

HEPiX Fall 2013 Workshop



Report of Contributions

Contribution ID: 0

Type: **Oral presentation**

Solving Small Files Problem in Enstore

Thursday, 31 October 2013 14:00 (30 minutes)

Enstore is a tape based Mass Storage System originally designed for Run II Tevatron experiments at FNAL (CDF, D0). Over the years it has proven to be reliable and scalable data archival and delivery solution, which meets diverse requirements of variety of applications including US CMS Tier 1, High Performance Computing, Intensity Frontier experiments as well as data backups. Data intensive experiments like CDF, D0 and US CMS Tier 1 generally produce huge amount of data stored in files with the average size of few Gigabytes, which is optimal for writing and reading data to/from tape. In contrast, much of the data produced by Intensity Frontier experiments, Lattice QCD and Cosmology is sparse, resulting in accumulation of large amounts of small files.

Reliably storing small files on tape is inefficient due to file marks writing which takes significant amount of the overall file writing time (few seconds). There are several ways of improving data write rates, but some of them are unreliable, some are specific to the type of tape drive and still do not provide transfer rates adequate to rates offered by tape drives (20% of the drives potential rate). In order to provide good rates for small files in a transparent and consistent manner, the Small File Aggregation (SFA) feature has been developed to provide aggregation of files into containers which are subsequently written to tape. The file aggregation uses reliable internal Enstore disk buffer. File grouping is based on policies using file metadata and other user defined steering parameters.

If a small file, which is a part of a container, is requested for read, the whole container is staged into internal Enstore read cache thus providing a read ahead mechanism in anticipation of future read requests for files from the same container. SFA is provided as service implementing file aggregation and staging transparently to user.

The SFA is has been successfully used since April 2012 by several experiments. Currently we are preparing to scale up write/read SFA cache.

This paper describes Enstore Small Files Aggregation feature and discusses how it can be scaled in size and transfer rates.

Primary author: Dr MOIBENKO, Alexander (Fermi NAtiona Accelerator Laboratoy)

Co-authors: Mr KULYAVTSEV, Alexander (FNAL); LITVINTSEV, Dmitry (FNAL); Dr OLEYNIK, Gene (Fermilab); HENDRY, John (Fermi National Accelerator Laboratory); NAYMOLA, Stan (Fermi Naitional Accelerator Laboratory)

Presenter: Dr MOIBENKO, Alexander (Fermi NAtiona Accelerator Laboratoy)

Session Classification: Storage and file systems

Track Classification: Storage & Filesystems

Contribution ID: 1

Type: **Oral presentation**

Fermilab Site Report - Fall 2013 HEPiX

Tuesday, 29 October 2013 09:45 (15 minutes)

Fermilab site report - Fall 2013 HEPiX.

Primary author: Dr CHADWICK, Keith (Fermilab)

Presenter: Dr CHADWICK, Keith (Fermilab)

Session Classification: Site reports

Track Classification: Site reports

Contribution ID: 2

Type: **Oral presentation**

Scientific Linux current status update

Thursday, 31 October 2013 16:45 (25 minutes)

This presentation will provide an update on the current status of Scientific Linux, descriptions for some possible future goals, and allow a chance for users to provide feedback on its direction.

Summary

This presentation will provide an update on the current status of Scientific Linux, descriptions for some possible future goals, and allow a chance for users to provide feedback on its direction.

Primary authors: RIEHECKY, Pat (Fermilab); SIEH, connie (Fermilab)

Presenter: RIEHECKY, Pat (Fermilab)

Session Classification: IT End User Services

Track Classification: End-user IT Services & Operating Systems

Contribution ID: 3

Type: **Oral presentation**

Puppet at Fermilab- Managing a Large Heterogeneous Environment

Wednesday, 30 October 2013 16:00 (25 minutes)

Puppet has been in use by the Fermilab Experiments Facilities department to support computing for a variety of experiments for the last several years. This presentation will discuss our experience deploying, refining, and upgrading Puppet to scale to thousands of systems, across different experiments, servers, batch nodes, and workstations. This presentation will describe our efforts to build an understandable and maintainable Puppet repository, that is flexible and usable by multiple administrators.

Primary author: SIMMONDS, Edward (Fermilab)

Co-authors: HARRINGTON, Jason (Fermilab); PARSONS, Tyler (Fermilab)

Presenters: SIMMONDS, Edward (Fermilab); PARSONS, Tyler (Fermilab)

Session Classification: Basic IT services

Track Classification: Basic IT Services

Contribution ID: 4

Type: **Oral presentation**

dCache news

Thursday, 31 October 2013 14:30 (30 minutes)

This presentation is intended to bring HEP storage administrators up to speed with ongoing dCache developments and activities.

In the context of WLCG, we will report on our collaboration with the xRootd folks in terms of federated storage and monitoring,

our efforts to support the strict separation of CMS between disk and tape storage endpoints, and we hope to have the first results on direct NFS 4.1/pNFS access of WLCG GRID jobs to the dCache storage elements at DESY.

On a related topic we will briefly describe a new module, allowing dCache to handle files, being too small to be exchanged with a tertiary storage system individually.

In general terms we will elaborate on interesting developments in dCache, triggered by our involvement in the German “Large Scale Data Management and Analysis, LSDMA” project. This includes, but is not limited to integrating dCache into federated identity infrastructures and, in collaboration with students of the HTW Berlin, providing standard cloud interfaces.

Last but not least we will give some insight into the current and future funding structure of the dCache collaboration; release policy and channels and our first attempt to provide a concrete cloud storage service.

Summary

This presentation is intended to bring HEP storage administrators up to speed with ongoing dCache developments and activities.

Primary author: Dr FUHRMANN, Patrick (DESY)

Co-authors: Dr ROSSI, Albert (FNAL); PETERSEN, Antje (DESY); BERNARDT, Christian (Deutsches Elektronen-Synchrotron (DE)); LITVINTSEV, Dmitry (FNAL); Dr BEHRMANN, Gerd (Nordunet); SCHWANK, Karsten; MILLAR, Paul (Deutsches Elektronen-Synchrotron (DE)); Mr MKRTCHYAN, Tigran Mkrtchyan (Deutsches Elektronen-Synchrotron DESY)

Presenter: Dr FUHRMANN, Patrick (DESY)

Session Classification: Storage and file systems

Track Classification: Storage & Filesystems

Contribution ID: 5

Type: **Oral presentation**

CERN Site report

Tuesday, 29 October 2013 10:00 (15 minutes)

News from CERN since the Bologna Workshop.

