

HEPiX Spring 2018 Workshop

Report of Contributions

Contribution ID: **89**

Type: **not specified**

DESY Site Report

Tuesday, May 15, 2018 9:00 AM (20 minutes)

News about what happened at DESY during the last months

Desired length

12

Primary author: ESSIGKE, Timm (DESY)

Presenter: ESSIGKE, Timm (DESY)

Session Classification: Site reports

Track Classification: Site Reports

Contribution ID: 90

Type: **not specified**

Fast Distributed Image Reconstruction using CUDA/MPI

Wednesday, May 16, 2018 5:10 PM (20 minutes)

In this work, we present a fast implementation for analytical image reconstruction from projections, using the so-called “backprojection-slice theorem” (BST). BST has the ability to reproduce reliable image reconstructions in a reasonable amount of time, before taking further decisions. The BST is easy to implement and can be used to take fast decisions about the quality of the measurement, i.e., sample environment, beam-line conditions, among others. A synchrotron facility able to measure a three-dimensional dataset Y within few seconds, needs a fast reconstruction algorithm able to provide a fast “preview” of the tomography within the same amount of time. If the experimental conditions are not satisfactory, the quality of the reconstruction will decrease, and the researcher can decide either to make another scan, or to process later the data using advanced reconstruction algorithms or even high quality segmentation methods. The difficulty here is that inversion algorithms depends on the backprojection operator, which is defined as an average through all the x-rays passing at a given pixel. Backprojection presents a high computational complexity of $O(N^3)$ for an image of N^2 pixels. The brute-force approach to compute the backprojection operator can be made extremely slow, even using a GPU implementation. Sophisticated ray-tracing strategies can also be used to make the running time faster and others analytical strategies reduce the backprojection complexity to $O(N^2 \log N)$. The BST approach have the same low complexity of $O(N^2 \log N)$ although easier to implement than his competitors, producing less numerical artifacts and following a more traditional “gridding strategy”.

Desired length

20

Primary authors: Dr MIQUELES, Eduardo (LNLS/CNPEM); Mr MARTINEZ JR., Gilberto (LNLS/CNPEM); Mr FURUSATO, Fernando (LNLS/CNPEM)

Presenter: Mr FURUSATO, Fernando (LNLS/CNPEM)

Session Classification: Computing and batch systems

Track Classification: Computing & Batch Services

Contribution ID: 91

Type: **not specified**

Cyberinfrastructure and China Science and Technology Cloud Plan in Chinese Academy of Sciences

Tuesday, May 15, 2018 2:40 PM (20 minutes)

Chinese Academy of Sciences has 104 research institutes, 12 branch academies, three universities and 11 supporting organizations in 23 provincial-level areas throughout the country. These institutions are home to more than 100 national key labs and engineering centers as well as nearly 200 CAS key labs and engineering centers. Altogether, CAS comprises 1,000 sites and stations across the country.

As the science research methods develops, we are coming to the fourth paradigm of the science research—Data intensive science research. Data, compute and the link of them, network has played an more important role in science research. All the institutes have various demands in cyberinfrastructure.

China Science and Technology Cloud(CSTC) was constructed in order to meet the needs of the research institutes under the Chinese Academy of Sciences and even the whole scientific and technological community in China. It is an IT-based resources management and cloud service platform with smart resource dispatching and user self-service. It constructs a new-generation information infrastructure which is at high speed, dynamic and self-adaptive; speed up the state-level high-performance computing environment development, and achieve one-stop service for scientific computing. It integrates cloud computing and cloud storage facilities to enhance data recovery capabilities of the whole academy's scientific data assets and application systems. It also integrates and gather various scientific and technological information resources, and sets up a smart cloud service platform to provide scientific and technological resources and information services.

CSTC has maintained long-term partnerships with world-class research organizations such as U.S. National Center for Supercomputing Applications and Forschungszentrum Jülich. The predecessor of CSTC, China Science Technology Network (CSTNET) was one of the founding organizations of the Global Ring Network for Advanced Application Development (GLORIAD) which connect the North American with 10Gb/s bandwidth.

Based on CSTC, we build an expandable basic environment which can carry Big Data resources and support Big Data analysis and processing, realizing management and on-line processing of massive scientific data; targeting fields of relevant disciplines, as well as major research projects and special projects of the state and the academy, to deploy a batch of Big Data driven scientific research and application service in the fields of astronomy, biology, high-energy physics, etc..

The CSTC plans to build a test bed for network research and big science research, for example a Dynamic Virtual Dedicat Network for VLBI research, a DMZ for Advanced Light Source, an Open Network Environment for LHC. We are looking forward for further cooperation with global science research institutes.

Desired length

20

Primary authors: Dr WANG, YANG (Computer Network Information Center, Chinese Academy

of Sciences); Prof. XUEHAI, Hong (Institute of Computing Technology, Chinese Academy of Sciences)

Co-author: Mr JINGJING, Li (Computer Network Information Center, Chinese Academy of Sciences)

Presenter: Dr WANG, YANG (Computer Network Information Center, Chinese Academy of Sciences)

Session Classification: Networking and security

Track Classification: Networking & Security

Contribution ID: 92

Type: **not specified**

Scientific Linux update

Monday, May 14, 2018 2:20 PM (20 minutes)

Updates on the status of Scientific Linux

Desired length

20m

Primary author: RIEHECKY, Pat (Fermilab)

Co-author: KING, Bonnie (Fermilab)

Presenter: KING, Bonnie (Fermilab)

Session Classification: End-user services and operating systems

Track Classification: End-User IT Services & Operating Systems

Contribution ID: 93

Type: **not specified**

Computer Security Update

Monday, May 14, 2018 4:50 PM (20 minutes)

This presentation provides an update on the global security landscape since the last HEPiX meeting. It describes the main vectors of risks to and compromises in the academic community including lessons learnt, presents interesting recent attacks while providing recommendations on how to best protect ourselves. It also covers security risks management in general, as well as the security aspects of the current hot topics in computing and around computer security.

This talk is based on contributions and input from the CERN Computer Security Team.

Desired length

20

Primary author: LUEDERS, Stefan (CERN)**Presenter:** LUEDERS, Stefan (CERN)**Session Classification:** Networking and security**Track Classification:** Networking & Security

Contribution ID: 94

Type: **not specified**

CERN Site Report

Tuesday, May 15, 2018 9:20 AM (20 minutes)

News from CERN since the HEPiX Fall 2017 workshop at KEK, Tsukuba, Japan.

Desired length

12

Primary author: DUMITRU, Andrei (CERN)

Presenter: DUMITRU, Andrei (CERN)

Session Classification: Site reports

Track Classification: Site Reports

Contribution ID: 95

Type: **not specified**

SDN implementation plan in China Science and Technology Network

Tuesday, May 15, 2018 4:05 PM (20 minutes)

Scientific activities generate huge data and need to transfer them to some places to research. Traditional networking infrastructure has a defined architecture and can not satisfy such real-time and high-quality transferring requirements.

China Science and Technology Network(CSTNet) was constructed in order to meet the needs of the research institutes under the Chinese Academy of Sciences and even the whole scientific and technological community in China. CSTNet has planned to construct a new-generation infrastructure using some new technology such as SDN, NFV and etc.

CSTNet has started to build a new NOC running system to achieve real-time measurements, monitoring and management of the network flows. The new NOC running system can provide an user interface to enable user to submit networking requirements dynamic and self-adaptive, some network configure task can become effective on time no need to connect to console board.

Dynamic network management will accelerate the integrating of cloud computing and cloud storage facilities because of faster data transfer command delivering and implementing to service research in the fields of astronomy, biology, high-energy physics, etc.

