

HEPiX Spring 2011 Workshop

Report of Contributions

Contribution ID: 0

Type: **not specified**

GSI site report

Monday, May 2, 2011 10:00 AM (15 minutes)

New developments at GSI

Author: SCHON, Walter (GSI)

Presenter: SCHON, Walter (GSI)

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 1

Type: **not specified**

Selecting a new batch system at CC-IN2P3

Monday, May 2, 2011 4:30 PM (30 minutes)

2 years ago, CC-IN2P3 decided to give up the home made batch system (BQS) for a new product. This presentation will expose the process we set up to make the selection and will explain our choice.

Author: Mr CHAMBON, Bernard (CC-IN2P3 /CNRS)

Presenter: Mr CHAMBON, Bernard (CC-IN2P3 /CNRS)

Session Classification: Computing

Track Classification: Computing & Batch Services

Contribution ID: 2

Type: **not specified**

Grid Engine setup at CC-IN2P3

Monday, May 2, 2011 5:00 PM (30 minutes)

As you know, we chose Grid Engine as the next batch system for CC-IN2P3.

This presentation will focus on 2 aspects we have examined during last months :

1/ Scalability and robustness testing

2/ Specific requirements at CC-IN2P3 : problems and solutions

Author: Mr CHAMBON, Bernard (CC-IN2P3 /CNRS)

Presenter: Mr CHAMBON, Bernard (CC-IN2P3 /CNRS)

Session Classification: Computing

Track Classification: Computing & Batch Services

Contribution ID: 3

Type: **not specified**

IHEP Site Report

Tuesday, May 3, 2011 9:30 AM (15 minutes)

Report of status of computing system of IHEP

Summary

The computing center of IHEP is to provide computing services for BESIII, the experiment on Beijing Electron Positron Collider. The computing center is also the tier-2 site for ATLAS and CMS. This report will discuss the status of the computing system at IHEP.

Author: Dr CHEN, Gang (Institute of High Energy Physics (IHEP))

Presenter: Dr CHEN, Gang (Institute of High Energy Physics (IHEP))

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 4

Type: **not specified**

Foundation of the EOFS - support the lustre development

Thursday, May 5, 2011 4:30 PM (30 minutes)

The European Open File Systems society is a non profit organisation to coordinate the future development of lustre. Founding members of the organisation are Universities, Supercomputing Centers and partners from industry. The next release of a lustre version is scheduled for summer 2011.

Author: SCHON, Walter (GSI)

Presenter: SCHON, Walter (GSI)

Session Classification: Storage & FileSystems

Track Classification: Storage & Filesystems

Contribution ID: 5

Type: **not specified**

FermiGrid Scalability and Reliability Improvements

Thursday, May 5, 2011 11:45 AM (30 minutes)

The Fermilab Campus Grid (FermiGrid) is a meta-facility that provides grid infrastructure for scientific computing at Fermilab. It provides highly available centralized authorization and authentication services, a site portal for Globus job submission, coordination for interoperability among the various stakeholders, and grid-enabled mass storage interfaces. We currently support approximately 25000 batch processing slots. This presentation will describe the current structure of FermiGrid and recent improvements in scalability and reliability of our authorization and authentication services. These improvements include orders of magnitude improvement in our web services based Site Authorization service (SAZ). We will also describe recent enhancements to the information system and matchmaking algorithm of our site job gateway. Finally we will describe the FermiGrid HA2 project currently under way which distributes our services across two buildings, making us resilient in the case of major building outages.

Author: Dr CHADWICK, Keith (Fermilab)

Presenter: Dr CHADWICK, Keith (Fermilab)

Session Classification: Cloud, grid and virtualization

Track Classification: Grid, cloud and virtualization

Contribution ID: 6

Type: **not specified**

Fermilab Site Report - Spring 2011 HEPiX

Monday, May 2, 2011 10:15 AM (20 minutes)

The Fermilab site report for the Spring 2011 HEPiX

Author: Dr CHADWICK, Keith (Fermilab)

Presenter: Dr CHADWICK, Keith (Fermilab)

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 7

Type: **not specified**

Evolution of CERN's Computing Facilities

Wednesday, May 4, 2011 2:45 PM (30 minutes)

CERN is currently evolving its computing facilities through a number of projects. This presentation will give an overview of the various projects and their current status.

