

HEPiX Spring 2009

Monday 25 May 2009 - Friday 29 May 2009

Umeå University

Book of Abstracts

Contents

Scientific Linux Status Report and Plenary Discussion	1
Tools and techniques for managing virtual machine images	1
Network Information and Monitoring Infrastructure (NIMI)	1
A security survey of the e-mail services at INFN	2
Virtualization for high availability	2
XEN and KVM in INFN production systems and a comparison between them	3
HA Cluster using Open Sharedroot	3
The OpenNebula Engine for on-Demand Resource Provisioning	3
StratusLab: Running a Grid Site in the Cloud	4
LCLS Online and Offline Computing	4
The CernVM Project	5
New HPSS architecture at SLAC	5
First experiences with dovecot at DESY	5
Operation of the Portuguese Tier2 (Storage Element component)	6
CERN Lustre Evaluation	6
Simple Linux Utility for Resource Management	7
Fermi Linux STS - SL 6	7
SMS based OTP system for SSH logins	7
The Academic Cloud - Virtualized Worker Nodes in The Grid	7
News from the Storage Working Group	8
AFS/OSD: massive production experience and work in progress	8
Lustre over WAN	9
Performance of 10G network cards in file servers for LHC computing use cases	9

Some SSD facts	9
iSCSI at CERN	9
CERN site report	9
Benchmarking up-to-date x86 processors	10
WiFi security appliance for authentication solution during conferences and seminars	10
LAL and GRIF Site Report	10
Next generation of virtual infrastructure with Hyper-V	10
Site Report PSI (CH)	11
NDGF Site Report	11
Virtual Machine CPU Benchmarking the HEPiX way	11
RAL Site Report	12
An Update on the new Computer Building at RAL	12
Integrating Quattor and virtualisation technologies	12
Cybersecurity Update	12
Virtualization and/vs security	13
Benchmarking current PC server hardware	13
Packaging Grid software for Linux distributions	13
GridKa Site Report	14
CINBAD - CERN Investigation of Network Behaviour Anomaly Detection	14
Playing with Puppets instead of managing your systems	15
Virtualization on the Track I	15
TRIUMF Site Report	15
A Highly Versatile Virtual Data Center Ressource Pool	15
SLAC Site Report	16
DESY site report	16
HPC2N Data Center tour groups A + B	16
HPC2N Data Center tour groups E + F	16
HPC2N Data Center tour groups C & D	17
A new Data Centre for CC-IN2P3	17
The CNAF Data Centre migration	17

Impact of Filesystems on Application Performance in an HNEP Environment.	17
State of Nordic R&E Networking	18
Cloud Security	18
Opening remarks	18
SouthGrid and Oxford Site Report	18
Virtualization on the Track II	18
INFN T1 (NAF)	19
ASGC (Taipei)	19
NERSC	19
Wrap-up & Conclusions	19
Virtualisation discussion (continued)	19

Operating Systems & Applications II / 0**Scientific Linux Status Report and Plenary Discussion****Author:** Troy Dawson¹¹ *FERMILAB***Corresponding Author:** dawson@fnal.gov

Progress of Scientific Linux over the past 6 months. What we are currently working on. What we see in the future for Scientific Linux.

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

Virtualisation I / 1**Tools and techniques for managing virtual machine images****Author:** Andreas Unterkircher¹¹ *CERN***Corresponding Author:** andreas.unterkircher@cern.ch

Virtual machines can have many different deployment scenarios and therefore may require generation of multiple VM images. We report on work that has been done in a collaboration between CERN's Grid Deployment group and openlab to address several issues with image generation. Libs-fimage is a standalone application which generates VM images with a rich selection of Linux distribution. OSFarm provides a user interface to libs-fimage. In order to optimize generation of images, a layered copy-on-write image structure is used and an image cache ensures that identical images are not regenerated. For distributing images, content-based transfer has been investigated.

