

HEPiX Fall 2012 Workshop

Monday, October 15, 2012 - Friday, October 19, 2012

Institute of High Energy Physics



Book of Abstracts

Contents

Fermilab Site Report	1
The High-Availability FermiCloud Infrastructure-as-a-Service Facility	1
GSI site report	1
Scientific Linux current status update	2
ZNeTS: log your network traffic !	2
LCG-BEIJING Site Status	3
Network Traffic Analysis using HADOOP Architecture	3
The Lustre file system at IHEP	4
OpenStack Chances and Practice in IHEP	4
IPv6 deployment in IHEP	4
Bringing cloud storage to your desk with Mucura	5
CERN site report	5
Virtualisation working group progress report	6
DYNES: Building a distributed networking instrument	6
Cyber security update	6
Mobile web development, and CERN mobile web site	7
Australia Site Report	7
Networking Tools for Sysadmins	7
Setting up the CSP mode in a Gridengine production cluster	8
CERN Batch System, Monitoring and Accounting	8
CERN Agile Infrastructure , Road to Production	8
Lync - Phone, voice mailbox, instant messaging... Get access to all of them from any place in the world.	9
Integration Lemon/LAS monitoring with the CERN Agile Infrastructure	9

scientific data cloud infrastructure and services in Chinese Academy of Sciences	9
CERN Cloud Storage Evaluation	10
NDGF site report	10
Data Center Network changes and extension to Wigner	10
IPv6 deployment status at CERN	11
Global Accounting in the Grid and Cloud	11
Lustre at GSI	11
ITIL at CC-IN2P3	12
The HEPiX IPv6 Working Group	12
Federated Identity Management for HEP	12
RAL Tier1 Disk only Storage Project status and plans	13
Quattor update - Integrating Aquilon into a grid site	13
STFC Scientific Computing Department Cloud computing	13
LRMS Migration at GridKa	14
RAL Site Report	14
JASMINE/CEMS and EMERALD	14
The ATLAS Great Lakes Tier-2 (AGLT2) Site Report	14
Workshop wrap-up	15
News from HEPiX Storage Working Group	15
Service Provisioning and Security Guarantee in CSTNET	15
Testing SLURM batch system for a grid farm: functionalities, scalability, performance and how it works in a GRID environment	16
EGI Federated Cloud Infrastructure	16
A decade of Condor experience at Fermilab	17
Welcome address	17
Workshop logistics	17
Oracle Grid Engine at CC-IN2P3 - report after One year	17
Slurm Experiences for WLCG in the Nordics	18
LAL and GRIF site report	18
A new shared datacenter in Orsay	18

Alternatives to Posix. Lessons with S3 Compatible Storage Systems 19

DESY Site Report 19

Selecting a Business-Process-Management-System in conjunction with an Identity-and-
Access-Management-System 19

CERN and Helix Nebula, the Science Cloud 19

Site Report of ASGC 20

Scientific Linux Infastructure Improvements 20

FZU IPv6 testbed updates 20

Report from batch system BOF 21

Site reports / 0**Fermilab Site Report****Author:** Steven Timm¹**Co-author:** Keith Chadwick¹¹ *Fermilab***Corresponding Author:** timmsteve@yahoo.com

We present recent developments in the Scientific Computing Facilities at Fermilab. We will discuss continued improvements in site networking and wide area networking. We will show significant developments in physical facilities. We will present an overview of major computing and organizational activities in support of the scientific program.

Grid, Cloud and Virtualisation / 1**The High-Availability FermiCloud Infrastructure-as-a-Service Facility****Author:** Steven Timm¹¹ *Fermilab***Corresponding Author:** timmsteve@yahoo.com

FermiCloud is an Infrastructure-as-a-Service private cloud built for the support of scientific computing at Fermilab. Within the past year we have deployed a facility capable of providing 24x7 service. We will present significant advances in monitoring and visualization, accounting, security, authorization, and user interface. We will also present our current plans for multi-cloud interoperability.

Summary:

FermiCloud has recently updated to a new version of OpenNebula and deployed the X.509 authentication features which were our contribution to OpenNebula. We have split our hardware between two different buildings such that we can survive the failure of either one.

We have deployed new monitoring and visualization of system usage and availability. We have a proof-of-principle of XACML-based callouts for X.509-based authorization as well.

Site reports / 2**GSI site report****Author:** Walter Schoen¹¹ *GSI***Corresponding Author:** w.schoen@gsi.de

GSI site report

IT Infrastructure / 3**Scientific Linux current status update**

Authors: Pat Riehecky¹; connie sieh¹

¹ *Fermilab*

Corresponding Author: riehecky@fnal.gov

This presentation will provide an update on the current status of Scientific Linux, descriptions for some possible future goals, and allow a chance for users to provide feedback on its direction.