Author: Mr SALTER, Wayne (CERN)

Presenter: Mr SALTER, Wayne (CERN)

Session Classification: IT Infrastructure

Track Classification: IT Infrastructure and Services

Contribution ID: 8

Type: **not specified**

Signing (and encrypted) message handling and implications for admins.

Monday, May 2, 2011 2:00 PM (30 minutes)

In grid computing we use an X509 PKI security infrastructure. This infrastructure is used to enable secure connections between hosts to deliver payload. This often leads to scalability and reliability issues. This talk presents the alternative approach of signing messages for asynchronous handling, allowing authentication of the payload rather than the connection.

The implications of this approach will be illustrated showing how service interdependency can be reduced, and clustering simplified. AMQP (RabbitMQ) will be used as a transport mechanism in this talk to illustrate these concepts. Both the openssl command line and a python library can be used to authenticate signed messages making scalable secure authentication between sites resources practical for administrators.

Summary

How authentication of signing (and encrypted) message handling makes life easier.

Author: SYNGE, Owen (DESY (HH))

Presenter: SYNGE, Owen (DESY (HH))

Session Classification: Networking & Security

Track Classification: Security & Networking

Contribution ID: 9

Type: **not specified**

GridKa Site Report

Monday, May 2, 2011 10:35 AM (15 minutes)

Current status and latest news at GridKa, e.g.:

- Hardware status
- Batch system issues
- CPU performance

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

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

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 10

Type: **not specified**

Moving virtual machines images securely between sites.

Thursday, May 5, 2011 9:15 AM (30 minutes)

A Grid service allows applications to run on many sites without modification. Virtualization provides the potential for deploying the same customized operating system at many sites.

This talk will present one of many possible security infrastructures and models to allow sharing and deployment of virtual machines images that meets the objectives of secure non-repudiation of images, auditing and fault tolerance that has been developed within the HEPiX virtualization working group. The talk focus will be made on the meta data describing the virtual machines, how to share this meta data, how to share the image described and to verify an image and its meta data, presenting packages and deployment, and how to audit the approach.

Summary

Non-repudiation of meta data messages describing images to be deployed without increasing interdependency.

Author: SYNGE, Owen (DESY (HH))

Presenter: SYNGE, Owen (DESY (HH))

Session Classification: Cloud, grid and virtualization

Track Classification: Grid, cloud and virtualization

Contribution ID: 11

Type: **not specified**

Virtualization at CERN: a status report

Thursday, May 5, 2011 9:45 AM (30 minutes)

Virtualization at CERN: a status report

We present updates to the virtualization services provided by CERN IT.

- CERNs internal cloud has been moved into full production mode in December 2010, and has been running since then providing virtualized batch resources. We will report on operational experiences, as well as further developments made since the last meeting in Cornell, including benchmark results, OpenNebula and ISF experiences, a first view on SLC6.
- the CVI Self-Service continues to grow rapidly (>1200 VMs on >200 hypervisors), and so are the use requirements. We describe the service evolution of CVI 2, with a focus on Linux VMs.

We will also present the plans to evaluate Openstack at CERN.

Author: Dr SCHWICKERATH, Ulrich (CERN)

Presenter: Dr SCHWICKERATH, Ulrich (CERN)

Session Classification: Cloud, grid and virtualization

Track Classification: Grid, cloud and virtualization

Contribution ID: 13

Type: **not specified**

NDGF Site Report

Tuesday, May 3, 2011 9:45 AM (15 minutes)

Site update from NDGF, recent developments, neat things, etc.

Author: WADENSTEIN, Mattias (NDGF)

Presenter: WADENSTEIN, Mattias (NDGF)

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 14

Type: **not specified**

Implementing Service Management processes with Service-Now

Wednesday, May 4, 2011 3:15 PM (30 minutes)

The choice of Service-Now as a tool for handling the request fulfillment and incident management ITIL processes in the IT and the General Services Departments at CERN has created several months of intensive development. Besides the implementation of these two standardized ITIL process it has been a very interesting task to model CERN Service catalogue in the tool. The integration with third party systems and workflows, as SSO, GGUS, organization data, knowledge base, has started and will be a running task for the next couple of years. The biggest challenge will be the transition of existing non-ITIL processes implemented in other tools into Service-Now.

