

# HEPiX Fall 2007



## Report of Contributions

Contribution ID: 0

Type: **not specified**

## Scientific Linux Status Report

*Friday, 9 November 2007 09:00 (30 minutes)*

Progress of Scientific Linux over the past 6 months.  
What we are currently working on.  
What we see in the future for Scientific Linux

**Primary author:** DAWSON, Troy (FNAL)

**Co-author:** SIEH, Connie (FNAL)

**Presenter:** DAWSON, Troy (FNAL)

**Session Classification:** General II

**Track Classification:** General

Contribution ID: 1

Type: **not specified**

## Large-scale remote management via IPMI

*Thursday, 8 November 2007 16:30 (30 minutes)*

In the recent years the Intelligent Platform Management Interface (IPMI) has matured and become widespread enough that even commodity hardware can now be reliably and securely managed using its functionality.

We report on the successful large-scale deployment of secure remote power control of machines and the integration of IPMI Serial-Over-LAN functionality into the existing remote console infrastructure.

**Primary author:** Mr HORVATH, Andras (CERN)

**Presenter:** Mr HORVATH, Andras (CERN)

**Session Classification:** Data Centers III

**Track Classification:** Data Centers

Contribution ID: 3

Type: **not specified**

## **DAPNIA Site report**

*Monday, 5 November 2007 09:30 (20 minutes)*

What is new since the last HEPiX fall meeting in DAPNIA Saclay.

**Primary author:** Mr MICOUT, Pierrick (CEA DAPNIA Saclay)

**Presenter:** Mr MICOUT, Pierrick (CEA DAPNIA Saclay)

**Session Classification:** Site Reports I

**Track Classification:** Site Reports

Contribution ID: 4

Type: **not specified**

## How INFN is moving out SI2K

*Thursday, 8 November 2007 12:05 (25 minutes)*

INFN Computing Committee is looking for a modern benchmark to replace the SPEC Int 2000 for Worker Node evaluation. Performance results from several HEP experiments have been compared with SPEC Int 2000 and Int 2006 to find the best agreement.

**Primary author:** Dr MICHELOTTO, michele (INFN Padova)

**Presenter:** Dr MICHELOTTO, michele (INFN Padova)

**Session Classification:** Benchmarking I

**Track Classification:** Benchmarking

Contribution ID: 5

Type: **not specified**

## SLAC Site Report

*Monday, 5 November 2007 09:50 (20 minutes)*

SLAC Site Report

**Primary author:** Mr BOEHEIM, Charles (SLAC)

**Presenter:** Mr BOEHEIM, Charles (SLAC)

**Session Classification:** Site Reports I

**Track Classification:** Site Reports

Contribution ID: 6

Type: **not specified**

## Thinking Inside the Box

*Wednesday, 7 November 2007 12:00 (30 minutes)*

SLAC recently acquired a Sun Project Black Box for expansion of its compute clusters. I will describe our experiences with installing and operating the box and its contents, and go over some of the economic tradeoffs in expansion strategies.

**Primary author:** Mr BOEHEIM, Charles (SLAC)

**Presenter:** Mr BOEHEIM, Charles (SLAC)

**Session Classification:** Data Centers II

**Track Classification:** Data Centers

Contribution ID: 7

Type: **not specified**

## High Density Visualizations

*Friday, 9 November 2007 10:00 (30 minutes)*

“GUIs don’t scale” has long been one of my mantras. I will demonstrate some recent work I have done in trying to overturn this limitation. By using very high density display techniques such as treemapping and pixel charts, I have been able to create some displays that show the status and history of thousands of machines without abstracting all of the meaningful detail out of them.

**Primary author:** Mr BOEHEIM, Charles (SLAC)

**Presenter:** Mr BOEHEIM, Charles (SLAC)

**Session Classification:** General II

**Track Classification:** General



Contribution ID: 8

Type: **not specified**

## **LAL / GRIF Site report**

*Monday, 5 November 2007 10:10 (20 minutes)*

Site report about LAL and GRIF (LCG T2)

**Primary author:** Mr JOUVIN, Michel (LAL / IN2P3)

**Presenter:** Mr JOUVIN, Michel (LAL / IN2P3)

**Session Classification:** Site Reports I

**Track Classification:** Site Reports

Contribution ID: 9

Type: **not specified**

## Experience with Integrated Site Security (ISS)

*Friday, 9 November 2007 11:00 (30 minutes)*

Integrated Site Security for Grids (ISSeG) is a project co-funded by the European Commission involving CERN, STFC and FZK. Integrated Site Security (ISS) coordinates improvements so that policies, rules, awareness and training evolve in step with technological developments.

