

HEPiX Spring 2010 Workshop

Monday, 19 April 2010 - Friday, 23 April 2010

LNEC

Book of Abstracts

Contents

Petersburg Nuclear Physics Institute (PNPI) status report	1
Progress Report 2010 for HEPiX Storage Working Group	1
Evaluation of NFS v4.1 (pNFS) with dCache	1
Virtualisation for Oracle databases and application servers	1
GEMSS: Grid Enabled Mass Storage System for LHC experiments	2
Lavoisier : a way to integrate heterogeneous monitoring systems.	2
CERN site report	3
Virtualization at CERN: a status report	4
Scientific Linux Status Report and Plenary Discussion	4
Fermilab Site Report	4
LCLS Data Analysis Facility	5
virtual machines over PBS	5
Lustre-HSM binding	5
BNL RHIC/ATLAS Computing Facility Site Report	6
CERN Lustre evaluation and storage outlook	6
LCLS Data Analysis Facility	6
RAL Site Report	7
CESGA Experience with the Grid Engine batch system	7
Building up a high performance data centre with commodity hardware	8
Scientific Computing: first quantitative methodologies for a production environment	8
Windows 7 Deployment at Cern	8
Jefferson Lab Site Report	9
NDGF Site Report	9

DESY site report	9
OpenAFS Performance Improvements: Linux Cache Manager and Rx RPC Library	10
TWiki at CERN: Past Present and Future	10
RAL Tier1 Quattor experience and Quattor outlook	10
Lessons Learned from a Site-Wide Power Outage	11
Site report from PDSF	11
Spacewalk and Koji at Fermilab	11
Virtual Network and Web Services (An Update)	11
Optimizing analysis dataflow for the ARC case	12
update on computer security	12
Hyperthreading influence on CPU performance	12
An Adaptive Batch Environment for Clouds	13
INFN Tier1 site report	13
PIC Site Report	13
T2/T3 discussion	14
Preliminary Measurements of Hep-Spec06 on the new multicore processor	14
The Portuguese WLCG Tier2: status and issues	14
The new Generations of CPU-GPU Clusters for e-Science	14
CERN Grid Data Management Middleware plan for 2010	15
SLAC Site Report	15
Management Information Systems and Information Systems Management: Two Sides of the Same Coin	15
Update on HEPiX Working Group on Virtualisation	16
Closing Remarks	16
Virtualization in the gLite Grid Middleware software process. Use-cases, technologies and future plans.	16
Introduction	16
LIP and Grid in Portugal	16
WG meeting (http://indico.cern.ch/conferenceDisplay.py?confId=92038)	17
WG meeting (http://indico.cern.ch/conferenceDisplay.py?confId=92038)	17

EGEE Site Deployment: The UMinho-CP case study	17
--	----

Site Reports / 0**Petersburg Nuclear Physics Institute (PNPI) status report**

Author: Andrey Shevel¹

¹ *Petersburg Nuclear Physics Institute (PNPI)*

Corresponding Author: andrey.shevel@cern.ch

In this short presentation I plan to describe the responsibility for Computing Systems Department (CSD): LAN (400 hosts), mail service for the Institute, other centralized servers, computing cluster. The computing needs for small physics laboratory, experience with the cluster at PNPI HEPD and plan for cluster future in connection to LHC data analysis are discussed in more details.

Summary:

Status report for computing facility at small physics laboratory

Storage and Filesystems / 2**Progress Report 2010 for HEPiX Storage Working Group**

Author: Andrei Maslennikov¹

¹ *CASPUR*

Corresponding Author: andrei.maslennikov@caspur.it

Storage and Filesystems / 3**Evaluation of NFS v4.1 (pNFS) with dCache**

Authors: Martin Gasthuber¹; Patrick Fuhrmann¹; Tigran Mkrtchyan¹; Yves Kemp¹

¹ *DESY*

Corresponding Authors: patrick.fuhrmann@desy.de, kofemann@gmail.com, yves.kemp@desy.de

The talk will be covering experience accumulated with NFS4.1/dCache at DESY.