Desired length

15

Primary author: LI, JINGJING

Co-author: Dr WANG, YANG (Computer Network Information Center, Chinese Academy of Sciences)

Presenter: LI, JINGJING

Session Classification: Networking and security

Track Classification: Networking & Security

Contribution ID: 96

Type: **not specified**

WLCG Archival Storage group report

Wednesday, May 16, 2018 9:00 AM (20 minutes)

The group has been formed to tackle two main themes

- establish a knowledge-sharing community for those operating archival storage for WLCG
- understand how to monitor usage of archival systems and optimise their exploitation by experiments

I will report on the recent activities of this group.

Desired length

10

Primary authors: KEEBLE, Oliver (CERN); BAHYL, Vladimir (CERN)

Presenter: BAHYL, Vladimir (CERN)

Session Classification: Storage and file systems

Track Classification: Storage & Filesystems

Contribution ID: 97

Type: **not specified**

INFN-T1 Site report

Tuesday, May 15, 2018 9:40 AM (20 minutes)

A brief update on INFN-T1 site, what is our current status and what is still to be done to reach 100% functionality

Desired length

12

Primary authors: CHIERICI, Andrea (INFN-CNAF); DAL PRA, Stefano (INFN)

Presenter: DAL PRA, Stefano (INFN)

Session Classification: Site reports

Track Classification: Site Reports

Contribution ID: 98

Type: **not specified**

Batch on EOS Extra Resources moving towards production

Wednesday, May 16, 2018 2:40 PM (20 minutes)

At the last HEPix meeting we described the results of a proof of concept study to run batch jobs on EOS disc server nodes. By now we have moved forward towards a production level configuration and the first pre-production nodes have been setup. Beside the relevance for CERN this is also a more general step towards a hyper-converged infrastructure.

Desired length

15

Primary authors: SMITH, David (CERN); SCHULZ, Markus (CERN); KIRYANOV, Andrey; JONES, Ben (CERN); MCCANCE, Gavin (CERN); LAMANNA, Massimo (CERN); ROUSSEAU, Herve (CERN)

Presenter: SCHULZ, Markus (CERN)

Session Classification: Computing and batch systems

Track Classification: Computing & Batch Services

Contribution ID: 99

Type: **not specified**

HSF-WLCG Cost and Performance Modeling Working Group

Wednesday, May 16, 2018 12:10 PM (20 minutes)

The working group has been established and is now working towards a cost and performance model that allows to quantitatively estimate the computing resources needed for HL-LHC and map them towards the cost at specific sites.

The group has defined a short and medium term plan and identified the main tasks. Around the tasks teams with members from experiments and sites have formed and started concrete work. We will report on the goals and status of the working group.

Desired length

15

Primary authors: SCIABA, Andrea (CERN); SCHULZ, Markus (CERN); FLIX MOLINA, Jose (Centro de Investigaciones Energéticas Medioambientales y Tecnol)

Presenter: FLIX MOLINA, Jose (Centro de Investigaciones Energéticas Medioambientales y Tecnol)

Session Classification: Computing and batch systems

Track Classification: Computing & Batch Services

Contribution ID: **100**Type: **not specified**

TRIDENT Tool for collecting and understanding performance hardware counters

Monday, May 14, 2018 3:00 PM (20 minutes)

Trident, a tool to use low level metrics derived from hardware counters to understand Core, Memory and I/O utilisation and bottlenecks. The collection of time series of these low level counters does not induce significant overhead to the execution of the application.

The Understanding Performance team is investigating on a new node characterisation tool, 'Trident', that can look at various low level metrics with respect to the Core, Memory and I/O. Trident uses a three pronged approach to analysing node's utilisation and understand the stress on different parts of the node based on the given job. Currently core metrics such as memory bandwidth, core utilization, active processor cycles, etc., are being collected. Interpretation of this data is often non intuitive. The tool preprocesses the data to make the data usable by developers and site managers without the need of in-depths expertise of CPU and systems architecture details.

Desired length

15

Primary authors: MURALIDHARAN, Servesh (CERN); SMITH, David (CERN)

Presenter: MURALIDHARAN, Servesh (CERN)

Session Classification: End-user services and operating systems

Track Classification: End-User IT Services & Operating Systems

Contribution ID: **101**Type: **not specified**

ExDeMon: a new scalable monitoring tool for the growing CERN infrastructure

Thursday, May 17, 2018 9:00 AM (20 minutes)

When monitoring an increasing number of machines, infrastructure and tools need to be rethought. A new tool, ExDeMon, for detecting anomalies and raising actions, has been developed to perform well on this growing infrastructure. Considerations of the development and implementation will be shared.

Daniel has been working at CERN for more than 3 years as Big Data developer, he has been implementing different tools for monitoring the computing infrastructure in the organisation.

Desired length

20

Primary author: LANZA GARCIA, Daniel (CERN)

Presenter: LANZA GARCIA, Daniel (CERN)

Session Classification: IT facilities

Track Classification: IT Facilities & Business Continuity

Contribution ID: **102**Type: **not specified**

HNSciCloud Status Report

Friday, May 18, 2018 9:20 AM (20 minutes)

The Helix Nebula Science Cloud (HNSciCloud) Horizon 2020 Pre-Commercial Procurement project (<http://www.hnscicloud.eu/>) brings together a group of 10 research organisations to procure innovative cloud services from commercial providers to establish a cloud platform for the European research community.

This 3 year project has recently entered its final phase which will deploy two pilots with a combined capacity of 20,000 cores and 2 PB of storage integrated with the GEANT network at 40Gbps. This presentation will provide an overview of the project, the pilots, the applications being deployed and lessons learned to-date.

Desired length

20

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

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

Session Classification: Clouds, virtualisation, grids

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: **103**Type: **not specified**

The Software Defined Online Storage System at the GridKa WLCG Tier-1 Center

Wednesday, May 16, 2018 9:20 AM (20 minutes)

The computing center GridKa is serving the ALICE, ATLAS, CMS and LHCb experiments as one of the biggest WLCG Tier-1 centers world wide with compute and storage resources. It is operated by the Steinbuch Centre for Computing at Karlsruhe Institute of Technology in Germany. In April 2017 a new online storage system was put into operation. In its current stage of expansion it offers the HEP experiments a capacity of 23 Petabytes of online storage distributed over 16 redundant storage servers with 3900 disks and 50TB SSDs. The storage is connected via two redundant infiniband fabrics to 44 file servers which in turn are connected each via 40Gbit/s and several 100Gbit/s ethernet uplinks to the GridKa backbone network. The whole storage is partitioned into few large file systems, one for each experiment, using IBM Spectrum Scale as software-defined-storage base layer. The system offers a combined read-write performance of 70Gbyte/s. It can be scaled transparently both in size and performance allowing to fulfill the growing needs especially of the LHC experiments for online storage in the coming years.

In this presentation we discuss the general architecture of the storage system and present first experiences with the performance of the system in production use. In addition we present the current plans for expansion of the system.

Desired length

20

Primary author: SUNDERMANN, Jan Erik (Karlsruhe Institute of Technology (KIT))

Presenter: SUNDERMANN, Jan Erik (Karlsruhe Institute of Technology (KIT))

Session Classification: Storage and file systems

Track Classification: Storage & Filesystems

Contribution ID: **104**Type: **not specified**

Run the latest software on a stable environment - A simpler way

Monday, May 14, 2018 2:00 PM (20 minutes)

What do our users want?

One group wants the latest version of foo, but the stable version of bar.

The other group wants the latest version of bar, but the old version of foo.

What have we tried?

SCL

SCL's are great in theory. But in practice they are hard for the packagers. They also make the developers have to jump through several hoops. If something was developed in an SCL environment, it often wouldn't translate into a non-SCL environment.

Containers

Containers are also great in theory. They are especially great for allowing people to run their code in the exact environment on different machines. But for developers, they still have to jump through many hoops to develop on SCL's. And often we restrict what version of foo and bar we have in those containers.

