

HEPiX Spring 2009

Report of Contributions

Contribution ID: 0

Type: **not specified**

Scientific Linux Status Report and Plenary Discussion

Thursday, May 28, 2009 3:30 PM (30 minutes)

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.

Primary author: Mr DAWSON, Troy (FERMILAB)

Presenter: Mr DAWSON, Troy (FERMILAB)

Session Classification: Operating Systems & Applications II

Track Classification: Operating Systems & Applications

Contribution ID: 1

Type: **not specified**

Tools and techniques for managing virtual machine images

Wednesday, May 27, 2009 2:30 PM (30 minutes)

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. Libsimage is a standalone application which generates VM images with a rich selection of Linux distribution. OSFarm provides a user interface to libsimage. 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.

Primary author: Mr UNTERKIRCHER, Andreas (CERN)

Presenter: Mr UNTERKIRCHER, Andreas (CERN)

Session Classification: Virtualisation I

Track Classification: Virtualisation

Contribution ID: 2

Type: **not specified**

Network Information and Monitoring Infrastructure (NIMI)

Friday, May 29, 2009 9:30 AM (30 minutes)

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.

Primary author: Mr MANDRICHENKO, Igor (FNAL)

Co-authors: PETRAVICK, Donald (FNAL); FROMM, James (FNAL); KLEMENCIC, Joseph (FNAL); LEININGER, Mark (FNAL); REITZ, Randolph (FNAL); BRAVOV, Vladimir (FNAL)

Presenter: Mr DAWSON, Troy (FERMILAB)

Session Classification: Miscellaneous Talks II

Track Classification: Security and Networking

Contribution ID: 3

Type: **not specified**

A security survey of the e-mail services at INFN

Tuesday, May 26, 2009 10:30 AM (30 minutes)

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.

Primary author: PINAZZA, Ombretta (INFN)

Co-authors: SPITALERI, Alessia (INFN); LEANZA, Diego (INFN); FERRERA, Francesco (INFN); BRASOLIN, Franco (INFN); COSTA, Fulvia (INFN); CECCHINI, Roberto (INFN)

Presenter: PINAZZA, Ombretta (INFN)

Session Classification: Security and Networking I

Track Classification: Security and Networking

Contribution ID: 4

Type: **not specified**

Virtualization for high availability

Wednesday, May 27, 2009 1:40 PM (25 minutes)

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.

Primary author: Dr CALZOLARI, Federico (Scuola Normale Superiore - INFN Pisa)

Co-authors: Dr CIAMPA, Alberto (INFN Pisa); Dr MAZZONI, Enrico (INFN Pisa); Dr AREZZINI, Silvia (INFN Pisa)

Presenter: Dr CALZOLARI, Federico (Scuola Normale Superiore - INFN Pisa)

Session Classification: Virtualisation I

Track Classification: Virtualisation

Contribution ID: 5

Type: **not specified**

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

Wednesday, May 27, 2009 2:05 PM (25 minutes)

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.

Primary author: Dr VERALDI, Riccardo (INFN)

Co-author: CHIERICI, Andrea (INFN-CNAF)

Presenter: Dr VERALDI, Riccardo (INFN)

Session Classification: Virtualisation I

Track Classification: Virtualisation

Contribution ID: 6

Type: **not specified**

HA Cluster using Open Sharedroot

Thursday, May 28, 2009 4:00 PM (30 minutes)

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

Primary author: STEINBERGER, Klaus (LMU München)

Presenter: STEINBERGER, Klaus (LMU München)

Session Classification: Operating Systems & Applications II

Track Classification: Operating Systems & Applications

Contribution ID: 7

Type: **not specified**

The OpenNebula Engine for on-Demand Resource Provisioning

Thursday, May 28, 2009 11:00 AM (45 minutes)

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.

