

HEPiX Fall 2010 Workshop

Monday 01 November 2010 - Friday 05 November 2010

The Statler Hotel

Book of Abstracts

Contents

A scheme for defining and deploying trusted Virtual Machines to Grid Sites using Configuration Management Systems	1
ASGC site report	1
ASSETS at LEPP - our FLOSS Inventory and Monitoring server	1
ATLAS Analysis on ARC	2
Access Grid via Web	2
BIRD: Batch Infrastructure Resource at DESY	3
BNL Site Report	3
BNL storage experiences	3
Board Summary and Meeting Wrap-Up	4
CASTOR development status and deployment experience at CERN	4
CC-IN2P3 Infrastructure Improvements	4
CC-IN2P3 Site report	4
CERN Computer Centre Status and Proposed Upgrade	5
CERN IT Facility Planning and Procurement	5
CERN site report	5
CERNs image distribution system for the internal cloud	5
CloudCRV - Cluster Deployment and Configuration Automation on the Cloud	6
Current storage status and plans at IN2P3	7
DESY Site Report	7
Deployment of Exchange 2010 mail platform	7
Digital Library and Conferencing update	7
FermiCloud - Current Status	8
Fermilab Site Report - Fall 2010 HEPiX	8

First results from the WLCG NFS4.1 Demonstrator	8
HEP and IPv6	8
High Performance Storage Pools for LHC	9
INFN-T1 site report	9
IPV6 @ INFN	9
JLab HPC Upgrades - GPU and Lustre Experiences	9
Jefferson Lab Site Report	10
KISTI - GSDC site report	10
LAL + GRIF Site Report	10
LEPP Site Report	11
Lessons learnt from Large LSF scalability tests	11
Lustre Consortium BOF	11
Magellan at NERSC: A Testbed to Explore Cloud Computing for Science	11
Measurement of HS06 on Intel Westmere and AMD Magny-Cours processors	12
NDGF Site Report	12
NERSC/PDSF Status report - A Year of Changes	12
New Tools Used by the S.L. Team	13
New network architecture at IN2P3-CC	13
Plans for a Single Kerberos Service at CERN	14
Prague Institute of Physics	14
Progress Report 4.2010 for HEPiX Storage Working Group	14
Quattor Update	14
RAL Site Report	15
Rapid web application design for silicon detector measurements	15
Registration	15
Report from the Virtualisation Working Group	15
SLAC Site Report	16
Saclay (IRFU) site report	16
Scientific Linux Status Report and Plenary Discussion	16
Site Report GSI	17

Status update of the CERN Virtual Infrastructure	17
Storage at FNAL: state and outlook	17
StratusLab, mixing grid and clouds	17
The CERN internal cloud infrastructure: a status report	18
Update on Windows 7 at CERN & Remote Desktop Gateway	18
Update on computer security	19
Update on the CERN Search Engine	19
Update on the anti spam system at CERN	19
VOMS/VOMRS Convergence	20
Virtualization at CERN: an overview	20
Welcome	21
cvmfs - a caching filesystem for software distribution	21
distcc at CERN	21

Virtualization / 25**A scheme for defining and deploying trusted Virtual Machines to Grid Sites using Configuration Management Systems**Yushu Yao¹¹ *Lawrence Berkeley National Lab. (LBNL)***Corresponding Author(s):** yao.yushu@gmail.com

Sites (service providers) and VOs (service users/functionality providers) are debating how to define trusted VO-provided VM images. We present an alternative scheme that, using a Configuration Management System (CMS), does not require Sites to trust VO-provided images. It gives VO the freedom to design and customize functionalities, while letting Sites retain full control over instances.

Configuration Management Systems (CMS, such as Puppet/Cfengine, have been widely used in computing centers to automatically manage resources. They provide high level languages to define desired states of target systems, such as installing software packages, running services, enforcing firewall, etc.

We propose a two level scheme. Both VO and Site start from a well trusted base image (e.g. base installation of SL5 with no customization). At the VO-level, VO experts customize the base image to perform VO specific tasks (e.g. ATLAS Condor Worker). These customizations are not committed to the image, instead, they are defined in CMS language and stored in a SVN repository. At the Site-level, site experts define site-specific configurations and security policies in CMS language.

When deploying, VO needs no privileged access in instances. Site manager starts the base-images that contains CMS clients, which apply VO-level definitions then Site-level definitions. At this time, the instance is ready to perform VO defined tasks. Since Site-level definitions are applied at last, the VMs are ensured to comply with Site policies.

This scheme eliminates the problem of trusting VM images. It is more flexible, more reliable, and more secure.

