

# HEPiX Fall 2012 Workshop



## Report of Contributions

Contribution ID: 0

Type: **Presentation**

## Fermilab Site Report

*Monday 15 October 2012 11:45 (20 minutes)*

We present recent developments in the Scientific Computing Facilities at Fermilab. We will discuss continued improvements in site networking and wide area networking. We will show significant developments in physical facilities. We will present an overview of major computing and organizational activities in support of the scientific program.

**Author:** TIMM, Steven (Fermilab)**Co-author:** Dr CHADWICK, Keith (Fermilab)**Presenter:** TIMM, Steven (Fermilab)**Session Classification:** Site reports**Track Classification:** Site Reports

Contribution ID: 1

Type: **Presentation**

# The High-Availability FermiCloud Infrastructure-as-a-Service Facility

*Thursday 18 October 2012 09:00 (30 minutes)*

FermiCloud is an Infrastructure-as-a-Service private cloud built for the support of scientific computing at Fermilab. Within the past year we have deployed a facility capable of providing 24x7 service. We will present significant advances in monitoring and visualization, accounting, security, authorization, and user interface. We will also present our current plans for multi-cloud interoperability.

## Summary

FermiCloud has recently updated to a new version of OpenNebula and deployed the X.509 authentication features which were our contribution to OpenNebula. We have split our hardware between two different buildings such that we can survive the failure of either one.

We have deployed new monitoring and visualization of system usage and availability. We have a proof-of-principle of XACML-based callouts for X.509-based authorization as well.

**Author:** TIMM, Steven (Fermilab)

**Presenter:** TIMM, Steven (Fermilab)

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

**Track Classification:** Grid, cloud and virtualization

Contribution ID: 2

Type: **Presentation**

## GSI site report

*Monday 15 October 2012 10:20 (15 minutes)*

GSI site report

**Author:** Dr SCHOEN, Walter (GSI)

**Presenter:** Dr SCHOEN, Walter (GSI)

**Session Classification:** Site reports

**Track Classification:** Site Reports

Contribution ID: 3

Type: **Presentation**

## Scientific Linux current status update

*Tuesday 16 October 2012 10:00 (30 minutes)*

This presentation will provide an update on the current status of Scientific Linux, descriptions for some possible future goals, and allow a chance for users to provide feedback on its direction.

### Summary

This presentation will provide an update on the current status of Scientific Linux, descriptions for some possible future goals, and allow a chance for users to provide feedback on its direction.

**Authors:** RIEHECKY, Pat (Fermilab); SIEH, connie (Fermilab)

**Presenter:** RIEHECKY, Pat (Fermilab)

**Session Classification:** IT Infrastructure

**Track Classification:** IT Infrastructure, Services, Business Continuity

Contribution ID: 4

Type: **Presentation**

## ZNeTS: log your network traffic !

*Wednesday 17 October 2012 09:25 (25 minutes)*

ZNeTS is an acronym for “The Network Traffic Supervisor”. It is a tool for monitoring and recording machines traffic during months.

ZNeTS is a network tool for network introspection and a response to the legal need, in France, to store one year traffic traces.

ZNeTS is very easy to deploy whatever the architecture of your network.

ZNeTS identifies compromised local machines (by virus, trojans, abusive or illegal usage, DNS or Mac spoofing, etc...).

ZNeTS graphical interface is intuitive and ergonomic. Integrated metrology features offer two levels of details. Alerts are simple and relevant.

Over the last 6 month, the tool has been successfully deployed as an appliance into all the IN2P3 laboratories (the french national research institute in physics) and we gave very positive feedbacks from the System administrators.

### Summary

ZNeTS is a powerful and easy tool for monitoring LAN, parts of LAN and machines. It has been developed for the CNRS (the french national research institute).

The purpose is:

- The acquisition and conservation of inbound and outbound network flows (during many months or years... even on very high speed networks)
- Research and filter data, with an integrated search engine
- The detection of anomalies causing the generation of alert and optional email sending
- Metrology, with the calculation and visualization of hourly and daily statistics of the overall traffic and detailed traffic (for each subnet, and machine on the LAN).

ZNeTS is very easy to deploy. It was developed in C++ and includes a web server that implements HTTP/1.1 standard (RFC2616). Authentication by login / password or certificate-X509 is also possible.

The web interface based on the Dojo framework is particularly ergonomic and allows the interpretation and visualization of data. ZNeTS is suitable for all network architectures. It is able to acquire not only the data from NetFlow probe (netflow mode), but also directly from a physical interface (sniffer mode). It is able to decode most versions of NetFlow and IPFIX. It supports IPv4 and IPv6. Packages have been built for most Linux distributions.

ZNeTS is not just a collector. It uses optimized flows that are re-aggregated during a adjustable period of time (from 1 minute to 1 hour). ZNeTS may even send its flows to another collector (and then behaves as a NetflowV9 probe)

The graphical interface is intuitive and ergonomic. Metrology offers two levels of detail. Alerts are simple and relevant (based on counting algorithms). Access to the network flows is easy, the selection forms are pre-filled automatically after each interaction. 2 clicks on charts or on an alert report are enough to see the corresponding flows.

**Author:** Mr DESCOMBES, Thierry (CNRS IN2P3)

**Presenter:** Mr DESCOMBES, Thierry (CNRS IN2P3)

**Session Classification:** Security and Networking

**Track Classification:** Security & Networking

Contribution ID: 5

Type: **Presentation**

## LCG-BEIJING Site Status

*Monday 15 October 2012 09:30 (15 minutes)*

The presentation will show the current status of BEIJING LCG Site and our next plan.