This session will cover the resources available on [www.isseg.eu](http://www.isseg.eu):

- Risk assessment questionnaire
- Security recommendations
- Checklist for system administrators
- Checklist for developers
- Training/advice for general users
- Advice and material for managers

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

20-30

**Primary author:** Mr SILVERMAN (SPEAKER), Alan (CERN)

**Co-author:** ISSEG PROJECT (AUTHOR), Various team members (CERN)

**Presenter:** Mr SILVERMAN (SPEAKER), Alan (CERN)

**Session Classification:** General III

**Track Classification:** General

Contribution ID: **10**

Type: **not specified**

## **SL Plenary Discussion**

*Friday, 9 November 2007 09:30 (30 minutes)*

Feedback to and input for the SL developers from the HEPiX community. This may influence upcoming decisions e.g. on distribution lifecycles.

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

30

**Primary author:** DAWSON, Troy (FNAL)

**Presenter:** DAWSON, Troy (FNAL)

**Session Classification:** General II

**Track Classification:** General

Contribution ID: 11

Type: **not specified**

## TRIUMF Site Report

*Monday, 5 November 2007 11:00 (20 minutes)*

TRIUMF site report

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

15

**Primary author:** Dr MCDONALD, Steven (TRIUMF)

**Presenter:** Dr MCDONALD, Steven (TRIUMF)

**Session Classification:** Site Reports II

**Track Classification:** Site Reports

Contribution ID: 12

Type: **not specified**

## **RAL Tier-1 Site Report**

*Monday, 5 November 2007 11:20 (20 minutes)*

Developments in the Tier-1 at RAL.

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

15

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

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

**Session Classification:** Site Reports II

**Track Classification:** Site Reports

Contribution ID: 13

Type: **not specified**

## Expansion plans for the Brookhaven computer center

*Wednesday, 7 November 2007 09:30 (30 minutes)*

Brookhaven is the primary computing center for RHIC, and it is also a Tier 1 computing center for the ATLAS experiment at CERN. The growing computing needs from both RHIC and ATLAS has exceeded the capacity of the current computer center, and a long-term expansion plan is now underway to meet future requirements. This presentation discusses our expansion plans as well as possible strategies that emphasize efficient growth of the computing facility resources.

### **Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

30 minutes

**Primary authors:** Dr CHAN, Antonio (Brookhaven National Lab); ERNST, Michael (Brookhaven National Lab)

**Presenter:** Dr CHAN, Antonio (Brookhaven National Lab)

**Session Classification:** Data Centers I

**Track Classification:** Data Centers

Contribution ID: 14

Type: **not specified**

## **Grid Security Update**

*Thursday, 8 November 2007 11:00 (30 minutes)*

This talk will describe recent progress on Grid Security in WLCG and EGEE covering technical, operational and policy issues.

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

30 minutes

**Primary author:** Dr KELSEY, David (RAL)

**Presenter:** Dr KELSEY, David (RAL)

**Session Classification:** Grid Management and Monitoring II

**Track Classification:** Grid Management and Monitoring

Contribution ID: 15

Type: **not specified**

## Windows Vista and Office 2007 deployments at CERN

*Monday, 5 November 2007 14:00 (30 minutes)*

Windows Vista has been released by Microsoft early this year and Office 2007 even at the end of the last year.

First, in March 2007, Office 2007 pilot has been announced at CERN. This new Office suite version is now pre-installed on each new computer at CERN. In addition users can upgrade to Office 2007 their existing PCs. With such deployment schema we currently have Office 2007 installed on 15% of Windows PCs.

Then, in August, Windows Vista pilot started. It took us some time and effort to test this new operating system and tune to our needs before we could give it to the end users. Currently we have more than 100 Vista computers in production at CERN.

### **Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

30

**Primary author:** Mr OTTO, Rafal (CERN)

**Presenter:** Mr OTTO, Rafal (CERN)

**Session Classification:** Windows I

**Track Classification:** Windows



Contribution ID: 16

Type: **not specified**

## **CERN Alerter – RSS based system for the information broadcast to all CERN offices**

*Friday, 9 November 2007 11:30 (30 minutes)*

Nearly every large organization uses a tool to broadcast messages and information across the internal campus (messages like alerts announcing interruption in services or just information about upcoming events). These tools typically allow administrators (operators) to send “targeted” messages which are sent only to specific groups of users or computers, e/g only those located in a specified building or connected to a particular computing service.

CERN has a long history of such tools. The last one - NICE Alerter used on all Windows-based computers had to be phased out as a consequence of phasing out NNTP at CERN. The new solution continues to provide the service based on cross-platform technologies, hence minimizing custom developments and relying on commercial software as much as possible. The new system, called CERN Alerter, is based on RSS (Really Simple Syndication) for the transport protocol and uses Microsoft SharePoint as the backend for database and posting interface. The windows-based client relies on Internet Explorer 7.0 with custom code to trigger the window pop-ups and the notifications for new events. Linux and Mac OS X clients could also rely on any RSS readers to subscribe to targeted notifications. The presentation will cover the architecture and implementation aspects of the new system.