Summary:

Comments to Reviewers:

The approach is based on the fact that, with the help of Configuration management systems, the VM can be customized/contextualized to meet the requirements of both VO and Site at boot time. So that root access is not required for VOs on deployed instances.

Benefits of this approach:

1. Eliminate the need for sites to trust an image provided by an individual.
2. VO needs no root access on deployed VMs (running instances).
3. Easier to expire, revoke or catalog a set of definitions than a VM image.
4. Versioning control of VO customizations and Site policies.
5. Give VO the freedom to customize to meet their needs
6. Give Sites the ease to enforce security policies

Site Reports / 57**ASGC site report**Hung Te Lee^{None}**Datacenter and Monitoring / 66****ASSETS at LEPP - our FLOSS Inventory and Monitoring server**James Pulver¹

¹ *Cornell University*

Corresponding Author(s): jmp242@cornell.edu

Inventory and monitoring are two areas that have become more important at LEPP over the last few years in the IT area. We've had increased requirements to track what computers we buy, who gets them "out of stock" and various status changes, from location changes through to disposal.

Knowing where a computer is located is only part of the battle. Important servers and systems must be tracked as to whether they are available, are necessary services running and what remediation is necessary in down events.

LEPP has implemented a Scientific Linux 5 x64 server running three FLOSS products to manage our needs in these areas.

Automatic Inventory is achieved on Windows and Linux nodes via OCSNG and some software deployments are also managed via that tool.

GLPI imports the automatically gathered data for additional manual inventory information such as ownership, stock management, ongoing historical notes etc...

Zenoss provides continuous monitoring and performance graphing for our important infrastructure. Events are stored locally for reference, and alerts are generated for specific events of interest and delivered via e-mail.

Grids and Clouds / 28

ATLAS Analysis on ARC

Erik Mattias Wadenstein¹

¹ *NDGF*

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

An overview and in places detailed case study on how atlas analysis jobs work on ARC CEs currently, and a look at future developments. With some efficiency and cache utilization numbers.

Grids and Clouds / 19

Access Grid via Web

Daniele Licari¹ ; Federico Calzolari²

¹ *University of Pisa*

² *Scuola Normale - INFN*

Corresponding Author(s): federico.calzolari@cern.ch

L-GRID is a light portal to access Grid infrastructure via Web browser, allowing users to submit their jobs in a few minutes, without any knowledge about the Grid infrastructure.

The portal is intended to be a helpful tool to access Grid resources shared all around the world via a simple Web interface, using whatever operating system and browser. It provides the control over the complete lifecycle of a Grid Job, from its submission and status monitoring, to the output retrieval. The end user needs only her/his own X.509 personal certificate, issued from a Certification Authority.

The system, implemented as client-server architecture, is based on the gLite Grid middleware. The client side application is based on a java applet, running both on Windows, Linux and Mac operating systems; it only needs a Web browser connected to the Internet. The server relies on a gLite User Interface with Web portal provided by an Apache/Tomcat server. The main differences with respect to a native gLite User Interface are the extreme ease of use and the no-need of the user registration.

L-GRID provides the typical operations involved in a Grid environment: certificate conversion, job submission, job status monitoring, and output retrieval. It provides also a JDL editor. The system is user-friendly, secure (it uses SSL protocol, mechanism for dynamic delegation and identity creation in public key infrastructures), highly customizable, open source, and easy to install - the installation requires a few MB. The X.509 personal certificate does not get out from the local machine, strictly compliant to the Certification Authority policies, and the Grid commands are splitted into client and server, increasing the security level.

An extra security improvement has been achieved by the inclusion of the MyProxy server, responsible for the dynamic delegation in long term proxy certificates, on the server side portal. It allows to reduce the time spent for the job submission, granting at the same time a higher efficiency and a better security level in proxy delegation and management.

The first running prototype is hosted at the moment at the High Performance Computing Center of the Scuola Normale Superiore, Pisa, Italy.

The results obtained encourage future developments. Further steps are represented by the integration with a LDAP Kerberos AAI Authentication Authorization Infrastructure, and the customization for LHC and Theophys Virtual Organizations.

Summary:

L-GRID is a light portal to access Grid infrastructure via Web browser.

Datacenter and Monitoring / 4

BIRD: Batch Infrastructure Resource at DESY

Thomas Finnern¹

¹ *DESY*

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

The BIRD cluster is a multi-core batch computing facility based on the Grid Engine Software. It provides resources for compute intensive applications running under Scientific Linux. The talk covers the basic design as well as the implementation of advanced features like afs/kerberos integration, parallel environments and interactive queues. Special demands for big jobs with up to 64 GByte memory and 250 GByte scratch space have been realized.