**Author:** Dr SHI, Jingyan (IHEP)

**Presenter:** Dr SHI, Jingyan (IHEP)

**Session Classification:** Site reports

**Track Classification:** Site Reports



Contribution ID: 6

Type: **Presentation**

## Network Traffic Analysis using HADOOP Architecture

*Wednesday 17 October 2012 09:00 (25 minutes)*

This report introduced a network traffic analysis tool using HADOOP architecture. By collecting the traffic information of the egress router in a campus or an institute, the network traffic analysis tool stored the traffic information which includes start time, end time, source IP, destination IP, Byte, Packet, Flow and etc. to HDFS which is a distributed file system as well as the RRD. In the frontend, the tool using rrdtool graph to draw graphs of the network flow trend chart, to get the details of the traffic information just click on the trend chart and there will be a detailed graph on the network flow information drawing by highstock which read data from HDFS, the user can also select a timeslot or a time window to get the netflow information which is calculated by the Map-reduce program running background. Meanwhile, by providing the IP specially related one HEP experiment, the tool can give the traffic information related to the specially HEP experiment, now it is used in collecting the network traffic information of DYB, YBJ, CMS and ATLAS, the network traffic information can be shown in realtime as well as the historic record, and once you put the mouse on the graph, the timeslot and the netflow traffic information will be shown on the graph.

**Author:** Ms ZENG, Shan (IHEP)**Presenter:** Ms ZENG, Shan (IHEP)**Session Classification:** Security and Networking**Track Classification:** Security & Networking

Contribution ID: 7

Type: **Presentation**

## The Lustre file system at IHEP

*Tuesday 16 October 2012 14:00 (30 minutes)*

Lustre has been selected as the main distributed file system solution in IHEP for more than four years. The Lustre File System at IHEP is currently at a scale of 2.2 PB capacity, 50 OSSs and 500+ OSTs, running Lustre 1.8.6. The file system which was built on top of commodity disk arrays, servers and 10 Gbit Ethernet, provides 24 GB/s bandwidth for five high energy physics experiments.

The presentation will mainly report the status of the Lustre file system at IHEP. It includes three parts: 1) an overview of the Lustre File system at IHEP, including the deployment history, current configuration and real performance captured during production usage; 2) the I/O pattern of High energy physics computing, including the file size, the read/write extent size and offset size and performance optimization according to this pattern; 3) management experience abstracted from 4 years' production run.

**Author:** Ms WANG, Lu (IHEP)**Presenter:** Ms WANG, Lu (IHEP)**Session Classification:** Storage and Filesystems**Track Classification:** Storage & Filesystems

Contribution ID: 8

Type: **Presentation**

## OpenStack Chances and Practice in IHEP

*Monday 15 October 2012 14:00 (30 minutes)*

OpenStack is a global collaboration of developers and cloud computing technologists producing the open standard cloud computing platform for both public and private clouds. The project aims to deliver solutions for all types of clouds by being simple to implement, massively scalable, and feature rich. The technology consists of a series of interrelated projects delivering various components for a cloud infrastructure solution. This talk will present the status of the project from the user's and developer's perspectives and also show our practice on using OpenStack to build our cloud computing environment in data center. Then a series of experiments will be conducted and the performance comparison with the current solution will be also demonstrated. Some interesting research and development what we concern on OpenStack will be stated, including private cloud store for HEPiX users, operation monitoring, etc. Finally, the future developments of OpenStack such as auto deployment, network planning and high availability will also be discussed.

**Author:** HAIBO, li**Co-authors:** Dr CHENG, Yaodong (Institute of High Energy Physics, Chinese Academy of Sciences); Mr QINGBAO, hu (IHEP)**Presenters:** Dr CHENG, Yaodong (Institute of High Energy Physics, Chinese Academy of Sciences); Mr QINGBAO, hu (IHEP)**Session Classification:** IT Infrastructure**Track Classification:** IT Infrastructure, Services, Business Continuity

Contribution ID: 9

Type: **Presentation**

## IPv6 deployment in IHEP

*Wednesday 17 October 2012 09:50 (25 minutes)*

Description of the IPV6 deployment in IHEP

### Summary

IHEP has finished the implementation of the IPv6 network infrastructure in the campus network, and the network is running in the dual stack mode. This talk will present the current status of IHEP network and some related system or software deployed in IHEP, including the IPv6 user management system, access control system, etc., the IPv6 related plan, applications and projects in IHEP also will be introduced in this presentation.

**Author:** FAZHI, Qi (IHEP)

**Presenter:** FAZHI, Qi (IHEP)

**Session Classification:** Security and Networking

**Track Classification:** Security & Networking

Contribution ID: 10

Type: **Presentation**

## Bringing cloud storage to your desk with Mucura

*Thursday 18 October 2012 16:00 (30 minutes)*

In this contribution we present our experience building a prototype of an open source software system for operating online file repositories of extensible capacity.

Built on the well-understood client-server architecture model, the system can be used by computing centers looking at solutions for providing online storage services for their individual users. The client-side component runs on the end-user's personal computer and provides both command-line and graphical user interfaces. It supports a deliberately limited set of operations on remote files, namely store, retrieve, organize and share remote files.

Mucura exposes the same HTTP-based standard API supported by Amazon S3 and extends it to support the certificate-based authentication mechanism used by production grid computing platforms such as WLCG. As a consequence, personal file repositories based on Mucura can be seamlessly accessed both from the user's personal computer and from grid jobs running on the user's behalf. This integration allows researchers to use their individual online storage space as a personal storage element conveniently managed from their personal computer.

At the core of the system there are components for managing file metadata and for secure storage of the files' contents, implemented on top of highly available, distributed, persistent and scalable key-value stores.

We will present a detailed architectural view of the system, the status of development and the perspectives for the months to come.