Miscellaneous Talks II / 2**Network Information and Monitoring Infrastructure (NIMI)****Author:** Igor Mandrichenko¹**Co-authors:** Donald Petravick¹; James Fromm¹; Joseph Klemencic¹; Mark Leininger¹; Randolph Reitz¹; Vladimir Bravov¹¹ *FNAL***Corresponding Authors:** dawson@fnal.gov, ivm@fnal.gov

Fermilab is a high energy physics research lab that maintains a highly dynamic network which typically supports around 15,000 active nodes. Due to the open nature of the scientific research conducted at FNAL, the portion of the network used to support open scientific research requires high bandwidth connectivity to numerous collaborating institutions around the world, and must facilitate convenient access by scientists at those institutions. Network Information and Monitoring Infrastructure (NIMI) is a framework built to help network management personnel and the computer security team monitor and manage the FNAL network. This includes the portions of the network used to support open scientific research as well as the portions for more tightly controlled administrative and

scientific support activities. As an infrastructure, NIMI has been used to build such applications as Node Directory, Network Inventory Database and Computer Security Issue Tracking System (TIssue). These applications have been successfully used by FNAL Computing Division personnel to manage local network, maintain necessary level of protection of LAN participants against external threats and promptly respond to computer security incidents. The article will discuss NIMI structure, functionality of major NIMI-based applications, history of the project, its current status and future plans.

Security and Networking I / 3

A security survey of the e-mail services at INFN

Author: Ombretta Pinazza¹

Co-authors: Alessia Spitaleri ¹; Diego Leanza ¹; Francesco Ferrera ¹; Franco Brasolin ¹; Fulvia Costa ¹; Roberto Cecchini ¹

¹ INFN

The INFN computing board has recently appointed a working group to investigate the status of the IT services provided in the different laboratories and divisions, with particular aim at security.

It's a legal requirement in Italy that public organizations provide monitoring and auditing procedures to protect users and personnel data.

For the e-mail service, the Security and Mailing working groups have collaborated to organize a survey and report to local administrators on vulnerabilities, misconfigurations and weaknesses in their systems.

The task has turned out as an opportunity to improve the technological knowledge of the system administrators as well as the quality and the strengthness of the service.

This work will describe strategies, tools and techniques and the main results.

Virtualisation I / 4

Virtualization for high availability

Author: Federico Calzolari¹

Co-authors: Alberto Ciampa ²; Enrico Mazzoni ²; Silvia Arezzini ²

¹ Scuola Normale Superiore - INFN Pisa

² INFN Pisa

A high availability service is one of the main problems for a data center. Till now high availability was achieved by host per host redundancy, a highly expensive method in terms of hardware and human costs. A new approach to the problem can be offered by virtualization.

Using virtualization, it is possible to achieve a redundancy system for all the services running on a data center. This new approach to high availability allows to share the running virtual machines over the servers up and running, by exploiting the features of the virtualization layer: start, stop and move virtual machines between physical hosts.

The system (3RC) is based on a finite state machine with hysteresis, providing the possibility to restart each virtual machine over any physical host, or reinstall it from scratch. A complete infrastructure has been developed to install operating system and middleware in a few minutes. To virtualize the main servers of a data center, a new procedure has been developed to migrate physical to virtual hosts.

The whole Grid data center SNS-PISA is running at the moment in virtual environment under the high availability system. As extension of the 3RC architecture, several storage solutions have been tested to store and centralize all the virtual disks, from NAS to SAN, to grant data safety and access

from everywhere.

Exploiting virtualization and ability to automatically reinstall a host, we provide a sort of host on-demand, where the action on a virtual machine is performed only when a disaster occurs.

Virtualisation I / 5

XEN and KVM in INFN production systems and a comparison between them

Author: Riccardo Veraldi¹

Co-author: Andrea Chierici²

¹ *INFN*

² *INFN-CNAF*

Corresponding Author: riccardo.veraldi@cnaf.infn.it

Virtualization allow multiple virtual machines to run on a bare bone single hardware.

There are many benefits in using Virtualization, fast disaster recovery, OS and software testing, maximization of hardware resources, server consolidation. INFN is using open source virtualization in different contexts mainly with Xen, but KVM is an interesting emerging technology ready for production systems. We will show how Xen is used inside INFN, and how KVM may be used as well, comparing the two different solutions.

