

HEPiX Workshop

Monday 26 October 2009 - Friday 30 October 2009

Lawrence Berkeley National Laboratory

Book of Abstracts

Contents

A Vision for Virtualisation in WLCG	1
AMD Roadmap	1
Adopting Quattor for managing the UK Tier 1 fabric at RAL	1
Batch virtualization project at CERN	1
Benchmarking at LIP	2
Benchmarking of CPU servers	2
CC-IN2P3 Site Report	2
CERN site report	3
Cyber security update	3
Deploying and Using the Lustre Monitoring Tool	3
ESnet: Networking for Science	3
Evolution of virtual infrastructure with Hyper-V	4
Experiences with StoRM and Lustre at an Atlas Tier-2 site	4
First exercises with PROOF on NFS v4.1/pNFS	4
Grid Security Update	5
HEP-SPEC06 Measurement on Nehalem and Istanbul	5
HPSS in the Extreme Scale Era	5
IN2P3 HPSS Migration (v5.1 to 6.2) report	6
INFN Site Report	6
INFN-T1 status report	6
IRFU Site report	7
ITIL at CERN	7
ITIL's roles and tools from a perspective of a Scientific Computing Centre	7

Intel HPC environment for Silicon Design and Key Learnings	8
JLab Site Report	8
Keynote Speech by Kathy Yelick	8
LIP Site Report	8
Linux Desktop Management with old school NFS-root boot	9
Monitoring CC-IN2P3 services with Nagios	9
Monitoring tape drives and medias at CC-IN2P3	9
NDGF Site Report	10
NIKHEF site report	10
Network Performance Tuning	10
On-demand Virtualization and Grid/Cloud Integration	10
Open Source Solution for Monitoring of Grid Services in WLCG	11
Optimizing tape data access	11
Options for expanding CERN's computing capacity without a new Building	12
Performance of Hadoop file system on a HEPiX storage workgroup testbed	12
Prague Tier2 site report	12
RAL Site Report	13
Review of Desktop Computing Support	13
SLAC Site Report	13
Scientific Linux Status Report and Plenary Discussion	13
ScotGrid, A UK ATLAS Tier 2 Approaching Readiness for Data Taking	14
Security aspects of the WLCG infrastructure	14
Site Report GSI	14
Site report from PDSF	15
Storage R&D at CERN	15
TRIUMF site report	15
The CDCE Project @ BNL	15
The FermiGrid Software Acceptance Process	16
The Magellan Cloud Computing Project at NERSC	16
The WLCG Technical Forum and HEPiX	16

U.S. Government Supports OpenAFS	17
Unified Performance and Environment Monitoring using Nagios,Ganglia and Cacti	17
Update on Version Control Services at CERN	17
Virtualization within FermiGrid	18
Web application security	18
Windows 7 horizons	18
lustre@gsi: A Petabyte file system for the analysis farm - status and outlook	19

Virtualization II / 7**A Vision for Virtualisation in WLCG**Tony Cass¹¹ *CERN***Corresponding Author(s):** tony.cass@cern.ch

This talk presents a possible roadmap for the use of virtualisation across WLCG sites to deliver improved computing services for the experiments and users.

Benchmarking II / 49**AMD Roadmap**Jeff Underhill¹¹ *AMD*

HPC AMD Roadmap and Benchmarking

Monitoring Infrastructure and Tools II / 52**Adopting Quattor for managing the UK Tier 1 fabric at RAL**Ian Collier¹¹ *UK Tier1 Centre*

The UK Tier 1 Centre at RAL will increase in size significantly in the coming year. The need for better automation of both system deployment and ongoing configuration management has prompted a survey of possible solutions. Since deciding to adopt Quattor earlier this year, we have successfully deployed our new SL5 batch service using the system and are already seeing better consistency and easier management. The talk will discuss the options we considered and planning and carrying out the deployment of a complex management system while avoiding disruption of the running farm.

