

HEPiX Fall 2011 Workshop

Monday 24 October 2011 - Friday 28 October 2011

Hosted by TRIUMF, SFU and the University of Victoria at the Harbour
Center - Downtown Vancouver

Book of Abstracts

Contents

| | |
|---|---|
| ”A year in the life of Eucalyptus” | 1 |
| 20 years of AFS service | 1 |
| A highly distributed, petascale migration from dCache to HDFS | 1 |
| AMD’s New 16-core Opteron Processors | 2 |
| ASGC site report | 2 |
| ATLAS Great Lakes Tier-2 Site Report | 2 |
| An overview of computing hardware changes from 1991 to 2011 | 3 |
| BNL Site Report | 3 |
| CASTOR and EOS status and plans | 4 |
| CERN Computing Facilities Evolution Update | 4 |
| CERN site report | 4 |
| CPU Benchmarking at GridKa (Update 10/2011) | 4 |
| CVMFS Production Status Update | 5 |
| CloudMan and VMIC projects overview | 5 |
| Computer Security update | 5 |
| Configuration Management at GSI | 6 |
| DESY site report | 6 |
| Deska: Maintaining your computing center | 6 |
| Distributing images with image lists | 6 |
| EMI, the second year. | 7 |
| Experience with IPv6 deployment in FZU ASCR in Prague | 7 |
| Fermilab Site Report - Fall 2011 HEPiX | 8 |
| GRIF, LAL (Orsay) and IRFU (Saclay) site report | 8 |

| | |
|--|----|
| HEPiX From the Beginning | 8 |
| Hardware failures at CERN | 8 |
| Hepi-X-Perience | 9 |
| INFN Status Report | 9 |
| IPv6 - If Networking is Ready, Is Cyber Security Ready? | 9 |
| IPv6 deployment at CERN | 10 |
| Jefferson Lab Site Report | 10 |
| NDGF Site Report | 10 |
| Network connectivity for WLCG: the LHCONE | 11 |
| Networking Retrospective | 11 |
| News from HEPiX Storage Working Group | 11 |
| Open discussion on HEP-SPEC06 | 12 |
| OpenShift | 12 |
| OpenStack: The OpenSource Cloud's Application in High Energy Physics | 12 |
| PDSF Site Report | 13 |
| Prague Tier2 site report | 13 |
| RAL Site Report | 13 |
| Report from the HEPiX IPv6 Working Group | 14 |
| SLAC National Accelerator Laboratory Site Report | 14 |
| Scientific Linux Status | 14 |
| Secure file storage and transfer | 14 |
| Site Report GSI | 15 |
| Storage Status and Experiences at TRIUMF | 15 |
| TACC 10 PFLOP System | 15 |
| TRIUMF Site Report | 15 |
| TSM Monitoring at CERN | 16 |
| Update on activities at IHEP | 16 |
| Usage of OCS Inventory for Hardware and Software Inventory at CERN | 16 |
| Virtualisation Working Group Status | 17 |
| Virtualisation and Clouds at RAL | 17 |

| | |
|--|----|
| Wrap-Up | 17 |
| lxcloud infrastructure - status and lessons learned | 17 |
| perfSONAR or: How I Learned to Stop Worrying and Love Network Performance Verification | 18 |

Grid, Cloud & Virtualisation / 38**”A year in the life of Eucalyptus”****Author(s):** Iwona Sakrejda^{None}**Co-author(s):** Lavanya Ramakrishnan¹ ; Shane Canon²¹ *LBNL*² *NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER***Corresponding Author(s):** isakrejda@lbl.gov

Funded by the American Recovery and Reinvestment Act (Recovery Act) through the U.S. Department of Energy (DOE), the Magellan project was charged with a task of evaluating if cloud computing could meet specialized needs of scientists.

Split between two DOE centers: the National Energy Research Scientific Computing Center (NERSC) in California and the Argonne Leadership Computing Facility (ALCF) in Illinois Magellan built based on midrange hardware a testbed spanning both sites.

