

HEPiX Fall 2015 Workshop

Monday 12 October 2015 - Friday 16 October 2015

Brookhaven National Laboratory



Book of Abstracts

Contents

INFN-T1 Site report	1
Proposal for a new Data Center at BNL	1
CERN CentOS 7 and community update	1
ITIL Service models in Detector DAQ Computing	2
The HP IT Data Center Transformation Journey	2
Energy Services Performance Contracting Abstract - HEPIX Fall 2015	3
Ceph object storage at RAL	3
PIC Site Report	4
Energy efficiency upgrades at PIC	4
Ceph Based Storage Systems at the RACF	5
NVM Express (NVMe): An Overview and Performance Study	5
Non-traditional workloads at the RACF	6
An initial evaluation of Docker at the RACF	6
CERN Site Report	7
CERN Cloud Status	7
Optimisations of the Compute Resources in the CERN Cloud Service	7
Update on Configuration Management at CERN	8
CvmFS deployment status and trends	8
Software collaboration tools as a stack of services	9
An Easy HTCondor Configuration to Support all Workloads (for some definition of all)	9
News from the HEPiX IPv6 Working Group	10
Future home directory at CERN	10
Accelerating High Performance Cluster Computing Through the Reduction of File System Latency	10

Document oriented database infrastructure for monitoring HEP data systems applications	11
Benchmarking commercial cloud resources	12
Status of DESY Batch Infrastructures	12
NDGF Site Report	12
Self service kiosk for Mac and mobiles	13
Scientific Data Storage at FNAL	13
The Open Science Grid: Physics Experiments to Campus Researchers	14
Simulating 5 Trillion Events on the Open Science Grid	14
Oxford and SouthGrid Site Status Report	15
First experiences with Mesos at RAL	15
PDSF Site Report and Relocation	16
Space usage monitoring for distributed heterogeneous data storage systems.	16
Computer Security update	17
Quality of Service in storage and the INDIGO-DataCloud project.	17
RAL Site Report	18
Fermilab Site Report	18
Scientific Linux Status Update	19
A comparison of performance between KVM and Docker instances in OpenStack	19
Using VPLS for VM mobility	20
Network infrastructure for the CERN Datacentre	20
NERSC Plans for the Computational Research and Theory Building	20
ELK at NERSC	21
University of Wisconsin Madison CMS T2 site report	21
Status on the IPv6 OSG Software stack tests	22
Growing our own systems administrators	22
CERN Monitoring Status Update	23
Future of Batch Processing at CERN	23
Asset management in CERN data centres	23
Improving IaaS resources to accommodate scientific applications	24
ATLAS Great Lakes Tier-2 (AGLT2) Update	25

Update on WLCG/OSG perfSONAR Infrastructure	25
Quattor Update	26
WLCG Network and Transfer Metrics WG After One Year	26
Running ATLAS, CMS, ALICE Workloads on the NERSC Cray XC30	27
HTCondor Recent Enhancement and Future Directions	27
HTCondor-CE: Managing the Grid with HTCondor	28
Jefferson Lab Scientific and High Performance Computing	28
Australia-ATLAS site report	28
BEIJING Site Report	29
Building a large scale Security Operations Centre	29
Using ownCloud to provide mobile access to existing storage	30
Automated performance testing framework	30
Supervised Machine Learning	30
SLAC Scientific Computing Services	31
Foreman - A full automation solution for a growing facility	31
Gitlab and its CI - an intelligent solution for hosting Git repos	32
Upgrade to UGE 8.2 : Positives effects at IN2P3-CC	32
Site report from KEK	33
Wigner Datacenter - Site report	33
Update on benchmarking	33
USATLAS dCache storage system at BNL	34
KIT Site Report	34
computing resources virtualization in IHEP	34
Host deployment and configuration technologies at SCC.	35
CSC.fi - Site Report	35
Monitoring with InfluxDB and Grafana	36
Welcome to BNL	36
BNL Site Report	36
LAL Site Report	37
Introduction to GDB	37

DESY Site Report	37
HEPIX Wrap-up	38

Site Reports / 0**INFN-T1 Site report**

Author: Andrea Chierici¹

¹ *INFN-CNAF*

Corresponding Author: chierici@cnaf.infn.it

Update on Italian Tier1 center status.

Length of presentation (max. 20 minutes):

15

IT Facilities and Business Continuity / 1**Proposal for a new Data Center at BNL**

Author: Imran Latif¹

Co-authors: David Free ²; John McCarthy ²; Shigeki Misawa ²; Tony Wong ¹

¹ *Brookhaven National Laboratory*

² *Brookhaven National Lab*

Corresponding Author: ilatif@bnl.gov

The methods and techniques of scientific research at Brookhaven National Laboratory are increasingly dependent upon the ability to acquire, analyze, and store vast quantities of data. The needs for data processing equipment and supporting infrastructure are anticipated to grow significantly in the next few years, soon exceeding the capacity of existing data center resources, currently located in Building 515.

The Core Facility Revitalization –Phase 1 (CFR-1) project anticipates the partial renovation and revitalization of Building 725 (NSLS-I) for use as a new state of the art computing facility which will house HPC equipment for RHIC-ATLAS, NSLS-II, CFN and CSC and also provide expansion space for future science programs. This presentation summarizes the technical discussions held so far and the status of the project.

Length of presentation (max. 20 minutes):

20

End-user Services and Operating Systems / 2**CERN CentOS 7 and community update**

Author: Thomas Oulevey¹

¹ *CERN*

Corresponding Author: arne.wiebalck@cern.ch

In this talk we will present a brief status update on CERN's work on CentOS 7, the uptake by the various IT services, and the interaction with the upstream CentOS community.

We will talk about the SIGs status, new projects and work done over the last months, presenting a list of our contributions and feedback about the experience.

Length of presentation (max. 20 minutes):

20

Basic IT Services / 3

ITIL Service models in Detector DAQ Computing

Author: Bonnie King¹

Co-authors: Pat Riehecky¹; Rennie Scott¹; connie sieh¹; kevin hill

¹ *Fermilab*

Corresponding Author: bonniek@fnal.gov

Fermilab has moved from the era of two large multi-decade experiments to hosting several smaller experiments with a shorter lifecycle. Improvements in microcontroller performance have moved computers closer to the experiment Data Acquisition systems where custom electronics have previously been used. There are also efforts to standardize DAQ software into reusable products in alignment with offline analysis tools.

Nearly all DAQ systems in Fermilab experiments contain Linux computers, with or without special hardware, to collect data, build events, generate online monitoring and ship data to offline. The challenge is to productionize DAQ computers (or test stands) that have often been set up by grad students or post docs and have not always been subject to rigorous best practices.

The Experiment Computing Facilities/Scientific Linux Architecture and Management group has had success in onboarding experiment DAQ computer systems to its Experiment Online service, which has reduced duplicate effort, improved resiliency, and introduced industry best practices. Legacy, running and future experiments have adopted our services, all of them presenting different challenges. Additionally, these systems require sensitivity to users' needs for elevated access and agility in development.

We discuss our successes and lessons learned in standing up the Experiment Online service, efficiencies gained from the ITIL service model, and upcoming innovations for Linux-based DAQ computers.

Length of presentation (max. 20 minutes):

20 minutes

IT Facilities and Business Continuity / 4

The HP IT Data Center Transformation Journey

Author: https://me.yahoo.com/a/tjwhXIwVwsfK.M1MXtI.xbOM3f9Sj5F3_FtFolDk Rotheroe^{None}

HP IT consolidated from 85 ancient data centers running old IT technology to six new mega data centers with modern IT running transformed applications in the late 2000's. That achievement resulted in (literally) billions saved through a meshing of significantly more efficient data centers, utilization of then-current IT technology, and application transformation. Since achieving that, HP IT has set and achieved higher goals several times, including deploying ultra-efficient modular data centers with containerized infrastructure, cloud and low-power solutions, and further application modernization. The author, the strategist and technology for HP IT data centers globally, will discuss how by executing a plan that encompasses the data center facility, modern IT equipment, and application modernization together, the total benefits can be further increased over the significant benefits from individually addressing each, using examples from HP IT's own transformation.