Primary author: Dr WIEBALCK, Arne (CERN)

Presenter: Dr WIEBALCK, Arne (CERN)

Session Classification: Site reports

Track Classification: Site reports

Contribution ID: 6

Type: **Oral presentation**

Building an organic block storage service at CERN with Ceph

Thursday, 31 October 2013 15:30 (25 minutes)

This is a report on CERN IT's 3 PetaByte Ceph pre-production cluster set up in the past couple of months which will initially serve as a storage backend for OpenStack images and volumes (backends for AFS or NFS servers are options we will explore as well). In addition to a discussion of the architecture and configuration of the cluster, we will present results of our functionality tests, some performance numbers as well as best practices learned.

Primary authors: Dr WIEBALCK, Arne (CERN); Dr VAN DER STER, Daniel (CERN)

Presenter: Dr WIEBALCK, Arne (CERN)

Session Classification: Storage and file systems

Track Classification: Storage & Filesystems

Contribution ID: 7

Type: **Oral presentation**

OpenAFS and IPv6: Follow-up on the Bologna discussions / HEPiX survey

Thursday, 31 October 2013 11:15 (15 minutes)

This will be a follow-up of the discussions about OpenAFS and IPv6 we had in Bologna, in particular summarizing input from potential developers on timelines/prices/development models, as well as conclusions from the survey conducted to understand the needs of the HEPiX community regarding the lack of IPv6 in OpenAFS.

Primary authors: Dr WIEBALCK, Arne (CERN); Mr VAN DER REEST, Peter (DESY)

Presenter: Dr WIEBALCK, Arne (CERN)

Session Classification: Storage and file systems

Track Classification: Storage & Filesystems

Contribution ID: 9

Type: **Oral presentation**

Automatic server registration and burn-in framework

Monday, 28 October 2013 16:25 (25 minutes)

This talk provides an overview of CERN IT's automatic server registration and burn-in framework. It will mainly focus on the reasons behind the development of such framework, and the implementation details. A detailed walkthrough of the process stages will be presented, along with the first results of the acceptance of about 1,500 servers. Finally, the talk will underline the benefits of the server registration and certification process automation.

Primary author: XAFI, Afroditi (CERN)

Co-authors: BONFILLOU, Eric (CERN); VALSAN, Liviu (CERN); BARRING, Olof (CERN)

Presenter: XAFI, Afroditi (CERN)

Session Classification: Basic IT services

Track Classification: Basic IT Services

Contribution ID: 10

Type: **Oral presentation**

Batch System Development and Evolution at RACF

Tuesday, 29 October 2013 14:00 (30 minutes)

Scheduling jobs with heterogeneous resource requirements to a pool of computers with heterogeneous resources is a challenging task and Condor is beginning to tackle this in the most generic form. Integrating so-called partitionable slots (a batch resource able to be sliced along a variety of dimensions, from RAM to CPUs, to disks or even GPUs) with the rest of Condor's accounting and matchmaking mechanisms is a significant challenge. At BNL we began working closely with the Condor team to integrate Condor's Hierarchical Accounting Groups with partitionable slots in a consistent and correct manner. The results of that work are presented in this talk.

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

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

Session Classification: Computing and batch systems

Track Classification: Computing & Batch Services

Contribution ID: 11

Type: **Oral presentation**

CERN Cloud Status

Friday, 1 November 2013 09:00 (25 minutes)

This talk will provide an update on the CERN private infrastructure-as-a-service cloud which is now in production based on OpenStack Grizzly. Along with the current status, we will give plans for the next developments and evolution of the service in areas such as block storage using ceph and Netapp, Kerberos and X.509 support and scaling to multiple cells across CERN's two data centres.

Primary authors: OULEVEY, Thomas (CERN); BELL, Tim (CERN)

Presenter: OULEVEY, Thomas (CERN)

Session Classification: Grids, clouds, virtualisation

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: 12

Type: **Oral presentation**

Building RPM packages at CERN with Koji

Thursday, 31 October 2013 16:20 (25 minutes)

This talk will show how we used Koji, to give the different IT teams flexibility. It will also cover the building of Redhat packages for Scientific Linux Cern and other Redhat addons. Finally, we will state the limitations of Koji and some workaround we found.

Primary author: OULEVEY, Thomas (CERN)

Presenter: OULEVEY, Thomas (CERN)

Session Classification: IT End User Services

Track Classification: End-user IT Services & Operating Systems

Contribution ID: 13

Type: **Oral presentation**

INFN-T1 site report

Monday, 28 October 2013 09:30 (15 minutes)

An update on what's going on at INFN-T1 center

Primary author: CHIERICI, Andrea (INFN-CNAF)

Presenter: CHIERICI, Andrea (INFN-CNAF)

Session Classification: Site reports

Track Classification: Site reports

Contribution ID: 14

Type: **Oral presentation**

USCMS T1 and LPC Data Storage Challenges and Solutions

Thursday, 31 October 2013 10:15 (30 minutes)

The CMS T1 facility at Fermilab manages many tens of petabytes of data for CMS. This talk will present some historical information on the solutions used to store this data as well as information on the new solutions we are in the process of implementing and how we got to where we are now.

Primary author: GIACCHETTI, Lisa Ann (Fermi National Accelerator Lab. (US))

Presenter: GIACCHETTI, Lisa Ann (Fermi National Accelerator Lab. (US))

Session Classification: Storage and file systems

Track Classification: Storage & Filesystems

Contribution ID: 15

Type: **Oral presentation**

Unified Communications, the new IP telephony @ CERN

Wednesday, 30 October 2013 16:50 (25 minutes)

The CERN Telephone service is moving towards unified communications.

New IP Phone devices connected to Lync IP Phone Service enhance the classic telephony by adding many features like IM, presence, voice mailbox, call delegation, etc. The complete integration with Exchange allows the mailbox to be used as call log history, voice mailbox, etc.

In addition, connecting the IP Phone to your computer via USB makes the range of possibilities wider by profiting from the integration of the device with the operating system: click to call from web pages or emails, etc.

The service is also available from anywhere with software clients on many platforms (desktops & laptops, Smartphones) allowing calls as if you were in your office, conference call meetings, instant setup in any meeting room, etc.