One of many services offered within this project was Eucalyptus, an open-source implementation of Amazon's popular EC2 cloud platform. Eucalyptus' interfaces are designed to replicate the APIs used on EC2. This includes implementing many of the capabilities of EC2 including Elastic Block, S3, Elastic IPs, etc. We ran Eucalyptus services on the NERSC portion of Magellan for the past year and we are going to share high(and low)lights from both the admin and the user perspective.

HEPiX Past and Future / 24**20 years of AFS service****Rainer Toebbicke**¹¹ *CERN*

Almost 20 years ago, the AFS service was born at CERN alongside a paradigm shift away from mainframe computing towards clusters. The scalable and manageable networked file system offered easy, ubiquitous access to files and greatly contributed to making this shift a success. Take a look back, with a smile rather than raised eyebrows, at how pre-Linux, pre-iPad, MegaByte and Megahertz technology faced the headwind of technological evolution and adapted over decades: AFS did a good job and continues to defeat, not always best in class but usually flexible.

Storage & File Systems / 8**A highly distributed, petascale migration from dCache to HDFS****Ajit Mohapatra**¹ ; **Dan Bradley**¹ ; **Sridhara Dasu**¹ ; **William Maier**¹¹ *University of Wisconsin (US)***Corresponding Author(s):** william.maier@cern.ch

The University of Wisconsin CMS Tier-2 center serves nearly a petabyte of storage and tens of thousands of hours of computation each day to the global CMS community. After seven years, the storage cluster had grown to 250 commodity servers running both the dCache distributed filesystem and the Condor batch scheduler. This multipurpose, commodity approach had quickly and efficiently scaled to meet growing analysis and production demands. By 2010, when alternatives to dCache became available in the CMS community, the center was ready to test alternatives that might be a better fit for its hybrid model. HDFS had become widely accepted in the web world and was designed to run in a similarly mixed storage and execution environment. In early evaluations, it performed as well as dCache while also reducing the operational burden. So, in the spring of 2011, the center successfully migrated all of its production data to HDFS with only a few hours downtime.

This migration was one of the largest to date within the CMS community. A unique and highly distributed mechanism was developed to complete the migration while maximizing availability of data to the thousands of jobs that run at Wisconsin each day. This talk presents the migration technique and evaluates its strengths, weaknesses and wider applicability as peers within the CMS community embark on their own migrations.

Computing / 52

AMD's New 16-core Opteron Processors

David (John) Cownie¹

¹ AMD

An overview of the architecture and power efficiency features of the latest 16-core processors from AMD including benchmark results for the HEP-SPEC suite -- showing the performance improvements over the current 12-core and older 6-, quad-, and dual-core processors. AMD's newest Opteron processors feature the "Bulldozer" x86 core pair compute module which is especially well-suited for modern C++ and object-based language workloads.

Site Reports / 55

ASGC site report

Hung-Te Lee¹

¹ Academia Sinica (TW)

Corresponding Author(s): felix@twgrid.org

Site report of ASGC.

Site Reports / 35

ATLAS Great Lakes Tier-2 Site Report

Author(s): Shawn McKee¹

Co-author(s): Ben Meekhof² ; Philippe Alain Luc Laurens³ ; Raymond Brock⁴ ; Robert Ball⁵ ; Tom Rockwell⁴

¹ University of Michigan ATLAS Group

² *University of Michigan*

³ *Michigan State University (US)*

⁴ *Michigan State University*

⁵ *High Energy Physics*

Corresponding Author(s): smckee@umich.edu

We will report on the ATLAS Great Lakes Tier-2 (AGLT2), one of five US ATLAS Tier-2 sites, providing a brief overview of our experiences planning, deploying, testing and maintaining our infrastructure to support the ATLAS distributed computing model. AGLT2 is one of the larger WLCG Tier-2s worldwide with 2.2 PB of dCache storage and 4500 job-slots, so we face a number of challenges in monitoring, managing and maintaining our site. Many of those challenges are related to storage, data-management and I/O capabilities.