### **Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

30

**Primary author:** Mr OTTO, Rafal (CERN)**Presenter:** Mr OTTO, Rafal (CERN)**Session Classification:** General III**Track Classification:** General

Contribution ID: 17

Type: **not specified**

## New WLCG Grid Monitoring Displays

*Thursday, 8 November 2007 09:30 (30 minutes)*

We will demonstrate two new tools for monitoring grid services within WLCG. The first is a set of grid probes, designed to test EGEE and OSG services at a site, integrated into nagios. The second is a new visualisation tool - GridMap (<http://gridmap.cern.ch>). This uses heatmaps to compactly display site availability for the entire WLCG Grid on a single page.

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

30

**Primary author:** CASEY, James (CERN)

**Presenter:** CASEY, James (CERN)

**Session Classification:** Grid Management and Monitoring I

**Track Classification:** Grid Management and Monitoring

Contribution ID: 18

Type: **not specified**

## Running EGEE services and worker nodes using virtual machines

*Wednesday, 7 November 2007 15:30 (20 minutes)*

New scientific communities, attracted by successful deployment of Grids by various scientific groups, bring new requirements on Grid middleware, such as fast turn-around for short jobs, access prioritisation, and need for various execution environment optimized to application specific requirements.

We will argue that virtualized Grids, based on idea that different virtual worker nodes, running on and sharing physical cluster nodes, may help to implement some of these features. In this presentation, we will present our effort in virtualization of the national grid infrastructure

in Czech Republic and we will describe a system called Magrathea, which we have developed to allow Grid job scheduling system to deal with several virtual machines (VMs) running concurrently on a single computer.

With help of two HEP and EGEE related deployment scenarios (service consolidation of EGEE services and sharing of worker nodes between EGEE and MetaCentre grids) we will demonstrate how virtual machines can be already used to fulfill some of new requirements.

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

15

**Primary author:** Mr RUDA, Miroslav (CESNET)

**Co-authors:** SVEC, Jan (CESNET); Dr MATYSKA, Ludek (CESNET)

**Presenter:** SVEC, Jan (CESNET)

**Session Classification:** Virtualization II

**Track Classification:** Virtualization

Contribution ID: 19

Type: **not specified**

## **CPU Benchmarking at GridKa - Update November 2007**

*Thursday, 8 November 2007 13:30 (30 minutes)*

I'll continue with my series of talks about CPU benchmarking issues.

New topics are:

- SPECint2006 measurements, what's the difference to SPECint2000?
- SPECfp2006 measurements.

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

30

**Primary author:** ALEF, Manfred (Forschungszentrum Karlsruhe)

**Presenter:** ALEF, Manfred (Forschungszentrum Karlsruhe)

**Session Classification:** Benchmarking II

**Track Classification:** Benchmarking

Contribution ID: 20

Type: **not specified**

## INFN Site Report

*Monday, 5 November 2007 11:40 (20 minutes)*

an overall report about updates of INFN computing infrastructure

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

15

**Primary author:** Dr GOMEZEL, Roberto (INFN)

**Presenter:** Dr GOMEZEL, Roberto (INFN)

**Session Classification:** Site Reports II

**Track Classification:** Site Reports

Contribution ID: 21

Type: **not specified**

## HA and Virtualization team @ INFN

*Wednesday, 7 November 2007 14:10 (20 minutes)*

Virtualization can enhance the functionality and ease the management of current and future Grids by enabling on-demand creation of services and virtual clusters with customized environments, dynamic provisioning and policy-based resource allocation, as well as enhance high availability and load balancing techniques.

In this work, we consider the work done in the last year in different INFN sections in both fields. We present a summary of the activity reports sent by each team.

At first we provide statistics and metrics of virtualization solutions for Tier-2 core services implemented in production; then a prototype of the use of provisioning system in a Grid data-center environment, allowing for classic OS provisioning, virtual machine partitioning, embedded monitoring and clustered configuration management; at last a prototype of highly available and load balanced Grid services on a virtualized architecture.

Most of these activities include advanced usage and customization of mostly open-source tools as Xen para-virtualizer, Nagios monitoring system, Cfengine and Puppet configuration management systems, SAN implementations through iSCSI and AoE (both software and hardware based).

### **Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

20

**Primary author:** Mr NEBIOLO, Federico (INFN)

**Presenter:** Mr NEBIOLO, Federico (INFN)

**Session Classification:** Virtualization I

**Track Classification:** Virtualization

Contribution ID: 22

Type: **not specified**

## Monitoring and Metrics within FermiGrid

*Thursday, 8 November 2007 09:00 (30 minutes)*

We will present the tools which are used to monitor the grid services within the Fermilab Campus Grid (FermiGrid).