Summary:

In particular, will be presented:

- 1) information on client kernel and distribution on how to get NFS 4.1 to work;
- 2) information on the status of security;
- 3) information on stability of the current test bed (DESY grid lab);
- 4) information on client behaviour on pool crashes;
- 5) performance compared with other file systems.

Virtualization / 4

Virtualisation for Oracle databases and application servers

Author: Carlos Garcia Fernandez¹

¹ CERN

Corresponding Author: carlos.garcia.fernandez@cern.ch

With the growing number of Oracle database instances and Oracle application servers, and the need to control the necessary resources in terms of manpower, electricity and cooling, virtualisation is a strategic direction that is being seriously considered at CERN for databases and application servers.

Oracle VM is the Oracle certified and supported server virtualisation solution for x86/x86_64 platforms.

In order to be able to use Oracle VM in large scale environments, it must be prepared so that it can be installed and configured in a very short time and without any human interaction as well as maintained in a central way. At CERN, this is done with the CERN ELFs (Extremely Large Fabric management system).

This presentation will explain the steps and work done to install and maintain 2 versions of Oracle VM: 2.1.5 and 2.2. The presentation will also cover the tests of Databases and Application servers running in this virtual environment.

Storage and Filesystems / 5

GEMSS: Grid Enabled Mass Storage System for LHC experiments

Authors: Alberto Forti¹; Alessandro Cavalli¹; Andrea Prosperini¹; Barbara Martelli¹; Daniele Gregori¹; Elisabetta Ronchieri¹; Luca Dell'Agnello¹; Luigi Li Gioi¹; Pier Paolo Ricci¹; Riccardo Zappi¹; Stefano Dal Pra¹; Vincenzo Vagnoni²; Vladimir Sapunenko¹

¹ CNAF

² INFN/BO

New Grid Enabled MSS solution based on IBM's products (GPFS and TSM) has been developed and commissioned into production for LHC experiments in Tier1 at CNAF. Architectural view as well as installation and configuration details will be presented. We will also present technical details of data migration from CASTOR to GEMSS for CMS, LHCb and ATLAS experiments.

Monitoring & Infrastructure tools / 6

Lavoisier : a way to integrate heterogeneous monitoring systems.

Author: Cyril L'Orphelin¹

Co-author: Sylvain Reynaud¹

¹ CNRS/CCIN2P3

Corresponding Author: cyril.lorphelin@cc.in2p3.fr

Lavoisier is an Open Source Tool developed at CC-IN2P3 used in the project EGEE/ LCG by the Operations Portal. Lavoisier is an extensible service designed to provide an unified view of data collected

from multiple heterogeneous data sources (e.g. database, Web Service,LDAP, text files, XML files...). This unified view is represented as XML documents, which can be queried using standard languages, such as XSLT (and soon: XPath, XQuery).

The access to data views can be done through different standard ways : SOAP, REST/JSON or REST/XML.

Lavoisier also improves performance by caching the generated data views. The caching mechanism can be tuned according to the characteristics of the data views and of the data sources, and according to the constraints of the use-cases. This tuning has no impact, neither on plug-ins, nor on user's code.

The global architecture has four layers: the legacy data sources, the adapters used to generate the data views, the cache mechanism with flexible refresh triggering rules, and the service exposing both SOAP and REST (with XML or JSON) interfaces.

Moreover the Lavoisier architecture permits a clear separation of roles and changes related to a role have no impact on other roles :

- the plug-in developer adds support for new technologies by developing adapters.
- the service administrator configures the adapters and cache management according to the characteristics of the data views and data sources, and according to the constraints of the use-cases.
- the user queries the data views.

This last year an effort has been put to ease the integration of data for service administrators by developing web interfaces which permits to :

- see the list of views and their characteristics
- see the status of these views and the eventual errors linked with these views
- see the configuration file
- evaluate the dependencies between views .