This work is inspired not only by the increasing number of commercial services available nowadays to individuals for their personal storage needs (backup, file sharing, synchronization, ...) such as Amazon S3, Dropbox, SugarSync, bitcasa, etc., but also by several efforts in the same area in the academic and research worlds (NASA, SDSC, etc.). We are persuaded that the level of flexibility offered to individuals by this kind of systems adds value to the day-to-day work of scientists.

**Author:** Mr HERNANDEZ, Fabio (IN2P3/CNRS Computing Center and IHEP Computing Center)

**Presenter:** Mr HERNANDEZ, Fabio (IN2P3/CNRS Computing Center and IHEP Computing Center)

**Session Classification:** Storage and Filesystems

**Track Classification:** Storage & Filesystems

Contribution ID: **11**

Type: **Presentation**

## **CERN site report**

*Monday 15 October 2012 10:00 (20 minutes)*

News from CERN since the previous meeting

**Author:** Dr MEINHARD, Helge (CERN)

**Co-author:** WIEBALCK, Arne (CERN)

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

**Session Classification:** Site reports

**Track Classification:** Site Reports

Contribution ID: 12

Type: **Presentation**

## Virtualisation working group progress report

*Thursday 18 October 2012 11:00 (30 minutes)*

An update on the work of the virtualisation working group.

**Author:** CASS, Tony (CERN)

**Presenter:** CASS, Tony (CERN)

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

**Track Classification:** Grid, cloud and virtualization

Contribution ID: 13

Type: **Presentation**

## DYNES: Building a distributed networking instrument

*Tuesday 16 October 2012 09:00 (30 minutes)*

This presentation will discuss the challenges of efficiently provisioning and deploying switch and host OS configurations enabling our collaboration to monitor, access, and repair the distributed instrument. Additionally we will cover some of the ongoing challenges post-deployment as regards configuration tracking, service verification, and monitoring the overall status of the DYNES instrument as well as enabling DYNES use for domain specific applications (like the LHC).

### Summary

The Dynamic Network System (DYNES) is a distributed networking instrument that creates point-to-point, variable-bandwidth circuits between DYNES sites. This requires hardware and software infrastructure deployment to all participating sites. Currently the deployment stands at approximately 40 sites in the US requiring a storage system, network switch, and IDC (inter-domain-controller) to enable network reservations between sites.

**Author:** MEEKHOF, Benjeman Jay (University of Michigan (US))

**Presenter:** MEEKHOF, Benjeman Jay (University of Michigan (US))

**Session Classification:** IT Infrastructure

**Track Classification:** IT Infrastructure, Services, Business Continuity



Contribution ID: 14

Type: **Presentation**

## Cyber security update

*Thursday 18 October 2012 15:00 (30 minutes)*

This talk gives an update on security trends, and issues affecting computers, software applications and networks during the last months. It includes information on emerging types of vulnerabilities and recent attack vectors, and provides an insight into the cyber-security world of 2012. New security tools developed at CERN will also be presented.

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

**Author:** Mr LOPIENSKI, Sebastian (CERN)

**Presenter:** Mr LOPIENSKI, Sebastian (CERN)

**Session Classification:** Security and Networking

**Track Classification:** Security & Networking

Contribution ID: 15

Type: **Presentation**

## Mobile web development, and CERN mobile web site

*Friday 19 October 2012 10:00 (30 minutes)*

Mobile computing is clearly on the rise –but developing mobile applications, especially for multiple platforms, is a considerable effort. Fortunately, there is an alternative to native apps: web sites that are optimized for mobile devices and touch screens. In this presentation, I will discuss both solutions, and will present a hybrid approach. The presentation will also include a brief introduction to technologies such as jQuery, jQuery Mobile and PhoneGap.

Additionally, as an example of a mobile web application, CERN mobile web site (<http://m.cern.ch>) will be presented.

**Author:** Mr LOPIENSKI, Sebastian (CERN)

**Presenter:** Mr LOPIENSKI, Sebastian (CERN)

**Session Classification:** Miscellaneous

**Track Classification:** Miscellaneous

Contribution ID: **16**

Type: **Presentation**

## Australia Site Report

*Monday 15 October 2012 09:45 (15 minutes)*

Details of the upgrades and changes that have recently been made at the Australian Centre of Excellence for Particle Physics at the Terascale.

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

**Co-author:** CROSBY, Sean Christopher (University of Melbourne (AU))

**Presenters:** BOLAND, Lucien Philip (University of Melbourne (AU)); CROSBY, Sean Christopher (University of Melbourne (AU))

**Session Classification:** Site reports

**Track Classification:** Site Reports

Contribution ID: 17

Type: **Presentation**

## Networking Tools for Sysadmins

*Friday 19 October 2012 11:30 (30 minutes)*

The talk wants to introduce netdisco, which is used at DESY for network inventory and has been enhanced slightly. A command line interface for the netdisco DB developed at DESY should be covered as well.

The concept of using netflow data to understand and analyze network traffic should be discussed, which is especially suitable for sites with high traffic. Programs that can deal with netflow data, such as the recently released version 5 of ntop, nfdump to record netflows and nfsen to visualize the data will be presented, as well as some plugins for nfsen.

All the tools covered in the talk are IPv6 enabled.

**Author:** Dr FRIEBEL, Wolfgang (Deutsches Elektronen-Synchrotron (DE))

**Presenter:** Dr FRIEBEL, Wolfgang (Deutsches Elektronen-Synchrotron (DE))

**Session Classification:** Security and Networking

**Track Classification:** Security & Networking

Contribution ID: **18**Type: **Presentation**

## Setting up the CSP mode in a Gridengine production cluster

*Wednesday 17 October 2012 15:00 (30 minutes)*

All the currently available Gridengine implementations don't provide any authenticated access with the default setup. This opens a big and easily exploitable security hole which might be considered severe especially in multi-community clusters.