As part of this report we will focus on our recent work in updating, configuring and monitoring our storage systems. In addition to describing new hardware like SSDs and multi-10GE storage nodes we will report on using such tools as pCache and LSM (Local Site Mover) and a new “site-aware” dCache configuration which have helped to remove some bottlenecks in our infrastructure.

Because AGLT2 utilizes a central syslog host, we are able to track the behavior of all our worker nodes in staging files in and out via LSM logging. We have constructed a system based upon a custom-built MySQL database which tracks our local resources and merges in information from the central syslog host and the dCache billing DB to allow us to better understand and optimize our site’s storage system behaviors. The last part of our report will show some results from using this new system.

Summary:

A site report from AGLT2 summarizing the current status and focusing on recent storage related efforts to find and remove bottlenecks in the infrastructure.

HEPiX Past and Future / 25

An overview of computing hardware changes from 1991 to 2011

Corrie Kost¹

¹ *TRIUMF*

An overview of computing hardware changes from 1991 to 2011 is given from a TRIUMF perspective. Aspects discussed are Moore’s law from speed, power consumption, and cost perspectives as well as how networks and commoditization, have influenced hardware. Speculation into the near and distant future nature of computing hardware is provided.

Site Reports / 1

BNL Site Report

Tony Wong¹

¹ *Brookhaven National Laboratory*

Corresponding Author(s): tony@bnl.gov

This presentation will cover recent developments and operational news of the RHIC-ATLAS computing facility at BNL.

Storage & File Systems / 42**CASTOR and EOS status and plans**

Giuseppe Lo Presti¹

¹ *CERN*

Corresponding Author(s): giuseppe.lopresti@cern.ch

[Still to be confirmed]

The Data and Storage Services (DSS) group at CERN develops and operates two storage solutions for the CERN Physics data, targeting both Tier0 central data recording and preservation, and user-space physics analysis. In this talk we present the current status of the two systems, CASTOR and EOS, and the foreseen evolution in the medium term.

IT Infrastructure / 18**CERN Computing Facilities Evolution Update**

Author(s): Wayne Salter¹

Co-author(s): Vincent Dore¹

¹ *CERN*

Corresponding Author(s): wayne.salter@cern.ch

There are a number of projects currently underway to improve and extend the CERN computing facilities which have been reported at previous HEPiX meetings. An update will be given on the current status of these projects and particular emphasis will be placed on efficiency improvements that have been made in the CERN Computer Centre and the resulting energy, and hence cost, savings.

Site Reports / 39**CERN site report**

Helge Meinhard¹

¹ *CERN*

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

News from CERN since last meeting

Computing / 34**CPU Benchmarking at GridKa (Update 10/2011)**

Manfred Alef¹

¹ *Karlsruhe Institute of Technology (KIT)*

Corresponding Author(s): manfred.alef@kit.edu

Report on current benchmark issues, e.g.:
- New release of SPEC CPU2006 available
- New processor architectures

Storage & File Systems / 50

CVMFS Production Status Update

Ian Collier¹

¹ *UK Tier1 Centre*

Corresponding Author(s): ian.peter.collier@cern.ch

The CernVM-FS has matured very quickly into a production quality tool for distributing VO software to grid sites. CVMFS is now in production use at a number of sites. This talk will recap the technology behind CVMFS and discuss the production status of the infrastructure.

Grid, Cloud & Virtualisation / 5

CloudMan and VMIC projects overview

Belmiro Daniel Rodrigues Moreira¹

¹ *CERN*

Corresponding Author(s): belmiro.moreira@cern.ch

CERN is developing a different set of tools to improve and agile the Cloud Computing infrastructure. Currently there are two active important projects: CloudMan is a project developed in collaboration with BARC institute and the VMIC project is developed in collaboration with ASGC.

The CloudMan project consists in the development of an Enterprise Graphical Management tool for IT resources and the VMIC project is a software tool that is used to manage in a trustful way the virtual images provided by different sites.

