

HEPiX Spring 2015 Workshop



Report of Contributions

Contribution ID: 0

Type: **not specified**

CERN Site Report

Tuesday, March 24, 2015 9:30 AM (15 minutes)

News from CERN since the Lincoln meeting.

Primary author: Dr WIEBALCK, Arne (CERN)

Presenter: Dr WIEBALCK, Arne (CERN)

Session Classification: Site Reports

Track Classification: Site reports

Contribution ID: 1

Type: **not specified**

Accelerating Scientific Analysis with SciDB

Monday, March 23, 2015 2:00 PM (25 minutes)

SciDB is an open-source analytical database for scalable complex analytics on very large array or multi-structured data from a variety of sources, programmable from Python and R. It runs on HPC, commodity hardware grids, or in a cloud and can manage and analyze terabytes of array-structured data and do complex analytics in-database.

We present an overall description of the SciDB framework and describe its implementation at NERSC at Lawrence Berkeley National Laboratory. A case study using SciDB to analyze data from the LUX dark matter detector is described. LUX is a 370 kg liquid xenon time-projection chamber built to directly detect galactic dark matter in an underground laboratory 1 mile under the Black Hills in South Dakota, USA. In the 2013 initial data run, LUX collected 86 million events and wrote 32 TB of data of which only 160 events are retained for final analysis. The data rate for the new dark matter run starting in 2014 is expected to exceed 250 TB / year. We describe how SciDB is used to dramatically streamline the data collection and analysis, and discuss future plans for a large parallel SciDB array at NERSC.

Primary author: GERHARDT, Lisa (LBNL)

Co-author: Mr YAO, Yushu (LBNL)

Presenters: GERHARDT, Lisa (LBNL); Mr YAO, Yushu (LBNL)

Session Classification: End-user Services and Operating Systems

Track Classification: End-User IT Services & Operating Systems

Contribution ID: 2

Type: **not specified**

Jefferson Lab Scientific and High Performance Computing

Monday, March 23, 2015 11:55 AM (15 minutes)

Current high performance and experimental physics computing environment updates: core exchanges between USQCD and Experimental Physics clusters for load balancing, job efficiency, and 12GeV data challenges; Nvidia K80 GPU experiences and updated Intel MIC environment; update on locally developed workflow tools and write-through to tape cache filesystem; status of LTO6 integration into our MSS; ZFS on Linux production environment, and status of our Lustre 2.5 update. It will also indicate JLab's interest in upcoming Ceph BOF sessions and workshops.

Primary author: PHILPOTT, Sandy (JLAB)

Presenter: PHILPOTT, Sandy (JLAB)

Session Classification: Site Reports

Track Classification: Site reports

Contribution ID: 3

Type: **not specified**

Overview of operational issues at CERN in the recent past

Monday, March 23, 2015 5:20 PM (25 minutes)

Many of you are aware of the power incident we had on the 16th October during the last HEPiX workshop. I will give a detailed explanation of what happened, the impact on IT services as well as the actions taken to recover from the incident. I will also note some improvements that will be implemented as a result of this incident. I will then go on to discuss other operations incidents that we have had in the recent past as well as a summary of the hardware failures.

Primary author: SALTER, Wayne (CERN)

Presenter: SALTER, Wayne (CERN)

Session Classification: IT Facilities and Business Continuity

Track Classification: IT Facilities & Business Continuity

Contribution ID: 4

Type: **not specified**

Update on software collaboration services at CERN

Tuesday, March 24, 2015 10:55 AM (20 minutes)

An update will be given on the status of collaborative tools for software developers, Version Control Services (Git and SVN), Issue Tracking (JIRA), Integration (Jenkins) and documentation (TWiki)

The presentation will focus on collaborative aspects for software developers and report on progress since the fall meeting.

Summary

An update will be given on the status of collaborative tools for software developers, Version Control Services (Git and SVN), Issue Tracking (JIRA), Integration (Jenkins) and documentation (TWiki)

Primary authors: LOSSENT, Alexandre (CERN); GONZALEZ ALVAREZ, Alvaro (CERN); APARICIO COTARELO, Borja (CERN); ANDERSEN, Terje (CERN)

Co-authors: HOIMYR, Nils (CERN); JONES, Pete (CERN)

Presenter: HOIMYR, Nils (CERN)

Session Classification: End-user Services and Operating Systems

Track Classification: End-User IT Services & Operating Systems

Contribution ID: 5

Type: **not specified**

Computer Security update

Tuesday, March 24, 2015 2:00 PM (25 minutes)

This presentation gives an overview of the current computer security landscape. It describes the main vectors of compromises in the academic community including lessons learnt, and reveal inner mechanisms of the underground economy to expose how our resources are exploited by organised crime groups, as well as recommendations to protect ourselves. By showing how these attacks are both sophisticated and profitable, the presentation concludes that the only mean to adopt and appropriate response is to build a tight international collaboration and trusted information sharing mechanisms within the community.

Primary author: Mr WARTEL, Romain (CERN)

Presenter: Mr WARTEL, Romain (CERN)

Session Classification: Security and Networking

Track Classification: Security & Networking

Contribution ID: 6

Type: **not specified**

Remote evaluation of hardware

Wednesday, March 25, 2015 2:25 PM (25 minutes)

The RHIC-ATLAS Computing Facility (RACF) at BNL has traditionally evaluated hardware on-site, with physical access to the systems. The effort to request evaluation hardware, shipping, set-up and testing has consumed an increasing amount of time and the process has become less productive over the years. To regain past productivity and shorten the evaluation process, BNL has started a pilot project to evaluate hardware remotely in 2015. This presentation discusses the status of the remote evaluation project and its future prospects.

Primary authors: Mr ZAYTSEV, Alexandr (Brookhaven National Laboratory (US)); HOLLOWELL, Christopher (Brookhaven National Laboratory); Dr WONG, Tony (Brookhaven National Laboratory); STRECKER-KELLOGG, William (Brookhaven National Lab)

Presenter: Dr WONG, Tony (Brookhaven National Laboratory)

Session Classification: Computing and Batch Systems

Track Classification: Computing & Batch Services

Contribution ID: 7

Type: **not specified**

Status of volunteer computing at CERN

Tuesday, March 24, 2015 11:40 AM (25 minutes)

Status of LHC@home, volunteer computing at CERN and for the LHC experiments. The presenter will give an update on the volunteer computing strategy for HEP and different scenarii for use of volunteer cloud computing or other lightweight cloud infrastructures to run experiment code under CernVM on available computing resources. Furthermore, the current status of the CERN BOINC server infrastructure and recent experience from ATLAS and CMS as well as the outlook for the service will be described.

Summary

Use of BOINC at CERN for LHC@home, Virtualization support with BOINC and CernVM allows for running HEP software under BOINC, and LHC@home will be extended to include applications from ATLAS, CMS and LHCb. A Description of BOINC application and server infrastructure at CERN is given.

Primary authors: HOIMYR, Nils (CERN); JONES, Pete (CERN); ASP, Tomi Juhani (University of Jyvaskyla (FI))

Co-authors: GONZALEZ ALVAREZ, Alvaro (CERN); FIELD, Laurence (CERN); Dr MARQUINA, Miguel (CERN)

Presenter: HOIMYR, Nils (CERN)

Session Classification: End-user Services and Operating Systems

Track Classification: End-User IT Services & Operating Systems

Contribution ID: 8

Type: **not specified**

DRMAA2 - An Open Standard for Job Submission and Cluster Monitoring

Thursday, March 26, 2015 12:20 PM (25 minutes)

- Introduction - DRMAA2 in a Nutshell
- The C Interface - Data Types, Monitoring Sessions, Job Sessions, Working with Jobs, Job Templates, Error Handling and Dealing with Enhancements
- Getting started with DRMAA2
- Example Applications - Job Monitoring Applications and Simple Multi-Clustering

Summary

Distributed Resource Management Application API 2 (DRMAA2) is an open standard defined by the Open Grid Forum. It standardizes a set of functions for job submission, job workflow management and cluster monitoring. This talk is about the design and functionality of DRMAA2 and its applications.

Daniel Gruber is Senior Software Engineer at Univa. Before joining Univa in 2011, he worked at Sun Microsystems in the Sun Grid Engine development team and later at Oracle in the Oracle Grid Engine development team. Daniel is co-chair of the DRMAA working group and helped Univa Grid Engine to be the first product to support the DRMAA2 standard.

Primary author: GRUBER, Daniel (U)

Presenter: GRUBER, Daniel (U)

Session Classification: Computing and Batch Systems

Track Classification: Computing & Batch Services

Contribution ID: 9

Type: **not specified**

Towards a modernisation of CERN's telephony infrastructure

Thursday, March 26, 2015 11:30 AM (25 minutes)

IP-based voice telephony (VoIP) and the SIP protocol are clear examples of disruptive technologies that have revolutionised a previously settled market. In particular, open-source solutions now have the ascendancy in the traditional Private Branch eXchange(PBX) market.