Summary:

This presentation will provide an update on the current status of Scientific Linux, descriptions for some possible future goals, and allow a chance for users to provide feedback on its direction.

Security and Networking / 4**ZNeTS: log your network traffic !**

Author: Thierry DESCOMBES¹

¹ *CNRS IN2P3*

Corresponding Author: thierry.descombes@lpsc.in2p3.fr

ZNeTS is an acronym for “The Network Traffic Supervisor”. It is a tool for monitoring and recording machines traffic during months.

ZNeTS is a network tool for network introspection and a response to the legal need, in France, to store one year traffic traces.

ZNeTS is very easy to deploy whatever the architecture of your network.

ZNeTS identifies compromised local machines (by virus, trojans, abusive or illegal usage, DNS or Mac spoofing, etc...).

ZNeTS graphical interface is intuitive and ergonomic. Integrated metrology features offer two levels of details. Alerts are simple and relevant.

Over the last 6 month, the tool has been successfully deployed as an appliance into all the IN2P3 laboratories (the french national research institute in physics) and we gave very positive feedbacks from the System administrators.

Summary:

ZNeTS is a powerful and easy tool for monitoring LAN, parts of LAN and machines. It has been developed for the CNRS (the french national research institute).

The purpose is:

- The acquisition and conservation of inbound and outbound network flows (during many months or years... even on very high speed networks)
- Research and filter data, with an integrated search engine
- The detection of anomalies causing the generation of alert and optional email sending
- Metrology, with the calculation and visualization of hourly and daily statistics of the overall traffic and detailed traffic (for each subnet, and machine on the LAN).

ZNeTS is very easy to deploy. It was developed in C++ and includes a web server that implements HTTP/1.1 standard (RFC2616). Authentication by login / password or certificate-X509 is also possible.

The web interface based on the Dojo framework is particularly ergonomic and allows the interpretation and visualization of data. ZNeTS is suitable for all network architectures. It is able to acquire not only the data from NetFlow probe (netflow mode), but also directly from a physical interface (sniffer mode). It is able to decode most versions of NetFlow and IPFIX. It supports IPv4 and IPv6. Packages have been built for most Linux distributions.

ZNeTS is not just a collector. It uses optimized flows that are re-aggregated during a adjustable period of time (from 1 minute to 1 hour). ZNeTS may even send its flows to another collector (and then behaves as a NetflowV9 probe)

The graphical interface is intuitive and ergonomic. Metrology offers two levels of detail. Alerts are simple and relevant (based on counting algorithms). Access to the network flows is easy, the selection forms are pre-filled automatically after each interaction. 2 clicks on charts or on an alert report are enough to see the corresponding flows.

Site reports / 5

LCG-BEIJING Site Status

Author: Jingyan Shi¹

¹ *IHEP*

Corresponding Author: shijy@ihep.ac.cn

The presentation will show the current status of BEIJING LCG Site and our next plan.

Security and Networking / 6

Network Traffic Analysis using HADOOP Architecture

Author: Shan Zeng¹

¹ *IHEP*

Corresponding Author: zengshan@ihep.ac.cn

This report introduced a network traffic analysis tool using HADOOP architecture. By collecting the traffic information of the egress router in a campus or an institute, the network traffic analysis tool stored the traffic information which includes start time, end time, source IP, destination IP, Byte, Packet, Flow and etc. to HDFS which is a distributed file system as well as the RRD. In the frontend, the tool using rrdtool graph to draw graphs of the network flow trend chart, to get the details of the traffic information just click on the trend chart and there will be a detailed graph on the network flow information drawing by highstock which read data from HDFS, the user can also select a timeslot or a time window to get the netflow information which is calculated by the Map-reduce program running background. Meanwhile, by providing the IP specially related one HEP experiment, the tool can give the traffic information related to the specially HEP experiment, now it is used in collecting the network traffic information of DYB, YBJ, CMS and ATLAS, the network traffic information can be shown in realtime as well as the historic record, and once you put the mouse on the graph, the timeslot and the netflow traffic information will be shown on the graph.

Storage and Filesystems / 7**The Lustre file system at IHEP****Author:** Lu Wang¹¹ *IHEP*

Lustre has been selected as the main distributed file system solution in IHEP for more than four years. The Lustre File System at IHEP is currently at a scale of 2.2 PB capacity, 50 OSSs and 500+ OSTs, running Lustre 1.8.6. The file system which was built on top of commodity disk arrays, servers and 10 Gbit Ethernet, provides 24 GB/s bandwidth for five high energy physics experiments.

