

HEPiX Spring 2016 Workshop

Monday 18 April 2016 - Friday 22 April 2016

DESY Zeuthen



Book of Abstracts

Contents

A Network Security self-service platform (NSSP) of IHEP	1
A virtual private network based on software-defined network architecture for high energy physics scientific data exchange	1
ASAP3: Status update and activities for XFEL	2
Authorization extension for the secure use of Elasticsearch and Kibana	2
Automating operational procedures with Rundeck	2
BNL RHIC/ATLAS Computing Facility Site Report	3
CERN Cloud Status Update	3
CERN Site Report	3
CPU Benchmarking at GridKa (Update April 2016)	4
Chef@GSI revisited	4
Computer Security update	5
Computing and Storage for Life Science at MDC	5
Computing on Low Power Architectures	5
Consolidating Scientific Computing Services at BNL	6
DCD - Desktop Chromodynamics, or: Linux on DESY Desktops	6
DESY Hamburg infrastructure	7
DESY site report	7
Deploying services with Mesos at a WLCG Tier 1	7
Evolution of CERN Printing Services	8
Fermilab HEP Cloud: an elastic computing facility for High Energy Physics.	8
Fermilab Site Report	9
Fermilab's Scientific Computing Storage Architecture	9
First Experiences with Container Orchestration in the CERN Cloud	10

ForHLR - New Energy Efficient HPC System at KIT with Warm Water Cooling	10
GSI Site Report	11
GreenITCube - the new Data Center for FAIR & GSI	11
Grid Computing System in the KEK Central Computer System	11
HEPiX AGLT2 Site Report Spring 2016	12
HTCondor European Workshop summary	12
HTCondor at DESY	13
IBM Spectrum Scale support for technical workflows	13
IRFU Site report	13
Ideas for a journal around "Computing and Software for data-intensive physics"	14
Identity Federation for HEP – what are the benefits?	14
Jefferson Lab Scientific and High Performance Computing	15
KR-KISTI-GSDC-01 Tier-1 Site Reports	15
Logistics	15
Monitoring at LHCb: Migrating to Icinga2,Puppet, Hiera and Foreman Stack for Monitor- ing.	16
Monitoring at scale: a needle in the haystack	16
NDGF Site Report	17
NERSC collection architecture	17
PDSF Site Report	17
RAL Site Report	18
Recent work of the HEPiX IPv6 Working Group	18
Running virtualized Hadoop, does it make sense?	18
SCD Cloud at STFC	19
Scientific Linux Update	19
Security and CSIRT in KEK	20
Server gerontology in LHCb	20
Status and new developments for lustre@GSI	20
Status report of TReqS	21
Storage at CERN: towards the Chamäleon	21

The Further Adventures of the NERSC Data Collect	22
The HELIX NEBULA Science Cloud project	22
The OSiRIS Project: Meeting the Multi-Institutional Data Collaboration Challenge	22
The State of OpenAFS	23
The Status of IHEP Site	23
The next generation system of KEKCC	24
US ATLAS SWT2 Site Report	24
Update from Database Services	24
Using Containers for HPC Workloads	25
Using Docker container virtualization in DESY HPC environment	25
Using Ganeti for running highly available virtualized services	25
Virtual Cluster Computing in IHEPCloud	26
Welcome to DESY Zeuthen	26
Why so Sirius? Ceph backed storage at the RAL Tier-1	27
Workshop wrap-up	27
ownCloud	27
perfSONAR Status in WLCG/OSG	27
xBatch: Extending the CERN Batch Service Into the Public Cloud	28

Security and networking / 43**A Network Security self-service platform (NSSP) of IHEP****Author:** Fazhi Qi¹**Co-authors:** Tian Yan² ; zhihui sun³¹ *Chinese Academy of Sciences (CN)*² *Institution of High Energy Physics, Chinese Academy of Science*³ *Institute of High Energy Physics Chinese Academy of Sciences***Corresponding Author:** sunzh@ihep.ac.cn

Network security has been progressively coming to the attention of the high energy physics (HEP) community. More and more HEP users and system administrators keep worrying about the security of their hosts. In order to help users getting rid of their host vulnerability, we developed and deployed a network security self-service platform (NSSP) in Institute of High Energy Physics (IHEP), China. This NSSP system can present straightforward result of quantized fuzzy evaluation of host security risk. The result is obtained by analytic hierarchy process and cloud model theory. We improved the multi-level index system of the analytic process by dynamic weight method, which increased the adaptability and objectivity of the index system. Moreover, the assessment follows the national classification standard on information security protection.

The system utilizes a distributed architecture based on the message queue. Its efficiency is optimized by the buffer and multi-thread technology. The system is scalable because we can bring in a new function model by adding new plug-in. Benefiting from the B/S architecture, end users can easily carry out self-evaluation from Web client.

This platform has been deployed at IHEP since April 2014. During the past two years of production, hundreds of users enhanced their host security with the help of this platform. And we are continue developing new features for this platform.

Length of presentation (minutes, max. 20):15 **Summary:****Security and networking / 42****A virtual private network based on software-defined network architecture for high energy physics scientific data exchange****Authors:** Fazhi Qi¹ ; zhihui sun²**Co-author:** Tao Cui³¹ *Chinese Academy of Sciences (CN)*² *Institute of High Energy Physics Chinese Academy of Sciences*³ *IHEP(Institute of High Energy Physics, CAS,China)***Corresponding Author:** sunzh@ihep.ac.cn

This presentation will detail a software defined virtual private network serving the massive data exchange of HEP(high energy physics) , and introduce a software-defined network crossing different locations among the collaborative members of HEP experiments. An intelligent network route algorithm was also designed to exploit the ipv6 resources for HEP scientific data transfer. The algorithm also solved the network bandwidth problem between the collaborative members for the HEP experiments.

Length of presentation (minutes, max. 20):15 **Summary:**

Storage and file systems / 46**ASAP3: Status update and activities for XFEL**

Author: Stefan Dietrich¹

Co-authors: Janusz Malka¹ ; Manuela Kuhn¹ ; Martin Gasthuber¹ ; Uwe Ensslin¹

¹ *DESY*

Corresponding Author: stefan.dietrich@desy.de

Since April 2015 we have been running our new storage infrastructure based on GPFS for the data acquisition and analysis of PETRA III.

