operating Systems

**Track Classification:** End-User IT Services & Operating Systems

Contribution ID: 61

Type: **not specified**

## Adopting Red Hat Satellite 6 for Lifecycle Management

*Thursday, October 20, 2016 4:10 PM (25 minutes)*

An overview of results and lessons learned from the Fermilab Scientific Linux and Architecture Management (SLAM) group's Satellite 6 Lifecycle Management Project. The SLAM team offers a portfolio of diverse system management service offerings with a small staff. Managing the risk of resource scarcity involves implementing tools and processes that will facilitate standardization, reduce complexity, and increase efficiency whenever possible. This short talk will give a brief overview of our experience and the results and the future of migrating to Satellite 6.1 as our new base for System Management.

**Primary author:** SCOTT, Rennie (Fermilab)

**Co-author:** RIEHECKY, Pat (Fermilab)

**Presenter:** SCOTT, Rennie (Fermilab)

**Session Classification:** Basic IT Services

**Track Classification:** Basic IT Services

Contribution ID: 64

Type: **not specified**

## Profiling data intensive workflows on Genepool and PDSF clusters at NERSC.

*Wednesday, October 19, 2016 12:00 PM (25 minutes)*

NERSC is well known for its user friendly, large-scale computing environment. Along with the large Cray systems (Edison and Cori), NERSC also supports data intensive workflows of the Joint Genome Institute, HEP and material science community via its Genepool, PDSF and Matgen clusters. These clusters are all provisioned from a single backend cluster, Mendel. This talk will briefly outline the workflows in Mendel and provide a comparative profile of its various applications. It will also summarize various user and system incidents over the last few years of its service. A deeper analysis of the bio-informatics workflow on the Genepool compute cluster, and a plan for testing workflows on a Mendel testbed with Cori-like environment will be discussed. Finally, a prospective plan for future evolution of Genepool part of Mendel will also be outlined.

**Primary author:** Dr THAKUR, Bhupender (NERSC, Lawrence Berkeley National Lab)

**Presenter:** Dr THAKUR, Bhupender (NERSC, Lawrence Berkeley National Lab)

**Session Classification:** Computing and Batch Services

**Track Classification:** Computing & Batch Services

Contribution ID: 65

Type: **not specified**

## AFS phaseout at CERN

*Thursday, October 20, 2016 9:30 AM (25 minutes)*

(Open)AFS has been used at CERN as general purpose filesystem for Linux homedirectories and project space for over 20 years. It has an excellent track record, but is showing its age. It is now slowly being phased out due to concerns on the project's long-term viability. The talk will briefly explain CERN's reasons for phasing out, give an overview of the process, introduce the migration targets for the various use cases (primarily EOS-FUSE), and highlight the challenges (and opportunities) of this migration.

**Primary author:** IVEN, Jan (CERN)**Presenter:** IVEN, Jan (CERN)**Session Classification:** Storage and Filesystems**Track Classification:** Storage & Filesystems

Contribution ID: 66

Type: **not specified**

## OSiRIS: One Year Update

*Wednesday, October 19, 2016 5:25 PM (25 minutes)*

The OSiRIS (Open Storage Research Infrastructure) project started in September 2015, funded under the NSF CC\*DNI DIBBs program (NSF grant #1541335). This program seeks solutions to the challenges many scientific disciplines are facing with the rapidly increasing size, variety and complexity of data they must work with. As the data grows, scientists are challenged to manage, share and analyze that data and become diverted from a focus on their scientific research to data-access and data-management concerns. Even more problematic is determining how to support many scientists sharing and accessing this ever increasing amount of data across multiple institutions.

We will describe the progress made during the OSiRIS project's first year. OSiRIS has fully deployed and benchmarked its initial multi-institutional Ceph deployment. To do this involved developing, deploying and configuring a number of tools to support consistent provisioning, monitoring and management of the distributed OSiRIS infrastructure. We will cover those details and discuss our initial science engagements and near-term plans for our hardware, Ceph, Authentication/Authorization and Software Defined Networking as well as the longer term plans for this 5-year project.