The presentation will mainly report the status of the Lustre file system at IHEP. It includes three parts: 1) an overview of the Lustre File system at IHEP, including the deployment history, current configuration and real performance captured during production usage; 2) the I/O pattern of High energy physics computing, including the file size, the read/write extent size and offset size and performance optimization according to this pattern; 3) management experience abstracted from 4 years' production run.

IT Infrastructure / 8**OpenStack Chances and Practice in IHEP****Author:** li haibo^{None}**Co-authors:** Yaodong CHENG¹; hu qingbao²¹ *Institute of High Energy Physics, Chinese Academy of Sciences*² *IHEP***Corresponding Author:** lihaibo@ihep.ac.cn

OpenStack is a global collaboration of developers and cloud computing technologists producing the open standard cloud computing platform for both public and private clouds. The project aims to deliver solutions for all types of clouds by being simple to implement, massively scalable, and feature rich. The technology consists of a series of interrelated projects delivering various components for a cloud infrastructure solution. This talk will present the status of the project from the user's and developer's perspectives and also show our practice on using OpenStack to build our cloud computing environment in data center. Then a series of experiments will be conducted and the performance comparison with the current solution will be also demonstrated. Some interesting research and development what we concern on OpenStack will be stated, including private cloud store for HEPiX users, operation monitoring, etc. Finally, the future developments of OpenStack such as auto deployment, network planning and high availability will also be discussed.

Security and Networking / 9**IPv6 deployment in IHEP****Author:** Qi Fazhi¹¹ *IHEP***Corresponding Author:** qfz@ihep.ac.cn

Description of the IPV6 deployment in IHEP

Summary:

IHEP has finished the implementation of the IPv6 network infrastructure in the campus network, and the network is running in the dual stack mode. This talk will present the current status of IHEP network and some related system or software deployed in IHEP, including the IPv6 user management system, access control system, etc., the IPv6 related plan, applications and projects in IHEP also will be introduced in this presentation.

Storage and Filesystems / 10**Bringing cloud storage to your desk with Mucura**

Author: Fabio Hernandez¹

¹ *IN2P3/CNRS Computing Center and IHEP Computing Center*

In this contribution we present our experience building a prototype of an open source software system for operating online file repositories of extensible capacity.

Built on the well-understood client-server architecture model, the system can be used by computing centers looking at solutions for providing online storage services for their individual users. The client-side component runs on the end-user's personal computer and provides both command-line and graphical user interfaces. It supports a deliberately limited set of operations on remote files, namely store, retrieve, organize and share remote files.

Mucura exposes the same HTTP-based standard API supported by Amazon S3 and extends it to support the certificate-based authentication mechanism used by production grid computing platforms such as WLCG. As a consequence, personal file repositories based on Mucura can be seamlessly accessed both from the user's personal computer and from grid jobs running on the user's behalf. This integration allows researchers to use their individual online storage space as a personal storage element conveniently managed from their personal computer.

At the core of the system there are components for managing file metadata and for secure storage of the files' contents, implemented on top of highly available, distributed, persistent and scalable key-value stores.

We will present a detailed architectural view of the system, the status of development and the perspectives for the months to come.

This work is inspired not only by the increasing number of commercial services available nowadays to individuals for their personal storage needs (backup, file sharing, synchronization, ...) such as Amazon S3, Dropbox, SugarSync, bitcasa, etc., but also by several efforts in the same area in the academic and research worlds (NASA, SDSC, etc.). We are persuaded that the level of flexibility offered to individuals by this kind of systems adds value to the day-to-day work of scientists.

Site reports / 11**CERN site report**

Author: Helge Meinhard¹

Co-author: Arne Wiebalck¹

¹ *CERN*

Corresponding Author: helge.meinhard@cern.ch

News from CERN since the previous meeting

Grid, Cloud and Virtualisation / 12

Virtualisation working group progress report

Author: Tony Cass¹

¹ *CERN*

Corresponding Author: tony.cass@cern.ch

An update on the work of the virtualisation working group.

IT Infrastructure / 13

DYNES: Building a distributed networking instrument

Author: Benjeman Jay Meekhof¹

¹ *University of Michigan (US)*

Corresponding Author: benjeman.jay.meekhof@cern.ch

This presentation will discuss the challenges of efficiently provisioning and deploying switch and host OS configurations enabling our collaboration to monitor, access, and repair the distributed instrument. Additionally we will cover some of the ongoing challenges post-deployment as regards configuration tracking, service verification, and monitoring the overall status of the DYNES instrument as well as enabling DYNES use for domain specific applications (like the LHC).