Lavoisier has proven effective in increasing maintainability of the Operations portal, by making its code independent from the technologies used by the data sources and from the data cache management policy. Its design and interfaces made easier writing reusable code, and good performances are easily obtained by tuning the cache mechanisms in an absolutely transparent way for the portal code.

Indeed, the different components work in a standardized way through the output of the Lavoisier Web Service. The translation of resource information into this standardized output is provided by different plug-ins.

Recently, a huge joint effort has been recently put into the configuration of Lavoisier, the structure of its caches and the rules of refreshing to have efficient, scalable and reliable data handling. Indeed, all information has been structured around the base component of the Operational Model: the site. We retrieve the global information about primary sources like GGUS (official ticketing system for EGEE) , GOC DB (official database for sites in EGEE), SAM (Monitoring System) , or Nagios (another Monitoring system) and we organize it by sites.

The main idea is to construct a summary of the different available information for a site: firstly, this organization permits to continue to work with the caches, even if a primary source is unavailable; then you access only information you need on the web page. Your information is structured around a synoptic view of the site and you don't access hundreds of times the primary sources but a subset of them with the site view.

Finally, we refresh the data sources only when we need it and only when an action has been triggered. Last but not least, it is very easy to add a new data source in this model. In the configuration file of Lavoisier you add the access to the primary source and also the split of this information per site.

This information is then readily available in the synoptic view of the site.

Site Reports / 7

CERN site report

Author: Helge Meinhard¹

¹ *CERN-IT*

Corresponding Author: helge.meinhard@cern.ch

Report on news at CERN-IT since the Fall meeting in Berkeley

Virtualization / 8

Virtualization at CERN: a status report

Author: Ulrich Schwickerath¹

Co-authors: Belmiro Rodrigues Moreira ¹; Ewan Rosch ¹; Sebastien Goasguen ²

¹ *CERN*

² *Clemson*

Corresponding Author: ulrich.schwickerath@cern.ch

During the HEPiX meeting 2009 in Berkeley several reports related to virtualization efforts at CERN were given. This presentation will be a status report of these projects, with the focus on service consolidation and batch virtualization. Recent developments like the evaluation of the Infrastructure Sharing Facility (ISF) from Platform computing, issues seen with the tools we use and the technical implementations, as well as the general status of the different projects will be reviewed, and first experiences in an operational mode will be presented.

Operating Systems and Applications / 9

Scientific Linux Status Report and Plenary Discussion

Author: Troy Dawson¹

Co-author: Connie Sieh ¹

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.
Also we will have a Plenary discussion to get feedback to and input for the Scientific Linux developers from the HEPiX community. This may influence upcoming decisions e.g. on distribution lifecycles, and packages added to the distribution.

Site Reports / 10

Fermilab Site Report

Author: Chadwick Keith¹

¹ *Fermilab*

Corresponding Author: chadwick@fnal.gov

Site Report from Fermilab

Summary:

Site Report from Fermilab

11

LCLS Data Analysis Facility

Author: Alf Wachsmann¹

¹ SLAC

Corresponding Author: alfw@slac.stanford.edu

I will present the experience and lessons learned from the first few month of real users taking real data at LCLS. From those lessons, a new offline data analysis facility was designed and will be installed before the next set of users arrive on May 6.

Virtualization / 12

virtual machines over PBS

Author: Marc Rodriguez Espadamala¹

Co-author: Christian Neissner²

¹ Marc Rodriguez

² Christian Neissner

Corresponding Author: marc.rodriguez@cern.ch

Since PIC is supporting the data management of several scientific projects it is faced to requirements that might not be possible to enable in a homogeneous infrastructure. When looking for an alternative approach we have chosen to export the specific requirements into a virtualized environment. In order to minimize the impact to our physical computing infrastructure it was necessary to develop tools that could be incorporated into our native batch system. A first solution is already put in production and can provide almost completely arbitrary software environments in a transparent manner to the end-user.

