

HEPiX Spring 2014 Workshop

Monday, 19 May 2014 - Friday, 23 May 2014

LAPP



Book of Abstracts

Contents

Fermilab Site Report - Spring 2014 HEPiX	1
New Windows security at CEA and IRFU	1
IRFU Saclay site report	1
INFN-T1 site report	1
CERN Site Report	2
Experiences with ATLAS and LHCb jobs in Vac virtual machines	2
Cluster Consolidation at NERSC	2
FermiCloud On-demand Services: Data-Intensive Computing on Public and Private Clouds	3
Intel Ivybridge vs. AMD Opteron: performance and power implications	4
Ceph at CERN: one year on	4
Future of Batch Processing at CERN	4
Support for Linux Control Groups	5
Shared datacenter in Orsay University : first results	5
GRIF and LAL Site Report	5
Batch System Review	6
CC IN2P3 experience with Univa Grid Engine	6
ASGC site report	6
BNL RACF Site Report	7
QMUL site report	7
IHEP Site Report	7
Scheduling of multicore jobs	7
Open Compute at CERN	8
from quattor to puppet, a T2 point of view	8

Big Data Transfer over Internet	8
A Year of Condor at the RAL Tier 1	9
Quattor Update	9
Security update	9
RAL Tier 1 Cloud & Virtualisation	10
Managing secrets	10
Emergency suspension list in WLCG	10
Scaling Agile Infrastructure, development and change management	11
Agile Infrastructure Monitoring	11
Ceph at the UK Tier 1	12
Agile Infrastructure: an updated overview of IaaS at CERN	12
Enabling multi-cloud resources at CERN within the Helix Nebula project	12
KIT Site Report	13
DESY Site Report	13
Beyond HS06 - Toward a New HEP CPU Benchmark	13
Scientific Linux Status and Futures	14
Field Experience in the Agile Infrastructure	14
Windows 8 Integration	14
Issue Tracking & Version Control Services status update	14
Status report from Tokyo Tier-2 for the one year operation after whole scale system upgrade	15
Update on CERN tape status	15
PIC Site Report	16
Evaluation of avoton CPU	16
Lesson learned after our recent cooling problem	16
IPv6 Deployment at FZU in Prague	17
Lavoisier : a data aggregation framework	17
Jefferson Lab Site Report	18
RAL Site Report	18
WIGNER Datacenter - Introduction	18

WIGNER Datacenter - Operational experience	19
Next Linux version at CERN	19
WLCG perfSONAR-PS Update	19
Nebraska Site Report	20
The Art of Running HTCondor as a batch system	20
Batch system data locality via managed caches	20
Measuring WLCG data streams at batch job level	21
Update on the bit-preservation Working Group	21
AGLT2 Site Update	22
Workshop wrap-up	22
LAPP welcome	22
Workshop logistics	22
IPv6 status and perfsonar testing in the UK	23
HEPiX configuration management working group update	23
The DESY Big Data Cloud Service	23
Business Continuity at DESY	24
10 Years of Scientific Linux	24
CentOS and Red Hat	24
Discussion about future OS for HEP	24
Announcement	24

Site reports / 1**Fermilab Site Report - Spring 2014 HEPiX**

Author: Keith Chadwick¹

¹ *Fermilab*

Corresponding Author: chadwick@fnal.gov

Fermilab Site Report - Spring 2014 HEPiX.

Security and networking / 2**New Windows security at CEA and IRFU**

Author: joel surget¹

¹ *CEA/Saclay*

Corresponding Author: joel.surget@cea.fr

In 2013/2014 the CEA has decided to change dramatically the security of the Windows PC and the way to manage them. I'll explain the new philosophy of the security based on two levels:

- Lateral security
- Escalade security

I'll explain the problematic for the end-users and also for the IT team.

Site reports / 3**IRFU Saclay site report**

Author: Pierrick Micout¹

¹ *Unknown*

What is new in the IRFU Saclay site?

Summary:

IRFU Saclay site report.

Site reports / 4**INFN-T1 site report**

Author: Andrea Chierici¹

¹ *INFN-CNAF*

Corresponding Author: chierici@cnaf.infn.it

INFN-T1 site update

Summary:

INFN-T1 site update

Site reports / 5

CERN Site Report

Author: Arne Wiebalck¹

¹ CERN

Corresponding Author: arne.wiebalck@cern.ch

News from CERN since the last workshop.

Grids, clouds, virtualisation / 6

Experiences with ATLAS and LHCb jobs in Vac virtual machines

Author: Andrew McNab¹

¹ University of Manchester (GB)

Corresponding Author: andrew.mcnab@cern.ch

We present experiences with running ATLAS and LHCb production jobs in virtual machines at Manchester and other sites in the UK using Vac. Vac is a self-contained VM management system in which individual hypervisor hosts act as VM factories to provide VMs contextualized for experiments, and offers an alternative to conventional CE/Batch systems and Cloud interfaces to resources. In the Vacuum model implemented by Vac, VMs appear spontaneously at sites, with contextualizations provided by the sites using templates provided by the experiments. This system takes advantage of the pilot job frameworks for managing jobs and cvmfs for managing software distribution, which together lead to these contextualizations being extremely simple in practice.