Summary:

The Dynamic Network System (DYNES) is a distributed networking instrument that creates point-to-point, variable-bandwidth circuits between DYNES sites. This requires hardware and software infrastructure deployment to all participating sites. Currently the deployment stands at approximately 40 sites in the US requiring a storage system, network switch, and IDC (inter-domain-controller) to enable network reservations between sites.

Security and Networking / 14

Cyber security update

Author: Sebastian Lopienski¹

¹ *CERN*

Corresponding Author: sebastian.lopienski@cern.ch

This talk gives an update on security trends, and 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-security world of 2012. New security tools developed at CERN will also be presented.

This talk is based on contributions and input from the CERN Computer Security Team.

Miscellaneous / 15

Mobile web development, and CERN mobile web site

Author: Sebastian Lopienski¹

¹ *CERN*

Corresponding Author: sebastian.lopienski@cern.ch

Mobile computing is clearly on the rise – but developing mobile applications, especially for multiple platforms, is a considerable effort. Fortunately, there is an alternative to native apps: web sites that are optimized for mobile devices and touch screens. In this presentation, I will discuss both solutions, and will present a hybrid approach. The presentation will also include a brief introduction to technologies such as jQuery, jQuery Mobile and PhoneGap.

Additionally, as an example of a mobile web application, CERN mobile web site (<http://m.cern.ch>) will be presented.

Site reports / 16

Australia Site Report

Author: Lucien Philip Boland¹

Co-author: Sean Christopher Crosby¹

¹ *University of Melbourne (AU)*

Corresponding Author: lucien.boland@cern.ch

Details of the upgrades and changes that have recently been made at the Australian Centre of Excellence for Particle Physics at the Terascale.

Security and Networking / 17

Networking Tools for Sysadmins

Author: Wolfgang Friebel¹

¹ *Deutsches Elektronen-Synchrotron (DE)*

Corresponding Author: wolfgang.friebel@desy.de

The talk wants to introduce netdisco, which is used at DESY for network inventory and has been enhanced slightly. A command line interface for the netdisco DB developed at DESY should be covered as well.

The concept of using netflow data to understand and analyze network traffic should be discussed, which is especially suitable for sites with high traffic. Programs that can deal with netflow data, such

as the recently released version 5 of ntop, nfdump to record netflows and nfsen to visualize the data will be presented, as well as some plugins for nfsen.

All the tools covered in the talk are IPv6 enabled.

Computing / 18

Setting up the CSP mode in a Gridengine production cluster

Author: Andreas Haupt¹

¹ *Deutsches Elektronen-Synchrotron (DE)*

Corresponding Author: andreas.haupt@desy.de

All the currently available Gridengine implementations don't provide any authenticated access with the default setup. This opens a big and easily exploitable security hole which might be considered severe especially in multi-community clusters.

This talk will describe in detail the attack vector available in such setups. It will furthermore give a step-by-step guide to activate the certificate-based authentication in Gridengine (the so called "CSP mode") based on the experience at DESY.

Computing / 19

CERN Batch System, Monitoring and Accounting

Author: Jerome Belleman¹

¹ *CERN*

Corresponding Author: jerome.belleman@cern.ch

The CERN batch service runs a 60k CPU core cluster using Platform LSF. We present some of the challenges of running a service at this scale, and describe the current planning of how we aim to evolve the current system to a more dynamic, larger scale service.

As part of this, we recently undertook a project of developing new monitoring tools and upgrading the batch accounting system; we present the current state of development in this area.

IT Infrastructure / 20

CERN Agile Infrastructure , Road to Production

Author: Steve Traylen¹

Co-author: Jan van Eldik ¹

¹ *CERN*

Corresponding Author: steve.traylen@cern.ch

The CERN Agile Infrastructure aims to redesign the work flow of machine and configuration management within CERN IT. As the AI project approaches

production the main software components , OpenStack , Puppet, Foreman have now been deployed with several iterations of scale and stability. We present the current status and next steps for the project.

IT Infrastructure / 21

Lync - Phone, voice mailbox, instant messaging... Get access to all of them from any place in the world.

Author: Pawel Grzywaczewski¹

¹ CERN

Corresponding Author: pawel.grzywaczewski@cern.ch

Use your device (computer or portable device) as your main tool for unified communication. Check presence of your colleagues, make phone calls, get call notifications, listen to your voice mailbox, answer mails and send instant messages. Do it from any place around the world with internet access.

The presentation will summarize our experience in integrating Microsoft Lync, Alcatel PBX and Microsoft Exchange. The goal was to provide a system that integrates VoIP phone/mail/instant messaging and presence to enhance communication capabilities.