What about the future plans? Lync 2013 and its integration with Skype will increase the already big possibilities that this solution offers.

Primary author: MORENO PASCUAL, Fernando (CERN)

Presenter: MORENO PASCUAL, Fernando (CERN)

Session Classification: Basic IT services

Track Classification: End-user IT Services & Operating Systems

Contribution ID: 16

Type: **Oral presentation**

BNL RACF Site Report

Monday, 28 October 2013 11:00 (15 minutes)

Brookhaven National Lab (BNL) will present the site report for the RHIC-ATLAS Computing Facility (RACF)

Primary authors: Dr RIND, Ofer (BROOKHAVEN NATIONAL LABORATORY); Dr WONG, Tony (Brookhaven National Laboratory)

Presenters: Dr RIND, Ofer (BROOKHAVEN NATIONAL LABORATORY); Dr WONG, Tony (Brookhaven National Laboratory)

Session Classification: Site reports

Track Classification: Site reports

Contribution ID: 17

Type: **Oral presentation**

Building a Puppet Infrastructure @ DESY

Monday, 28 October 2013 16:00 (25 minutes)

Starting in 2012, DESY has been extending it's IT systems infrastructure to make use of the Puppet configuration management system developed by Puppetlabs. The main focus of this talk is to share the experience gained in this program of work and to summarize the current status and outlook of the Puppet infrastructure at the DESY site.

Summary

Status of Puppet at the DESY site.

Primary author: ENGELS, Jan (Deutsches Elektronen-Synchrotron (DE))

Co-authors: DIETRICH, Stefan (Deutsches Elektronen-Synchrotron (DE)); STERNBERGER, Sven (Deutsches Elektronen-Synchrotron (DE))

Presenter: ENGELS, Jan (Deutsches Elektronen-Synchrotron (DE))

Session Classification: Basic IT services

Track Classification: Basic IT Services

Contribution ID: 18

Type: **Oral presentation**

PanDA Beyond ATLAS: Workload Management for Data Intensive Science

Friday, 1 November 2013 10:30 (25 minutes)

The PanDA Production ANd Distributed Analysis system has been developed by ATLAS to meet the experiment's requirements for a data-driven workload management system for production and distributed analysis processing capable of operating at LHC data processing scale. After 7 years of impressively successful PanDA operation in ATLAS there are also other experiments which can benefit from PanDA in the Big Data challenge, with several at various stages of evaluation and adoption. The new project "Next Generation Workload Management and Analysis System for Big Data" is extending PanDA to meet the needs of other data intensive scientific applications in HEP, astro-particle and astrophysics communities, bio-informatics and other fields as a general solution to large scale workload management. PanDA can utilize dedicated or opportunistic computing resources such as grids, clouds, and High Performance Computing facilities, and is being extended to leverage next generation intelligent networks in automated workflow management and brokerage. This presentation will provide an overview, the current status and future plans of the Big PanDA project.

Primary authors: VANIACHINE, Alexandre (ANL); KLIMENTOV, Alexei (Brookhaven National Laboratory (US)); PETROSYAN, Artem (Joint Inst. for Nuclear Research (RU)); OLEYNIK, Danila (Joint Inst. for Nuclear Research (RU)); YU, Dantong (BROOKHAVEN NATIONAL LABORATORY); SCHOVANCOVA, Jaroslava (Brookhaven National Laboratory (US)); DE, Kaushik (University of Texas at Arlington (US)); NILSSON, Paul (University of Texas at Arlington (US)); PANITKIN, Sergey (Brookhaven National Laboratory (US)); MAENO, Tadashi (Brookhaven National Laboratory (US)); WENAUS, Torre (Brookhaven National Laboratory (US))

Presenter: SCHOVANCOVA, Jaroslava (Brookhaven National Laboratory (US))

Session Classification: Grids, clouds, virtualisation

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: 19

Type: **Oral presentation**

The CSC Kajaani datacentre

Monday, 28 October 2013 14:00 (25 minutes)

Early 2013 CSC (the it centre for science in Finland) powered up the new datacentre on Kajaani, 600 km north of the offices. The datacentre focuses of energy efficiency and cost-cutting, but the design and implementation was not easy.

Summary

Globalization created free space in factories in smaller cities of Finland. This has created a insurge of datacenters using cheap cooling, stable political systems and stable hydroelectric power.

Currently the datacentre hosts a Cray XE30 supercomputer and a supercluster/cloud infrastructure in addition to the common storage for both systems.

Primary author: TIGERSTEDT, Ulf (CSC Oy)

Presenter: TIGERSTEDT, Ulf (CSC Oy)

Session Classification: IT facilities and business continuity

Track Classification: IT Facilities & Business Continuity

Contribution ID: 20

Type: **Oral presentation**

The Caltech CMS Tier2

Monday, 28 October 2013 10:00 (15 minutes)

The Caltech Tier2 is a major site providing substantial and reliable computational and storage resources to CMS, combining production processing of simulated events, support for US CMS physics analysis, and computing, software systems, and network developments. Caltech continues to lead key several areas of the LHC computing and software aimed at enabling grid-based data analysis, as well as global-scale networking. An update on the status of the Caltech Tier 2 is given together with the synergistic activities in networking operations and R&D at Caltech.

Summary

The Caltech Tier2 is a major site providing substantial and reliable computational and storage resources to CMS, combining production processing of simulated events, support for US CMS physics analysis, and computing, software systems, and network developments. Caltech continues to lead key several areas of the LHC computing and software aimed at enabling grid-based data analysis, as well as global-scale networking. An update on the status of the Caltech Tier 2 is given together with the synergistic activities in networking operations and R&D at Caltech.

Primary authors: BARCZYK, Artur Jerzy (California Institute of Technology (US)); MUGHAL, Azher (California Institute of Technology (US)); Dr KCIRA, Dorian (California Institute of Technology (US)); NEWMAN, Harvey (California Institute of Technology (US)); CURY SIQUEIRA, Samir (California Institute of Technology (US))

Presenter: Dr KCIRA, Dorian (California Institute of Technology (US))

Session Classification: Site reports

Track Classification: Site reports

Contribution ID: 21

Type: **Oral presentation**

Batch system status at the RAL Tier 1

Tuesday, 29 October 2013 14:30 (30 minutes)

The RAL Tier 1 maintains a batch farm with close to 10000 job slots that is used by all the LHC VOs as well as a number of smaller users. We have increasingly found that our existing batch system is unable to cope with the demands placed on it by our users. During the past year work has been carried out evaluating alternative technologies to our existing Torque/Maui batch system and preparing a new system for use in production. This talk will discuss why we decided to migrate to a new batch system, the selection process, testing carried out, and our experiences with the new batch system so far.