Length of presentation (max. 20 minutes):

20

IT Facilities and Business Continuity / 5

Energy Services Performance Contracting Abstract - HEPiX Fall 2015

Author: Michael Ross¹

¹ HP

Corresponding Author: michael.l.ross@hp.com

In today's Federal Information Technology world you are face with many challenges. To name a few Cyber-attacks, un-funded mandates to consolidate data centers, new IT acquisitions laws requiring de-duplication with agency level CIO oversight implemented by the Federal Information Technology Acquisition Reform Act (FITARA), aging equipment and data center infrastructure. Each of these challenges requires, in many cases a significant dollar investment agencies unfortunately do not have.

There is a potential answer that can provide the funding to meet these challenges. The Energy Savings Performance Contract administered by Department of Energy and the US Army Corp of Engineers allows agencies to payback the cost of traditional building infrastructure and now IT projects that results in energy savings. Simply put, as energy savings are realized these dollar savings are used to pay back the cost of the project. The holders of the ESPC or Energy Service Companies (ESCO) provide the up-front dollars and are paid back through these savings.

Using the ESPC is a win...win...win solution. It allows for the creation of jobs, no appropriations required, and once the project is paid off the savings are realized by the federal agency. We hope you attend HEPiX Fall 2015 to learn more about using the ESPC.

Length of presentation (max. 20 minutes):

20 Mins

Storage and Filesystems / 6

Ceph object storage at RAL

Author: George Vasilakakos¹

Co-authors: Alastair Dewhurst²; Bruno Canning³; Ignacy Debicki⁴; James Adams⁴; Shaun de Witt¹

¹ STFC

² STFC - Rutherford Appleton Lab. (GB)

³ RAL

⁴ STFC RAL

Corresponding Author: george.vasilakakos@stfc.ac.uk

RAL is currently developing storage services powered by a Ceph object storage system. We review the test results and experiences of the newly-installed 5 PB cluster at RAL, as well as our plans for it. Since the aim is to provide large scale storage for experimental data with minimal space overhead, we focus on testing a variety of erasure coding techniques and schemes. We look at functionality and performance testing of XrootD and GridFTP servers that have been adapted to use a Ceph backend, as well as, RADOS gateways providing an S3-compatible interface. We also look at the effects of erasure coding on the stability and resiliency of a highly distributed system at this scale.

Length of presentation (max. 20 minutes):

20

Site Reports / 7

PIC Site Report

Author: Jose Flix Molina¹

¹ *Centro de Investigaciones Energ. Medioambientales y Tecn. - (ES)*

Corresponding Author: jose.flix.molina@cern.ch

We will be revising the status of PIC by Fall 2015. News from Oxford meeting will be reported.

Length of presentation (max. 20 minutes):

15

IT Facilities and Business Continuity / 8

Energy efficiency upgrades at PIC

Authors: Adrià Herbera¹; Jordi Hernández¹; Jose Flix Molina²; Manuel Delfino Reznicek³; Vanessa Acin Portella³

¹ *Universitat Autònoma de Barcelona*

² *Centro de Investigaciones Energ. Medioambientales y Tecn. - (ES)*

³ *Universitat Autònoma de Barcelona (ES)*

Corresponding Author: jose.flix.molina@cern.ch

Energy consumption is an increasing concern for data centers. This contribution summarizes recent energy efficiency upgrades at the Port d'Informació Científica (PIC) in Barcelona, Spain which have considerably lowered energy consumption. The upgrades were particularly challenging, as they involved modifying the already existing machine room, which is shared by PIC with the general IT services of the local University (Universitat Autònoma de Barcelona, UAB), with all the services in

full operation, as well as the introduction of “free cooling” techniques in a location 20 km from the Mediterranean Sea.

Length of presentation (max. 20 minutes):

15

Storage and Filesystems / 9

Ceph Based Storage Systems at the RACF

Author: Alexandr Zaytsev¹

Co-authors: Christopher Hollowell²; Hironori Ito¹; Ofer Rind³; Tejas Rao⁴; Tony Wong²

¹ Brookhaven National Laboratory (US)

² Brookhaven National Laboratory

³ BROOKHAVEN NATIONAL LABORATORY

⁴ Brookhaven National Laboratory (BNL)

Corresponding Author: alezayt@bnl.gov

We give a report on the status of Ceph based storage systems deployed in 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 front end) layers of Ceph storage system. The hardware deployment and software upgrade procedures performed in the year of 2015, and the resulting improvements of our Ceph clusters performance and functionality are also summarized, and the usage statistics for the main clients of our Ceph clusters, including the ATLAS Event Service, is shown. The future prospects for the Ceph based storage systems in RACF are also discussed.

Length of presentation (max. 20 minutes):

15

Computing and Batch Services / 10

NVM Express (NVMe): An Overview and Performance Study

Authors: Alexandr Zaytsev¹; Christopher Hollowell²; Costin Caramarcu¹; Shawn Hoose²; Tejas Rao²

¹ Brookhaven National Laboratory (US)

² Brookhaven National Laboratory

Corresponding Author: hollowec@bnl.gov

NVMe (Non-Volatile Memory express) is a leading-edge SSD technology where drives are directly attached to the PCI-e bus. Typical SAS/SATA controllers are optimized for use with traditional rotating hard drives, and as such can increase latency, and reduce the bandwidth available to attached SSDs. Since NVMe drives bypass a SAS/SATA controller, they can help minimize/eliminate performance reductions often seen when using these controllers with solid state storage. NVMe may become a

key building block in future applications that require high-performance disk storage. This presentation describes the testbeds used to evaluate this technology in the context of a data center with a heterogeneous user community and a variety of applications, and summarizes performance and the pros and cons in the RACF environment.

Length of presentation (max. 20 minutes):

20

Computing and Batch Services / 11

Non-traditional workloads at the RACF

Authors: Alexandr Zaytsev¹; Christopher Hollowell²; Costin Caramarcu¹; Tejas Rao²; William Strecker-Kellogg³

¹ *Brookhaven National Laboratory (US)*

² *Brookhaven National Laboratory*

³ *Brookhaven National Lab*

Corresponding Authors: william.edward.strecker-kellogg@cern.ch, willsk@bnl.gov

The RACF is a key component in BNL's new Computational Science Initiative (CSI). One of CSI's goals is to leverage the RACF's expertise to shorten the time and effort needed to archive, process and analyze data from non-traditional fields at BNL. This presentation describes a concrete example of how the RACF has helped non-traditional workloads run in the RACF computing environment, and how this helps establish a roadmap for future collaborative efforts.

Length of presentation (max. 20 minutes):

20

Grid, Cloud and Virtualization / 12

An initial evaluation of Docker at the RACF

Authors: Alexandr Zaytsev¹; Christopher Hollowell²; Costin Caramarcu¹; Joseph Zuhusky²; Tejas Rao²; Tony Wong²; William Strecker-Kellogg³

¹ *Brookhaven National Laboratory (US)*

² *Brookhaven National Laboratory*

³ *Brookhaven National Lab*

Corresponding Author: hollowec@bnl.gov

Application containers have become a competitive alternative to virtualized servers. Containers allow applications to be written once, distributed across a heterogeneous environment (ie, cloud, remote data centers) and executed transparently on multiple platforms without the performance overhead commonly found on virtual systems. We present an initial evaluation of Docker, along with a description of the testbeds and the sample applications

used to simulate a distributed computing environment, and we examine the pros and cons of Docker containers in the RACF environment.

Length of presentation (max. 20 minutes):

20

Site Reports / 13

CERN Site Report

Author: Arne Wiebalck¹

¹ *CERN*

Corresponding Author: arne.wiebalck@cern.ch

News from CERN since the Oxford workshop.

Length of presentation (max. 20 minutes):

15

Grid, Cloud and Virtualization / 14

CERN Cloud Status

Author: Arne Wiebalck¹

¹ *CERN*

Corresponding Author: arne.wiebalck@cern.ch

This presentation will provide an update of the activities in the CERN Cloud service since the Oxford workshop.

Length of presentation (max. 20 minutes):

20

Grid, Cloud and Virtualization / 15

Optimisations of the Compute Resources in the CERN Cloud Service

Author: Arne Wiebalck¹

Co-authors: Sean Crosby²; Tim Bell¹; Ulrich Schwickerath¹

¹ *CERN*

² *University of Melbourne (AU)*

Corresponding Author: arne.wiebalck@cern.ch

This talk will summarise our activities related to the optimisation of the virtualised compute resources in our OpenStack-based infrastructure. In particular, we will discuss some of the issues we've encountered and the various optimisations we've applied to bring the virtual resources as close as possible to bare-metal performance.

