

Hepix 2006

Report of Contributions

Contribution ID: 0

Type: **not specified**

BNL

Monday, October 9, 2006 9:30 AM (20 minutes)

Presenter: WITHERS, Alexander (BNL)

Session Classification: Site Reports

Contribution ID: 1

Type: **not specified**

CERN

Monday, October 9, 2006 9:50 AM (20 minutes)

Presenter: MEINHARD, Helge (CERN)

Session Classification: Site Reports

Contribution ID: 2

Type: **not specified**

Fermilab

Monday, October 9, 2006 10:10 AM (20 minutes)

Presenter: GIACCHETTI, Lisa (Fermilab)

Session Classification: Site Reports

Contribution ID: 3

Type: **not specified**

GridKa

Monday, October 9, 2006 11:00 AM (20 minutes)

Presenter: ALEF, Manfred (GridKa)

Session Classification: Site Reports

Contribution ID: 4

Type: **not specified**

Jlab

Monday, October 9, 2006 11:20 AM (20 minutes)

Presenter: PHILPOTT, Sandy (Jlab)

Session Classification: Site Reports

Contribution ID: 5

Type: **not specified**

TRIUMF

Monday, October 9, 2006 11:40 AM (20 minutes)

Presenter: KOST, Corrie (Triumf)

Session Classification: Site Reports

Contribution ID: 6

Type: **not specified**

Using RT to Manage Installation Workflow

Tuesday, October 10, 2006 2:00 PM (30 minutes)

We had a need to more formally manage the workflow of installation tasks, because there had gotten to be so many happening simultaneously that confusion was resulting. We modeled the workflow using RT, the Request Tracking system that we use for user requests. The result is a relatively lightweight and flexible system that gives planners a “dashboard” of the status of all active projects, and the information they need to execute the task.

Presenter: BOEHEIM, Chuck (SLAC)

Session Classification: Core Services and Infrastructure

Contribution ID: 7

Type: **not specified**

Service Level Status - A Real-time status Display for IT

Tuesday, October 10, 2006 11:30 AM (30 minutes)

Nowadays, IT departments provide, and people use many various computing services of more and more heterogeneous nature. And there is a growing need of having a common display that groups these different services and reports about their status and availabilities in a uniform way. At CERN, it led to launching the SLS project.

Service Level Status Overview (SLS) is a web-based tool that dynamically shows availability, basic information and statistics about various IT services, as well as dependencies between them.

The presentation starts with a short description of the project, its goals, architecture, and users. Then, the concepts of subservices, metaservices, dependencies, service availability etc. are introduced, followed by a demonstration of the system and an explanation of how to add a service to SLS. The talk ends with a information on how SLS could be used by other HEP institutes.

Presenter: LOPIENSKI, Sebastian (CERN)

Session Classification: Core Services and Infrastructure

Contribution ID: 9

Type: **not specified**

Network Security Monitoring with Sguil

Wednesday, October 11, 2006 4:15 PM (30 minutes)

Most mid- or large-sized organizations conduct some sort of network monitoring for security purposes. Traditional Intrusion Detection Systems (IDS) tell only part of the story, leaving analysts to perform complex and time-consuming data-mining operations from multiple sources just to answer simple questions about IDS alerts. This talk presents a more efficient model that uses the open source Sguil software to optimize the process for analyst time and efficiency.

Presenter: BIANCO, David (Jlab)

Session Classification: Cyber Security/Authentication

Contribution ID: 10

Type: **not specified**

Support of Kerberos 5 Authenticated Environment by TORQUE

Wednesday, October 11, 2006 9:30 AM (30 minutes)

TORQUE is a successor of the OpenPBS batch queuing system, available as an Open Source product. Despite the wide spread usage of TORQUE as Job Management System on computational farms and LHC grid installations, this batch system does not support any advanced authentication mechanisms.

We show two possibilities, how to redesign the existing source code in order to add Kerberos 5 authentication support for batch jobs.

The first way uses local server-client RPC connections while the second one makes use of the Authenticated Remote Control tool (ARCV2).

The described modifications have been successfully deployed in the local computing infrastructure of the H1 Collaboration at DESY. This provides on identical environment for batch jobs and user desktop processes.