Operating Systems & Applications II / 6

HA Cluster using Open Sharedroot

Author: Klaus Steinberger¹

¹ *LMU München*

Corresponding Author: klaus.steinberger@physik.uni-muenchen.de

At the LMU München, Faculty of Physics we currently deploy a storage and server cluster, which should serve home directories and other services to the whole faculty. The cluster member servers run from shared storage using Open Sharedroot as a single system image.

Virtualisation IV / 7

The OpenNebula Engine for on-Demand Resource Provisioning

Authors: Ignacio Martin Llorente¹; Rubén Santiago Montero¹

¹ *DSA-Research, Universidad Complutense de Madrid*

Corresponding Author: llorente@dacya.ucm.es

OpenNebula is an open source virtual infrastructure engine that enables the dynamic placement of VMs on a pool of physical resources. It provides a powerful and agile CLI and API for monitoring and controlling large scale VM deployments, including networking and image management, and a flexible and generic framework to define new policies for capacity provision. Additionally, OpenNebula provides plugins to access commercial clouds (Amazon EC2 and ElasticHosts) to supplement local resources with cloud resources to satisfy peak or fluctuating demands in the service workload. OpenNebula extends the benefits of virtualization platforms (hypervisors) from a single physical resource to a pool of resources, decoupling the server not only from the physical infrastructure but also from the physical location. In computing environments, such separation of resource provisioning from job execution management provides several benefits: (1) elastic site capacity, the capacity of the site can be modified by deploying (or shutting down) virtual worker nodes on an on-demand basis, either in local physical resources or in remote resources; (2) cluster partitioning, the physical resources of the site could be used to execute worker nodes bound to different virtual computing clusters, and thus isolating their workloads and partitioning the performance assigned to each virtual cluster; and (3) heterogeneous configurations, the virtual worker nodes of a virtual cluster can have multiple (even conflicting) software configurations.

Virtualisation IV / 8

StratusLab: Running a Grid Site in the Cloud

Authors: Cal Loomis¹; Michel Jouvin¹

¹ CNRS/LAL

Cloud technologies have matured quickly over the last couple of years and now provide an interesting platform on which to host grid services. The dynamic nature of these resources could ease life-cycle management for system administrators and could provide customized resources for users. However, questions remain about how these resources can meet the grid's security and operational policies. This presentation explains the challenges raised by using cloud resources for a EGEE grid site.

StratusLab (<http://www.stratuslab.org/wiki/>), an informal collaboration between CNRS/LAL, GRNET, SixSq Sàrl, and UCM, aims to determine how mature and robust cloud resources are by running a full grid site within the Amazon cloud. This will also show how compatible cloud resources are with grid technology and with standard system administrator management tools. The first steps have already uncovered both administrative and technical problems in using cloud resources with the grid. This presentation will describe those problems, the current state of the experiment, and the future directions of the collaboration.

Summary:

The StratusLab collaboration investigates the maturity and robustness of cloud resources by trying to run a full grid site in the Amazon cloud. The presentation describes the results from this experiment and future directions of the collaboration.

Miscellaneous Talks 0 / 9

LCLS Online and Offline Computing

Author: Alf Wachsmann¹

¹ SLAC

Corresponding Author: alfw@slac.stanford.edu

I will present the architectural design of the online and offline computing system for SLAC's LCLS X-ray laser facility. I will briefly show the DAQ system to illustrate the data flow from the various experiments to the system's boundary. Presenting the offline data management system will be the main focus of the talk.

Virtualisation II / 10

The CernVM Project

Author: Carlos Aguado Sanchez¹

Co-authors: Jakob Blomer¹; Predrag Buncic¹

¹ CERN

Corresponding Author: carlos.aguado.sanchez@cern.ch

The CERNVM project addresses two aspects of software distribution, namely platform compatibility and general usability. It proposes a new approach based on extensive use of virtualization and distribution of contents over the network. Thus, two main components are envisioned in this new paradigm: the CERNVM virtual machine and the CERNVM File System (CVMFS). The former, based on rPath Linux, leverages the position of virtualization as a new technology enabler letting the user choose his own preferred hardware platform. The latter builds an efficient and highly distributable Content Delivery Network on top of the well-known HTTP protocol. Using the CVMFS implements aggressive caching policies to enable disconnected mode of operation. This talk will be focused on the current implementation aspects of the underlying infrastructure and the choices for its essential building blocks. That is, computing virtualization (VMWare vs. XEN), storage virtualization (hardware based vs. software based), provisioning interfaces (Virtual Infrastructure) and content switching (switching L4 vs. L7). In addition, a brief comparison of performance between this approach and the existing ones under common scenarios will be outlined with some benchmarks.