Primary author: LAHIFF, Andrew David (STFC - Science & Technology Facilities Council (GB))

Co-authors: DEWHURST, Alastair (STFC - Science & Technology Facilities Council (GB)); COLLIER, Ian (UK Tier1 Centre); KELLY, John (STFC)

Presenter: LAHIFF, Andrew David (STFC - Science & Technology Facilities Council (GB))

Session Classification: Computing and batch systems

Track Classification: Computing & Batch Services

Contribution ID: 22

Type: **Oral presentation**

Future of Batch Processing at CERN

Tuesday, 29 October 2013 11:30 (30 minutes)

The CERN Batch System is comprised of 4000 worker nodes, 60 queues and various types of large user communities. In light of the recent developments driven by the Agile Infrastructure and the more demanding processing needs, the Batch System will be faced with increasingly challenging scalability and flexibility needs.

So as to prepare for these high expectations, the CERN Batch Team have been designing a framework to be able to easily subject the Batch System to different types of strains and scalability tests. This framework has been the testbed for a number of candidate batch systems, one of which will foster the future of batch processing at CERN. So far, SLURM, Condor and Grid Engine have been under evaluation.

In this talk, we present the design of this test framework and the initial results of our evaluation of the aforementioned batch systems, from a scalability and an administrative perspective.

Primary author: BELLEMAN, Jerome (CERN)

Presenter: BELLEMAN, Jerome (CERN)

Session Classification: Computing and batch systems

Track Classification: Computing & Batch Services

Contribution ID: 23

Type: **Oral presentation**

IN2P3-CC IAAS cloud status

Friday, 1 November 2013 09:25 (25 minutes)

The talk presents some past and ongoing testings made around IAAS cloud technologies at IN2P3-CC. Openstack has been deployed for some years now, involved in a variety of projects. The presentation tackles:

- Openstack: the CMP of choice
- use cases with test and production services, computing, community cloud
- implemented features, hardware and services
- impacts on the datacenter
- how to implement high availability with Openstack
- Openstack experience (pros/cons)
- what's next ?

Primary author: PUEL, Mattieu (CNRS)

Presenter: PUEL, Mattieu (CNRS)

Session Classification: Grids, clouds, virtualisation

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: 24

Type: **Oral presentation**

Mobility at CERN

Tuesday, 29 October 2013 16:00 (30 minutes)

With the dramatic increase of wireless-capable devices, Wi-Fi connectivity has become an essential network service on a par with the traditional cabled network. The evolution of the CERN Wi-Fi infrastructure will be presented, including the BYOD strategy and the integration of eduroam. The deployment considerations for large conference rooms and underground facilities will also be addressed.

Primary author: CEUTERICKX, Sebastien (CERN)

Presenter: CEUTERICKX, Sebastien (CERN)

Session Classification: Security and networking

Track Classification: Security & Networking

Contribution ID: 25

Type: **Oral presentation**

Latest changes on CERN networks

Tuesday, 29 October 2013 16:30 (30 minutes)

The latest changes on the CERN network infrastructure will be presented. This includes the deployment of IPv6, the network connectivity for the Data Centre extension at Wigner and the upgrade of the network infrastructure for Business Continuity. The second part of the talk will give an overview of the implementation of a new, safety-related wireless network (TETRA).

Primary author: CEUTERICKX, Sebastien (CERN)

Co-authors: PASCAL, Aurelie (CERN); KISHIMOTO BISBE, Carles (CERN); GUTIERREZ RUEDA, David (CERN); MARTELLI, Edoardo (CERN)

Presenter: CEUTERICKX, Sebastien (CERN)

Session Classification: Security and networking

Track Classification: Security & Networking

Contribution ID: 26

Type: **Oral presentation**

The HEPiX IPv6 Working Group

Wednesday, 30 October 2013 09:00 (30 minutes)

An update on the activities of the group in IPv6 testing and planning since the Bologna meeting.

Primary author: KELSEY, Dave (STFC - Science & Technology Facilities Council (GB))

Presenter: KELSEY, Dave (STFC - Science & Technology Facilities Council (GB))

Session Classification: Security and networking

Track Classification: Security & Networking

Contribution ID: 27

Type: **Oral presentation**

Federated Identity Management for HEP

Wednesday, 30 October 2013 11:00 (30 minutes)

There is much activity in the area of identity management for research communities. This talk will present the current status of this work and explore possible future options for WLCG and HEP more generally.

Primary author: KELSEY, Dave (STFC - Science & Technology Facilities Council (GB))

Presenter: KELSEY, Dave (STFC - Science & Technology Facilities Council (GB))

Session Classification: Security and networking

Track Classification: Security & Networking

Contribution ID: 28

Type: **Oral presentation**

PDSF at NERSC - Site Report HEPiX Fall 2013

Monday, 28 October 2013 12:15 (15 minutes)

PDSF (Parallel Distributed Systems Facility) has been in continuous operation at NERSC since 1996 on dedicated and ever changing hardware, supporting a broad user base in the high energy physics community. We will describe recent and ongoing changes in the underlying architecture of PDSF. We are moving to a model where the PDSF cluster will consist of dedicated front end servers and a set of compute nodes drawn from a backend cluster, unified by a batch system. User storage will be migrated from a dedicated GPFS cluster to a site wide NERSC Global File system, facilitating the use of other NERSC clusters which have a faster and lower latency interconnect for parallel jobs. These changes should help minimize downtime when the cluster is moved to a new building at Lawrence Berkeley Laboratory in 2014/15.

Primary author: BOTTTS, James (L)

Co-authors: SAKREJDA, Iwona; PEZZAGLIA, Larry (LBNL)

Presenter: BOTTTS, James (L)

Session Classification: Site reports

Track Classification: Site reports

Contribution ID: 29

Type: **Oral presentation**

Using control system tools for operation and debugging

Wednesday, 30 October 2013 14:00 (30 minutes)

When investigating a problem, one typically needs to gather and correlate information from disparate sources: operating system, batch system, data transfer tools, and so on. We investigate the use of Control System Studio to gather information from multiple places. When the appropriate hooks are created, this should allow the end user to create ad-hoc ways to mix and match data without the need of writing programs or web pages.