This talk will describe in detail the attack vector available in such setups. It will furthermore give a step-by-step guide to activate the certificate-based authentication in Gridengine (the so called "CSP mode") based on the experience at DESY.

**Author:** HAUPT, Andreas (Deutsches Elektronen-Synchrotron (DE))

**Presenter:** HAUPT, Andreas (Deutsches Elektronen-Synchrotron (DE))

**Session Classification:** Computing

**Track Classification:** Computing & Batch Services

Contribution ID: **19**Type: **Presentation**

## CERN Batch System, Monitoring and Accounting

*Tuesday 16 October 2012 11:30 (30 minutes)*

The CERN batch service runs a 60k CPU core cluster using Platform LSF. We present some of the challenges of running a service at this scale, and describe the current planning of how we aim to evolve the current system to a more dynamic, larger scale service.

As part of this, we recently undertook a project of developing new monitoring tools and upgrading the batch accounting system; we present the current state of development in this area.

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

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

**Session Classification:** Computing

**Track Classification:** Computing & Batch Services

Contribution ID: 20

Type: **Presentation**

## CERN Agile Infrastructure , Road to Production

*Monday 15 October 2012 14:30 (30 minutes)*

The CERN Agile Infrastructure aims to redesign the work flow of machine and configuration management within CERN IT. As the AI project approaches production the main software components , OpenStack , Puppet, Foreman have now been deployed with several iterations of scale and stability. We present the current status and next steps for the project.

**Author:** TRAYLEN, Steve (CERN)**Co-author:** VAN ELDIK, Jan (CERN)**Presenter:** TRAYLEN, Steve (CERN)**Session Classification:** IT Infrastructure**Track Classification:** IT Infrastructure, Services, Business Continuity

Contribution ID: 21

Type: **Presentation**

## **Lync - Phone, voice mailbox, instant messaging... Get access to all of them from any place in the world.**

*Monday 15 October 2012 16:30 (30 minutes)*

Use your device (computer or portable device) as your main tool for unified communication. Check presence of your colleagues, make phone calls, get call notifications, listen to your voice mailbox, answer mails and send instant messages. Do it from any place around the world with internet access.

The presentation will summarize our experience in integrating Microsoft Lync, Alcatel PBX and Microsoft Exchange. The goal was to provide a system that integrates VoIP phone/mail/instant messaging and presence to enhance communication capabilities.

**Author:** GRZYWACZEWSKI, Pawel (CERN)

**Presenter:** GRZYWACZEWSKI, Pawel (CERN)

**Session Classification:** IT Infrastructure

**Track Classification:** IT Infrastructure, Services, Business Continuity



Contribution ID: 22

Type: **Presentation**

## Integration Lemon/LAS monitoring with the CERN Agile Infrastructure

*Monday 15 October 2012 15:00 (30 minutes)*

The Agile Infrastructure (AI) project will deliver a solution for the CERN Computer Centre resources management.

Part of the solution will consist in a new monitoring infrastructure of which the LHC Era Monitoring (Lemon) system an early adopters.

Lemon is a client/server based monitoring system, covering performance, application, environment and facilities (e.g. temperature, power consumption, cooling efficiency, etc.) monitoring.

The Lemon Alarming System, a Lemon extension, is used at CERN for notifying the operator about error situations.

This talk covers the migration strategy to the new infrastructure as well as support for the non-Quattor environment (e.g. Puppet).

**Author:** FEDORKO, Ivan (CERN)

**Presenter:** FEDORKO, Ivan (CERN)

**Session Classification:** IT Infrastructure

**Track Classification:** IT Infrastructure, Services, Business Continuity

Contribution ID: 23

Type: **Presentation**

## scientific data cloud infrastructure and services in Chinese Academy of Sciences

*Thursday 18 October 2012 09:30 (30 minutes)*

In order to solve big data challenge in scientific research, a scientific data cloud has been planed to build in Chinese Academy of Sciences, which constitutes 12 data centers and one data archive center, and will provide big data online storage, data backup ,data archive and data intensive analysis services.This talk will introduce the infrastructure, key technology, including distributed file system, virtulization,monitoring, etc, and services.

**Author:** Dr LI, jianhui (no 4, south 4 street, zhongguancun, haidian district, beijing, china)

**Co-authors:** Mr ZHANG, Bo (no 4, south 4th street, zhongguancun, haidian district, beijing china); Dr ZHANG, Haiming (no 4, south 4th street, zhongguancun, haidian district, beijing, china)

**Presenter:** Dr LI, jianhui (no 4, south 4 street, zhongguancun, haidian district, beijing, china)

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

**Track Classification:** Grid, cloud and virtualization

Contribution ID: 24

Type: **Presentation**

## CERN Cloud Storage Evaluation

*Tuesday 16 October 2012 16:00 (30 minutes)*

Currently there is a growing interest in the area of cloud based infrastructures (either private or public) to implement data centres in a more scalable and manageable way and to include external resources in a more flexible way. In this context the CERN DSS group together with participation from IHEP and Huawei have investigated several cloud storage implementation with respect to their stability, performance and scalability.

In this presentation we will summarise the possible motivations to use cloud storage components, their potential roles within the HEP context and the performance results achieved in a PB size test system. We will described the operational experience during several months of test activity and conclude with a tests plan for the next evaluation phase including replication studies between different distributed sites and different storage implementations.

**Author:** DUELLMANN, Dirk (CERN)**Presenter:** DUELLMANN, Dirk (CERN)**Session Classification:** Storage and Filesystems**Track Classification:** Storage & Filesystems

Contribution ID: 25

Type: **Presentation**

## NDGF site report

*Monday 15 October 2012 11:00 (15 minutes)*

Overview of new developments at the distributed NDGF tier-1. New hardware, new organization and new lessons learned will feature.