This presentation will give an overview over this two projects and the current status, and will explain how they fit into the new Cloud Computing model at CERN.

Network & Security / 7

Computer Security update

Romain Wartel¹

¹ *CERN*

Corresponding Author(s): 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 presents interesting

recent attacks. It also covers security risks management in general, as well as the security aspects of the current hot topics in computing, for example identity federation and virtualisation.

IT Infrastructure / 54

Configuration Management at GSI

Author(s): Christopher Huhn¹

Co-author(s): Bastian Neuburger²

¹ *GSI Darmstadt*

² *GSI*

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

GSI is successfully utilizing Cfengine for configuration management since almost a decade. Even though Cfengine is powerful as well as reliable we started to test the configuration management system Chef as a successor or complement to Cfengine to implement features we are lacking up to now.

Site Reports / 15

DESY site report

Dirk Jahnke-Zumbusch¹

¹ *DESY*

current information about DESY IT - both for the Hamburg and Zeuthen sites

IT Infrastructure / 43

Deska: Maintaining your computing center

Jan Kundrat¹

¹ *Unknown-Unknown-Unknown*

Corresponding Author(s): kundratj@fzu.cz

The proposed talk discusses the Deska project [1], our attempt at delivering an inventory database whose goal is to provide a central source of machine-readable information about one's computing center. We mention the motivation behind the project, describe the design choices we have made and talk about how the Deska system could help reduce maintenance effort on other sites.

Grid, Cloud & Virtualisation / 44

Distributing images with image lists

Owen Millington Synge^{None}

Corresponding Author(s): owen.synge@desy.de

The use of signed image lists, and updates the use of signed image lists.

Summary:

The use of signed image lists, and updates the use of signed image lists.

Storage & File Systems / 29

EMI, the second year.

Author(s): Patrick Fuhrmann¹

Co-author(s): Fabrizio Furano²; Jon Kerr Nilsen³; Michail Salichos²; Oliver Keeble²; Paul Millar¹; Ricardo Brito Da Rocha²; Riccardo Zappi⁴; Tigran Mkrtchyan Mkrtchyan⁵

¹ DESY

² CERN

³ University of Oslo (NO)

⁴ CNAF-INFN

⁵ Deutsches Elektronen-Synchrotron DESY

Corresponding Author(s): patrick.fuhrmann@desy.de

The European Middleware Initiative is now rapidly approaching its projects half-value period. Nearly all objectives of the first year of EMI-Data have been achieved and the feedback from the first EMI review has been very positive. Internet standards, like WebDAV and NFS4.1/pNFS have been integrated into the EMI set of storage elements, the already existing accounting record has been extended to cover storage and the synchronization of catalogues and storage elements has been designed and implemented within gLite. Furthermore, the close collaboration between EMI and EGI resulted in a very positive feedback from EGI and the subsequent creation of a set of new objectives focusing on the EGI acceptance of the EMI software distribution. This presentation will briefly describe the achieved goals of the first year but will primarily focus on the work EMI Data is facing for the rest of its projects lifetime, including the design of the common EMI data client libraries, the new gLite File Transfer Service (FTS3) and our efforts in consolidating http(s) and WebDAV.

Network & Security / 33

Experience with IPv6 deployment in FZU ASCR in Prague

Author(s): Lukas Fiala¹; Marek Elias¹; Tomas Kouba¹

Co-author(s): Jan Kundrat¹; Jan Svec¹; Jiri Chudoba¹; Jiri Horky¹

¹ Institute of Physics AS CR, v. v. i. (FZU)

We are facing exhaustion of IPv4 addresses and transition to IPv6 is becoming more and more urgent. In this contribution we describe our current problems with IPv4 and our special motivation for transition to IPv6. We present our current IPv6 setup and installation of core network services like DNS and DHCPv6. We describe our PXE installation testbed and results of our experiments with installing operating system with PXE through IPv6. We have tested native PXE implementations of our current hardware as well as open source network bootloader gPXE. Actualisation of CRL of certification authorities is a

service used by all components of gLite. We have prepared a web service for testing availability of revocation lists of certification authorities from lcg-CA bundle.