Primary authors: Prof. MARTIN LLORENTE, Ignacio (DSA-Research, Universidad Complutense de Madrid); Prof. SANTIAGO MONTERO, Rubén (DSA-Research, Universidad Complutense de Madrid)

Presenter: FONTAN, Javier

Session Classification: Virtualisation IV

Track Classification: Virtualisation

Contribution ID: 8

Type: **not specified**

StratusLab: Running a Grid Site in the Cloud

Thursday, May 28, 2009 10:30 AM (30 minutes)

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.

Primary authors: LOOMIS, Cal (CNRS/LAL); JOUVIN, Michel (CNRS/LAL)

Presenter: JOUVIN, Michel (CNRS/LAL)

Session Classification: Virtualisation IV

Track Classification: Virtualisation

Contribution ID: 9

Type: **not specified**

LCLS Online and Offline Computing

Tuesday, May 26, 2009 2:00 PM (30 minutes)

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.

Primary author: WACHSMANN, Alf (SLAC)

Presenter: WACHSMANN, Alf (SLAC)

Session Classification: Miscellaneous Talks 0

Contribution ID: 10

Type: **not specified**

The CernVM Project

Wednesday, May 27, 2009 4:00 PM (30 minutes)

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.

Primary author: Mr AGUADO SANCHEZ, Carlos (CERN)

Co-authors: Mr BLOMER, Jakob (CERN); Mr BUNCIC, Predrag (CERN)

Presenter: Mr AGUADO SANCHEZ, Carlos (CERN)

Session Classification: Virtualisation II

Track Classification: Virtualisation

Contribution ID: 11

Type: **not specified**

New HPSS architecture at SLAC

Wednesday, May 27, 2009 11:30 AM (20 minutes)

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.

Primary author: WACHSMANN, Alf (SLAC)

Presenter: WACHSMANN, Alf (SLAC)

Session Classification: Storage and Filesystems II

Track Classification: Storage

Contribution ID: 12

Type: **not specified**

First experiences with dovecot at DESY

Monday, May 25, 2009 11:30 AM (30 minutes)

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

Primary author: FRIEBEL, Wolfgang (Deutsches Elektronen-Synchrotron (DESY))

Co-author: JAHNKE-ZUMBUSCH, Dirk (Deutsches Elektronen-Synchrotron (DESY))

Presenter: FRIEBEL, Wolfgang (Deutsches Elektronen-Synchrotron (DESY))

Session Classification: Operating Systems & Applications

Track Classification: Operating Systems & Applications

Contribution ID: 13

Type: **not specified**

Operation of the Portuguese Tier2 (Storage Element component)

Monday, May 25, 2009 1:00 PM (20 minutes)

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.

Primary author: DAVID, Mario (LIP Laboratorio de Instrumentaco e Fisica Experimental de Particulas)

Co-authors: MANUEL, Carlos (LIP Laboratorio de Instrumentaco e Fisica Experimental de Particulas); BORGES, Goncalo (LIP Laboratorio de Instrumentaco e Fisica Experimental de Particulas); MARTINS, J. Paulo (LIP Laboratorio de Instrumentaco e Fisica Experimental de Particulas); GOMES, Jorge (LIP Laboratorio de Instrumentaco e Fisica Experimental de Particulas); OLIVEIRA, Miguel (LIP Laboratorio de Instrumentaco e Fisica Experimental de Particulas)

Presenter: DAVID, Mario (LIP Laboratorio de Instrumentaco e Fisica Experimental de Particulas)

Session Classification: Site reports I

Track Classification: Site reports

Contribution ID: 14

Type: **not specified**

CERN Lustre Evaluation

Wednesday, May 27, 2009 11:00 AM (30 minutes)

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.

Primary author: WIEBALCK, Arne (CERN)

Presenter: WIEBALCK, Arne (CERN)

Session Classification: Storage and Filesystems II

Track Classification: Storage

Contribution ID: 15

Type: **not specified**

Simple Linux Utility for Resource Management

Thursday, May 28, 2009 5:00 PM (30 minutes)

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.

Primary author: ANDERSSON, Pär (NSC, Linköpings universitet)

Presenter: ANDERSSON, Pär (NSC, Linköpings universitet)

Session Classification: Operating Systems & Applications II

Track Classification: Operating Systems & Applications

Contribution ID: 16

Type: **not specified**

Fermi Linux STS - SL 6

Monday, May 25, 2009 10:30 AM (30 minutes)

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.