Vac is implemented as a daemon, vacd, which runs on each hypervisor host. Each daemon uses a peer-to-peer UDP protocol to gather information from other Vac daemons at the site about what mix of experiment VMs are already running, and acts autonomously to decide which VMs to start using a policy given in its configuration file. The UDP protocol is also used to avoid starting VMs for experiments which have no work available, by detecting when a VM has been started recently and has stopped immediately because the pilot framework client could find no work.

Vac has been running LHCb production jobs since 2013 and in 2014 a suitable ATLAS VM contextualization was developed and has been used to run ATLAS production work too. We present some preliminary comparisons of the efficiency of running LHCb and ATLAS jobs on batch worker nodes and in virtual machines using the same hardware.

Basic IT services / 7

Cluster Consolidation at NERSC

Author: Larry Pezzaglia¹

¹ *LBNL*

Corresponding Author: lmpezzaglia@lbl.gov

This talk will provide a case study of cluster consolidation at NERSC.

In 2012, NERSC began deployment of “Mendel”, a 500+ node, Infiniband-attached, Linux “meta-cluster” which transparently expands NERSC production clusters and services in a scalable and maintainable fashion. The success of the software automation infrastructure behind the Mendel multi-clustering model encouraged investigation into even more aggressive consolidation efforts.

This talk will detail one such effort: under the constraints of a 24x7, disruption-sensitive environment, NERSC staff merged a 400-node legacy production cluster, consisting of multiple hardware generations and ad-hoc software configurations, into Mendel’s automation infrastructure. By leveraging the hierarchical management features of the xCAT software package in combination with other open-source and in-house tools, such as Cfengine and CHOS, NERSC abstracted the unique characteristics of both clusters away below a unified management interface. Consequently, both cluster components are now managed as a single, albeit complex, integrated system.

Additionally, this talk will provide an update on the PDSF system at NERSC, including improvements to trending data collection and ongoing CHOS development.

Grids, clouds, virtualisation / 8

FermiCloud On-demand Services: Data-Intensive Computing on Public and Private Clouds

Author: Steven Timm¹

Co-author: Gabriele Garzoglio²

¹ *Fermilab*

² *FERMI NATIONAL ACCELERATOR LABORATORY*

Corresponding Author: timmssteve@yahoo.com

The FermiCloud project exists to provide on-demand computing and data movement services to the various experiments at Fermilab. We face a dynamically changing demand for compute resources and data movement, which we meet by enabling users to run on our own site, remote grid sites, and cloud sites. We also instantiate on-demand data movement and web caching services to support this remote analysis. In this presentation we will summarize some of our recent research results and outline the challenges of our current research projects.

These include coordinated launches of compute nodes and data movement servers, interoperability with new commercial clouds, idle machine detection, and exploration of distributed storage models.

Summary:

We have successfully demonstrated several methods of launching compute nodes to the public cloud including the vcluster software, the GlideinWMS, and OpenNebula cloud bursting features. Using the

GlideinWMS we have submitted several workflows from the NOvA experiment to FermiCloud and the public Amazon EC2 cloud. Current research and discussion centers on tiered policy-based provisioning. A large input to the provisioning model is the cost in time and bandwidth that it takes to transfer virtual machine images and launch them, which we have modeled in detail. We have also published results on high-performance calculations using virtualized Infiniband adapters.

Current research is focused on coordinated provisioning of diverse cloud-based scientific workflow. We are also developing an automated conversion service to change between different virtual machine image formats. We are developing a set of tools to port applications to run both on external grid sites and on public clouds. We will summarize the best practices that we have found in running physics applications on public cloud as well as investigations of running using spot pricing. Finally we continue to evaluate other commercial and community clouds.

Computing and batch systems / 9

Intel Ivybridge vs. AMD Opteron: performance and power implications

Author: Tony Wong¹

Co-authors: Alexandr Zaytsev²; Christopher Hollowell¹; Costin Caramarcu²; William Strecker-Kellogg³

¹ *Brookhaven National Laboratory*

² *Brookhaven National Laboratory (US)*

³ *Brookhaven National Lab*

The RACF has evaluated the Intel Ivybridge and AMD Opteron cpu's before an anticipated purchase of Linux servers for its RHIC and USATLAS programs in 2014. Price performance considerations are no longer sufficient as we must consider long-term power, cooling and space capacities in the data center. This presentation describes how these long-term considerations are increasingly altering hardware acquisition cycles at BNL.

Storage and file systems / 10

Ceph at CERN: one year on

Author: Daniel van der Ster¹

¹ *CERN*

Corresponding Author: daniel.vanderster@cern.ch

Ceph was introduced at CERN in early 2013 as a potential solution to new use-cases (e.g. cloud block storage) while also providing a path toward a consolidated storage backend for other services including AFS, NFS, etc...

This talk will present the outcome of the past year of testing and production experience with Ceph. We will present our real operations experience and lessons-learned, and review the state of each of the object, block, and filesystem components of Ceph.

Finally, we will present our plans moving forward, with a discussion about the potential and non-potential of Ceph as a backend for our physics data stores.

Computing and batch systems / 11**Future of Batch Processing at CERN****Author:** Janos Daniel Pek¹¹ *CERN***Corresponding Author:** daniel.pek@cern.ch

The CERN Batch System is comprised of 4000 worker nodes. 60 queues offer a service for various types of large user communities. In light of the recent developments driven by the Agile Infrastructure and the more demanding processing requirements, the Batch System will be faced with increasingly challenging scalability and flexibility needs. Last year the CERN Batch Team has started to evaluate three candidate batch systems: SLURM, HTCondor and GridEngine. This year as we are reaching a conclusion, one of our candidates is HTCondor. In this talk we give a short reminder of our requirements and our preliminary results from last year. Then we'll focus on HTCondor, our experience with it thus far, our testing framework and the results of our performance tests. Finally, we give a summary of the foreseeable challenges we would have to face if we decide to migrate the CERN Batch Service to Condor.