Primary author: Mr CARCASSI, Gabriele (Brookhaven National Laboratory (US))

Presenter: Mr CARCASSI, Gabriele (Brookhaven National Laboratory (US))

Session Classification: Basic IT services

Track Classification: Basic IT Services

Contribution ID: 30

Type: **Oral presentation**

Jefferson Lab Site Report

Tuesday, 29 October 2013 09:15 (15 minutes)

An update of high performance and scientific computing activities since the Spring 2012 meeting.

Primary author: PHILPOTT, Sandy (JLAB)

Presenter: PHILPOTT, Sandy (JLAB)

Session Classification: Site reports

Track Classification: Site reports

Contribution ID: 31

Type: **Oral presentation**

OpenShift on your own cloud

Friday, 1 November 2013 11:45 (25 minutes)

OpenShift has three offerings, Origin, Online, and Enterprise. Now you can enjoy the benefits of Paas in the public cloud, or on your own cloud.

I will be showing OpenShift Origin, setup locally. What features does it have for both admin and user. How will that help both labs and experiments.

Summary

Get all the benefits of OpenShift on your own infrastructure.

Primary author: Mr DAWSON, Troy (Red Hat)

Presenter: Mr DAWSON, Troy (Red Hat)

Session Classification: Grids, clouds, virtualisation

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: 32

Type: **Oral presentation**

CernVM-FS - Beyond LHC Computing

Friday, 1 November 2013 11:20 (25 minutes)

In the last three years the CernVM Filesystem (CernVM-FS) has transformed the distribution of experiment software to WLCG grid sites. CernVM-FS removes the need for local installations jobs and performant software at sites, in addition it often improves performance at the same time. Furthermore the use of CernVM-FS standardizes the computing environment across the grid and removes the need for software tagging at sites.

Now established and proven to work at scale, CernVM-FS is beginning to perform a similar role for non-LHC computing.

We discuss the deployment of a non-LHC Stratum 0 'master' CernVM-FS repository at the RAL Tier 1 and the development of a network of Stratum 1 replicas somewhat modeled upon the infrastructure developed to support WLCG computing.

Primary author: COLLIER, Ian (UK Tier1 Centre)

Co-author: CONDURACHE, Catalin (STFC - Science & Technology Facilities Council (GB))

Presenter: COLLIER, Ian (UK Tier1 Centre)

Session Classification: Grids, clouds, virtualisation

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: 33

Type: **Oral presentation**

Grid Engine: One Roadmap

Tuesday, 29 October 2013 15:00 (30 minutes)

The presentation will cover the status and future of Grid Engine. Specifically a summary of what Univa offers Grid Engine users, what we have done since SGE 6.2.u5, and what is included in our soon to be released 8.2.0 version. Additionally several case studies of Grid Engine users from different industries will be presented to show the value Univa provides our users.

Primary author: BRUNNER, Cameron (U)

Presenter: BRUNNER, Cameron (U)

Session Classification: Computing and batch systems

Track Classification: Computing & Batch Services

Contribution ID: 34

Type: **Oral presentation**

Evolution of the OSG Authentication Model

Wednesday, 30 October 2013 11:30 (30 minutes)

The Open Science Grid (OSG) has undergone some changes in its authentication model for user job submission. This talk will outline changes already implemented as well as future plans as they stand now, together with the use cases that motivated them.

Summary

This will cover changes in the OSG in regards to:

- * OSG's CA transition for X.509 authentication.
- * CILogon basic credentials.
- * Job submission without x.509 credentials.

Primary author: HILL, kevin

Presenter: HILL, kevin

Session Classification: Security and networking

Track Classification: Security & Networking

Contribution ID: 35

Type: **Oral presentation**

Security update

Wednesday, 30 October 2013 09:30 (30 minutes)

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.

Primary author: Mr WARTEL, Romain (CERN)**Presenter:** Mr WARTEL, Romain (CERN)**Session Classification:** Security and networking**Track Classification:** Security & Networking

Contribution ID: 36

Type: **Oral presentation**

Technical security tips and techniques

Wednesday, 30 October 2013 12:00 (30 minutes)

This is an interactive technical session aimed at system administrators, and presenting several Linux tools and tips that can be used to both drastically reduce the chance of root compromise, as well as increase the amount information available during forensics.

(This session is not necessarily meant to be a general presentation for the audience and can also be a dedicated BoF.)

Primary author: Mr WARTEL, Romain (CERN)

Presenter: Mr WARTEL, Romain (CERN)

Session Classification: Security and networking

Track Classification: Security & Networking

Contribution ID: 37

Type: **Oral presentation**

Why specify and monitor hard disk drive workloads

Thursday, 31 October 2013 09:00 (45 minutes)

The reliability of hard disk drives (HDD) has been quantified historically by a mean time to failure (MTTF), or an annualized failure rate (AFR), defined at a specified operating temperature, and an assumed functional duty cycle. We provide justification for replacing the ambiguous concept of duty cycle with the readily quantifiable “workload”, which is defined as the total amount of data read from or written to the drive per unit time. This relatively subtle change leads to the conclusion that MTTF alone is insufficient to describe the field reliability of HDDs. This results from the fact that HDD failure rates are more tightly coupled to the total amount of data transferred rather than the total power-on-time. It immediately follows that a metric based on Mean Petabytes to Failure (MPbF) is most appropriate to quantify the intrinsic reliability of HDDs. Until MPbF is accepted as the critical measure of product quality, however, WD will specify both the MTTF and the maximum workload at which the product meets the MTTF requirement to unambiguously define the HDD reliability. Furthermore, WD will offer a new drive feature, the Drive Workload Monitor (DWM), that will facilitate retrieval of the total data transferred at any point of the HDD lifetime.

Primary authors: Dr CHATTOPADHYAY, Amit (Western Digital Corporation); Dr TYNDALL, George (Western Digital Corporation)

Presenter: Dr CHATTOPADHYAY, Amit (Western Digital Corporation)

Session Classification: Storage and file systems

Track Classification: Storage & Filesystems

Contribution ID: **38**

Type: **Oral presentation**

Nikhef Site Report

Monday, 28 October 2013 09:45 (15 minutes)

Fall 2013 site report

Primary author: KUIPERS, Paul (Nikhef)

Presenter: KUIPERS, Paul (Nikhef)

Session Classification: Site reports

Track Classification: Site reports

Contribution ID: 39

Type: **Oral presentation**

HS06 - A first look in the Performance per Watt

Tuesday, 29 October 2013 11:00 (30 minutes)

I started to make measurement of Power consumption when running the HEP-SPEC06 benchmark.
A few slides on the move from SL5 to SL6.