IT Infrastructure / 22

Integration Lemon/LAS monitoring with the CERN Agile Infrastructure

Author: Ivan Fedorko¹

¹ CERN

Corresponding Author: ivan.fedorko@cern.ch

The Agile Infrastructure (AI) project will deliver a solution for the CERN Computer Centre resources management.

Part of the solution will consist in a new monitoring infrastructure of which the LHC Era Monitoring (Lemon) system an early adopters.

Lemon is a client/server based monitoring system, covering performance, application, environment and facilities (e.g. temperature, power consumption, cooling efficiency, etc.) monitoring.

The Lemon Alarming System, a Lemon extension, is used at CERN for notifying the operator about error situations.

This talk covers the migration strategy to the new infrastructure as well as support for the non-Quattor environment (e.g. Puppet).

Grid, Cloud and Virtualisation / 23

scientific data cloud infrastructure and services in Chinese Academy of Sciences

Author: jianhui Li¹

Co-authors: Bo Zhang²; Haiming Zhang³

¹ no 4, south 4 street, zhongguancun, haidian district, beijing, china

² no 4, south 4th street, zhongguancun, haidian district, beijing china

³ no 4, south 4th street, zhongguancun, haidian district, beijing, china

Corresponding Author: lijh@cnic.cn

In order to solve big data challenge in scientific research, a scientific data cloud has been planned to build in Chinese Academy of Sciences, which constitutes 12 data centers and one data archive center, and will provide big data online storage, data backup, data archive and data intensive analysis services. This talk will introduce the infrastructure, key technology, including distributed file system, virtualization, monitoring, etc, and services.

Storage and Filesystems / 24

CERN Cloud Storage Evaluation

Author: Dirk Duellmann¹

¹ CERN

Corresponding Author: dirk.duellmann@cern.ch

Currently there is a growing interest in the area of cloud based infrastructures (either private or public) to implement data centres in a more scalable and manageable way and to include external resources in a more flexible way. In this context the CERN DSS group together with participation from IHEP and Huawei have investigated several cloud storage implementation with respect to their stability, performance and scalability.

In this presentation we will summarise the possible motivations to use cloud storage components, their potential roles within the HEP context and the performance results achieved in a PB size test system. We will describe the operational experience during several months of test activity and conclude with a tests plan for the next evaluation phase including replication studies between different distributed sites and different storage implementations.

Site reports / 25

NDGF site report

Author: Erik Mattias Wadenstein¹

¹ Unknown

Corresponding Author: mattias.wadenstein@cern.ch

Overview of new developments at the distributed NDGF tier-1. New hardware, new organization and new lessons learned will feature.

Security and Networking / 26

Data Center Network changes and extension to Wigner

Authors: Carles Kishimoto Bisbe¹; David Gutierrez Rueda¹; Edoardo Martelli¹

¹ CERN

Corresponding Author: david.gutierrez@cern.ch

The latest changes on CERN's Data Center Network will be presented, including the migration to High-End Brocade routers, tests and introduction of 100Gbps in the Core of the LCG and bandwidth increase on the firewall system. At the same time, a Network Architecture for the Data Center extension at Wigner will be discussed.

Security and Networking / 27

IPv6 deployment status at CERN

Authors: David Gutierrez Rueda¹; Edoardo Martelli¹

¹ CERN

Corresponding Author: david.gutierrez@cern.ch

An update on the IPv6 latest changes at CERN.

Grid, Cloud and Virtualisation / 28

Global Accounting in the Grid and Cloud

Author: John Gordon¹

Co-authors: Alison Packer²; Will Rogers³

¹ STFC - Science & Technology Facilities Council (GB)

² S

³ STFC

Corresponding Author: john.gordon@stfc.ac.uk

Running jobs all over the world requires a method of recording and aggregating usage of users and VOs to present the worldwide view of that usage. The APEL Accounting systema has done that successfully for the cpu usage in the worldwide LHC Computing Grid since 2004. This presentation will cover the evolution of APEL and all the other systems who helped collect the data. It will also report on accounting of other types of usage (eg storage), how usage from the cloud can be incorporated and future evolution planned or required.

Storage and Filesystems / 29

Lustre at GSI

Author: Thomas Roth¹

Co-author: Walter Schon

¹ GSI

Corresponding Author: t.roth@gsi.de

Since March 2012, GSI is running a second Lustre file system. The younger cluster introduced a host of new technologies, problems and challenges. Step-by-step migration of both data and hardware from the old to the newer system is under way. This younger installation is the current work horse for GSI HPC, but also provides experience and knowledge base for future projects coming with FAIR. Related is the TeraLink project connecting compute clusters at neighboring sites to GSI's Lustre. A first link has been established to the CSC supercomputer at Frankfurt University.