Computing and batch systems / 12**Support for Linux Control Groups****Author:** Daniel Gruber¹¹ *Univa*

Current Linux distributions including support for a new kernel enhancement called control groups (cgroups). This talk is about how Univa Grid Engine integrates the Linux cgroup subsystems for better resource isolation, utilization, and limitation in the job execution and resource allocation context. Example configurations and use cases for today's NUMA compute nodes are discussed.

IT facilities and business continuity / 13**Shared datacenter in Orsay University : first results****Author:** Michel Jouvin¹¹ *Universite de Paris-Sud 11 (FR)***Corresponding Author:** jouvin@lal.in2p3.fr

As presented at past HEPiX, 8 labs in Orsay region/university started 2 years ago a project to build a new datacenter aimed to replace the existing inefficient computing rooms. This project has been delivered on-time and is in production since last October. This presentation will summarize the needs that motivated the project, the design choices, the building phase experience and gives an early feedback after 6 months of operations. It will also present the future directions for this project and the other related initiatives.

Site reports / 14

GRIF and LAL Site Report

Author: Michel Jouvin¹

¹ *Universite de Paris-Sud 11 (FR)*

Corresponding Author: jouvin@lal.in2p3.fr

This site report will cover GRIF grid site and LAL internal computing

Computing and batch systems / 15

Batch System Review

Author: Michel Jouvin¹

¹ *Universite de Paris-Sud 11 (FR)*

Corresponding Author: jouvin@lal.in2p3.fr

WLCG GDB organized a meeting last March about batch systems. With an audience mostly from grid sites, it has been a successful review of the main batch systems used in the community by sites with concrete experience. This presentation will summarize what was presented and the main conclusions of this meeting.

Computing and batch systems / 16

CC IN2P3 experience with Univa Grid Engine

Authors: Nadia Lajili¹; Suzanne Poulat²

¹ *Centre de Calcul IN2P3*

² *Centre de calcul IN2P3*

Corresponding Author: suzanne@in2p3.fr

After 20 years using a home made batch system named BQS (Batch Queuing System), CC-IN2P3 decided to move to Grid Engine in order to offer the scalability and robustness needed for multi-experiment production, HEP et non HEP.

The site migrated from BQS to Oracle Grid Engine in 2011, then switched to Univa's version after only two years, in June 2013.

The talk presents the assessment of the change from Oracle to Univa and gives an overview of the configuration equilibrating user requirements and constraints of the site's infrastructure, especially for multi-core jobs. Finally, plans for the deployment of new features are shown and requests to Univa are explained.

Site reports / 17

ASGC site report

Author: Hung-Te Lee¹

¹ *Academia Sinica (TW)*

Corresponding Author: felix@twgrid.org

Site report for ASGC

Site reports / 18

BNL RACF Site Report

Author: Ofer Rind¹

¹ *BROOKHAVEN NATIONAL LABORATORY*

Corresponding Author: rind@bnl.gov

A summary of developments at BNL's RHIC/ATLAS Computing Facility since the last HEPiX meeting.

Site reports / 19

QMUL site report

Author: Christopher John Walker¹

¹ *University of London (GB)*

Corresponding Author: christopher.john.walker@cern.ch

Site report from Queen Mary University of London.

Site reports / 20

IHEP Site Report

Author: Jingyan Shi¹

Co-author: Fazhi QI

¹ *IHEP*

The site report will give a summary of IHEP site status including local cluster, EGI site. Besides, it will talk about improvement of distribution computing we have done.

Computing and batch systems / 21

Scheduling of multicore jobs

Author: Nathalie Rauschmayr¹

¹ CERN

Corresponding Author: nathalie.rauschmayr@cern.ch

Nowadays, the Worldwide LHC Computing Grid consists of multi and manycore CPUs. A lot of work is undertaken by the experiments and the HEP community, in order to use these resources more efficiently. As a result, the parallelization of applications has been the main goal so far in order to allow a parallel execution of jobs. However, experiments must also consider how to schedule multicore jobs within the Computing Grid. Taking into account the trend of going towards manycore architectures, tasks might not scale sufficiently well on large number of cores. Since non linear speedup can drastically decrease overall throughput, a scheduler must define the best degree of parallelism for each job.

The aim of the presentation is to define the scheduling problem and to present algorithms to solve it. Related problems, like estimation of job runtimes, will be also discussed.

IT facilities and business continuity / 22

Open Compute at CERN

Authors: Marco Guerri¹; Olof Barring¹

¹ CERN

Corresponding Author: olof.barring@cern.ch

The Open Compute Project, OCP (<http://www.opencompute.org/>), was launched by Facebook in 2011 with the objective of building efficient computing infrastructures at lowest possible cost. The technologies are released as open hardware design, with the goal to develop servers and data centers following the model traditionally associated with open source software projects. In order to try out the hardware we acquired two OCP twin servers (<http://hyvesolutions.com/resources/docs/2013Hyve1500Datasheet.pdf>) in 2013. The servers have been tested and compared with our production hardware. Some results from this testing will be presented as well as the future plans for a possible larger deployment.