Site Reports / 68

BNL Site Report

Corresponding Author(s): tony@bnl.gov

Storage and File Systems / 65

BNL storage experiences

Maurice Askinazi¹ ; Ofer Rind¹ ; Tony Wong¹

¹ *BROOKHAVEN NATIONAL LAB*

Corresponding Author(s): tony@bnl.gov

BNL will provide a brief update on its storage choices (hardware and software) in the context of the evolving historical nature of its storage needs and requirements. A brief outlook on testbeds and possible future choices will also be presented.

Wrap-up / 56

Board Summary and Meeting Wrap-Up

Storage and File Systems / 43

CASTOR development status and deployment experience at CERN

Lukasz Janyst¹

¹ *CERN*

Corresponding Author(s): ljanyst@cern.ch

In the presentation we will give an overview of recent CASTOR developments focused on further consolidating the system and lowering its deployment cost. We will outline the release plan for the medium term and give an update on the operational experience gained with CASTOR during the last year of LHC running.

Datacenter and Monitoring / 48

CC-IN2P3 Infrastructure Improvements

philippe olivero¹

¹ *CC-IN2P3*

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

CC-In2p3 is currently building an additional machine room to face storage and computing requests up to year 2020. A status of both the current and the next computing room will be presented, reporting latest enhancements in the first one, and some considerations about the new one.

Site Reports / 0

CC-IN2P3 Site report

philippe olivero¹

¹ *CC-IN2P3*

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

Report hardware and software updates since one year

Datacenter and Monitoring / 46

CERN Computer Centre Status and Proposed Upgrade

Wayne Salter¹

¹ *CERN*

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

This presentation will quickly summarise the current status of the CERN Computer Centre in terms of available/used power and cooling. It will then go on to describe a project which is currently underway to increase the available capacity as well as to address a number of long standing issues.

Datacenter and Monitoring / 47

CERN IT Facility Planning and Procurement

Olof Barring¹

¹ *CERN*

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

The talk covers some aspects of the planning and procurement of server and storage hardware for installation in the CERN IT facility. The current arrangements for warranty services will also be discussed as well as some its perceived limitations.

Site Reports / 10

CERN site report

Helge Meinhard¹

¹ *CERN-IT*

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

Site report

Virtualization / 32

CERNs image distribution system for the internal cloud

Belmiro Moreira¹; Ewan Roche¹; Juan Manuel Gujarro¹; Romain Wartel¹; Sebastien Goasguen²; Ulrich Schwickerath¹

¹ CERN

² Clemson University

Corresponding Author(s): romain.wartel@cern.ch

In any large-scale virtualized environment, the image distribution system plays a central role. It must not only deliver up-to-date images to the hypervisor with appropriate performances, but it should also establish and maintain a sufficient level of trust in the actual images that are being distributed. These issues can be addressed by using a peer-to-peer-based image distribution system, depending on a trusted and signed index providing an up-to-date list of images that hypervisors can trust. This presentation describes the technologies used to establish such an infrastructure at CERN, the challenges encountered in their deployment and in their integration in a production environment, as well as the future developments, in particular to enable virtual image sharing with other sites.

Grids and Clouds / 26

CloudCRV - Cluster Deployment and Configuration Automation on the Cloud

Yushu Yao¹

¹ Lawrence Berkeley National Lab. (LBNL)

Corresponding Author(s): yao.yushu@gmail.com

With the development of virtualization technology and IaaS cloud, it is much easier for users to obtain large number of computing (virtual) resource. However, customizing these resources as a computing clusters remains a difficult task that requires in-depth IT knowledge and complex user-specific customizations. We develop a tool called CloudCRV (Cloud-Cluster-Role-VMs) to help users design, distribute and deploy a secure, functional cluster on the allocated resources.

We believe a predefined cluster in a whole can be distributed as a product to perform certain task (e.g. an ATLAS Tier3 cluster), we call this kind of product a Virtual Cluster Appliance (As extension to Virtual Appliance). The purpose of CloudCRV is to help the Cluster Designer to design such a product and to help the Cluster Managers to deploy it.

Most clusters can be abstracted to a set of Roles (e.g. a NFS server, or a Condor Head) and their relations (Condor Head depends on NFS server). Cluster Designer's work is to define the Roles and their relations. The Roles are defined with the help of configuration management systems such as Puppet or Cfengine. Once designed, the Virtual Cluster Appliance can be deployed at multiple sites by local Cluster Managers onto physical or virtual resources. CloudCRV provide interfaces to both Cloud Providers (such as EC2 and Nimbus), and to physical computers and libvirt based clusters via gPXE remote booting and image deployment.