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

30

**Primary author:** Dr CHADWICK, Keith (Fermilab)

**Presenter:** Dr CHADWICK, Keith (Fermilab)

**Session Classification:** Grid Management and Monitoring I

**Track Classification:** Grid Management and Monitoring

Contribution ID: 23

Type: **not specified**

## FermiGrid-HA

*Wednesday, 7 November 2007 15:50 (20 minutes)*

We will discuss the deployment of Linux Virtualisation within the Fermilab Campus Grid (Fermi-Grid) infrastructure to provide highly available Grid services.

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

30

**Primary author:** Dr CHADWICK, Keith (Fermilab)

**Presenter:** Dr CHADWICK, Keith (Fermilab)

**Session Classification:** Virtualization II

**Track Classification:** Virtualization



Contribution ID: 24

Type: **not specified**

## The Genome Sequencing Center's Data Center Plans

*Wednesday, 7 November 2007 10:00 (30 minutes)*

The Washington University School of Medicine is building a new data center with the needs of the Genome Sequencing Center in mind. It is being constructed to have full power and cooling redundancy for 120 racks of equipment, with an anticipated mixture of 2/3 8 kW storage racks and 1/3 25 kW compute racks. We will discuss its design, the problems we encountered along the way, our current thinking on delivering power to our racks and our best projections for how we will grow into that space.

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

30

**Primary author:** STIEHR, Gary (The Genome Sequencing Center at Washington University in St. Louis)

**Presenter:** STIEHR, Gary (The Genome Sequencing Center at Washington University in St. Louis)

**Session Classification:** Data Centers I

**Track Classification:** Data Centers

Contribution ID: 25

Type: **not specified**

## Introduction

*Wednesday, 7 November 2007 13:30 (10 minutes)*

**Primary author:** Dr LEFEBURE, Veronique (CERN)

**Presenter:** Mrs LEFEBURE, Veronique (CERN Staff)

**Session Classification:** Virtualization I

**Track Classification:** Virtualization

Contribution ID: 26

Type: **not specified**

## **Operational issues related to the usage of Virtualisation**

*Wednesday, 7 November 2007 13:40 (30 minutes)*

**Primary author:** Mr IRIBARREN, Alex (CERN)

**Presenter:** Mr IRIBARREN, Alex (CERN)

**Session Classification:** Virtualization I

Contribution ID: 27

Type: **not specified**

## GridKa Site Report

*Monday, 5 November 2007 12:00 (20 minutes)*

Brief presentation of the current status at the Grid Computing Centre Karlsruhe.

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

15

**Primary author:** ALEF, Manfred (Forschungszentrum Karlsruhe)

**Presenter:** ALEF, Manfred (Forschungszentrum Karlsruhe)

**Session Classification:** Site Reports II

**Track Classification:** Site Reports

Contribution ID: 29

Type: **not specified**

## **Virtualisation Infrastructure in Karlsruhe (University & Research Center)**

*Wednesday, 7 November 2007 14:30 (25 minutes)*

**Primary author:** Mr OBERST, Oliver (FZK)

**Presenter:** Mr OBERST, Oliver (FZK)

**Session Classification:** Virtualization I

Contribution ID: 32

Type: **not specified**

## Virtualisation usage at CERN

*Wednesday, 7 November 2007 16:10 (20 minutes)*

One may say that virtualisation techniques were present since some years. But this year it seems to be that virtualisation surged ahead of “Web 2.0” as the buzziest of buzzwords in the sector of information technology.

In the first half of March 2007 the FIO group launched its virtualisation project. The main object is to use the availability of cheaper and more powerful machines to consolidate low utilisation services. Besides this, there is an increasing demand for test and development machines, which are usually not heavily loaded and thus could be easily virtualised.

Half a year later I'd like to share experiences the project team made during the integration of a Xen based virtual machine environment into the existing infrastructures of CERNs Computing Centre. This covers issues regarding configuration, installation, monitoring, managing and modelling of the virtual machine environment.

Since FIO is not the only group using virtualisation technologies I will also report about two other use cases. One is the Certification and Test of the gLite middleware within the EGEE project. The other one is the ETICS project which uses virtual machines in their build and test system.

**Primary author:** Mr MICHAEL, Jan

**Co-authors:** Mr UNTERKIRCHER, Andreas (CERN); Dr ZUREK, Marian (CERN)

**Presenter:** Mr MICHAEL, Jan

**Session Classification:** Virtualization II

Contribution ID: 33

Type: **not specified**

## Virtualisation with the Globus Toolkit

*Wednesday, 7 November 2007 16:30 (20 minutes)*

**Primary author:** Dr KEAHAY, Kate

**Presenter:** Dr KEAHAY, Kate

**Session Classification:** Virtualization II

Contribution ID: 34

Type: **not specified**

## Application level monitoring for the LHC experiments with the Experiment Dashboard monitoring system

*Thursday, 8 November 2007 10:00 (30 minutes)*

The presentation will give an overview of the Experiment Dashboard monitoring applications, providing LHC users with the possibility to monitor analysis and production activities on the distributed infrastructure. Experiment Dashboard monitoring applications are working in a transparent way across several middleware platforms used by the LHC experiments (LCG/gLite,OSG,NorduGrid). The presentation will cover production monitoring for ATLAS and CMS experiments and task monitoring for the CMS analysis users.