We present a possible architecture for the modernisation of CERN's fixed telephony network, highlighting the technical challenges to be addressed and the critical services that must be maintained and then describing how the introduction of open-source call routers based on the SIP protocol and Session Border Controllers (SBC) could foster the introduction of new services and increase the agility of our telephony network to adapt to new communication standards.

Primary author: VALENTIN VINAGRERO, Francisco (CERN)

Co-authors: HESNAUX, Anthony Gerard (Group ESAIP Enseignement Superieur (FR)); SIERRA, Rodrigo (CERN)

Presenter: VALENTIN VINAGRERO, Francisco (CERN)

Session Classification: Basic IT Services

Track Classification: Basic IT Services

Contribution ID: **10**

Type: **not specified**

BNL RACF Site Report

Monday, March 23, 2015 11:10 AM (15 minutes)

Brookhaven National Lab (BNL) will present the site report for the RHIC-ATLAS Computing Facility (RACF), covering developments over the past 6 months.

Primary author: STRECKER-KELLOGG, William (Brookhaven National Lab)

Presenter: STRECKER-KELLOGG, William (Brookhaven National Lab)

Session Classification: Site Reports

Track Classification: Site reports

Contribution ID: 11

Type: **not specified**

Effects of packet loss and delay on TCP performance

Tuesday, March 24, 2015 12:20 PM (25 minutes)

Following an incident with a slow database replication between CERN's data centers, we discovered that even a very low rate packet loss in the network (order of 0.001%) can induce significant penalties to long distance single stream TCP transfers. We explore the behaviour of multiple TCP congestion control algorithms in a controlled loss and delay environment in order to understand what is the achievable throughput of TCP data transfers between CERN's remote data centres.

Primary author: KRAJEWSKI, Adam Lukasz (Warsaw University of Technology (PL))

Presenter: KRAJEWSKI, Adam Lukasz (Warsaw University of Technology (PL))

Session Classification: Security and Networking

Track Classification: Security & Networking

Contribution ID: 12

Type: **not specified**

Configuration management at CERN: Status and directions

Thursday, March 26, 2015 9:25 AM (25 minutes)

CERN's experience of migrating a large site to a Puppet-based and more dynamic Configuration Service will be presented. The presentation will review some of the challenges encountered along the way and describe future plans for how to scale the service and improve the overall automation of operations on the site.

Primary author: RODRIGUEZ PEON, Alberto (Universidad de Oviedo (ES))

Co-author: MCCANCE, Gavin (CERN)

Presenter: RODRIGUEZ PEON, Alberto (Universidad de Oviedo (ES))

Session Classification: Basic IT Services

Track Classification: Basic IT Services

Contribution ID: 13

Type: **not specified**

CSC - It Center for Science

Monday, March 23, 2015 9:15 AM (15 minutes)

- CSC General HPC updates
- New Haswell Hardware for supercluster and supercomputer
- Slurm/Lustre
- taito-shell.csc.fi - a slurm/sshd/iptables-based interactive shell load balancer. Replaced large memory (1TB) interactive nodes.
- DDN SFA12k
- Using ELK stack (Elasticsearch Logstash Kibana)
- sort through dCache logs
- Search through auditd logs - anybody used shred? root logins from bad IPs?
- Updates at T2_FI_HIP - SE and CE for Helsinki Institute of Physics
- dCache to 2.10 / postgresql to 9.3 update
- To virtualize dCache door/admin and other nodes (bdii, xrootd AAA)
- Will use Ansible for configuration
- HP 5x DL360 G7 + 26 x D2600 has been very stable. (7 Seagate 2TB disks replaced over 3 years, two P411 controllers)
- Next hardware renewal hopefully during 2015, probably HP SL4540
- Decommission of jade.hip.fi - resources moved into one of CSC's openstack
- Converted unused IB port in dCache pool servers to 10GbE with QSFP-to-SFP+ adapters.

Summary

- CSC Updates
- Using ELK stack
- T2_FI_HIP updates

Primary author: GULDMYR, Johan Henrik (Helsinki Institute of Physics (FI))

Presenter: GULDMYR, Johan Henrik (Helsinki Institute of Physics (FI))

Session Classification: Site Reports

Track Classification: Site reports

Contribution ID: 14

Type: **not specified**

Ceph operations at CERN

Wednesday, March 25, 2015 11:30 AM (25 minutes)

Ceph has become over time a key component of CERN's Agile Infrastructure by providing storage for the Openstack service.

In this talk, we will briefly introduce Ceph's concepts, our current cluster and the services we provide such as NFS filers, Object Store for the Atlas experiment and Xroot-to-Ceph gateways. We will then talk about our experience running Ceph with some real-world examples (everyday operations, lessons learned, upgrades)

Finally we will try to have a look at what is coming for Ceph at CERN and how it could be used to secure data to a remote place or to consolidate storage (Castor).

Primary author: ROUSSEAU, Herve (CERN)

Presenter: ROUSSEAU, Herve (CERN)

Session Classification: Storage and File Systems

Track Classification: Storage & Filesystems

Contribution ID: 15

Type: **not specified**

Evolutions in the CERN Conferencing Services Landscape

Monday, March 23, 2015 3:15 PM (25 minutes)

A lot of visible and behind-the-scene actions have been taken in recent months to prepare CERN conferencing services (Indico, Vidyo, the webcast and conference rooms services) for challenges to come. These services will be described in terms of features and usage statistics. We will present their integration to the CERN layered cloud infrastructure, and with other IT base services. We will analyse the feedback gathered through recent user satisfaction surveys and describe the orientation taken for their development, including the challenges to open some of these services to communities outside CERN. Some security aspects will also be highlighted as well as the importance of cross-service integrations.

Summary

A description of recent evolutions in the CERN conferencing services, from the usage, features, security and management points of view.

Primary author: BARON, Thomas (CERN)

Presenter: BARON, Thomas (CERN)

Session Classification: End-user Services and Operating Systems

Track Classification: End-User IT Services & Operating Systems

Contribution ID: 16

Type: **not specified**

CERN Cloud Report

Thursday, March 26, 2015 2:25 PM (25 minutes)

This is a report on the current status and future plans of CERN's OpenStack-based Cloud Infrastructure.

Primary author: BOMPASTOR, Bruno (CERN)

Presenter: BOMPASTOR, Bruno (CERN)

Session Classification: Grids, Clouds and Virtualisation

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: 17

Type: **not specified**

The Vacuum model for running jobs in VMs

Thursday, March 26, 2015 3:15 PM (25 minutes)

The Vacuum model provides a method for managing the lifecycle of virtual machines based on their observed success or failure in finding work to do for their experiment. In contrast to centrally managed grid job submission and cloud VM instantiation systems, the Vacuum model gives resource providers direct control over which experiments' VMs or jobs are created and in what proportion. This model also leads to a highly decentralised, feedback-based infrastructure, in which the responsibility of providing VMs for the same experiment can be undertaken by a mix of sites, local groups, national teams, and the experiment's central operations staff. This mixture is well matched to the variety of entities which need to act as a virtual resource provider, due to differences in available funding and objectives. We present three implementations of this model developed by GridPP: Vac which manages VMs on autonomous hypervisor machines; Vcycle which manages VMs on IaaS systems such as OpenStack; and a system for managing VMs at an HTCondor site.

The Pilot VM architecture originally developed by LHCb is particularly suitable to the lightweight Vacuum approach, and involves the configuration of an environment similar to conventional WLCG batch worker nodes within the VM, and then the execution of the pilot framework client as with conventional grid jobs. We present Pilot VM designs used to run production jobs for LHCb, ATLAS, CMS, and GridPP DIRAC, where the same VM design is able to run on all three implementations of the Vacuum model for VM management.

Primary author: MCNAB, Andrew (University of Manchester (GB))

Co-author: LAHIFF, Andrew David (STFC - Rutherford Appleton Lab. (GB))

Presenter: MCNAB, Andrew (University of Manchester (GB))

Session Classification: Grids, Clouds and Virtualisation

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: **18**

Type: **not specified**

INFN-T1 Site report

Monday, March 23, 2015 10:15 AM (15 minutes)

Update on INFN-T1

Primary author: CHIERICI, Andrea (INFN-CNAF)

Co-author: MISURELLI, Giuseppe (Unknown)

Presenter: MISURELLI, Giuseppe (Unknown)

Session Classification: Site Reports

Track Classification: Site reports

Contribution ID: 19

Type: **not specified**

Batch Processing at CERN

Friday, March 27, 2015 9:00 AM (25 minutes)

The CERN Batch System comprises 4000 worker nodes, 60 queues and offers a service for various types of large user communities. In light of the developments driven by the Agile Infrastructure and the more demanding processing requirements, it is faced with increasingly challenging scalability and flexibility needs.

This production cluster currently runs IBM/Platform LSF. Over the last few months, an increasing number of large-scale interventions had to take place, betraying some critical limitations we will need to overcome in the future. We have started working on a project helping us implementing work flows to help use face these problems.