Storage and Filesystems II / 11

New HPSS architecture at SLAC

Author: Alf Wachsmann¹

¹ SLAC

Corresponding Author: alfw@slac.stanford.edu

SLAC's HPSS hardware and software is reaching its end of support by 2010. We have purchased two Sun SL8500 robots to replace our 6 Powderhorns and we are in the process of migrating from HPSS version 5.1 to version 6.2. In the same project, we are changing our HPSS architecture to now include disk caches. I will outline the entire project, the new hardware and software architecture, and the current status.

Operating Systems & Applications / 12

First experiences with dovecot at DESY

Author: Wolfgang Friebel¹

Co-author: Dirk Jahnke-Zumbusch¹

¹ *Deutsches Elektronen-Synchrotron (DESY)*

Corresponding Author: wolfgang.friebel@desy.de

DESY recently planned to migrate some of its UNIX based IMAP servers to dovecot. The latest releases of dovecot (version 1.2) include support for several mailbox formats, quota handling, compressed folders, ACLs and allows for a comfortable configuration of namespaces. The additional sieve component and the managesieve protocol does enable users to configure server side filtering and managing sieve scripts. A web interface has been made available to let users compose and install syntactically correct filter rules without knowing the sieve language. One of the dovecot installations is in the pilot phase with several users, the other one is being set up currently

Site reports I / 13

Operation of the Portuguese Tier2 (Storage Element component)

Author: Mario David¹

Co-authors: Carlos Manuel¹; Goncalo Borges¹; J. Paulo Martins¹; Jorge Gomes¹; Miguel Oliveira¹

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

Corresponding Author: david@lip.pt

The Portuguese LCG federated Tier-2 supports the ATLAS and CMS experiments and is composed of three sites. Two of these sites (Lisbon and Coimbra) are already in production and the third one is currently being setup (LNEC/FCCN).

We will describe briefly the Tier2 and then concentrate on the Storage Element component. In the middle of 2008, it was decided to use the Storm SRM with SUN's Lustre has backend filesystem.

The report will focus on the operational issues and proceed with the monitoring, and performance of the system as part of the ATLAS and CMS experiments.

The integration of the Lustre specific sensors in the Nagios monitoring framework, as well as Ganglia metrics will be also described.

Storage and Filesystems II / 14

CERN Lustre Evaluation

Author: Arne Wiebalck¹

¹ *CERN*

Corresponding Author: arne.wiebalck@cern.ch

CERN is evaluating the cluster file system Lustre as a potential consolidated storage solution for project space, home directories, analysis space and HSM. Rather than on performance or on scalability, the main focus of the evaluation will be on operational questions. This talk will cover the various aspects to be looked at during the survey and report on the initial thoughts and findings as of the first months.

Operating Systems & Applications II / 15**Simple Linux Utility for Resource Management****Author:** Pär Andersson¹¹ *NSC, Linköpings universitet***Corresponding Author:** paran@nsc.liu.se

We have been using SLURM on new clusters since 2007. I will give an introduction to SLURM and our experiences with it. Features in the upcoming 1.4 release will be discussed, and also how SLURM can be used with grid software.

Operating Systems & Applications / 16**Fermi Linux STS - SL 6****Author:** Troy Dawson¹¹ *FERMILAB***Corresponding Author:** dawson@fnal.gov

Fermi Linux Short Term Support
This project helped fix a problem at Fermilab for bleeding edge hardware. It also helped the Scientific Linux development team look at what might be possible with Scientific Linux 6.
This presentation will show what we were able to do with Fermi Linux STS, and give some insight into what might be coming with Scientific Linux 6.