**Author:** WADENSTEIN, Erik Mattias (Unknown)

**Presenter:** WADENSTEIN, Erik Mattias (Unknown)

**Session Classification:** Site reports

**Track Classification:** Site Reports

Contribution ID: 26

Type: **Presentation**

## Data Center Network changes and extension to Wigner

*Thursday 18 October 2012 14:00 (30 minutes)*

The latest changes on CERN's Data Center Network will be presented, including the migration to High-End Brocade routers, tests and introduction of 100Gbps in the Core of the LCG and bandwidth increase on the firewall system. At the same time, a Network Architecture for the Data Center extension at Wigner will be discussed.

**Authors:** KISHIMOTO BISBE, Carles (CERN); GUTIERREZ RUEDA, David (CERN); MARTELLI, Edoardo (CERN)

**Presenter:** GUTIERREZ RUEDA, David (CERN)

**Session Classification:** Security and Networking

**Track Classification:** Security & Networking

Contribution ID: 27

Type: **Presentation**

## IPv6 deployment status at CERN

*Wednesday 17 October 2012 11:00 (30 minutes)*

An update on the IPv6 latest changes at CERN.

**Authors:** GUTIERREZ RUEDA, David (CERN); MARTELLI, Edoardo (CERN)

**Presenter:** GUTIERREZ RUEDA, David (CERN)

**Session Classification:** Security and Networking

**Track Classification:** Security & Networking

Contribution ID: 28

Type: **Presentation**

## Global Accounting in the Grid and Cloud

*Thursday 18 October 2012 11:30 (30 minutes)*

Running jobs all over the world requires a method of recording and aggregating usage of users and VOs to present the worldwide view of that usage. The APEL Accounting systema has done that successfully for the cpu usage in the worldwide LHC Computing Grid since 2004. This presentation will cover the evolution of APEL and all the other systems who helped collect the data. It will also report on accounting of other types of usage (eg storage), how usage from the cloud can be incorporated and future evolution planned or required.

**Author:** Dr GORDON, John (STFC - Science & Technology Facilities Council (GB))

**Co-authors:** PACKER, Alison (S); ROGERS, Will (STFC)

**Presenter:** Dr GORDON, John (STFC - Science & Technology Facilities Council (GB))

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

**Track Classification:** Grid, cloud and virtualization

Contribution ID: 29

Type: **Presentation**

## Lustre at GSI

*Tuesday 16 October 2012 14:30 (30 minutes)*

Since March 2012, GSI is running a second Lustre file system. The younger cluster introduced a host of new technologies, problems and challenges. Step-by-step migration of both data and hardware from the old to the newer system is under way. This younger installation is the current work horse for GSI HPC, but also provides experience and knowledge base for future projects coming with FAIR. Related is the TeraLink project connecting compute clusters at neighboring sites to GSI's Lustre. A first link has been established to the CSC supercomputer at Frankfurt University.

**Author:** ROTH, Thomas (GSI)**Co-author:** SCHON, Walter**Presenter:** ROTH, Thomas (GSI)**Session Classification:** Storage and Filesystems**Track Classification:** Storage & Filesystems



Contribution ID: 30

Type: **Presentation**

## ITIL at CC-IN2P3

*Friday 19 October 2012 09:00 (30 minutes)*

IN2P3 Computing Center cares about the quality of its services and tries to improve processes and tools using ITIL best practices.

In this talk, I'll describe what we are doing on quality. I'll show the different ongoing work : the ticketing system, the CMDB, the service catalog, the business continuity plan, the identity management, ... .

I'll take some time to go deeper into the change of our ticketing system to OTRS : why change ? what software were evaluated and how ? what changes are expected in daily work ? what more than ticketing could we do with it ?

**Author:** AZEVEDO, Frédéric (CC-IN2P3)

**Presenter:** AZEVEDO, Frédéric (CC-IN2P3)

**Session Classification:** IT Infrastructure

**Track Classification:** IT Infrastructure, Services, Business Continuity

Contribution ID: **31**

Type: **Presentation**

## The HEPiX IPv6 Working Group

*Wednesday 17 October 2012 12:00 (30 minutes)*

This talk will provide an update on the activities of the IPv6 working group since the Prague meeting.

**Author:** Dr KELSEY, David (STFC - Science & Technology Facilities Council (GB))

**Presenter:** Dr KELSEY, David (STFC - Science & Technology Facilities Council (GB))

**Session Classification:** Security and Networking

**Track Classification:** Security & Networking

Contribution ID: 32

Type: **Presentation**

## Federated Identity Management for HEP

*Thursday 18 October 2012 14:30 (30 minutes)*

This talk will present an update on the activities in Federated Identity since the last HEPiX meeting.

**Author:** Dr KELSEY, David (STFC - Science & Technology Facilities Council (GB))

**Presenter:** Dr KELSEY, David (STFC - Science & Technology Facilities Council (GB))

**Session Classification:** Security and Networking

**Track Classification:** Security & Networking

Contribution ID: 33

Type: **Presentation**

## RAL Tier1 Disk only Storage Project status and plans

*Tuesday 16 October 2012 15:00 (30 minutes)*

A working group has been investigating alternatives to Castor for disk only storage at the RAL Tier 1. Requirements have been gathered and we are now deploying test instances of a number of technologies. This talk will discuss both the reasons for the project, requirements and the current status and findings.