Presenter: LOBODZINSKI, Bogdan (DESY)

Session Classification: Compute Clusters/Storage

Contribution ID: 11

Type: **not specified**

RACF's PXE Installation Management System

Wednesday, October 11, 2006 9:00 AM (30 minutes)

The BNL RHIC/ATLAS Computing Facility (RACF) Central Analysis/Reconstruction Server (CAS/CRS) Farm is a large scale computing cluster currently consisting of ~2000 multiprocessor hosts running Scientific Linux. Besides providing for computation, the CAS/CRS systems' local disk drives are used by network distributed data systems such as dCache, ROOTD and XROOTD to store considerable amounts of data (presently ~400 TB). The sheer number of systems in the farm, combined with our distributed storage model complicates network installation management. This presentation describes the system developed at RACF to fully automate and simplify management of the PXE installation process.

Presenter: HOLLOWELL, Christopher (BNL)

Session Classification: Compute Clusters/Storage

Contribution ID: 12

Type: **not specified**

Grid Security in WLCG and EGEE

Thursday, October 12, 2006 4:00 PM (30 minutes)

This talk will present the current status, plans and issues for Grid Security in WLCG and EGEE. This will include Authentication, Authorization, Policy and Operational Security.

Presenter: KELSEY, David (CCLRC/RAL)

Session Classification: Grid Projects

Contribution ID: 13

Type: **not specified**

GridX1: A Canadian Computational grid for HEP Applications

Thursday, October 12, 2006 9:00 AM (30 minutes)

GridX1 is a Canadian computational grid which combines the shared resources of several Canadian research institutes for the primary purpose of executing HEP applications. With more than two years of production experience, GridX1 has demonstrated the successful application of Globus Toolkit (GT) v.2 cluster gatekeepers managed by a Condor-G resource brokering system. A novel feature of the project was a resource brokering interface to the LHC Compute Grid, which was used during Data Challenge 2 to route ATLAS jobs to the Canadian resources without having dedicated Compute Elements at each cluster. Further, independent Condor-G resource brokers have been implemented to manage the Canadian ATLAS and BaBar MC production systems. Finally, our recent efforts have been directed toward building a service-oriented grid using GT4, including a WS-MDS registry service and WS-GRAM enabled metaschedulers built upon Condor and GridWay.

Presenter: GABLE, I. (University of Victoria/HEPnet Canada)

Session Classification: Grid Projects

Contribution ID: 14

Type: **not specified**

From a Spark in Vacuum to Sparking the Vacuum

Tuesday, October 10, 2006 6:00 PM (30 minutes)

Presenter: DYLLA, Fred (Jlab)

Session Classification: Networking Dinner at Newport News City Center Marriott

Contribution ID: 15

Type: **not specified**

NERSC Site Report

Monday, October 9, 2006 1:30 PM (20 minutes)

Presenter: WHITNEY, Cary (NERSC)

Session Classification: Site Reports II

Contribution ID: 16

Type: **not specified**

Scientific Linux Update

Tuesday, October 10, 2006 10:30 AM (30 minutes)

In this talk, we will present the status of Scientific Linux, focusing on relevant changes in the past six months. Next, we will also present current projects with SL, focusing on SL 5.x and scientific applications. To conclude, we will talk about future enhancements.

Presenter: FROMM, Jim (Fermilab)

Session Classification: Core Services and Infrastructure

Contribution ID: 17

Type: **not specified**

Experiences with SpamCop

Tuesday, October 10, 2006 9:00 AM (30 minutes)

Spamcop is a popular tool on the internet for reporting “spammers” to the ISP’s. Several HEP sites have signed up SpamCop as a method of detecting spam. Unfortunately, the way Fermilab processes bounced spam email it can appear to Spamcop that Fermilab is an initiator of spam. This has occurred several times in the past year. To resolve the last incident, we requested several sites add Fermilab to their servers whitelists. In addition, we adjusted our mail gateways to eliminate bounced spam messages whenever possible. We are also looking into improving our spam filtering systems to minimize any spam that might get through and subsequently forwarded to another site. We would like to discuss getting together a list of other HEP sites and their email servers, and sharing this list so that we don’t inadvertently block each others email transmissions.