This presentation will show the current state of ASAP3, experiences from the first run period in production and current activities for XFEL.

Length of presentation (minutes, max. 20):

20 **Summary:**

Basic IT services / 32**Authorization extension for the secure use of Elasticsearch and Kibana**

Author: Wataru Takase¹

¹ *KEK*

Although Elasticsearch and Kibana bring great monitoring platform, they lack access control feature by default. This means any user who can access to Kibana can retrieve any information from Elasticsearch. In CERN cloud service, a homemade Elasticsearch plugin has been deployed to restricts data access based on cloud user. It enables each user to have a separated dashboard for cloud usage. Based on the solution, we integrated Kerberos authentication to Kibana and Elasticsearch. Our solution enables user/role based Elasticsearch access control and Kibana dashboards separation. The integration and deployment was completed in CC-IN2P3.

Length of presentation (minutes, max. 20):

20 **Summary:**

Basic IT services / 12**Automating operational procedures with Rundeck**

Author: Daniel Fernandez Rodriguez¹

¹ *Universidad de Oviedo (ES)*

Corresponding Author: daniel.fernandez@cern.ch

During the past two years, CERN Cloud Infrastructure has been using an open source tool called Rundeck for automating routine operational procedures. The aim of this project was to provide the team with a common place for implemented workflows and jobs. Thanks to Rundeck we were able to delegate internal tasks to other teams without exposing internal procedures or credentials. In addition to this, it has been possible to automate some tedious and repetitive tasks and minimize the

amount of human prone errors.

Currently, Rundeck is being used by four different groups at CERN and the jobs run on it cover diverse tasks such as HW interventions, OpenStack Project Creations and Quota Updates, tickets resolution, retirement of physical hardware, etc.

In this talk it will be presented what we have been doing to adapt Rundeck to our needs and how we are using it.

Length of presentation (minutes, max. 20):

20 **Summary:**

Site reports / 6

BNL RHIC/ATLAS Computing Facility Site Report

Author: Christopher Hollowell¹

¹ *Brookhaven National Laboratory*

Corresponding Author: hollowec@bnl.gov

Presentation of recent developments at Brookhaven National Laboratory's (BNL) RHIC/ATLAS Computing Facility (RACF).

Length of presentation (minutes, max. 20):

15

Grids, clouds, virtualisation / 9

CERN Cloud Status Update

Author: Arne Wiebalck¹

¹ *CERN*

Corresponding Author: arne.wiebalck@cern.ch

We'll give an update on the status of our cloud, highlighting some of the recently added features (such as federation or container support).

Summary:

Length of presentation (minutes, max. 20):

20

Site reports / 8

CERN Site Report

Author: Arne Wiebalck¹

¹ *CERN*

Corresponding Author: arne.wiebalck@cern.ch

News from CERN since the BNL workshop.

Length of presentation (minutes, max. 20):

15

Computing and batch systems / 10

CPU Benchmarking at GridKa (Update April 2016)

Author: Manfred Alef¹

¹ *Karlsruhe Institute of Technology (KIT)*

Corresponding Author: manfred.alef@kit.edu

Presentation of the latest CPU benchmarking results at GridKa:

- Scaling of HS06 with HEP applications
- First suggestions for a fast benchmark

Length of presentation (minutes, max. 20):

20 **Summary:**

Basic IT services / 34

Chef@GSI revisited

Author: Christopher Huhn¹

Co-author: Dennis Klein¹

¹ *GSI*

Corresponding Author: c.huhn@gsi.de

At Hepix Fall 2011 at Vancouver I gave a presentation about GSI's starting migration from CFengine to Chef configuration management.

This migration was a bumpier ride than initially expected (as usual?).
So now, 5 years later, I'd like to

- take a look back at our intentions for the migration,
- the difficulties we encountered,
- the current situation and issues still to be solved,
- the lessons we learned,
- and the possibilities we still see to improve or workflow with Chef.

Length of presentation (minutes, max. 20):

20 **Summary:**

Security and networking / 19

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 compromise in the academic community including lessons learnt, presents interesting recent attacks and security vulnerabilities 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 to implement trusted information sharing mechanisms within the community.

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

Length of presentation (minutes, max. 20):

20 **Summary:**

Computing and batch systems / 0

Computing and Storage for Life Science at MDC

Author: Alf Wachsmann¹

¹ *Max Delbrück Center for Molecular Medicine (MDC)*

Corresponding Author: awachs@gmail.com

I will introduce the Max Delbrück Center for Molecular Medicine (MDC; Berlin, Germany) with special focus on high-performance computing and storage. I will present the challenges recent developments in gene sequencing and imaging equipment pose for IT.

Length of presentation (minutes, max. 20):

20 **Summary:**

Computing and batch systems / 48

Computing on Low Power Architectures

Author: Michele Michelotto¹

¹ *Universita e INFN, Padova (IT)*

Corresponding Author: michele.michelotto@cern.ch

Low power architectures and SoC processor are still immature to build a computing farm for HEP, but they are still capable of running HEP-SPEC06 and HEP applications. The performances are not at the level of the x86 architectures however the HS06/watt is much better.

Length of presentation (minutes, max. 20):

20 **Summary:**

IT infrastructure / 4

Consolidating Scientific Computing Services at BNL

Authors: Alexandr Zaytsev¹ ; Christopher Hollowell² ; Costin Caramarcu¹ ; Shigeki Misawa² ; Tony Wong² ; William Strecker-Kellogg³

¹ *Brookhaven National Laboratory (US)*

² *Brookhaven National Laboratory*

³ *Brookhaven National Lab*

BNL is undergoing a re-organization of scientific computing services with the RACF as its core. This presentation describes the motivation, plans, current status and future plans of this consolidation, and the implications to the scientific community served by BNL.