Storage and Filesystems / 13

Lustre-HSM binding

Author: Thomas LEBOVICI¹

¹ CEA

Corresponding Author: thomas.leibovici@cea.fr

Lustre-HSM binding is a collaborative effort between Sun and CEA.
This presentation will introduce the first version of Lustre-HSM:
HSM support model, its features, components, development status, and roadmap.
It will also give examples of importing files from an existing HPSS namespace, archiving and restoring data, releasing disk space...

Summary:

Overview, design and features, status, use-cases, demo

Site Reports / 15

BNL RHIC/ATLAS Computing Facility Site Report

Author: Christopher Hollowell¹

¹ *Brookhaven National Laboratory*

Corresponding Author: hollowec@bnl.gov

New developments at the RHIC/ATLAS Computing Facility (RACF) at Brookhaven National Laboratory.

Storage and Filesystems / 16

CERN Lustre evaluation and storage outlook

Authors: Arne Wiebalck¹; Tim Bell¹

¹ *CERN*

Corresponding Authors: tim.bell@cern.ch, arne.wiebalck@cern.ch

The CERN evaluation of Lustre as a potential solution for home directories, project space and physics storage has now completed. This experience will be presented along with outlook for bulk data storage and life cycle management in a capacity hungry and performance sensitive environment.

Storage and Filesystems / 18

LCLS Data Analysis Facility

Author: Alf Wachsmann¹

¹ *SLAC*

Corresponding Author: alfw@slac.stanford.edu

I will present the experience and lessons learned from the first few month of real users taking real data at LCLS. From those lessons, a new offline data analysis facility was designed and will be installed before the next set of users arrive on May 6.

Site Reports / 20**RAL Site Report**

Author: Martin Bly¹

¹ STFC-RAL

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

Developments at RAL

Grid and WLCG / 21**CESGA Experience with the Grid Engine batch system**

Authors: Alvaro Simon Garcia¹; Carlos Fernandez Sanchez¹; Esteban Freire Garcia¹; Javier Lopez Cacheiro¹; Pablo Rey Mayo¹; Ruben Diez Lazaro¹; Sergio Diaz Montes¹

¹ CESGA

Grid Engine (GE) is an open source batch system with a complete documentation and support for advanced features like the possibility to configure a shadow master host for failover purposes, support for up to 10.000 nodes per master server, application level checkpointing, array jobs, DRMAA, fully integrated MPI support, a complete administration GUI, a web-based accounting tool (ARCo), etc.

CESGA has been using GE in its systems during more than 5 years. Currently it is the only batch system used at CESGA both for the supercomputers and for grid clusters.

One of the last challenges was the integration of GE with the FINIS TERRAE supercomputer installed at CESGA, it is formed by 142 Itanium nodes of 16 CPUs and 128GB of memory each one. Several stress tests were done to check SGE behaviour in a large cluster, results will be shown. It also will be explained some special SGE configurations like:

- HP-MPI and Intel MPI integration: most of the jobs that run in Finis Terrae are MPI jobs.
- Checkpointing. It has been configured in the queue configuration.
- Exclusivity: possibility to request a node exclusively.
- SSH-less node configuration: remote connections between nodes are done using qssh instead of ssh. This is very important in the case of MPI jobs to avoid jobs expanding outside the nodes they have been assigned or trying to use more resources in a node than the ones assigned to the job.
- Interactive jobs: jobs run interactively using a shell interface.
- Application integration with GE: for example Gaussian, Gaussian is used only under the batch system, the batch job requirements are taken accordingly to Gaussian input.

In order to manage special user requirements, CESGA has developed a qsub wrapper to implement some additional functionalities like special resources, i.e. the possibility to request additional resources for the jobs than the established limits (memory, CPU time, number of processors, space on scratch disk, ...)

Some kind of jobs (challenges, priority agreements, ...) need to be prioritized, it will be explained how these jobs are prioritized using GE functionalities.

Thanks to the efforts of IC, LIP and CESGA currently GE is fully supported by gLite middleware.

In a standard configuration it requires a Computer Element (CE) for each grid infrastructure. CESGA has done several modifications to the standard configuration to support different grid projects (EGEE,

EELA, int.eu.grid, Ibergrid, and other regional grid projects) using just one centralized batch system. This way resources can be shared among different grids with minimal system administration overhead.