Author: TOTEVA, Zhechka (CERN)

Presenter: TOTEVA, Zhechka (CERN)

Session Classification: IT Infrastructure

Track Classification: IT Infrastructure and Services

Contribution ID: 15

Type: **not specified**

Performance Comparison of Multi and Many-Core Batch Nodes

Tuesday, May 3, 2011 3:15 PM (30 minutes)

The compute power of batch nodes is measured in units of HEP-SPEC06 which is based on the industry standard SPEC CPU2006 benchmark suite.

In this talk I will compare the HEP-SPEC06 scores of multi-core worker nodes with accounting data taken from the batch system.

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

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

Session Classification: Computing

Track Classification: Computing & Batch Services

Contribution ID: 16

Type: **not specified**

Drupal at CERN

Tuesday, May 3, 2011 11:45 AM (30 minutes)

Drupal is an open source content management platform used worldwide. CERN has chosen to use Drupal for building multilingual, content managed web sites and applications. The infrastructure is based on cluster of Apache web servers, MySQL database servers and storage servers. The setup uses SLC6 operating system as a platform. The high availability configuration is achieved with Red Hat Cluster Suite. The talk will present the details of the Drupal configuration at CERN, current status of the project and integration with existing CERN services: e-groups, CERN Authentication and CERN Document Server.

Author: Mr SUCIK, Juraj (CERN)

Co-author: Mr POLOK, Jarek (CERN)

Presenter: Mr SUCIK, Juraj (CERN)

Session Classification: IT Infrastructure

Track Classification: IT Infrastructure and Services

Contribution ID: 17

Type: **not specified**

OpenMP Performance on Virtual Machines

Tuesday, May 3, 2011 2:15 PM (30 minutes)

Virtualization technology has been applied to a variety of areas including server consolidation, High Performance Computing, as well as Grid and Cloud computing. Due to the fact that applications do not run directly on the hardware of a host machine, virtualization generally causes a performance loss for both sequential and parallel applications.

This talk studies the OpenMP applications running on a virtualized multicore machine. It shows the overhead of parallelization and compares the parallel performance on virtual machines with the performance of native executions. An interesting scenario is that one application runs much slower in parallel than the sequential runs. A performance analysis tool is then applied to investigate the cause of such abnormal behavior. The talk demonstrates the performance optimization and the results based on the analysis.

Author: Dr TAO, Jie**Presenter:** Dr TAO, Jie**Session Classification:** Computing**Track Classification:** Computing & Batch Services

Contribution ID: **18**

Type: **not specified**

CC-IN2P3 Site-Report update

Tuesday, May 3, 2011 9:15 AM (15 minutes)

I'll present a short update of the CC-In2p3 Site-report, specifying the main changes occurred since last hepix.

Author: Mr OLIVERO, philippe (CC-IN2P3)

Presenter: Mr OLIVERO, philippe (CC-IN2P3)

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 19

Type: **not specified**

Computer security update

Monday, May 2, 2011 2:30 PM (30 minutes)

This presentation provides an update of the security landscape since the last meeting. It describes the main vectors of compromises in the academic community and discusses security risks management in general, as well as the security aspects of the current hot topics in computing, for example identity federation and virtualisation.

Author: Mr WARTEL, Romain (CERN)

Presenter: Mr WARTEL, Romain (CERN)

Session Classification: Networking & Security

Track Classification: Security & Networking

Contribution ID: 20

Type: **not specified**

StratusLab Marketplace for Sharing Virtual Machine Images

Thursday, May 5, 2011 10:15 AM (30 minutes)

StratusLab (<http://stratuslab.eu/>) provides a complete, open-source solution for deploying an “Infrastructure as a Service” cloud infrastructure. Use of the cloud requires the use of prepared machine and disk images, yet preparing correct, secure images remains difficult and represents a significant barrier to the adoption of cloud technologies.

The StratusLab Marketplace is an image registry, containing cryptographically signed metadata associated with shared images. It simultaneously allows: end-users to search for existing images, image creators to publicize their images, and cloud administrators to evaluate the trustworthiness of an image. The image files themselves are stored elsewhere—either in cloud storage or in web-accessible repositories.

The StratusLab Marketplace facilitates the sharing of images and use of IaaS cloud infrastructures, allowing users access to a diverse set of existing images and providing cloud administrators with the confidence to allow them to run. Its integration with the StratusLab distribution makes use of registered images easy, further reducing barriers to adoption.