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

30

**Primary authors:** Dr GAIDIOZ, Benjamin (CERN); Mrs MAIER, Gerhild (CERN); SIDOROVA, Irina (CERN); Mrs ANDREEVA, Julia (CERN); Mr SAIZ, Pablo (CERN); Mr BRITO DA ROCHA, Ricardo (CERN); Mr WAKEFIELD, Stuart (Imperial College)

**Presenter:** Dr GAIDIOZ, Benjamin (CERN)

**Session Classification:** Grid Management and Monitoring I

**Track Classification:** Grid Management and Monitoring



Contribution ID: 35

Type: **not specified**

## DESY Site Report

*Tuesday, 6 November 2007 09:00 (20 minutes)*

DESY site report - Update on recent activities at DESY

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

15

**Primary author:** FRIEBEL, Wolfgang (DESY)

**Presenter:** FRIEBEL, Wolfgang (DESY)

**Session Classification:** Site Reports III

**Track Classification:** Site Reports

Contribution ID: 36

Type: **not specified**

## NDGF Site report

*Tuesday, 6 November 2007 09:20 (20 minutes)*

Current status at the distributed tier1 “site” NDGF.

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

15

**Primary author:** WADENSTEIN, Mattias (NDGF)

**Presenter:** WADENSTEIN, Mattias (NDGF)

**Session Classification:** Site Reports III

**Track Classification:** Site Reports

Contribution ID: 37

Type: **not specified**

## Procuring a New Machine Room

*Wednesday, 7 November 2007 11:30 (30 minutes)*

We describe the processes surrounding the procurement of a new machine room for the RAL site.

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

25

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

**Co-author:** Dr ROBINSON, Graham (STFC - RAL)

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

**Session Classification:** Data Centers II

**Track Classification:** Data Centers

Contribution ID: 38

Type: **not specified**

## NERSC Data Center Presentation

*Wednesday, 7 November 2007 11:00 (30 minutes)*

The abstract will be updated ASAP.

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

30

**Primary author:** Dr SAKREJDA, Iwona (LBNL/NERSC)

**Co-author:** Mr WHITNEY, Cary (LBNL/NERSC)

**Presenter:** Dr SAKREJDA, Iwona (LBNL/NERSC)

**Session Classification:** Data Centers II

**Track Classification:** Data Centers

Contribution ID: 39

Type: **not specified**

## **PDSF Site Report.**

*Tuesday, 6 November 2007 09:40 (20 minutes)*

Recent changes in PDSF organization. Benefits and pitfalls of a shared resource.

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

15

**Primary author:** Dr SAKREJDA, Iwona (LBNL/NERSC)

**Presenter:** Dr SAKREJDA, Iwona (LBNL/NERSC)

**Session Classification:** Site Reports III

**Track Classification:** Site Reports

Contribution ID: 40

Type: **not specified**

## Introduction & Results of HEPiX Survey

*Wednesday, 7 November 2007 09:00 (30 minutes)*

I will give a general introduction to the problems faced by today's data centres and summarise the responses to the survey of HEPiX sites carried out prior to this meeting.

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

30

**Primary author:** Dr CASS, Tony (CERN)

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

**Session Classification:** Data Centers I

**Track Classification:** Data Centers

Contribution ID: 41

Type: **not specified**

## Jefferson Lab Site Report

*Tuesday, 6 November 2007 10:00 (20 minutes)*

Update on JLab computing since the last HEPiX meeting.

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

15

**Primary author:** PHILPOTT, Sandy (JLab)

**Presenter:** PHILPOTT, Sandy (JLab)

**Session Classification:** Site Reports III

**Track Classification:** Site Reports

Contribution ID: 42

Type: **not specified**

## **GSI site report**

*Tuesday, 6 November 2007 11:00 (20 minutes)*

Latest IT developments at GSI

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

15

**Primary author:** Dr SCHOEN, Walter (GSI)

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

**Session Classification:** Site Reports IV

**Track Classification:** Site Reports



Contribution ID: 43

Type: **not specified**

## lustre at GSI

*Monday, 5 November 2007 16:30 (40 minutes)*

Evaluation of a clustre file system for the GSI data file system

- architecture of the test setup
- RAID controller/disk array performance
- OST I/O performance
- client I/O performance
- clustre file system performance
- trunking/bonding