The batch system is shared using one single GE qmaster server and a shadow qmaster for fault tolerance purposes. Jobs are submitted from different sources but all jobs are collected in a single batch server that distributes them between all the available WN.

Storage and Filesystems / 22

Building up a high performance data centre with commodity hardware

Author: Andreas Haupt¹

¹ *Deutsches Elektronen-Synchrotron (DESY)*

High data throughput is a key feature in today's LHC analysis centres. This talk will present the approach of the DESY site Zeuthen to build up a high performance dCache out of commodity hardware. Additional optimisations to achieve further throughput enhancements will be mentioned, as well. By presenting current benchmarks we show competitive results compared to more expensive installations.

Monitoring & Infrastructure tools / 23

Scientific Computing: first quantitative methodologies for a production environment

Author: Alberto Ciampa¹

Co-author: Enrico Mazzoni¹

¹ *INFN, Sezione di Pisa*

In Data Center the necessity for the accounting of resources used by different groups is growing. In this talk (and the correspondent note) the context and the scientific computing activity will be defined. Taking as an example a medium size Data Center (Pisa INFN) an industrial production approach will be followed.

A first methodology for the quantitative evaluation of the following data will be proposed:

- production level
- resource utilization efficiency
- cost distribution.

The evaluation of the resource utilization and the cost accounting are based on the electric power consumption: it will be described that this value can be considered as an practical example. The same methodology can be extended to different kind of costs and resources.

As practical application the results of a survey regarding the INFN-Pisa are reported.

Operating Systems and Applications / 24

Windows 7 Deployment at Cern

Author: Michal Budzowski¹

Co-author: Andreas Wagner ¹

¹ *Cern*

Corresponding Author: michal.budzowski@cern.ch

Windows 7 is officially supported at CERN since the end of March. The official Windows 7 service was preceded by pilot project that was started in December 2009. We will present our experience gathered during pilot phase as well as trends and deployment plans for Windows 7 at CERN.

Site Reports / 25

Jefferson Lab Site Report

Author: Sandy Philpott¹

¹ *JLAB*

Corresponding Author: sandy.philpott@jlab.org

An update on JLab's Scientific Computing facilities, including the new ARRA cpu and gpu clusters for USQCD, Lustre, experimental physics' data analysis farm, data storage, and networking.

26

NDGF Site Report

Author: Mattias Wadenstein¹

¹ *NDGF*

Corresponding Author: mattias.wadenstein@cern.ch

Overview and recent news in NDGF and the associated sub-sites.

Site Reports / 27

DESY site report

Author: Wolfgang Friebe¹

¹ *Deutsches Elektronen-Synchrotron (DESY)*

Corresponding Author: wolfgang.friebe@desy.de

DESY site report

Summary:

DESY site report

Storage and Filesystems / 29**OpenAFS Performance Improvements: Linux Cache Manager and Rx RPC Library**

Authors: Jeffrey Altman¹; Simon Wilkinson¹

¹ *Your File System Inc.*

Corresponding Author: jaltman@your-file-system.com

This talk will focus on two areas of performance improvement that will be shipping with the OpenAFS 1.6 series this coming Summer.

Significant work on improving performance has been completed on the OpenAFS 1.5 branch. Linux memory management has been heavily overhauled, removing a number of deadlock conditions, and significantly speeding access to data which is already held in the local cache. We now make correct use of the Linux page cache, and use data cached in memory in preference to pulling it from local disk. Performance in real world applications (a web server with its working set in the local cache) has been more than doubled by these changes.

The Rx RPC library has also been a focus of considerable attention with funding from the U.S. Dept of Energy. A wide variety of implementation errors that resulted in lock contention, lost packets, and extended timeouts were identified and corrected.