Tarballs and Zip files

We admit it, that's what our developers are really doing. They are pulling down what they need, from who knows where, writing their code around it, and then asking us (or you) to support it. This is very bad for security, as well as for administrators trying to duplicate the environment on another machine.

Modules - The simpler way

The easiest way to explain modules is they are like yum groups, done right.

Using dnf admins are able to install nodejs 6, or nodejs 9. They aren't installed over in /opt/, they are installed in their usual place in /usr/.

Are you ready to move from python36 to python37? Just change the version of the python module and dnf with change it all.

Users will only be allowed to have one version of python, or nodejs, just like normal python or nodejs. But the developers (or their code) won't have to do anything special to use them.

This presentation will go through our up and downs as we've worked on getting a new technology created to help our users.

<https://docs.pagure.org/modularity/>

Desired length

20

Primary author: DAWSON, Troy

Presenter: DAWSON, Troy

Session Classification: End-user services and operating systems

Track Classification: End-User IT Services & Operating Systems

Contribution ID: 105

Type: **not specified**

Automatic for the People: Containers for LIGO software development on the Open Science Grid and other diverse computing resources

Friday, May 18, 2018 11:30 AM (20 minutes)

Distributed research organizations are faced with wide variation in computing environments to support. LIGO has historically resolved this problem by providing RPM/DEB packages for (pre-)production software and coordination between clusters operated by LIGO-affiliated facilities and research groups. This has been largely successful although it leaves a gap in operating system support and in the development process prior to formal point releases.

We describe early developments in LIGO's use of GitLab, GitHub, and DockerHub to continuously deploy researcher-maintained containers for immediate use on all LIGO clusters, the Open Science Grid, and user workstations. Typical latencies are below an hour, dominated by the build-time of the software itself and the client refresh rate of the CernVM File System

Desired length

15 minutes

Primary author: Dr DOWNES, Thomas (University of Wisconsin-Milwaukee)

Presenter: Dr DOWNES, Thomas (University of Wisconsin-Milwaukee)

Session Classification: Clouds, virtualisation, grids

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: **106**Type: **not specified**

Evolution of the Hadoop and Spark platform for HEP

Thursday, May 17, 2018 3:20 PM (20 minutes)

The interest in using Big Data solutions based on Hadoop ecosystem is constantly growing in HEP community. This drives the need for increased reliability and availability of the central Hadoop service and underlying infrastructure provided to the community by the CERN IT department. This contribution will report on the overall status of the Hadoop platform and the recent enhancements and features introduced in many areas including the service configuration, availability, alerting, monitoring and data protection, in order to meet the new requirements posed by the users community.

Desired length

20

Primary author: BARANOWSKI, Zbigniew (CERN)**Presenter:** BARANOWSKI, Zbigniew (CERN)**Session Classification:** Basic IT services**Track Classification:** Basic IT Services

Contribution ID: **107**Type: **not specified**

Monitoring Infrastructure for the CERN Data Centre

Thursday, May 17, 2018 2:20 PM (20 minutes)

Since early 2017, the MONIT infrastructure provides services for monitoring the CERN data centre, together with the WLCG grid resources, and progressively replaces in-house technologies, such as LEMON and SLS, using consolidated open source solutions for monitoring and alarms.

The infrastructure collects data from more than 30k data centre hosts in Meyrin and Wigner sites, with a total volume of 3 TB/day and a rate of 65k documents/sec. It includes OS and hardware metrics, as well as specific IT service metrics. Logs and metrics collection is deployed by default in every machine of the data centre, together with alert reporting. Each machine has a default configuration that can be extended for service-specific data (e.g. for specifically monitoring a database server). Service managers can send custom metrics and logs from their applications to the infrastructure through generic endpoints, and they are provided with an out-of-the-box discovery and visualization interface, data analysis tools and integrated notifications.

The infrastructure stack relies on open source technologies, developed and widely used by the industry and research leaders. Our architecture uses collectd for metric collection, Flume and Kafka for transport, Spark for stream and batch processing, Elasticsearch, HDFS and InfluxDB for search and storage, Kibana and Grafana for visualization, and Zeppelin for analytics. The modularity of collectd provides flexibility to the infrastructure users to configure default and service-specific monitoring, and allows to develop and deploy custom plugins.

This contribution is an updated overview of the monitoring service for CERN data centre. We present our main use cases for collection of metrics and logs. Given that the proposed stack of technologies is widely used, and the MONIT architecture is well consolidated, a main objective is to share the lessons learned and find common monitoring solutions within the community.

Desired length

20

Primary authors: AGUADO CORMAN, Asier (Universidad de Oviedo (ES)); AIMAR, Alberto (CERN); ANDRADE, Pedro (CERN); BRUNDU, Simone (CERN); DELGADO FERNANDEZ, Javier (CERN); FERNANDEZ ALVAREZ, Luis (CERN); GARRIDO BEAR, Borja (CERN); Dr KARAVAKIS, Edward (CERN); KULIKOWSKI, Dominik Marek (Wroclaw University of Science and Technology (PL)); MAGNONI, Luca (CERN)

Presenter: AGUADO CORMAN, Asier (Universidad de Oviedo (ES))

Session Classification: Basic IT services

Track Classification: Basic IT Services

Contribution ID: **109**Type: **not specified**

Benchmarking Working Group. An update

Wednesday, May 16, 2018 11:50 AM (20 minutes)

The benchmarking working group holds biweekly meeting. we are focusing on the health of HS06, fast benchmark and study of a new benchmark to replace HS06 since SPEC has moved to a new family of benchmark

Desired length

20

Primary authors: MICHELOTTO, Michele (Università e INFN, Padova (IT)); ALEF, Manfred (Karlsruhe Institute of Technology (KIT)); GIORDANO, Domenico (CERN)

Presenter: MICHELOTTO, Michele (Università e INFN, Padova (IT))

Session Classification: Computing and batch systems

Track Classification: Computing & Batch Services

Contribution ID: 111

Type: **not specified**

Swiss HPC Tier-2 @ CSCS

Wednesday, May 16, 2018 4:30 PM (20 minutes)

For the past 10 years, CSCS has been providing computational resources for the ATLAS, CMS, and LHCb experiments on a standard commodity cluster.

The High Luminosity LHC upgrade (HL-LHC) presents new challenges and demands with a predicted 50x increase in computing needs over the next 8 to 10 years. High Performance Computing capabilities could help to equalize the computing demands due to their ability to provide specialized hardware and economies of scale. For the past year, CSCS has been running the Tier-2 workload for these experiments on the flagship system Piz Daint, a Cray XC system.

Desired length

20

Primary authors: Mr CONCIATORE, Dino (CSCS (Swiss National Supercomputing Centre)); GILA, Miguel (CSCS (Swiss National Supercomputing Centre)); Mr PETRUSIC, Dario (CSCS (Swiss National Supercomputing Centre)); Mr FERNANDEZ, Pablo (CSCS (Swiss National Supercomputing Centre)); Mr RICCIARDI, Gianni Mario (CSCS (Swiss National Supercomputing Centre))

Presenter: Mr CONCIATORE, Dino (CSCS (Swiss National Supercomputing Centre))

Session Classification: Computing and batch systems

Track Classification: Computing & Batch Services

Contribution ID: 112

Type: **not specified**

A fully High-availability logs/metrics collector @ CSCS

Thursday, May 17, 2018 2:00 PM (20 minutes)

As the complexity of systems increases and the scale of these systems increases, the amount of system level data recorded increases.

Managing the vast amounts of log data is a challenge that CSCS solved with the introduction of a centralized log and metrics infrastructure based on Elasticsearch, Graylog, Kibana, and Grafana. This is a fundamental service at CSCS that provides easy correlation of events bridging the gap from the computation workload to nodes enabling failure diagnosis.

Currently, the Elasticsearch cluster at CSCS is handling more than 22'000'000'000 online documents (one year) and another 20'000'000'000 archived. The integrated environment from logging to graphical representation enables powerful dashboards and monitoring displays.

