

HEPiX Spring 2011 Workshop

Monday, 2 May 2011 - Friday, 6 May 2011

GSI Helmholtzzentrum für Schwerionenforschung GmbH

Book of Abstracts

Contents

GSI site report	1
Selecting a new batch system at CC-IN2P3	1
Grid Engine setup at CC-IN2P3	1
IHEP Site Report	1
Foundation of the EOFS - support the lustre development	2
FermiGrid Scalability and Reliability Improvements	2
Fermilab Site Report - Spring 2011 HEPiX	2
Evolution of CERN's Computing Facilities	3
Signing (and encrypted) message handling and implications for admins.	3
GridKa Site Report	3
Moving virtual machines images securely between sites.	3
Virtualization at CERN: a status report	4
NDGF Site Report	4
Implementing Service Management processes with Service-Now	5
Performance Comparison of Multi and Many-Core Batch Nodes	5
Drupal at CERN	5
OpenMP Performance on Virtual Machines	6
CC-IN2P3 Site-Report update	6
Computer security update	6
StratusLab Marketplace for Sharing Virtual Machine Images	7
Scientific Linux Status Report + Discussion	7
Evaluation of gluster file system at IHEP	7
RAL Site Report	8

Indico - Present and future	8
Benefits of a virtualized approach to mass storage system	8
Petersburg Nuclear Physics Institute (PNPI) status report	8
Open Source at Oracle	9
CERN site report	9
Version Control Services at CERN	9
CASTOR status and development	10
Operating a distributed IaaS Cloud for BaBar MC production and user analysis	10
Batch Monitoring and Testing	10
An introduction to Oracle Linux	11
Discussion	11
Overview of the new Computing room at CC-IN2P3	11
FAIR 3D Tier-0 Green-IT Cube	11
DESY Site Report	12
The DESY Grid-Lab, a detailed 'local access protocol' evaluation	12
Lustre at GSI	12
Invenio at CERN	13
CernVM-FS Production Service and Deployment	13
Nikhef site report	13
Adopting Infrastructure as Code to run HEP applications	13
CMS 64bit transition and multicore plans	14
HEPiX IPv6 Working Group	14
SLAC Site Report	14
Discussion	15
DLS site report	15
Host based intrusion detection with OSSEC	15
BNL Site report	15
ASGC site report	16
PSI - Site report	16
Evaluation of distributed file systems using trace and replay mechanism	16

HEPiX VWG Status Report	16
The National Grid Service Cloud	17
Registration	17
Welcome Address	17
Wrap-Up	17

Site Reports / 0**GSI site report**

Author: Walter Schon¹

¹ *GSI*

Corresponding Author: w.schoen@gsi.de

New developments at GSI

Computing / 1**Selecting a new batch system at CC-IN2P3**

Author: Bernard CHAMBON¹

¹ *CC-IN2P3 /CNRS*

Corresponding Author: chambon@cc.in2p3.fr

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.

Computing / 2**Grid Engine setup at CC-IN2P3**

Author: Bernard CHAMBON¹

¹ *CC-IN2P3 /CNRS*

Corresponding Author: chambon@cc.in2p3.fr

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

Site Reports / 3**IHEP Site Report**

Author: Gang Chen¹

¹ *Institute of High Energy Physics (IHEP)*

Corresponding Author: gang.chen@cern.ch

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.

Storage & FileSystems / 4**Foundation of the EOFS - support the lustre development**

Author: Walter Schon¹

¹ *GSI*

Corresponding Author: w.schoen@gsi.de

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.

Cloud, grid and virtualization / 5**FermiGrid Scalability and Reliability Improvements**

Author: Keith Chadwick¹

¹ *Fermilab*

Corresponding Author: chadwick@fnal.gov

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.

Site Reports / 6**Fermilab Site Report - Spring 2011 HEPiX**

Author: Keith Chadwick¹

¹ *Fermilab*

Corresponding Author: chadwick@fnal.gov

The Fermilab site report for the Spring 2011 HEPiX

IT Infrastructure / 7**Evolution of CERN's Computing Facilities****Author:** Wayne Salter¹¹ CERN**Corresponding Author:** wayne.salter@cern.ch

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.

Networking & Security / 8**Signing (and encrypted) message handling and implications for admins.****Author:** Owen Synge¹¹ DESY (HH)**Corresponding Author:** owen.synge@desy.de

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.