Summary:

Length of presentation (minutes, max. 20):

20

End-user services and operating systems / 62

DCD - Desktop Chromodynamics, or: Linux on DESY Desktops

Author: Yves Kemp¹

Co-author: Stephan Wiesand¹

¹ *Deutsches Elektronen-Synchrotron (DE)*

Corresponding Authors: yves.kemp@cern.ch, stephan.wiesand@desy.de

In recent years, DESY has discussed within IT and with users the Linux Desktop strategy. This presentation will give insight why this discussion was necessary, which arguments came up, which solutions were implemented, and what the experience is after some months of running the latest "Ubuntu green desktop" at the Hamburg site, as well as its main features (and rationale behind):

- Ubuntu LTS, support life limited to one year after availability of successor
- Fully centrally managed (e.g. with puppet)
- No root rights for users, but ability to install software from approved repositories
- Local \$HOME, but access to network filesystems where feasible

Length of presentation (minutes, max. 20):

15

IT infrastructure / 64

DESY Hamburg infrastructure

Author: Martin Koch¹

¹ *DESY Hamburg*

Corresponding Author: martin.koch@desy.de

At the DESY location in Hamburg a distant cooling ring has been built and for the future growth of the computing resources a new cooling distribution was put into operation in the data center which will be accompanied by a new electrical power infrastructure soon. This presentation describes the motivation, plans, current status and future plans of these projects.

Length of presentation (minutes, max. 20):

15

Site reports / 54

DESY site report

Authors: Peter van der Reest¹ ; Yves Kemp²

¹ *DESY*

² *Deutsches Elektronen-Synchrotron (DE)*

Corresponding Authors: peter.van.der.reest@desy.de, yves.kemp@cern.ch

News from DESY since the Fall 2015 meeting

Length of presentation (minutes, max. 20):

20 **Summary:**

Grids, clouds, virtualisation / 27

Deploying services with Mesos at a WLCG Tier 1

Author: Andrew David Lahiff¹

Co-author: Ian Peter Collier¹

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

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

Container orchestration is rapidly emerging as a means of gaining many potential benefits compared to a traditional static infrastructure, such as increased resource utilisation through multi-tenancy, the ability to handle changing loads due to elasticity, and improved availability as a result of self-healing.

Whilst many large organisations are using this technology, in some cases for many years, it is not yet common in the scientific community.

At the RAL Tier-1 we have been investigating migration of services to an Apache Mesos cluster running on bare metal. In this architecture the whole concept of individual machines is abstracted away and services are run on the cluster in ephemeral Docker containers. Instead of the standard approach of manually placing long-running services on specific hosts, services are managed by a scheduler. This means that any host or application failures, as well as procedures such as rolling starts or upgrades, can be handled automatically and no longer require human intervention. Similarly, the number of instances of applications can be scaled automatically in response to changes in load. Even though there are these clear benefits, a number of new challenges arise, such as how monitoring, logging and in particular service discovery are dealt with in such a dynamic environment where services are no longer tied to specific hosts. In addition, an important question is whether it is even possible to run traditional grid middleware in this type of environment.

This talk will describe the Mesos infrastructure which has been deployed at RAL, the testing we have done, our progress towards migrating production services and discuss our future plans.

Length of presentation (minutes, max. 20):

20 **Summary:**

End-user services and operating systems / 20

Evolution of CERN Printing Services

Authors: Natalie Kane¹ ; Vincent Nicolas Bippus¹

¹ *CERN*

Corresponding Author: natalie.kane@cern.ch

This talk will summarise the evolution of the CERN Print Services and related infrastructure over recent years from both the Service Management and Technical viewpoints. We will discuss some of the issues we have encountered and present the solutions we have found to facilitate the end-user experience of using the Print Services at CERN. This includes streamlining support contracts and lease agreements, optimising costs, enabling easy printer installation, reducing printer numbers and offering a simple print solution for visitors.

Length of presentation (minutes, max. 20):

20 **Summary:**

Grids, clouds, virtualisation / 39

Fermilab HEP Cloud: an elastic computing facility for High Energy Physics.

Author: Anthony Tiradani¹

¹ *Fermilab*

The need for computing in the HEP community follows cycles of peaks and valleys mainly driven by holiday schedules, conference dates and other factors. Because of this, the classical method of provisioning these resources at providing facilities has drawbacks such as potential overprovisioning. As the appetite for computing increases, however, so does the need to maximize cost efficiency by developing a model for dynamically provisioning resources only when needed.

To address this issue, the HEP Cloud project was launched by the Fermilab Scientific Computing Division at Fermilab in June 2015. Its goal is to develop a facility that provides a common interface to a variety of resources, including local clusters, grids, high performance computers, and community and commercial Clouds.

Initially targeted communities include CMS and NOvA, as well as other Fermilab stakeholders. In its first phase, the project has demonstrated the use of the “elastic” provisioning model offered by commercial clouds, such as Amazon Web Services. In this model, resources are rented and provisioned automatically over the Internet upon request. Cost was contained by the use of the Amazon Spot Instance Market, a rental model that allows Amazon to sell their overprovisioned capacity at a fraction of the regular price. Data access was made to scale in terms of volume and cost through a variety of techniques, including autoscaling data caching services.

In January 2016, the project demonstrated the ability to increase the total amount of global CMS resources by 58,000 cores from 150,000 cores - a 25 percent increase. This burst of resources was used in preparation for the Recontres de Moriond conference to generate and reconstruct Monte Carlo events. At the same time, the NOvA experiment has also run data-intensive computations through HEP Cloud, readily provisioning 7,500 cores on Amazon to process Monte Carlo and reconstructed detector data for Neutrino conferences. NOvA is using the same familiar services they use for local computations such as data handling and job submission.

This talk will discuss the architecture used, lessons learned along the way, and some of the next steps in the evolution of the Fermilab HEPCloud Facility.

Length of presentation (minutes, max. 20):

20 **Summary:**

Site reports / 1

Fermilab Site Report

Authors: Anthony Richard Tiradani¹ ; Anthony Tiradani² ; Rennie Scott²

¹ *Fermi National Accelerator Lab. (US)*

² *Fermilab*

Corresponding Authors: rennie@fnal.gov, tiradani@fnal.gov, anthony.tiradani@cern.ch

News and updates from Fermilab since the Fall HEPiX Workshop.