Length of presentation (max. 20 minutes):

20

Basic IT Services / 16

Update on Configuration Management at CERN

Author: Alberto Rodriguez Peon¹

Co-authors: Akos Hencz²; Andrea Giardini³; Ben Jones²; Gavin McCance²; Nacho Barrientos Arias²; Steve Traylen²

¹ *Universidad de Oviedo (ES)*

² *CERN*

³ *Universita e INFN, Bologna (IT)*

Corresponding Author: alberto.rodriguez.peon@cern.ch

An update on CERN's Configuration Service will be presented. This presentation will review the current status of the infrastructure and describe some of the ongoing work and future plans, with a particular focus on automation and continuous integration. Recent effort to scale and accommodate a higher number of puppet clients will also be mentioned.

Length of presentation (max. 20 minutes):

20

Storage and Filesystems / 17

CvmFS deployment status and trends

Author: Alberto Rodriguez Peon¹

Co-author: Steve Traylen²

¹ *Universidad de Oviedo (ES)*

² *CERN*

Corresponding Author: alberto.rodriguez.peon@cern.ch

CvmFS is a network file system based on HTTP and optimised to deliver experiment software in a fast, scalable, and reliable way. This presentation will review the status of the stratum 0 deployment at CERN, mentioning some of the challenges faced during its migration to ZFS as the underlying file system.

Length of presentation (max. 20 minutes):

20

End-user Services and Operating Systems / 18

Software collaboration tools as a stack of services

Author: Borja Aparicio Cotarelo¹

Co-authors: Alexandre Lossent¹; Alvaro Gonzalez Alvarez¹; Carolina Lindqvist²; Marta Panuszewska³; Nils Hoimyr¹; Pete Jones¹; Terje Andersen¹

¹ CERN

² University of Helsinki (FI)

³ AGH University of Science and Technology (PL)

Corresponding Author: borja.aparicio.cotarelo@cern.ch

The Version Control and Issue Tracking services team at CERN, is facing a transition and consolidation of the services provided in order to fulfill the needs of the CERN community as well as maintain the current deployments that are used.

Software and services development is a key activity at CERN that currently is widely carried out using Agile methodologies. Code hosting and review, version control, documentation, issue tracking and continuous integration and delivery are common requirements that each working group applies in daily basis.

As platform and engineering services providers, we are deploying a stack of tools and procedures around GitLab EE for version control based on Git and code review, Atlassian JIRA for issue tracking, Jenkins for continuous integration and TWiki for documentation, that can integrate to each other and allow each working group to build up their own development environment.

On the other hand, important software projects at CERN, such as LHC accelerator controls, CMS on-line and ATLAS off-line software rely on the legacy SVN service. This poses a challenge for the team to keep both environments and frameworks maintained and up to date in parallel.

Length of presentation (max. 20 minutes):

20 minutes

Computing and Batch Services / 19

An Easy HTCondor Configuration to Support all Workloads (for some definition of all)

Author: William Strecker-Kellogg¹

¹ Brookhaven National Lab

Corresponding Authors: william.edward.strecker-kellogg@cern.ch, willsk@bnl.gov

Scheduling jobs with heterogeneous requirements to a heterogeneous pool of computers is a challenging task. HTCondor does a great job supporting such a general-purpose setup with features like Hierarchical Group

Quotas and Partitionable Slots. At BNL we have a model, configuration, and software to handle the administration of such a pool, and in this talk we will share our experience building and running this setup along with the software we've used to make it easy to administer.

Length of presentation (max. 20 minutes):

20

Security and Networking / 20

News from the HEPiX IPv6 Working Group

Author: Dave Kelsey¹

¹ STFC - Rutherford Appleton Lab. (GB)

Corresponding Author: david.kelsey@stfc.ac.uk

This talk will present a status update from the IPv6 working group, including recent testing and the deployment of (some) dual-stack services and monitoring in WLCG.

Length of presentation (max. 20 minutes):

20 minutes

Storage and Filesystems / 21

Future home directory at CERN

Author: Alberto Pace¹

¹ CERN

Corresponding Author: alberto.pace@cern.ch

Several discussions are ongoing at CERN concerning the future of AFS and a possible replacement using the existing CERNBOX service based on OwnCloud and the CERN disk storage solution developed for the LHC Computing (EOS).

The talk will mainly review the future plans for the CERNBOX Services and the various use cases that have been identified.

Length of presentation (max. 20 minutes):

20

Storage and Filesystems / 23

Accelerating High Performance Cluster Computing Through the Reduction of File System Latency

Author: David Fellingner¹

¹ *DataDirect Networks, Inc*

Corresponding Author: dave.fellinger@gmail.com

The acceleration of high performance computing applications in large clusters has primarily been achieved with a focus on the cluster itself. Lower latency interconnects, more efficient message passing structures, higher performance processors, and general purpose graphics processing units have been incorporated in recent cluster designs. There has also been a great deal of study regarding processing techniques such as symmetric multi-processing versus efficient message passing to accomplish true parallel processing.

There has been, however, only incremental changes in parallel file system technology. Clusters perform input/output operations through gateway servers and a file creation infers locking operations in all parallel file systems. In fact, a file creation is a serial process which locks and assigns V-nodes, I-nodes and extent lists through one server to complete the operation.

For years, web users have explored parallel methods of moving data to get around network connection limitations. Applications such as Napster and Bit Torrent have used the technology of a Distributed Hash Table to effectively allow true parallel file operations where “pieces” can be placed to, or gathered from, a number of service nodes arranged in a redundant fashion.

This paper will explore the use of a Distributed Hash Table technology to service the data needs of a large scale cluster allowing the same parallelism in data mobility as is assumed in processing. This new paradigm will displace the concept of the gateway server and will allow data intensive operations in a “non-blocking” construct.

Length of presentation (max. 20 minutes):

20

Basic IT Services / 24

Document oriented database infrastructure for monitoring HEP data systems applications

Author: Carlos Fernando Gamboa¹

¹ *Brookhaven National Laboratory (US)*

Corresponding Author: cgamboa@bnl.gov

Within the High Energy Physics (HEP) software infrastructure a diverse Data storage and Distribution software technologies are used. Despite their heterogeneity they all provide a capability to trace application event data to be able to troubleshoot problems related to the software stack or software usage. The subsequent data is written once and stored in the majority of cases in a file, commonly known as log file, for posterior analysis. An open source infrastructure to collect, store in a document oriented type database, and visualize data processed from the application log files is explored in this presentation. Experience in deploying the ELK (Elasticsearch, Logstash and Kibana) framework applied to a data storage technology is summarized.

Length of presentation (max. 20 minutes):

20

Grid, Cloud and Virtualization / 25**Benchmarking commercial cloud resources****Author:** Domenico Giordano¹**Co-authors:** Alessandro Di Girolamo¹; Cristovao Cordeiro¹; Laurence Field¹; Luis Villazon Esteban²¹ *CERN*² *Universidad de Oviedo (ES)***Corresponding Author:** domenico.giordano@cern.ch

Performance measurements and monitoring are essential for the efficient use of computing resources as they allow selecting and validating the most effective resources for a given processing workflow. In a commercial cloud environment an exhaustive resource profiling has additional benefits due to the intrinsic variability of a virtualised environment. In this context resource profiling via initial benchmarking quickly allows to identify issues and mitigate them. Ultimately it provides additional information to compare the presumed performance of invoiced resources and the actual delivered performance as perceived at the client level.

In this report we will discuss the experience acquired in benchmarking commercial cloud resources during production activities such as the recent ATLAS Monte Carlo production in Helix Nebula cloud providers. The workflow put in place to collect and analyse performance metrics will be described. Results of the comparison study among commonly used benchmark metrics will be also reported. Those common benchmarks span from generic open-source benchmarks (encoding algorithm and kernel compilers) to experiment specific benchmarks (ATLAS KitValidation) and fast benchmarks based on random number generation.

Length of presentation (max. 20 minutes):

20

Computing and Batch Services / 26**Status of DESY Batch Infrastructures****Author:** Thomas Finnern¹¹ *DESY***Corresponding Author:** thomas.finnern@desy.de

This presentation will provide information on the status of the batch systems at DESY Hamburg. This includes the clusters for GRID, HPC and local batch purposes showing the current state and the activities for upcoming enhancements.