Site Reports / 0

Fermilab Site Report - Fall 2011 HEPiX

Keith Chadwick¹

¹ *Fermilab*

Corresponding Author(s): chadwick@fnal.gov

Fermilab Site Report - Fall 2011 HEPiX

Summary:

Fermilab Site Report - Fall 2011 HEPiX

Site Reports / 10

GRIF, LAL (Orsay) and IRFU (Saclay) site report

Pierrick Micout¹

¹ *CEA*

Corresponding Author(s): micout@cea.fr, jouvin@lal.in2p3.fr

Site report of GRIF/LAL and GRIF/Irfu

HEPiX Past and Future / 3

HEPiX From the Beginning

Alan Silverman¹

¹ *CERN*

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

HEPiX is 20 years old this year and this talk will try to summarise some of the significant events of those 20 years. The speaker will also try to answer the question - is HEPiX worth the money?

IT Infrastructure / 23

Hardware failures at CERN

Author(s): Olof Barring¹

Co-author(s): Wayne Salter ¹

¹ *CERN*

Corresponding Author(s): olof.barring@cern.ch

A detailed study of approximately 4000 vendor interventions for hardware failures experienced in CERN IT computing facility in 2010-2011 will be presented. The rates of parts replacements are compared for different components and as expected disk failures are dominating with approximately 1% quarterly replacement rate. When plotting the variation with age a higher rate is seen in the first year after deployment whereas there is no significant sign of wear-out at the end of the 3 years warranty.

HEPiX Past and Future / 14

Hepi-X-Perience

Thomas Finnern¹

¹ *DESY*

Corresponding Author(s): thomas.finnern@desy.de

This is a personal retrospective view on 18 years of membership in the HEPiX community. Starting in 1993, it was associated with my career as a computer system engineer, the progression of high performance computing, and shifts of paradigm. The talk gives some spot lights on my own and community aspects during this time by recalling personal projects and events.

Site Reports / 17

INFN Status Report

Author(s): Michele Michelotto¹

Co-author(s): Roberto Gomez ²

¹ *INFN Padua & CMS*

² *INFN*

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

Istituto Nazionale di Fisica Nucleare - Site Report

Network & Security / 2

IPv6 - If Networking is Ready, Is Cyber Security Ready?

Bob Cowles¹

¹ *SLAC*

Corresponding Author(s): bob.cowles@gmail.com

The coming of IPv6 represents the introduction of a new protocol stack, rich in features and, if the past is any guide, an interesting set of challenges for cyber security. The talk will cover both current recommendations for IPv6 configuration and open issues requiring further discussion and investigation.

Summary:

The coming of IPv6 represents the introduction of a new protocol stack, rich in features and, if the past is any guide, an interesting set of challenges for cyber security. IPv6 is already being used to bypass web proxies, to get around torrent rate limiting by ISPs, and it integrated into the Metasploit Framework. Most sites will be running IPv4 and IPv6 for the foreseeable future, increasing the complexity of network traffic and creating more opportunities for attackers. Current cyber security practices (e. g. ping all addresses to find active hosts) and tools will require rethinking and reengineering. The talk will cover both current recommendations for IPv6 configuration and open issues requiring further discussion and investigation.

Network & Security / 13

IPv6 deployment at CERN

Edoardo Martelli¹

¹ *CERN*

Corresponding Author(s): edoardo.martelli@cern.ch

Description of the CERN IPV6 deployment project: service definition, features, implementation plan

Summary:

CERN has already started the implementation of all the tools that will allow to deploy the IPv6 network over the existing IPv4 network. This talk will present the characteristics of the IPv6 service, the implementation plan and the current status.

Site Reports / 27

Jefferson Lab Site Report

Sandy Philpott¹

¹ *JLAB*

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

An update of scientific computing activities at JLab since the Cornell meeting, including Lustre and GPU status.

Site Reports / 21

NDGF Site Report

Erik Mattias Wadenstein¹

¹ *Unknown*

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

Update of recent and current events at NDGF.