Length of presentation (minutes, max. 20):

20 **Summary:**

Storage and file systems / 52

Fermilab’s Scientific Computing Storage Architecture

Author: Gerard Bernabeu Altayo¹

Co-author: Oliver Gutsche²

¹ *Fermilab*

² *Fermi National Accelerator Lab. (US)*

High Energy Physics experiments record and simulate very large volumes of data and the trend in the future is only going up. All this data needs to be archived, and accessed by central processing workflows as well as a diverse group of scientists to extract physics results. Fermilab supports a wealth of storage technologies for the experiments for very different tasks, from NFS mounted appliances to mass storage systems handling transparent access to files on tape through large disk caches. To prepare for the future, Fermilab recently developed a Scientific Computing Storage Architecture to consolidate and to prepare for a future evolution of the supported storage forms. A big aspect was to remove POSIX access to the storage systems from the worker nodes, necessary to be able to transparently support workflows on different resources like the local batch systems, grids and commercial clouds.

In this presentation, we will discuss the architecture and describe how experts categorizes workloads accessing storage, files and access patterns. We are currently in the process of implementing this strategy together with Fermilab's experiments and projects and will report on the progress.

Length of presentation (minutes, max. 20):

15 **Summary:**

Grids, clouds, virtualisation / 15

First Experiences with Container Orchestration in the CERN Cloud

Author: Bertrand Noel¹

¹ *Ministere des affaires etrangeres et europeennes (FR)*

Corresponding Author: bertrand.noel@cern.ch

We recently added initial container support to the CERN private cloud service. After a brief recap of what container orchestration is, we will discuss what the service offers in terms of cluster managers (Kubernetes, Docker Swarm, Mesos), describe some of the use cases, and show how we integrate with OpenStack and other general CERN services.

Length of presentation (minutes, max. 20):

20 **Summary:**

IT infrastructure / 61

ForHLR - New Energy Efficient HPC System at KIT with Warm Water Cooling

Author: Rudolf Lohner¹

¹ *Karlsruhe Institute of Technology (KIT)*

Corresponding Author: rudolf.lohner@kit.edu

A new HPC-System has been installed at Steinbuch Centre for Computing (SCC) of Karlsruhe Institute of Technology (KIT) delivering about one Petaflops of computing power. For this system a new data center has been built featuring an innovative and very energy efficient warm water cooling. The water temperature level of 40°C inlet and 45°C outlet allows free cooling with dry coolers all over the year even in hot summer conditions as well as reuse of the waste heat for heating office buildings in the colder seasons.

This talk will present the features of this new system, details of the innovative warm water cooling, and further interesting aspects.

Length of presentation (minutes, max. 20):

15

Site reports / 44

GSI Site Report

Author: Walter Schön¹

¹ *GSI*

Corresponding Author: w.schoen@gsi.de

GSI site report

Length of presentation (minutes, max. 20):

20

IT infrastructure / 25

GreenITCube - the new Data Center for FAIR & GSI

Author: Jan Trautmann¹

¹ *GSI Darmstadt*

Corresponding Author: j.trautmann@gsi.de

We will give an overview of the construction phase of the building and will present facts and technical details including the cooling system and function test.

Other topics will be the migration of clusters from the old data center to the GreenITCube and the current status of the infrastructure monitoring.

Length of presentation (minutes, max. 20):

20 **Summary:**

Basic IT services / 36

Grid Computing System in the KEK Central Computer System

Author: Go Iwai¹

¹ *KEK*

Corresponding Author: go.iwai@kek.jp

High Energy Accelerator Research Organization (KEK) plays a key role in particle physics experiments, as well as supporting the communities in Japanese universities. In order to ensure those important missions, KEK has two large-scale computer systems: the Supercomputer System (KEKSC) and the Central Computer System (KEKCC).

The KEKSC is mainly used by collaborative researches in theoretical elementary particle and nuclear physics, condensed matter physics, as well as for accelerator simulations. The system is composed of two different systems: Hitachi SR16000 model M1 (System A) and IBM Blue Gene/Q (System B).

The KEKCC caters to the research demands of particle physics, nuclear physics, the photon factory, neutron science, accelerator development, theory computation, and various fields in the science. In order to meet the demand, KEKCC consists of several subsystems: Data Analysis System, Grid Computing System (EMI/iRODS), and common IT Services such as Mail System, Web System, and so on. Grid Computing System is operated under the Worldwide LHC Computing Grid (WLCG) project. The Belle II, T2K, ILC, and Kagra experiments do their data analysis using the Grid computing infrastructure to manage large amount of experimental data.

The KEKCC is totally replaced every 4 or 5 years according to Japanese government procurement policy for computer system. Current KEKCC has been in operation since April 2012 and will be shutdown in August 2016.

We would like to share our experiences and challenges in the security, the operation, and experiment-specific applications, as well as requirements for storage and computing resources particularly focusing on Grid Computing System through nearly 4 years operation of the current KEKCC. In addition we then discuss future prospects for the next KEKCC system, which will be newly introduced in September 2016.

Length of presentation (minutes, max. 20):

15 **Summary:**

Site reports / 22

HEPiX AGLT2 Site Report Spring 2016

Author: Shawn Mc Kee¹

Co-authors: Philippe Alain Luc Laurens² ; Raymond Brock² ; Robert Ball¹

¹ *University of Michigan (US)*

² *Michigan State University (US)*

Corresponding Author: shawn.mckee@cern.ch

We will present an update on our site since the Fall 2015 report and cover our work with various storage technologies (Lustre, dCache, ZFS and Ceph), ATLAS Muon Calibration, our use of the ELK stack for central syslogging and our experiences with using Check_mk(RAW) as our preferred “OMD” implementation.

We will also report on our recent hardware purchases for 2016 as well as the status of our new networking reconfiguration incorporating a Mellanox SN2700 100G switch. Personnel changes will also be covered. We conclude with a summary of what has worked and what problems we encountered and indicate directions for future work.

Length of presentation (minutes, max. 20):

20

Computing and batch systems / 28

HTCondor European Workshop summary

Author: Michel Jouvin¹

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

Corresponding Author: jouvin@lal.in2p3.fr

The second HTCondor European workshop took place in Barcelona beginning of March (Feb. 29 - March 4). This presentation will present the main topics discussed and the status of the European HTCondor community.