**Primary author:** MC KEE, Shawn (University of Michigan (US))

**Co-author:** MEEKHOF, Benjeman Jay (University of Michigan (US))

**Presenter:** MC KEE, Shawn (University of Michigan (US))

**Session Classification:** Storage and Filesystems

**Track Classification:** Storage & Filesystems



Contribution ID: 67

Type: **not specified**

## Irfu site report

*Tuesday, October 18, 2016 12:00 PM (15 minutes)*

- Windows10 migration
- network : IPV6
- infra : monitoring
- new H2020 call EOSF

**Primary author:** FERRY, Sophie

**Presenter:** FERRY, Sophie

**Session Classification:** Site Report

**Track Classification:** Site Reports

Contribution ID: 68

Type: **not specified**

## Further Adventures in Container Orchestration at RAL

*Friday, October 21, 2016 11:10 AM (25 minutes)*

We provide an update on our continued experiments with container orchestration at the RAL Tier 1.

**Primary authors:** LAHIFF, Andrew David (STFC - Rutherford Appleton Lab. (GB)); COLLIER, Ian (STFC - Rutherford Appleton Lab. (GB))

**Presenter:** COLLIER, Ian (STFC - Rutherford Appleton Lab. (GB))

**Session Classification:** Grid, Cloud and Virtualisation

**Track Classification:** Grid, Cloud & Virtualisation

Contribution ID: 69

Type: **not specified**

## Running HEP Workloads on the NERSC HPC Systems

*Friday, October 21, 2016 10:45 AM (25 minutes)*

Running HEP workloads on a Cray system can be challenging since these systems typically don't look very much like a standard Linux system. This presentation will describe several tools NERSC has deployed to enhance HEP and other data intensive computing: Shifter, a docker-like container technology developed at NERSC, the Burst Buffer, a super fast IO layer, and a software defined network that allows high speed connections to the outside world. We will give an overview of the software and hardware architecture, deployment, and performance of these services.

**Primary authors:** PAUL, D. (LBL); BARD, Deborah (LBL); JACOBSEN, Doug (LBL); BOTTS, James (LBNL); GERHARDT, Lisa (LBNL); CANNON, Shane (LBL); DECLERCK, T (LBL); QUAN, Tony (LBL); BHIMJI, Wahid (Lawrence Berkeley National Lab. (US))

**Presenter:** QUAN, Tony (LBL)

**Session Classification:** Grid, Cloud and Virtualisation

**Track Classification:** Grid, Cloud & Virtualisation

Contribution ID: 70

Type: **not specified**

## Update from Database Services

*Thursday, October 20, 2016 9:05 AM (25 minutes)*

With the terabytes of data stored in databases and Hadoop at CERN and great number of critical applications relying on them, the database service is evolving and the Hadoop service is expanding to adapt to changing needs and requirements of its users. The demand is high and the scope is broad. This presentation gives an overview of current state of databases services and new technologies approaching in Hadoop Service to make better use of latest hardware developments. Update to Database-On-Demand management model and technologies (MySQL, PostgreSQL) will also be provided.

**Primary author:** DZIEDZINIEWICZ-WOJCIK, Katarzyna Maria (CERN)

**Presenter:** DZIEDZINIEWICZ-WOJCIK, Katarzyna Maria (CERN)

**Session Classification:** Storage and Filesystems

**Track Classification:** Storage & Filesystems

Contribution ID: 71

Type: **not specified**

## Wigner Datacenter - Site report

*Tuesday, October 18, 2016 12:15 PM (15 minutes)*

We give an update on the infrastructure, Tier-0 hosting services, Cloud services and other recent developments at the Wigner Datacenter.