Primary author: BELLEMAN, Jerome (CERN)

Presenter: BELLEMAN, Jerome (CERN)

Session Classification: Computing and Batch Systems

Track Classification: Computing & Batch Services

Contribution ID: 20

Type: **not specified**

Running ATLAS at scale on Amazon EC2

Thursday, March 26, 2015 4:05 PM (25 minutes)

Beginning in September 2014, the RACF at Brookhaven National Lab has been collaborating with Amazon's scientific computing group in a pilot project. The goal of this project is to demonstrate the usage of Amazon AWS (EC2, S3, etc.) for real-world ATLAS production. This will prove the practical and economic feasibility of ATLAS beginning to leverage commercial cloud computing to optimize resource provisioning. This project includes all elements needed for at-scale work: networking, data management, capacity provisioning, VM image creation and management, and alterations/additions to ATLAS compute infrastructure to leverage AWS. Each of these elements is established for all three AWS regions (east, west-1, and west-2). The work done so far, the results of initial scaling tests, and plans for the near future will be presented in detail.

Summary

An overview of an ATLAS pilot project to demonstrate full-scale usage of Amazon AWS for scientific data analysis for ATLAS.

Primary author: HOVER, John (Brookhaven National Laboratory (BNL)-Unknown-Unknown)

Presenter: HOVER, John (Brookhaven National Laboratory (BNL)-Unknown-Unknown)

Session Classification: Grids, Clouds and Virtualisation

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: **21**

Type: **not specified**

Nikhef site-report

Monday, March 23, 2015 9:30 AM (15 minutes)

Spring 2015 site report

Primary author: KUIPERS, Paul (Nikhef)

Presenter: KUIPERS, Paul (Nikhef)

Session Classification: Site Reports

Track Classification: Site reports

Contribution ID: 22

Type: **not specified**

OpenStack Heat @ CERN

Thursday, March 26, 2015 4:55 PM (25 minutes)

Heat, the Openstack orchestration service, is being deployed at CERN. We will be presenting the overall architecture and features included in the project, our deployment challenges and future plans.

Primary author: BOMPASTOR, Bruno (CERN)

Presenter: BOMPASTOR, Bruno (CERN)

Session Classification: Grids, Clouds and Virtualisation

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: 23

Type: **not specified**

Status Report on Ceph Based Storage Systems at the RACF

Wednesday, March 25, 2015 11:55 AM (25 minutes)

We review various functionality, performance, and stability tests performed at the RHIC and ATLAS Computing Facility (RACF) at Brookhaven National Laboratory (BNL) in 2014-2015. Tests were run on all three (object storage, block storage and file system) levels of Ceph, using a range of hardware platforms and networking solutions, including 10/40 Gbps Ethernet and IPoIB/4X FDR Infiniband. We also report the current status of a 1 PB scale Ceph-based object storage system, provided with Amazon S3 compliant RADOS Gateway interfaces, that was built for the RACF in 2014 and went through a major hardware upgrade in the beginning of 2015. We present performance measurements of this system and discuss the experience gained while operating it as ATLAS event service storage for the past 8 months.

Primary authors: Mr ZAYTSEV, Alexandr (Brookhaven National Laboratory (US)); ITO, Hironori (Brookhaven National Laboratory (US)); Dr RIND, Ofer (BROOKHAVEN NATIONAL LABORATORY)

Co-authors: WONG, Antonio (Brookhaven National Laboratory (US)); HOLLOWELL, Christopher (Brookhaven National Laboratory); Mr RAO, Tejas (Brookhaven National Laboratory (US))

Presenter: Dr RIND, Ofer (BROOKHAVEN NATIONAL LABORATORY)

Session Classification: Storage and File Systems

Track Classification: Storage & Filesystems

Contribution ID: 24

Type: **not specified**

Subtlenoise: sonification of distributed computing activity

Thursday, March 26, 2015 11:05 AM (25 minutes)

The dominant monitoring system used in distributed computing consists of visually rich time-series graphs and notification systems for alerting operators when metrics fall outside of accepted values. For large systems this can quickly become overwhelming. In this contribution a different approach is described using the sonification of monitoring messages with an architecture which fits easily within existing infrastructures. The benefits of this approach are described in the context of various computing operations.

Primary author: LOVE, Peter (Lancaster University (GB))

Presenter: LOVE, Peter (Lancaster University (GB))

Session Classification: Basic IT Services

Track Classification: Basic IT Services

Contribution ID: 25

Type: **not specified**

Operation and benchmarking of a commercial datacentre

Wednesday, March 25, 2015 2:00 PM (25 minutes)

This contribution describes the usage and benchmarking of a commercial data centre running Openstack. Different cloud provisional tools are described highlighting the pros and cons of each system. A comparison is made between this facility and a standard grid T2 site in terms of job throughput and availability. Usage of the centre's local object store is also described.

Primary author: LOVE, Peter (Lancaster University (GB))

Presenter: LOVE, Peter (Lancaster University (GB))

Session Classification: Computing and Batch Systems

Track Classification: Computing & Batch Services

Contribution ID: 26

Type: **not specified**

Status of Centralized Config Management at the RACF

Thursday, March 26, 2015 9:00 AM (25 minutes)

It's simple enough to instantiate a new process in an existing environment; it can be much more challenging to foster acceptance of such a process in IT environments and cultures that are traditionally stagnant and resistant to change, and to maintain and optimize that process to ensure it continues to realize optimal benefit. To enhance our computing facility, we've already taken considerable strides toward simplifying, optimizing, and automating our technical deployment and maintenance procedures by analyzing, adopting, and implementing policies and technologies. As our facility evolves, so does technology change: we continuously revisit our environment and needs, evaluate our current tools and processes, and watch the horizon of the IT landscape for new and more optimal technology and solutions.

Our configuration management core relies in part upon Puppet: Puppet Labs has developed a new server model, which we have tested and evaluated against our existing Puppet deployment. We've developed an automated testing process based upon Jenkins CI, a continuous integration tool that validates pending changes before they can be pushed into our production environment. We've begun evaluating MCollective, an orchestration framework that may prove useful to our current automation processes by adding functionality such as resource grouping and reporting. We're working with other organizations at our site that share our interest in configuration management to share ideas and refine solutions.

In this talk, we present an overview of the current state of configuration management environment in our facility, the technical challenges we currently face, the technology we're evaluating and using to address those challenges, and the direction in which we plan to steer our future efforts.

Summary

An overview of our Puppet deployment at the RACF.

Primary authors: SMITH, Jason Alexander (Brookhaven National Laboratory (US)); STRECKER-KELLOGG, William (Brookhaven National Lab)

Co-author: DE STEFANO JR, John Steven (Brookhaven National Laboratory (US))

Presenter: STRECKER-KELLOGG, William (Brookhaven National Lab)

Session Classification: Basic IT Services

Track Classification: Basic IT Services

Contribution ID: 27

Type: **not specified**

Future of Batch Processing at CERN

Friday, March 27, 2015 11:30 AM (25 minutes)

While we are taking measures to face the limitations discussed earlier on in our IBM/Platform LSF cluster, we have been working on setting up a new batch system based on HTCondor. There has been some progress with the pilot service which we described last HEPiX. We also went on investigating some of the more advanced functions which will lead up to the production state of the new CERN Batch Service.

Primary author: BELLEMAN, Jerome (CERN)

Presenter: BELLEMAN, Jerome (CERN)

Session Classification: Computing and Batch Systems

Track Classification: Computing & Batch Services

Contribution ID: 28

Type: **not specified**

Site Report - Prague

Monday, March 23, 2015 10:00 AM (15 minutes)

We will give an overview of the site and will share experience with these topics: migration of virtualized servers to a new infrastructure, migration from cfengine to puppet and spacewalk as new systems management solution, procurement of a new hardware (worker nodes and storage servers).

Summary

Site report

Primary authors: ADAMOVA, Dagmar (Acad. of Sciences of the Czech Rep. (CZ)); CHUDOBA, Jiri (Acad. of Sciences of the Czech Rep. (CZ)); ADAM, Martin (UJF Rez); ŘÍKAL, Václav

Co-authors: SVEC, Jan (Acad. of Sciences of the Czech Rep. (CZ)); SOBOTKA, Zdenek (Czech Technical University (CZ))

Presenters: ADAM, Martin (UJF Rez); ŘÍKAL, Václav

Session Classification: Site Reports

Track Classification: Site reports

Contribution ID: 29

Type: **not specified**

PIC Tier-1 Spring 2015 report

Tuesday, March 24, 2015 9:00 AM (15 minutes)

We will be revising the status of PIC Tier-1 by Spring 2015. The typical site report which is reported in HEPiX.