Virtualization I / 32**Batch virtualization project at CERN****Author(s):** Ulrich Schwickerath¹**Co-author(s):** Sebastien Goasguen² ; Tony Cass¹¹ *CERN*² *Clemson University*

Corresponding Author(s): ulrich.schwickerath@cern.ch

Between March and August 2009 a project has been set up at CERN with the aim to evaluate possibilities to use virtualization at a large scale, with the focus on batch computing. Two key issues have been identified for this specific application: the placing of virtual machines on an appropriate hyper-visor, and the selection of an appropriate image which should be driven by the actual demand. Both commercial and free software solutions exist which are able to solve the placing issue.

The virtual machine orchestrator, VMO, a commercial solution by Platform computing, and the free software solution OpenNebula have been evaluated during the project. For VMO, the vendor provided a first implementation of an algorithm for selecting the image to be deployed, which is driven by user requirements of pending batch jobs. For OpenNebula, an external mechanism needs to be developed to perform this task. In the presentation, the basic concepts of the project and lessons learned will be presented. Further visions and possible implications for services offered at CERN will be described.

Benchmarking II / 13

Benchmarking at LIP

Gonçalo Borges¹ ; Jorge Gomes¹ ; João Martins¹ ; Miguel Oliveira¹ ; Mário David¹

¹ *LIP - Laboratório de Instrumentação e Física Experimental de Partículas*

Benchmarking is a key activity on all computer centers not only for tender procedures but also to optimize resources. At LIP we underwent recently major upgrades of all sites and deployed a new one.

We report on CPU HEP-SPEC and storage benchmarking results.

Benchmarking I / 34

Benchmarking of CPU servers

Ulrich Schwickerath¹

¹ *CERN*

Corresponding Author(s): ulrich.schwickerath@cern.ch

The recent generation of Intel XEON CPUs comes with support for symmetric multi processing, formally known as hyperthreading. In addition, a new CPU feature has been added, the Intel Turbo mode. The influence of these features on the system performance has been tested using the HEP-SPEC06 benchmark suite, with enhanced statistics, and a study of the scaling behaviour has been done.

The results of these tests are presented, and consequences for certain applications as well as for procurement procedures themselves are discussed.

Site Reports I / 5

CC-IN2P3 Site Report

philippe olivero¹

¹ *CC-IN2P3*

Corresponding Author(s): philippe.olivero@cc.in2p3.fr

News, changes, upgrades occurred last year at CC-IN2P3.

Site Reports I / 28

CERN site report

Helge Meinhard¹

¹ *CERN-IT*

Corresponding Author(s): helge.meinhard@cern.ch

Summary of important news at CERN since the spring 2009 meeting

Network, Security I / 10

Cyber security update

Sebastian Lopienski¹

¹ *CERN*

Corresponding Author(s): sebastian.lopienski@cern.ch

This talk gives an update on security issues affecting computers, software applications and networks during the last months. It includes information on emerging types of vulnerabilities and recent attack vectors, and provides an insight into the cyber-crime economy of 2009. This talk is based on contributions and input from the CERN Computer Security Team.

Monitoring Infrastructure and Tools II / 54

Deploying and Using the Lustre Monitoring Tool

Andrew Uselton¹

¹ *NERSC/LBNL*

The Lustre Monitoring Tool (LMT) provides a useful view of the server-side behavior of the Lustre parallel file system. This talk presents a brief overview of the architecture of the tool and explores several use cases including tracking system health, server-side performance tuning, applications-side performance tuning, and incident evaluation, among others.

Network, Security III / 48**ESnet: Networking for Science**Joe Burrencia¹¹ *ESnet*

This talk will cover an introduction to the Energy Sciences Network (ESnet). It will outline both ESnet as a program in DOE's Office of Science and the current network implementation known as ESnet4. The roles of both components of ESnet4, the IP network and Science Data Network (SDN) will be discussed. The presentation will also touch upon several international network collaborations that ESnet is taking a key role in to advance the usability of R&D networks to support science. Lastly, the talk will cover ESnet's participation in ARRA funded research into the demonstration of 100Gbs wide area networks.