Presenter: FROMM, Jim (Fermilab)

Session Classification: Core Services and Infrastructure

Contribution ID: **18**

Type: **not specified**

NIKHEF Site Report

Monday, October 9, 2006 1:50 PM (20 minutes)

Presenter: KUIPERS, Paul (NIKHEF)

Session Classification: Site Reports II

Contribution ID: **19**

Type: **not specified**

CCLRC-RAL Site Report

Monday, October 9, 2006 2:10 PM (20 minutes)

Presenter: BLY, Martin (CCLRC-RAL)

Session Classification: Site Reports II

Contribution ID: **20**Type: **not specified**

TWiki at CERN

Tuesday, October 10, 2006 11:00 AM (30 minutes)

The Database and Engineering Services (DES) Group of the IT Department at CERN supports and maintains a CERN TWiki.

This presentation will cover the history of TWiki at CERN, facts about the system, the technical setup, problems we face and our plans for finding a solution to them.

Presenter: HANSBAKK, Hege (CERN)

Session Classification: Core Services and Infrastructure

Contribution ID: 21

Type: **not specified**

INFN-CNAF Site Report

Monday, October 9, 2006 4:20 PM (20 minutes)

Presenter: CHIERICI, Andrea (INFN-CNAF)

Session Classification: Site Reports II

Contribution ID: 22

Type: **not specified**

NGF NERSC's Global Filesystem and PDSF

Wednesday, October 11, 2006 11:30 AM (30 minutes)

I would like to explain a bit about our global filesystem and it's use on PDSF. Also about how this filesystem can be extended to other sites/labs. Our filesystem is GPFS, but the concept can also be extended to Lustre or other cluster filesystems.

Presenter: LANGLEY, Tom (NERSC)

Session Classification: Compute Clusters/Storage

Contribution ID: 23

Type: **not specified**

INFN Site Report

Monday, October 9, 2006 2:30 PM (20 minutes)

Presenter: GOMEZEL, Roberto (INFN)

Session Classification: Site Reports II

Contribution ID: 25

Type: **not specified**

The Stakkato Intrusions

Wednesday, October 11, 2006 3:30 PM (45 minutes)

During 15 months, from late 2003 until early 2005, hundreds of supercomputing sites, universities and companies worldwide were hit by a series of intrusions, with the perpetrator leapfrogging from site to site using stolen ssh passwords. These are collectively known as the Stakkato intrusions, and include the Teragrid Incident and the Cisco IOS source code theft, both of which received widespread attention from the media. This talk will cover case studies of performed intrusions, an analysis of why Stakkato could be so successful, and the story of how the suspect was finally tracked down and caught.

Presenter: NIXON, Leif

Session Classification: Cyber Security/Authentication

Contribution ID: 27

Type: **not specified**

GSI/Darmstadt Site Report

Monday, October 9, 2006 3:20 PM (20 minutes)

Presenter: SCHOEN, Walter (GSI/Darmstadt)

Session Classification: Site Reports II

Contribution ID: 28

Type: **not specified**

High Availability Methods at GSI

Tuesday, October 10, 2006 2:30 PM (30 minutes)

This presentation gives an overview about the methods used to ensure the high availability of important services such as data base, web service, central file server a. o. Apart from commercial products for certain systems (Oracle, Exchange) different open source linux tools (heartbeat, drbd, mon) are combined with monitoring and hardware methods and adapted to our special needs.

Presenter: MIERS, Karin (GSI/Darmstadt)

Session Classification: Core Services and Infrastructure

Contribution ID: 29

Type: **not specified**

Using Quattor to manage a grid (EGEE) Fabric

Tuesday, October 10, 2006 3:30 PM (30 minutes)

Deploying grid services means managing a potentially large number of machines that partially share their configuration. A tool is needed not only to install but to maintain such a configuration. Quattor, developed as part of EDG, is such a tool. This talk will focus on the LCG/gLite support in Quattor.

Presenter: JOUVIN, Michel (LAL / IN2P3)

Session Classification: Core Services and Infrastructure

Contribution ID: 30

Type: **not specified**

Managing system history and problem tracking with SVN/Trac

Tuesday, October 10, 2006 1:30 PM (30 minutes)

This talk will present LAL experience to address the need to track system configuration changes and link this with an issue tracker, using a combination of Subversion and Trac.