**Author:** COLLIER, Ian (UK Tier1 Centre)**Co-author:** DE WITT, Shaun (STFC - Science & Technology Facilities Council (GB))**Presenter:** COLLIER, Ian (UK Tier1 Centre)**Session Classification:** Storage and Filesystems**Track Classification:** Storage & Filesystems

Contribution ID: 34

Type: **Presentation**

## Quattor update - Integrating Aquilon into a grid site

*Monday 15 October 2012 16:00 (30 minutes)*

Aquilon is a Quattor configuration database and management broker developed by an investment to meet the needs of their large worldwide grid. Providing much better relational integrity to the Quattor configuraton database and a workflow that is both more agile and more disciplined, Aquilon can transform the use of Quattor to manage sites.

This talk will discuss RAL Tier 1 experiencing deployed an Aquilon instance outside its original commercial environment and beginning to use it for managing a grid site.

**Author:** COLLIER, Ian (UK Tier1 Centre)

**Presenter:** COLLIER, Ian (UK Tier1 Centre)

**Session Classification:** IT Infrastructure

**Track Classification:** IT Infrastructure, Services, Business Continuity

Contribution ID: 35

Type: **Presentation**

## **STFC Scientific Computing Department Cloud computing**

*Thursday 18 October 2012 12:00 (30 minutes)*

A report on various projects investigating and using cloud computing technologies across the Scientific Computing Department.

**Author:** COLLIER, Ian (UK Tier1 Centre)

**Presenter:** COLLIER, Ian (UK Tier1 Centre)

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

**Track Classification:** Grid, cloud and virtualization

Contribution ID: 36

Type: **Presentation**

## LRMS Migration at GridKa

*Tuesday 16 October 2012 12:00 (30 minutes)*

This talk describes the scheduled migration to another LRMS at GridKa:

- Problems and limitations of the LRMS which is currently used at GridKa
- Selection and tests of a new one
- Configuration details, e.g. fair-share configurations, and experiences with a first sub-cluster which is already managed by the new LRMS

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

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

**Session Classification:** Computing

**Track Classification:** Computing & Batch Services

Contribution ID: 37

Type: **Presentation**

## RAL Site Report

*Monday 15 October 2012 11:15 (15 minutes)*

News from RAL

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

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

**Session Classification:** Site reports

**Track Classification:** Site Reports



Contribution ID: **38**

Type: **Presentation**

## **JASMINE/CEMS and EMERALD**

*Friday 19 October 2012 09:30 (30 minutes)*

Details of the new e-Infrastructure South services at RAL, including the 4.5PB Panasas installation and the GPU service.

**Author:** OLIVER, Peter (STFc-RAL)

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

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

**Session Classification:** IT Infrastructure

**Track Classification:** IT Infrastructure, Services, Business Continuity

Contribution ID: 39

Type: **Presentation**

## The ATLAS Great Lakes Tier-2 (AGLT2) Site Report

*Monday 15 October 2012 11:30 (15 minutes)*

We will present an update on our site since the last report and cover our work with VMware, dCache and perfSONAR-PS. In addition we will discuss our new denser storage system from Dell, recent networking changes and describe how we are integrating these into our site. We will conclude with a summary of what has worked and what problems we encountered and indicate directions for future work.

### Summary

Update on AGLT2 including changes in software, hardware and site configurations and summary of status and future work.

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

**Co-authors:** Mr MEEKHOF, Ben (University of Michigan); LAURENS, Philippe Alain Luc (Michigan State University (US)); BROCK, Raymond (Michigan State University (US)); BALL, Robert (University of Michigan (US)); ROCKWELL, Thomas (Michigan State University (US))

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

**Session Classification:** Site reports

**Track Classification:** Site Reports

Contribution ID: 40

Type: **Presentation**

## Workshop wrap-up

*Friday 19 October 2012 12:00 (30 minutes)*

Closing comments

**Author:** Dr MEINHARD, Helge (CERN)

**Co-author:** PHILPOTT, Sandy (JLAB)

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

**Session Classification:** Miscellaneous

**Track Classification:** Miscellaneous

Contribution ID: 41

Type: **Presentation**

## News from HEPiX Storage Working Group

*Thursday 18 October 2012 16:30 (30 minutes)*

**Presenter:** MASLENNIKOV, Andrei (CASPUR/CINECA)

**Session Classification:** Storage and Filesystems

**Track Classification:** Storage & Filesystems

Contribution ID: 42

Type: **Presentation**

## Service Provisioning and Security Guarantee in CSTNET

*Friday 19 October 2012 11:00 (30 minutes)*

CSTNET is the non-profitable, nationwide academic network in China, with the aim of providing Internet service and applications for the demand of scientific research and constructing an innovative network environment for the future ICT. The service provisioning and security guarantee are key factors to ensure the success of an operational network.

In this talk, we will first give an overview of CSTNET, followed by the introduction of its infrastructure. And then, we will demonstrate the advanced scientific network services and applications provided by CSTNET. The main services include the network management cloud service, network security cloud service, unified communication service, Duckling collaboration working environment service, and network research and experimentation service. The main network applications include light-path provisioning for eVLBI and its tracking for Chang'E-1&2 Lunar Mission, massive data transmission for IHEP-NERSC Daya Bay neutrino experiment, and integrated services for ITER.

Finally, we will present how to guarantee the security in CSTNET, which is provided by the network security infrastructure and the security cloud platform. The network security cloud delivers a clear and achievable path for network administrators using SAAS paradigm to achieve the centralized and unified monitoring and to provide multitenant, on-demand, location-independent network security services. Based on the security facilities and security operating center, both personalized special services and cloud-based general services are provided for 100 institutions in CAS, including security monitoring and situation awareness service, malicious code in-depth analysis service, emergency response on information security service, security assessment and reinforcement service, and safety training service.