Operating Systems and Applications / 30**TWiki at CERN: Past Present and Future**

Author: Pete Jones¹

¹ *CERN*

Corresponding Author: peter.l.jones@cern.ch

TWiki was introduced at CERN in 2003 following a request from a small group of software developers. The service soon grew in popularity and today there are over 7000 registered editors of 60000 topics.

This presentation takes a look at the current service, the problems experienced and a look at future developments for TWiki at CERN.

Monitoring & Infrastructure tools / 31**RAL Tier1 Quattor experience and Quattor outlook**

Author: Ian Peter Collier¹

¹ *STFC/RAL*

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

The RAL Tier1 is continuing to benefit from bringing more of the Tier1 systems under Quattor control. That experience will be presented along with a report on recent and planned developments with Quattor across the wider grid and non-grid environments.

Miscellaneous / 32**Lessons Learned from a Site-Wide Power Outage****Author:** John Bartelt¹¹ SLAC**Corresponding Author:** bartelt@slac.stanford.edu

On January 19th, 2010, a series of powerful thunderstorms knocked out all power to the SLAC site (and surrounding areas). While the SLAC Facilities personnel worked to restore power to the computer center and the rest of the site, Computing Division planned for restoring services. This included setting priorities for such things as business systems needed to prepare the payroll and servers needed by the Fermi Gamma-Ray Space Telescope to process their data in a timely manner. I will describe problems that arose, what solutions were found, and what lessons were learned to minimize the impact of futures and streamline recovery from them.

Site Reports / 34**Site report from PDSF****Author:** Jay Srinivasan¹¹ Lawrence Berkeley National Lab. (LBNL)-Unknown-Unknown**Corresponding Author:** jay@nersc.gov

We will present a report on the current status of the PDSF cluster at NERSC/LBNL.

Monitoring & Infrastructure tools / 35**Spacewalk and Koji at Fermilab****Author:** Troy Dawson¹¹ FERMILAB**Corresponding Author:** dawson@fnal.gov

Spacewalk is an open source Linux systems management solution. It is the upstream community project from which the Red Hat Network Satellite product is derived.

Koji is Fedora's build platform.

Fermilab is currently testing Spacewalk for various roles. We are testing it for desktop and server monitoring, system configuration, and errata information. Although we currently are still in the testing stage, we have been impressed with it's functionality and development.

Fermilab's Scientific Linux development team is looking at Koji for the building of SL6 and possibly SL5. Many of it's built in functionality's make building packages easier, and it's reporting tools are very nice. But there is a steep learning curve when installing and setting up Koji.

Virtualization / 37

Virtual Network and Web Services (An Update)

Author: Thomas Finnern¹

¹ DESY

Corresponding Author: thomas.finnern@desy.de

This talk describes techniques for providing highly available network services with help of a hardware based load balancing cluster handled by a sophisticated control software. An abstract representation of web, mail and other services is implemented on base of application specific routing on network layers 3 to 7. Combining application, network and security issues in one point you may change server configurations or enhance availability and performance without service interruption. This is an update talk with some emphasis on the application and server side.

38

Optimizing analysis dataflow for the ARC case

Author: Mattias Wadenstein¹

¹ NDGF

Corresponding Author: mattias.wadenstein@cern.ch

The huge amount of data needed for analysis jobs made the naive data caching approach used by the ARC grid manager fall over. This talk gives details on the solutions implemented and under implementation to make an efficient distributed analysis facility in the ATLAS NDGF cloud.

39

update on computer security

Authors: Romain Wartel¹; Ulrich Schwickerath¹

¹ CERN

Corresponding Author: romain.wartel@cern.ch

In the recent years, High Energy Physics sites have significantly improved their collaboration and are providing services to users from a growing number of locations. The resulting attack surface, along with the increased sophistication of the attacks, has been a decisive change to encourage all the involved security teams to cooperate very closely together. News challenges in the security area have also appeared, including a more noticeable interest from the press in security incident handling. This presentation provides an outlook of these evolutions, along with several upcoming challenges and security risks that the community will need to deal with.