Primary author: Dr MICHELOTTO, Michele (Universita e INFN (IT))

Presenter: Dr MICHELOTTO, Michele (Universita e INFN (IT))

Session Classification: Computing and batch systems

Track Classification: Computing & Batch Services

Contribution ID: 40

Type: **Oral presentation**

HPC Activities at CERN

Tuesday, 29 October 2013 12:00 (30 minutes)

While the majority of physics computing needs are covered sufficiently by the CERN batch services, some other applications have some special requirement and therefore need special treatment. In this presentation we will discuss which these applications are, the methodology we used in order to gather more in depth statistics and we will present some preliminary results. Finally some future steps and thoughts will be presented.

Primary author: AGTZIDIS, Ioannis (CERN)**Co-author:** Mr HUSEJKO, Michal (CERN)**Presenter:** AGTZIDIS, Ioannis (CERN)**Session Classification:** Computing and batch systems**Track Classification:** Computing & Batch Services

Contribution ID: 41

Type: **Oral presentation**

OpenAFS Status Report

Thursday, 31 October 2013 11:30 (30 minutes)

A status report on OpenAFS with a focus on:

- 2013 Security Vulnerabilities

. OPENAFS-SA-2013-001

Buffer overflows in OpenAFS fileserver

. OPENAFS-SA-2013-002

Buffer overflow in OpenAFS ptserver

. OPENAFS-SA-2013-003

Brute force DES attack permits compromise of AFS cell

. OPENAFS-SA-2013-004

vos -encrypt doesn't encrypt connection data

- latest OS platform support including Windows 8.1, OSX Maverick, and Linux kernels
- Foundation status

Primary author: BRASHEAR, Derrick (Y)

Presenter: BRASHEAR, Derrick (Y)

Session Classification: Storage and file systems

Track Classification: Storage & Filesystems

Contribution ID: 42

Type: **Oral presentation**

Introducing YFS 1.0

Thursday, 31 October 2013 12:00 (30 minutes)

YFS is a Software Defined Storage solution for secure private, public and hybrid cloud storage deployments. YFS 1.0 clients and servers are dual protocol stack providing next generation file system capabilities while maintaining backward compatibility with IBM AFS 3.6 and OpenAFS clients and servers.

This talk will highlight the enhanced capabilities of YFS 1.0 vs OpenAFS 1.6.5 including the new high performance IPv6-capable RX implementation, security improvements, improved Windows / OSX application compatibility, Year 2038 safety, and other protocol improvements.

Primary authors: BRASHEAR, Derrick (Y); ALTMAN, Jeffrey (Your File System Inc.)

Presenters: BRASHEAR, Derrick (Y); ALTMAN, Jeffrey (Your File System Inc.)

Session Classification: Storage and file systems

Track Classification: Storage & Filesystems

Contribution ID: 43

Type: **Oral presentation**

FermiCloud update - enabling scientific workflows in the cloud

Friday, 1 November 2013 09:50 (25 minutes)

In 2010, Fermilab initiated the FermiCloud project to deliver a dynamic and scalable Infrastructure-as-a-Service (IaaS) capability using open source cloud computing frameworks to support the needs of the Fermilab scientific communities. A collaboration of personnel from Fermilab and the Korea Institute of Science and Technology Information (KISTI) has focused significant work over the past 18 months to deliver improvements to the applicability and robustness of FermiCloud, together with specific accomplishments with respect to direct and indirect support of science.

The opportunities, challenges and successes of cloud computing at Fermilab will be presented, including:

- GridBursting to FermiCloud: extending the grid through the cloud
- Idle VM suspension
- FermiCloud interoperability

Primary author: BERNABEU ALTAYO, Gerard (F)

Presenter: BERNABEU ALTAYO, Gerard (F)

Session Classification: Grids, clouds, virtualisation

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: 44

Type: **Oral presentation**

Puppet at USCMS-T1 and FermiLab

Wednesday, 30 October 2013 16:25 (25 minutes)

Over the past year, the USCMS-T1 project has decided to jump head-first into Puppet as our primary configuration management tool. We would like to talk about what's worked, what hasn't worked, and how we've been able to work with the other Fermilab teams to share our experiences without necessarily sharing a code base.

Summary

This will be a fairly informal presentation explaining why we chose puppet, what's happened with puppet in the near past, how we're collaborating with the other Fermi teams that use puppet, and the benefits of having a central puppet-users@fnal.gov list has been helpful for everyone involved.

Primary author: SKIRVIN, Timothy Michael (F)

Presenter: SKIRVIN, Timothy Michael (F)

Session Classification: Basic IT services

Track Classification: Basic IT Services

Contribution ID: 45

Type: **Oral presentation**

HEPiX Bit Preservation Working Group

Thursday, 31 October 2013 09:45 (30 minutes)

The goal of the HEPiX Bit Preservation Working Group is to share ideas, practices and experience on bit stream preservation activities across sites providing long-term and large-scale archive services. Different aspects should be covered like: technology used for long-term archiving, definition of reliability, mitigation of data loss risks, monitoring/verification of the archive contents, procedures for recovering unavailable and/or lost data, procedures for archive migration to new-generation technology.

The Working Group is producing a survey on existing practices across HEPiX and WLCG sites responsible for large-scale long-term archiving, and will present its initial findings at the HEPiX Fall'2013 workshop.

Primary authors: OZEROV, Dmitry (DESY); CANCIO MELIA, German (CERN)

Presenter: CANCIO MELIA, German (CERN)

Session Classification: Storage and file systems

Track Classification: Storage & Filesystems

Contribution ID: 46

Type: **Oral presentation**

NDGF Site Report

Monday, 28 October 2013 11:15 (15 minutes)

Overview of recent developments in the distributed NDGF Tier1. Might include a closer look at running Atlas computing on a couple of different HPC resources.