Primary author: FLIX MOLINA, Jose (Centro de Investigaciones Energ. Medioambientales y Tecn. - (ES)

Presenter: FLIX MOLINA, Jose (Centro de Investigaciones Energ. Medioambientales y Tecn. - (ES)

Session Classification: Site Reports

Track Classification: Site reports

Contribution ID: **30**

Type: **not specified**

WLCG Cloud Traceability Working Group

Tuesday, March 24, 2015 2:25 PM (25 minutes)

Report on the initial activities of the WLCG Cloud Traceability Working Group

Primary author: COLLIER, Ian Peter (STFC - Rutherford Appleton Lab. (GB))

Co-author: Mr WARTEL, Romain (CERN)

Presenter: COLLIER, Ian Peter (STFC - Rutherford Appleton Lab. (GB))

Session Classification: Security and Networking

Track Classification: Security & Networking

Contribution ID: 31

Type: **not specified**

Two years of HTCondor at the RAL Tier-1

Friday, March 27, 2015 11:05 AM (25 minutes)

After running Torque/Maui for many years, the RAL Tier-1 migrated to HTCondor during 2013 in order to benefit from improved reliability, scalability and additional functionality unavailable in Torque. This talk will discuss the deployment of HTCondor at RAL, our experiences and the evolution of our pool over the past two years, as well as our future plans.

Primary author: LAHIFF, Andrew David (STFC - Rutherford Appleton Lab. (GB))

Co-authors: DEWHURST, Alastair (STFC - Rutherford Appleton Lab. (GB)); COLLIER, Ian Peter (STFC - Rutherford Appleton Lab. (GB)); KELLY, John (STFC - Rutherford Appleton Lab. (GB))

Presenter: LAHIFF, Andrew David (STFC - Rutherford Appleton Lab. (GB))

Session Classification: Computing and Batch Systems

Track Classification: Computing & Batch Services

Contribution ID: 32

Type: **not specified**

HTCondor within the European Grid and in the Cloud

Friday, March 27, 2015 11:55 AM (25 minutes)

With the increasing interest in HTCondor in Europe, an important question for sites considering migrating to HTCondor is how well it integrates with the standard grid middleware, in particular integration with the information system and APEL accounting. Also, with the increasing interest and usage of private clouds, how easily a batch system can be integrated with a private cloud is another important question. This talk will discuss both of these topics based on the experience of the RAL Tier-1.

Primary author: LAHIFF, Andrew David (STFC - Rutherford Appleton Lab. (GB))

Presenter: LAHIFF, Andrew David (STFC - Rutherford Appleton Lab. (GB))

Session Classification: Computing and Batch Systems

Track Classification: Computing & Batch Services

Contribution ID: 33

Type: **not specified**

STAR Experience with Automated High Efficiency Grid Based Data Production Framework at KISTI/Korea

Thursday, March 26, 2015 5:20 PM (25 minutes)

In statistically hungry science domains, data taking data deluges can be both a blessing and a curse. They allow the winnowing out of statistical errors from known measurements, open the door to new scientific opportunities as the physics program matures but are also a testament to the efficiency of the experiment and accelerator and skill of its operators. However, the data samples need to be dealt with and in experiments like those at RHIC, the planning for computer resources do not allow huge increases in computing capacity. A standard strategy has then been to share resources across multiple experiments at a given facility. Another has been to use middleware that “glues” resources across the world so they are able to locally run the experimental software stack (either natively or virtually).

In this presentation, we will describe a framework STAR has successfully used to reconstruct a ~400TB of data consisting of over 100,000 jobs submitted to a remote site in Korea from its Tier0 facility at the Brookhaven National Laboratory. The framework automates the full path taking raw data files from tape and writing Physics ready output back to tape without operator or remote site intervention. Through hardening we have demonstrated an efficiency of 99%, over a period of 7 months of operation. The high performance is attributed to finite state checking with retries to encourage resilience in the system over capricious and fallible infrastructure.

Primary authors: Dr LAURET, Jerome (BROOKHAVEN NATIONAL LABORATORY); Mr HAJDU, Levente (Brookhaven National Laboratory); DIDENKO, Lidia (Brookhaven National Laboratory)

Presenter: Mr HAJDU, Levente (Brookhaven National Laboratory)

Session Classification: Grids, Clouds and Virtualisation

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: 34

Type: **not specified**

News from the HEPiX IPv6 Working Group

Tuesday, March 24, 2015 4:15 PM (25 minutes)

This talk will present an update from the HEPiX IPv6 Working Group. This will include details of recent testing activities and plans for the deployment of dual-stack data services and monitoring on (at least some of) the WLCG infrastructure.

Primary author: KELSEY, Dave (STFC - Rutherford Appleton Lab. (GB))

Presenter: KELSEY, Dave (STFC - Rutherford Appleton Lab. (GB))

Session Classification: Security and Networking

Track Classification: Security & Networking

Contribution ID: 35

Type: **not specified**

Recent experiences in operational security: incident prevention and incident handling in the EGI and WLCG infrastructure

Tuesday, March 24, 2015 2:50 PM (25 minutes)

The European Grid Infrastructure (EGI) and Worldwide Large Hadron collider Grid (WLCG) infrastructure largely overlap and share the majority of security activities. A lot of security related activity goes on behind the scenes concerning such a large scale distributed computing infrastructure. Security incident prevention takes up the larger amount of effort, and this is carried out via Security Policy definition, Software vulnerability handling and monitoring to ensure the infrastructure is as secure as is practical. This will describe some of the recent experiences concerning security, from handling vulnerabilities (including some of the high profile ones) to the types of incidents which have occurred in recent times. This will also describe some of the changes needed to handle new technologies and trust models which are coming into use.

Primary author: CORNWALL, Linda Ann (STFC - Rutherford Appleton Lab. (GB))

Presenter: CORNWALL, Linda Ann (STFC - Rutherford Appleton Lab. (GB))

Session Classification: Security and Networking

Track Classification: Security & Networking

Contribution ID: 36

Type: **not specified**

Looking for a fast benchmark

Wednesday, March 25, 2015 4:05 PM (25 minutes)

The WLCG community has requested a fast benchmark to quickly assess the performances of a worker node. A good candidate is a python script used in LHCb

Primary author: Dr MICHELOTTO, Michele (INFN Padua & CMS)

Co-author: ALEF, Manfred (Karlsruhe Institute of Technology (KIT))

Presenter: Dr MICHELOTTO, Michele (INFN Padua & CMS)

Session Classification: Computing and Batch Systems

Track Classification: Computing & Batch Services

Contribution ID: 37

Type: **not specified**

Testing dCache and IPv6

Tuesday, March 24, 2015 4:40 PM (15 minutes)

A view back on testing IPv6 and different versions of dCache as it has evolved from 2.6 to 2.12 and barely-working to well working.

Summary

A few years back I got tasked to testing dCache and IPv6 since “somebody” needed to do it and noone else was brave enough to try it.

Every version of dCache has then improved significantly for each new release, reaching maturity at the brand new 2.12.

Primary author: TIGERSTEDT, Ulf Bobson Severin (Helsinki Institute of Physics (FI))

Presenter: TIGERSTEDT, Ulf Bobson Severin (Helsinki Institute of Physics (FI))

Session Classification: Security and Networking

Track Classification: Security & Networking

Contribution ID: 38

Type: **not specified**

Evaluation of Memory and CPU usage via Cgroups of ATLAS workloads running at a Tier-2

Wednesday, March 25, 2015 2:50 PM (25 minutes)

Modern Linux Kernels include a feature set that enables the control and monitoring of system resources, called Cgroups. Cgroups have been enabled on a production HTCondor pool sited at the Glasgow site of the UKI-SCOTGRID distributed Tier-2. A system has been put in place to collect and aggregate metrics extracted from Cgroups on all worker nodes within the Condor pool. From this aggregated data, memory and CPU usage footprints are extracted.

From the extracted footprints the resource usage for each type of ATLAS workload can be obtained and studied. This system has been used to identify broken payloads, real-world memory usage, job efficiencies etc. Additionally work has begun on near-real-time tracking of running jobs with a goal to proactively identify and stop broken payloads from consuming unnecessary CPU time and resources.

Primary authors: QIN, Gang (University of Glasgow (GB)); ROY, Gareth (University of Glasgow)

Co-authors: Prof. BRITTON, David (University of Glasgow (GB)); CROOKS, David (University of Glasgow (GB)); Dr STEWART, Gordon (University of Glasgow); Dr SKIPSEY, Samuel Cadellin

Presenter: QIN, Gang (University of Glasgow (GB))

Session Classification: Computing and Batch Systems

Track Classification: Computing & Batch Services

Contribution ID: 39

Type: **not specified**

A recent view of OSSEC and Elasticsearch at Scotgrid Glasgow

Tuesday, March 24, 2015 3:15 PM (25 minutes)

OSSEC, the popular HIDS (Host Intrusion Detection System), has been widely used for a number of years. More recently, tools like Elasticsearch, Logstash and Kibana (ELK) have become popular in visualising and working with data such as that aggregated by OSSEC. We report on a recent implementation of OSSEC, coupled to an ELK instance, at the Glasgow site of the UKI-SCOTGRID distributed Tier-2. In particular, we report on our experience of the installation and use of these tools in a puppet deployment context. We cover installation, additional utility scripts deployed as well as the configuration workflow. We broadly cover the specific Grid related rules that have been implemented thus far. This presentation is particularly relevant for sysadmins and security officers interested in a recent view of the installation of this software and our experience with it.