Miscellaneous Talks I / 18**SMS based OTP system for SSH logins****Author:** Jonas Lindemann¹¹ *Lund University***Corresponding Author:** jonas.lindemann@lunarc.lu.se

To strengthen the security at Lunarc, SMS based OTP system has been implemented for SSH logins. The system is currently used in production and can be used by both Linux, Windows and Mac OS X based clients. The solution is based on a commercial available java-based OTP server.

Virtualisation III / 19**The Academic Cloud - Virtualized Worker Nodes in The Grid****Author:** Owen Synge¹

¹ DESY

Corresponding Author: owen.synge@cern.ch

Cloud computing provides commodity Virtualized computers on demand, Grid computing has underlying batch queues which may be integrated with the Virtualization technology can give the same predictable run time environment we know of from the Cloud Computing world.

Two groups have achieved Virtualized worker nodes using two different batch queues. Details of the two implementations using similar methods will be presented. While the technology is simple and it's flexibility maybe attractive for grid end users, administrators have justified security concerns that may limit real world expectations.

Summary:

Admins can implement "The academic cloud - Virtualized Worker Nodes in The Grid" easily, this talk shows how two simple solutions work.

Storage and Filesystems I / 20

News from the Storage Working Group

Author: Andrei Maslennikov¹

¹ CASPUR

Storage and Filesystems I / 22

AFS/OSD: massive production experience and work in progress

Author: Hartmut Reuter¹

¹ Rechenzentrum Garching

AFS/OSD was presented two years ago on the HEPiX Spring 2007 in Hamburg. At this time the results of the R&D project sponsored by CERN and ENEA were presented which showed that performance of AFS/OSD scales linearly with the number of OSDs used. For RZG, however, the goal was to replace MR-AFS which for the last 12 years had offered HSM-features to AFS. Therefore the concept of "archival OSDs" and "wipeable OSDs" has been introduced. Files in OSDs automatically get copies on archival OSDs with the double purpose to protect the files against loss of a disk system and to allow wiping of files in order to free disk space. In 2007 these features have been added to AFS/OSD and at the begin of 2008 MR-AFS has been replaced by AFS/OSD in place (without moving data on tapes). In Garching TSM-HSM is used as an underlying HSM system, but any filesystem based HSM system can be used. In cooperation with DESY also an interface to dCache/Chimera has been developed. The AFS cell at RZG today contains nearly 300 TB 80% of which are stored in OSDs. Felix Frank from DESY added the policy support which allows to specify for a volume or for a directory which files should go into OSDs, should be striped over multiple OSDs, or get copies in multiple OSDs. Policies can be based on file name patterns or size.

Additional work is in progress to use cluster filesystems such as Lustre or GPFS for AFS. This technique offers fastest access to data stored in AFS inside the cluster while world wide access with normal AFS speed is still possible. It also adds HSM features Lustre.

Storage and Filesystems I / 23**Lustre over WAN**

Author: Pierre-Francois Honore¹

¹ *CEA*

Storage and Filesystems I / 24**Performance of 10G network cards in file servers for LHC computing use cases**

Author: Artem Trunov¹

¹ *FZK*

Storage and Filesystems II / 25**Some SSD facts**

Author: Tim Polland¹

¹ *Texas Memory Systems / Takan*

Storage and Filesystems I / 27**iSCSI at CERN**

Author: Andras Horvath¹

¹ *CERN*

We report on the ongoing evaluation of the iSCSI technology regarding stability, manageability on the large scale and performance.