In this contribution we demonstrates the process of designing and deploying such a Virtual Cluster Appliance with the help of CloudCRV.

Summary:

Comments to the Reviewers:

We define the following three concepts:

VM:

A running virtual or physical computer that has a supported OS.

Role:

a set of requirements on a VM that once met, will let the VM provide certain functionality. E.g. NFS server, Condor Head, Condor Worker are all Roles. One Role can depend on another Role, e.g. Condor Worker depend on Condor Head.

Roles are defined by Configuration Management System Languages (e.g. Puppet/CFengine).

Cluster:

A collection of Roles.

CloudCRV contains two parts, the CloudCRV Designer will help a Cluster Designer (e.g. the Tier3 Working group of ATLAS) to design the cluster, test its functionalities, and package it. As a Cluster Designer, one has to have certain level of system administration knowledge.

The CloudCRV Manager, on the other hand, will help a Cluster Manager to customize and deploy a pre-defined cluster. The Cluster Manager needs to have access to the hardware resources (e.g. and EC2 account, or a physical cluster), and little system administration knowledge is needed.

The key is design once, deploy multiple times. This gives a normal cloud user the power of running a cluster with least effort.

Storage and File Systems / 13

Current storage status and plans at IN2P3

Pierre-Emmanuel Brinette¹

¹ *IN2P3*

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

Site Reports / 21

DESY Site Report

Peter Van Der Reest¹

¹ *DESY*

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

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

Operating Systems and Applications / 41

Deployment of Exchange 2010 mail platform

Pawel Grzywaczewski¹

¹ *CERN*

Corresponding Author(s): pawel.grzywaczewski@cern.ch

CERN is in process of deploying new version of mail system - Microsoft Exchange 2010. The talk gives an overview of new features introduced in Exchange 2010 and about the deployment process.

Miscellaneous / 14

Digital Library and Conferencing update

Tim Smith¹

¹ *CERN*

Corresponding Author(s): tim.smith@cern.ch

A successor to the venerable SPIRES has been prepared by CERN, DESY, FNAL and SLAC; called INSPIRE. I will describe all the digital library services this will provide to the HEP community and wider. I will also give an update on Indico, its recent and planned developments.

Grids and Clouds / 3

FermiCloud - Current Status

Keith Chadwick¹

¹ *Fermilab*

Corresponding Author(s): chadwick@fnal.gov

The current status of FermiCloud will be presented together with the experience and “lessons learned”.

Site Reports / 2

Fermilab Site Report - Fall 2010 HEPiX

Keith Chadwick¹

¹ *Fermilab*

Corresponding Author(s): chadwick@fnal.gov

Fall 2010 Fermilab Site Report

Storage and File Systems / 12

First results from the WLCG NFS4.1 Demonstrator

Patrick Fuhrmann¹

¹ *DESY*

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

Security and Networking / 22

HEP and IPv6

David Kelsey¹

¹ *RAL*

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

This session will include a report on the answers given to a recent questionnaire on IPv6 status. Discussion as to what (if any) coordination is required. Should we create a HEPiX group on this topic?

Storage and File Systems / 42

High Performance Storage Pools for LHC

Lukasz Janyst¹

¹ *CERN*

Corresponding Author(s): ljanyst@cern.ch

The Data and Storage Services Group at CERN is continuing its strategy to provide highly scalable storage components to support LHC analysis and data production. This contribution will summarize the recent EOS developments which are currently being tested with the experiment users. We will give an overview of the EOS system, the results from tests in a 1PB prototype pool and describe the future work-plan to evaluate the new system.

Site Reports / 69

INFN-T1 site report

Andrea Chierici¹

¹ *INFN-CNAF*

Corresponding Author(s): chierici@cnaf.infn.it

updates at INFN Tier1 center

Security and Networking / 24

IPV6 @ INFN

Francesco Prelz¹

¹ *INFN - Sezione di Milano*

Corresponding Author(s): francesco.prelz@mi.infn.it

The current understanding of the timeline and constraints of INFN-wide deployment of (native) IPv6 is presented. This talk is meant primarily as input for discussion.

Datacenter and Monitoring / 31**JLab HPC Upgrades - GPU and Lustre Experiences**Sandy Philpott¹¹ *JLAB***Corresponding Author(s):** sandy.philpott@jlab.org

JLab's HPC environment for Lattice QCD has recently been upgraded, including the additions of GPUs and Lustre. We are now running two GPU-enabled clusters with NVIDIA GeForce GTX-285, GTX-480, and Tesla C2050 capabilities, in addition to our 3 IB clusters. We are also running a Lustre filesystem on Amax storage servers. This talk will share our experiences integrating these new technologies into our environment.