**Primary author:** Mr SZABO, Domokos (Wigner Datacenter)

**Co-authors:** BAGO, Balazs (Hungarian Academy of Sciences (HU)); SZABÓ, Domokos

**Presenter:** Mr SZABO, Domokos (Wigner Datacenter)

**Session Classification:** Site Report

**Track Classification:** Site Reports

Contribution ID: 72

Type: **not specified**

## GreenITCube - Status & Monitoring

*Thursday, October 20, 2016 12:00 PM (25 minutes)*

The GreenITCube is in production for half a year now. We want to present our experience so far, what we have learned about the system and give an outlook for the next couple of months.

As a second part of the talk, we want to give a detailed overview of the infrastructure monitoring. The focus will be on the different systems, we have in work and how we put all monitoring data together.

**Primary author:** Mr TRAUTMANN, Jan (GSI Darmstadt)

**Presenter:** Mr TRAUTMANN, Jan (GSI Darmstadt)

**Session Classification:** IT Facilities and Business Continuity

**Track Classification:** IT Facilities & Business Continuity

Contribution ID: 73

Type: **not specified**

## CSNS Computing Environment Based on OpenStack

*Friday, October 21, 2016 11:35 AM (25 minutes)*

OpenStack is an open source software for creating private and public clouds. It controls large pools of compute, storage, and networking resources throughout a datacenter, managed through a dashboard or via the OpenStack API. Hundreds of the world's largest brands rely on OpenStack to run their businesses every day, reducing costs and helping them move faster.

We are applying this computing mode to the China Spallation Neutron Source (CSNS) computing environment. So from the research and practice aspects, firstly, the application status of cloud computing science in High Energy Physics Experiments and the special requirements of CSNS are introduced in this paper. Secondly, our design and practice of cloud computing platform based on OpenStack are mainly demonstrated from the aspects of cloud computing system framework, some improvements to openstack network, Storage architecture and so on. Finally, some future prospects of CSNS cloud computing environment are discussed in the ending of this paper.

**Primary author:** LI, Yakang (ihep)

**Presenter:** LI, Yakang (ihep)

**Session Classification:** Grid, Cloud and Virtualisation

**Track Classification:** Grid, Cloud & Virtualisation

Contribution ID: 74

Type: **not specified**

## The future of AFS family file systems in research computing

*Thursday, October 20, 2016 9:55 AM (25 minutes)*

Since the introduction of Transarc AFS in 1991, the AFS family of file systems have played a role in research computing around the globe.

This talk will discuss the resurgence in development of the AFS family of file systems. A summary of recent development for several family members will be presented including:

- AuriStor File System suite of clients and servers
- kAFS, the Linux in-tree client and the associated AF\_RXRPC socket interface
- OpenAFS clients and servers

The talk will describe the potential uses of the /afs file namespace as a persistent storage solution for Containers.

Finally, the talk will discuss the Tennessee Open Research storage Cloud (TORC) proposal that was submitted to the U.S. National Science Foundation for funding as part of the Cyber Infrastructure initiative. If funded, TORC will provide a wide-area, high-performance and interoperable storage infrastructure designed for scalable, multi-level federation under cooperative management. TORC will combine the global, federated /afs file namespace and the multi-level security and privacy provided by the AuriStor File System with the high performance, scalability and reliability of L-Store and the Internet Backplane Protocol.

**Primary author:** Mr ALTMAN, Jeffrey

**Presenter:** Mr ALTMAN, Jeffrey

**Session Classification:** Storage and Filesystems

**Track Classification:** Storage & Filesystems



Contribution ID: 75

Type: **not specified**

## The advances in IHEP Cloud facility

*Friday, October 21, 2016 9:50 AM (25 minutes)*

In IHEP, more large scientific facilities requests more computing resources. Management of large scale resources requests efficient and flexible system architecture. Virtual computing through cloud technical is an approach. IHEPCloud is a private LaaS cloud which supports multi-users and multi-projects to achieve virtual computing. In this paper, we describe the infrastructure of virtual computing cluster in IHEP and discuss the work we done. We also show the performance testing for BES job. IHEPCloud has been online since Nov 2014 and works well. The performance penalty is also acceptable.