Desired length

20

Primary author: Mr CONCIATORE, Dino (CSCS (Swiss National Supercomputing Centre))

Presenter: Mr CONCIATORE, Dino (CSCS (Swiss National Supercomputing Centre))

Session Classification: Basic IT services

Track Classification: Basic IT Services

Contribution ID: 113

Type: **not specified**

Purdue University CMS T2 site report

Monday, May 14, 2018 9:30 AM (20 minutes)

Through participation in the Community Cluster Program of Purdue University, our Tier-2 center has for many years been one of the most productive and reliable sites for CMS computing, providing both dedicated and opportunistic resources to the collaboration. In this report we will present an overview of the site, review the successes and challenges of the last year of operation, and outline the perspectives and plans for future developments.

Desired length

12

Primary authors: PIPEROV, Stefan (Purdue University (US)); GOUGH, Erik (Purdue University (US)); ARABGOL, Majid (Purdue University (US)); HACKER, Thomas (Purdue University (US)); NEUMEIS-TER, Norbert (Purdue University (US))

Presenter: PIPEROV, Stefan (Purdue University (US))

Session Classification: Site reports

Track Classification: Site Reports

Contribution ID: 115

Type: **not specified**

Data To Network: building balanced throughput storage in a world of increasing disk sizes

Wednesday, May 16, 2018 10:50 AM (20 minutes)

The ever-decreasing cost of high capacity spinning media has resulted in a trend towards very large capacity storage ‘building blocks’. Large numbers of disks - with up to 60 drives per enclosure being more-or-less standard –indeed allow for dense solutions, maximizing storage capacity in terms of floor space, and can in theory be packed almost exclusively with disks. The result are building blocks with a theoretical gross capacity of about 180 TByte per unit height when employing 12 TByte disks. This density comes at a cost, though: getting the data to and from the disks, via front-end storage management software and through a network, has different scaling characteristics than the gross storage density, and as a result maintaining performance in terms of throughput per storage capacity is an ever more complex challenge. At Nikhef, and for the NL-T1 service, we aim to maintain 12MiB/s/TiB combined throughput, supporting at least 2 read- and 1 write stream per 100TiB netto storage, from any external network source down to the physical disks. Especially this combined read-write operational pattern poses challenges not usually found in commercial deployments. Yet this is the pattern most commonly seen for our scientific applications in the Dutch National e-Infrastructure.

In this study we looked at each of the potential bottlenecks in such a mixed-load storage system: network throughput, limitations in the system bus between CPU, network card, and disk subsystem, at different disk configuration models (JBOD, erasure-encodings, hardware, and software RAID) and the effect on processor load in different CPU architectures. We present the results of different disk configurations and show the limitations of commodity redundancy technologies and how they affect processor load in both x86-64 and PowerPC systems, and how the corresponding system bus design impacts overall throughput.

Combining network and disk performance optimizations we show how high-density commodity components can be combined to build a cost-cutting system without bottlenecks –offering constant-throughput multi-stream performance with over 700TiB netto in just 10U and able to keep a 100Gbps network link full –as a reference architecture for everything from a single Data Transfer Node down to a real distributed storage cluster.

Desired length

20 minutes

Primary author: SUERINK, Tristan (Nikhef National institute for subatomic physics (NL))**Presenter:** SUERINK, Tristan (Nikhef National institute for subatomic physics (NL))**Session Classification:** Storage and file systems**Track Classification:** Storage & Filesystems

Contribution ID: **116**

Type: **not specified**

PIC site report

Tuesday, May 15, 2018 10:00 AM (20 minutes)

News from PIC since the HEPiX Fall 2017 workshop at KEK, Tsukuba, Japan.

Desired length

12

Primary author: FLIX MOLINA, Jose (Centro de Investigaciones Energéti cas Medioambientales y Tecno)

Presenter: FLIX MOLINA, Jose (Centro de Investigaciones Energéti cas Medioambientales y Tecno)

Session Classification: Site reports

Track Classification: Site Reports

Contribution ID: 117

Type: **not specified**

Evolution of technology and markets

Thursday, May 17, 2018 11:05 AM (20 minutes)

A short review of how technology and markets have evolved in areas relevant for HEP computing

Desired length

20 minutes (see comment)

Primary authors: PANZER-STEINDEL, Bernd (CERN); MEINHARD, Helge (CERN)

Presenter: MEINHARD, Helge (CERN)

Session Classification: IT facilities

Track Classification: IT Facilities & Business Continuity

Contribution ID: 118

Type: **not specified**

Proposal for a technology watch WG

Thursday, May 17, 2018 11:25 AM (20 minutes)

Following up from abstract #117, a proposal to form a working group dedicated to technology watch

Desired length

20 minutes (see comment)

Primary authors: MEINHARD, Helge (CERN); PANZER-STEINDEL, Bernd (CERN)

Presenter: MEINHARD, Helge (CERN)

Session Classification: IT facilities

Track Classification: IT Facilities & Business Continuity

Contribution ID: 119

Type: **not specified**

Next generation of large-scale storage services at CERN

Wednesday, May 16, 2018 9:40 AM (20 minutes)

CERN IT Storage (IT/ST) group leads the development and operation of large-scale services based on EOS for the full spectrum of use-cases at CERN and in the HEP community. IT/ST group also provides storage for other internal services, such as Open Stack, using a solution based on Ceph. In this talk we present current operational status, ongoing development work and future architecture outlook for next generation storage services for the users based on EOS —a technology developed and integrated at CERN.

EOS is the home for all physics data-stores for LHC and non-LHC experiments (at present 250PB storage capacity). It is designed to operate at high data rates for experiment data-taking while running concurrent complex production work-loads. EOS also provides a flexible distributed storage back-end and architecture with plugins for tape archival (CTA - evolution and replacement for CASTOR), synchronization&sharing services (CERNBox) and general-purpose filesystem access for home directories (FUSE for Linux and SMB Gateways for Windows and Mac).

CERNBox is the cloud storage front-end for desktop, mobile and web access focused on personal user files, general-purpose project spaces and smaller physics datasets (at present 12K user accounts and 500M files). CERNBox provides simple and uniform access to storage on all modern devices and operating systems. CERNBox is also hub for integration with other services: Collaborative editing —MS Office365 and alternatives: Collabora and OnlyOffice; Web-based analysis —SWAN Jupyter Notebooks with access to computational resources via Spark and Batch; and software distribution via CVMFS.

This storage service ecosystem is designed to provide “total data access”: from end-user devices to geo-aware data lakes for WLCG and beyond. It also provides a foundation for strategic partnerships (AARNet, JRC, ...), new communities such as CS3 (Cloud Storage and Synchronization Services) and new application projects such as Up2University (cloud storage ecosystem for education). CERN Storage technology has been showcased to work with commercial cloud providers such as Amazon, T-Systems (Helix Nebula) or COMTRADE (Openlab) and there is an increasing number of external sites testing the CERN storage service stack in their local computing centers.

This strategy proves very successful with the users and as a result storage services at CERN see exponential growth: CERNBox alone has grown by 450% in 2017. Growing overall demand drive the evolution of the service design and implementation of the full ecosystem: EOS core storage as well as CERNBox and SWAN. Recent EOS improvements include new distributed namespace to provide scaling and high-availability; new robust FUSE module providing client-side caching, lower latency and more IOPs; new workflow engine and many more. CERNBox is moving to micro-service oriented architecture and SWAN is tested with Kubernetes container orchestration.

New developments come together with a constant effort to streamline QA, testing and documentation as well as reduce manual configuration and operational effort for managing large-scale storage services.