Primary author: Mr DAWSON, Troy (FERMILAB)

Presenter: Mr DAWSON, Troy (FERMILAB)

Session Classification: Operating Systems & Applications

Track Classification: Operating Systems & Applications

Contribution ID: **18**

Type: **not specified**

SMS based OTP system for SSH logins

Tuesday, May 26, 2009 5:00 PM (30 minutes)

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.

Primary author: Dr LINDEMANN, Jonas (Lund University)

Presenter: Dr LINDEMANN, Jonas (Lund University)

Session Classification: Miscellaneous Talks I

Track Classification: Security and Networking

Contribution ID: 19

Type: **not specified**

The Academic Cloud - Virtualized Worker Nodes in The Grid

Thursday, May 28, 2009 9:00 AM (30 minutes)

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.

Primary author: SYNGE, Owen (DESY)

Presenter: SYNGE, Owen (DESY)

Session Classification: Virtualisation III

Track Classification: Virtualisation

Contribution ID: 20

Type: **not specified**

News from the Storage Working Group

Wednesday, May 27, 2009 8:00 AM (30 minutes)

Primary author: MASLENNIKOV, Andrei (CASPUR)

Presenter: MASLENNIKOV, Andrei (CASPUR)

Session Classification: Storage and Filesystems I

Contribution ID: 22

Type: **not specified**

AFS/OSD: massive production experience and work in progress

Wednesday, May 27, 2009 8:30 AM (30 minutes)

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.

Primary author: REUTER, Hartmut (Rechenzentrum Garching)

Presenter: REUTER, Hartmut (Rechenzentrum Garching)

Session Classification: Storage and Filesystems I

Track Classification: Storage

Contribution ID: 23

Type: **not specified**

Lustre over WAN

Wednesday, May 27, 2009 9:00 AM (20 minutes)

Primary author: HONORE, Pierre-Francois (CEA)

Presenter: HONORE, Pierre-Francois (CEA)

Session Classification: Storage and Filesystems I

Contribution ID: 24

Type: **not specified**

Performance of 10G network cards in file servers for LHC computing use cases

Wednesday, May 27, 2009 9:20 AM (20 minutes)

Primary author: TRUNOV, Artem (FZK)

Presenter: TRUNOV, Artem (FZK)

Session Classification: Storage and Filesystems I

Contribution ID: 25

Type: **not specified**

Some SSD facts

Wednesday, May 27, 2009 10:30 AM (30 minutes)

Primary author: POLLAND, Tim (Texas Memory Systems / Takan)

Presenter: POLLAND, Tim (Texas Memory Systems / Takan)

Session Classification: Storage and Filesystems II

Contribution ID: 27

Type: **not specified**

iSCSI at CERN

Wednesday, May 27, 2009 9:40 AM (20 minutes)

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

Primary author: HORVATH, Andras (CERN)

Presenter: HORVATH, Andras (CERN)

Session Classification: Storage and Filesystems I

Contribution ID: **28**

Type: **not specified**

CERN site report

Monday, May 25, 2009 1:20 PM (20 minutes)

Summary of changes at CERN since the last meeting

Primary author: Dr MEINHARD, Helge (CERN-IT)

Presenter: Dr MEINHARD, Helge (CERN-IT)

Session Classification: Site reports I

Track Classification: Site reports

Contribution ID: 29

Type: **not specified**

Benchmarking up-to-date x86 processors

Tuesday, May 26, 2009 2:30 PM (30 minutes)

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

Primary author: Dr MEINHARD, Helge (CERN-IT)

Presenter: Dr MEINHARD, Helge (CERN-IT)

Session Classification: Miscellaneous Talks 0

Track Classification: Misc Talks

Contribution ID: 30

Type: **not specified**

WiFi security appliance for authentication solution during conferences and seminars

Tuesday, May 26, 2009 11:00 AM (30 minutes)