Virtualization II / 22**Evolution of virtual infrastructure with Hyper-V**Juraj Sucik¹¹ *CERN***Corresponding Author(s):** juraj.sucik@cern.ch

The Internet services group provides the infrastructure and sophisticated management tools for virtual machine provisioning based on Hyper-V, Microsoft Virtual Machine Manager, management SOAP web services and a user web interface. This virtualisation service has already confirmed its reliability and efficiency by wide range of satisfied users. The infrastructure which was presented at the last HEPiX meeting has undergone significant improvements, which allowed us to provide new features - live migration, rapid provisioning and better Linux support. This talk will present these important updates of our infrastructure and summarize the experience gained from running the Linux operating system in the virtual machines.

Storage I / 36**Experiences with StoRM and Lustre at an Atlas Tier-2 site**Christopher J Walker¹¹ *Queen Mary, University of London***Corresponding Author(s):** c.j.walker@qmul.ac.uk

Queen Mary, University of London has been using the StoRM SRM in front of a Lustre filesystem. We present the results of benchmarks on the Lustre filesystem, and the throughput from simulated analysis using the hammercloud framework.

Storage II / 60**First exercises with PROOF on NFS v4.1/pNFS**

Peter van der Reest¹

¹ *DESY*

Corresponding Author(s): peter.van.der.reest@desy.de

At DESY we have installed a testbed for Grid and storage related issues. One of the first trials we did was running a ALICE Proof job against data held in a distributed NFS 4.1 service based on dCache. Also trails were made against an industrial system, of which I am not allowed to speak yet. Comparisons we made with regard to scaling in PROOF and IO performance in comparison to NFS3 To make it more savvy, we have an accompanying movie which shows the running computation and the IO received.

Network, Security I / 21

Grid Security Update

David Kelsey¹

¹ *RAL*

Corresponding Author(s): d.p.kelsey@rl.ac.uk

An update on Grid security in WLCG, EGEE and EGI, concentrating on progress in operational and policy issues.

Benchmarking I / 9

HEP-SPEC06 Measurement on Nehalem and Istanbul

Michele Michelotto¹

¹ *Univ. + INFN*

Corresponding Author(s): michele.michelotto@cern.ch

Performances of the last intel DP processor Nehalem 55xx and last amd DP processor Istanbul 24xx using the HEP-SPEC06 benchmark.

Storage II / 38

HPSS in the Extreme Scale Era

Jason Hick¹

¹ *LBNL*

Corresponding Author(s): jhick@lbl.gov

The High Performance Storage System (HPSS) has served the DOE community for high performance archival storage for the past fifteen years. It specifically serves the HEP community at LBNL/NERSC by providing archival storage from the PDSF system. This presentation will provide a brief overview

of how HPSS works, what its current more unique features are, and what our plans are for our next major release (8.1) and thoughts on preparing for Extreme Scale (2018-2020).

Storage II / 24

IN2P3 HPSS Migration (v5.1 to 6.2) report

Pierre Emmanuel Brinette¹

¹ CNRS-CCIN2P3

Corresponding Author(s): pbrinette@cc.in2p3.fr

IN2P3 Computing Center has been using HPSS as a Mass Storage System since 1999. There has been no major system upgrade since 2005 and IN2P3 still runs HPSS 5.1. This version is no more supported by IBM and doesn't include the T10K-B drive support.

In June 2009, the system has been upgraded to HPSS 6.2.2.2. This upgrade implies major changes (DCE removal, DB2 and systems upgrade).

This presentation will expose all the operations done to upgrade the system during a 3 day downtime, and the issues encountered.

- HPSS new 6.2 features and changes.
- Operation planning.
- Systems preparation.
- Core server Migration.
- Metadata Migration.
- Issues encountered.