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

40 minutes

**Primary author:** Dr SCHOEN, Walter (GSI)

**Co-authors:** Mr HUHN, Christo (GSI); Dr MIERS, Karin (GSI); Dr HALLER, Stefan (GSI); Dr ROTH, Thomas (GSI)

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

**Session Classification:** General I

**Track Classification:** Storage

Contribution ID: 44

Type: **not specified**

## CERN site report

*Tuesday, 6 November 2007 11:20 (30 minutes)*

News on CERN

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

30 minutes

**Primary author:** Dr MEINHARD, Helge (CERN-IT)

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

**Session Classification:** Site Reports IV

**Track Classification:** Site Reports

Contribution ID: 45

Type: **not specified**

## Introduction to benchmarking session

*Thursday, 8 November 2007 12:00 (5 minutes)*

Just a few words to get the session kicked off...

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

10

**Primary author:** Dr MEINHARD, Helge (CERN-IT)

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

**Session Classification:** Benchmarking I

**Track Classification:** Benchmarking

Contribution ID: 46

Type: **not specified**

## Topics around benchmarking at CERN

*Thursday, 8 November 2007 14:00 (30 minutes)*

A collection of remarks on recent procurements and our experience with including power consumption, SPECpower etc.

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

30 minutes

**Primary author:** Dr MEINHARD, Helge (CERN-IT)

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

**Session Classification:** Benchmarking II

**Track Classification:** Benchmarking

Contribution ID: 47

Type: **not specified**

## Single-Sign-On, Identity and Access Management at CERN

*Friday, 9 November 2007 12:00 (30 minutes)*

In order to streamline access to applications, CERN recently deployed Single-Sign-On (SSO) mechanisms along with changes in identity management and access management.

We will present the selected SSO solution and how it integrates with existing applications, as well as related changes in user accounts and roles.

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

30

**Primary author:** Mr LOSSENT, Alexandre (CERN)

**Co-authors:** Mr PACE, Alberto (CERN); Mr ORMANCEY, Emmanuel (CERN)

**Presenter:** Mr LOSSENT, Alexandre (CERN)

**Session Classification:** General III

**Track Classification:** General

Contribution ID: 48

Type: **not specified**

## Sharepoint 2007 deployment at CERN

*Monday, 5 November 2007 16:00 (30 minutes)*

CERN recently introduced a new type of web sites labelled “Collaboration Workspaces”, based on SharePoint 2007. It aims at facilitating a number of tasks like web publishing, information gathering, and collaborative work on documents.

We will present the use cases of SharePoint at CERN, brief technical details about the deployment and how we integrated SharePoint with existing web sites and applications.

### **Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

30

**Primary author:** Mr LOSSENT, Alexandre (CERN)**Presenter:** Mr LOSSENT, Alexandre (CERN)**Session Classification:** General I**Track Classification:** General

Contribution ID: 49

Type: **not specified**

## **Performance of Intel and AMD Processors on Lattice QCD Codes**

*Thursday, 8 November 2007 14:30 (30 minutes)*

We will present lattice QCD applications performance data on current Intel and AMD processors and platforms, including the new quad-core AMD processors ("Barcelona").

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

30

**Primary author:** Dr HOLMGREN, Don (Fermilab)

**Presenter:** Dr HOLMGREN, Don (Fermilab)

**Session Classification:** Benchmarking II

**Track Classification:** Benchmarking

Contribution ID: 50

Type: **not specified**

## Experiences with Microsoft Key Management Server

*Monday, 5 November 2007 14:30 (30 minutes)*

Experiences with Microsoft Key Management Server

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

30

**Primary author:** Mr BAISLEY, Wayne (Fermilab)

**Presenter:** Mr BAISLEY, Wayne (Fermilab)

**Session Classification:** Windows I

**Track Classification:** Windows



Contribution ID: 51

Type: **not specified**

## Using WSUS v3.0

*Monday, 5 November 2007 13:30 (30 minutes)*

Using WSUS v3.0

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

30

**Primary author:** Mr BAISLEY, Wayne (Fermilab)

**Presenter:** Mr BAISLEY, Wayne (Fermilab)

**Session Classification:** Windows I

**Track Classification:** Windows

Contribution ID: 53

Type: **not specified**

## Centralized Logging for the Open Science Grid

*Thursday, 8 November 2007 11:30 (15 minutes)*

The DOE SciDAC Center for Enabling Distributed Petascale Science (CEDPS) has been working with the Open Science Grid to design and deploy tools to facilitate troubleshooting problems on the Grid. This includes tools for central log collection, and advice on what should be logged. This talk will give an overview of progress to date and near-term plans for the future.