Author: LOOMIS, Cal (CNRS/LAL)

Presenter: LOOMIS, Cal (CNRS/LAL)

Session Classification: Cloud, grid and virtualization

Track Classification: Grid, cloud and virtualization

Contribution ID: 21

Type: **not specified**

Scientific Linux Status Report + Discussion

Wednesday, May 4, 2011 4:15 PM (45 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.

Author: Mr DAWSON, Troy (FERMILAB)

Presenter: Mr DAWSON, Troy (FERMILAB)

Session Classification: IT Infrastructure

Track Classification: IT Infrastructure and Services

Contribution ID: 22

Type: **not specified**

Evaluation of gluster file system at IHEP

Wednesday, May 4, 2011 9:15 AM (30 minutes)

GlusterFS is an open source, clustered file system capable of scaling to several petabytes and handling thousands of clients. At IHEP, we setup a testbed to evaluate the file system, including functionality, performance and current status. The advantages and disadvantages for HEP data processing are also discussed.

Author: Dr CHENG, Yaodong (Institute of High Energy Physics, Chinese Academy of Sciences)

Presenter: Dr CHENG, Yaodong (Institute of High Energy Physics, Chinese Academy of Sciences)

Session Classification: Storage & FileSystems

Track Classification: Storage & Filesystems

Contribution ID: 23

Type: **not specified**

RAL Site Report

Monday, May 2, 2011 12:15 PM (20 minutes)

Update on activities at RAL Tier1

Author: BLY, Martin (STFC-RAL)

Presenter: BLY, Martin (STFC-RAL)

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 24

Type: **not specified**

Indico - Present and future

Tuesday, May 3, 2011 12:15 PM (30 minutes)

Indico (Integrated Digital Conference: <http://indico.cern.ch>) is a web-based, multi-platform conference lifecycle management system and agenda. It has also become the long term archiving tool for documents and metadata related to all kinds of events that take place at CERN. The software is used in production at CERN (hosting >114.000 events, > 580.000 presentations, > 770.000 files and around 10.000 visitors per day) and installed in more than 90 institutes world-wide.

Indico has been changing a lot in the last 3 years, therefore we will review all these changes and new features along this period, and we will also give an overview of the future for Indico.

Author: Mr GONZALEZ LOPEZ, Jose Benito (CERN)

Presenter: Mr GONZALEZ LOPEZ, Jose Benito (CERN)

Session Classification: IT Infrastructure

Track Classification: IT Infrastructure and Services

Contribution ID: 25

Type: **not specified**

Benefits of a virtualized approach to mass storage system

Wednesday, May 4, 2011 9:45 AM (20 minutes)

The talk will show the benefits of grouping a number of heterogeneous tape libraries into one virtual container of tape media and drives. The backup and archive applications send their data to this huge container which has all necessary mechanisms to control and access the tape resources(cartridges, drives, physical libraries).The implementation is based on an ibm software "Enterprise Removable Media Manager"

Author: LOBONTU, Dorin Daniel (Karlsruhe Institute of Technology (KIT))

Presenter: LOBONTU, Dorin Daniel (Karlsruhe Institute of Technology (KIT))

Session Classification: Storage & FileSystems

Contribution ID: 26

Type: **not specified**

Petersburg Nuclear Physics Institute (PNPI) status report

Tuesday, May 3, 2011 10:00 AM (15 minutes)

It is planned to describe the updated status for the computing infrastructure of High Energy Physics Division (HEPD): LAN (400 hosts), mail service for the Institute, other centralized servers, computing cluster. A number of updated topics is observed: security and SPAM, cluster virtualization, WiFi, video conferences.

Author: Mr SHEVEL, Andrey (Petersburg Nuclear Physics Institute (PNPI))

Presenter: Mr SHEVEL, Andrey (Petersburg Nuclear Physics Institute (PNPI))

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 27

Type: **not specified**

Open Source at Oracle

Thursday, May 5, 2011 2:40 PM (30 minutes)

In this presentation, Oracle will go over its major open source products, and their future directions.

Summary

In this presentation, Oracle will go over its major open source products, and their future directions.