Site Reports IV / 17

INFN Site Report

Roberto Gomezell¹

¹ INFN

Corresponding Author(s): roberto.gomezell@ts.infn.it

An overview about INFN computing and networking activities

Site Reports III / 14

INFN-T1 status report

Andrea Chierici¹

¹ INFN-CNAF

I will present the status report of the Italian Tier1 site

Summary:

I will present the status report of the Italian Tier1 site

Site Reports II / 2

IRFU Site report

Pierrick Micout¹

¹ *CEA IRFU*

Corresponding Author(s): pierrick.micout@cea.fr

What is new at IRFU Saclay.

Summary:

Presentation on the news related to computing in IRFU Saclay:

New computing room

Evolution of the GRIF IRFU site

Evolution of the analysis facility

Windows

ITIL / 8

ITIL at CERN

Tony Cass¹

¹ *CERN*

Corresponding Author(s): tony.cass@cern.ch

This talk will cover progress at CERN to improve service organisation through adoption of ITIL principles.

ITIL / 43

ITIL's roles and tools from a perspective of a Scientific Computing Centre

Achim Grindler¹ ; Holger Marten¹ ; Tobias Koenig¹

¹ *Karlsruhe Institute of Technology (KIT)*

Corresponding Author(s): tobias.koenig@kit.edu

The Karlsruhe Institute of Technology (KIT) was founded at 1st October 2009, merging the University of Karlsruhe and the Forschungszentrum Karlsruhe. As a first new organizational unit of the KIT the

Steinbuch Centre for Computing (SCC) was established, combining the former Institute for Scientific Computing of Forschungszentrum Karlsruhe and the Computing Centre of the Technical University Karlsruhe. The KIT merge process directly influenced the SCC which has to cover two locations with a distance of 10 km.

The IT service management, according to the industrial standard "IT Infrastructure Library (ITIL), was selected by the SCC as a strategic element to support the merging of the two existing computing centres. The service support processes of ITIL like: Incident, Problem, Change, Configuration and Release Management are the basis of SCC's first class IT Services. The talk explains the different roles and tools of each ITIL support process and it puts attention to the special needs of GridKa, the German Tier-1 centre of the WLCG infrastructure, hosted at the SCC.

Other (O/S, Appls., Data centers/Facilities) I / 58

Intel HPC environment for Silicon Design and Key Learnings

Corresponding Author(s): shesha.krishnapura@intel.com

Silicon design technical complexity is increasing every year due to several new features and process technology shrinks. Additionally, the business drivers such as shorter product development time, reduced headcount, and lower cost is increasing pre-silicon verification, high degree of design automation, and global multi-site design teams. These two factors (technological and business) are astronomically increasing demand on computing and storage driving design computing to be engineered as an HPC environment. This presentation will cover Intel HPC design compute environment, generational improvements, and realized value in the areas of compute clusters, very high large memory servers, optimal network, and parallel storage.

Site Reports I / 35

JLab Site Report

Sandy Philpott¹

¹ *JLAB*

Corresponding Author(s): sandy.philpott@jlab.org

Status of Scientific Computing at JLab, including experimental physics and high performance computing for Lattice QCD.

Keynote Speech / 57

Keynote Speech by Kathy Yelick

Site Reports III / 12

LIP Site Report

Gonçalo Borges¹ ; Jorge Gomes¹ ; João Martins¹ ; Miguel Oliveira¹ ; Mário David¹

¹ *LIP - Laboratório de Instrumentação e Física Experimental de Partículas*

Site report for all sites (LIP-Lisbon,LIP-Coimbra,NCG-INGRID-PT) and activities at LIP.

Desktop Management I / 46

Linux Desktop Management with old school NFS-root boot

Christopher Huhn¹

¹ *GSI Helmholtzzentrum für Schwerionenforschung GmbH*

Corresponding Author(s): c.huhn@gsi.de

The standard Linux desktops at GSI received their operating system via NFS and not from the local hard drive since the last millennium.