IT Infrastructure / 30

ITIL at CC-IN2P3

Author: Frédéric AZEVEDO¹

¹ *CC-IN2P3*

Corresponding Author: frederic.azevedo@cc.in2p3.fr

IN2P3 Computing Center cares about the quality of its services and tries to improve processes and tools using ITIL best practices.

In this talk, I'll describe what we are doing on quality. I'll show the different ongoing work : the ticketing system, the CMDB, the service catalog, the business continuity plan, the identity management, ...

I'll take some time to go deeper into the change of our ticketing system to OTRS : why change ? what software were evaluated and how ? what changes are expected in daily work ? what more than ticketing could we do with it ?

Security and Networking / 31

The HEPiX IPv6 Working Group

Author: David Kelsey¹

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

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

This talk will provide an update on the activities of the IPv6 working group since the Prague meeting.

Security and Networking / 32

Federated Identity Management for HEP

Author: David Kelsey¹

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

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

This talk will present an update on the activities in Federated Identity since the last HEPiX meeting.

Storage and Filesystems / 33

RAL Tier1 Disk only Storage Project status and plans

Author: Ian Collier¹

Co-author: Shaun De Witt²

¹ *UK Tier1 Centre*

² *STFC - Science & Technology Facilities Council (GB)*

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

A working group has been investigating alternatives to Castor for disk only storage at the RAL Tier 1. Requirements have been gathered and we are now deploying test instances of a number of technologies. This talk will discuss both the reasons for the project, requirements and the current status and findings.

IT Infrastructure / 34

Quattor update - Integrating Aquilon into a grid site

Author: Ian Collier¹

¹ *UK Tier1 Centre*

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

Aquilon is a Quattor configuration database and management broker developed by an investment to meet the needs of their large worldwide grid. Providing much better relational integrity to the Quattor configuraton database and a workflow that is both more agile and more disciplined, Aquilon can transform the use of Quattor to manage sites.

This talk will discuss RAL Tier 1 experiencing deployed an Aquilon instance outside its original commercial environment and beginning to use it for managing a grid site.

Grid, Cloud and Virtualisation / 35

STFC Scientific Computing Department Cloud computing

Author: Ian Collier¹

¹ *UK Tier1 Centre*

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

A report on various projects investigating and using cloud computing technologies across the Scientific Computing Department.

Computing / 36**LRMS Migration at GridKa**

Author: Manfred Alef¹

¹ *Karlsruhe Institute of Technology (KIT)*

Corresponding Author: manfred.alef@kit.edu

This talk describes the scheduled migration to another LRMS at GridKa:
- Problems and limitations of the LRMS which is currently used at GridKa
- Selection and tests of a new one
- Configuration details, e.g. fair-share configurations, and experiences with a first sub-cluster which is already managed by the new LRMS

Site reports / 37**RAL Site Report**

Author: Martin Bly¹

¹ *STFC-RAL*

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

News from RAL

IT Infrastructure / 38**JASMINE/CEMS and EMERALD**

Author: Peter Oliver¹

Co-author: Martin Bly²

¹ *STFC-RAL*

² *STFC-RAL*

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

Details of the new e-Infrastructure South services at RAL, including the 4.5PB Panasas installation and the GPU service.

Site reports / 39**The ATLAS Great Lakes Tier-2 (AGLT2) Site Report**

Author: Shawn Mc Kee¹

Co-authors: Ben Meekhof²; Philippe Alain Luc Laurens³; Raymond Brock³; Robert Ball¹; Thomas Rockwell³

¹ *University of Michigan (US)*

² *University of Michigan*

³ *Michigan State University (US)*

Corresponding Author: shawn.mckee@cern.ch

We will present an update on our site since the last report and cover our work with VMware, dCache and perfSONAR-PS. In addition we will discuss our new denser storage system from Dell, recent networking changes and describe how we are integrating these into our site. We will conclude with a summary of what has worked and what problems we encountered and indicate directions for future work.

Summary:

Update on AGLT2 including changes in software, hardware and site configurations and summary of status and future work.

Miscellaneous / 40

Workshop wrap-up

Author: Helge Meinhard¹

Co-author: Sandy Philpott²

¹ *CERN*

² *JLAB*

Corresponding Author: helge.meinhard@cern.ch

Closing comments

Storage and Filesystems / 41

News from HEPiX Storage Working Group

Corresponding Author: andrei.maslennikov@caspur.it

Security and Networking / 42

Service Provisioning and Security Guarantee in CSTNET

Author: Yulei Wu¹

¹ *CSTNET*

Corresponding Author: wuyulei@cstnet.cn

CSTNET is the non-profitable, nationwide academic network in China, with the aim of providing Internet service and applications for the demand of scientific research and constructing an innovative network environment for the future ICT. The service provisioning and security guarantee are key factors to ensure the success of an operational network.