Desired length

20

Primary author: MOSCICKI, Jakub (CERN)

Co-authors: BOCCHI, Enrico (CERN); CONTESCU, Cristian (CERN); GONZALEZ LABRADOR, Hugo (CERN); LAMANNA, Massimo (CERN); MASCETTI, Luca (CERN); PETERS, Andreas Joachim (CERN); LO PRESTI, Giuseppe (CERN); ROUSSEAU, Herve (CERN); VALVERDE CAMESELLE, Roberto (Universidad de Oviedo (ES))

Presenter: MOSCICKI, Jakub (CERN)

Session Classification: Storage and file systems

Track Classification: Storage & Filesystems

Contribution ID: 120

Type: **not specified**

Operating a large scale distributed XRootd cache across Caltech and UCSD

Wednesday, May 16, 2018 11:10 AM (20 minutes)

After the successful adoption of the CMS Federation an opportunity arose to cache xrootd requests in Southern California. We present the operational challenges and the lessons learned from scaling a federated cache (a cache composed of several independent nodes) first at UCSD and the scaling and network challenges to augment it to include the Caltech Tier 2 Site. In which would be a first of a kind multisite Xrootd cache which could potentially ease the data management of CMS.

Desired length

20 minutes

Primary authors: FAJARDO HERNANDEZ, Edgar (Univ. of California San Diego (US)); TADEL, Matevz (Univ. of California San Diego (US)); MRAK TADEL, Alja (Univ. of California San Diego (US)); MARTIN, Terrence; WUERTHWEIN, Frank (UCSD)

Presenter: FAJARDO HERNANDEZ, Edgar (Univ. of California San Diego (US))

Session Classification: Storage and file systems

Track Classification: Storage & Filesystems

Contribution ID: 121

Type: **not specified**

Status update of the CERN private cloud

Friday, May 18, 2018 9:00 AM (20 minutes)

CERN runs a private OpenStack Cloud with ~300K cores, ~3000 users and a number of OpenStack services. CERN users can built services using a pool of compute and storage resources using the OpenStack APIs like Ironic, Nova, Magnum, Cinder and Manila, on the other hand CERN cloud operators face some operational challenges at scale in order to offer them. In this talk, you will learn about the status of the CERN cloud, new services and plans for expansion.

Desired length

20

Primary author: TRIGAZIS, Spyridon (CERN)**Presenter:** TRIGAZIS, Spyridon (CERN)**Session Classification:** Clouds, virtualisation, grids**Track Classification:** Grid, Cloud & Virtualisation

Contribution ID: 122

Type: **not specified**

Baremetal provisioning in the CERN cloud

Friday, May 18, 2018 9:40 AM (20 minutes)

Virtual machines is the technology that formed the modern clouds - private and public - however the physical machine are back in a more cloudy way. Cloud providers are offering APIs for baremetal server provisioning on demand and users are leveraging containers for isolation and reproducible deployments. In this talk, I will be presenting one of the newest services at the CERN cloud, Ironic, the Baremetal service of OpenStack. You will learn how the Cloud team improves its operational workflow and accounting and how users can use the same tooling they are used to when working with virtual machines. Finally, you will hear about the recent integration efforts between the container and baremetal services.

Desired length

20

Primary author: TRIGAZIS, Spyridon (CERN)**Presenter:** TRIGAZIS, Spyridon (CERN)**Session Classification:** Clouds, virtualisation, grids**Track Classification:** Grid, Cloud & Virtualisation

Contribution ID: 123

Type: **not specified**

Nikhef Site Report

Tuesday, May 15, 2018 10:50 AM (20 minutes)

Site report from Nikhef

Desired length

12

Primary author: VAN DER WAL, Bart (Nikhef)

Co-author: KUIPERS, Paul (Nikhef)

Presenter: VAN DER WAL, Bart (Nikhef)

Session Classification: Site reports

Track Classification: Site Reports

Contribution ID: 124

Type: **not specified**

Changing Compute Landscape at Brookhaven

Wednesday, May 16, 2018 2:00 PM (20 minutes)

Computing is changing at BNL, we will discuss how we are restructuring our Condor pools, integrating them with new tools like Jupyter notebooks, and other resources like HPC systems run with Slurm.

Desired length

15-20

Primary author: STRECKER-KELLOGG, William Edward (Brookhaven National Laboratory (US))

Presenter: STRECKER-KELLOGG, William Edward (Brookhaven National Laboratory (US))

Session Classification: Computing and batch systems

Track Classification: Computing & Batch Services

Contribution ID: 125

Type: **not specified**

A smorgasbord of tools around Linux at DESY

Thursday, May 17, 2018 3:00 PM (20 minutes)

In the past, we have developed lots of smaller and larger tools to help in various aspects of Linux administration at DESY.

We present (some) of them in this talk.

An incomplete list is:

- Two-Factor-Authentication
- Timeline repositories
- Making Kernel upgrade notifications (more) audit safe
- Fail2ban

Desired length

20

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

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

Session Classification: Basic IT services

Track Classification: Basic IT Services

Contribution ID: 126

Type: **not specified**

Techlab benchmarking web portal

Wednesday, May 16, 2018 3:00 PM (20 minutes)

Techlab, a CERN IT project, is a hardware lab providing experimental systems and benchmarking data for the HEP community.

Techlab is constantly on the lookout for new trends in HPC, cutting-edge technologies and alternative architectures, in terms of CPUs and accelerators.

We believe that in the long run, a diverse offer and a healthy competition in the HPC market will serve science in particular, computing in general, and everyone in the end.

For this reason, we encourage the use of not-quite-there-yet alternatives to the standard x86 quasi-monopoly, in the hope that in the near future, such alternative architectures can proudly compete, on an equal footing.

We buy hardware, set it up, test and benchmark it, then make it available to members of the HEP community for porting and testing their scientific applications and algorithms. On a best-effort basis, we try and help users make the best out of the hardware we provide.

To serve as basis for hardware choice, we run extensive benchmarks on all the systems we can get our hands on, and share the results to help others make fully informed choices when buying hardware that will fit their computing needs. As a means to achieve this, we developed a benchmarking web portal, open to everyone in the HEP community, to upload and publish data about all kinds of hardware. It was built with security in mind, and provides fine-grained access control to encourage even people working on yet-unreleased hardware to contribute.

As Techlab cannot possibly buy and test everything, it is our hope that this portal gets used by other HEP labs, and the database we build together becomes the 'one-stop shop' for benchmarking.

This presentation both gives an overview of Techlab's benchmarking web portal —what and whom it is designed for, what we hope to achieve with it —and delves into the technology choices of the implementation.

Desired length

20

Primary author: REIS, Maxime (CERN)**Presenter:** REIS, Maxime (CERN)**Session Classification:** Computing and batch systems**Track Classification:** Computing & Batch Services

Contribution ID: 127

Type: **not specified**

Integration of OpenStack and Amazon Web Service into local batch job system

Friday, May 18, 2018 11:10 AM (20 minutes)

Cloud computing enables flexible resource provisioning on demand. Through the collaboration with National Institute of Informatics (NII) Japan, we have been integrating our local batch job system with clouds for expanding its computing resource and providing heterogeneous clusters dynamically. In this talk, we will introduce our hybrid batch job system which can dispatch jobs to provisioned instances on on-premise OpenStack and Amazon Web Service as well as local servers. We will also report some performance test results conducted for investigation of the scalability.

Desired length

20

Primary author: TAKASE, Wataru (High Energy Accelerator Research Organization (JP))

Co-authors: NAKAMURA, Tomoaki (KEK); MURAKAMI, Koichi (KEK); SASAKI, Takashi (KEK)

Presenter: TAKASE, Wataru (High Energy Accelerator Research Organization (JP))

Session Classification: Clouds, virtualisation, grids

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: 128

Type: **not specified**

News from the DESY batch-clusters

Wednesday, May 16, 2018 2:20 PM (20 minutes)

The batch facilities at DESY are currently enlarged significantly while at the same time partly migrated from SGE to HTCondor.