Since then this policy has been enhanced to work with shared read-only OS images that provide advanced security and fast OS deployment and upgrades. Image generation and configuration management of NFS-root desktops is completely integrated into our infrastructure for standalone servers.

My talk will give an overview of the currently used techniques, the shortcomings of this approach and an outlook on our future plans.

Monitoring Infrastructure and Tools II / 33

Monitoring CC-IN2P3 services with Nagios

marc hausard¹

¹ *CC-IN2P3*

Corresponding Author(s): marc.hausard@cc.in2p3.fr

At CC-IN2P3, Nagios have taken over from the previous system to become the main monitoring tool used by the operation. This presentation will introduce its configuration in a Tier-1 environment and will present various extra features developed at CC-IN2P3 in order to customize the notification system and to provide multi-user development and failover mechanism.

Monitoring Infrastructure and Tools I / 23

Monitoring tape drives and medias at CC-IN2P3

Frédéric AZEVEDO¹

¹ *CC-IN2P3*

Corresponding Author(s): frederic.azevedo@cc.in2p3.fr

Due to the continuous load and intensive usage on our robotics, we regularly face some hardware issues with tapes and tape drives. A recurrent issue concerns possible data loss which leads to go through a long recovery process.

In order to improve our reliability, we have studied commercial solutions to avoid permanent write/read errors, or at least foresee occurring errors. We've tested two products (one month period each) and purchased the one that met our requirements best.

In this talk I'll expose the criteria to select the product, our daily usage after 4 months of use and finally what we expect to do with/around it in the near future.

Site Reports IV / 44

NDGF Site Report

Mattias Wadenstein¹

¹ *NDGF*

Corresponding Author(s): mattias.wadenstein@cern.ch

News and overview of the whats happening in the NDGF region, as well as some small updates on previously covered topics in NDGF-related matters.

Site Reports II / 61

NIKHEF site report

Paul Kuipers¹

¹ *NIKHEF*

NIKHEF site report

Network, Security III / 55

Network Performance Tuning

Brian Tierney¹

¹ *NERSC/LBNL*

Corresponding Author(s): bltierney@lbl.gov

Virtualization II / 15

On-demand Virtualization and Grid/Cloud Integration

Andrea Chierici¹

¹ *INFN-CNAF*

INFN-T1 implemented a solution to get Worker-nodes on demand using virtualization technology. This solution is allowing us extreme flexibility providing dynamic virtual execution environments and integrating seamlessly into our production grid. Currently 200 VMs slots are available and will increase further.

Summary:

INFN-T1 implemented a solution to get Worker-nodes on demand using virtualization technology. This solution is allowing us extreme flexibility providing dynamic virtual execution environments and integrating seamlessly into our production grid. Currently 200 VMs slots are available and will increase further.

Monitoring Infrastructure and Tools I / 41

Open Source Solution for Monitoring of Grid Services in WLCG

Wojciech Lapka¹

¹ *Unknown*

Corresponding Author(s): wojciech.lapka@cern.ch

Since 2005 Worldwide LHC Computing Grid (WLCG) services have been monitored by the Service Availability Monitoring (SAM) system which has been the main source of information for the monthly WLCG availability and reliability calculations.

During this time SAM framework gained popularity amongst site and service managers and was very useful in building robust grid infrastructure.

Experience with this monitoring tool as well as preparation to the evolution of the European grid infrastructure from EGEE to national grid initiatives (NGI) led to design of the enhanced and distributed model for monitoring grid services. Nagios has been adopted as a monitoring framework and messaging technology (ActiveMq) has been chosen as a transport mechanism.

This talk covers the architecture of the new system.

Storage II / 45

Optimizing tape data access

Jonathan Schaeffer¹

¹ *CCIN2P3*

Corresponding Author(s): jonathan.schaeffer@cc.in2p3.fr

TReqS is our Tape Request Scheduler. Based on BNLBatch, its goal is to get between dCache and HPSS and to reorder the files requests.