In this talk, we will first give an overview of CSTNET, followed by the introduction of its infrastructure. And then, we will demonstrate the advanced scientific network services and applications provided by CSTNET. The main services include the network management cloud service, network security cloud service, unified communication service, Duckling collaboration working environment service, and network research and experimentation service. The main network applications include light-path provisioning for eVLBI and its tracking for Chang'E-1&2 Lunar Mission, massive data transmission for IHEP-NERSC Daya Bay neutrino experiment, and integrated services for ITER.

Finally, we will present how to guarantee the security in CSTNET, which is provided by the network security infrastructure and the security cloud platform. The network security cloud delivers a clear and achievable path for network administrators using SAAS paradigm to achieve the centralized and unified monitoring and to provide multitenant, on-demand, location-independent network security services. Based on the security facilities and security operating center, both personalized special services and cloud-based general services are provided for 100 institutions in CAS, including security monitoring and situation awareness service, malicious code in-depth analysis service, emergency response on information security service, security assessment and reinforcement service, and safety training service.

Computing / 43

Testing SLURM batch system for a grid farm: functionalities, scalability, performance and how it works in a GRID environment

Author: Giacinto Donvito¹

¹ INFN-Bari

We will show all the work done in order to install and configure the batch system itself together with the security configuration needed. In this presentation we will show the results of the deep testing that we have done on SLURM, in order to be sure that it will cover all the needed functionalities like: priorities, fairshare, limits, QoS, failover capabilities and others. We will report also on the possibility of exploiting this batch system within a complex mixed farm environment where grid job, local job and interactive activities are managed exploiting the same batch system. From a point of view of the scalability we will show how the SLURM batch system is able to deal with the increasing number of node, CPU and jobs served. We will also show the performance achieved with several client accessing the same batch server. We also will make some comparison with other available open source batch system both in terms of performance and functionalities. We will also provide feedback on mixed configuration with SLURM and MAUI as job scheduling. We will also describe the work done in order to support SLURM in a EGI grid environment.

Summary:

As the grid computing farm are increasing in size in terms of nodes but even more in terms of CPU slots available, it become of great interest to have a scheduler solution that could scale up to tens of thousands of CPU slots and hundreds of nodes. In order to try to keep the Total Cost of Ownership as low as possible it will be preferred to have an easy to use and open source solution. SLURM is able to fulfil all those requirements and it looks promising also in terms of community that is supporting it, as it is used in several of the TOP500 supercomputing. For this reason we deeply tested the SLURM batch system in order to prove if it could be a suitable solution. In the work we will present the result of all the test executed on SLURM batch system and the results of the development activity carried on in order to provide the possibility to use SLURM in a grid computing environment.

Grid, Cloud and Virtualisation / 44

EGI Federated Cloud Infrastructure

Author: Matteo Turilli¹

Co-author: Ian Collier²

¹ *University of Oxford*

² *UK Tier1 Centre*

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

Follow up to talk at last HEPiX - describing recent developments & roadmap.

Detailed abstract to follow.

Computing / 45

A decade of Condor experience at Fermilab

Author: Steven Timm¹

¹ *Fermilab*

Corresponding Author: timmssteve@yahoo.com

The Condor Batch System has been used at Fermilab for a decade in the Run II Reprocessing and Analysis, the USCMS Tier 1 facility, and the FermiGrid General Purpose Grid Cluster. In this talk I present an overview of the operational stability, the scalability, and the best practices we have learned to build a 27,000 job slot campus grid using the Condor system.

Miscellaneous / 46

Welcome address

Author: Yifang Wang¹

¹ *IHEP*

Miscellaneous / 47

Workshop logistics

Corresponding Author: gang.chen@cern.ch

Computing / 48

Oracle Grid Engine at CC-IN2P3 - report after One year

Author: philippe olivero¹

¹ *CC-IN2P3*

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

CC-IN2P3 has been running OGE for more than one year now. After describing the current context, I will report the difficulties encountered, solved or not, and the new enhancements we would like to get.

Computing / 49

Slurm Experiences for WLCG in the Nordics

Author: Erik Mattias Wadenstein¹

¹ *Unknown*

Corresponding Author: mattias.wadenstein@cern.ch

Many compute clusters in the nordics run Slurm, this includes the grid connected ones. This talk looks at the experience, which parts works well, what could use improvements, and some comparisons to other batch systems.