Basic IT services / 23

from quattor to puppet, a T2 point of view

Author: Frederic Schaer¹

¹ CEA

Corresponding Author: frederic.schaer@cea.fr

The IRFU site, member of the GRIF wLCG T2 site, decided to move from quattor to puppet in 2012. The migration was almost complete early april 2014.

This talk will focus mainly on the goals, the ways to achieve what we achieved, the manpower that was required, what we gained with puppet and the new challenges that we must now face as a T2 with this management tool.

Grids, clouds, virtualisation / 24

Big Data Transfer over Internet

Author: Andrey SHEVEL¹

Co-authors: Oleg SADOV¹; Sergey Khoruzhnikov¹; Vladimir Grudin¹

¹ *University of Information Technology, Mechanics, and Optics*

In many cases where Big Data phenomenon is taken place there is the need to transfer the Big Data from one point of computer network to another point. Quite often those points are far away from each other. The transfer time is significant factor to transfer the Big Data. During this time the features of the data link might be changed drastically including interruptions of channel operation once or more times during data transfer. There are a number of known utilities/systems which are used for Big Data transfer. The authors investigate which utilities/systems are more suitable for Big Data transfer and which are most important architecture features for such the systems. It is of interest the comparison study of the data transfer methods. The testbed is developed to compare the data transfer utilities and study how Software Defined Networks (SDN) approach affects the Big Data transfer.

Computing and batch systems / 25

A Year of Condor at the RAL Tier 1

Author: Ian Peter Collier¹

Co-author: Andrew David Lahiff¹

¹ *STFC - Rutherford Appleton Lab. (GB)*

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

It's been almost a year since we first started running ATLAS and CMS production jobs at RAL using HTCondor, and 6 months since we fully migrated from Torque/Maui. This talk will discuss our experience so far and future plans.

Basic IT services / 26

Quattor Update

Author: Ian Peter Collier¹

Co-author: James Adams²

¹ *STFC - Rutherford Appleton Lab. (GB)*

² *STFC RAL*

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

A report on the status of the Quattor toolset, with particular emphasis on recent developments in both the user and development communities.

Security and networking / 27

Security update

Author: Vincent Brillault¹

¹ *CERN*

Corresponding Author: vincent.brillault@cern.ch

This presentation provides an update of the security landscape since the last meeting. It describes the main vectors of compromises in the academic community and presents interesting recent attacks. It also covers security risks management in general, as well as the security aspects of the current hot topics in computing, for example identity federation and virtualisation.

Grids, clouds, virtualisation / 28

RAL Tier 1 Cloud & Virtualisation

Author: Ian Peter Collier¹

¹ *STFC - Rutherford Appleton Lab. (GB)*

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

The RAL Tier 1 is now working deploying a production quality private cloud, to meet the emerging needs of both the Tier 1 and STFCs Scientific Computing Department. This talk will describe the work so far and the roadmap for the coming year. We will also discuss other virtualisation developments.

Basic IT services / 29

Managing secrets

Author: Sven Sternberger¹

¹ *D*

Corresponding Author: sven.sternberger@desy.de

The talk will discuss the problems which arise from managing and distributing secrets like root passwords, keytabs, certificates in a large site.

Secrets are needed in the process of installing and administrating of compute and storage systems. They should be accessible by authorized admins and from the system they belong to. There should be a way to audit the information to enforce the policies from your security department. For example quality and lifetime of passwords.

In the presentation we will describe the workflows at DESY/Hamburg and show systems we use today and their deficits. Then we will describe our upcoming solution, and threats we still see.

Security and networking / 30

Emergency suspension list in WLCG

Author: Vincent Brillault¹

¹ *CERN*

Corresponding Author: vincent.brillault@cern.ch

The Emergency suspension list (also known as central banning list) is finally getting deployed in WLCG, allowing quick automated responses to incidents. This short presentation will present the goal of this new features, the technology behind this system and details about the current deployment.

Basic IT services / 31

Scaling Agile Infrastructure, development and change management

Author: Ben Jones¹

¹ *CERN*

Corresponding Author: ben.dylan.jones@cern.ch

As the Agile Infrastructure project scaled from being a development effort of a few people that could sit in one meeting room to a production service for the CERN computer centre, what changes were needed to our service, tools and workflow? We will look at the technical challenges scaling the puppet infrastructure, scaling the development effort of puppet code, and also the procedural changes to QA, change management and continuous delivery.

How well does puppet scale to thousands of nodes and how much work is involved? What tools have been useful to manage an agile workflow? How can we fit a fast moving development pipeline to different groups with different expectations of the speed of change?

Basic IT services / 32

Agile Infrastructure Monitoring

Author: Pedro Andrade¹

Co-authors: Benjamin Fiorini¹; Luis Pigueiras²; Lukasz Starakiewicz³; Miguel Coelho dos Santos¹; Susie Murphy¹

¹ *CERN*

² *Universidad de Oviedo (ES)*

³ *T*

Corresponding Author: pedro.andrade@cern.ch

The agile infrastructure monitoring team is working on new solutions to modernise and improve how monitoring and analytics is done at CERN. We will give an update on these activities, in particular the recent progress on testing and adopting different open source technologies (e.g. hadoop, elasticsearch, flume, kibana) for the various monitoring architecture layers. We will report on the efforts to build a monitoring and analytics community with participants from different areas: service managers, users, security, management, etc. We will present concrete examples on how this community is using these different solutions to improve their daily activities.