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

Primary author: Dr VERALDI, Riccardo (INFN)

Presenter: Dr VERALDI, Riccardo (INFN)

Session Classification: Security and Networking I

Track Classification: Security and Networking

Contribution ID: **31**

Type: **not specified**

LAL and GRIF Site Report

Monday, May 25, 2009 2:20 PM (20 minutes)

LAL and GRIF Site report

Primary author: JOUVIN, Michel (LAL / IN2P3)

Presenter: JOUVIN, Michel (LAL / IN2P3)

Session Classification: Site reports I

Track Classification: Site reports

Contribution ID: 32

Type: **not specified**

Next generation of virtual infrastructure with Hyper-V

Wednesday, May 27, 2009 3:30 PM (30 minutes)

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.

Primary author: KWIATEK, Michal (CERN)

Presenter: KWIATEK, Michal (CERN)

Session Classification: Virtualisation II

Track Classification: Virtualisation

Contribution ID: **34**

Type: **not specified**

Site Report PSI (CH)

Monday, May 25, 2009 1:40 PM (20 minutes)

Scientific Linux 5
Configuration Management with Puppet
New Spam Filter

Primary author: GASSER, Marc (PSI)

Presenter: GASSER, Marc (PSI)

Session Classification: Site reports I

Track Classification: Site reports

Contribution ID: 35

Type: **not specified**

NDGF Site Report

Monday, May 25, 2009 2:00 PM (20 minutes)

Recent development and news from the NDGF sphere.

Primary author: WADENSTEIN, Mattias (NDGF)

Presenter: WADENSTEIN, Mattias (NDGF)

Session Classification: Site reports I

Track Classification: Site reports

Contribution ID: 36

Type: **not specified**

Virtual Machine CPU Benchmarking the HEPiX way

Thursday, May 28, 2009 8:30 AM (30 minutes)

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.

Primary authors: GABLE, Ian (University of Victoria); ALEF, Manfred (FZK Karlsruhe / GridKa)

Presenter: GABLE, Ian (University of Victoria)

Session Classification: Virtualisation III

Track Classification: Virtualisation

Contribution ID: 37

Type: **not specified**

RAL Site Report

Monday, May 25, 2009 2:40 PM (20 minutes)

Update on activities at the RAL Tier1

Primary author: Mr BLY, Martin (STFC-RAL)

Presenter: Mr BLY, Martin (STFC-RAL)

Session Classification: Site reports I

Track Classification: Site reports

Contribution ID: **38**

Type: **not specified**

An Update on the new Computer Building at RAL

Thursday, May 28, 2009 1:30 PM (30 minutes)

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.

Primary author: Mr BLY, Martin (STFC-RAL)

Presenter: Mr BLY, Martin (STFC-RAL)

Session Classification: Data centres

Track Classification: Data centres

Contribution ID: 39

Type: **not specified**

Integrating Quattor and virtualisation technologies

Wednesday, May 27, 2009 4:30 PM (30 minutes)

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.

Primary author: ROCHE, Ewan (CERN)

Presenter: ROCHE, Ewan (CERN)

Session Classification: Virtualisation II

Track Classification: Virtualisation

Contribution ID: 40

Type: **not specified**

Cybersecurity Update

Tuesday, May 26, 2009 11:30 AM (30 minutes)

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.

Primary author: IVEN, Jan (CERN)

Presenter: IVEN, Jan (CERN)

Session Classification: Security and Networking I

Track Classification: Security and Networking

Contribution ID: 41

Type: **not specified**

Virtualization and/vs security

Wednesday, May 27, 2009 5:00 PM (30 minutes)

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.

Primary author: IVEN, Jan (CERN)

Presenter: IVEN, Jan (CERN)

Session Classification: Virtualisation II

Track Classification: Virtualisation

Contribution ID: 42

Type: **not specified**

Benchmarking current PC server hardware

Tuesday, May 26, 2009 4:30 PM (30 minutes)

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.