Author: Mr GRAVIER, Gilles (Oracle)

Presenter: Mr GRAVIER, Gilles (Oracle)

Session Classification: Oracle

Track Classification: Miscellaneous

Contribution ID: **28**

Type: **not specified**

CERN site report

Monday, May 2, 2011 11:35 AM (20 minutes)

News from CERN since last meeting

Author: Dr MEINHARD, Helge (CERN-IT)

Presenter: Dr MEINHARD, Helge (CERN-IT)

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 29

Type: **not specified**

Version Control Services at CERN

Friday, May 6, 2011 9:15 AM (30 minutes)

CERN offers three Version Control Services, one using SVN and two older services using CVS. The older CVS service is to be closed by 2Q 2011 and will be merged into the high availability CVS service on AFS where the performance has been improved to suite the needs of all users. The main SVN service has expanded a great deal, in users, commits and repositories, since it started in 2009. Our future plans include new tools for users, internal software upgrades, improved statistics and monitoring.

Author: Mr GONZALEZ ALVAREZ, Alvaro (CERN)

Presenter: Mr GONZALEZ ALVAREZ, Alvaro (CERN)

Session Classification: IT Infrastructure

Track Classification: IT Infrastructure and Services

Contribution ID: 30

Type: **not specified**

CASTOR status and development

Wednesday, May 4, 2011 10:05 AM (30 minutes)

We will present the performance achieved during data taking for the 2010 LHC run, including heavy ion run. The operational benefits reaped from the deployed improvement as well as the roadmap for further developments to consolidate the system and lower its deployment cost will be introduced. Our performance assessment for the new generation of Oracle tape drives –T10000C –will also be shown.

Author: CANO, Eric (CERN)

Presenter: CANO, Eric (CERN)

Session Classification: Storage & FileSystems

Track Classification: Storage & Filesystems

Contribution ID: 31

Type: **not specified**

Operating a distributed IaaS Cloud for BaBar MC production and user analysis

Thursday, May 5, 2011 11:15 AM (30 minutes)

In the last year we have established a system which replicates a standard Condor HTC environment across multiple distinct IaaS clouds of different types including EC2, Nimbus and Eucalyptus. Users simply submit batch jobs to a Condor queue containing a custom attribute which is a pointer to the Virtual Machine Image they would like booted to service their job. The system automatically boots instances of the requested machine type on one of the available clouds and contextualizes them to connect to the batch system. The system is being used on a continual basis for astronomy and HEP jobs. We report on our experience operating this system which has booted over 30 000 VMs and completed over 250 000 jobs.

Author: GABLE, Ian (University of Victoria)

Co-author: SOBIE, Randall (University of Victoria)

Presenter: GABLE, Ian (University of Victoria)

Session Classification: Cloud, grid and virtualization

Track Classification: Grid, cloud and virtualization

Contribution ID: 32

Type: **not specified**

Batch Monitoring and Testing

Monday, May 2, 2011 4:00 PM (30 minutes)

In order to improve its batch service for local and Grid users, development is ongoing at CERN to design a batch monitoring system and set up a test instance. The goal is to enhance the batch service by investigating new scheduler features, fine-tuning the already used ones and decreasing the time spent in problem identification and fault resolution.

Author: Mr BELLEMAN, Jerome (CERN)

Presenter: Mr BELLEMAN, Jerome (CERN)

Session Classification: Computing

Track Classification: Computing & Batch Services

Contribution ID: 33

Type: **not specified**

An introduction to Oracle Linux

Thursday, May 5, 2011 2:00 PM (40 minutes)

In this presentation, Lenz will provide an overview about Oracle Linux, Oracle's Enterprise Linux distribution and the Oracle Unbreakable Enterprise Kernel (UEK). The session will cover the technical highlights and improvements as well as the support offerings that complement it.

Author: Mr GRIMMER, Lenz (Oracle)

Presenter: Mr GRIMMER, Lenz (Oracle)

Session Classification: Oracle

Track Classification: Miscellaneous

Contribution ID: **34**

Type: **not specified**

Discussion

Thursday, May 5, 2011 3:10 PM (50 minutes)

Session Classification: Oracle

Contribution ID: 35

Type: **not specified**

Overview of the new Computing room at CC-IN2P3

Wednesday, May 4, 2011 2:15 PM (30 minutes)

The presentation will give an overview of the new commissioned infrastructure and computing room at CC-IN2P3.

The presentation will give an update of technical infrastructure project and improvement. It will focus on major achievement and will spotlight and describe the future capacity offered up to 2019
Topics to be reviewed: Building, Cooling system, Power distribution and Confined Racks, Future capacity, projects and scheduling.