Storage and file systems / 33**Ceph at the UK Tier 1****Author:** George Ryall¹**Co-authors:** James Adams²; Shaun De Witt³¹ *STFC*² *STFC RAL*³ *STFC - Rutherford Appleton Lab. (GB)***Corresponding Author:** george.ryall@stfc.ac.uk

We are trialling the use of Ceph both as a file-system and as a cloud storage back end. I will present our experiences so far.

Basic IT services / 34**Agile Infrastructure: an updated overview of IaaS at CERN****Author:** Stefano Zilli¹¹ *CERN***Corresponding Author:** stefano.zilli@cern.ch

The CERN private cloud has been in production since July 2013 and has grown steadily to 60000 cores, hosting more than 5500 Virtual Machines for 370 users and 140 shared projects. New features have been made available this year like block storage and IPv6. This presentation will provide an overview of the current status of the infrastructure and of the plans for the next developments and evolution of the services. During this talk, different topics will be dealt with such as the successful migration from OpenStack Grizzly to Havana, the imminent upgrade to IceHouse, IPv6 ready machines and the metering infrastructure.

Grids, clouds, virtualisation / 35**Enabling multi-cloud resources at CERN within the Helix Nebula project****Author:** Domenico Giordano¹**Co-authors:** Alessandro Di Girolamo¹; Bob Jones¹; Cristovao Jose Domingues Cordeiro²; Laurence Field¹¹ *CERN*² *ADI Agencia de Inovacao (PT)***Corresponding Author:** domenico.giordano@cern.ch

Helix Nebula – the Science Cloud is a European public-private-partnership between leading scientific research organisations (notably CERN, EMBL and ESA) and European IT cloud providers. Its goal is to establish a Cloud Computing Infrastructure for the European Research Area and the Space Agencies, serving as a platform for innovation and evolution of a federated cloud framework for e-Science.

CERN contributes to the Helix Nebula initiative by providing a flagship use case: the exploitation

of cloud resources within the workload management system of the ATLAS and CMS experiments at the Large Hadron Collider.

This contribution will summarize the CERN experience in Helix Nebula during the past two years and the lessons learned in deploying applications from ATLAS and CMS with several commercial providers. The integration with the experiment framework will also be explained.

Site reports / 36

KIT Site Report

Author: Manfred Alef¹

¹ *Karlsruhe Institute of Technology (KIT)*

Corresponding Author: manfred.alef@kit.edu

KIT Site Report

Site reports / 37

DESY Site Report

Author: Yves Kemp¹

¹ *Deutsches Elektronen-Synchrotron (DE)*

Corresponding Author: yves.kemp@cern.ch

DESY site report for Spring 2014 HEPiX workshop

Computing and batch systems / 38

Beyond HS06 - Toward a New HEP CPU Benchmark

Author: Manfred Alef¹

Co-author: Michele Michelotto²

¹ *Karlsruhe Institute of Technology (KIT)*

² *Universita e INFN (IT)*

Corresponding Author: manfred.alef@kit.edu

The HEPiX Benchmarking Working Group is preparing for the deployment of a successor of the widely used HS06 benchmark.

- Why we are looking for a replacement of HS06
- Summary of discussions at GDB
- Requirements
- Benchmark candidates
- Volunteers

IT end user services and operating systems / 39**Scientific Linux Status and Futures**

Author: Connie Sieh¹

Co-authors: Bonnie King¹; Patrick Riehecky¹

¹ *Fermilab*

Corresponding Author: csieh@fnal.gov

Status of Scientific Linux and Futures

Basic IT services / 40**Field Experience in the Agile Infrastructure**

Author: Jerome Belleman¹

¹ *CERN*

Corresponding Author: jerome.belleman@cern.ch

As the Agile Infrastructure is moving forwards at CERN, more and more services are migrating to it. New tools are put in place to get the most out of its strengths while we learn lessons from the problems we hit when converting services from Quattor. At CERN, a number of services have made some significant progress in the migration to the new infrastructure; the batch service, several interactive services, CEs and VOMS are but a few examples.

In this talk, we will describe some aspects of the migration process, such as virtualisation, Puppet configuration and alarming. We will discuss the strengths of the Agile Infrastructure which make running services easier, and expose the problematic areas for which we will present some of the future projects which are to address them.

IT end user services and operating systems / 41**Windows 8 Integration**

Author: Sebastien Dellabella¹

Co-author: Guillaume Metral¹

¹ *CERN*

Corresponding Author: sebastien.dellabella@cern.ch

In this presentation we will talk about Windows 8 and how we are integrating it at CERN. What are the issues we met and how we solved them for our users community.

We will focus on issues, customization and deployment.

If you have to start a Windows 8 pilot project in your organization, you must be there.

IT end user services and operating systems / 42

Issue Tracking & Version Control Services status update

Author: Alvaro Gonzalez Alvarez¹

Co-authors: Anna Trzcinska²; Borja Aparicio Cotarelo³; David Asbury¹; Georgios Koloventzos¹; Nils Hoimyr¹

¹ CERN

² Warsaw University of Technology (PL)

³ Universidad de Oviedo (ES)

Corresponding Authors: anna.trzcinska@cern.ch, alvaro.gonzalez.alvarez@cern.ch

The current efforts on the issue tracking and version control services at CERN will be presented. Special attention to the new central git service, the integration between issue tracking and version control and future service deployments.