Site Reports / 30**Jefferson Lab Site Report**Sandy Philpott¹¹ *JLAB***Corresponding Author(s):** sandy.philpott@jlab.org

Update since the spring meeting report at LIP.

Site Reports / 27**KISTI - GSDC site report**Christophe Bonnaud¹¹ *KiSTi Korea Institute of Science & Technology Information (KiSTi)***Corresponding Author(s):** cbonnaud@kisti.re.kr

Presentation of Global Science Data Center (GSDC) project at KISTI: status of activities, system infrastructure and futur plans.

Site Reports / 34**LAL + GRIF Site Report**Michel Jouvin¹¹ *LAL / IN2P3***Corresponding Author(s):** jouvin@lal.in2p3.fr

Site report about GRIF and LAL.

Site Reports / 62**LEPP Site Report**Devin Bougie¹¹ *Cornell University***Corresponding Author(s):** devin.bougie@cornell.edu

Introduction to the research at Cornell's Laboratory for Elementary-Particle Physics and the computing behind it.

Datacenter and Monitoring / 7**Lessons learnt from Large LSF scalability tests****Author(s):** Ulrich Schwickerath¹**Co-author(s):** Gavin McCance¹ ; Ricardo Silva¹¹ *CERN***Corresponding Author(s):** ulrich.schwickerath@cern.ch

During summer 2010, a large LSF test cluster infrastructure was put in place to allow scalability tests of the batch software (LSF) at a scale which exceeds the production instance by up to a factor 5.

The response time of several central commands was measured as a function of the number of worker nodes and the number of batch nodes in the farm.

Several issues which were found during the tests were fixed on the fly by the vendor. This way, it was possible to go up to 15,000 virtual worker nodes, and more than 400,000 jobs in the system. Some results from these scalability tests will be presented, lessons learned during the tests, and possible consequences for planning will be discussed.

Storage and File Systems / 73**Lustre Consortium BOF****Grids and Clouds / 71****Magellan at NERSC: A Testbed to Explore Cloud Computing for Science****Author(s):** Shane Canon¹**Co-author(s):** Iwona Sakrejda¹ ; Lavanya Ramakrishnan¹¹ *LBNL/NERSC***Corresponding Author(s):** isakrejda@lbl.gov

Cloud computing is gaining a foothold in the business world, but can clouds meet the specialized needs of scientists? That is the question NERSC's Magellan cloud computing test bed is exploring.

Funded by the American Recovery and Reinvestment Act (Recovery Act) through the U.S. Department of Energy (DOE), the system is distributed between DOE centers: the National Energy Research Scientific Computing Center (NERSC) in California and the Argonne Leadership Computing Facility (ALCF) in Illinois. Research efforts range from evaluating what applications work well on today's commercial cloud offerings to studying how jobs can be distributed across multiple DOE clouds and exploring emerging programming models like MapReduce. We will provide an overview of the project and present some of our findings to date. This includes recent results of the performance of scientific applications running on commercial cloud offerings compared with traditional systems.

Benchmarking / 33

Measurement of HS06 on Intel Westmere and AMD Magny-Cours processors

Michele Michelotto¹

¹ *Univ. + INFN*

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

I'll report measurements performed in Padova in the CSN5 experiment HEPMARK. HS06 has been measured on Intel Westmere 5650 and a couple of the new AMD Magny-Cours processor.

Site Reports / 29

NDGF Site Report

Erik Mattias Wadenstein¹

¹ *NDGF*

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

Current status and recent developments in NDGF and the sites making up the distributed tier1.

Site Reports / 72

NERSC/PDSF Status report - A Year of Changes

Elizabeth Bautista¹ ; Iwona Sakrejda¹ ; Jay Srinivasan¹

¹ *LBNL/NERSC*

Corresponding Author(s): isakrejda@lbl.gov

PDSF is a networked, distributed computing environment used to meet the detector, simulation, and data analysis requirements of physics (large-scale, high-energy physics, and astrophysics) and nuclear science investigations.

Since our last report two years ago, the cluster has been upgraded significantly. We retired older nodes, expanded compute and storage capacities, and fortified network connectivity. While eight-core AMD compute nodes were retained, older nodes were replaced dual quadcore- (Intel Nehalem) and dual hexcore- (Intel Westmere) based Dell PowerEdge R410 systems. Total cluster capacity also grew by 50 percent, from 800 to 1200 job slots. Total storage capacity has been expanded to about 1TB (comprised of a combination of GPFS and xRootd file systems deployed atop mostly Dell SAS-based units). Networking infrastructure, although still 1 GigE-based, has been fortified with new switches (which couple the PDSF cluster to ESnet at 10 Gbps) and additional data transfer nodes with 10GigE network connectivity.