Since May of this year, a first implementation hit our production system.

We will present here :

- the problematic of tape access for LHC experiments
- the solution we implemented
- TReqS in its production environment, our 5 month experiment with the beast

Summary:

All LHC experiments use dCache as a front end to file access at most T1 sites.

At CCIN2P3, dCache is interfaced with HPSS. While writing to the mass storage system is effective, reading from it has shown to be most unsatisfying.

The major cause of this issue is that dCache has no knowledge about the files location on tape and the reading requests are rather chaotic from the HPSS point of view. It generates a lot of useless tape mount and dismount and the average dCache activity impacts the HPSS system badly.

We present here TReqS, a method for scheduling the file requests and submit them to HPSS to ensure a staging order and minimize the access latency due to the tape movements in the library. TReqS is based on BNLBatch, itself based on OakRidge batch Scheduler.

We will present the global concept and show the results and the conclusions of 6 month of tape scheduling.

Other (O/S, Applns., Data centers/Facilities) I / 37

Options for expanding CERN's computing capacity without a new Building

Alan Silverman¹

¹ *CERN*

Corresponding Author(s): alan.silverman@cern.ch

CERN is approaching the limit of what can be housed in its Computer Centre but there is no clear consensus about the provision of new capacity. While discussions continue, CERN has decided to take two interim measures, partly to satisfy immediate and medium-term needs and partly to gain experience in these domains, namely external hosting and the acquisition and operation of container-based solutions. I will describe the options considered for each of these and expose our current plans.

Storage I / 16

Performance of Hadoop file system on a HEPiX storage workgroup testbed

Artem Trunov¹

¹ *Karlsruhe Institute of Technology*

Corresponding Author(s): artem.trunov@cern.ch

This work is continuation of storage solution testing performed by HEPiX storage workgroup on it's testbed at FZK. Hadoop, an Apache project, offers a cluster file system called HDFS inspired by Google File System and designed to run on commodity hardware. It has gained some popularity in OSG, where it has become a supported storage solution, and is currently in production at a few T2 sites.

In this series of tests we used HEPiX testbed worker nodes' hard drives as a underlying storage, without using external file servers or storage arrays. We used a standard HEPiX storage application suite to evaluate performance of this solution. We present obtained results in this paper.

Site Reports IV / 47

Prague Tier2 site report

Jan Svec¹

¹ *site administrator at Prague Tier2*

Corresponding Author(s): svecj@fzu.cz

Prague Tier2, its current status and plans for nearest future.

Site Reports II / 6

RAL Site Report

Martin Bly¹

¹ *STFC-RAL*

Corresponding Author(s): martin.bly@stfc.ac.uk

Latest news from the RAL Tier1

Desktop Management I / 59

Review of Desktop Computing Support

We will review the Desktop Computing Support across some HEP sites.

Site Reports IV / 40

SLAC Site Report

John Bartelt¹

¹ *SLAC*

Corresponding Author(s): bartelt@slac.stanford.edu

Report on new personnel and projects, and status of IT and HPC at SLAC.

Other (O/S, Applns., Data centers/Facilities) II / 0

Scientific Linux Status Report and Plenary Discussion

Author(s): Troy Dawson¹

Co-author(s): Connie Sieh¹

¹ *FERMILAB*

Corresponding Author(s): 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 II / 42

ScotGrid, A UK ATLAS Tier 2 Approaching Readiness for Data Taking

Author(s): Douglas McNab¹

Co-author(s): Graeme Stewart¹; Michael Kenyon¹; Sam Skipsey¹; Stuart Purdie¹

¹ *University of Glasgow*

This presentation will provide an overview of the very successful UK ATLAS Tier-2, ScotGrid, as we fast approach LHC data taking and data analysis. This will cover a variety of topics, ranging from an overview of the fabric and middleware, the site's current readiness, success during STEP, storage issues, site optimisations all the way through to disaster planning and site security. The presentation will conclude with a look into the future and how we can retain our position as one of the most successful ATLAS Tier-2 centres.