Primary author: WADENSTEIN, Erik Mattias (Unknown)

Presenter: WADENSTEIN, Erik Mattias (Unknown)

Session Classification: Site reports

Track Classification: Site reports

Contribution ID: 47

Type: **Oral presentation**

DESY Site report

Monday, 28 October 2013 11:30 (15 minutes)

Fall 2013 DESY Site report

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

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

Session Classification: Site reports

Track Classification: Site reports

Contribution ID: 49

Type: **Oral presentation**

RAL Tier1 Site Report

Monday, 28 October 2013 11:45 (15 minutes)

Update form RAL

Primary author: BLY, Martin (STFC-RAL)

Presenter: BLY, Martin (STFC-RAL)

Session Classification: Site reports

Track Classification: Site reports

Contribution ID: 50

Type: **Oral presentation**

The migration from ELFms/Quattro to Agile Infrastructure

Wednesday, 30 October 2013 15:00 (30 minutes)

The life cycle of the ELFSms (extremely large fabric management system) is reaching its end. This set of tools provided to manage machines in the CERN Computer Centre has reached the end-of-life and a new Configuration Management System is going to take its place.

The new Configuration Management System is going to change drastically the way we manage machines in the CERN Computer Centre and a migration is necessary to be implemented. To ensure a smooth migration the CERN Configuration Management System Team has been doing several activities to anticipate the switch off of the ELFms.

In this talk we will introduce the new Configuration Management System and its principal components. We will compare it with ELFSms/Quattro system and provide inner details how the migration from ELFSms/Quattro is being implemented. The details will include of what have been done so far and what are the plans for the future.

Primary author: GOMES GOUVEIA, Vitor Emanuel (CERN)

Presenter: GOMES GOUVEIA, Vitor Emanuel (CERN)

Session Classification: Basic IT services

Track Classification: Basic IT Services

Contribution ID: 51

Type: **Oral presentation**

UK GridPP Tier 2s Status Report

Tuesday, 29 October 2013 09:30 (15 minutes)

An update from the UK GridPP Tier 2s

Primary author: Dr BREW, Chris (STFC - Science & Technology Facilities Council (GB))

Presenter: Dr BREW, Chris (STFC - Science & Technology Facilities Council (GB))

Session Classification: Site reports

Track Classification: Site reports

Contribution ID: 52

Type: **Oral presentation**

Deploying perfSONAR-PS for WLCG: An Overview

Tuesday, 29 October 2013 17:00 (30 minutes)

The WLCG infrastructure has evolved from its original restrictive network topology, based on the MONARC model, to a more interconnected system, where data movement between regions or countries does not necessarily need to involve T1 centers. While this evolution brought obvious advantages, especially in terms of flexibility for the LHC experiment's data management systems, it also raises the question of how to monitor and troubleshoot the increasing number of possible network paths, in order to provide a global, reliable network service.

The perfSONAR network monitoring system (specifically the perfSONAR-PS implementation) has been 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 discuss our activity to deploy a perfSONAR-PS based network monitoring infrastructure, in the scope of the WLCG Operations Coordination initiative: we will motivate the main choices we agreed in terms of configuration and management, describe the additional tools we developed to complement the standard packages and present the status of the deployment, together with the possible future evolution.

Summary

Provide an overview of perfSONAR-PS monitoring for WLCG

Primary authors: MC KEE, Shawn (University of Michigan (US)); Dr CAMPANA, Simone (CERN)

Co-author: FOR THE, WLCG perfSONAR Deployment Task Force (WLCG)

Presenter: MC KEE, Shawn (University of Michigan (US))

Session Classification: Security and networking

Track Classification: Security & Networking

Contribution ID: 53

Type: **Oral presentation**

Experience with dynamically provisioned worker nodes at the RAL Tier 1

Friday, 1 November 2013 10:55 (25 minutes)

Even with the growing interest in cloud computing, grid-based submission to traditional batch systems is still the primary way for the experiments to run jobs at WLCG sites. Integrating a batch system with virtualised worker nodes on a cloud potentially offers sites many benefits. At RAL we have recently investigated making opportunistic use of a private StratusLab cloud when it has unused resources and there are idle jobs in the batch system. Our ability to do this is greatly simplified due to our migration of the batch system to HTCondor, currently in progress. Here we describe the work that has been done so far, present preliminary results and discuss some of the issues raised by the testing, including virtualisation overheads, fairshares, virtual machine lifetimes, and monitoring requirements for dynamic environments.

Primary author: LAHIFF, Andrew David (STFC - Science & Technology Facilities Council (GB))

Co-authors: COLLIER, Ian (UK Tier1 Centre); ALEXANDROV, Orlin (STFC - Science & Technology Facilities Council (GB))

Presenter: LAHIFF, Andrew David (STFC - Science & Technology Facilities Council (GB))

Session Classification: Grids, clouds, virtualisation

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: 54

Type: **Oral presentation**

The University of Wisconsin Madison CMS T2 site report

Monday, 28 October 2013 12:00 (15 minutes)

As a major WLCG/OSG T2 site, the University of Wisconsin Madison CMS T2 has provided very productive and reliable services for CMS MonteCarlo production/processing, and large scale global CMS physics analysis using high throughput computing, highly available storage system, and scalable distributed software systems. The close integration of the CMS specific T2 resources with that of the UW campus grid (GLOW), and the strong collaboration with the UW Condor Team has resulted in many significant contributions to the CMS computing and the global grid users community. UW CMS T2 continues to lead and contribute to several key areas of the LHC computing/software aimed at providing CMS data access to users transparently and round the clock, and serving as a testbed for expanding the CMS MC production by successfully integrating commercial cloud resources into the T2 for trial use. An update on the current status of and activities at the UW Tier-2 will be presented.

Primary authors: MOHAPATRA, Ajit Kumar (University of Wisconsin (US)); BRADLEY, Daniel Charles (University of Wisconsin (US)); Prof. DASU, Sridhara (University of Wisconsin (US)); SARANGI, Tapas (University of Wisconsin (US))

Presenter: MOHAPATRA, Ajit Kumar (University of Wisconsin (US))

Session Classification: Site reports

Track Classification: Site reports

Contribution ID: 55

Type: **Oral presentation**

Logstash and ElasticSearch deployment scenario at GSI

Wednesday, 30 October 2013 14:30 (30 minutes)

This talk will present some use cases for Logstash and ElasticSearch. In particular we will show how we are using these tools to collect, parse, index and analyze logs of different services: Grid Engine, SSH, Apache and Cisco Firewalls, among the others. Moreover, the introduction of the version three of Kibana, as an HTML plus Javascript interface, has improved the analytics capabilities of ElasticSearch.