Primary author: CROOKS, David (University of Glasgow (GB))

Co-authors: Prof. BRITTON, David (University of Glasgow (GB)); QIN, Gang (University of Glasgow (GB)); ROY, Gareth Douglas (University of Glasgow (GB)); Dr STEWART, Gordon (University of Glasgow); Dr SKIPSEY, Samuel Cadellin

Presenter: CROOKS, David (University of Glasgow (GB))

Session Classification: Security and Networking

Track Classification: Security & Networking

Contribution ID: 40

Type: **not specified**

Oxford University Site Report

Tuesday, March 24, 2015 9:15 AM (15 minutes)

A site report from the University of Oxford focusing on the integration challenges between the various systems.

Summary

A site report from the University of Oxford.

Primary author: Dr BRISBANE, Sean (University of Oxford)

Presenter: Dr BRISBANE, Sean (University of Oxford)

Session Classification: Site Reports

Track Classification: Site reports

Contribution ID: 41

Type: **not specified**

Ceph development update

Wednesday, March 25, 2015 12:20 PM (25 minutes)

The Ceph storage system is an open source, highly scalable, resilient data storage service providing object, block and file interfaces. This presentation will introduce what is new in the latest Ceph release, codenamed *Hammer*, and describe the ongoing development activities around CephFS, the Ceph filesystem.

An intermediate level of familiarity with large scale storage systems will be assumed.

Primary author: Mr JOHN, Spray (Red Hat, Inc.)

Presenter: Mr JOHN, Spray (Red Hat, Inc.)

Session Classification: Storage and File Systems

Track Classification: Storage & Filesystems

Contribution ID: 42

Type: **not specified**

Condor Workshop Summary

Friday, March 27, 2015 10:15 AM (25 minutes)

I propose to give a summary of the Condor workshop, held at CERN mid-December.

Primary author: Mr JOUVIN, Michel (Laboratoire de l'Accelérateur Lineaire (FR))

Presenter: Mr JOUVIN, Michel (Laboratoire de l'Accelérateur Lineaire (FR))

Session Classification: Computing and Batch Systems

Track Classification: Computing & Batch Services

Contribution ID: 43

Type: **not specified**

Update on WLCG/OSG perfSONAR Infrastructure

Tuesday, March 24, 2015 3:50 PM (25 minutes)

WLCG relies on the network as a critical part of its infrastructure and therefore needs to guarantee effective network usage and prompt detection and resolution of any network issues, including connection failures, congestion and traffic routing. The WLCG Network and Transfer Metrics working group was established to ensure sites and experiments can better understand and fix networking issues. In addition, it aims to integrate and combine all network-related monitoring data collected by the WLCG infrastructure from both network and transfer systems. This has been facilitated by the already existing network of the perfSONAR instances that is being commissioned to operate in full production.

Recently, several higher level services were developed to help bring perfSONAR network to its full potential. This includes a Web-based mesh configuration system, which allows to centrally schedule and manage all the network tests performed by the instances; a network datastore (esmond), which collects, stores and provides interfaces to access all the network monitoring information from a single place as well as perfSONAR infrastructure monitoring, which ensures that the current perfSONAR instances are configured and operated correctly.

In this presentation we will provide an update on the status of perfSONAR in WLCG and OSG as well as highlight the future plans for the commissioning and production operations of the perfSONAR in the scope of the WLCG Network and Transfer Metrics working group. We will also present details on the specific higher level services and provide a summary of the current status of the pilot projects to integrate network and transfer metrics.

Summary

An overview of perfSONAR status and plans in WLCG and OSG.

Primary author: BABIK, Marian (CERN)

Co-author: Dr MCKEE, Shawn (University of Michigan ATLAS Group)

Presenter: BABIK, Marian (CERN)

Session Classification: Security and Networking

Track Classification: Security & Networking

Contribution ID: 44

Type: **not specified**

Dust sensors for long term data preservation

Wednesday, March 25, 2015 9:00 AM (25 minutes)

CERN Computer Center (CC) is a large building that integrates several kilometers of fibers, copper cables, pipes and several complex installations (UPSes, water cooling, heat exchangers...). This evolving building is a large theater with numerous actors:

- contractors, performing construction work, building maintenance or hardware replacement
- engineers and technicians, debugging hardware and performing critical upgrades on the servers
- visitors, touring inside the computing facilities

All those external activities are not without consequences on the hosted IT equipment.

Last year, the CERN CC tape infrastructure was impacted by an isolated air contamination incident which affected around 125 files on a dozen tapes and two drives.

Since long term data preservation is one of our missions we need to take this new environmental parameter into account.

This presentation will expose the problems and challenges we are facing, and more importantly the solutions we developed to better monitor CERN CC environment around our tape libraries and strategies to limit the impact of airborne particles on the LHC data.

Primary author: LEDUC, Julien (CERN)

Presenter: LEDUC, Julien (CERN)

Session Classification: IT Facilities and Business Continuity

Track Classification: IT Facilities & Business Continuity

Contribution ID: 45

Type: **not specified**

Grid Engine at GridKa

Friday, March 27, 2015 9:25 AM (25 minutes)

The Grid Computing Centre Karlsruhe (GridKa) is using the Grid Engine batch system since 2011. In this presentation I will talk about the experiences with this batch system, including multi-core job support, and first experiences with cgroups.

Primary author: ALEF, Manfred (Karlsruhe Institute of Technology (KIT))

Presenter: ALEF, Manfred (Karlsruhe Institute of Technology (KIT))

Session Classification: Computing and Batch Systems

Track Classification: Computing & Batch Services

Contribution ID: 46

Type: **not specified**

Beyond HS06: Toward a New HEP CPU Benchmark - Update March 2015

Wednesday, March 25, 2015 3:15 PM (25 minutes)

In this talk we will provide information about the current status of the preliminary work to re-launch the HEPiX Benchmarking Working Group which will develop the next release of the HEP CPU benchmark.

Primary author: ALEF, Manfred (Karlsruhe Institute of Technology (KIT))

Co-author: MICHELOTTO, Michele (Universita e INFN (IT))

Presenter: ALEF, Manfred (Karlsruhe Institute of Technology (KIT))

Session Classification: Computing and Batch Systems

Track Classification: Computing & Batch Services

Contribution ID: 47

Type: **not specified**

Quattor in 2015

Thursday, March 26, 2015 10:15 AM (25 minutes)

The Quattor community has been maintaining Quattor for over ten years and having recently held our 19th community workshop the pace of development continues to increase.

This talk will demonstrate why Quattor is more than just a configuration management system, report on recent developments and provide some notable updates and experiences from sites.

Primary author: ADAMS, James (STFC RAL)

Presenter: ADAMS, James (STFC RAL)

Session Classification: Basic IT Services

Track Classification: Basic IT Services

Contribution ID: 48

Type: **not specified**

The IPV6 post office: labeling and sorting everywhere.

Tuesday, March 24, 2015 4:55 PM (25 minutes)

Probably the most prominent change that IPv6 introduces in the semantics of internet protocol applications is the need to *always* deal with multiple addresses (possibly both IPv4 and IPv6) associated to each network endpoint. A quick overview of how and where addresses are categorised, ordered and preferred is presented, both from the system administrator and the developer viewpoint. A few not too obvious practical consequences of RFC3484 are also shown, in an attempt to tame its subtleties.

Primary author: PRELZ, Francesco (Università degli Studi e INFN Milano (IT))

Co-authors: DEWHURST, Alastair (STFC - Rutherford Appleton Lab. (GB)); Dr SCIABA, Andrea (CERN); HOEFT, Bruno Heinrich (KIT - Karlsruhe Institute of Technology (DE)); WALKER, Christopher John (University of London (GB)); GRIGORAS, Costin (CERN); KELSEY, Dave (STFC - Rutherford Appleton Lab. (GB)); RAND, Duncan (Imperial College Sci., Tech. & Med. (GB)); MARTELLI, Edoardo (CERN); LOPEZ MUNOZ, Fernando (Universitat Autònoma de Barcelona (ES)); BERNIER, Jerome (IN2P3 Computing Center); CHUDOBA, Jiri (Acad. of Sciences of the Czech Rep. (CZ)); OHRENBERG, Kars (DESY); Dr CHADWICK, Keith (Fermilab); ELIAS, Marek (Acad. of Sciences of the Czech Rep. (CZ)); NANDAKUMAR, Raja (STFC - Rutherford Appleton Lab. (GB)); VOICU, Ramiro (California Institute of Technology (US)); FAYER, Simon (Imperial College Sci., Tech. & Med. (GB)); Dr CAMPANA, Simone (CERN); FINNERN, Thomas (DESY); Mr IDICULLA, Tiju (STFC - Rutherford Appleton Laboratory); Dr WILDISH, Tony (Princeton University (US)); TIGERSTEDT, Ulf (CSC Oy)

Presenter: PRELZ, Francesco (Università degli Studi e INFN Milano (IT))

Session Classification: Security and Networking

Track Classification: Security & Networking

Contribution ID: 49

Type: **not specified**

IHEP Site Report

Monday, March 23, 2015 9:45 AM (15 minutes)

The status of IHEP site and the improvement we've got and what is in our plan this year.