Site reports I / 28**CERN site report**

Author: Helge Meinhard¹

¹ *CERN-IT*

Corresponding Author: helge.meinhard@cern.ch

Summary of changes at CERN since the last meeting

Miscellaneous Talks 0 / 29**Benchmarking up-to-date x86 processors**

Author: Helge Meinhard¹

¹ *CERN-IT*

Corresponding Author: helge.meinhard@cern.ch

Reports on benchmarks (both performance and power) of Nehalem and Shanghai processors

Security and Networking I / 30**WiFi security appliance for authentication solution during conferences and seminars**

Author: Riccardo Veraldi¹

¹ *INFN*

Corresponding Author: riccardo.veraldi@cnaf.infn.it

It's possible to build a security appliance to set up a custom WiFi authentication solution based both on 802.1x and Captive Portal. I developed a solution inside INFN which allows users during Workshops or seminars to authenticate a WiFi session using 802.1x or Captive portal authentication. The system is a portable device, it is a small Soekris device with 4 network ports. It implements VLAN tagging and Access Points and can be attached directly to the soekris ports. All the security appliance implementation is based on customized OpenBSD 4.5 distribution. Configuration allows multiple VLANs and multiple SSIDs. Since the authentication is based on radius server proxying, it can easily be integrated into EDUROAM

Site reports I / 31**LAL and GRIF Site Report**

Author: Michel Jouvin¹

¹ *LAL / IN2P3*

Corresponding Author: jouvin@lal.in2p3.fr

LAL and GRIF Site report

Virtualisation II / 32**Next generation of virtual infrastructure with Hyper-V**

Author: Michal Kwiatek¹

¹ CERN

Corresponding Author: michal.kwiatek@cern.ch

One of the important benefits of virtualisation is the increased flexibility of hardware provisioning. Internet Services group provides the hardware and OS layers for services operated by different CERN organizational units. Within this framework, a user can ask for a new virtual server from a web browser and have it ready within 15 minutes. The latest release of the virtual infrastructure is based on Hyper-V, Microsoft Virtual Machine Manager, management SOAP web services and user web interface. This talk will present this approach and will discuss the evolution of the Microsoft Hypervisor from the perspective of both the end-user and the system administrator.

Site reports I / 34

Site Report PSI (CH)

Author: Marc Gasser¹

¹ PSI

Corresponding Author: marc.gasser@psi.ch

Scientific Linux 5
Configuration Management with Puppet
New Spam Filter

Site reports I / 35

NDGF Site Report

Author: Mattias Wadenstein¹

¹ NDGF

Corresponding Author: mattias.wadenstein@cern.ch

Recent development and news from the NDGF sphere.

Virtualisation III / 36

Virtual Machine CPU Benchmarking the HEPiX way

Authors: Ian Gable¹; Manfred Alef²

¹ University of Victoria

² FZK Karlsruhe / GridKa

Corresponding Author: igable@uvic.ca

We employ the HEP-SPEC06 benchmark, developed by the HEPiX CPU benchmarking working group, to evaluate CPU performance of a number of virtual machines configurations for the highly CPU loaded HEP worker node. Benchmarks are performed on 8 different models of AMD and Intel CPU spanning 2003 to 2008 generation architectures. We demonstrate that multi-core worker node

can run n VMs where n is the number of cores without suffering significant CPU performance penalties. Our focus is primarily on Xen, but we have some preliminary KVM results.

Summary:

We employ the HEP-SPEC06 benchmark, developed by the HEPiX CPU benchmarking working group, to evaluate CPU performance of a number of virtual machines configurations for the highly CPU loaded HEP worker node.

Site reports I / 37

RAL Site Report

Author: Martin Bly¹

¹ *STFC-RAL*

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

Update on activities at the RAL Tier1

Data centres / 38

An Update on the new Computer Building at RAL

Author: Martin Bly¹

¹ *STFC-RAL*

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

An update on progress towards occupation of the new computer building at RAL, with notes on the building completion, new installations of hardware, and migration of the RAL Tier1 to the new building.

Virtualisation II / 39

Integrating Quattor and virtualisation technologies

Author: Ewan Roche¹

¹ *CERN*

Corresponding Author: ewan.roche@cern.ch

The integration of virtualisation technologies and the Quattor system management toolkit raises a number of challenges. We present the approaches taken so far with their limitations and successes and discuss current projects for development.

Security and Networking I / 40**Cybersecurity Update**

Author: Jan Iven¹

¹ *CERN*

Corresponding Author: jan.iven@cern.ch

An update on recent computer security issues and vulnerabilities affecting Windows, Linux and Mac platforms. This talk is based on contributions and input from a range of colleagues both within and outside CERN. It covers clients, servers and control systems.