**Primary authors:** CUI, Tao (IHEP(Institute of High Energy Physics, CAS,China)); CHENG, Yaodong (IHEP)

**Presenter:** CUI, Tao (IHEP(Institute of High Energy Physics, CAS,China))

**Session Classification:** Grid, Cloud and Virtualisation

**Track Classification:** Grid, Cloud & Virtualisation

Contribution ID: 76

Type: **not specified**

## SDN Implementation in IHEP

*Monday, October 17, 2016 5:05 PM (25 minutes)*

High energy physics experiments produce huge amounts of raw data, while because of the sharing characteristics of the network resources, there is no guarantee of the available bandwidth for each experiment which may cause link competition problems. On the other side, with the development of cloud computing technologies, IHEP have established a cloud platform based on OpenStack which can ensure the flexibility of the computing and storage resources, and more and more computing applications have been moved to this platform, however, under the traditional network architecture, network capability become the bottleneck of restricting the flexible application of cloud computing.

This report introduces the SDN implementation in IHEP to solve the above problems, we built a dedicated and elastic network platform based on the data center SDN technologies and network virtualization technologies, meanwhile the SDN@WAN solution in IHEP will also be introduced. In the end, the test results and future works will be shared and analyzed.

**Primary author:** Mrs ZENG, SHAN (IHEP)

**Presenter:** Mrs ZENG, SHAN (IHEP)

**Session Classification:** Security & Networking

**Track Classification:** Security & Networking

Contribution ID: 77

Type: **not specified**

## Caltech Site Report

*Monday, October 17, 2016 5:45 PM (15 minutes)*

Caltech site report (USCMS Tier 2 site)

**Primary author:** HENDRICKS, Wayne (California Institute of Technology (US))

**Presenter:** HENDRICKS, Wayne (California Institute of Technology (US))

**Session Classification:** Site Report

**Track Classification:** Site Reports

Contribution ID: 78

Type: **not specified**

## SLAC Site Report

*Monday, October 17, 2016 5:30 PM (15 minutes)*

### **Update on SLAC Scientific Computing Service**

SLAC's Scientific Computing Services team provide long-term storage and midrange compute capability for multiple science projects across the lab. The team is also responsible for core enterprise (non-science) unix infrastructure. Sustainable hardware lifecycle is a key part of the central computing strategy. We continue to push the idea of business models for computing services as an alternative to one-time hardware investments. Seamless cloud bursting for high-throughput batch compute is under development using OpenStack and AWS with VPN.

**Presenter:** ADESANYA, Yemi

**Session Classification:** Site Report

**Track Classification:** Site Reports

Contribution ID: 79

Type: **not specified**

## ASGC Site Report

*Tuesday, October 18, 2016 9:00 AM (15 minutes)*

report on facility deployment, recent activities, collaborations and plans

**Primary authors:** YEN, Eric (Academia Sinica Grid Computing); LEE, Felix.hung-te (Academia Sinica (TW))

**Presenters:** YEN, Eric (Academia Sinica Grid Computing); LEE, Felix.hung-te (Academia Sinica (TW))

**Session Classification:** Site Report

**Track Classification:** Site Reports

Contribution ID: 80

Type: **not specified**

## Effective and non-intrusive security within NERSC's Open Science HPC environment

*Thursday, October 20, 2016 5:25 PM (25 minutes)*

Providing effective and non-intrusive security within NERSC's Open Science HPC environment introduces a number of challenges for both researchers and operational personnel. As what constitutes HPC expands in scope and complexity, the need for timely and accurate decision making about user activity remains unchanged. This growing complexity is balanced against a backdrop of routine user and application attacks, which remain surprisingly effective over time.

This presentation describes current efforts at NERSC to maintain system integrity without getting in the way of the science being done here. These efforts include network monitoring, 2 factor authentication as well as ssh and host based data analysis”

**Primary author:** SINGER, Abe (Lawrence Berkeley Lab)

**Presenter:** SINGER, Abe (Lawrence Berkeley Lab)

**Session Classification:** Security & Networking

**Track Classification:** Security & Networking