Length of presentation (max. 20 minutes):

20

Site Reports / 27**NDGF Site Report****Author:** Erik Mattias Wadenstein¹¹ *University of Umeå (SE)***Corresponding Author:** mattias.wadenstein@cern.ch

Update on recent events in the Nordic countries

Length of presentation (max. 20 minutes):

15

End-user Services and Operating Systems / 28**Self service kiosk for Mac and mobiles****Author:** Vincent Nicolas Bippus¹**Co-authors:** Michal Kwiatek¹; Tim Bell¹¹ *CERN***Corresponding Author:** tim.bell@cern.ch

CERN has recently deployed a Mac self service portal to allow users to easily select software and perform standard configuration steps.

This talk will review the requirements, product selection and potential evolution for mobile device management.

Length of presentation (max. 20 minutes):

20

Storage and Filesystems / 29**Scientific Data Storage at FNAL****Authors:** Dmitry Litvintsev¹; Gerard Bernabeu Altayo²¹ *FNAL*² *Fermilab***Corresponding Author:** gerard1@fnal.gov

Fermilab stores more than 110PB of data employing different technologies (dCache, EOS, BlueArc) to address a wide variety of use cases and application domains. This presentation captures present state of data storage at Fermilab and maps out future directions in storage technology choices at the lab.

Length of presentation (max. 20 minutes):

20 minutes

Grid, Cloud and Virtualization / 30

The Open Science Grid: Physics Experiments to Campus Researchers

Author: Chander Sehgal¹

Co-authors: Bodhitha Jayatilaka¹; Lauren Michael²; Robert William Gardner Jr³; Tanya Levshina⁴

¹ *Fermilab*

² *University of Wisconsin (US)*

³ *University of Chicago (US)*

⁴ *FERMILAB*

Corresponding Author: cssehgal@fnal.gov

The Open Science Grid (OSG) ties together computing resources from a broad set of research communities, connecting their resources to create a large, robust computing grid; this computing infrastructure started with large HEP experiments such as ATLAS, CDF, CMS, and Dzero and it has evolved so now it is also enabling the scientific computation of many non-physics researchers. OSG has been funded by the Department of Energy Office of Science and National Science Foundation since 2006 to meet the US LHC community's computational needs and support other science communities. The large physics experiments played a key role in evolving the job submission methods in OSG resulting in the implementation of pilot based architectures that improved the ease of job submission while improving site utilization. An important aspect of these overlay systems is the change in the trust relationship for sites which now depend more on the science communities (i.e. Virtual Organizations) instead of the individual researcher. The OSG has further extended this trust model to remove many dependencies on PKI certificates without reducing security and job traceability; today, many researchers use OSG without having a PKI certificate. OSG continues to evolve these access methods as part of its drive to improve the access to Distributed High Throughput Computing (DHTC) for US researchers and increase delivery of DHTC to users from a wide variety of disciplines at campus and research institutions. Presently, the OSG Open Facility delivers over 150 million computing wall hours annually to researchers who are not already owners of the OSG computing sites; this is primarily accomplished by harvesting and organizing the temporarily unused capacity (i.e. opportunistic cycles) from the various sites in the OSG. We present an overview of the infrastructure developed to accomplish this from access methods to harvesting opportunistic resources to providing user support to a diverse set of researchers. In addition, we discuss the architectures that have been developed for small research communities and individual researchers to use the OSG for science and to share their own compute clusters with the broader OSG community. We believe that expanded access to DHTC is an essential tool for scientific innovation and OSG continues to expand these services for researchers from all science domains.

Length of presentation (max. 20 minutes):

20

Grid, Cloud and Virtualization / 31

Simulating 5 Trillion Events on the Open Science Grid

Author: Martin Lothar Purschke¹

¹ *Brookhaven National Laboratory (US)*

Corresponding Author: purschke@bnl.gov

We report on a simulation effort using the Open Science Grid which utilized a large fraction of the available OSG resources for about 13 weeks in the first half of 2015.

sPHENIX is a proposed upgrade of the PHENIX experiment at the Relativistic Heavy Ion Collider. We have collected large data sets of proton-proton collision data in 2012 and 2013, and plan to carry out a similar study with upgraded sPHENIX detector further into the 2020s. One important aspect of the study is to understand the different contributions for the forward production of muons that are measured in the detector. This requires a large-scale PYTHIA simulation that matches the integrated luminosity of the PHENIX 2013 data set (about 5000 Billion collisions). Since there is no way to pre-select the different parton-level collision types, the simulation selects minimum-bias collisions and then discards those which are not of the desired topology, leading to an acceptable I/O-to-CPU ratio. This made this simulation ideal to be run on the Grid.

The Open Science Grid opened up the possibility of carrying out this simulation in a reasonable time frame. The project used about 5 million CPU hours total. We will report on the experience running the jobs on the OSG, and describe the steps there were taken to automate the production in such a way that the entire simulation could be run by one person.

Length of presentation (max. 20 minutes):

15

Site Reports / 32

Oxford and SouthGrid Site Status Report

Author: Peter Gronbech¹

¹ *University of Oxford (GB)*

Corresponding Author: pete.gronbech@physics.ox.ac.uk

Update on Oxford University Particle Physics group computing setup, including short updates from the other member sites of SouthGrid.

Length of presentation (max. 20 minutes):

15

Grid, Cloud and Virtualization / 33

First experiences with Mesos at RAL

Author: Andrew David Lahiff¹

¹ *STFC - Rutherford Appleton Lab. (GB)*

Corresponding Author: andrew.david.lahiff@cern.ch

Running services in containers managed by a scheduler offers a number of potential benefits compared to traditional infrastructure, such as increased resource utilisation through multi-tenancy, the ability to have elastic services, and improved site availability due to self-healing. At the RAL Tier-1 we have been investigating migration of services to an Apache Mesos cluster running on bare metal. In this model the whole concept of individual machines is abstracted away and services are run on the cluster in ephemeral Docker containers. This talk will describe the setup of our test Mesos infrastructure, including how we deal with service discovery, monitoring, and logging in such a dynamic environment. The challenges faced running grid middleware in this way and future plans will be discussed.

Length of presentation (max. 20 minutes):

20

Site Reports / 34

PDSF Site Report and Relocation

Author: Tony Quan¹

Co-authors: Iwona Sakrejda ; James Botts ²; Lisa Gerhardt ³

¹ *LBL*

² *LBNL*

³ *NERSC*

Corresponding Author: twquan@lbl.gov

PDSF, the Parallel Distributed Systems Facility, has been in continuous operation since 1996 serving high energy physics research. It is currently a tier-1 site for Star, a tier-2 site for Alice and a tier-3 site for Atlas.

The PDSF cluster will move early next year from its current site at Oakland to a new building on the LBNL campus. Several racks have already been installed at the new site. This site report will describe recent updates to the system, upcoming modifications, and how the PDSF cluster will operate between the two locations.

Length of presentation (max. 20 minutes):

20 Minutes

Storage and Filesystems / 35

Space usage monitoring for distributed heterogeneous data storage systems.

Author: Natalia Ratnikova¹

¹ *Fermilab*

Corresponding Author: natasha@fnal.gov

Prior to LHC Run 2 CMS collected over 100 PB of physics data on the distributed storage facilities outside CERN, and the storage capacity will considerably increase in the next years. During the Run 2 the amount of storage allocated to individual users analysis data will reach up to 40% of the total space pledged by the CMS sites. CMS Space Monitoring system is developed to give a comprehensive view of storage usage across distributed CMS sites, including centrally managed experiment data and individual users data not accounted centrally, as one common namespace.

We discuss general architecture and the components of the system, as well as the challenges we met in the process of deploying and operating the system at CMS Tier-1 and Tier-2 sites.

Length of presentation (max. 20 minutes):

20

Security and Networking / 36

Computer Security update

Author: Liviu Valsan¹

¹ CERN

Corresponding Author: liviu.valsan@cern.ch

This presentation provides an update on the global security landscape since the last HEPiX meeting. It describes the main vectors of compromises in the academic community including lessons learnt, presents interesting recent attacks while providing recommendations on how to best protect ourselves. It also covers security risks management in general, as well as the security aspects of the current hot topics in computing. By showing how the attacks we are facing are both sophisticated and profitable, the presentation concludes that the only mean to adopt an appropriate response is to build a tight international collaboration and trusted information sharing mechanisms within the community.