Site reports / 43

Status report from Tokyo Tier-2 for the one year operation after whole scale system upgrade

Author: Tomoaki Nakamura¹

Co-authors: Hiroshi Sakamoto¹; I Ueda¹; Nagataka Matsui¹; Tetsuro Mashimo¹

¹ University of Tokyo (JP)

Corresponding Author: tomoaki.nakamura@cern.ch

The Tokyo Tier-2, which is located at International Center for Elementary Particle Physics (ICEPP) in the University of Tokyo, was established as a regional analysis center in Japan for the ATLAS experiment. The official operation with WLCG was started in 2007 after the several years' development since 2002. In December 2012, we have replaced almost all hardware as the third system upgrade to deal with analysis for further growing data of the ATLAS experiment. The number of CPU cores are increased by a factor of two from the previous system (9984 cores including CPUs for service instance), and the performance of individual CPU core is improved by 20 % according to the HEPSPEC06 benchmark test at 32bit compile mode. The score is estimated as 18.03 per core under Scientific Linux 6 by using Intel Xeon E5-2680 2.70 GHz. As of February 2013, 2560 CPU cores and 2.00 PB of disk storage have already been deployed for ATLAS. They have been operated stably with 95% availability in one year operation after the system upgrade. Since the number of CPU cores in the new worker node was increased from 8 cores to 16 cores, the local I/O performance for the data staging area might become a possible bottleneck for the job throughput. We have evaluated the performance by making a special worker node, which have a SSD for the local storage, at the mixture situation of running real ATLAS production jobs and analysis jobs. In consequence, we could confirm that SAS-HDD attached with nominal worker nodes at Tokyo Tier-2 is not a bottleneck for the long batch type jobs at least for the situation of 16 jobs running concurrently in one node. In this report, we would like to introduce several results of the evaluation of the local I/O performance with some experiences on the site operation.

Storage and file systems / 44

Update on CERN tape status

Author: German Cancio Melia¹

Co-author: Dmitry Ozerov²

¹ CERN

² DESY

Corresponding Author: german.cancio.melia@cern.ch

CERN stores over 100PB of data on tape via CASTOR and TSM. This talk will present the current status of the CERN tape infrastructure, with a particular focus on tape performance and efficiency and the status of the large media repacking exercise.

Site reports / 45

PIC Site Report

Author: Jose Flix Molina¹

¹ *Centro de Investigaciones Energ. Medioambientales y Tecn. - (ES)*

Corresponding Author: jose.flix.molina@cern.ch

Spring 2014 PIC Tier 1 site report covering recent updates, improving energy efficiency by means of free-cooling techniques, and preparation for new challenges.

Computing and batch systems / 46

Evaluation of avoton CPU

Author: Andrea Chierici¹

¹ *INFN-CNAF*

Corresponding Author: chierici@cnaf.infn.it

At INFN-T1 we are facing the problem of TCO of computing nodes, which count for the bigger part of our electricity bill.

Intel recently introduced the Avoton SOC, targeted on the microserver, entry communication infrastructure and cloud storage market. We benchmarked this CPU and evaluated the possible adoption of this technology in our computing farm.

Summary:

At INFN-T1 we are facing the problem of TCO of computing nodes, which count for the bigger part of our electricity bill.

Intel recently introduced the Avoton SOC, targeted on the microserver, entry communication infrastructure and cloud storage market. We benchmarked this CPU and evaluated the possible adoption of this technology in our computing farm.

IT facilities and business continuity / 47

Lesson learned after our recent cooling problem

Author: Andrea Chierici¹

Co-author: Andrea Ferraro²

¹ *INFN-CNAF*

² *Unknown*

Corresponding Author: chierici@cnaif.infn.it

In march we had a major cooling problem in our computing center and we had to completely shut the center down.

We learnt a lot from this problem and would like to share the experience within the community.

Summary:

In march we had a major cooling problem in our computing center and we had to completely shut the center down.

We learnt a lot from this problem and would like to share the experience within the community.

Security and networking / 48

IPv6 Deployment at FZU in Prague

Authors: Marek Elias¹; Tomas Kouba¹

Co-authors: Jiri Chudoba ¹; Lukas Fiala ¹

¹ *Institute of Physics ASCR (FZU)*

At FZU we are continuing with deployment of IPv6 in our testbed as well as the production network. On dual stack, we are currently running several subclusters of worker nodes and our DPM storage system.

Production data transfers from DPM to dualstack worker nodes using lcg-cp are currently running via IPv6. We present our experience with this deployment, new nagios sensors needed in this situation, results of our tests with IPv6 transfers using webdav, and several news from our IPv6 testbed.

Basic IT services / 49

Lavoisier : a data aggregation framework

Authors: Lequeux Olivier¹; Sylvain Reynaud¹

¹ *CNRS*

Corresponding Author: sreynaud@in2p3.fr

Many of us need tools for service monitoring adapted to our site specificities, or tools to do custom processing on user data. Regardless of the use-cases, we have to develop (or get developed) applications that aggregate, process and format data from heterogeneous data sources.

Lavoisier (<http://software.in2p3.fr/lavoisier>) is a framework, which enables building such applications by assembling reusable software components (i.e. plugins). These applications can then be used through a RESTful web service API, a web interface or a command line interface with little effort.