Site reports / 50

LAL and GRIF site report

Author: Michel Jouvin¹

¹ *Universite de Paris-Sud 11 (FR)*

Corresponding Author: jouvin@lal.in2p3.fr

LAL and GRIF changes since 1 year

51

A new shared datacenter in Orsay

Author: Michel Jouvin¹

¹ *Universite de Paris-Sud 11 (FR)*

Corresponding Author: jouvin@lal.in2p3.fr

Fundamental research labs in southern part of Paris, known as P2IO, have decided to build a common computing facility shared between labs and designed to be energy efficient. This involves the creation of a new datacenter as the first step. This project started one year ago is about to enter its construction phase. This presentation will detail the requirements and technical design and introduce the current ideas for the future computing platform.

Storage and Filesystems / 52**Alternatives to Posix. Lessons with S3 Compatible Storage Systems****Author:** James Hughes¹¹ *Huawei*

There are three major areas that of this presentation. First, a trend in storage systems away from tape and Posix to alternative systems that favor scale, latency tolerance, and the integration of storage backup processes including device and system trends. Second, how the features of the OpenStack/Swift and Huawei storage system meet these trends. Third, a little about the testing that has been completed in CERN OpenLab.

Site reports / 53**DESY Site Report****Author:** Peter van der Reest¹¹ *DESY***Corresponding Author:** peter.van.der.reest@desy.de

The site report will discuss changes and developments at the Hamburg and Zeuthen sites since Spring.

IT Infrastructure / 54**Selecting a Business-Process-Management-System in conjunction with an Identity-and-Access-Management-System****Author:** Dirk Jahnke-Zumbusch¹¹ *DESY*

Business processes are integral parts of every day (non-technical) administrative tasks. Many of these tasks at DESY are still paper-bound. A joint project of DESY-Administration and High-Energy-Physics department was started to provide the organisational and technical prerequisites for establishing electronic workflows by using business process management systems (BPMS) as well as an identity and access management (IAM) system. Thus processes will be handled faster, be traceable better and executed in a uniform manner. The presentation will show aspects of the procedures in choosing a BPM-system as well as an IAM-system and underlying requirements.

Grid, Cloud and Virtualisation / 55**CERN and Helix Nebula, the Science Cloud****Authors:** Daniel van der Ster¹; Fernando Harald Barreiro Megino¹; Ramon Medrano Llamas¹; Rodney Walker²

¹ CERN

² Ludwig-Maximilians-Univ. Muenchen (DE)

Corresponding Author: fernando.harald.barreiro.megino@cern.ch

Helix Nebula, the Science Cloud, is a collaborative effort of several European organizations, including CERN, ESA and EMBL, to engage with European industry in public-private partnerships to build a European cloud infrastructure capable of supporting the missions of these organisations.

During the initial pilot phase of Helix Nebula, the ATLAS experiment at CERN was selected as one of the flagship projects and a proof-of-concept phase was defined in order to demonstrate the feasibility of integrating commercial cloud facilities into the ATLAS distributed computing infrastructure.

This talk will outline the status of Helix Nebula and present the results of the ATLAS use-case in particular. Three commercial cloud providers with varied infrastructures were tested; all were successfully able to run ATLAS simulation jobs, though the paths to success at some providers were more difficult than others. We will give an insight into the lessons learned, the technical recommendations for the supply side and some of the future work in the Helix Nebula partnership.

Site reports / 56

Site Report of ASGC

Author: YEN Eric¹

¹ ASGC

Corresponding Author: eric.yen@twgrid.org

Update of e-Science infrastructure at ASGC will be reported. Focus will cover the international networking, computing and storage infrastructure, overall continuous operation improvement including the data center, e-Science application and virtual research framework, as well as the development of distributed cloud.

IT Infrastructure / 57

Scientific Linux Infrastructure Improvements

Authors: Pat Riehecky¹; connie sieh¹

¹ Fermilab

Corresponding Author: riehecky@fnal.gov

The underlying infrastructure of Scientific Linux is starting to change. These changes should make the environment more stable and higher performing with a greater feature set. This presentation will detail some of the plans and progress thus far.

Security and Networking / 58

FZU IPv6 testbed updates

Author: Marek Elias¹

Co-authors: Jiri Chudoba ¹; Lukas Fiala ¹; Tomas Kouba ¹

¹ *Institute of Physics AS CR (FZU)*

we present updates since Vancouver about our IPv6 testbed at FZU. We have setup nagios, smokeping several middleware services and several computing centre management procedures.

Computing / 59

Report from batch system BOF

Author: Aresh Vedae¹

¹ *CC-IN2P3 - Centre de Calcul (FR)*

Report from the BOF session the day before