Benchmarking / 40

Hyperthreading influence on CPU performance

Author: Joao Marttins¹

Co-authors: Goncalo Borges ¹; Jorge Gomes ¹; Mario David ¹

¹ *LIP*

Corresponding Author: martinsj@lip.pt

In this report we present the hyperthreading influence on CPU performance when running the HEP-SPEC2006 benchmark suite in a 2 quadcore CPU shared memory system (HP BL460c G6). This study was performed as a function of the number of running instances (from eight to sixteen), and with and without SPEC RATE. We concluded that the elapsed application run time can be clearly reduced if hyperthreading is on. The effect becomes stronger as the number of running instances increases reaching up to 30% for the maximum number of tested running instances. We intend to extend this work for other shared memory systems, and try to define an upper limit for the CPU performance increase as a function of running instances.

Virtualization / 42

An Adaptive Batch Environment for Clouds

Author: Ian Gable¹

¹ *University of Victoria*

Corresponding Author: igable@uvic.ca

We have developed a method using Condor and a simple software component we call Cloud Scheduler to run High Throughput Computing workloads on multiple clouds situated at different sites. Cloud Scheduler is able to instantiate user-customized VMs to meet the requirements of jobs in Condor queues. Users supply Condor job attributes to indicate the type of VM required to complete their job. Condor then matches jobs with the appropriate VMs. Cloud Scheduler supports Nimbus and Amazon EC2, in addition we have preliminary support for Eucalyptus, and we plan support for OpenNebula 1.4. We present our early experiences with two separate projects using this method.

Site Reports / 43

INFN Tier1 site report

Author: Vladimir Sapunenko¹

¹ *INFN-CNAF*

Corresponding Author: vladimir.sapunenko@cern.ch

News from INFN-T1

44

PIC Site Report

Author: Francisco Martinez Ramirez¹

¹ *PIC*

Recent development and news from the PIC site, as well as an overview of the status of our installations, deployed software and current topics we are working on.

47

T2/T3 discussion

Benchmarking / 48

Preliminary Measurements of Hep-Spec06 on the new multicore processor

Author: Michele Michelotto¹

¹ *Univ. + INFN*

Corresponding Author: michele.michelotto@cern.ch

We had access to a remote server with 2 x 6174 amd processor.
Preliminary measurements will be show
Some early measurements on an Intel 5600 worker node will be exposed.

50

The Portuguese WLCG Tier2: status and issues

Author: Mario David¹

¹ *LIP Laboratorio de Instrumentaco e Fisica Experimental de Particulas*

Corresponding Author: david@lip.pt

The Portuguese WLCG Tier 2 for Atlas and CMS will be presented.
The status, performance and issues will be discussed

Keynote Speech / 51

The new Generations of CPU-GPU Clusters for e-Science

Author: Alberto Proença¹

¹ *Dep. Informática, Universidade do Minho*

Corresponding Author: aproenca@di.uminho.pt

The third paradigm of research known today as e-Science requires that the computational sciences community rely heavily on HPC clusters to run their simulations. At the core of each cluster lie

commodity devices with increasingly parallel capabilities. The two most successful processing devices today complement each other: the multicore CPUs and the manycore GPUs as accelerator vector devices. The main focus of this communication will be on an overview of some of the most promising approaches in this new breed of processing devices (and on the CUDA environment) to build the new generations of HPC clusters and how the scientific community has reacted to this challenge.

Grid and WLCG / 53

CERN Grid Data Management Middleware plan for 2010

Author: Oliver Keeble¹

¹ CERN

Corresponding Author: oliver.keeble@cern.ch

The 2010 program of work for the CERN Grid Data Management team is presented, describing the planned development activity for 2010. The plan for FTS, DPM/LFC and gfal/lcg_util is described along with longer term perspectives.

Site Reports / 54

SLAC Site Report

Author: Randy Melen¹

¹ SLAC

Corresponding Author: randym@slac.stanford.edu

SLAC site update since the Fall 2009 HEPiX meeting.