This talk is based on contributions and input from the CERN Computer Security Team.

Length of presentation (max. 20 minutes):

20

Storage and Filesystems / 37

Quality of Service in storage and the INDIGO-DataCloud project.

Authors: Giacinto Donvito¹; Marcus Hardt²; Patrick Fuhrmann³; Paul Millar^{None}

¹ INFN-Bari

² Karlsruhe Institute of Technology

³ Deutsches Elektronen-Synchrotron Hamburg and Zeuthen (DE)

Corresponding Author: patrick.fuhrmann@cern.ch

The pressure to provide cheap, reliable and unlimited cloud storage space in the commercial area has provided science with affordable storage hardware and open source storage solutions with low maintenance costs and tuneable performance and durability properties, resulting in different cost models per storage unit. Those models, already introduced by WLCG a decade ago (disk vs tape) are now offered by large companies like Amazon (Bock Storage, S3 and Glacier) or Google (Standard, Durable Reduced, Cloud Storage Nearline). Industry appliances or software stacks (e.g. HPSS) offer similar storage properties for your locally installed storage.

However, other than with SRM for WLCG, those offered storage quality properties don't follow a common description or specification and are hard to compare programatically. Moreover they aren't even close to a common way of been negotiated between the requesting client and the providing storage technology, which would be a prerequisite for federating different public and private storage services.

To fill this gap, the INDIGO-DataCloud project is proposing a process to agree on common semantics in describing QoS attributes in storage in a consistent way, independently of the used API or protocol.

The process involves gathering uses-cases from scientific communities and creating working groups in international organisations, like RDA, OGF and SNIA to further discuss possible solutions with other interested parties like EUDAT and EGI.

In a second step, based on feedback received, INDIGO will propose an implementation of the defined semantics to steer quality of service in storage as an extension to an existing industry standard, e.g. CDMI. As a proof of concept INDIGO will implement the proposed solution in storage systems used within the INDIGO project, like dCache, StoRM and some typical industry products, as a reference for other systems.

We are presenting our work at HEPiX, in order to receive feedback from HEP communities on our plans and because we regard the outcome of our work as beneficial for HEP in general and for WLCG experiments in particular, as with the proposed end of the usage of SRM in WLCG, experiments are again left alone with the necessary steering of data quality of service attributes (e.g. disk, tape or both) at the different storage endpoints.

This presentation will report on the goals achieved so far and on our next steps.

Length of presentation (max. 20 minutes):

20

Site Reports / 38

RAL Site Report

Author: Martin Bly¹

¹ *STFC-RAL*

Corresponding Author: martin.bly@stfc.ac.uk

Update from RAL Tier1.

Length of presentation (max. 20 minutes):

15

Site Reports / 39

Fermilab Site Report

Author: Rennie Scott¹

¹ *Fermilab*

Corresponding Author: rennie@fnal.gov

News and updates from Fermilab since the Spring HEPiX Workshop.

Length of presentation (max. 20 minutes):

15

End-user Services and Operating Systems / 40

Scientific Linux Status Update

Author: Rennie Scott¹

Co-authors: Bonnie King ¹; Pat Riehecky ¹; connie sieh ¹; kevin hill

¹ *Fermilab*

Corresponding Author: bonniek@fnal.gov

Scientific Linux status and news.

Length of presentation (max. 20 minutes):

15

Grid, Cloud and Virtualization / 42

A comparison of performance between KVM and Docker instances in OpenStack

Author: Wataru Takase¹

¹ *KEK*

Corresponding Author: wataru.takase@kek.jp

Cloud computing enables a flexible provisioning of computing resources by utilizing virtual machines on demand, and can provide an elastic data analysis environment. At KEK we plan to integrate cloud-computing technology into our batch cluster system in order to provide heterogeneous clusters dynamically, to support various kinds of data analyses, and to enable elastic resource allocation among the various projects supporting at KEK.

Container-based virtualization provides an alternative approach to virtual machines with a reduction in the overhead associated with the hardware and operating system resources on the host machine. We have evaluated the performance for virtual machines and containers in a cloud environment and have also investigated the performance of collocated instances on the same host machine. This talk will present a comparison of performance between KVM and Docker instances in OpenStack by using Rally.

Length of presentation (max. 20 minutes):

20

Security and Networking / 43

Using VPLS for VM mobility

Author: Carles Kishimoto Bisbe¹

Co-author: David Gutierrez Rueda¹

¹ *CERN*

Corresponding Author: carles.kishimoto@cern.ch

With the virtualization of the data centre, there is a need to move virtual machines transparently across racks when the physical servers are being decommissioned. We will present the solution being tested at CERN using VPLS in an MPLS network

Length of presentation (max. 20 minutes):

15

Security and Networking / 44

Network infrastructure for the CERN Datacentre

Author: Eric Sallaz¹

¹ *CERN*

Corresponding Author: eric.sallaz@cern.ch

With the evolution of transmission technologies, going above 10Gb Ethernet requires a complete renewal of the fibre infrastructure.

In the last year the CERN Datacentre has evolved to deal with the expansion of the physical infrastructure inside and outside the local site, to support higher speeds like 40GbE and 100GbE, to be ready for any other future requirement. We will explain the choice we made at CERN in terms of network infrastructure for our datacentre.

Length of presentation (max. 20 minutes):

20

IT Facilities and Business Continuity / 45

NERSC Plans for the Computational Research and Theory Building

Author: Elizabeth Bautista¹

¹ *Lawrence Berkeley National Lab*

Corresponding Author: ejbautista@lbl.gov

The NERSC facility is transitioning to a new building housed at the LBNL main campus in the 2015 timeframe. This state of the art facility is energy efficient, providing year-round free air and water cooling, is initially provisioned for 12.5 MW power and capable of up to 42 MW power, has two office floors, a 20K square foot HPC floor with seismic isolation and a mechanical floor. Substantial completion was scheduled for May 2015.

As more supercomputing facilities are investigating water cooling as an alternative, the Computational Research and Theory building (CRT) demonstrates the viability of energy efficiency using alternative environmental cooling and leverages the Bay Area's climate to maintain temperature throughout the HPC floor and offices. The seismic isolation floor ensures that the systems are able to withstand earthquakes that are known to occur in this geographic area. The backup power system uses gasoline and can sustain energy to maintain the critical infrastructure of the facility up to 24 hours and can be refilled. Environmental data is collected throughout the facility to monitor not only temperature but also humidity, dust, electrical data from the power distribution units and sub panels.

The move will occur through several months in order to maximize system availability, preserve users' data, minimize the number of outages and minimize move costs.

Length of presentation (max. 20 minutes):

20 min

Basic IT Services / 46

ELK at NERSC

Authors: Cary Whitney¹; Tom Davis¹

¹ *LBNL*

Corresponding Author: clwhitney@lbl.gov

The ELK (Elastic, Logstash, Kibana) stack has been chosen to be one of the key component for our new centerwide monitoring project. I'll discuss our overall philosophy on monitoring and how ELK fits in. The current structure and how it is performing.

Define centerwide: Everything in the data centers. All hosts, filesystems, most applications, power, cooling, water flow, temperature, particulate counts, starting from the substation and cooling tower to the node.

Length of presentation (max. 20 minutes):

20

Site Reports / 47

University of Wisconsin Madison CMS T2 site report

Authors: Ajit Kumar Mohapatra¹; Carl Vuosalo¹; Daniel Charles Bradley¹; Sridhara Dasu¹; Tapas Sarangi¹

¹ *University of Wisconsin (US)*

Corresponding Author: ajit.kumar.mohapatra@cern.ch

The University of Wisconsin Madison CMS T2 is a major WLCG/OSG T2 site. It has consistently delivered highly reliable and productive services for CMS MC production/processing, and large scale CMS physics analysis using high throughput computing (HTCondor), highly available storage system (Hadoop), efficient data access using xrootd/AAA, and scalable distributed software systems (CVMFS). The site is a member of the LHCONE community with a 100Gb WAN connectivity, and supports IPv6 networking. An update on the current status of and activities (since the last report at the Nebraska meeting) will be presented.

Length of presentation (max. 20 minutes):

20

Security and Networking / 49

Status on the IPv6 OSG Software stack tests

Author: Edgar Fajardo Hernandez¹

¹ *Univ. of California San Diego (US)*

Corresponding Author: edgar.mauricio.fajardo.hernandez@cern.ch

This talk will present the latest results from the IPv6 compatibility tests performed on the OSG Software stack.