This is a short overview of what is going on on site in terms of GRID-, local- and HPC cluster development.

Desired length

15

Primary authors: Dr HARTMANN, Thomas (DESY); FINNERN, Thomas (DESY); STERNBERGER, Sven (DESY); BEYER, Christoph

Presenter: BEYER, Christoph

Session Classification: Computing and batch systems

Track Classification: Computing & Batch Services

Contribution ID: 129

Type: **not specified**

Deployment of IPv6 on WLCG - an update from the HEPiX IPv6 working group

Monday, May 14, 2018 4:10 PM (20 minutes)

For several years the HEPiX IPv6 Working Group has been testing WLCG services to ensure their IPv6 compliance. The transition of WLCG central and storage services to dual-stack IPv4/IPv6 is progressing well, thus enabling the use of IPv6-only CPU resources as agreed by the WLCG Management Board and presented by us at previous HEPiX meetings.

By April 2018, all WLCG Tier 1 data centres have provided access to their services over IPv6. The LHC experiments have requested all WLCG Tier 2 centres to provide dual-stack access to their storage by the end of LHC Run 2. The working group, driven by the requirements of the LHC VOs to be able to use IPv6-only opportunistic resources, continues to encourage wider deployment of dual-stack services and has been monitoring the transition. We will present the progress of the transition to IPv6.

Desired length

20 minutes

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

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

Session Classification: Networking and security

Track Classification: Networking & Security

Contribution ID: 130

Type: **not specified**

First Impressions of Saltstack and Reclass as our new Configuration Management System

Thursday, May 17, 2018 2:40 PM (20 minutes)

In the Autumn of 2016 the Nikhef data processing facility (NDPF) found itself at a junction on the road of configuration management. The NDPF was one of the early adopters of Quattor, which served us well since the early days of the Grid. But where grid deployments were uniquely complex to require the likes of Quattor then, nowadays a plethora of configuration systems have cropped up to fulfill the needs of the booming industry of cloud orchestration.

Faced with the choice of an overhaul of our Quattor installation to bring our site up-to-date, or an investment to adopt a brand new system, we opted for the latter. And among the many candidates like Chef, Puppet, and Ansible, we chose Saltstack and partnered it with Reclass.

This led to hours of discussions designing the flows and processes of the new system, as well figuring out how to do the transition from the old to the new. And many more were spent trying out these ideas, pioneering as it were to find a way forward that would work for us.

We like to present the insights we've gained about Saltstack and configuration management in general, and our design choices in particular. This is very much a work in progress, as we are getting to know our new system while we are implementing more and more of the moving parts. We are gaining invaluable experience with tasks that are otherwise rarely among the day-to-day business of the system administrator, and we like to share it with our peers.

Desired length

20

Primary author: VAN DOK, Dennis**Co-author:** Dr PICKFORD, Andrew (Nikhef)**Presenter:** VAN DOK, Dennis**Session Classification:** Basic IT services**Track Classification:** Basic IT Services

Contribution ID: 132

Type: **not specified**

Logistics and announcements

Monday, May 14, 2018 9:15 AM (15 minutes)

Session Classification: Miscellaneous

Track Classification: Miscellaneous

Contribution ID: 133

Type: **not specified**

Welcome to University of Wisconsin-Madison

Monday, May 14, 2018 9:00 AM (15 minutes)

Session Classification: Miscellaneous

Track Classification: Miscellaneous

Contribution ID: 134

Type: **not specified**

News from the world of Federated Identity Management and AAI

Tuesday, May 15, 2018 11:50 AM (20 minutes)

There are many ongoing activities related to the development and deployment of Federated Identities and AAI (Authentication and Authorisation Infrastructures) in research communities and cyber Infrastructures including WLCG and others. This talk will give a high-level overview of the status of at least some of the current activities in FIM4R, AARC, WLCG and elsewhere.

Desired length

20 minutes

Primary authors: KELSEY, Dave (STFC - Rutherford Appleton Lab. (GB)); SHORT, Hannah (CERN)

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

Session Classification: Networking and security

Track Classification: Networking & Security

Contribution ID: 135

Type: **not specified**

What's new in HTCondor? What is upcoming?

Wednesday, May 16, 2018 3:20 PM (20 minutes)

The goal of the HTCondor team is to develop, implement, deploy, and evaluate mechanisms and policies that support High Throughput Computing (HTC) on large collections of distributively owned computing resources. Increasingly, the work performed by the HTCondor developers is being driven by its partnership with the High Energy Physics (HEP) community.

This talk will present recent changes and enhancements to HTCondor, including details on some of the enhancements created for the forthcoming HTCondor v8.8.0 release, as well as changes created on behalf of the HEP community. It will also discuss the upcoming HTCondor development roadmap, and seek to solicit feedback on the roadmap from HEPiX attendees.

Desired length

15 min, but I am flexible shorter or longer

Primary author: TANNENBAUM, Todd (University of Wisconsin Madison (US))

Presenter: TANNENBAUM, Todd (University of Wisconsin Madison (US))

Session Classification: Computing and batch systems

Track Classification: Computing & Batch Services

Contribution ID: 136

Type: **not specified**

BNL Site Report

Monday, May 14, 2018 9:50 AM (20 minutes)

Updates from BNL since KEK meeting

Desired length

15

Primary author: YU, David (BNL)

Presenter: YU, David (BNL)

Session Classification: Site reports

Track Classification: Site Reports

Contribution ID: 137

Type: **not specified**

FZU site report

Tuesday, May 15, 2018 11:30 AM (20 minutes)

Recently we deployed new cluster with worker nodes with 10 Gbps network connection and new disk servers for DPM and xrootd. I will also discuss migration from Torque/Maui to HTCondor batch system.

Desired length

12

Primary author: CHUDOBA, Jiri (Acad. of Sciences of the Czech Rep. (CZ))

Presenter: CHUDOBA, Jiri (Acad. of Sciences of the Czech Rep. (CZ))

Session Classification: Site reports

Track Classification: Site Reports

Contribution ID: 138

Type: **not specified**

BNL New Data Center - Status and Plans

Thursday, May 17, 2018 9:20 AM (20 minutes)

BNL is planning a new on-site data center for its growing portfolio of programs in need of scientific computing support. This presentation will provide an update on the status and plans for this new data center.

Desired length

20

Primary authors: WONG, Tony (Brookhaven National Laboratory); ZAYTSEV, Alexandr (Brookhaven National Laboratory (US)); Dr MISAWA, Shigeki (Brookhaven National Laboratory); LANCON, Eric Christian (BNL)

Presenter: WONG, Tony (Brookhaven National Laboratory)

Session Classification: IT facilities

Track Classification: IT Facilities & Business Continuity

Contribution ID: 139

Type: **not specified**

PDSF - Current Status and Migration to Cori

Wednesday, May 16, 2018 4:10 PM (20 minutes)

PDSF, the Parallel Distributed Systems Facility, has been in continuous operation since 1996 serving high energy physics research. It is currently a tier-1 site for Star, a tier-2 site for Alice and a tier-3 site for Atlas. We are in the process of migrating PDSF workload from commodity cluster to the Cori a Cray XC40 system. The process will involve preparing containers that will allow PDSF community to effectively move workloads between systems and other HPC centers. We are in the discovering process of optimizing serial jobs in a parallel environment. The goal is to minimize service interruptions as we shut down the existing PDSF by the second half of 2019.