Network & Security / 12

Network connectivity for WLCG: the LHCONE

Edoardo Martelli¹

¹ *CERN*

Corresponding Author(s): edoardo.martelli@cern.ch

LHCONE (LHC Open Network Environment) is the network which will give dedicated bandwidth for LHC data transfer to Tier2s and Tier3s

Summary:

The LHCOPN has been successfully distributing LHC data from CERN to the remote Tier1 centres across the globe. However, the increasing data volume and computational power requires an extended infrastructure to optimise the way data is made available to physicists at Tier2 and Tier3 centres. The LHCONE is a new network infrastructure built on the concept of Open Exchange Points and Interdomain Circuit Provisioning which will provide the LHC community with high bandwidth on a world-wide scale.

HEPiX Past and Future / 4

Networking Retrospective

Les Cottrell¹

¹ *SLAC*

At the inauguration of HEPiX in 1991, mainframes (and HEPVM) were on their way out with their bus & tag cables, channels with 3270 emulators and channel attached Ethernets. DEC/VMS and DECnet were still a major player in the scientific world. Mainframes and to a lesser extent VMS hosts were being replaced by Unix hosts with native TCP stacks running on thin and thicknet shared media, the phone system was still a separate entity, wireless networking was very much a niche. The wide area network consisted of a multitude of networks and protocols such as DECnet, SNA, XNS, Color books, Bitnet/EARN and the emerging (soon to die off) OSI. All these were soon to “pass like tears in rain”* and be displaced by TCP/IP and what we know as the Internet today. We will look back at the evolution of the local area network, home networking and the wide area network over the last 30-40 years in particular noting the state and changes since HEPiX was formed 20 years ago.

- Blade Runner

Storage & File Systems / 30

News from HEPiX Storage Working Group

Andrei Maslennikov¹

¹ *CASPUR*

Computing / 48

Open discussion on HEP-SPEC06

Michele Michelotto¹

¹ *INFN Padua & CMS*

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

The HEP-SPEC06 benchmark was designed by a working group born during the HEPPIX meeting in JLAB. The HS06 is now the standard for measuring computing power in HEP and also in other scientific areas that make use of Computing GRID.

The goal of this discussion is to understand how the HEPPIX community sees the future of HS06.

Grid, Cloud & Virtualisation / 20

OpenShift

Troy Dawson¹

¹ *Red Hat*

OpenShift is a collection of cloud services forming a solid redefining Platform-as-a-Service (PaaS) for developers who build apps on open source technologies.

Grid, Cloud & Virtualisation / 56

OpenStack: The OpenSource Cloud's Application in High Energy Physics

Author(s): Neil Johnston¹

Co-author(s): Christopher MacGown² ; Joshua McKenty³

¹ *Piston Cloud Computing*

² *Piston Cloud Computing, Inc.*

³ *Piston Cloud Computing, Inc. and OpenStack PPB*

Corresponding Author(s): neil@pistoncloud.com

OpenStack's mission is "To produce the ubiquitous Open Source cloud computing platform that will meet the needs of public and private cloud providers regardless of size, by being simple to implement and massively scalable." This talk will review the implications of this vision to meet the storage and compute needs of data intensive research projects, then examine OpenStack's potential as a largely common hardware-agnostic platform for the federation of resources across sites and organizations.

Site Reports / 28**PDSF Site Report****Author(s):** Iwona Sakrejda^{None}**Co-author(s):** Jay Srinivasan¹¹ *Lawrence Berkeley National Lab. (LBNL)-Unknown-Unknown***Corresponding Author(s):** isakrejda@lbl.gov

PDSF is a networked distributed computing cluster designed primarily to meet the detector simulation and data analysis requirements of Physics, Astrophysics and Nuclear Science collaborations. Located at NERSC and benefiting from excellent network and storage infrastructure cluster constantly changes to keep up with computing requirements of physics experiments like ATLAS, Alice and Daya Bay. Over the past year - as it is usual to us - we replaced a big fraction of hardware (both computing and storage), moved towards diskless installs with xCAT and were deploying and supporting in addition to GPFS, also XRootD and CVMFS. We also considered migrating to a new batch system and we are going to share our motivation behind the decision to continue with the Univa Grid Engine.