Length of presentation (minutes, max. 20):

20 **Summary:**

Computing and batch systems / 16

HTCondor at DESY

Author: Christoph Beyer¹

Co-author: Thomas Finnern¹

¹ *DESY*

Corresponding Authors: christoph.beyer@desy.de, thomas.finnern@desy.de

After running SOGE, Torque and MYsched for many years DESY HH is preparing to migrate grid and local batch to HTCondor during 2016 in order to benefit from improved reliability and scalability. The talk discusses some essential differences between HTCondor and the queue oriented batchscheduler models, the experience with the running pilot service and the future migration scenario at DESY.

Length of presentation (minutes, max. 20):

0 **Summary:**

Storage and file systems / 33

IBM Spectrum Scale support for technical workflows

Author: Ulf Troppens¹

¹ *IBM*

Corresponding Author: ulf.troppens@gmail.com

IBM Spectrum Scale (formerly known as IBM GPFS) is a feature-rich clustered file system. This talk will cover selected Spectrum Scale features and directions which are in particular relevant for data ingest, data analytics and data management of huge amounts of measured data.

Length of presentation (minutes, max. 20):

20 **Summary:**

Site reports / 49

IRFU Site report

Author: joel surget¹

¹ *CEA/Saclay*

Corresponding Author: joel.surget@cea.fr

Site report of CEA IRFU

Length of presentation (minutes, max. 20):

10

Miscellaneous / 63

Ideas for a journal around "Computing and Software for data-intensive physics"

Author: Yves Kemp¹

¹ *Deutsches Elektronen-Synchrotron (DE)*

Corresponding Author: yves.kemp@cern.ch

In this presentation, we present first ideas about a journal around topics in "Computing and Software for data-intensive physics"

- Why a place for publications?
- For whom to publish?
- Which topics?
- Comparison to other HEP computing related events?
- Who is behind?
- Status?

... waiting for input and ideas from the community - YOU!

Summary:

Length of presentation (minutes, max. 20):

10

Security and networking / 13

Identity Federation for HEP – what are the benefits?

Authors: Dave Kelsey¹ ; Hannah Short² ; Romain Wartel²

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

² *CERN*

Corresponding Author: hannah.short@cern.ch

Activity in the area of Federated Identity Management has been accelerating. 38 national federations have now joined eduGAIN - the interederation service will shortly encircle the globe, facilitating collaboration worldwide. There are clear benefits, but do those benefits outweigh the risks and do they make sense for HEP? We will discuss the need for Federated Identity Management and the progress

made to date within the HEP community. Specific topics will include an overview of eduGAIN, the integration of eduGAIN within WLCG, and security incident response within federations.

Length of presentation (minutes, max. 20):

20 **Summary:**

Site reports / 53

Jefferson Lab Scientific and High Performance Computing

Author: Sandy Philpott^{None}

Corresponding Author: sandy.philpott@jlab.org

JLab high performance and experimental physics computing environment updates since the fall 2015 meeting, including upcoming hardware procurements for Broadwell compute nodes, Pascal and/or Intel KNL accelerators, and Supermicro storage; our Lustre 2.5.3 upgrade; 12GeV computing; and Data Center modernization.

Length of presentation (minutes, max. 20):

15

Site reports / 37

KR-KISTI-GSDC-01 Tier-1 Site Reports

Author: Jeongheon Kim^{None}

Corresponding Author: jkim8104@gmail.com

We will present the latest status of the GSDC. Migration plan of administrative system will be presented.

Length of presentation (minutes, max. 20):

10

Miscellaneous / 65

Logistics

Author: Wolfgang Friebel¹

¹ *Deutsches Elektronen-Synchrotron (DE)*

Corresponding Author: wolfgang.friebel@desy.de

Workshop logistics

Summary:

Length of presentation (minutes, max. 20):

10

Basic IT services / 5**Monitoring at LHCb: Migrating to Icinga2, Puppet, Hiera and Foreman Stack for Monitoring.**

Authors: Balazs Voneki¹ ; Christophe Haen¹ ; Francesco Sborzacchi² ; Hristo Umaru Mohamed³ ; Loic Brarda¹ ; Mohamed Chebbi¹ ; Mohammed Daoudi¹ ; Niko Neufeld¹ ; Rainer Schwemmer¹

¹ CERN

² Istituto Nazionale Fisica Nucleare Frascati (IT)

³ University of Cincinnati (US)

Corresponding Author: hristo.mohamed@cern.ch

The LHCb experiment operates a large computing infrastructure with more than 2000 servers, 300 virtual machines and 400 embedded systems. Many of the systems are operated diskless from NFS or iSCSI root-volumes. They are connected by more than 200 switches and routers. A large fraction of these systems are mission critical for the experiment and as such need to be constantly monitored.

The main part of the monitoring infrastructure is done by tightly integrated instances of Icinga2, Foreman, Hiera and Puppet, which allow for dynamic and automatic generation of configuration files and removal of phased out hosts. We will discuss the steps that were taken and the problems encountered in implementing this integration in an SLC6 dominated environment. We will also touch on our experience with monitoring Windows and FreeNAS hosts, as well as our experience with FreeNAS reports. Furthermore we will show our successful usage of nsca running on SLC6 in our Icinga2 infrastructure.

Summary:

Length of presentation (minutes, max. 20):

15

Basic IT services / 38**Monitoring at scale: a needle in the haystack**

Author: Fabien Wernli¹

¹ CCIN2P3

Corresponding Author: wernli@in2p3.fr

Many of today's opensource monitoring tools have grown to distributed, horizontally scaling solutions.

When designing a new infrastructure, choosing and configuring the right software stack to analyze and record logs and metrics can admittedly still be a challenge, but we are no longer restricted to the vertically scaling rrdtool-type timeseries storage.

The real challenge is the amount of data a monitored system can produce, and the difficulty to process it without classifying and tagging it appropriately.