Site Reports / 9**GridKa Site Report****Author:** Manfred Alef¹¹ Karlsruhe Institute of Technology (KIT)

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

- Hardware status
- Batch system issues
- CPU performance

Cloud, grid and virtualization / 10**Moving virtual machines images securely between sites.**

Author: Owen Synge¹

¹ DESY (HH)

Corresponding Author: owen.synge@desy.de

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.

Cloud, grid and virtualization / 11**Virtualization at CERN: a status report**

Author: Ulrich Schwickerath¹

¹ CERN

Corresponding Author: ulrich.schwickerath@cern.ch

Virtualization at CERN: a status report

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

- CERN's 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.

Site Reports / 13**NDGF Site Report**

Author: Mattias Wadenstein¹

¹ NDGF

Corresponding Author: mattias.wadenstein@cern.ch

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

IT Infrastructure / 14

Implementing Service Management processes with Service-Now

Author: Zhechka Toteva¹

¹ CERN

Corresponding Author: zhechka.toteva@cern.ch

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.

Computing / 15

Performance Comparison of Multi and Many-Core Batch Nodes

Author: Manfred Alef¹

¹ Karlsruhe Institute of Technology (KIT)

Corresponding Author: manfred.alef@iwr.fzk.de

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.

IT Infrastructure / 16

Drupal at CERN

Author: Juraj Sucik¹

Co-author: Jarek Polok¹

¹ CERN

Corresponding Author: juraj.sucik@cern.ch

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.

Computing / 17

OpenMP Performance on Virtual Machines

Author: Jie Tao^{None}

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.

Site Reports / 18

CC-IN2P3 Site-Report update

Author: philippe olivero¹

¹ CC-IN2P3

Corresponding Author: philippe.olivero@cc.in2p3.fr

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

Networking & Security / 19

Computer security update

Author: Romain Wartel¹

¹ CERN

Corresponding Author: romain.wartel@cern.ch

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.

Cloud, grid and virtualization / 20

StratusLab Marketplace for Sharing Virtual Machine Images

Author: Cal Loomis¹

¹ CNRS/LAL

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.

IT Infrastructure / 21

Scientific Linux Status Report + Discussion

Author: Troy Dawson¹

¹ FERMILAB

Corresponding Author: dawson@fnal.gov

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

Storage & FileSystems / 22

Evaluation of gluster file system at IHEP

Author: Yaodong CHENG¹

¹ Institute of High Energy Physics, Chinese Academy of Sciences

Corresponding Author: chyd@ihep.ac.cn

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.

Site Reports / 23

RAL Site Report

Author: Martin Bly¹

¹ *STFC-RAL*

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

Update on activities at RAL Tier1

IT Infrastructure / 24

Indico - Present and future

Author: Jose Benito Gonzalez Lopez¹

¹ *CERN*

Corresponding Author: jose.benito.gonzalez@cern.ch

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.

Storage & FileSystems / 25

Benefits of a virtualized approach to mass storage system

Author: Dorin Daniel Lobontu¹

¹ *Karlsruhe Institute of Technology (KIT)*

Corresponding Author: dorin-daniel.lobontu@kit.edu

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"

Site Reports / 26**Petersburg Nuclear Physics Institute (PNPI) status report**

Author: Andrey Shevel¹

¹ *Petersburg Nuclear Physics Institute (PNPI)*

Corresponding Author: andrey.shevel@cern.ch

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.

Oracle / 27**Open Source at Oracle**

Author: Gilles Gravier¹

¹ *Oracle*

Corresponding Author: gilles.gravier@oracle.com

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.

Site Reports / 28**CERN site report**

Author: Helge Meinhard¹

¹ *CERN-IT*

Corresponding Author: helge.meinhard@cern.ch

News from CERN since last meeting

IT Infrastructure / 29**Version Control Services at CERN**

Author: Alvaro Gonzalez Alvarez¹

¹ *CERN*

Corresponding Author: alvaro.gonzalez.alvarez@cern.ch

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.

Storage & FileSystems / 30

CASTOR status and development

Author: Eric Cano¹

¹ CERN

Corresponding Author: eric.cano@cern.ch

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.

Cloud, grid and virtualization / 31

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

Author: Ian Gable¹

Co-author: Randall Sobie¹

¹ University of Victoria

Corresponding Author: igable@uvic.ca

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.

Computing / 32

Batch Monitoring and Testing

Author: Jerome Belleman¹

¹ CERN

Corresponding Author: jerome.belleman@cern.ch

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.