Length of presentation (max. 20 minutes):

15

Miscellaneous / 50

Growing our own systems administrators

Author: Ian Peter Collier¹

¹ *STFC - Rutherford Appleton Lab. (GB)*

Corresponding Author: ian.peter.collier@cern.ch

Introducing a non-graduate recruitment path in scientific computing at the Rutherford Appleton Laboratory.

Recruiting and retaining high quality staff is an increasing challenge at STFC. We traditionally recruit people with relevant degrees and/or industry experience. But this becomes increasingly difficult, as does recruiting to our graduate recruitment program. At the same time steep rises in tuition fees mean that young people face the prospect of building many tens of thousands of pounds of debt to gain a degree. This creates an opportunity for a different approach. We describe the development of our new computing apprenticeship scheme which combines part time study toward s degree qualification with practical work experience in software development and systems administration.

Length of presentation (max. 20 minutes):

20

Basic IT Services / 51**CERN Monitoring Status Update****Author:** Miguel Coelho dos Santos¹**Co-authors:** Benjamin Fiorini¹; Luis Pigueiras²; Pedro Andrade¹; Susie Murphy¹¹ *CERN*² *Universidad de Oviedo (ES)***Corresponding Author:** miguel.coelho.santos@cern.ch

This presentation will provide an update of the activities concerning IT Monitoring. This includes the monitoring of data centre infrastructure, hardware monitoring, host monitoring and application monitoring; as well as the tools being used or tested.

Length of presentation (max. 20 minutes):

20

Computing and Batch Services / 52**Future of Batch Processing at CERN****Author:** Jerome Belleman¹¹ *CERN***Corresponding Author:** jerome.belleman@cern.ch

So as to have our Batch Service at CERN answer increasingly challenging scalability and flexibility needs, we have chosen to set up a new batch system based on HTCondor. We have set up a Grid-only pilot service and major LHC experiments have started trying it out. While the pilot is slowly becoming production-ready, we're laying out a plan for our next major milestone: to run local jobs too, providing all the service dependencies they need.

Length of presentation (max. 20 minutes):

20

IT Facilities and Business Continuity / 53**Asset management in CERN data centres****Author:** Eric Bonfillou¹¹ *CERN*

Corresponding Author: eric.bonfillou@cern.ch

This presentation will give an overview of the recent efforts to establish an accurate and consistent inventory and stock management for CERN data centre assets. The underlying tool, Infor EAM (<http://www.infor.com/solutions/eam/>), was selected because of its wider usage over many years in other areas at CERN. The presentation will focus on the structuring of the IT assets data and how it is prepared for the integration with the spare parts stock management. The talk will also summarize the experience and some lessons learnt from the long and manual process for achieving an accurate recording of the legacy assets. Finally the integration with an on-going undertaking of global assets management at CERN will be described.

Length of presentation (max. 20 minutes):

15

Grid, Cloud and Virtualization / 54

Improving IaaS resources to accommodate scientific applications

Authors: Alvaro Lopez Garcia¹; Giacinto Donvito²; Patrick Fuhrmann³

¹ *Unknown*

² *INFN-Bari*

³ *DESY*

Corresponding Author: chierici@cnaf.infn.it

INDIGO-DataCloud aims at developing a data and computing platform targeted at scientific communities, integrating existing standards and open software solutions. INDIGO proposes:

- to build-up a PaaS solution leveraging existing resources and e-Infrastructures, since the mere access to IaaS resources has been demonstrated as not being a realistic option for most Research Communities
- to improve the virtualization layer provided by already diffed IaaS

In order to build such a Scientific oriented PaaS, several improvements are needed at the IaaS layer. However, the approach of the project is to ensure that those improvements are beneficial enough by their own and not focused on just enabling the upper PaaS layer. Therefore, we find that they can be interesting for the HEP community and for the WLCG in particular.

In this presentation we will try to give an overview of the developments that we consider of interest for the HEPiX community, namely:

- Container support in Cloud stacks: INDIGO will incorporate specific drivers into the most used cloud management frameworks (i.e. OpenNebula and OpenStack) to support the deployment and execution of containers as first-class resources on the IaaS.
- Scheduling improvements: Currently scheduling in most cloud middlewares is based in first-come, first-served policies. The project will develop new intra-cloud scheduling algorithms that are more appropriate for scientific scenarios. Namely we foresee two different and complementary approaches: a fair-share based scheduling with a queuing mechanism similar to the current batch scheduling and a spot-instances based approach for opportunistic usage.
- IaaS orchestration through standards: We will provide IaaS orchestration using a common standard language (TOSCA) in both OpenStack and OpenNebula, easing the deployment and management of (computing, storage and network) resources at the IaaS layer.

- Container integration in batch systems: We will study the possibility of executing containers through batch systems, with the aim of providing the access to specialized hardware (such as Infiniband networks and GPGPUs).

Length of presentation (max. 20 minutes):

20

Site Reports / 55

ATLAS Great Lakes Tier-2 (AGLT2) Update

Author: Shawn Mc Kee¹

Co-authors: Benjeman Jay Meekhof¹; Philippe Alain Luc Laurens²; Raymond Brock²; Richard Drake³; Robert Ball¹

¹ *University of Michigan (US)*

² *Michigan State University (US)*

³ *Michigan State Univesity*

Corresponding Author: shawn.mckee@cern.ch

We will present an update on our site since the Fall 2014 report and cover our work with various storage technologies, ATLAS Muon Calibration and our use of the ELK stack for central syslogging. We will also report on our recent hardware purchases for 2015 as well as the status of our new networking configuration and 100G connection to the WAN. I conclude with a summary of what has worked and what problems we encountered and indicate directions for future work.

Length of presentation (max. 20 minutes):

15+5

Security and Networking / 56

Update on WLCG/OSG perfSONAR Infrastructure

Authors: Marian Babik¹; Shawn Mc Kee²

¹ *CERN*

² *University of Michigan (US)*

Corresponding Author: shawn.mckee@cern.ch

WLCG relies on the network as a critical part of its infrastructure and therefore needs to guarantee effective network usage and prompt detection and resolution of any network issues, including connection failures, congestion and traffic routing. The WLCG Network and Transfer Metrics working group was established to ensure sites and experiments can better understand and fix networking issues. In addition, it aims to integrate and combine all network-related monitoring data collected by the WLCG infrastructure from both network and transfer systems. This has been facilitated by the already existing network of the perfSONAR instances that is being commissioned to operate in full production.

Recently, several higher level services were developed to help bring perfSONAR network to its full potential. This includes a Web-based mesh configuration system, which allows to centrally schedule and manage all the network tests performed by the instances; a network datastore (esmond), which collects, stores and provides interfaces to access all the network monitoring information from a single place as well as perfSONAR infrastructure monitoring, which ensures that the current perfSONAR instances are configured and operated correctly.

In this presentation we will provide an update on how to use and benefit from perfSONAR, including new features that are included in the recent release of version 3.5 of the Toolkit. We will also cover the status of our WLCG/OSG deployment and provide some information on our future plans.

Length of presentation (max. 20 minutes):

20

Basic IT Services / 57

Quattor Update

Author: James Adams¹

¹ *STFC RAL*

Corresponding Author: james.adams@stfc.ac.uk

An update of developments and activities in the Quattor community over the last six months.

Length of presentation (max. 20 minutes):

15

Security and Networking / 58

WLCG Network and Transfer Metrics WG After One Year

Authors: Marian Babik¹; Shawn Mc Kee²

¹ *CERN*

² *University of Michigan (US)*

Corresponding Author: shawn.mckee@cern.ch

It has been approximately one year since the WLCG Network and Transfer Metrics working group was initiated and we would like provide a summary what has been achieved during this first year and discuss future activities planned for the group.

The working group as chartered had a number of objectives:

- Identify and make continuously available relevant network and transfer metrics
- Document those metrics and their use
- Facilitate their integration in the middleware and/or experimental tool chains
- Coordinate commissioning and maintenance of WLCG network monitoring

We will report on the status of these objectives and summarize the work that has been done in a number of areas, concluding with our vision for where the working group is likely to go in the future.