Site Reports / 49**Prague Tier2 site report****Author(s):** Jan Svec¹**Co-author(s):** Jan Kundrat²; Jiri Chudoba¹; Jiri Horky³; Marek Elias⁴; Tomas Kouba⁵¹ *Acad. of Sciences of the Czech Rep. (CZ)*² *Unknown-Unknown-Unknown*³ *Institute of Physics-Acad. of Sciences of the Czech Rep. (ASCR)*⁴ *FZU ASCR*⁵ *Institute of Physics - Acad. of Sciences of the Czech Rep. (ASCR)***Corresponding Author(s):** jan.svec@cern.ch

Main computing and storage facilities for LHC computing in the Czech Republic are situated at Prague Tier-2 site. We participate in grid activities since the beginning of European Data Grid. The recent years were significant in the growth of our computing and storage capacities. In this talk, we will present the current state of our site, its history and plans for the near future.

Site Reports / 11**RAL Site Report**Martin Bly¹¹ *STFC-RAL***Corresponding Author(s):** martin.bly@stfc.ac.uk

The latest from the RAL Tier1.

Network & Security / 9**Report from the HEPiX IPv6 Working Group**

David Kelsey¹

¹ *STFC - Science & Technology Facilities Council (GB)*

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

This new working group was formed earlier in 2011. There have been several meetings, sub-topics have been planned and work is now well underway. This talk will present the current status and plans for the future.

Site Reports / 31**SLAC National Accelerator Laboratory Site Report**

Neal Adams¹

¹ *SLAC*

Site report for SLAC.

IT Infrastructure / 47**Scientific Linux Status**

connie sieh¹

¹ *Fermilab*

Corresponding Author(s): csieh@fnal.gov

Current status of Scientific Linux
Futures of Scientific Linux

IT Infrastructure / 16**Secure file storage and transfer**

Author(s): Ivan Fedorko¹ ; Jan Dudzic²

Co-author(s): Veronique Lefebure¹

¹ *CERN*

² *Warsaw University of Technology*

Corresponding Author(s): veronique.lefebure@cern.ch

SINDES, Secure INFORMATION DELIVERY System, is a tool aimed at ensuring enough level of privacy in storing and delivering confidential files. Initially written at CERN in 2005, SINDES is now being

rewritten in order to have its user-interface, flexibility and maintainability improved: access control granularity, logging, file modifications, history, machine upload, unattended installations and support for different OS, are examples of points that are improved in the new version, which is based on Kerberos authentication.

Site Reports / 53

Site Report GSI

Walter Schon^{None}

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

Site report GSI

Storage & File Systems / 22

Storage Status and Experiences at TRIUMF

Author(s): Simon Liu¹

Co-author(s): Reda Tafirout ¹

¹ *TRIUMF (CA)*

Corresponding Author(s): simonliu@triumf.ca

The ATLAS Tier1 data centre at TRIUMF provides a highly efficient and scalable storage components to support LHC data analysis and production. This contribution will describe and review the storage infrastructure and configuration currently deployed at the Tier-1 data centre at TRIUMF for both disk and tape, as well sharing of past experiences. A brief outlook on test beds and future expansion will also be presented.

Computing / 46

TACC 10 PFLOP System

Roger Goff¹

¹ *DELL*

To be defined

Site Reports / 32

TRIUMF Site Report

Denice Deatrigh^{None}

Corresponding Author(s): deatrigh@triumf.ca

Updates on the status of the Tier-1 and other TRIUMF computing news.