Oracle / 33

An introduction to Oracle Linux

Author: Lenz Grimmer¹

¹ Oracle

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.

Oracle / 34

Discussion

IT Infrastructure / 35

Overview of the new Computing room at CC-IN2P3

Author: Pascal Trouve¹

¹ CC-IN2P3

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.

IT Infrastructure / 36

FAIR 3D Tier-0 Green-IT Cube

Author: Volker Lindenstruth¹

¹ FIAS, GSI

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 overall FAIR computing concept is presented as well as the FAIR Tier-0 data center architecture.

Site Reports / 37

DESY Site Report

Author: Wolfgang Friebel¹

¹ *Deutsches Elektronen-Synchrotron (DESY)-Unknown-Unknown*

Corresponding Author: wolfgang.friebel@desy.de

Computing at DESY news

Storage & FileSystems / 38

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

Authors: Dmitry Ozerov¹; Yves Kemp¹

¹ *DESY*

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.

Storage & FileSystems / 39

Lustre at GSI

Author: Thomas Roth¹

¹ *GSI*

Corresponding Author: t.roth@gsi.de

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.

IT Infrastructure / 40**Invenio at CERN****Author:** Jerome Caffaro¹¹ CERN**Corresponding Author:** jerome.caffaro@cern.ch

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.

IT Infrastructure / 41**CernVM-FS Production Service and Deployment****Author:** Ian Peter Collier¹¹ STFC RAL Tier1**Corresponding Author:** ian.peter.collier@cern.ch

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.

Site Reports / 42**Nikhef site report****Author:** Paul Kuipers¹¹ nikhef

Nikhef site report

Cloud, grid and virtualization / 43**Adopting Infrastructure as Code to run HEP applications**

Author: Mykhaylo Zynovyev¹

¹ *GSI*

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

Computing / 44

CMS 64bit transition and multicore plans

Author: Giulio Eulisse¹

¹ *FERMILAB*

Corresponding Author: giulio.eulisse@cern.ch

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.

Networking & Security / 45

HEPiX IPv6 Working Group

Author: David Kelsey¹

¹ *RAL*

Corresponding Author: d.p.kelsey@rl.ac.uk

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.

Site Reports / 46

SLAC Site Report

Author: Randy Melen¹

¹ *SLAC National Accelerator Laboratory*

Corresponding Author: randy@slac.stanford.edu

Update on activities at SLAC

Storage & FileSystems / 47

Discussion

Site Reports / 48

DLS site report

Author: Tina Friedrich¹

Co-authors: Frederik Ferner¹; Greg Matthews¹; Max Von Seibold¹; Nick Rees¹

¹ *Diamond Light Source Ltd*

Corresponding Author: tina.friedrich@diamond.ac.uk

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

Networking & Security / 49

Host based intrusion detection with OSSEC

Author: Bastian Neuburger¹

¹ *GSI*

Corresponding Author: b.neuburger@gsi.de

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.

Site Reports / 50

BNL Site report

Corresponding Author: rind@bnl.gov

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.

Site Reports / 51

ASGC site report

Author: Felix Lee¹

¹ *Academia Sinica*

Corresponding Author: felix@twgrid.org

ASGC current status.

Site Reports / 53

PSI - Site report

Author: Derek Feichtinger¹

¹ *PSI*

Corresponding Author: derek.feichtinger@cern.ch

Site report for the Paul Scherrer Institut.

Storage & FileSystems / 54

Evaluation of distributed file systems using trace and replay mechanism

Author: Jiri Horky¹

¹ *Institute of Physics of Acad. of Sciences of the Czech Rep. (ASCR)*

Corresponding Author: jiri.horky@cern.ch

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.

Cloud, grid and virtualization / 55

HEPiX VWG Status Report

Author: Tony Cass¹

¹ *CERN*

Corresponding Author: tony.cass@cern.ch

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.

Cloud, grid and virtualization / 56

The National Grid Service Cloud

Authors: David Fergusson¹; David Wallom²; Matteo Turilli²; Steve Thorn¹

¹ *University of Edinburgh*

² *University of Oxford*

Corresponding Author: steve.thorn@ed.ac.uk

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.

Introduction / 58

Registration

Introduction / 59

Welcome Address

Corresponding Author: k.langanke@gsi.de

Wrap-up / 60

Wrap-Up

Corresponding Author: jouvin@lal.in2p3.fr