Primary author: SHI, Jingyan (IHEP)

Presenter: SHI, Jingyan (IHEP)

Session Classification: Site Reports

Track Classification: Site reports

Contribution ID: 50

Type: **not specified**

T2_US_Nebraska Site Report

Monday, March 23, 2015 11:40 AM (15 minutes)

Site report covering the status of T2_US_Nebraska and changes / updates since the Fall 2014 meeting.

Primary author: ATTEBURY, Garhan (University of Nebraska (US))

Presenter: ATTEBURY, Garhan (University of Nebraska (US))

Session Classification: Site Reports

Track Classification: Site reports

Contribution ID: 51

Type: **not specified**

HEP Software Foundation

Monday, March 23, 2015 2:25 PM (25 minutes)

The HEP Software Foundation (HSF) is a one year old initiative to foster collaboration in software development in the HEP community and related scientific communities. Launched by a kick-off meeting at CERN in April 2014, the first year has been spent to better define what HSF should be. An HSF workshop was held in January at SLAC and HSF is now entering its “implementation phase”. This talk will present the reasons behind HSF, the goals and organizational model.

Primary author: Mr JOUVIN, Michel (Laboratoire de l'Accelérateur Lineaire (FR))

Presenter: Mr JOUVIN, Michel (Laboratoire de l'Accelérateur Lineaire (FR))

Session Classification: End-user Services and Operating Systems

Track Classification: End-User IT Services & Operating Systems

Contribution ID: 52

Type: **not specified**

Updates from Database Services at CERN

Thursday, March 26, 2015 11:55 AM (25 minutes)

CERN has a great number of applications that rely on a database for their daily operations. From physics related databases to the administrative, sector there is a high demand to have a database system appropriate to the users' needs and requirements. This presentation gives a summary of the current state of the Database Services at CERN, the work done during LS1 and some insights into the evolution of our services.

Primary author: DUMITRU, Andrei (CERN)

Presenter: DUMITRU, Andrei (CERN)

Session Classification: Basic IT Services

Track Classification: Basic IT Services

Contribution ID: 53

Type: **not specified**

ASAP3: New data taking and analysis infrastructure for PETRA III

Wednesday, March 25, 2015 9:50 AM (25 minutes)

PETRA III is DESY's largest ring accelerator and the most brilliant storage-ring-based X-ray radiation source in the world.

With its recent extension, new and faster detectors are used for the data acquisition.

They exceed previous detectors in terms of data rate and volume; this is highly demanding for the underlying storage system.

This talk will present the challenges we faced, the new infrastructure and services based on IBM's GPFS and our first experiences with it.

Primary author: DIETRICH, Stefan (DESY)

Co-authors: Mrs KUHN, Manuela (DESY); Mr STRUTZ, Marco (DESY); Mr GASTHUBER, Martin (DESY); Mr ENSSLIN, Uwe (DESY)

Presenter: DIETRICH, Stefan (DESY)

Session Classification: Storage and File Systems

Track Classification: Storage & Filesystems

Contribution ID: 54

Type: **not specified**

Best Practices for Big Science in the Cloud

Thursday, March 26, 2015 4:30 PM (25 minutes)

On the heels of discussing the BNL RACF Group's Proof Of Concept on AWS, this session will share best practices on some of the most common AWS services used by Big Science, such as EC2, VPC, S3, and complex hybrid networking and routing. We will also provide an overview of the AWS Scientific Computing Group which was created to help Global Scientific collaborations develop and ecosystem supporting the long-tail of science, research and engineering and to educate these communities on the role AWS can play in making Science successful in the cloud.

Summary

On the heels of discussing the BNL RACF Group's Proof Of Concept on AWS, this session will share best practices on some of the most common AWS services used by Big Science, such as EC2, VPC, S3, and complex hybrid networking and routing. We will also provide an overview of the AWS Scientific Computing Group which was created to help Global Scientific collaborations develop and ecosystem supporting the long-tail of science, research and engineering and to educate these communities on the role AWS can play in making Science successful in the cloud.

Primary author: Mr RIVERA, Dario (Amazon Web Services)

Presenter: Mr RIVERA, Dario (Amazon Web Services)

Session Classification: Grids, Clouds and Virtualisation

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: 55

Type: **not specified**

DrainBoss: A Drain Rate Controller for ARC/HTCondor

Friday, March 27, 2015 12:20 PM (25 minutes)

This talk describes DrainBoss, which is a proportional integral (PI) controller with conditional logic that strives to maintain the correct ratio between single-core and multi-core jobs in an ARC/HTCondor cluster. DrainBoss can be used instead of the HTCondor DEFrag Daemon.

Primary author: JONES, Stephen (Liverpool University)

Presenter: JONES, Stephen (Liverpool University)

Session Classification: Computing and Batch Systems

Track Classification: Computing & Batch Services

Contribution ID: 56

Type: **not specified**

Site Report GSI

Monday, March 23, 2015 10:55 AM (15 minutes)

Site Report GSI

Primary author: SCHON, Walter

Presenter: SCHON, Walter

Session Classification: Site Reports

Track Classification: Site reports

Contribution ID: 57

Type: **not specified**

Evaluation of low power Systems on Chip for scientific computing

Wednesday, March 25, 2015 4:30 PM (25 minutes)

Systems on Chip (SoCs), originally targeted for mobile and embedded technology, are becoming attractive for HEP and HPC scientific communities, given their low cost, huge worldwide shipments, low power consumption and increasing processing power - mostly associated with their GPUs.

A variety of development boards are currently available, making it foreseeable to use these power-efficient platforms in a standard computing center configuration.

We discuss hardware limitations and programming constraints of low power solutions, namely ARM-based SoCs and Intel (Avoton) SoCs. Finally, we present some preliminary evaluations of the performances of SoC architectures running non-trivial scientific applications exploiting both the CPU and GPU available on the low power chip: results get particularly encouraging when considering workloads with low memory requirements and/or image processing.

Primary author: Dr MORGANTI, Lucia (INFN)

Co-authors: Dr FERRARO, Andrea (INFN); CESINI, Daniele (Universita e INFN, Roma I (IT)); Dr CALORE, Enrico (INFN); MICHELOTTO, Michele (Universita e INFN (IT)); Dr ALFIERI, Roberto (INFN); Dr SCHIFANO, fabio (INFN)

Presenter: Dr MORGANTI, Lucia (INFN)

Session Classification: Computing and Batch Systems

Track Classification: Computing & Batch Services

Contribution ID: 58

Type: **not specified**

CERN Search and Social for the Enterprise Web experience

Monday, March 23, 2015 2:50 PM (25 minutes)

 Status of CERN Web Services Overview Web Site Life Cycle Management Web Analytics CERN's Enterprise Social Networking System Motivation purpose Feature overview: microblogging, profiles, social networking, suggestion systems and discussion forums CERN Search Status Enterprise vs Web search New features

Summary

The CERN Web Services that provides a range of technologies and platforms to meet the various needs of CERN's user communities for effective and flexible Web content and Web application creation. We will summarise the status and evolution of the CERN Web Services, including Life cycle management of Web sites and Web site analytics. Furthermore, as two key elements for enhancing CERN's Enterprise Web Experience in order to facilitate information and knowledge creation and sharing, we will present our recently introduced Enterprise Social Networking platform and the evolution of our CERN enterprise Search solution.

CERN's Enterprise Social Networking (ESN) system provides the framework for effective and dynamic communication. ESN systems add value to an organization by encouraging information sharing, capturing knowledge, enabling action and empowering people. Another important aspect for the Web experience is the CERN enterprise search solution that provides access not only to public Web pages but furthermore to a wide range of protected information and CERN document repositories such as CDS, Indico, Twiki, SharePoint, JIRA and EDMS. We will present the recent and upcoming changes and enhancements to the CERN Search solution.

Primary authors: Mr WAGNER, Andreas (CERN); SILVA DE SOUSA, Bruno (CERN)

Presenter: Mr WAGNER, Andreas (CERN)

Session Classification: End-user Services and Operating Systems

Track Classification: End-User IT Services & Operating Systems

Contribution ID: 59

Type: **not specified**

PDSF Site Report and Relocation

Monday, March 23, 2015 11:25 AM (15 minutes)

PDSF, the Parallel Distributed Systems Facility, has been serving high energy physics and been in continuous operation at NERSC since 1996. It is currently a tier-1 site for Star, tier-2 for Alice and tier-3 for Atlas. This site report will describe recent updates to the system and upcoming modifications. PDSF will move this year from its current site to a new building on the LBNL campus and particular focus will given on how this will affect computing during the transition.