Virtualisation II / 41**Virtualization and/vs security**

Author: Jan Iven¹

¹ *CERN*

Corresponding Author: jan.iven@cern.ch

Virtualization technology is generating a strong interest from the HEP community. The talk will look at some of the security issues around this - both newly-added concerns as well as potential benefits for using virtualization.

Miscellaneous Talks I / 42**Benchmarking current PC server hardware**

Author: Götz Waschk¹

Co-authors: Dirk Pleiter¹; Frank Winter¹; Hubert Simma¹; Peter Wegner¹

¹ *DESY*

Corresponding Author: goetz.waschk@desy.de

HEP-SPEC06 is the standard measurement for computing power in the LCG community. The computing requirements of groups like Lattice QCD at DESY need specialized benchmarks to evaluate new hardware. The theoretical particle physics group uses the form benchmark, which solves symbolical equations. The DD-HMC and the Chroma benchmarks include kernels of high-performance parallel Lattice QCD applications. Both parallel benchmarks were run on the up to 8 cores of a compute node. The three mentioned benchmarks were used to evaluate the current server hardware with the latest CPUs by Intel and AMD.

Operating Systems & Applications II / 43

Packaging Grid software for Linux distributions

Author: Mattias Ellert¹

Co-authors: Anders Wäänänen²; Steffen Möller³

¹ *Uppsala Universitet & NDGF*

² *Niels Bohr Institute*

³ *Universität zu Lübeck*

Corresponding Author: mattias.ellert@fysast.uu.se

High energy physics experiments today utilise computing grids to access computing resources in order to fulfil their needs for processing power and storage of their ever increasing datasets. However, grid middleware so far has not been part of the mainstream Linux distributions used by resource providers and users. The installation and maintenance of the grid middleware therefore imposes additional burdens on both computing centres and users wishing to participate in the processing of the experimental data. In this talk we will present our efforts to bring the Globus toolkit, a basic building block in many grid middlewares, into the Debian and Fedora Linux distributions.

Site reports II / 44

GridKa Site Report

Author: Manfred Alef¹

¹ *Forschungszentrum Karlsruhe*

Corresponding Author: manfred.alef@iwr.fzk.de

Current status of GridKa, e.g.
- new hardware
- problem discussion

Security and Networking II / 45

CINBAD - CERN Investigation of Network Behaviour Anomaly Detection

Authors: Milosz Marian Hulboj¹; Ryszard Erazm Jurga¹

¹ *CERN*

Corresponding Author: ryszard.jurga@cern.ch

The CINBAD (CERN Investigation of Network Behaviour and Anomaly Detection) project was launched in 2007 in collaboration with ProCurve Networking by HP. The project mission is to understand the behaviour of large computer networks (10,000 or more nodes) in the context of high performance computing and large campus installations such as CERN, whose network counts today roughly 70,000 Gigabit user ports. The goals of the project are to be able to detect traffic anomalies in such systems, perform trend analysis, automatically take counter measures and provide post-mortem analysis facilities. This talk will present the main project principles, data sources, data collection and analysis approaches as well as the initial findings.

Operating Systems & Applications / 46**Playing with Puppets instead of managing your systems****Author:** Thomas Bellman¹¹ *National Supercomputer Centre, Sweden***Corresponding Author:** bellman@nsc.liu.se

As a way of automating and structuring the administration of our clusters, we have started using Puppet, which is a system for configuring and administering Unix machines, similar to Cfengine. We here present our experiences, and show some of Puppet's strengths and weaknesses.

Virtualisation I / 47**Virtualization on the Track I****Author:** Thomas Finnern¹¹ *DESY***Corresponding Author:** thomas.finnern@desy.de

A brief introduction to virtualization.

Summary:

Virtualization becomes more and more a core technology for today's computing. A full day of talks on this topic should give information about the variety of implementations and use cases, showing the benefits, the limits and the trends of this technology.