We explain the necessity to attach relevant metadata to monitored events in order to offer a solution to the needle-in-the-haystack problem

that affects large datasets. Through practical ideas and use-cases at CCIN2P3 we underline the capital importance of metadata to leverage

the power to query the underlying indexing backend. Aggregating metrics and querying logs against technical or business-oriented key-value pairs

are a powerful way to answer questions, and provide high-level alerts.

We present the current solution of managing log events at CCIN2P3 as well as the upcoming metric solution. The primary focus is on the tool stack and experience we gathered during the last decade.

The current monitoring stack is based on facter (puppet), collectd, riemann, syslog-ng and elastic-search and has successfully been used in production at CCIN2P3 on its 2 datacenters with roughly 1500 monitored nodes and 12'000 events per second on average.

Length of presentation (minutes, max. 20):

45 **Summary:**

Site reports / 58

NDGF Site Report

Author: Erik Mattias Wadenstein¹

¹ *University of Umeå (SE)*

Corresponding Author: mattias.wadenstein@cern.ch

Update on recent developments in the NDGF Tier1 site and surroundings.

Length of presentation (minutes, max. 20):

10

IT infrastructure / 57

NERSC collection architecture

Author: Thomas Davis¹

¹ *LBNL/NERSC*

Corresponding Author: tadavis@lbl.gov

An overview of environmental and system information collection at NERSC using virtual machines, containers, python, elasticsearch, logstash, rabbitmq, and web based interfaces.

Some tools that will be covered are elasticsearch, logstash, rabbitmq, kibana, graphana, nagios, librenms, oxidized.

Length of presentation (minutes, max. 20):

20 **Summary:**

Site reports / 69

PDSF Site Report

Author: James Botts¹

¹ *LBNL*

Corresponding Author: jfbotts@lbl.gov

The relocation of PDSF to a new building at LBNL has mostly completed. The lessons learned during the moving process will be described. A new petabyte storage system using EOS has been put on line for the ALICE collaboration. Like many aspects of system administration, deploying new software takes much longer than treading a familiar path and we will describe what we would do for the next deployment.

Length of presentation (minutes, max. 20):

15 **Summary:**

Site reports / 35

RAL Site Report

Author: Martin Bly¹

¹ *STFC-RAL*

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

An update on activities at RAL.

Length of presentation (minutes, max. 20):

15

Security and networking / 24

Recent work of 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 the work of the HEPiX IPv6 working group since the October 2015 HEPiX meeting. Driven by the ATLAS experiment representative, work has included planning for more production dual-stack services to allow for the support of IPv6-only worker nodes/virtual machines in 2017. Guidance for best practices in IPv6 security is also being prepared.

Length of presentation (minutes, max. 20):

20 **Summary:**

Storage and file systems / 21

Running virtualized Hadoop, does it make sense?

Authors: Kacper Surdy¹ ; Zbigniew Baranowski¹

¹ *CERN*

Corresponding Author: kacper.surdy@cern.ch

Public and private clouds based on VMs are a modern approach for deploying computing resources. Virtualisation of computer hardware allows additional optimizations in the utilisation of computing resources compared to the traditional HW deployment model. A price to pay when running virtual machines on physical hypervisors is an additional overhead. This is an area of concern in the context of high throughput computing and big data analytics where distributed data processing frameworks typically push hardware capabilities to their limit.

This presentation reports on our tests and experience with the Hadoop components running on fully virtualized hardware using CERN OpenStack infrastructure. Pros and cons of running Hadoop on VMs vs. physical machines will be discussed as well as performance aspects when running CERN data analytics workloads on a virtual stack.

Length of presentation (minutes, max. 20):

15 **Summary:**

Grids, clouds, virtualisation / 67

SCD Cloud at STFC

Author: Alexander Dibbo¹

¹ *STFC RAL*

Corresponding Author: alexander.dibbo@stfc.ac.uk

An update on the cloud deployment in the Scientific Computing Department at RAL. I will describe our OpenNebula deployment and the use cases we have online including LOFAR. Our OpenNebula deployment has served us well, however new requirements mean that we are looking at OpenStack again. I will describe how we are deploying OpenStack as a replacement for OpenNebula and the work done to get this online.

Length of presentation (minutes, max. 20):

15

End-user services and operating systems / 2

Scientific Linux Update

Authors: Gerard Bernabeu¹ ; Rennie Scott²

Co-authors: Bonnie King² ; Pat Riehecky² ; connie sieh²

¹ *Fermi National Accelerator Lab. (US)*

² *Fermilab*

Corresponding Authors: rennie@fnal.gov, gerard.bernabeu@cern.ch

This talk will present recent updates to Scientific Linux. It will cover the current and future roadmap, new features, and the changes to the customization architectural structure beginning with SL7.2

Summary:

Length of presentation (minutes, max. 20):

20

Security and networking / 47

Security and CSIRT in KEK

Author: Tadashi Murakami¹¹ KEK**Corresponding Author:** tadashi.murakami@kek.jp

The number of cyber security threats are increasing and it gets harder to protect. This presentation will introduce our efforts against the cyber security threats of CSIRT activities, security infrastructures, cultural difficulties, and so on.

Length of presentation (minutes, max. 20):20 **Summary:****Basic IT services / 29**

Server gerontology in LHCb

Authors: Francesco Sborzacchi¹ ; Hristo Umaru Mohamed² ; Joel Closier³ ; Loic Brarda³ ; Mohammed Daoudi³ ; Niko Neufeld³ ; Rainer Schwemmer³¹ *Istituto Nazionale Fisica Nucleare Frascati (IT)*² *University of Cincinnati (US)*³ *CERN***Corresponding Author:** niko.neufeld@cern.ch

In the LHCb Online system we keep systems significantly beyond the warranty period, in some cases up to 7 or more years. We also have upgraded systems in large numbers with third party components (disks for instance). In this contribution give an overview of the various problems we encountered and how we overcome them. We discuss hardware problems, inhouse repairs and related load on the admin team.

Length of presentation (minutes, max. 20):15 **Summary:****Storage and file systems / 45**

Status and new developments for lustre@GSI

Author: Walter Schon^{None}**Corresponding Author:** w.schoen@gsi.de

Status and recent developments for the lustre file system at GSI.
New method to analyse log file, measurements and experience with ZFS
as base system for ZFS and a new project: Interfacing lustre with the
TSM tape robot.