Primary author: BOTTTS, James (LBNL)

Co-authors: SAKREJDA, Iwona; GERHARDT, Lisa (LBNL)

Presenter: BOTTTS, James (LBNL)

Session Classification: Site Reports

Track Classification: Site reports

Contribution ID: 60

Type: **not specified**

Cloud @ RAL, an update

Thursday, March 26, 2015 2:00 PM (25 minutes)

The STFC Scientific computing department has been developing an OpenNebula based cloud underpinned by Ceph block storage. I will describe some of our use cases, our set up, and give a demonstration of our development VM on demand service. I will go on to explore some of the problems we have overcome to reach this point. Finally, I will present the work we are doing to use spare capacity on this cloud to run virtual worker nodes

Primary author: RYALL, George (STFC - Rutherford Appleton Lab.)

Co-authors: LAHIFF, Andrew David (STFC - Rutherford Appleton Lab. (GB)); ALEXANDER, Dibbo (STFC - RAL); BARNESLEY, Frazer (STFC - RAL); COLLIER, Ian Peter (STFC - Rutherford Appleton Lab. (GB))

Presenter: RYALL, George (STFC - Rutherford Appleton Lab.)

Session Classification: Grids, Clouds and Virtualisation

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: 61

Type: **not specified**

Building large storage systems with small units: How to make use of disks with integrated network and CPU

Wednesday, March 25, 2015 9:25 AM (25 minutes)

Recent advances in both, hard-disks and system-on-a-chip (SoC) designs enabled the development of a novel form of hard-disk: a disk that includes a network interface and an additional ARM processor, not involved in low level disk operations. This setup allows those disks to run an operating system and to communicate with other nodes autonomously using wired Ethernet. No additional hardware or infrastructure is required.

The HGST laboratory provided us, the dCache team, with early access to those devices. We investigated how such devices might be used in grid and cloud environments. By deploying dCache software and observing the system's behavior we evaluated how a petabyte-sized storage infrastructure, based on these disks, might be build, including possible changes to the dCache software.

We will present two realistic deployment scenarios for those new disks and compare them to existing deployments at the Deutsches Elektronen-Synchrotron (DESY) research centre, where direct attached RAID systems are in use. The results of our initial investigations are presented along with an outline of future work.

Primary author: KEMP, Yves (Deutsches Elektronen-Synchrotron (DE))

Presenter: KEMP, Yves (Deutsches Elektronen-Synchrotron (DE))

Session Classification: Storage and File Systems

Track Classification: Storage & Filesystems

Contribution ID: **62**

Type: **not specified**

BeeGFS at DESY

Wednesday, March 25, 2015 10:15 AM (25 minutes)

The presentation will present:

- History and current status of the BeeGFS project (formerly know as FhGFS, originating from Fraunhofer)
- Design and technology decisions made by BeeGFS developers
- BeeGFS setup and operational experience as InfiniBand based high-performance cluster file system serving as scratch space for the DESY HPC system
- Discussion of future usage scenarios and comparison with other products in use at DESY

Primary author: KEMP, Yves (Deutsches Elektronen-Synchrotron (DE))

Co-author: STERNBERGER, Sven (DESY)

Presenter: KEMP, Yves (Deutsches Elektronen-Synchrotron (DE))

Session Classification: Storage and File Systems

Track Classification: Storage & Filesystems

Contribution ID: **63**

Type: **not specified**

DLS site report

Tuesday, March 24, 2015 9:45 AM (15 minutes)

Diamond Light Source site report

Primary author: FRIEDRICH, Tina (Diamond Light Source Ltd)

Presenter: FRIEDRICH, Tina (Diamond Light Source Ltd)

Session Classification: Site Reports

Track Classification: Site reports

Contribution ID: 64

Type: **not specified**

KISTI-GSDC Site Report

Tuesday, March 24, 2015 10:00 AM (15 minutes)

The status of KISTI-GSDC Tier-1 site will be present including a brief of history of the KISTI-GSDC Site, system summary (configuration management), PBS batch issues, Tier-1 operations and future plan.

Primary author: AHN, Sang Un (KISTi Korea Institute of Science & Technology Information (KR))

Co-author: Mr KIM, Jin (KISTI Korea Institute of Science & Technology Information (KR))

Presenter: AHN, Sang Un (KISTi Korea Institute of Science & Technology Information (KR))

Session Classification: Site Reports

Track Classification: Site reports

Contribution ID: 65

Type: **not specified**

A look beyond x86: OpenPOWER8 & AArch64

Wednesday, March 25, 2015 4:55 PM (25 minutes)

x86 is the uncontested leader for server platforms in terms of market share and is currently the architecture of choice for High Energy Physics applications. But as more and more importance is given to power efficiency, physical density and total cost of ownership we are seeing new processor architectures emerging and some existing ones becoming more open. With the introduction of AArch64, ARM's 64-bit architecture, coupled with the adoption of industry standards such as UEFI, ACPI and SMBIOS, ARM has for the first time in history the opportunity of becoming a real contender in the server space. Sharing some similarities with the Open Computer Project and with a strong industry backing the OpenPOWER Foundation aims to create an open ecosystem around the enterprise centric POWER architecture. I will present the specificities of each of these alternative architectures highlighting the differentiating server features of each. Performance and power profiling of the following uni-socket platforms will be presented: Intel Xeon E3-1200 v3 (Haswell), Intel Atom (Avoton), Applied Micro X-Gene and OpenPOWER 8.

Primary author: VALSAN, Liviu (CERN)**Presenter:** VALSAN, Liviu (CERN)**Session Classification:** Computing and Batch Systems**Track Classification:** Computing & Batch Services

Contribution ID: **66**

Type: **not specified**

RAL Site Report

Tuesday, March 24, 2015 10:15 AM (15 minutes)

Latest updates for the RAL Tier-1.

Primary author: BLY, Martin (STFC-RAL)

Presenter: BLY, Martin (STFC-RAL)

Session Classification: Site Reports

Track Classification: Site reports

Contribution ID: 67

Type: **not specified**

Ceph storage at RAL

Wednesday, March 25, 2015 11:05 AM (25 minutes)

RAL is currently exploring the possibilities offered by Ceph. This talk will describe two of these projects. The first project aims to provide large scale, high throughput storage for experimental data. This will initially be used by the WLCG VOs. A prototype cluster built from old hardware has been in testing since October 2014. The WLCG VOs will continue to need to access their data via familiar protocols. The cluster is therefore accessed via xrootd and GridFTP gateways. We describe the current state of the project, the testing that has already been carried out as well as future plans for the cluster.

The second project provides Ceph RBD storage for the cloud infrastructure at RAL. This cluster is on hardware optimised for Ceph and provides low latency data access. This talk will describe the current status of the project as well as our experience running the cluster.

Primary author: DEWHURST, Alastair (STFC - Rutherford Appleton Lab. (GB))

Co-authors: CANNING, Bruno (RAL); RYALL, George (STFC); Mr VASILAKAKOS, George (STFC); BYRNE, Tom (STFC)

Presenter: DEWHURST, Alastair (STFC - Rutherford Appleton Lab. (GB))

Session Classification: Storage and File Systems

Track Classification: Storage & Filesystems

Contribution ID: 68

Type: **not specified**

Ceph vs Local Disk For Virtual Machines

Thursday, March 26, 2015 2:50 PM (25 minutes)

The Scientific Computing Department at the STFC has been developing a Ceph block storage backed OpenNebula cloud. We have carried out a quantitative evaluation of the performance characteristics of virtual machines which have been instantiated with a variety of different storage configurations (using both Ceph and local disks). I will describe our motivations for this testing, our methodology and present our initial results.

Primary author: DIBBO, Alexander (urn:Google)

Co-authors: LAHIFF, Andrew David (STFC - Rutherford Appleton Lab. (GB)); BARNSELEY, Frazer (STFC - RAL); RYALL, George (STFC); COLLIER, Ian Peter (STFC - Rutherford Appleton Lab. (GB))

Presenter: DIBBO, Alexander (urn:Google)

Session Classification: Grids, Clouds and Virtualisation

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: 69

Type: **not specified**

Deployment and usage of MCollective in production

Thursday, March 26, 2015 9:50 AM (25 minutes)

Marionette Collective, also known as MCollective, is a framework for building server orchestration, monitoring, and parallel job execution.

MCollective uses a modern “Publish Subscribe Middleware” for a scalable and fast execution environment.

It is a powerful tool in combination with Puppet, due to the good integration.

However it can be a challenging task to configure and deploy MCollective in a reliable and secure way.

In this talk, we will give an overview of our DESY MCollective setup, the design decisions, intended use and encountered problems.