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

15 min

**Primary author:** TIERNEY, Brian (LBNL)

**Co-authors:** GUNTER, Dan (LBNL); SAKREJDA, Iwona (LBNL)

**Presenter:** SAKREJDA, Iwona (LBNL)

**Session Classification:** Grid Management and Monitoring II

**Track Classification:** Grid Management and Monitoring

Contribution ID: 54

Type: **not specified**

## FSWG Progress Report

*Tuesday, 6 November 2007 13:35 (25 minutes)*

The talk will be covering the work done during the second phase of HEPiX/IHEPCCC File Systems /Storage Working Group. The plans for the third and final phase will also be presented.

**Primary author:** MASLENNIKOV, Andrei (CASPUR)

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

**Session Classification:** Storage I

**Track Classification:** Storage

Contribution ID: 55

Type: **not specified**

## **NERSC storage update. What's happening and where are we going.**

*Tuesday, 6 November 2007 14:00 (30 minutes)*

NERSC storage update. What's happening and where are we going.

**Primary author:** Dr SAKREJDA, Iwona (LBNL/NERSC)

**Presenter:** Dr SAKREJDA, Iwona (LBNL/NERSC)

**Session Classification:** Storage I

**Track Classification:** Storage

Contribution ID: 56

Type: **not specified**

## Lustre HSM Project

*Tuesday, 6 November 2007 14:30 (30 minutes)*

Lustre is a high performance parallel filesystem. It manages Petabytes of storage and allows hundreds of GB/s. Lustre HSM project will bring to Lustre ILM functionalities. This talk will describe the design of Lustre HSM.

**Primary author:** Mr LAFOUCRIERE, Jacques-Charles (CEA)

**Presenter:** Mr LAFOUCRIERE, Jacques-Charles (CEA)

**Session Classification:** Storage I

**Track Classification:** Storage

Contribution ID: 57

Type: **not specified**

## Update on Chimera and NFS v4.1

*Tuesday, 6 November 2007 16:00 (30 minutes)*

With the start of the Large Hadron Collider, presumably in 2008, the largest share of the produced and generated physics data outside CERN will be managed by the dCache storage element. Beside the challenge to handle storage in the order of tens of Petabyte in a single installation, the number of file operations per second is certainly even more challenging. Two years ago we identified PNFS, the dCache file system component, as to be the part which may become a bottleneck at the point in time when the global LHC data chain will be operated at its full capacity. There are various reasons for this assumption. Due to the fact, that pnfs existed before dCache was designed, dCache and pnfs interact by means of the nfs2 protocol, which introduces a nonessential layer of complexity. Moreover, although pnfs is using the Postgres database, it doesn't make use of enhanced modern database functionalities. Another striking issue is certainly the fact that single read/write locks are used internally, protecting large collection of file which leads to unnecessary queuing of file operation requests. To circumvent those and many other issues we designed and implemented Chimera, a file system name space engine based on the Java programming language as well as on modern database technologies. Furthermore, Chimera is an inevitable prerequisite in dCache to support the NFS 4.1 protocol. NFS 4.1 is currently in its final state of specification. As to our current plans NFS 4.1 will be supported in dCache in the foreseeable future. This presentation will give more insight in the advantages of Chimera and the NFS 4.1 plans of the dCache team.

**Primary author:** Dr FUHRMANN, Patrick (DESY)

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

**Session Classification:** Storage II

**Track Classification:** Storage

Contribution ID: 58

Type: **not specified**

## Native Infiniband Storage

*Tuesday, 6 November 2007 15:30 (30 minutes)*

Storage devices with the native Infiniband interface are gaining more and more popularity. I will discuss the available Infiniband storage access protocols, and will then use one of our recent products as an example.

**Primary author:** JOSEPHAKIS, John (Data Direct Networks)

**Presenter:** JOSEPHAKIS, John (Data Direct Networks)

**Session Classification:** Storage II

**Track Classification:** Storage

Contribution ID: 59

Type: **not specified**

## OpenAFS Status and Futures

*Tuesday, 6 November 2007 16:30 (30 minutes)*

OpenAFS has completed seven years as an open source project. During that time OpenAFS has successfully tackled many challenges including new platforms, multi-processor systems, and network environments consisting of mobile clients and partitioning. OpenAFS Gatekeepers, Derrick Brashear and Jeffrey Altman, will provide updates on the latest OpenAFS offerings, the most common usage models, and describe a vision of what OpenAFS can achieve with the support of the HEP community. Time will be left to answer questions and receive feedback from current users and administrators and those considering deploying AFS in support of their applications and research projects.