Length of presentation (minutes, max. 20):

20 **Summary:**

Storage and file systems / 31

Status report of TReqS

Author: Bernard CHAMBON¹

¹ *CC-IN2P3*

Corresponding Author: bernard.chambon@cc.in2p3.fr

I will give a status report of TReqS, a software companion of HPSS, the HSM we are using at the CC-IN2P3.

TReqS, which stands for Tape Requests Scheduler, is intended to provide regulation and optimization of the staging requests to HPSS.

TReqS is used at the CC-IN2P3 for several years now from DCACHE and XROOTD, but since fall 2015, we have started a full rewrite of the software, based on a new architecture and implementation with, as main guideline, stability, scalability, new features but also configurable and portable software.

In this presentation I will focus on the benefit of the new version so called TReqS-2, the current status of development (server and client) and the plan for the next months (aim of deployment of production instance for 2016/Q3)

Length of presentation (minutes, max. 20):

20 **Summary:**

Storage and file systems / 17

Storage at CERN: towards the Chamäleon

Author: Xavier Espinal Curull¹

Co-authors: Alberto Pace¹; Dan van der Ster¹; Herve Rousseau¹; Jan Iven¹; Luca Mascetti¹; Massimo Lamanna¹

¹ *CERN*

Corresponding Author: xavier.espinal@cern.ch

Tailoring storage services for the growing community requirements demands high flexibility in our systems. Huge volumes of data coming from the detectors need to be quickly available in a highly scalable mode for data processing and in parallel guarantee high throughput for long term storage. These activities are radically different in terms of storage QoS but all of them are critical to comply with the timings of the experiments' workflows.

Different storage services at CERN cover the needs of our community: EOS and CASTOR as Large Scale Storage Services, CERNBox for community storage, CEPH providing support for Openstack virtual images and attached storage volumes, Filers for very specific filesystem needs and AFS which replacement is being evaluated. Trends in the usage of storage systems show the need for fast adapting to the changing community demands. Behaving more like Chamäleons rather than Elephants.

Length of presentation (minutes, max. 20):

20 **Summary:**

IT infrastructure / 56

The Further Adventures of the NERSC Data Collect

Author: Cary Whitney¹

Co-author: Thomas Davis¹

¹ *LBNL*

Corresponding Author: clwhitney@lbl.gov

Discuss the data pipeline in more details, (logstash, RabbitMQ, collectd, filebeats, Elasticsearch and Kibana). Showing the current data ingest rates and some early results.

Length of presentation (minutes, max. 20):

20 **Summary:**

Grids, clouds, virtualisation / 50

The HELIX NEBULA Science Cloud project

Author: Helge Meinhard¹

¹ *CERN*

Corresponding Author: helge.meinhard@cern.ch

HEP is only one of many sciences with sharply increasing compute requirements that cannot be met by profiting from Moore's law alone. Commercial clouds potentially allow for realising larger economies of scale. While some small-scale experience requiring dedicated effort has been collected, European science has not ramped up to significant scale yet; in addition, public cloud resources have not been integrated yet with the standard workflows of science organisations in their private data centres. The HELIX NEBULA Science Cloud project, partly funded by the European Commission, addresses these points. Ten organisations under CERN's leadership, covering particle physics, bioinformatics, photon science and other sciences, have joined to procure public cloud resources as well as dedicated development efforts towards this integration. The contribution will give an overview of the project, explain the findings so far, and provide an outlook into the future.

Length of presentation (minutes, max. 20):

20 **Summary:**

Storage and file systems / 23

The OSiRIS Project: Meeting the Multi-Institutional Data Collaboration Challenge

Author: Shawn Mc Kee¹

Co-author: Benjeman Jay Meekhof¹

¹ *University of Michigan (US)*

Corresponding Author: shawn.mckee@cern.ch

The OSiRIS (Open Storage Research Infrastructure) project started in September 2015, funded under the NSF CC*DNI DIBBs program. 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 how the OSiRIS project is tackling this challenge using a combination of Ceph, software-defined storage, various open-source management, security and monitoring components and software-defined networking to enable an infrastructure that supports multi-institutional access for scientists working to collaboratively extract scientific results from large, distributed or diverse data. The presentation will cover the current status of OSiRIS as we bring up it's initial implementation, describe the technical details and choices we have made, and summarize our longer term goals and plans for this 5-year project.

Length of presentation (minutes, max. 20):

20 **Summary:**

Storage and file systems / 55

The State of OpenAFS

Author: Stephan Wiesand¹

¹ *Deutsches Elektronen-Synchrotron (DE)*

Corresponding Author: stephan.wiesand@desy.de

What's going on in OpenAFS development, and what are the major challenges, from the Release Manager's perspective.

Length of presentation (minutes, max. 20):

15 **Summary:**

Site reports / 40

The Status of IHEP Site

Author: Jingyan Shi¹

¹ *IHEP*

Corresponding Author: shijy@ihep.ac.cn

The report talks about both hardware and software upgrade that IHEP site has done. It discusses about the problem the site suffered during the last half year. It also gives a brief introduction about the current status of IHEP monitoring system and cloud computing. At last, it shows some new user services the site provided.

Length of presentation (minutes, max. 20):

15

Site reports / 7**The next generation system of KEKCC****Author:** Tomoaki Nakamura¹**Co-authors:** Go Iwai¹ ; Hiroyuki Matsunaga² ; Koichi Murakami ; Takashi Sasaki² ; Wataru Takase¹¹ *KEK*² *High Energy Accelerator Research Organization (J-P)***Corresponding Author:** tomoaki.nakamura@kek.jp

The next generation system of KEK Central Computer (KEKCC) is just in the construction phase striving toward the start of operation from September 2016. In this talk, the detailed configuration of hardware and expected improvement on the performance of new KEKCC will be reported.