Author: Mr TROUVE, Pascal (CC-IN2P3)

Presenter: Mr TROUVE, Pascal (CC-IN2P3)

Session Classification: IT Infrastructure

Track Classification: IT Infrastructure and Services

Contribution ID: 36

Type: **not specified**

FAIR 3D Tier-0 Green-IT Cube

Tuesday, May 3, 2011 4:15 PM (1 hour)

FAIR computing presents computing requirements for the first level processing of the experiment data exceeding those at CERN. All computing resources, including the first level event selectors will be hosted in one data center, which is currently being planned. It sets new standards with respect to energy density, implementing more than 100 kW/sqm and energy efficiency by requiring less than 10% for the data center cooling, while allowing the use of general purpose computer servers. The over all FAIR computing concept is presented as well as the FAIR Tier-0 data center architecture.

Author: Prof. LINDENSTRUTH, Volker (FIAS, GSI)

Presenter: Prof. LINDENSTRUTH, Volker (FIAS, GSI)

Session Classification: IT Infrastructure

Track Classification: IT Infrastructure and Services

Contribution ID: 37

Type: **not specified**

DESY Site Report

Monday, May 2, 2011 11:55 AM (20 minutes)

Computing at DESY news

Author: Dr FRIEBEL, Wolfgang (Deutsches Elektronen-Synchrotron (DESY)-Unknown-Unknown)

Presenter: Dr FRIEBEL, Wolfgang (Deutsches Elektronen-Synchrotron (DESY)-Unknown-Unknown)

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 38

Type: **not specified**

The DESY Grid-Lab, a detailed 'local access protocol' evaluation

Wednesday, May 4, 2011 11:15 AM (30 minutes)

Since mid of 2010, DESY IT is operating a performance evaluation facility of the size of a small gLite Tier II, the DESY Grid-Lab. Regular gLite software is deployed, allowing to execute commonly used LHC analysis jobs as well as applications provided by other communities. This presentation focuses on the comparison of different implementations of XROOTD, dCap as well as of the NFS4.1/pNFS dCache implementation. The evaluation scenarios include real world analysis jobs of LHC VO's including standard Hammercloud jobs, I/O intensive jobs provided by the ROOT team and examples from non HEP communities.

Authors: OZEROV, Dmitry (DESY); KEMP, Yves (DESY)

Presenter: OZEROV, Dmitry (DESY)

Session Classification: Storage & FileSystems

Contribution ID: 39

Type: **not specified**

Lustre at GSI

Wednesday, May 4, 2011 11:45 AM (30 minutes)

Lustre has been employed with great success as the general purpose distributed file system for all experimental and theory groups at GSI.

Currently there are 100 mio files stored on Lustre, and between batch nodes, interactive nodes and desktops there are ca. 500 clients with access to Lustre.

Past issues with stability have been overcome by running the Lustre 1.8 version. Hardware upgrades of metadata servers and OSS are under way. The total file space will increase to > 2 PB soon.

Author: ROTH, Thomas (GSI)

Presenter: ROTH, Thomas (GSI)

Session Classification: Storage & FileSystems

Track Classification: Storage & Filesystems

Contribution ID: 40

Type: **not specified**

Invenio at CERN

Tuesday, May 3, 2011 12:45 PM (30 minutes)

Invenio <http://invenio-software.org/> is a software suite enabling to run a digital library or document repository on the web. The technology offered by the software covers all aspects of digital library management from document ingestion through classification, indexing, and curation to dissemination. Invenio has been originally developed at CERN to run the CERN document server (CDS), managing over 1,000,000 bibliographic records in high-energy physics since 2002, covering articles, books, journals, photos, videos, and more. Invenio is nowadays co-developed by an international collaboration comprising institutes such as CERN, DESY, EPFL, FNAL, SLAC and is being used by about thirty scientific institutions worldwide.

The presentation will focus on the current and future usage of Invenio at CERN: integration with other CERN IT services (Drupal, GRID, Indico, MediaArchive, AIS, etc.) as well as other HEP-related information systems, newly introduced features and workflows, usage statistics, etc. The software development strategy, including future planned developments as well as insight into the underlying technologies will be covered.

Author: Mr CAFFARO, Jerome (CERN)

Presenter: Mr CAFFARO, Jerome (CERN)

Session Classification: IT Infrastructure

Track Classification: IT Infrastructure and Services

Contribution ID: 41

Type: **not specified**

CernVM-FS Production Service and Deployment

Friday, May 6, 2011 9:45 AM (30 minutes)

CernVM-FS is now a production service supported at CERN distributing VO software to sites/worker nodes. This talk will describe the production service as well as give details on the deployment and management required to use CVMFS at sites.