The Lavoisier framework is developed by CC-IN2P3 and used by several projects; the Operations Portal of the European Grid Infrastructure (EGI), the VAPOR portal, and some CC-IN2P3 internal tools.

The presentation will give an overview of Lavoisier, and explain how it can help to easily get a maintainable, performant, robust and secure data aggregation application, while focusing on business code.

Site reports / 50

Jefferson Lab Site Report

Author: Sandy Philpott¹

¹ *JLAB*

Corresponding Author: sandy.philpott@jlab.org

The JLab talk will cover our current high performance and experimental physics computing status, including node-sharing between clusters for the 12GeV data challenges, Puppet configuration management plans, our latest GPU and MIC environment, workflow tools, LTO6 integration into the mass storage system, initial results of XFS on Linux testing, and plans for a Lustre 2.5 update and LMDS reconfiguration.

Site reports / 51

RAL Site Report

Author: Martin Bly¹

¹ *STFC-RAL*

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

Update for UK Tier1 and RAL

Site reports / 52

WIGNER Datacenter - Introduction

Author: Szabolcs Hernath¹

¹ *Hungarian Academy of Sciences (HU)*

Corresponding Author: szabolcs.tamas.hernath@cern.ch

As newcomers to the HEPiX community, WIGNER Datacenter, the newly established scientific computing facility of the WIGNER Research Centre for Physics in Budapest, would like to give an introduction on its background, construction and mission and model of operation. Featuring a long-term sustainable, energy-efficient and high availability infrastructure, WIGNER Datacenter aims to provide a full range of computing services (including hosting, cluster and cloud based resources) to the scientific community.

IT facilities and business continuity / 53**WIGNER Datacenter - Operational experience****Author:** Szabolcs Hernath¹**Co-author:** Domokos Szabó²¹ *Hungarian Academy of Sciences (HU)*² *MTA WIGNER RCP***Corresponding Author:** szabolcs.tamas.hernath@cern.ch

In this talk we would like to give a summary on the experiences of the first year of operation of the WIGNER Datacenter. We will discuss the topics of infrastructure operations, facility management, energy efficiency and value added hosting services, with a special focus on the CERN@WIGNER project, the hosting of the external capacity of CERN Tier-0 resources. We will highlight some of the difficulties and pitfalls, along with insights and best practices we gathered during our first year of operation.

IT end user services and operating systems / 54**Next Linux version at CERN****Author:** Jarek Polok¹¹ *CERN***Corresponding Author:** jaroslaw.polok@cern.ch

CERN is maintaining and deploying Scientific Linux CERN since 2004. In January 2014 CentOS and Red Hat announced joining forces in order to provide common platform for open source community project needs. How is this merger affecting plans for future CERN Linux version ?

Security and networking / 55**WLCG perfSONAR-PS Update****Author:** Shawn Mc Kee¹¹ *University of Michigan (US)***Corresponding Author:** shawn.mckee@cern.ch

As reported at the last HEPiX meeting, the WLCG has been supporting the deployment of perfSONAR-PS Toolkit instances at all WLCG sites over the last year. The WLCG perfSONAR-PS Deployment Task Force has now wrapped up its work in April 2014.

The perfSONAR network monitoring framework was evaluated and agreed as a proper solution to cover the WLCG network monitoring use cases: it allows WLCG to plan and execute latency and bandwidth tests between any instrumented endpoint through a central scheduling configuration, it allows archiving of the metrics in a local database, it provides a programmatic and a web based interface exposing the tests results; it also provides a graphical interface for remote management operations.

In this presentation we will provide an update on the status of perfSONAR in WLCG and future plans for commissioning and maintaining perfSONAR in the scope of the WLCG Operations Coordination initiative and its role in supporting higher level services that are under development.

Summary:

An overview of perfSONAR status and plans in WLCG

Site reports / 56

Nebraska Site Report

Author: Brian Paul Bockelman¹

¹ *University of Nebraska (US)*

Corresponding Author: brian.bockelman@cern.ch

The Holland Computing Center at the University of Nebraska-Lincoln hosts the state's research computing resources. There are several grid-enabled clusters available to HEP for opportunistic computing and a CMS Tier-2 site.

In this presentation, we will cover the recent updates to site networking and the CMS Tier-2 cluster. Particular attention will be paid to:

1. The recent rollout of IPv6 to the production services.
2. Software upgrades, including the OSG software stack and planning for RHEL7.
3. Progress on the site's 100Gbps upgrade.

We will also give context in how the Nebraska Tier-2 updates fit in to the broader USCMS activities

Computing and batch systems / 57

The Art of Running HTCondor as a batch system

Author: Brian Paul Bockelman¹

¹ *University of Nebraska (US)*

Corresponding Author: brian.bockelman@cern.ch

HTCondor is a well known platform for distributed high-throughput computing and often resembles a the Swiss-Army-knife of computing - there's a bit of something for everyone. With a user manual weighing in at about 1,100 printed pages, there's no wonder that sysadmins can overlook some of the most exciting features.

This presentation will be dedicated to uncovering the hidden gems for running HTCondor as a batch system - useful features that are well-hidden, under-appreciated, or very recently added. This broad overview will include topics in worker node resource management, scripting, monitoring, deployment, and debugging the system.