Primary authors: ENGELS, Jan Dominik (Deutsches Elektronen-Synchrotron (DE)); DIETRICH, Stefan (DESY)

Co-author: STERNBERGER, Sven (DESY)

Presenter: DIETRICH, Stefan (DESY)

Session Classification: Basic IT Services

Track Classification: Basic IT Services

Contribution ID: 70

Type: **not specified**

Getting the most from the farm at the Sanger Institute

Monday, March 23, 2015 4:55 PM (25 minutes)

The Wellcome Trust Sanger Institute is a charitably funded genomic research centre. A leader in the Human Genome Project, it is now focused on understanding the role of genetics in health and disease. Large amounts of data is produced at the institute by next-generation sequencing machines. The data is then stored, processed and analysed on the institute's computing cluster. The main compute farm has 14,000 cores and has of the order of 20PB of storage in a mix of NFS, lustre and iRODS.

Two of the main challenges from a systems administration point-of-view is helping the users to get the best out of the computing resources available and to manage their storage use effectively. We present examples of how the human genetics informatics team is providing tools and reports to facilitate this effort.

On the compute side we generate a weekly usage report which is used as a point of discussion in a standup style meeting of the user community, and on the storage side we discuss tools we have developed which provide an easy visualisation of utilisation using treemaps.

Primary author: Mr JAMES, Emyr (Wellcome Trust, Sanger Institute)

Co-author: Dr RANDALL, Joshua (Wellcome Trust, Sanger Institute)

Presenter: Mr JAMES, Emyr (Wellcome Trust, Sanger Institute)

Session Classification: End-user Services and Operating Systems

Track Classification: End-User IT Services & Operating Systems

Contribution ID: 71

Type: **not specified**

Evaluation of distributed open source solutions in CERN database use cases

Tuesday, March 24, 2015 5:20 PM (25 minutes)

There are terabytes of data stored in a relational database (Oracle) at CERN which in fact does not need a relational model. Moreover, using a relational database management system very often brings a significant overhead in terms of resource utilization. The problem is notably observable for warehouse-type data sets. At the same time running analytical workloads on such data sets requires large amount of computing power combined with high storage throughput, a combination which can be achieved with a scalable database system. Introducing this kind of system will not only speed up processing but open new possibilities for data mining as well. This presentation will discuss advantages of using distributed architecture like Hadoop for scalable data processing of CERN data sets; such as stored in Oracle: DB LHC logging system, SCADA systems or even LHC experiments events data stored in Ntuples.

Primary authors: SURDY, Kacper (CERN); GRZYBEK, Maciej (Warsaw University of Technology (PL)); BARANOWSKI, Zbigniew (CERN)

Co-authors: LANZA GARCIA, Daniel (Univ. Extremadura, Cen. Uni. Merida (ES)); GRANCHER, Eric (CERN); CANALI, Luca (CERN)

Presenter: SURDY, Kacper (CERN)

Session Classification: Storage and File Systems

Track Classification: Storage & Filesystems

Contribution ID: 72

Type: **not specified**

SLURM update from the Nordics

Friday, March 27, 2015 9:50 AM (25 minutes)

An update on the current status of SLURM usage in the Nordics, as well as recent developments in improving support for LHC type jobs including tuning for efficient scheduling of multicore grid jobs. Also an overview of some remaining challenges will be given together with discussion on how to address them.

Primary author: WADENSTEIN, Erik Mattias (University of Umeå (SE))

Presenter: WADENSTEIN, Erik Mattias (University of Umeå (SE))

Session Classification: Computing and Batch Systems

Track Classification: Computing & Batch Services

Contribution ID: 73

Type: **not specified**

E-Mail-Migration: transition from Exchange and UNIX mail to Zimbra

Tuesday, March 24, 2015 11:15 AM (25 minutes)

After more than ten years of operations the game is over for Exchange 2003 at DESY. Now Zimbra has been set into production and data from both Exchange 2003 and the UNIX mail service is being migrated and consolidated gradually. The architecture of the Zimbra mail service, the migration procedures and some experiences will be presented. Finally we will look at some integration aspects of Zimbra in the existing DESY landscape.

Primary author: Mr JAHNKE-ZUMBUSCH, Dirk (DESY)

Presenter: Mr JAHNKE-ZUMBUSCH, Dirk (DESY)

Session Classification: End-user Services and Operating Systems

Track Classification: End-User IT Services & Operating Systems

Contribution ID: 75

Type: **not specified**

DESY site report

Monday, March 23, 2015 12:10 PM (15 minutes)

DESY site report

Primary author: Mr VAN DER REEST, Peter (DESY)

Co-author: KEMP, Yves (Deutsches Elektronen-Synchrotron (DE))

Presenter: Mr VAN DER REEST, Peter (DESY)

Session Classification: Site Reports

Track Classification: Site reports

Contribution ID: 76

Type: **not specified**

HEPSPEC analysis across modern architectures

Wednesday, March 25, 2015 5:20 PM (25 minutes)

The talk's coverage will include Xeon Haswell, ARM and Open Compute Platforms

Primary author: Mr POWER, David (Boston Ltd.)

Presenter: Mr POWER, David (Boston Ltd.)

Session Classification: Computing and Batch Systems

Track Classification: Computing & Batch Services

Contribution ID: 77

Type: **not specified**

CERN CentOS 7 Update

Monday, March 23, 2015 4:30 PM (25 minutes)

In this talk we will present a brief status update on CERN's work on CentOS 7, the uptake by the various IT services, and the interaction with the upstream CentOS community.

Primary author: OULEVEY, Thomas (CERN)

Co-author: Dr WIEBALCK, Arne (CERN)

Presenter: Dr WIEBALCK, Arne (CERN)

Session Classification: End-user Services and Operating Systems

Track Classification: End-User IT Services & Operating Systems

Contribution ID: 78

Type: **not specified**

Welcome Address

Monday, March 23, 2015 9:00 AM (10 minutes)

Welcome Address

Primary author: Prof. WHEATER, John (Oxford University)

Presenter: Prof. WHEATER, John (Oxford University)

Session Classification: Miscellaneous

Track Classification: Miscellaneous

Contribution ID: 79

Type: **not specified**

Workshop Logistics

Monday, March 23, 2015 9:10 AM (5 minutes)

Workshop Logistics

Primary author: GRONBECH, Peter (University of Oxford (GB))

Presenter: GRONBECH, Peter (University of Oxford (GB))

Session Classification: Miscellaneous

Track Classification: Miscellaneous

Contribution ID: **80**

Type: **not specified**

Workshop Wrap-Up

Friday, March 27, 2015 12:50 PM (15 minutes)

Usual summary and conclusions

Primary author: Dr MEINHARD, Helge (CERN)

Co-author: Dr WONG, Tony (Brookhaven National Laboratory)

Presenter: Dr MEINHARD, Helge (CERN)

Session Classification: Miscellaneous

Track Classification: Miscellaneous

Contribution ID: **81**

Type: **not specified**

Panel and BoF session: Ask the CEPH experts

Wednesday, March 25, 2015 1:00 PM (1 hour)

Presenter: Dr WIEBALCK, Arne (CERN)

Session Classification: Storage and File Systems

Contribution ID: **82**

Type: **not specified**

IPv6 tutorial for administrators

Tuesday, March 24, 2015 6:00 PM (1h 30m)

Presenters: KELSEY, Dave (STFC - Rutherford Appleton Lab. (GB)); Dr MCKEE, Shawn (University of Michigan ATLAS Group)

Session Classification: Security and Networking

Contribution ID: **83**

Type: **not specified**

Scientific Linux Current Status

Monday, March 23, 2015 4:05 PM (25 minutes)

Current Status of Scientific Linux

Summary

Status of Scientific Linux since Fall HEPix 2014

Primary author: SIEH, Connie (FNAL)

Co-authors: KING, Bonnie (FNAL); RIEHECKY, Pat (Fermilab)

Presenter: SIEH, Connie (FNAL)

Session Classification: End-user Services and Operating Systems

Track Classification: End-User IT Services & Operating Systems

Contribution ID: **84**

Type: **not specified**

FNAL site report

Tuesday, March 24, 2015 12:05 PM (15 minutes)

Site report from Fermilab

Primary authors: SCOTT, Rennie S. (FNAL); SIEH, connie (Fermilab)

Presenters: SCOTT, Rennie S. (FNAL); SIEH, connie (Fermilab)

Session Classification: Site Reports

Track Classification: Site reports

Contribution ID: 85

Type: **not specified**

CCIN2P3 Site Report

Monday, March 23, 2015 12:25 PM (15 minutes)

We will present the latest status of the IN2P3 Computer Center. Emphasis will be made to the infrastructure and system area.

Summary

We will present the latest status of the IN2P3 Computer Center. Emphasis will be made to the infrastructure and system area.

Primary author: Mr CARPENTIER, Julien (CCIN2P3)

Presenter: Mr CARPENTIER, Julien (CCIN2P3)

Session Classification: Site Reports

Track Classification: Site reports

Contribution ID: **86**

Type: **not specified**

A word on phishing

Friday, March 27, 2015 12:45 PM (5 minutes)

... at HEPiX

Primary author: Mr WARTEL, Romain (CERN)

Presenter: Mr WARTEL, Romain (CERN)

Session Classification: Miscellaneous

Track Classification: Miscellaneous