Length of presentation (max. 20 minutes):

20

Grid, Cloud and Virtualization / 59

Running ATLAS, CMS, ALICE Workloads on the NERSC Cray XC30

Authors: Doug Jacobsen¹; James Botts¹; Lisa Gerhardt¹

¹ *LBNL*

Corresponding Author: jfbotts@lbl.gov

This presentation will describe work that has been done to make NERSC Cray systems friendlier to data intensive workflows in anticipation of the availability of the NERSC-8 system this Autumn. Using shifter, a docker-like container technology developed at NERSC by Doug Jacobsen and Shane Canon, the process of delivering software stacks to Cray compute nodes has been greatly simplified. Some detail will be given to describe the method used to deliver CVMFS to the Cray compute nodes using shifter. We found unexpected increases in metadata performance though the use of image file systems. Results of tests with ATLAS G4 simulations and Analysis Software on the Cray XC30 system, edison, will be presented.

Length of presentation (max. 20 minutes):

20

Computing and Batch Services / 60

HTCondor Recent Enhancement and Future Directions

Author: Todd Tannenbaum¹

¹ *Univ of Wisconsin-Madison, Wisconsin, USA*

Corresponding Author: tannenba@cs.wisc.edu

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 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 the HEP community.

Length of presentation (max. 20 minutes):

20

Grid, Cloud and Virtualization / 61

HTCondor-CE: Managing the Grid with HTCondor

Author: Brian Lin¹

¹ *University of Wisconsin - Madison*

Corresponding Author: brian.h.lin@gmail.com

HTCondor-CE is a special configuration of HTCondor designed to connect compute resources to the wider grid. Leveraging the power of HTCondor, HTCondor-CE is able to provide built-in security measures, end-to-end job tracking, and better integration with overlay job systems. This talk will present an overview of the HTCondor-CE software, its deployment in the Open Science Grid (OSG), and upcoming development plans.

Length of presentation (max. 20 minutes):

20

Site Reports / 62

Jefferson Lab Scientific and High Performance Computing

Author: Sandy Philpott^{None}

Corresponding Author: sandy.philpott@jlab.org

Current high performance and experimental physics computing environment updates: core exchanges between USQCD and Experimental Physics clusters for load balancing, job efficiency, and 12GeV data challenges; Nvidia K80 GPU experiences and updated Intel MIC environment; update on locally developed workflow tools and write-through to tape cache filesystem; status of LTO6 integration into our MSS; ZFS on Linux production environment, and status of our Lustre 2.5 update.

Length of presentation (max. 20 minutes):

15

Site Reports / 63

Australia-ATLAS site report

Author: Lucien Philip Boland¹

Co-authors: Goncalo Borges ²; Sean Crosby ¹; Shunde Zhang ³

¹ *University of Melbourne (AU)*

² *University of Sydney (AU)*

³ *University of Adelaide (AU)*

Corresponding Author: lucien.boland@cern.ch

Update on activities at Australia's HEP Tier 2 grid facility.

Length of presentation (max. 20 minutes):

15

Site Reports / 64

BEIJING Site Report

Author: Qiulan Huang¹

¹ *Institute of High Energy of Physics Chinese Academy of Science*

Corresponding Author: qhuang2015@gmail.com

News and updates from IHEP since the Spring HEPiX Workshop. In this talk we will present a brief status of IHEP site including computing farm, Grid, data storage ,network and so on.

Length of presentation (max. 20 minutes):

20

Security and Networking / 65

Building a large scale Security Operations Centre

Author: Liviu Valsan¹

¹ *CERN*

Corresponding Author: liviu.valsan@cern.ch

The HEP community is facing an ever increasing wave of computer security threats, with more and more recent attacks showing a very high level of complexity. Having a centralised Security Operations Centre (SOC) in place is paramount for the early detection and remediation of such threats. Key components and recommendations to build an appropriate monitoring and detection Security Operation Centre will be presented, as well as means to obtain and share relevant and accurate threat intelligence information. The presentation concludes that the key to achieve an appropriate response is to both build an efficient security infrastructure and a tight international collaboration, enabling information to be shared globally with trusted partners, and in particular between the various HEP sites.

Length of presentation (max. 20 minutes):

20

End-user Services and Operating Systems / 66**Using ownCloud to provide mobile access to existing storage**

Author: Chris Brew¹

Co-authors: Federico Melaccio²; Ian Loader³; Kevin Dunford⁴

¹ *STFC - Rutherford Appleton Lab. (GB)*

² *STFC RAL PPD*

³ *STFC - RAL*

⁴ *Rutherford Appleton Laboratory (STFC)*

Corresponding Author: chris.brew@stfc.ac.uk

We present our work deploying on ownCloud gateway to our existing home file storage to provided access to mobile clients.

Length of presentation (max. 20 minutes):

20

Grid, Cloud and Virtualization / 67**Automated performance testing framework**

Author: Sean Crosby¹

Co-author: Arne Wiebalck²

¹ *University of Melbourne (AU)*

² *CERN*

Corresponding Author: sean.christopher.crosby@cern.ch

In the CERN Cloud Computing project, there is a need to ensure that the overall performance of hypervisors and virtual machines does not decrease due to configuration changes, or just because of the passage of time.

This talk will outline an automated performance framework currently being developed, which will allow performance of virtual machines and hypervisors to be graphed and linked.

Length of presentation (max. 20 minutes):

10

Storage and Filesystems / 68**Supervised Machine Learning**

Author: Amit Chattopadhyay¹

¹ *WD*

Corresponding Author: amit.chattopadhyay@wdc.com

Given the amount of inline data generated, large volume hard-drive manufacturing is an appropriate environment to employ contemporary 'Big Data' techniques. It can be used to generate a feedback loop, as well as a feed-forward path. The feedback loop helps fix what is broken. The feed-forward path can be used to predict drive health.

In this presentation, I will focus on some work we have started at Western Digital - using Supervised Machine Learning techniques to predict the health of a hard drive and develop a hierarchy of "health grades". As the technique matures, this could fundamentally influence the manufacturing environment as well as WD's customer engagement.

Length of presentation (max. 20 minutes):

20

Site Reports / 69

SLAC Scientific Computing Services

Author: Yemi Adesanya^{None}

Corresponding Author: yemi.oluwayomi@gmail.com

An update on SLAC's central Unix services in support of Scientific Computing and Core Infrastructure. New funding model for FY15 identifies indirect vs direct-funded effort. Socializing the concept of service and service lifecycle. Sustainable business models to address hardware lifecycle replacement: Storage-as-a-Service with GPFS, OpenStack for dev/test environments and cluster provisioning.

Length of presentation (max. 20 minutes):

20 mins

Basic IT Services / 70

Foreman - A full automation solution for a growing facility

Author: Mizuki Karasawa¹

¹ BNL

Corresponding Author: mizuki@bnl.gov

In a rapidly growing facility as NSLS-II, we use foreman as an automation tool that integrates to DNS, DHCP, TFTP, Puppet which makes installation & provisioning processes much easier and help to bring the service/server components online in a short timely manner. For those who use Puppet Enterprise as an paid version ENC, Foreman can also substitute of that. This talk will present the detail of the foreman architecture in NSLS-II and the work flow in a real computing facility.

Length of presentation (max. 20 minutes):

20 mins

Basic IT Services / 71**Gitlab and its CI - an intelligent solution for hosting Git repos****Author:** Mizuki Karasawa¹¹ *BNL***Corresponding Author:** mizuki@bnl.gov

Gitlab - a MIT licensed open source tool, that cooperates a set of rich features managing Git repositories, code reviews, issue tracking, activity feeds and wikis. The most powerful feature - CI for continuous integration makes the code developing much more efficient and cost saving, it's also a great tool to enhance the communication and collaboration. In NSLS-II, we have a great number of developers, currently hosting more than 300+ repos in mercurial, in a direction moving towards new git repo hosting solutions such as Gitlab, this presentation will discuss the details of Gitlab and the architecture view of Gitlab in NSLS-II.

Length of presentation (max. 20 minutes):

20 mins

Computing and Batch Services / 72**Upgrade to UGE 8.2 : Positives effects at IN2P3-CC****Authors:** Nadia Lajili¹; Suzanne Poulat²¹ *Centre de Calcul de l'IN2P3 (CC-IN2P3)*² *Centre de calcul IN2P3***Corresponding Author:** vanessa.hamar@cc.in2p3.fr

We are using Univa Grid Engine as BATCH scheduling system to our satisfaction since four years. We focus on the latest major version 8.2.1, which was deployed at IN2P3-CC 4 months ago, and provides further scalability improvements.