Author: Mr COLLIER, Ian Peter (STFC RAL Tier1)

Presenter: Mr COLLIER, Ian Peter (STFC RAL Tier1)

Session Classification: IT Infrastructure

Track Classification: IT Infrastructure and Services

Contribution ID: 42

Type: **not specified**

Nikhef site report

Monday, May 2, 2011 10:50 AM (15 minutes)

Nikhef site report

Author: Mr KUIPERS, Paul (nikhef)

Presenter: Mr KUIPERS, Paul (nikhef)

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 43

Type: **not specified**

Adopting Infrastructure as Code to run HEP applications

Thursday, May 5, 2011 12:15 PM (30 minutes)

GSI is a German national laboratory for heavy-ion beams, planning to build the new accelerator complex “Facility for Antiproton and Ion Research” (FAIR). In preparation for the Tier-0 computing center for FAIR different Infrastructure as a Service (IaaS) cloud technologies have been compared, to construct a private cloud. Simultaneously, effort has been taken to learn how to efficiently execute HEP applications in a virtual environment. The result is a private cloud testbed, called SCLab, build with the help of the OpenNebula toolkit. The concept Infrastructure as Code (IaC), based on the Chef configuration management system, has been adopted for the deployment and operation of HEP applications in clouds. Tools have been developed to start virtual clusters in any IaaS cloud on demand. The first successful applications are a completely virtual AliEn grid site for the ALICE experiment at LHC and simulations for radiation protection studies for FAIR. The talk will present the design decisions and the experience in running HEP applications in IaaS clouds

Author: ZYNOVYEV, Mykhaylo (GSI)

Presenter: ZYNOVYEV, Mykhaylo (GSI)

Session Classification: Cloud, grid and virtualization

Track Classification: Grid, cloud and virtualization

Contribution ID: 44

Type: **not specified**

CMS 64bit transition and multicore plans

Tuesday, May 3, 2011 2:45 PM (30 minutes)

CMS has ported its complete software stack to run natively on 64 bit linux and it's using it for all its computing workflows, from data acquisition to final analysis. In this talk we'll present our experience with such a transition, both in terms of deployment issues and actual performance gains. Moreover, we'll give an insight of what we consider our present and future challenges, focusing in particular on how we plan to exploit multi-core architectures.

Author: Mr EULISSE, Giulio (FERMILAB)

Presenter: Mr EULISSE, Giulio (FERMILAB)

Session Classification: Computing

Track Classification: Computing & Batch Services

Contribution ID: 45

Type: **not specified**

HEPiX IPv6 Working Group

Wednesday, May 4, 2011 12:45 PM (30 minutes)

A new working group on IPv6 in HEP was discussed and agreed at the previous HEPiX meeting. This new working group has recently been created and work is just starting. This talk will present the status and plans of the working group for the year ahead.

Author: Dr KELSEY, David (RAL)

Presenter: Dr KELSEY, David (RAL)

Session Classification: Networking & Security

Track Classification: Security & Networking

Contribution ID: 46

Type: **not specified**

SLAC Site Report

Monday, May 2, 2011 12:35 PM (25 minutes)

Update on activities at SLAC

Author: Mr MELEN, Randy (SLAC National Accelerator Laboratory)

Presenter: Mr MELEN, Randy (SLAC National Accelerator Laboratory)

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 47

Type: **not specified**

Discussion

Thursday, May 5, 2011 5:00 PM (30 minutes)

Session Classification: Storage & FileSystems

Contribution ID: 48

Type: **not specified**

DLS site report

Tuesday, May 3, 2011 10:15 AM (15 minutes)

Overview of computing systems at Diamond, including current status and planned future developments.

Author: Ms FRIEDRICH, Tina (Diamond Light Source Ltd)

Co-authors: Mr FERNER, Frederik (Diamond Light Source Ltd); Dr MATTHEWS, Greg (Diamond Light Source Ltd); Mr VON SEIBOLD, Max (Diamond Light Source Ltd); Dr REES, Nick (Diamond Light Source Ltd)

Presenter: Ms FRIEDRICH, Tina (Diamond Light Source Ltd)

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 49

Type: **not specified**

Host based intrusion detection with OSSEC

Monday, May 2, 2011 3:00 PM (30 minutes)

In this talk the open source host-based intrusion detection system OSSEC is described. Besides an overview of its features it will also be explained how to use it for non-security related monitoring and notifying. Furthermore several possible real life scenarios will be demonstrated and some of the current drawbacks will be elaborated.