Primary author: Mr WASCHK, Götz (DESY)

Co-authors: Dr PLEITER, Dirk (DESY); Mr WINTER, Frank (DESY); Dr SIMMA, Hubert (DESY); Dr WEGNER, Peter (DESY)

Presenter: Mr WASCHK, Götz (DESY)

Session Classification: Miscellaneous Talks I

Track Classification: Misc Talks

Contribution ID: 43

Type: **not specified**

Packaging Grid software for Linux distributions

Thursday, May 28, 2009 4:30 PM (30 minutes)

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.

Primary author: ELLERT, Mattias (Uppsala Universitet & NDGF)

Co-authors: WÄÄNÄNEN, Anders (Niels Bohr Institute); MÖLLER, Steffen (Universität zu Lübeck)

Presenter: ELLERT, Mattias (Uppsala Universitet & NDGF)

Session Classification: Operating Systems & Applications II

Track Classification: Operating Systems & Applications

Contribution ID: 44

Type: **not specified**

GridKa Site Report

Monday, May 25, 2009 3:30 PM (20 minutes)

Current status of GridKa, e.g.

- new hardware
- problem discussion

Primary author: Mr ALEF, Manfred (Forschungszentrum Karlsruhe)

Presenter: Mr ALEF, Manfred (Forschungszentrum Karlsruhe)

Session Classification: Site reports II

Track Classification: Site reports

Contribution ID: 45

Type: **not specified**

CINBAD - CERN Investigation of Network Behaviour Anomaly Detection

Tuesday, May 26, 2009 1:30 PM (30 minutes)

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.

Primary authors: Mr HULBOJ, Milosz Marian (CERN); Mr JURGA, Ryszard Erazm (CERN)

Presenter: Mr JURGA, Ryszard Erazm (CERN)

Session Classification: Security and Networking II

Track Classification: Security and Networking

Contribution ID: 46

Type: **not specified**

Playing with Puppets instead of managing your systems

Monday, May 25, 2009 11:00 AM (30 minutes)

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.

Primary author: Mr BELLMAN, Thomas (National Supercomputer Centre, Sweden)

Presenter: Mr BELLMAN, Thomas (National Supercomputer Centre, Sweden)

Session Classification: Operating Systems & Applications

Track Classification: Operating Systems & Applications

Contribution ID: 47

Type: **not specified**

Virtualization on the Track I

Wednesday, May 27, 2009 1:00 PM (15 minutes)

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.

Primary author: FINNERN, Thomas (DESY)

Presenter: FINNERN, Thomas (DESY)

Session Classification: Virtualisation I

Track Classification: Virtualisation

Contribution ID: 48

Type: **not specified**

TRIUMF Site Report

Monday, May 25, 2009 3:50 PM (20 minutes)

A summary of TRIUMF Site activities

Primary author: Dr MCDONALD, Steven (TRIUMF)

Presenter: Dr MCDONALD, Steven (TRIUMF)

Session Classification: Site reports II

Track Classification: Site reports

Contribution ID: 49

Type: **not specified**

A Highly Versatile Virtual Data Center Ressource Pool

Wednesday, May 27, 2009 1:15 PM (25 minutes)

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.

Primary author: BUJACK, Stefan (Deutsches Elektronen-Synchrotron DESY)

Presenter: BUJACK, Stefan (Deutsches Elektronen-Synchrotron DESY)

Session Classification: Virtualisation I

Track Classification: Virtualisation

Contribution ID: **50**

Type: **not specified**

SLAC Site Report

Monday, May 25, 2009 4:10 PM (20 minutes)

News from SLAC

Primary author: WACHSMANN, Alf (SLAC)

Presenter: WACHSMANN, Alf (SLAC)

Session Classification: Site reports II

Track Classification: Site reports

Contribution ID: 51

Type: **not specified**

DESY site report

Tuesday, May 26, 2009 8:50 AM (20 minutes)

DESY site report

Primary author: BEYER, Christoph (DESY HH)

Presenter: BEYER, Christoph (DESY HH)

Session Classification: Site reports III

Track Classification: Site reports

Contribution ID: 52

Type: **not specified**

HPC2N Data Center tour groups A + B

Tuesday, May 26, 2009 4:00 PM (30 minutes)

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.