Network, Security II / 27

Security aspects of the WLCG infrastructure

Maarten Litmaath¹

¹ *CERN*

Corresponding Author(s): maarten.litmaath@cern.ch

The Worldwide LHC Computing Grid (WLCG) infrastructure has been built up for the storage and analysis of the very large data volumes that will be recorded by the LHC experiments. Its existing security mechanisms and policies are foreseen to evolve in various respects, for example with an increasing use of virtual machines, pilot jobs, clouds, enhancements to data storage and access models, and potential integration with single sign-on campus-wide or federated identity management systems. To steer such evolution, input from the HEPiX community would be very desirable.

Site Reports I / 25

Site Report GSI

Walter Schoen¹

¹ *GSI*

Corresponding Author(s): w.schoen@gsi.de

Site report GSI

Site Reports II / 1

Site report from PDSF

Jay Srinivasan¹

¹ *Lawrence Berkeley National Lab. (LBNL)-Unknown-Unknown*

Corresponding Author(s): jay@nersc.gov

We present the current status of PDSF and updates since the last HEPiX meeting.

Storage II / 30

Storage R&D at CERN

Andras Horvath¹ ; Arne Wiebalck¹ ; Helge Meinhard¹

¹ *CERN-IT*

Corresponding Author(s): helge.meinhard@cern.ch

This talk will present an update on R&D activities around storage at CERN. The main focus will be various activities around iSCSI technology, but an update will also be presented on the Lustre evaluation project.

Site Reports III / 39

TRIUMF site report

Kelvin Raywood¹

¹ *TRIUMF*

Corresponding Author(s): kray@triumf.ca

An external review of TRIUMF computing took place and some changes have been recommended. The Tier-1 Center has completed its acquisition for the 2009/2010 upgrade. We continue to use a mix of Xen, OpenVZ and Hyper-V for virtualisation with OpenVZ being preferred for hosted servers. For Linux desktops, we provide a repository of TRIUMF rpms which customise a standard Scientific Linux installation.

Other (O/S, Applns., Data centers/Facilities) I / 19

The CDCE Project @ BNL

Eric Blum¹ ; Michael Ernst¹ ; Richard Hogue¹ ; Tony Chan¹

¹ *Brookhaven National Laboratory*

Corresponding Author(s): awchan@bnl.gov

This presentation will describe the expansion of the RHIC/ATLAS Computing Facility (RACF) to accommodate its commitments to the computational needs of the scientific programs at Brookhaven National Laboratory. The expansion has nearly tripled the footprint of the facility over the past 2+ years and allows the RACF to adequately meet our computing and storage requirements for the foreseeable future. The presentation will describe the challenges faced during the design, construction and commissioning phase of the project and will also provide an update on the current status and plans for the newly available floor space.

ITIL / 4

The FermiGrid Software Acceptance Process

Chadwick Keith¹

¹ *Fermilab*

Corresponding Author(s): chadwick@fnal.gov

The software acceptance process that is used by FermiGrid together with the operational experience of operating using this process will be presented.

Virtualization I / 51

The Magellan Cloud Computing Project at NERSC

Brent Draney¹

¹ *NERSC/LBNL*

NERSC and the Argonne LCF have been funded by DOE to acquire test systems to explore cloud computing technologies. We present an overview of the Cloud Computing Project at NERSC.

Other (O/S, Applns., Data centers/Facilities) II (contd.) / 29

The WLCG Technical Forum and HEPiX

Maarten Litmaath¹

¹ *CERN*

Corresponding Author(s): maarten.litmaath@cern.ch

The Worldwide LHC Computing Grid (WLCG) Technical Forum has been set up for discussions between WLCG stakeholders about middleware etc. in view of improving the reliability and efficiency of the WLCG infrastructures.