Length of presentation (minutes, max. 20):

15

Site reports / 26**US ATLAS SWT2 Site Report****Author:** Horst Severini¹**Co-authors:** Joel Snow² ; Mark Sosebee³ ; Patrick Mcguigan³¹ *University of Oklahoma (US)*² *Langston University*³ *University of Texas at Arlington (US)***Corresponding Author:** hs@nhn.ou.edu

We will present a site report of the US ATLAS SWT2 Computing Center, which consists of UT Arlington, Univ. of Oklahoma, and Langston U. We will give an update on hardware and grid middleware installations at each site, as well as the various opportunistic resources we have available for ATLAS production, and plans for the future.

Length of presentation (minutes, max. 20):

15

Storage and file systems / 11**Update from Database Services****Author:** Katarzyna Maria Dziejniewicz-Wojcik¹¹ *CERN***Corresponding Author:** katarzyna.maria.dziejniewicz@cern.ch

With the terabytes of data stored in relational databases at CERN and great number of critical applications relying on them, the database service is evolving 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 Oracle to make better use of

latest hardware developments. New database management model and technologies (MySQL, PostgreSQL) introduced with Database-On-Demand will also be described. Presenter will be present on Wednesday and Thursday morning.

Length of presentation (minutes, max. 20):

15 **Summary:**

Grids, clouds, virtualisation / 3

Using Containers for HPC Workloads

Author: Fritz Ferstl¹

¹ UNIVA

Containers have quite some history but Docker has helped to make them an exciting trend which has first penetrated DevOps and is now spreading out further in the IT industry. How can containers be utilized in an HPC environment and what benefits can be gained? This paper describes the status quo of container technology, analyzes benefits as well as disadvantages, discusses use case scenarios for HPC and provides detail about integrating container technology with state-of-the-art workload management technology.

Summary:

Length of presentation (minutes, max. 20):

20

Grids, clouds, virtualisation / 51

Using Docker container virtualization in DESY HPC environment

Author: Sergey Yakubov¹

Co-authors: Frank Schluenzen¹ ; Sven Sternberger¹ ; Yves Kemp¹

¹ DESY

Docker container virtualization provides an efficient and, after recent implementation of user namespaces, secure application portability across various environments and operating systems. An application inside a Docker container is packaged with all of its dependencies, has low overhead, can run on any infrastructure, whether it is a single machine, a cluster or a cloud. Container-based applications in high performance computing environment are not fully recognized yet and there is no native solution available.

The present talk describes implementation of container virtualization within DESY HPC cluster, including security issues, high-performance networking and I/O, integration with a resource management system.

Length of presentation (minutes, max. 20):

15 **Summary:**

Grids, clouds, virtualisation / 59

Using Ganeti for running highly available virtualized services

Authors: Erik Mattias Wadenstein¹ ; Ulf Bobson Severin Tigerstedt²

¹ *University of Umeå (SE)*

² *Helsinki Institute of Physics (FI)*

Corresponding Author: mattias.wadenstein@cern.ch

An overview of the virtual server management software stack Ganeti and how it is used at NDGF for running highly available services, like the dCache head nodes for both production and testing, but also some other samples of deployments.

Length of presentation (minutes, max. 20):

20 **Summary:**

Grids, clouds, virtualisation / 41

Virtual Cluster Computing in IHEPCloud

Author: Li Haibo¹

Co-authors: Cheng Yaodong¹ ; Cheng zhenjing¹ ; Cui Tao¹

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

Corresponding Author: lihaibo@ihep.ac.cn

With the rapid growth of high energy physics experimental data, the data processing system encounters many problems such as low resource utilization, migration complex and so on, which makes it urgent to enhance the data analysis system ability. Cloud computing which uses virtualization technology provides many advantages to solve these problems in a cost-effective way.

In this presentation, we will focus on some work we have done on virtual cluster computing. We will discuss the progress of the project, such as the virtual resource quota management and the support of virtual queue management system: VPBS and VCondor.

Length of presentation (minutes, max. 20):

15 **Summary:**

Miscellaneous / 66

Welcome to DESY Zeuthen

Author: Thomas Naumann¹

¹ *Deutsches Elektronen-Synchrotron (DE)*

Corresponding Author: thomas.naumann@desy.de

Welcome talk

Length of presentation (minutes, max. 20):

20

Storage and file systems / 68**Why so Sirius? Ceph backed storage at the RAL Tier-1**

Author: James Adams¹

¹ *STFC RAL*

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

For several years we have been investigating and running Ceph, we have recently reached the point where we are providing production level services underpinned by Ceph and are on the verge of deploying tens of petabytes of Ceph backed storage for large scale scientific data.

I will give an update on the state of our clusters and the various use cases and interfaces we are currently (and planning to) provide.

Length of presentation (minutes, max. 20):

20

Miscellaneous / 60**Workshop wrap-up**

Author: Helge Meinhard¹

¹ *CERN*

Corresponding Author: helge.meinhard@cern.ch

Final talk

Length of presentation (minutes, max. 20):

20

Storage and file systems / 70**ownCloud**

Author: Christian Schmitz¹

¹ *ownCloud Inc*

Corresponding Author: cschmitz@owncloud.com

This talk will provide a strategic outlook around ownCloud in Research and Education. It will start out with an overall ownCloud overview and touch on existing success stories. Furthermore it will focus on federations that allow independent sites to interoperate with regard to cloud-based storage.

Length of presentation (minutes, max. 20):

20 **Summary:**

Security and networking / 30**perfSONAR Status in WLCG/OSG****Author:** Shawn Mc Kee¹**Co-author:** Marian Babik²¹ *University of Michigan (US)*² *CERN***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 information on changes that are included in the recent release of version 3.5.1 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 (minutes, max. 20):20 **Summary:****Computing and batch systems / 14****xBatch: Extending the CERN Batch Service Into the Public Cloud****Author:** Jerome Belleman¹¹ *CERN***Corresponding Author:** jerome.belleman@cern.ch

For the last few years, the CERN Batch Service has been exclusively hosted on our internal cloud service. As procurement for cloud resources to augment the compute available in our computer centre becomes a reality, we are planning the extension of the HTCondor batch service into the public cloud. This talk will provide the initial strategy we are pursuing to configure, provision and manage public cloud resources, the technical and toolsets choice made, and the progress we've made so far.

Length of presentation (minutes, max. 20):

20 **Summary:**