Operating Systems and Applications / 50

New Tools Used by the S.L. Team

Author(s): Troy Dawson¹

Co-author(s): Connie Sieh¹

¹ *FERMILAB*

Corresponding Author(s): dawson@fnal.gov

The Scientific Linux Team has been testing new tools to make life easier for development, for site maintainers, admins, and end users. This presentation will talk about our work with Koji, Spacewalk, Revisor, and other tools.

Summary:

Koji, Spacewalk, Revisor, Tools, Scientific Linux Development Team

Security and Networking / 9

New network architecture at IN2P3-CC

Author(s): Guillaume Cessieux¹

Co-author(s): Gerard Drevon¹ ; Jérôme Bernier¹ ; Laurent Caillat-Vallet¹

¹ *CNRS/IN2P3-CC*

Corresponding Author(s): guillaume.cessieux@cc.in2p3.fr

The Computer Centre of the French National Institute for nuclear and particle physics (IN2P3-CC), located in Lyon, has recently rolled out a major network upgrade. The previous network architecture, nearly 4 years old, reached limits and an upgrade was necessary to face new challenges, particularly massive data transfers, virtualisation, heavy Grid computation and an upcoming additional computing room.

After a thorough analysis of current network devices (feature, topology, configuration, usage) and network behaviour (identifying traffic patterns, main areas of exchange, bottlenecks, major consumers and producers) a new architecture was designed. A key objective, besides removing bottlenecks, was to improve scalability of the network, especially by enabling seamless and non disruptive bandwidth upgrades in the future. Strong attention was paid to use configurations able to deliver wire speed.

Even with strong preliminary testing and anticipating all possible tasks (pre-wiring, creating new configurations, making checklists...) the deployment was done in September 2010 within a nightly scheduled maintenance during a 5 hours network intervention (not continuously service impacting).

We also used the maintenance window to upgrade software on 170 network devices, harmonising management and supported features.

Layout was completely re-organised to reduce as much as possible paths length for heavy exchanges. The new network architecture is built around a central redundant Cisco Nexus 7018 aggregating flows up to 60G from several key functional areas (storage, computing, WAN...). Hosts doing intensive exchanges are connected up to 10G directly through a distribution layer, mainly featuring 4900M and Catalyst 6500, while other consumers are offloaded onto an access layer. 80G are foreseen to connect the new computing room.

Summary:

The process leading to the new network infrastructure at IN2P3-CC, how it was deployed, benefits, lessons learnt and early feedback will be presented.

Security and Networking / 59

Plans for a Single Kerberos Service at CERN

Author(s): John Hefferman¹

Co-author(s): Lukasz Janyst¹

¹ *CERN*

Corresponding Author(s): john.hefferman@cern.ch

CERN IT is planning to merge CERN's two Kerberos services. The aim of this presentation is to provide an overview of:

The problems of having two Kerberos services and why this merger is being carried out

The planned 'post-merge' Kerberos infrastructure

The method which will be used to merge the two Kerberos realms and the infrastructure changes made

The project timeline and user involvement

Site Reports / 58

Prague Institute of Physics

Jan Kunderat^{None}

Corresponding Author(s): kunderatj@fzu.cz, kunderat@fzu.cz

Storage and File Systems / 11

Progress Report 4.2010 for HEPiX Storage Working Group

Andrei Maslennikov¹

¹ *CASPUR*

Corresponding Author(s): andrei.maslennikov@caspur.it

Datacenter and Monitoring / 37**Quattor Update**Ian Peter Collier¹¹ *STFC-RAL***Corresponding Author(s):** ian.peter.collier@cern.ch

The use of Quattor continues to have a positive impact on the operation of the RAL Tier 1. This talk will describe recent developments at RAL and in the Quattor Toolkit itself, and report on the 10th Quattor workshop, hosted at RAL in October 2010.

Site Reports / 16**RAL Site Report**Martin Bly¹¹ *STFC-RAL***Corresponding Author(s):** martin.bly@stfc.ac.uk

Latest Developments at RAL and the UK Tier1

Miscellaneous / 1**Rapid web application design for silicon detector measurements**Wolfgang Friebel¹¹ *Deutsches Elektronen-Synchrotron (DESY)-Unknown-Unknown***Corresponding Author(s):** wolfgang.friebel@desy.de

Developments of new silicon detectors come with a demand for comprehensive measurements of its characteristics. To allow access to the measured (and processed) data by all interested parties a central data repository combined with an adequate remote query mechanism is necessary. The talk will demonstrate how the development of a web application for this purpose can be achieved with minimal resources. By using the Open Source web framework Catalyst and a SQL database a very flexible and modular design of the entire system has been achieved. Changing requirements to the system such as DB schema changes are easy to handle.