Keynote Speech / 56

Management Information Systems and Information Systems Management: Two Sides of the Same Coin

Author: Miguel Mira da Silva¹

¹ IST

Corresponding Author: mms@ist.utl.pt

The relationship between information systems and management outside large organizations has traditionally been one way only, with information systems supporting management but management not supporting information systems.

However, with information systems becoming ever more critical and complex, the need to manage information systems is becoming pervasive in most, if not all, organizations. The success rate of information systems projects, which *decreased* last year, is just an example that management is desperately needed in the IT world.

In this talk I will remember what is information systems and management, and then discuss the two relationships that exist between them. I will pay special attention to ITIL – which has been (wrongly) proposed as the solution for all kinds of problems – and its role in the much more general and interesting IT Governance area. The talk will finish with a reminder that innovation, not technology, is the source of economic growth.

Virtualization / 59

Update on HEPiX Working Group on Virtualisation

Author: Tony Cass¹

¹ *CERN*

The HEPiX working group on Virtualisation started its activities since the last HEPiX meeting. This talk will present the current status and future plans.

Closing remarks / 60

Closing Remarks

Corresponding Authors: sandy.philpott@jlab.org, jouvin@lal.in2p3.fr

Virtualization / 61

Virtualization in the gLite Grid Middleware software process. Use-cases, technologies and future plans.

Author: Lorenzo Dini¹

¹ *CERN*

Corresponding Author: lorenzo.dini@cern.ch

During the past years, the Grid Deployment and Grid Technology groups at CERN have gained experience in using virtualization to support the gLite Grid Middleware software-engineering process. Some technologies such as VNode and VMLoader have been developed to leverage virtualization in supporting activities such as development, build, test, integration and certification of software. This know-how has been collected in use-case analysis, usage statistics and requirements to help the development of a common virtualization solution strategy.

63

Introduction

Corresponding Authors: david@lip.pt, sandy.philpott@jlab.org, jouvin@lal.in2p3.fr

Site Reports / 64**LIP and Grid in Portugal****Author:** Goncalo Borges¹¹ *LIP Laboratorio de Instrumentaco e Fisica Experimental de Particulas***Corresponding Author:** goncalo@lip.pt**Virtualization Working Group F2F / 65****WG meeting (<http://indico.cern.ch/conferenceDisplay.py?confId=92038>)****Virtualization Working Group F2F / 66****WG meeting (<http://indico.cern.ch/conferenceDisplay.py?confId=92038>)****Grid and WLCG / 67****EGEE Site Deployment: The UMinho-CP case study****Author:** Tiago Sá¹**Co-authors:** António Pina ¹; Bruno Oliveira ¹¹ *Uminho*

EGEE is a dynamic organism with requirements that constantly evolve over time.

The deployment of UMinho-CP - an EGEE site supporting Civil Protection related activities -, revealed new challenges, some of which were overcome till now, while others are planned for future work.

A great effort has been made to automate the distribution, installation and configuration of an EGEE site, for efficient deployment and maintenance, making each site an agile intervenient of the European-wide grid infrastructure.

Summary:

This presentation is a retrospective of the Uminho-CP deployment and administration, how new ideas were implemented, and how they were used to tackle the evolutionary nature of the EGEE infrastructure.

The suggested implementation covers all aspects of the deployment of an EGEE site, from hardware to software and human resources.

UMinho-CP is fully supported by a set of Virtual Machines, running ScientificLinux[Cern] as guest Operating System, replacing the original Rocks distribution, RHEL, to better support gLite middleware requirements.

The EGEE roll - a self-contained software package feature, based on Rocks Clusters -, was created in order to install and configure a cluster to be part of the EGEE infrastructure. The results of the continuous development of the EGEE roll are presented. While the early version was focused on the installation and initial configuration of an EGEE site, current work aims to arm administrators with easy and efficient

ways to perform the most common administration tasks. The transition to the new gLite 3.2 middleware and 64bit-only architecture is also under heavy development, to ease the transition effort that site administrators will have to undertake.

This EGEE deployment example shall be useful, both to existing site administrators, as also for new players of the European Grid Initiative community.