We are supporting about 200 groups and experiments running up to 17.000 jobs simultaneously. The requirements, in terms of computing resources, storage or network, differ widely from one group to another, and therefore the batch system must be particularly scalable and reliable which is what we get with this release.

The new read-only threads work independently from the SGE qmaster which removes the load previously induced by serving status requests.

In addition, this version introduces a way to limit user requests which avoids system overloads, it comes also with new information for job accounting and uses a 32-bit range for job ID's.

In this talk we will present the assessment of this upgrade and give an overview of the issues being resolved according to the operation of the service.

Finally, we will show the plans for the deployment of new features and the roadmap of Univa Grid Engine at the IN2P3-CC for the next few months.

Length of presentation (max. 20 minutes):

15 minutes

Site Reports / 73**Site report from KEK**

Author: Tomoaki Nakamura¹

Co-authors: Go Iwai²; Hiroyuki Matsunaga¹; Koichi Murakami¹; Soh Suzuki; Takashi Sasaki¹; Wataru Takase²

¹ *High Energy Accelerator Research Organization (JP)*

² *KEK*

Corresponding Author: tomoaki.nakamura@cern.ch

KEK computing research center supports various project of accelerator based science in Japan. Hadron and Neutrino experiments (T2K) at J-PARC have started with good rate after the recovery of earthquake damage at Fukushima. Belle II experiment is going to collect 100PB of raw data within the several years. In this talk, we would like to report the current status of our computing facility and near future plan of resource provisioning for the KEK experiments. The current situation of international network for Japan and our preparation of Grid system will be also introduced.

Length of presentation (max. 20 minutes):

15

Site Reports / 74**Wigner Datacenter - Site report**

Author: Szabolcs Hernath¹

¹ *Hungarian Academy of Sciences (HU)*

Corresponding Author: szabolcs.tamas.hernath@cern.ch

We give an update on the infrastructure, Tier-0 hosting services, Wigner Cloud and other recent developments at the Wigner Datacenter. We will also include a short summary on the Budapest WLCG Tier-2 site status as well.

Length of presentation (max. 20 minutes):

15 minutes

Computing and Batch Services / 75**Update on benchmarking**

Authors: Helge Meinhard¹; Manfred Alef²; Michele Michelotto³

¹ *CERN*

² *Karlsruhe Institute of Technology (KIT)*

³ *Universita e INFN, Padova (IT)*

Corresponding Authors: helge.meinhard@cern.ch, michele.michelotto@cern.ch

Status of the benchmarking working group and work going on in WLCG around benchmarking

Length of presentation (max. 20 minutes):

20 minutes

Storage and Filesystems / 76

USATLAS dCache storage system at BNL

Authors: Hironori Ito¹; Yingzi Wu²; Zhenping Liu¹

¹ *Brookhaven National Laboratory (US)*

² *Brookhaven National Laboratory*

Corresponding Author: zhliu@bnl.gov

As ATLAS tier-1 computing facility center in US, RHIC and ATLAS Computing Facility (RACF) at Brookhaven National Lab has been operating a very large scale disk storage system dCache with tape back-end to serve a geographically diverse, worldwide ATLAS scientific community. This talk will present the current state of USATLAS dCache storage system at BNL. It will describe structure, configuration, usage, data replication, tape-backend, client access, monitoring methods, and etc. Our near-future plan of reconfiguration from symmetric replication to asymmetric replication will also be discussed.

Length of presentation (max. 20 minutes):

20

Site Reports / 77

KIT Site Report

Author: Andreas Petzold¹

¹ *KIT - Karlsruhe Institute of Technology (DE)*

Corresponding Author: andreas.petzold@cern.ch

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

Length of presentation (max. 20 minutes):

15

Grid, Cloud and Virtualization / 78

computing resources virtualization in IHEP

Authors: Jingyan Shi¹; Tao Cui²; Yaodong Cheng¹

¹ *IHEP*

² *IHEP(Institute of High Energy Physics, CAS,China)*

Corresponding Author: qhuang2015@gmail.com

The report will introduce our plan of computing resources virtualization. We will discuss progress of this project such as Openstack-based infrastructure, virtual scheduling and measure. And we will share experience of our test-bed which deployed with Openstack icehouse.

Length of presentation (max. 20 minutes):

20

Basic IT Services / 79

Host deployment and configuration technologies at SCC.

Author: Dmitry Nilsen^{None}

Co-authors: Andreas Petzold¹; Pavel Weber¹

¹ *KIT - Karlsruhe Institute of Technology (DE)*

Corresponding Author: dmitry.nilsen@googlemail.com

Host deployment and configuration technologies at SCC.

The Steinbuch Centre for Computing(SCC) at Karlsruhe Institute Of Technology (KIT) serves a number of projects, including the WLCG Tier-1 GridKa, the Large Scale Data Facility (LSDF), and the Smart Data Innovation Lab (SDIL) using bare metal and virtual compute resources and provides a variety of storage and computing services to the end users. To take control of the provisioning and configuration process, a unified configuration management and deployment infrastructure was set up, based on a stack of open source solutions like Puppet, Foreman, GitLab, GitLab-CI and others. Because each project has its specific use cases, the infrastructure is build to be customizable. Underlying workflows and software setups ensure that they can be reused and are as uniform as possible. Workflows to provision, setup and manage hosts and services are defined. Users are able to access infrastructure using existing configuration templates, or provide custom solutions, sharing them with other projects.

To achieve control over changes everything is managed with GIT/GitLab and Continues Integration(CI), so that all production and testing services can be rolled back in case of failures.

This talk describes the evaluation of workflows and gives an overview of the components and their connections within the infrastructure.

Length of presentation (max. 20 minutes):

15 min

Site Reports / 80

CSC.fi - Site Report

Author: Johan Henrik Guldmyr¹

¹ *Helsinki Institute of Physics (FI)*

Corresponding Author: johan.guldmyr@cern.ch

- CSC
- Ansible and HPC
- Slurm as an interactive shell load balancer

Length of presentation (max. 20 minutes):

15

Basic IT Services / 81

Monitoring with InfluxDB and Grafana

Author: Andrew David Lahiff¹

Co-author: George Vasilakakos²

¹ *STFC - Rutherford Appleton Lab. (GB)*

² *STFC*

Corresponding Author: andrew.david.lahiff@cern.ch

At RAL we have been considering InfluxDB and Grafana as a possible replacement for Ganglia, in particular for application-specific metrics. Here we present our experiences with setting up monitoring for services such as Ceph, FTS3 and HTCondor, and discuss the advantages and disadvantages of InfluxDB and Grafana over Ganglia.

Length of presentation (max. 20 minutes):

15

Welcome to BNL / 82

Welcome to BNL

Author: Michael Ernst¹

¹ *Unknown*

Corresponding Author: mernst@bnl.gov

Length of presentation (max. 20 minutes):

15

Site Reports / 83

BNL Site Report

Author: Shigeki Misawa¹

¹ *Brookhaven National Laboratory*

Corresponding Author: misawa@bnl.gov

This site report will discuss the latest developments at the RHIC-ATLAS Computing Facility (RACF).

Length of presentation (max. 20 minutes):

20

Site Reports / 84

LAL Site Report

Author: Michel Jouvin¹

¹ *Laboratoire de l'Accelérateur Lineaire (FR)*

Corresponding Author: jouvin@lal.in2p3.fr

Changes at LAL and GRIF in the last 18 months.

Length of presentation (max. 20 minutes):

15

Grid, Cloud and Virtualization / 85

Introduction to GDB

Author: Michel Jouvin¹

¹ *Laboratoire de l'Accelérateur Lineaire (FR)*

Corresponding Author: jouvin@lal.in2p3.fr

This presentation is a general introduction, and it will also describe the opportunities for cooperation with HEPiX and how sites can participate in GDB.

Length of presentation (max. 20 minutes):

10

Site Reports / 86

DESY Site Report

Author: Wolfgang Friebel¹

¹ *Deutsches Elektronen-Synchrotron Hamburg and Zeuthen (DE)*

Corresponding Author: wolfgang.friebel@desy.de

Updates from DESY Zeuthen

Length of presentation (max. 20 minutes):

15

Closing and HEPiX Business / 87

HEPiX Wrap-up

Corresponding Authors: helge.meinhard@cern.ch, tony@bnl.gov