Site reports II / 48**TRIUMF Site Report****Author:** Steven McDonald¹¹ *TRIUMF***Corresponding Author:** mcdonald@triumf.ca

A summary of TRIUMF Site activities

Virtualisation I / 49**A Highly Versatile Virtual Data Center Ressource Pool****Author:** Stefan Bujack¹¹ *Deutsches Elektronen-Synchrotron DESY*

Corresponding Author: stefan.bujack@desy.de

Different virtualization techniques are in use at the DESY data centers, with Xen servers being the dominant hypervisor. Besides its open source implementations, DESY, location Hamburg, has set up a pool of DELL servers with NetApp storage and Xen Enterprise software to provide a highly available, scalable, and versatile environment for managing data center services on various Windows and Linux flavoured operating systems. The system currently supplies 90 paravirtualized servers on 18 dual CPU servers.

We will provide an overview of the implementation, its integration with installation services and storage systems, use of features like HA, Snapshotting, Ontap connection, and central administration.

We will present advantages and limitations of our approach and our recent upgrades to the pool.

Site reports II / 50

SLAC Site Report

Author: Alf Wachsmann¹

¹ SLAC

Corresponding Author: alfw@slac.stanford.edu

News from SLAC

Site reports III / 51

DESY site report

Author: Christoph Beyer¹

¹ DESY HH

Corresponding Author: christoph.beyer@desy.de

DESY site report

Miscellaneous Talks I / 52

HPC2N Data Center tour groups A + B

Tour for groups A and B of the HPC2N facilities.

Group registration on papers at the registration desk.

Those not in groups A or B have this slot free.

Miscellaneous Talks II / 54**HPC2N Data Center tour groups E + F**

Tour of the HPC2N facilities, groups E and F.

Group registration on papers at the registration desk.

Those not in groups E or F have this slot free.

Data centres / 55**HPC2N Data Center tour groups C & D****Data centres / 56****A new Data Centre for CC-IN2P3**

Author: Dominique Boutigny¹

¹ *CC-IN2P3*

57

The CNAF Data Centre migration

Author: Felice Rosso¹

¹ *INFN*

Miscellaneous Talks II / 58**Impact of Filesystems on Application Performance in an HNEP Environment.**

Author: Iwona Sakrejda¹

¹ *LBNL/NERSC*

Corresponding Author: isakrejda@lbl.gov

Efficiency of data intensive applications heavily depends on the performance of storage. A comparison of different solutions ranging from GPFS filesystems with varying configuration through serving files from dedicated sets of hosts and storage sharing resources with computing will be shown. Measurements are being done with a mix of applications characteristic for HENP. Total throughput is being used to evaluate results.

Impact of aging hardware and reality of coexisting generations of computing and storage will be discussed as well..

Security and Networking II / 59**State of Nordic R&E Networking**

Overview and state of the current and future Nordic Research and Education Networking, from Nordunet.

Virtualisation III / 60**Cloud Security**

Author: David Kelsey¹

¹ *RAL*

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

This talk will cover the various security implications of the use of Cloud resources attached to the HEP Grids. There are several issues related to trust, policy and operational security in addition to the more general security issues of virtualisation.

Opening remarks / 61**Opening remarks**

Author: Mattias Wadenstein¹

¹ *NDGF*

Opening remarks by Prof. Bo Kågström, director HPC2N.

Site reports III / 62**SouthGrid and Oxford Site Report**

Author: Peter Gronbech¹

¹ *Nuclear Physics Laboratory*

Corresponding Author: p.gronbech1@physics.ox.ac.uk

SouthGrid and Oxford Status report

Virtualisation IV / 63**Virtualization on the Track II**

Author: Thomas Finnern¹

¹ DESY

Corresponding Author: thomas.finnern@desy.de

A short plenary discussion.

Summary:

15 minutes to collect opinions and comments from the audience.

Site reports II / 65

INFN T1 (NAF)

Site reports III / 66

ASGC (Taipe)

Author: Jason Shih¹

¹ ASGC

Corresponding Author: hlshih@gate.sinica.edu.tw

Site reports III / 67

NERSC

Wrap-Up & Conclusions / 68

Wrap-up & Conclusions

Wrap-up & Conclusions

Data centres / 69

Virtualisation discussion (continued)