**Primary authors:** BRASHEAR, Derrick (Sine Nomine Associates and OpenAFS Gatekeeper); ALTMAN, Jeffrey (Secure Endpoints Inc. and OpenAFS Gatekeeper)

**Presenters:** BRASHEAR, Derrick (Sine Nomine Associates and OpenAFS Gatekeeper); ALTMAN, Jeffrey (Secure Endpoints Inc. and OpenAFS Gatekeeper)

**Session Classification:** Storage II

**Track Classification:** Storage



Contribution ID: 60

Type: **not specified**

## End to End Data Integrity in the Enterprise

*Thursday, 8 November 2007 15:30 (30 minutes)*

The loss of a business critical application resulting from data corruption can be crippling and require huge efforts to repair. All businesses hope they have the know how to recover, yet no business wants an unplanned real-time test. Data corruption prevention, rather than recovery is key. Storage vendors, including those designing adapters, fabrics, and storage arrays design component-level data integrity into their products. Sometimes though, these techniques have not been sufficient and there is a trend to provide a mechanism for end to end data integrity. This talk discusses the nature of data corruption and what the storage industry is doing to provide open and ubiquitous end to end data integrity.

**Primary author:** WILLIAMS, Jim (Oracle Corp.)

**Presenter:** WILLIAMS, Jim (Oracle Corp.)

**Session Classification:** Storage III

**Track Classification:** Storage

Contribution ID: **61**

Type: **not specified**

## **Genome Sequencing Center site report**

*Tuesday, 6 November 2007 11:50 (20 minutes)*

Information about the current IT infrastructure at the Genome Sequencing Center at Washington University in St. Louis and specifically how it has changed over the last year.

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

20

**Primary author:** STIEHR, Gary (The Genome Sequencing Center at Washington University in St. Louis)

**Presenter:** STIEHR, Gary (The Genome Sequencing Center at Washington University in St. Louis)

**Session Classification:** Site Reports IV

**Track Classification:** Site Reports

Contribution ID: 62

Type: **not specified**

## Advanced Group Policy Management

*Monday, 5 November 2007 15:30 (30 minutes)*

Group Policy Objects (GPOs) have a pole position in a domain network management. Unfortunately, the primary tool for managing Group Policy in an enterprise, the Group Policy Management Console (GPMC), doesn't provide a base for a workflow in any but the simplest management models.

Microsoft Advanced Group Policy Management (AGPM), the technology included in the Microsoft Desktop Optimization Pack for Software Assurance for the Microsoft Windows Vista, addresses the problem and adds the needed functionality. It increases the capabilities of the GPMC, providing offline editing of GPOs, role-based delegation of control, check-in/check-out capability, GPO templates, version control and rich reporting.

### **Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

30

**Primary author:** POPINSKI, Kazimierz (DESY)**Presenter:** POPINSKI, Kazimierz (DESY)**Session Classification:** Windows II**Track Classification:** Windows

Contribution ID: **63**

Type: **not specified**

## **Sanger Institute site report**

*Tuesday, 6 November 2007 12:10 (20 minutes)*

Sanger Institute site report

**Primary author:** Dr COATES, Guy (The Wellcome Trust Sanger Institute)

**Presenter:** Dr COATES, Guy (The Wellcome Trust Sanger Institute)

**Session Classification:** Site Reports IV

**Track Classification:** Site Reports

Contribution ID: 64

Type: **not specified**

## **Presentation of the storage agenda**

*Tuesday, 6 November 2007 13:30 (5 minutes)*

Presentation of the storage agenda

**Primary author:** MASLENNIKOV, Andrei (CASPUR)

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

**Session Classification:** Storage I

**Track Classification:** Storage

Contribution ID: 65

Type: **not specified**

## Welcome Message

*Monday, 5 November 2007 09:00 (30 minutes)*

**Primary author:** STIEHR, Gary (The Genome Sequencing Center at Washington University in St. Louis)

**Session Classification:** Welcome

Contribution ID: 66

Type: **not specified**

## HEP Applications with Globus Virtual Workspaces

*Wednesday, 7 November 2007 16:50 (10 minutes)*

HEP applications often have very specific OS requirements and shared resources available to HEP in Canada may not be able to meet the specific OS requirements of a particular application. We have investigated VMs as a mechanism for packaging a complete application and shipping it out to a remote grid site complete with all its requirements. We will discuss experiences using Globus Virtual Workspaces for remote deployment of VMs to Canadian Grid sites.

**Approximate Length (in minutes) (default is 15 minutes for Site Reports and 30 minutes for others, including questions)**

10

**Primary author:** GABLE, Ian (University of Victoria)

**Presenter:** GABLE, Ian (University of Victoria)

**Session Classification:** Virtualization II

**Track Classification:** Virtualization

Contribution ID: 67

Type: **not specified**

## HEPIX Fall 2007 Cyber Security Update

*Thursday, 8 November 2007 16:00 (30 minutes)*

HEPIX Fall 2007 Cyber Security Update

**Primary author:** Mr BUHRMASTER, Gary (SLAC)

**Presenter:** Mr BUHRMASTER, Gary (SLAC)

**Session Classification:** Security

**Track Classification:** General