Primary author: DESSALVI, Matteo (GSI)

Presenter: DESSALVI, Matteo (GSI)

Session Classification: Basic IT services

Track Classification: Basic IT Services

Contribution ID: 56

Type: **Oral presentation**

Identity Management in Future Scientific Collaborations

Wednesday, 30 October 2013 10:00 (30 minutes)

Scientific collaborations are evolving to a model where large, multi-dimensional data sets are analyzed in whole or in part by relatively small groups of researchers. These groups are often without the expertise and/or resources to develop and maintain a sophisticated IT infrastructure and represent the growing “long tail of science”. The presentation will discuss the evolving structures and new approaches, policies, and services needed to lower the barriers to providing Identity Management to future collaborations.

Primary author: COWLES, Bob (Indiana University / CACR)

Co-authors: JACKSON, Craig (Indiana University / CACR); WELCH, Von (Indiana University / CACR)

Presenter: COWLES, Bob (Indiana University / CACR)

Session Classification: Security and networking

Track Classification: Security & Networking

Contribution ID: 57

Type: **Oral presentation**

GridKa Site Report

Tuesday, 29 October 2013 09:00 (15 minutes)

Current status and latest news at GridKa, e.g.:

- Hardware status
- Storage systems
- Batch system

Primary author: PETZOLD, Andreas (KIT - Karlsruhe Institute of Technology (DE))

Presenter: PETZOLD, Andreas (KIT - Karlsruhe Institute of Technology (DE))

Session Classification: Site reports

Track Classification: Site reports

Contribution ID: 58

Type: **Oral presentation**

AGLT2 Site Report

Monday, 28 October 2013 10:15 (15 minutes)

Fall 2013 ATLAS Great Lakes Tier 2 site report covering recent network updates, work with AFS on ZFS, new provisioning with cobbler and cfengine, recent experiences with dCache issues, and the usual statistics/status information.

Primary author: Mr MEEKHOF, Benjeman (University of Michigan)

Co-authors: BALL, Robert (University of Michigan (US)); MC KEE, Shawn (University of Michigan (US))

Presenter: Mr MEEKHOF, Benjeman (University of Michigan)

Session Classification: Site reports

Track Classification: Site reports

Contribution ID: 59

Type: **Oral presentation**

Safety in the Data Center

Monday, 28 October 2013 14:50 (25 minutes)

We describe a recent rack installation incident at the RACF and its effects on facility operations. A draft proposal to address safety issues will also be discussed.

Primary author: Dr WONG, Tony (Brookhaven National Laboratory)

Presenter: Dr WONG, Tony (Brookhaven National Laboratory)

Session Classification: IT facilities and business continuity

Track Classification: IT Facilities & Business Continuity

Contribution ID: 60

Type: **Oral presentation**

Operating Dedicated Data Centers - Is it cost-effective?

Monday, 28 October 2013 14:25 (25 minutes)

The advent of cloud computing centers such as Amazon's EC2 and Google's Computing Engine has elicited comparisons with dedicated computing clusters. Discussions on appropriate usage of cloud resources (both academic and commercial) and costs have ensued. This presentation discusses a detailed analysis of the costs of operating and maintaining the RACF (RHIC and ATLAS Computing Facility) compute cluster at Brookhaven National Lab and compares them with the cost of cloud computing resources under various usage scenarios. An extrapolation of likely future cost effectiveness of dedicated computing resources is also presented.

Primary author: Dr WONG, Tony (Brookhaven National Laboratory)

Co-authors: Mr ZAYTSEV, Alexandr (Brookhaven National Laboratory (US)); HOLLOWELL, Christopher (Brookhaven National Laboratory); ERNST, Michael (Unknown); Mr HOGUE, Richard (Brookhaven National Laboratory); STRECKER-KELLOGG, William (Brookhaven National Lab)

Presenter: Dr WONG, Tony (Brookhaven National Laboratory)

Session Classification: IT facilities and business continuity

Track Classification: IT Facilities & Business Continuity

Contribution ID: **61**

Type: **not specified**

Welcome address

Monday, 28 October 2013 09:00 (20 minutes)

Primary author: NEAL, Homer (University of Michigan (US))

Presenter: NEAL, Homer (University of Michigan (US))

Session Classification: Miscellaneous

Track Classification: Miscellaneous

Contribution ID: **62**

Type: **not specified**

Workshop logistics

Monday, 28 October 2013 09:20 (10 minutes)

Primary author: Dr MC KEE, Shawn (University of Michigan (US))

Presenter: Dr MC KEE, Shawn (University of Michigan (US))

Session Classification: Miscellaneous

Track Classification: Miscellaneous

Contribution ID: **63**

Type: **not specified**

Workshop wrap-up

Friday, 1 November 2013 12:10 (15 minutes)

Primary authors: Dr MEINHARD, Helge (CERN); PHILPOTT, Sandy (JLAB)

Presenter: Dr MEINHARD, Helge (CERN)

Session Classification: Miscellaneous

Track Classification: Miscellaneous

Contribution ID: 64

Type: **Oral presentation**

Experiences with Ceph at the ATLAS Midwest Tier 2

Thursday, 31 October 2013 15:55 (25 minutes)

This talk will cover our deployment of Ceph, a highly-scalable next-generation distributed filesystem, at Midwest Tier 2 to back-end various projects. We'll talk about our experience deploying Ceph, performance benchmarks, and some thoughts about where we would like to go next.

Primary author: BRYANT, Lincoln (University of Chicago (US))

Presenter: BRYANT, Lincoln (University of Chicago (US))

Session Classification: Storage and file systems

Track Classification: Storage & Filesystems

Contribution ID: 65

Type: **not specified**

AI monitoring update

Monday, 28 October 2013 15:40 (20 minutes)

Summary

We will give an update on monitoring and analytics, in particular the recent progress on mix-and-matching different open source tools (e.g. hadoop, elasticsearch, flume, kibana) for the various 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 the different tools to improve their daily activities.

Presenter: PALADIN, Massimo (CERN)

Session Classification: Basic IT services

Track Classification: Basic IT Services

Contribution ID: **66**

Type: **not specified**

IHEP Site Report

Tuesday, 29 October 2013 10:15 (15 minutes)

Summary

The report covers the new activities on storage and scheduler of IHEP site this half year.

Presenter: SHI, Jingyan

Session Classification: Site reports

Track Classification: Site reports