Storage and file systems / 58**Batch system data locality via managed caches****Author:** Max Fischer¹¹ *KIT - Karlsruhe Institute of Technology (DE)***Corresponding Author:** max.fischer@cern.ch

Modern data processing solutions increasingly rely on data locality to achieve high data access rates and scalability. In contrast the common HEP system architectures emphasis uniform resource pools with minimal locality, allowing even for cross-site data access. The concept for the new High Performance Data Analysis (HPDA) Tier3 at KIT aims at introducing data locality to HEP batch systems. Coordinating dedicated cache drives on worker nodes, existing storage hierarchies are extended into the active batch system. The presentation will illustrate the considerations of extending the classic batch architecture and showcase the planned software and hardware architecture of the HPDA T3.

Security and networking / 59**Measuring WLCG data streams at batch job level****Author:** Eileen Kuhn¹¹ *KIT - Karlsruhe Institute of Technology (DE)***Corresponding Author:** eileen.kuhn@cern.ch

Batch system monitoring and related system monitoring tools allow tracking data streams at different levels. With the introduction of federated data access to the workflows of WLCG it is becoming increasingly important for data centers to understand specific data flows regarding storage element accesses, firewall configurations, or the scheduling of workflows themselves. For this purpose a proof of concept has been implemented at the GridKa Tier1 center for monitoring data streams of batch jobs. The approach aims for a direct integration into the existing batch system to enhance batch job statistics by adding continuous traffic profiles for WLCG jobs and pilots. The presentation will introduce the general concept of the developed tool and integration into the batch system as well as first results of measurements at GridKa.

Storage and file systems / 60**Update on the bit-preservation Working Group****Author:** German Cancio Melia¹**Co-author:** Dmitry Ozerov²¹ *CERN*² *DESY***Corresponding Author:** german.cancio.melia@cern.ch

In this talk, we will provide an update of bit-level preservation WG activities, notable on the ongoing work on a set of recommendations and on a model for estimating long-term (10-20-30 years) archiving cost outlooks.

Site reports / 61**AGLT2 Site Update**

Author: Shawn Mc Kee¹

Co-authors: Benjeman Jay Meekhof¹; James Donald Koll²; Philippe Alain Luc Laurens²; Robert Ball¹

¹ *University of Michigan (US)*

² *Michigan State University (US)*

Corresponding Author: shawn.mckee@cern.ch

I will present an update on our site since the last report and cover our work with dCache, perfSONAR-PS, VMWare and experience with Cobbler and CFengine3 as our node provisioning system. There will also be an overview of our recent networking changes including the status of our new 100G connection to the WAN. I conclude with a summary of what has worked and what problems we encountered and indicate directions for future work.

Summary:

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

Miscellaneous / 62**Workshop wrap-up**

Author: Helge Meinhard¹

¹ *CERN*

Corresponding Author: helge.meinhard@cern.ch

Wrap-up

Miscellaneous / 63**LAPP welcome**

Author: Nadine Neyroud¹

¹ *LAPP/CNRS*

Welcome

Miscellaneous / 64**Workshop logistics**

Author: Frederique Chollet¹

¹ *Centre National de la Recherche Scientifique (FR)*

Logistics

Security and networking / 65

IPv6 status and perpsonar testing in the UK

Author: Christopher John Walker¹

Co-author: Duncan Rand ²

¹ *University of London (GB)*

² *Imperial College Sci., Tech. & Med. (GB)*

Corresponding Author: christopher.john.walker@cern.ch

IPv6 rollout at UK sites varies from one site where nearly all services are dual stack (Imperial), to others without any IPv6 addresses. The current rollout status will be presented. In addition, results of IPv6 connectivity testing using perpsonar will be discussed.

Basic IT services / 66

HEPiX configuration management working group update

Author: Ben Jones¹

Co-author: Yves Kemp ²

¹ *CERN*

² *Deutsches Elektronen-Synchrotron (DE)*

Corresponding Author: ben.dylan.jones@cern.ch

A year ago we began working with other sites to see how we could best share knowledge and effort amongst sites migrating to puppet. This talk will present a reminder of the working group, its formation and mandate, and how puppet had been developing already amongst earlier adopters. We will discuss how puppet module development occurs in the wider puppet community, and what conventions the working group has agreed upon. The current state of HEP related development will be explained, along with areas of interest that have been discussed. Along with that we will look at areas in which there are either open questions or perceived barriers to development.

Storage and file systems / 67

The DESY Big Data Cloud Service

Authors: Patrick Fuhrmann¹; Peter van der Reest¹

¹ *DESY*

Corresponding Author: peter.van.der.reest@desy.de

DESY -IT- has implemented a cloud storage service on the basis of dCache.
The talk will describe architecture and service concepts.

IT facilities and business continuity / 68

Business Continuity at DESY

Author: Yves Kemp¹

Co-authors: Peter van der Reest¹; kars Ohrenberg

¹ *DESY*

Corresponding Author: peter.van.der.reest@desy.de

a collection of themes and thoughts on BC, covering among others measures, procedures and dependencies

IT end user services and operating systems / 69

10 Years of Scientific Linux

Author: Alan Silverman¹

¹ *CERN (retired)*

Corresponding Author: alan.silverman@cern.ch

SL was announced to the world at the Spring 2004 HEPiX meeting in Edinburgh so it seems a good moment to review its origins and how it became the preferred Linux of most HEP sites.

IT end user services and operating systems / 70

CentOS and Red Hat

Author: Karanbir Singh^{None}

Corresponding Author: kbsingh@karan.org

IT end user services and operating systems / 71

Discussion about future OS for HEP

Announcement