**Author:** Dr WU, Yulei (CSTNET)**Presenter:** Dr WU, Yulei (CSTNET)**Session Classification:** Security and Networking**Track Classification:** Security & Networking

Contribution ID: 43

Type: **Presentation**

## Testing SLURM batch system for a grid farm: functionalities, scalability, performance and how it works in a GRID environment

*Wednesday 17 October 2012 16:30 (30 minutes)*

We will show all the work done in order to install and configure the batch system itself together with the security configuration needed. In this presentation we will show the results of the deep testing that we have done on SLURM, in order to be sure that it will cover all the needed functionalities like: priorities, fairshare, limits, QoS, failover capabilities and others. We will report also on the possibility of exploiting this batch system within a complex mixed farm environment where grid job, local job and interactive activities are managed exploiting the same batch system. From a point of view of the scalability we will show how the SLURM batch system is able to deal with the increasing number of node, CPU and jobs served. We will also show the performance achieved with several client accessing the same batch server. We also will make some comparison with other available open source batch system both in terms of performance and functionalities. We will also provide feedback on mixed configuration with SLURM and MAUI as job scheduling. We will also describe the work done in order to support SLURM in a EGI grid environment.

### Summary

As the grid computing farm are increasing in size in terms of nodes but even more in terms of CPU slots available, it become of great interest to have a scheduler solution that could scale up to tens of thousands of CPU slots and hundreds of nodes. In order to try to keep the Total Cost of Ownership as low as possible it will be preferred to have an easy to use and open source solution. SLURM is able to fulfil all those requirements and it looks promising also in terms of community that is supporting it, as it is used in several of the TOP500 supercomputing. For this reason we deeply tested the SLURM batch system in order to prove if it could be a suitable solution. In the work we will present the result of all the test executed on SLURM batch system and the results of the development activity carried on in order to provide the possibility to use SLURM in a grid computing environment.

**Author:** Dr DONVITO, Giacinto (INFN-Bari)

**Presenter:** Dr DONVITO, Giacinto (INFN-Bari)

**Session Classification:** Computing

**Track Classification:** Computing & Batch Services

Contribution ID: 44

Type: **Presentation**

## EGI Federated Cloud Infrastructure

*Thursday 18 October 2012 10:00 (30 minutes)*

Follow up to talk at last HEPiX - describing recent developments & roadmap.

Detailed abstract to follow.

**Author:** Dr TURILLI, Matteo (University of Oxford)

**Co-author:** COLLIER, Ian (UK Tier1 Centre)

**Presenter:** COLLIER, Ian (UK Tier1 Centre)

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

**Track Classification:** Grid, cloud and virtualization

Contribution ID: 45

Type: **Presentation**

## A decade of Condor experience at Fermilab

*Tuesday 16 October 2012 11:00 (30 minutes)*

The Condor Batch System has been used at Fermilab for a decade in the Run II Reprocessing and Analysis, the USCMS Tier 1 facility, and the FermiGrid General Purpose Grid Cluster. In this talk I present an overview of the operational stability, the scalability, and the best practices we have learned to build a 27,000 job slot campus grid using the Condor system.

**Author:** TIMM, Steven (Fermilab)**Presenter:** TIMM, Steven (Fermilab)**Session Classification:** Computing**Track Classification:** Computing & Batch Services



Contribution ID: 46

Type: **Presentation**

## Welcome address

*Monday 15 October 2012 09:00 (20 minutes)*

**Author:** Prof. WANG, Yifang (IHEP)

**Presenter:** Prof. WANG, Yifang (IHEP)

**Session Classification:** Miscellaneous

**Track Classification:** Miscellaneous

Contribution ID: 47

Type: **Presentation**

## Workshop logistics

*Monday 15 October 2012 09:20 (10 minutes)*

**Presenter:** CHEN, Gang (Chinese Academy of Sciences (CN))

**Session Classification:** Miscellaneous

**Track Classification:** Miscellaneous

Contribution ID: 48

Type: **Presentation**

## Oracle Grid Engine at CC-IN2P3 - report after One year

*Wednesday 17 October 2012 14:30 (30 minutes)*

CC-IN2P3 has been running OGE for more than one year now. After describing the current context, I will report the difficulties encountered, solved or not, and the new enhancements we would like to get.

**Author:** OLIVERO, philippe (CC-IN2P3)

**Presenter:** OLIVERO, philippe (CC-IN2P3)

**Session Classification:** Computing

**Track Classification:** Computing & Batch Services

Contribution ID: 49

Type: **Presentation**

## Slurm Experiences for WLCG in the Nordics

*Wednesday 17 October 2012 16:00 (30 minutes)*

Many compute clusters in the nordics run Slurm, this includes the grid connected ones. This talk looks at the experience, which parts works well, what could use improvements, and some comparisons to other batch systems.

**Author:** WADENSTEIN, Erik Mattias (Unknown)**Presenter:** WADENSTEIN, Erik Mattias (Unknown)**Session Classification:** Computing**Track Classification:** Computing & Batch Services

Contribution ID: 50

Type: **Presentation**

## **LAL and GRIF site report**

*Monday 15 October 2012 12:05 (15 minutes)*

LAL and GRIF changes since 1 year

**Author:** JOUVIN, Michel (Universite de Paris-Sud 11 (FR))

**Presenter:** JOUVIN, Michel (Universite de Paris-Sud 11 (FR))

**Session Classification:** Site reports

**Track Classification:** Site Reports

Contribution ID: 51

Type: **Presentation**

## A new shared datacenter in Orsay

Fundamental research labs in southern part of Paris, known as P2IO, have decided to build a common computing facility shared between labs and designed to be energy efficient. This involves the creation of a new datacenter as the first step. This project started one year ago is about to enter its construction phase. This presentation will detail the requirements and technical design and introduce the current ideas for the future computing platform.