Desired length

15

Primary author: QUAN, Tony (LBL)**Co-authors:** BALEWSKI, Jan (NERSC); Mr GEORG, Rath (NERSC); BOTTS, James (LBNL); Mr LOCKWOOD, Glenn (NERSC); Mr HAZEN, Damian (NERSC)**Presenter:** QUAN, Tony (LBL)**Session Classification:** Computing and batch systems**Track Classification:** Computing & Batch Services

Contribution ID: **140**

Type: **not specified**

RAL Site Report

Tuesday, May 15, 2018 11:10 AM (20 minutes)

Update on activities at RAL

Desired length

12

Primary author: BLY, Martin (STFC-RAL)

Presenter: BLY, Martin (STFC-RAL)

Session Classification: Site reports

Track Classification: Site Reports

Contribution ID: 141

Type: **not specified**

WLCG/OSG Networking Update

Monday, May 14, 2018 3:50 PM (20 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 OSG Networking Area is a partner of the WLCG effort and is focused on being the primary source of networking information for its partners and constituents. We will report on the changes and updates that have occurred since the last HEPiX meeting.

The WLCG Network Throughput 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 OSG/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.

We will provide a status update on the LHCOPN/LHCONE perfSONAR infrastructure as well as cover recent changes in the higher level services due to reorganisation of the OSG. This will include details on the central service migrations, updates to the dashboards; updates and changes to the Web-based mesh configuration system and details on the newly established pipeline for processing perfSONAR results.

In addition, we will provide an overview of the recent major network incidents that were investigated with the help of perfSONAR infrastructure and provide information on changes that will be included in the next perfSONAR Toolkit version 4.1. We will also cover the status of our WLCG/OSG deployment and provide some information on our future plans.

Desired length

20

Primary authors: MC KEE, Shawn (University of Michigan (US)); BABIK, Marian (CERN)

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

Session Classification: Networking and security

Track Classification: Networking & Security

Contribution ID: 142

Type: **not specified**

Network Functions Virtualisation Working Group Update

Tuesday, May 15, 2018 2:00 PM (20 minutes)

High Energy Physics (HEP) experiments have greatly benefited from a strong relationship with Research and Education (R&E) network providers and thanks to the projects such as LHCOPN/LHCONE and REN contributions, have enjoyed significant capacities and high performance networks for some time. RENs have been able to continually expand their capacities to over-provision the networks relative to the experiments needs and were thus able to cope with the recent rapid growth of the traffic between sites, both in terms of achievable peak transfer rates as well as in total amount of data transferred. For some HEP experiments this has led to designs that favour remote data access where network is considered an appliance with almost infinite capacity. There are reasons to believe that the network situation will change due to both technological and non-technological reasons starting already in the next few years. Various non-technological factors that are in play are for example anticipated growth of the non-HEP network usage with other large data volume sciences coming online; introduction of the cloud and commercial networking and their respective impact on usage policies and securities as well as technological limitations of the optical interfaces and switching equipment.

As the scale and complexity of the current HEP network grows rapidly, new technologies and platforms are being introduced, collectively called Network Functions Virtualisation (NFV), ranging from software-based switches such as OpenVSwitch, Software Defined Network (SDN) controllers such as OpenDaylight up to full platform based open solutions such as Cumulus Linux. With many of these technologies becoming available, it's important to understand how we can design, test and develop systems that could enter existing production workflows while at the same time changing something as fundamental as the network that all sites and experiments rely upon. In this talk we'll give an update on the Network Functions Virtualisation (NFV) WG that was established at the last HEPiX meeting. We'll provide details on its mandate, objectives, organisation of work as well as areas of interest that were already discussed and plans for the near-term future.

Desired length

15

Primary authors: MC KEE, Shawn (University of Michigan (US)); BABIK, Marian (CERN)**Presenter:** MC KEE, Shawn (University of Michigan (US))**Session Classification:** Networking and security**Track Classification:** Networking & Security

Contribution ID: 143

Type: **not specified**

Data analysis as a service

Friday, May 18, 2018 10:50 AM (20 minutes)

We are seeing an increasingly wide variety of uses being made of Hybrid Cloud (and Grid!) computing technologies at STFC, this talk will focus on the services being delivered to end users and novel integrations with existing local compute and data infrastructure.

Desired length

20

Primary author: Mr ADAMS, James (STFC RAL)**Presenter:** Mr ADAMS, James (STFC RAL)**Session Classification:** Clouds, virtualisation, grids**Track Classification:** Grid, Cloud & Virtualisation

Contribution ID: 144

Type: **not specified**

Recent status of KEK network

Tuesday, May 15, 2018 2:20 PM (20 minutes)

The Belle II detector is already taking data by cosmic ray test and is about to record data by beam. The importance of the network connectivity becomes higher than all other experiments in KEK. It is not only for the data transfer but also for researchers who are watching the condition of detectors from off sites.

We will report the present status of the campus network and the upgrade plan in this summer.

Desired length

15minutes

Primary authors: SUZUKI, Soh; MURAKAMI, Tadashi (KEK); NAKAMURA, Tomoaki (KEK); YUASA, Fukuko (KEK); Mr NAKAMURA, Teiji (KEK); Mr HASHIMOTO, Kiyoharu (KEK); MANABE, Atsushi (High Energy Accelerator Research Organization (KEK))

Presenter: SUZUKI, Soh

Session Classification: Networking and security

Track Classification: Networking & Security

Contribution ID: 145

Type: **not specified**

AFS and Linux Containers

Wednesday, May 16, 2018 11:30 AM (20 minutes)

One future model of software deployment and configuration is containerization.

AFS has been used for software distribution for many decades. Its global file namespace, the @sys path component substitution macro which permits file paths to be platform-agnostic, and the atomic publication model ("vos release") have proven to be critical components of successful software distribution systems that scale to hundreds of thousands of systems and have survived multiple OS and processor architecture changes.

The AuriStorFS security model consisting of combined-identity authentication, multi-factor authorization, and mandatory security policies permits a global name space to be shared between internal, dmz and cloud; and to store a mix of open and restricted data.

The combination of Linux Containers, the global AFS namespace, and the AuriStorFS security model is powerful permitting the development of container based software deployments that can safely bridge internal, dmz and cloud with reduced risk of data leaks.

This session will discuss the most recent updates to AuriStorFS and the Linux kernel implementation of AF_RXRPC socket family and (k)AFS filesystem. A demonstration will be included consisting of:

- Containers with binary executable files stored in /afs
- Containers mounting private AFS volume for scratch space
- AuriStorFS and (k)AFS file system implementations running side-by-side
- Linux namespaces for /afs

AuriStorFS milestones since HEPiX Spring 2017 include:

1. Successful migration and replication of volumes exceeding 5.5TB. The largest production volume so far is 50TB with a 250TB volume
2. Deployment of a single AuriStorFS cell spanning an internal data center, AWS and GCP with more than 25,000 nodes for distribution of software and configuration data.
3. Meltdown and Spectre remediation. In response to nearly 30% performance hit from Meltdown and Spectre the AuriStor team optimized the Rx stack, Ubik database and fileserver to reduce the number of syscalls by more than 50%
4. AES-NI, SSSE3, AVX and AVX2 Intel processor optimization of AES256-CTS-HMAC-SHA1-96 cryptographic operations for kernel cache managers reduces computation time by 64%

AF_RXRPC and kAFS highlights:

- IPv6 support for AuriStorFS
- dynamic root mount -o dyn
- @sys and @cell support
- multipage read and write support
- local hero directory caching
- per file acls (for AuriStorFS)

- server failover and busy volume retries

Desired length

20 minutes

Primary authors: Mr ALTMAN, Jeffrey (AuriStor, Inc); Mr DIONNE, Marc (AuriStor, Inc); Mr HOWELLS, David (Red Hat)

Presenter: Mr ALTMAN, Jeffrey (AuriStor, Inc)

Session Classification: Storage and file systems

Track Classification: Storage & Filesystems

Contribution ID: 146

Type: **not specified**

RAL Cloud update

Friday, May 18, 2018 10:00 AM (20 minutes)

As our OpenStack cloud enters full production, we give an overview of the design and how it leverages the RAL Tier 1 infrastructure & support. We also present some of the new use cases and science being enabled by the cloud platform.