Session Classification: Miscellaneous Talks I

Contribution ID: 54

Type: **not specified**

HPC2N Data Center tour groups E + F

Friday, May 29, 2009 8:30 AM (30 minutes)

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.

Session Classification: Miscellaneous Talks II

Contribution ID: 55

Type: **not specified**

HPC2N Data Center tour groups C & D

Thursday, May 28, 2009 2:30 PM (30 minutes)

Session Classification: Data centres

Contribution ID: 56

Type: **not specified**

A new Data Centre for CC-IN2P3

Thursday, May 28, 2009 2:00 PM (30 minutes)

Primary author: BOUTIGNY, Dominique (CC-IN2P3)

Session Classification: Data centres

Contribution ID: 57

Type: **not specified**

The CNAF Data Centre migration

Primary author: ROSSO, Felice (INFN)

Contribution ID: 58

Type: **not specified**

Impact of Filesystems on Application Performance in an HNEP Environment.

Friday, May 29, 2009 9:00 AM (30 minutes)

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..

Primary author: SAKREJDA, Iwona (LBNL/NERSC)

Presenter: SAKREJDA, Iwona (LBNL/NERSC)

Session Classification: Miscellaneous Talks II

Track Classification: Misc Talks

Contribution ID: 59

Type: **not specified**

State of Nordic R&E Networking

Tuesday, May 26, 2009 1:00 PM (30 minutes)

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

Session Classification: Security and Networking II

Contribution ID: **60**

Type: **not specified**

Cloud Security

Thursday, May 28, 2009 9:30 AM (30 minutes)

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.

Primary author: Dr KELSEY, David (RAL)

Presenter: Dr KELSEY, David (RAL)

Session Classification: Virtualisation III

Track Classification: Virtualisation

Contribution ID: **61**

Type: **not specified**

Opening remarks

Monday, May 25, 2009 9:30 AM (30 minutes)

Opening remarks by Prof. Bo Kågström, director HPC2N.

Primary author: WADENSTEIN, Mattias (NDGF)

Session Classification: Opening remarks

Contribution ID: 62

Type: **not specified**

SouthGrid and Oxford Site Report

Tuesday, May 26, 2009 8:30 AM (20 minutes)

SouthGrid and Oxford Status report

Primary author: Mr GRONBECH, Peter (Nuclear Physics Laboratory)

Presenter: Mr GRONBECH, Peter (Nuclear Physics Laboratory)

Session Classification: Site reports III

Track Classification: Site reports

Contribution ID: 63

Type: **not specified**

Virtualization on the Track II

Thursday, May 28, 2009 11:45 AM (15 minutes)

A short plenary discussion.

Summary

15 minutes to collect opinions and comments from the audience.

Primary author: FINNERN, Thomas (DESY)

Presenter: FINNERN, Thomas (DESY)

Session Classification: Virtualisation IV

Track Classification: Virtualisation

Contribution ID: **65**

Type: **not specified**

INFN T1 (NAF)

Monday, May 25, 2009 4:30 PM (20 minutes)

Session Classification: Site reports II

Contribution ID: **66**

Type: **not specified**

ASGC (Taipe)

Tuesday, May 26, 2009 9:30 AM (20 minutes)

Primary author: Mr SHIH, Jason (ASGC)

Session Classification: Site reports III

Contribution ID: **67**

Type: **not specified**

NERSC

Tuesday, May 26, 2009 9:10 AM (20 minutes)

Session Classification: Site reports III

Contribution ID: **68**

Type: **not specified**

Wrap-up & Conclusions

Friday, May 29, 2009 10:00 AM (15 minutes)

Wrap-up & Conclusions

Session Classification: Wrap-Up & Conclusions

Contribution ID: **69**

Type: **not specified**

Virtualisation discussion (continued)

Thursday, May 28, 2009 1:00 PM (30 minutes)

Session Classification: Data centres