**Author:** JOUVIN, Michel (Universite de Paris-Sud 11 (FR))

**Presenter:** JOUVIN, Michel (Universite de Paris-Sud 11 (FR))

**Track Classification:** IT Infrastructure, Services, Business Continuity

Contribution ID: 52

Type: **Presentation**

## Alternatives to Posix. Lessons with S3 Compatible Storage Systems

*Tuesday 16 October 2012 16:30 (45 minutes)*

There are three major areas that of this presentation. First, a trend in storage systems away from tape and Posix to alternative systems that favor scale, latency tolerance, and the integration of storage backup processes including device and system trends. Second, how the features of the OpenStack/Swift and Huawei storage system meet these trends. Third, a little about the testing that has been completed in CERN OpenLab.

**Author:** HUGHES, James (Huawei)**Presenter:** HUGHES, James (Huawei)**Session Classification:** Storage and Filesystems**Track Classification:** Storage & Filesystems

Contribution ID: 53

Type: **Presentation**

## DESY Site Report

*Monday 15 October 2012 12:20 (15 minutes)*

The site report will discuss changes and developments at the Hamburg and Zeuthen sites since Spring.

**Author:** Mr VAN DER REEST, Peter (DESY)

**Presenter:** Mr VAN DER REEST, Peter (DESY)

**Session Classification:** Site reports

**Track Classification:** Site Reports



Contribution ID: 54

Type: **Presentation**

## Selecting a Business-Process-Management-System in conjunction with an Identity-and-Access-Management-System

*Tuesday 16 October 2012 09:30 (30 minutes)*

Business processes are integral parts of every day (non-technical) administrative tasks. Many of these tasks at DESY are still paper-bound. A joint project of DESY-Administration and High-Energy-Physics department was started to provide the organisational and technical prerequisites for establishing electronic workflows by using business process management systems (BPMS) as well as an identity and access management (IAM) system. Thus processes will be handled faster, be traceable better and executed in a uniform manner. The presentation will show aspects of the procedures in choosing a BPM-system as well as an IAM-system and underlying requirements.

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

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

**Session Classification:** IT Infrastructure

**Track Classification:** IT Infrastructure, Services, Business Continuity

Contribution ID: 55

Type: **Presentation**

## CERN and Helix Nebula, the Science Cloud

*Thursday 18 October 2012 17:00 (30 minutes)*

Helix Nebula, the Science Cloud, is a collaborative effort of several European organizations, including CERN, ESA and EMBL, to engage with European industry in public-private partnerships to build a European cloud infrastructure capable of supporting the missions of these organisations.

During the initial pilot phase of Helix Nebula, the ATLAS experiment at CERN was selected as one of the flagship projects and a proof-of-concept phase was defined in order to demonstrate the feasibility of integrating commercial cloud facilities into the ATLAS distributed computing infrastructure.

This talk will outline the status of Helix Nebula and present the results of the ATLAS use-case in particular. Three commercial cloud providers with varied infrastructures were tested; all were successfully able to run ATLAS simulation jobs, though the paths to success at some providers were more difficult than others. We will give an insight into the lessons learned, the technical recommendations for the supply side and some of the future work in the Helix Nebula partnership.

**Authors:** Dr VAN DER STER, Daniel (CERN); BARREIRO MEGINO, Fernando Harald (CERN); MEDRANO LLAMAS, Ramon (CERN); Dr WALKER, Rodney (Ludwig-Maximilians-Univ. Muenchen (DE))

**Presenter:** BARREIRO MEGINO, Fernando Harald (CERN)

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

**Track Classification:** Grid, cloud and virtualization

Contribution ID: 56

Type: **Presentation**

## Site Report of ASGC

*Wednesday 17 October 2012 10:15 (15 minutes)*

Update of e-Science infrastructure at ASGC will be reported. Focus will cover the international networking, computing and storage infrastructure, overall continuous operation improvement including the data center, e-Science application and virtual research framework, as well as the development of distributed cloud.

**Author:** Mr ERIC, YEN (ASGC)**Presenter:** Mr ERIC, YEN (ASGC)**Session Classification:** Site reports**Track Classification:** Site Reports

Contribution ID: 57

Type: **Presentation**

## Scientific Linux Infrastructure Improvements

*Monday 15 October 2012 17:00 (30 minutes)*

The underlying infrastructure of Scientific Linux is starting to change. These changes should make the environment more stable and higher performing with a greater feature set. This presentation will detail some of the plans and progress thus far.

**Authors:** RIEHECKY, Pat (Fermilab); SIEH, connie (Fermilab)

**Presenter:** RIEHECKY, Pat (Fermilab)

**Session Classification:** IT Infrastructure

**Track Classification:** IT Infrastructure, Services, Business Continuity

Contribution ID: 58

Type: **Presentation**

## FZU IPv6 testbed updates

*Wednesday 17 October 2012 11:30 (30 minutes)*

we present updates since Vancouver about our IPv6 testbed at FZU. We have setup nagios, smokeping several middleware services and several computing centre management procedures.

**Author:** ELIAS, Marek (Institute of Physics AS CR (FZU))

**Co-authors:** CHUDoba, Jiri (Institute of Physics AS CR (FZU)); FIALA, Lukas (Institute of Physics AS CR (FZU)); KOUBA, Tomas (Institute of Physics AS CR (FZU))

**Presenter:** ELIAS, Marek (Institute of Physics AS CR (FZU))

**Session Classification:** Security and Networking

**Track Classification:** Security & Networking

Contribution ID: 59

Type: **Presentation**

## Report from batch system BOF

*Wednesday 17 October 2012 14:00 (30 minutes)*

Report from the BOF session the day before

**Author:** VEDAEE, Aresh (CC-IN2P3 - Centre de Calcul (FR))

**Presenter:** VEDAEE, Aresh (CC-IN2P3 - Centre de Calcul (FR))

**Session Classification:** Computing

**Track Classification:** Computing & Batch Services