The framework is rather generic and has also successfully tested for other applications.

Introduction / 53**Registration**

Virtualization / 23**Report from the Virtualisation Working Group**Tony Cass¹¹ *CERN***Corresponding Author(s):** tony.cass@cern.ch

The presentation will cover the work of the Virtualisation Working Group over the past year.

Site Reports / 20**SLAC Site Report**Alf Wachsmann¹¹ *SLAC***Corresponding Author(s):** alfw@slac.stanford.edu

Summary of changes at SLAC over the last 6 months.

Site Reports / 5**Saclay (IRFU) site report**Pierrick Micout¹¹ *CEA IRFU***Corresponding Author(s):** micout@cea.fr

Site report of the IRFU Saclay site.

Summary:

Everything new on the IRFU Saclay site on computing since last year.

Operating Systems and Applications / 15**Scientific Linux Status Report and Plenary Discussion**Troy Dawson¹¹ *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 / 49

Site Report GSI

Walter Schon¹

¹ *GSI*

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

News from GSI

Virtualization / 51

Status update of the CERN Virtual Infrastructure

Author(s): Jan van Eldik¹

Co-author(s): Tim Bell ¹

¹ *CERN*

Corresponding Author(s): jan.van.eldik@cern.ch

During 2010, the number of Virtual Machines in the CERN Virtual Infrastructure has doubled to 600. These VMs are owned by a large number of users from very different CERN communities.

CVI is based in Microsoft's Hyper-V product, with a web-based self-service and a SOAP interface. We present details of the service architecture, its current implementation and usage, and our plans for future enhancements. Special emphasis will be given to the SLC5 Virtual Machines.

Storage and File Systems / 18

Storage at FNAL: state and outlook

Matt Crawford¹

¹ *FNAL*

Grids and Clouds / 35

StratusLab, mixing grid and clouds

Michel Jouvin¹

¹ LAL / IN2P3

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

StratusLab (StratusLab.eu) is a two-year European FP7 project (from 1 June 2010) that aims to provide a production-quality, open source cloud (“Infrastructure as a Service” or IaaS) distribution. StratusLab will integrate, distribute, and maintain this cloud distribution to bring cloud technology to end-users and resource providers of existing distributed computing infrastructures like EGI.

The StratusLab toolkit combines existing, cutting-edge, open source software with innovative service and cloud management technologies developed within the project. The project uses agile software development practices to ensure rapid evolution of the toolkit to meet end-user and system administrator needs. It demonstrates the production quality of the toolkit by running two grid resource centers on top of the toolkit and by quantitatively testing the performance of a spectrum of representative applications on the hybrid infrastructure.

Grid and cloud technologies complement one another. Existing grid middleware would continue to provide the glue to federate the distributed resources and the services for high-level job and data management. StratusLab will help to improve usability of distributed computing infrastructures. Providing a cloud API for the grid will attract the scientific and industrial users that have embraced the cloud computing provisioning model and thus expand the scope and interest of infrastructures like EGI.

This talk will detail StratusLab’s goals and present its two-year roadmap. It will explain the expected benefits for the e-Infrastructure ecosystem and the availability and current features of the StratusLab toolkit.

Grids and Clouds / 6

The CERN internal cloud infrastructure: a status report

Author(s): Ulrich Schwickerath¹

Co-author(s): Belmiro Moreira ¹ ; Ewan Roche ¹ ; Romain Wartel ¹ ; Sebastien Goasguen ² ; Tony Cass ¹

¹ CERN

² Clemson university

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

CERNs virtualization plans have been presented during the last HEPiX meetings. Since the last HEPiX meeting, the ideas have materialized in a prototype which was used to perform large scale scalability tests of the batch system LSF. This presentation will give an overview over the architecture of the new infrastructure, and report on the experiences and lessons we learned when growing the system up to 500 machines and 15,000 virtual machine slots. It is planned to release this infrastructure into production, with the batch service being the first and initially only user. Production deployment at a small scale is planned for the second half of November, and an outlook on possible future extensions will be given.

Summary:

This presentation will be a follow-up of presentations given on this topic at previous HEPiX meetings.

Operating Systems and Applications / 38

Update on Windows 7 at CERN & Remote Desktop Gateway

Andreas Wagner¹ ; Michal Budzowski¹ ; Michal Kwiatek¹ ; Tim Bell¹

¹ CERN

Corresponding Author(s): andreas.wagner@cern.ch

Windows 7 is officially supported at CERN since March 2010. We will present the status of the NICE Windows 7 service, that is offered for both 32 and 64 bit, and share our first months of experience with this latest Windows OS version. In addition we will outline our plans to phase-out the previous versions Windows Vista and Windows XP.