IT Infrastructure / 41

TSM Monitoring at CERN

Author(s): Daniele Francesco Kruse¹

Co-author(s): Giuseppe Lo Presti¹

¹ CERN

Corresponding Author(s): giuseppe.lopresti@cern.ch

The TSM server network at CERN - with its 17 TSM servers in production, 30 drives, ~1300 client nodes and ~4 PB of data - often needs an overwhelming amount of effort to be properly managed by the few TSM administrators.

Hence, the need for a central monitoring system able to cope with the increasing number of servers, client nodes and volumes. We will discuss our approach to this issue, focusing on TSMMS, a TSM Monitoring System developed in-house, able to give an effective view of the needs and status of the network and of the individual servers, as well as statistics and usage reports.

Avoiding repetitive error-prone manual checks, TSM admins are able to manage the whole TSM system just by looking at the periodic reports, and taking appropriate action.

TSMMS scales seamlessly with the enlargement of the network, thus saving the cost of additional administrative personnel.

Site Reports / 37

Update on activities at IHEP

Author(s): Jingyan Shi¹

Co-author(s): Fazhi Qi¹

¹ IHEP

Corresponding Author(s): shijy@ihep.ac.cn

Give an introduction about the update activities on IHEP computing site during the past half year.

IT Infrastructure / 40

Usage of OCS Inventory for Hardware and Software Inventory at CERN

Matthias Schroeder¹

¹ CERN

Corresponding Author(s): matthias.schroeder@cern.ch

CERN has started to use OCS Inventory for the HW and SW inventory of SLC nodes on site, and plans to do the same for the MacOS nodes.

I will report about the motivation for this, the setup used and the experience gained.

Grid, Cloud & Virtualisation / 26

Virtualisation Working Group Status

Tony Cass¹

¹ *CERN*

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

An update on the progress of the working group since the Spring HEPiX meeting at GSI.

Grid, Cloud & Virtualisation / 51

Virtualisation and Clouds at RAL

Ian Collier¹

¹ *UK Tier1 Centre*

Corresponding Author(s): ian.peter.collier@cern.ch

Status of work on virtualisation and cloud computing at the RAL Tier 1.

Wrap-Up / 57

Wrap-Up

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

Grid, Cloud & Virtualisation / 6

Ixcloud infrastructure - status and lessons learned

Belmiro Daniel Rodrigues Moreira¹

¹ *CERN*

Corresponding Author(s): belmiro.moreira@cern.ch

In December 2010 CERN moved parts of the batch resources into a cloud like infrastructure, and is running some of the batch resources in a fully virtualized infrastructure since then. This presentation will give an overview over the experiences learned from this exercise, the performance and results, impressions on operational overhead, and problems seen since the deployment of the infrastructure. On the development path, first experiences with SLC6 are shown.

In addition, CERN has opened some resources to special users from ATLAS and LHCb via OpenNebula's EC2 interface, which they have been using to test CERNVM generated images which directly

connect to the experiment frameworks to get their payload. First results from these tests are shown as well.

Network & Security / 36

perfSONAR or: How I Learned to Stop Worrying and Love Network Performance Verification

Jason Zurawski¹

¹ *Internet2*

Corresponding Author(s): zurawski@internet2.edu

Scientific innovation produced by Virtual Organizations (VOs) such as the LHC, demands high capacity and highly available network technologies to link remote data creation, storage, and processing facilities. Research and Education (R&E) networks are a vital cog in this supply chain, and offer advanced capabilities to this distributed scientific project. Network operations staff spend countless hours monitoring and assuring internal performance and traffic management needs, all to benefit local user communities. Often the “big picture” of end-to-end performance is forgotten, or cast aside, due to the relative complexity of multi-domain operational scenarios and the lack of human and technological resources.

Software designed to monitor and share network information between domains, developed by the perfSONAR-PS project, is available to help with end-to-end performance concerns. This framework, in use within the USATLAS project since 2007, and emerging on other collaborations including the Italian and Canadian ATLAS clouds, has been beneficial in identifying complex network faults while imposing minimal operational overhead on local administrators.

Summary:

This talk will introduce perfSONAR-PS as a performance tool, discuss use cases, and highlight deployment success stories within the LHC communities.