Presenter: JOUVIN, Michel (LAL / IN2P3)

Session Classification: Core Services and Infrastructure

Contribution ID: **31**Type: **not specified**

GridPP

Thursday, October 12, 2006 9:30 AM (30 minutes)

GridPP is a UK e-Science project which started in 2001 with the aim of developing and operating a production Grid for UK Particle Physicists. It is aligned with the EGEE infrastructure and the WLCG Project but also works with current running experiments and theorists. GridPP aims to provide an environment in which all UK particle physicists can do their analysis, share data, etc, and the UK can also contribute to the worldwide collaboration and activities of their experiments .

Presenter: GORDON, John (CCLRC-RAL)

Session Classification: Grid Projects

Contribution ID: 32

Type: **not specified**

Benchmark Updates

Wednesday, October 11, 2006 2:00 PM (30 minutes)

This talk will present the current state of the art of benchmarking at CERN. We will explain our benchmarking procedures, review our latest results and talk about where we are going from here. As part of the results review, we will comment on the current CPU trends and we will talk about the increasingly important power consumption.

Presenter: MEINHARD, Helge (CERN)

Session Classification: Compute Clusters/Storage

Contribution ID: 33

Type: **not specified**

The EGEE Grid Infrastructure

Thursday, October 12, 2006 10:30 AM (30 minutes)

The EGEE grid infrastructure is in constant production use with significant workloads, not only for High Energy Physics but for many other scientific applications. An overview of the EGEE project, the infrastructure itself, and how it is being used will be given. Several applications rely on a long term infrastructure being in place; the current ideas of how this may be achieved will be discussed.

Presenter: BIRD, Ian (CERN)

Session Classification: Grid Projects

Contribution ID: 34

Type: **not specified**

Issues and problems around Grid site management

Thursday, October 12, 2006 1:30 PM (1 hour)

The problems of grid site reliability and availability are becoming the biggest outstanding issue in building a reliable grid service. This is particularly important for WLCG where specific reliability targets are set. This talk will outline the scope of the problems that need to be addressed, and point out potential areas where HEPiX members can contribute, and will seek input on how we can address some of the problems.

Presenter: Dr BIRD, Ian (CERN)

Session Classification: Grid Projects

Contribution ID: 35

Type: **not specified**

Virtual Machines in a Distributed Environment

Thursday, October 12, 2006 11:00 AM (1 hour)

Presenter: TSUGAWA, Mauricio (University of Florida)

Session Classification: Grid Projects

Contribution ID: 36

Type: **not specified**

What is TRAC?

Tuesday, October 10, 2006 9:30 AM (30 minutes)

This talk will present Trac, a unique open source tool combining a wiki, an issue tracker, a Subversion client and a roadmap manager. More than a tool, Trac is an extensible framework based on plugins. LAL is currently using this tool both for software development and system administration.

Presenter: JOUVIN, Michel (LAL / IN2P3)

Session Classification: Core Services and Infrastructure

Contribution ID: 37

Type: **not specified**

Spam - Statistics and Fighting Methods

Tuesday, October 10, 2006 4:00 PM (30 minutes)

Presenter: SCHOEN, Walter (GSI/Darmstadt)

Session Classification: Core Services and Infrastructure

Contribution ID: **38**

Type: **not specified**

DAPNIA Site Report

Monday, October 9, 2006 3:40 PM (20 minutes)

Presenter: MICOUT, Pierrick (CEA DAPNIA Saclay)

Session Classification: Site Reports II

Contribution ID: **39**

Type: **not specified**

FermiGrid - Status and Plans

Thursday, October 12, 2006 2:30 PM (30 minutes)

FermiGrid is the Fermilab Campus Grid. This talk will discuss the current state of FermiGrid and plans for the upcoming year.

Presenter: CHADWICK, Keith (Fermilab)

Session Classification: Grid Projects

Contribution ID: 40

Type: **not specified**

Open Science Grid Progress and Vision

Thursday, October 12, 2006 3:30 PM (30 minutes)

This talk will detail recent Open Science Grid progress and outline the vision for the upcoming year.