Furthermore, we will present our Remote Desktop Gateway implementation that allows CERN users to connect to their on-site Desktop PCs in a secure manner from any offsite location.

Security and Networking / 17

Update on computer security

Romain Wartel¹

¹ CERN

Corresponding Author(s): romain.wartel@cern.ch

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

Miscellaneous / 39

Update on the CERN Search Engine

Andreas Wagner¹ ; Eduardo Alvarez Fernandez¹ ; Tim Bell¹

¹ CERN

Corresponding Author(s): andreas.wagner@cern.ch

The CERN Search engine facilitates access to a wide range of information such as the CERN Web pages, TWiki, CDS, Indico and the CERN Phonebook. This presentation will describe the necessary components of a enterprise wide search solution that indexes a range of heterogeneous information sources. We will present the recent work done to allow indexing of protected TWiki areas for the ATLAS and CMS experiments and we will outline the future plans for evolving the CERN Search solution.

Operating Systems and Applications / 40

Update on the anti spam system at CERN

Pawel Grzywaczewski¹

¹ CERN

Corresponding Author(s): pawel.grzywaczewski@cern.ch

In April 2010 a new email security system was deployed at CERN: Microsoft ForeFront Protection 2010 for Exchange servers. It provides both anti spam and anti virus functionalities. The talk gives an overview of the product itself and anti spam infrastructure at CERN.

Grids and Clouds / 70

VOMS/VOMRS Convergence

Author(s): Andrea Ceccanti¹

Co-author(s): Tanya Levshina²

¹ *INFN*

² *Fermilab*

Corresponding Author(s): andrea.ceccanti@cnafr.infn.it

The Grid community uses two well-established registration services, which allow users to be authenticated under the auspices of Virtual Organizations (VOs). The Virtual Organization Membership Service (VOMS), developed in the context of the Enabling Grid for E-science (EGEE) project, is an Attribute Authority service that issues attributes expressing membership information of a subject within a VO. VOMS allows to partition users in groups, assign them roles and free-form attributes which are then used to drive authorization decisions. The VOMS administrative application, VOMS-Admin, manages and populates the VOMS database with membership information.

The Virtual Organization Management Registration Service (VOMRS), developed at Fermilab, extends the basic registration and management functionalities present in VOMS-Admin. It implements a registration workflow that requires VO usage policy acceptance and membership approval by administrators. VOMRS supports management of multiple grid certificates, and handling users' request for group and role assignments, and membership status. VOMRS is capable of interfacing to local systems with personnel information (e.g. the CERN Human Resource Database) and of pulling relevant member information from them. VOMRS synchronizes the relevant subset of information with VOMS. The recent development of new features in VOMS-Admin raises the possibility of rationalizing the support and converging on a single solution by continuing and extending existing collaborations between EGEE and OSG. Such strategy is supported by WLCG, OSG, US CMS, US Atlas, and other stakeholders worldwide. In this presentation, we will give an update on the status of the convergence between the two products.

Virtualization / 8

Virtualization at CERN: an overview

Author(s): Helge Meinhard¹ ; Ulrich Schwickerath¹

Co-author(s): Ewan Roche¹ ; Juan Manuel Guijarro¹

¹ *CERN*

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

This presentation will give an overview over the various virtualization projects going on at CERN. Specifically, consolidation using virtualization has been moved into production mode, and already more than 100 machines are run on this infrastructure, covering a variety of services. The basic concepts of the setup will be presented, and future plans will be described. This presentation will also cover possible long term perspectives in virtualization at CERN, including a merge of the different project into a single internal cloud infrastructure.

Introduction / 54**Welcome****Virtualization / 36****cvmfs - a caching filesystem for software distribution**

Ian Peter Collier¹

¹ *STFC-RAL*

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

In common with other sites the RAL Tier 1 experiences performance problems on experiment software servers. CernVM-FS is a caching http based filesystem that may provide an elegant solution. This talk will describe cvmfs and RAL's experience in testing its scalability.

Operating Systems and Applications / 55**distcc at CERN**

Peter Kelemen¹

¹ *CERN*

Corresponding Author(s): peter.kelemen@cern.ch

Last year CERN has implemented GSSAPI authentication for (now Google's) "distcc", a wrapper around "gcc" for distributed C/C++ compilation and established a 128-core compile cluster available for all CERN users. The service is in production for six months, we will report on the implementation challenges and current utilization of the service.