Author: NEUBURGER, Bastian (GSI)

Presenter: NEUBURGER, Bastian (GSI)

Session Classification: Networking & Security

Track Classification: Security & Networking

Contribution ID: 50

Type: **not specified**

BNL Site report

Tuesday, May 3, 2011 10:30 AM (15 minutes)

A report on the current status of the RHIC/ATLAS Computing Facility at BNL with an emphasis on developments and updates since the last Fall Hepix meeting.

Presenter: Dr RIND, Ofer (BROOKHAVEN NATIONAL LABORATORY)

Session Classification: Site Reports

Contribution ID: 51

Type: **not specified**

ASGC site report

Tuesday, May 3, 2011 11:15 AM (15 minutes)

ASGC current status.

Author: Mr LEE, Felix (Academia Sinica)

Presenter: Mr LEE, Felix (Academia Sinica)

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 53

Type: **not specified**

PSI - Site report

Tuesday, May 3, 2011 11:30 AM (15 minutes)

Site report for the Paul Scherrer Institut.

Author: Dr FEICHTINGER, Derek (PSI)

Presenter: Dr FEICHTINGER, Derek (PSI)

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 54

Type: **not specified**

Evaluation of distributed file systems using trace and replay mechanism

Wednesday, May 4, 2011 12:15 PM (30 minutes)

Reliable benchmarking of file systems is a complex and time consuming task when one has to test against a production environment to achieve relevant results.

In case of the HEP community, this eventually leads to setting up a particular experiments' software environment, which could be a rather complicated task for a system administrator.

To simplify this task, we developed an application for exact replaying of IO requests to reliably replicate an IO behavior of the original applications without a need of installing the whole working environment.

Using the application, we present performance comparison of Lustre, GPFS and Hadoop file systems by replaying traces of LHCb, CMS and ATLAS jobs.

Author: Mr HORKY, Jiri (Institute of Physics of Acad. of Sciences of the Czech Rep. (ASCR))

Presenter: Mr HORKY, Jiri (Institute of Physics of Acad. of Sciences of the Czech Rep. (ASCR))

Session Classification: Storage & FileSystems

Track Classification: Storage & Filesystems

Contribution ID: 55

Type: **not specified**

HEPiX VWG Status Report

Friday, May 6, 2011 11:15 AM (30 minutes)

This presentation will give an update of the activities of the HEPiX Virtualisation Working Group over the past few months, describe the current status and give an outlook on future progress.

Author: CASS, Tony (CERN)

Presenter: CASS, Tony (CERN)

Session Classification: Cloud, grid and virtualization

Track Classification: Grid, cloud and virtualization

Contribution ID: 56

Type: **not specified**

The National Grid Service Cloud

Friday, May 6, 2011 10:45 AM (30 minutes)

The UK's National Grid Service is investigating how it can best make use of cloud technologies in the future. The focus is on users, not only those who want to perform computationally intensive research, but also others in the wider academic setting. The usefulness of Infrastructure as a Service clouds to this community is crucial in determining future cloud provision in this area. To examine this question, Eucalyptus-based clouds were deployed at the Universities of Edinburgh and Oxford to gain real experience from the users' perspective.

Authors: Dr FERGUSSON, David (University of Edinburgh); Dr WALLOM, David (University of Oxford); Dr TURILLI, Matteo (University of Oxford); Dr THORN, Steve (University of Edinburgh)

Presenter: Dr THORN, Steve (University of Edinburgh)

Session Classification: Cloud, grid and virtualization

Track Classification: Grid, cloud and virtualization

Contribution ID: 58

Type: **not specified**

Registration

Monday, May 2, 2011 9:15 AM (30 minutes)

Session Classification: Introduction

Contribution ID: 59

Type: **not specified**

Welcome Address

Monday, May 2, 2011 9:45 AM (15 minutes)

Presenter: Prof. LANGANKE, Karlheinz (GSI Darmstadt)

Session Classification: Introduction

Contribution ID: **60**

Type: **not specified**

Wrap-Up

Friday, May 6, 2011 11:45 AM (15 minutes)

Presenter: JOUVIN, Michel (LAL / IN2P3)

Session Classification: Wrap-up