Presenter: CHADWICK, Keith (Fermilab)

Session Classification: Grid Projects

Contribution ID: 41

Type: **not specified**

Planning for Hall D: The Hazards of Fast Tape Drives

Wednesday, October 11, 2006 10:30 AM (30 minutes)

The upgrade to Jefferson Lab will require a hardware refresh of the mass storage system in order to handle the higher volume of data from new experiments and simulations. The next generation, higher capacity tape drives are also significantly faster, a fact that has implications for almost all parts of the mass storage system. This talk examines the performance tuning required to make efficient use of these drives and underscores some of the particular needs of tape-based storage systems used by most experiments.

Presenter: HESS, Bryan (Jefferson Lab)

Session Classification: Compute Clusters/Storage

Contribution ID: 42

Type: **not specified**

Porting to and Running Applications on 64 Bit Platforms

Wednesday, October 11, 2006 11:00 AM (30 minutes)

The author describes his recent experience porting software packages to and running these packages on 64 bit machines with Solaris and Linux. Issues discussed include code modification, compiling, operating system requirements, and performance comparisons with 32 bit machines.

Presenter: TIMMER, Carl (Jefferson Lab)

Session Classification: Compute Clusters/Storage

Contribution ID: 43

Type: **not specified**

Storage Class : Problematic and Implementation at CCIN2P3

Wednesday, October 11, 2006 1:30 PM (30 minutes)

Storage Classes attempt to represent storage use cases for a given experiment. It is considered harmful to match the storage classes to real life storage system especially if the latter is based on path to get the storage configuration of a file.

This presentation aims to define the problematic of Storage Classes, explain one possible solution which is implemented at CCIN2P3 and discuss the pros and cons.

Presenter: SCHAEFFER, Jonathan (CC-IN2P3)

Session Classification: Compute Clusters/Storage

Contribution ID: 44

Type: **not specified**

Scientific Linux Inventory Project (SLIP)

Tuesday, October 10, 2006 4:30 PM (30 minutes)

This talk will discuss the effort to provide an inventory of all Linux machines at Fermilab. We will describe the motivation for the project, the package we selected, and the current state of the project.

Presenter: FROMM, Jim (Fermilab)

Session Classification: Core Services and Infrastructure

Contribution ID: 45

Type: **not specified**

IHEPCCC

Friday, October 13, 2006 10:00 AM (30 minutes)

Presenter: SOBIE, Randy

Contribution ID: 46

Type: **not specified**

Testing the UK Tier 2 Data Transfer and Storage Infrastructure

Friday, October 13, 2006 9:30 AM (30 minutes)

When the LHC experiments start taking data next year the Tier 2 sites in the UK (and elsewhere) will need to be able to receive and transmit data at unprecedented rates and reliabilities. We present the efforts in the UK to test the disk to disk transfer rates between Tier 2 sites along with some of the lessons learnt and results obtained.

Presenter: BREW, Chris (CCLRC - RAL)

Session Classification: Grid Projects II

Contribution ID: 47

Type: **not specified**

Welcome to JLab

Monday, October 9, 2006 9:00 AM (20 minutes)

Presenter: WHITNEY, Roy (Jefferson Lab)

Session Classification: Welcome

Contribution ID: 48

Type: **not specified**

Recent Fabric Management Improvements at CERN

Wednesday, October 11, 2006 2:30 PM (30 minutes)

This talk will describe some improvements to the monitoring and management of the storage and CPU services in the following areas

- use of SMART for disk monitoring
- integration of disk server monitoring and storage system management
- transmission of Grid job memory requirements to the local workload management

Presenter: CASS, Tony (CERN)

Session Classification: Compute Clusters/Storage

Contribution ID: 49

Type: **not specified**

SLAC Site Report

Monday, October 9, 2006 4:00 PM (20 minutes)

Presenter: BOEHEIM, Chuck (SLAC)

Session Classification: Site Reports II

Contribution ID: 50

Type: **not specified**

LAL Site Report

Monday, October 9, 2006 4:40 PM (20 minutes)

Presenter: Mr JOUVIN, Michel (LAL / IN2P3)

Session Classification: Site Reports II