HEPiX is a good venue for discussions pertaining to the operation, usage and evolution of computing and storage facilities, from the perspectives of WLCG sites as well as the LHC experiments. Some topics of interest: virtual machines, clouds, pilot jobs, efficient data access, security.

Storage I / 18

U.S. Government Supports OpenAFS

Jeffrey Altman¹

¹ *Your File System Inc.*

Corresponding Author(s): jaltman@your-file-system.com

The U.S. Department of Energy has awarded Your File System Inc. a US\$648,000 Small Business Innovative Research Phase II grant to support the development of a next generation globally distributed file system that is compatible with AFS. This talk will describe the technologies that Your File System Inc. will be implementing and contributing to OpenAFS through August 2011.

Monitoring Infrastructure and Tools I / 53

Unified Performance and Environment Monitoring using Nagios, Ganglia and Cacti

Thomas Davis¹

¹ *NERSC/LBNL*

We present a method of monitoring the environment and performance using open source tools such as Nagios, Ganglia and Cacti to collect and display performance data as well as availability information for various components of large computing systems in an integrated fashion. We will present information on how the data is collected, viewed and analyzed, with specific examples from NERSC's Cray system.

Other (O/S, Appls., Data centers/Facilities) II (contd.) / 31

Update on Version Control Services at CERN

Author(s): Giacomo Tenaglia¹

Co-author(s): Alvaro Gonzalez Alvarez ¹; Artur Wiecek ¹; Jonatan Hugo Hugosson ¹; Juan Manuel Guijarro ¹

¹ *CERN*

Corresponding Author(s): giacomo.tenaglia@cern.ch

CERN Central Subversion Service was started as a pilot project on January 2008, and since January 2009 is an official service offered by CERN IT to CERN users. In the long distance it is meant to replace CERN Central CVS Server.

This talk will present an overview of the CERN Version Control Services lifecycles, with an emphasis on community-driven Service Design, and Service Operation integrated with CERN IT infrastructure.

Virtualization I / 3

Virtualization within FermiGrid

Chadwick Keith¹

¹ *Fermilab*

Corresponding Author(s): chadwick@fnal.gov

The current virtualization infrastructure in use within FermiGrid and the operational experience will be presented.

Network, Security II / 11

Web application security

Sebastian Lopienski¹

¹ *CERN*

Corresponding Author(s): sebastian.lopienski@cern.ch

CERN hosts a large number of Web sites (CERN-related, but also private), both on central Web Services, as well as on machines managed by particular Web site owners. Some of these Web sites are actually interactive Web applications developed with languages like PHP, ASP, Java, Perl, Python etc. - and unavoidably a fraction of them have bugs making them vulnerable to attacks such as Cross Site Scripting (XSS), Code/SQL Injection, Cross Site Request Forgery (CSRF), and so on. To address this issue, several Web application vulnerability assessment tools have been evaluated at CERN, and chosen ones are used to find vulnerabilities before the attackers do. This talk will discuss the choice of tools, the findings, and suggestions how Web application security can be improved in large organizations.

Desktop Management I / 20

Windows 7 horizons

Author(s): Michal Kwiatek¹

Co-author(s): Juraj Sucik¹ ; Sebastien Dellabella¹

¹ *CERN*

Corresponding Author(s): michal.kwiatek@cern.ch

Windows 7, the next version of Windows OS, is scheduled to be available worldwide on 22 October 2009. CERN IT-IS group has been working with it ever since its beta release in January 2009. The purpose of this talk is to discuss this experience and to share the plans for deployment of Windows 7 at CERN.

Storage I / 26

lustre@gsi: A Petabyte file system for the analysis farm - status and outlook

Walter Schoen¹

¹ *GSi*

Corresponding Author(s): w.schoen@gsi.de

lustre@gsi: A Petabyte file system for the analysis farm - status and outlook