Desired length

20 minutes

Primary authors: COLLIER, Ian (Science and Technology Facilities Council STFC (GB)); DIBBO, Alexander (STFC RAL)

Presenter: COLLIER, Ian (Science and Technology Facilities Council STFC (GB))

Session Classification: Clouds, virtualisation, grids

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: 147

Type: **not specified**

CC-IN2P3 User Portal

Monday, May 14, 2018 2:40 PM (20 minutes)

CC-IN2P3 is one of the largest academic data centers in France. Its main mission is to provide the particle, astroparticle and nuclear physics community with IT services, including largescale compute and storage capacities. We are a partner for dozens of scientific experiments and hundreds of researchers that make a daily use of these resources. The CC-User Portal project's goal is to develop a web portal providing the users with a single-entry point to monitor their activity and incidents, to receive vital information and have the necessary links in order to access and use our services efficiently.

During HEPiX Fall 2017, we presented our first developments and exchanged with the community our thoughts on how to display the information to the users. With this presentation we would like to show what is now deployed in production and which features we are already developing to complete the web portal offer and meet the users' needs.

Desired length

20 minutes

Primary authors: MARCHETTI, Gino (CNRS); L'ORPHELIN, Cyril (CNRS/IN2P3); VERNET, Renaud (CC-IN2P3 - Centre de Calcul (FR))

Presenter: MARCHETTI, Gino (CNRS)

Session Classification: End-user services and operating systems

Track Classification: End-User IT Services & Operating Systems

Contribution ID: 148

Type: **not specified**

AGLT2 Site Update

Monday, May 14, 2018 10:10 AM (20 minutes)

We will present an update on our site since the Fall 2017 report, covering our changes in software, tools and operations.

Some of the details to cover include the enabling of IPv6 for all of our AGLT2 nodes, our migration to SL7, exploration of the use of Bro/MISP at the UM site, the use of Open vSwitch on our dCache storage and information about our newest hardware purchases and deployed middleware.

We conclude with a summary of what has worked and what problems we encountered and indicate directions for future work.

Desired length

12

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

Co-authors: Dr BALL, Robert (University of Michigan (US)); LAURENS, Philippe (Michigan State University (US))

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

Session Classification: Site reports

Track Classification: Site Reports

Contribution ID: 149

Type: **not specified**

Teraflops of Jupyter: A Notebook Based Analysis Portal at BNL

Thursday, May 17, 2018 11:45 AM (20 minutes)

The BNL Scientific Data and Computing Center (SDCC) has begun to deploy a user analysis portal based on Jupyterhub. The Jupyter interfaces have back-end access to the Atlas compute farm via Condor for data analysis, and to the GP-GPU resources on the Institutional Cluster via Slurm, for machine learning applications. We will present the developing architecture of this system, current use cases and results, and discuss future plans.

Desired length

20 minutes

Primary authors: RIND, Ofer; BENJAMIN, Doug (Duke University (US)); THROWE, Thomas (BNL); STRECKER-KELLOGG, William (Brookhaven National Lab)

Presenter: RIND, Ofer

Session Classification: Basic IT services

Track Classification: End-User IT Services & Operating Systems

Contribution ID: 150

Type: **not specified**

Network status at IHEP and LHCONE progress in China

Tuesday, May 15, 2018 3:45 PM (20 minutes)

Present the Network status at IHEP and LHCONE progress in China

Desired length

15

Primary author: ZENG, Shan (Chinese Academy of Sciences (CN))

Co-author: QI, Fazhi

Presenter: ZENG, Shan (Chinese Academy of Sciences (CN))

Session Classification: Networking and security

Track Classification: Networking & Security

Contribution ID: **151**

Type: **not specified**

Workshop wrap-up

Friday, May 18, 2018 12:10 PM (30 minutes)

Session Classification: Miscellaneous

Contribution ID: 152

Type: **not specified**

AFS Update: Spring 2018

Wednesday, May 16, 2018 10:00 AM (20 minutes)

Last May it was announced “AFS” was awarded the 2016 ACM System Software Award. This presentation will discuss the current state of the AFS file system family including:

- IBM AFS 3.6
- OpenAFS
- kAFS
- AuriStor File System

IBM AFS 3.6 is a commercial product no longer publicly available.

OpenAFS is fork from IBM AFS 3.6 available under the IBM Public License 1.0. The currently supported release branches are 1.6 and 1.8.

AuriStorFS is a commercial file system that is backward compatible with both IBM AFS 3.6 and OpenAFS clients and servers. AFS cells hosted on AuriStorFS servers benefit from substantial improvements in

kAFS is an AFS and AuriStorFS client distributed as part of the mainline Linux kernel distribution. kAFS shares no source code with IBM AFS or OpenAFS.

Desired length

20 minutes

Primary author: Mr ALTMAN, Jeffrey (AuriStor, Inc.)

Presenter: Mr ALTMAN, Jeffrey (AuriStor, Inc.)

Session Classification: Storage and file systems

Track Classification: Storage & Filesystems

Contribution ID: **153**Type: **not specified**

IHEP Site Report

Monday, May 14, 2018 11:40 AM (20 minutes)

The computing center of IHEP maintains a HTC cluster with 10,000 cpu cores and a site including about 15,000 CPU cores and more than 10PB storage. The presentation will talk about the its progress and next plan of IHEP Site.

Desired length

12

Primary author: SHI, Jingyan (IHEP)

Presenter: SHI, Jingyan (IHEP)

Session Classification: Site reports

Track Classification: Site Reports

Contribution ID: 154

Type: **not specified**

INFN-T1 flooding report

Thursday, May 17, 2018 10:00 AM (20 minutes)

On November 9 2017, a major flooding occurred in the computing rooms: this has turned into a down of all the services for a prolonged period of time.

In this talk we will go through all the issues we faced in order to recover the services in the quickest and most efficient way; we will analyze in detail the incident and all the steps made to recover the computing rooms, electrical power, network, storage and farming.

Moreover, we will discuss the hidden dependencies among services discovered during the recovery of the systems and will detail how we solved them.

Desired length

15

Primary authors: DAL PRA, Stefano (INFN); CHIERICI, Andrea (INFN-CNAF); DELL'AGNELLO, Luca (INFN)

Presenter: DAL PRA, Stefano (INFN)

Session Classification: IT facilities

Track Classification: IT Facilities & Business Continuity

Contribution ID: 155

Type: **not specified**

University of Nebraska CMS Tier2 Site Report

Monday, May 14, 2018 11:00 AM (20 minutes)

Updates from T2_US_Nebraska covering our experiences operating CentOS 7 + Docker/Singularity, random dabbling with SDN to better HEP transfers, involvement with the Open Science Grid, and trying to live the IPv6 dream.

Desired length

12

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

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

Session Classification: Site reports

Track Classification: Site Reports

Contribution ID: 156

Type: **not specified**

Planning new datacenter network architecture

Thursday, May 17, 2018 9:40 AM (20 minutes)

In scope of the Wigner Datacenter cloud project we are consolidating our network equipment. According to our plans we would like to purchase 100 Gbps datacenter switches in order to anticipate our current and future needs. We need automated, vendor neutral and easily operable network. This presentation highlights our requirements and design goals, candidates we have tested in our lab. We take the opportunity here to introduce our knowledge lab initiative where we can expand the scope of testing solutions.

Desired length

15

Primary author: Ms RACZ, Szilvia (Wigner Datacenter)**Presenter:** Ms RACZ, Szilvia (Wigner Datacenter)**Session Classification:** IT facilities**Track Classification:** IT Facilities & Business Continuity

Contribution ID: 157

Type: **not specified**

The OpenAFS Foundation

Tuesday, May 15, 2018 4:45 PM (20 minutes)

We would like to have one of the Board members of The OpenAFS Foundation, Inc, speak about this 501(c)(3), US-based, non-profit organization dedicated to fostering the stability and growth of OpenAFS, an open source implementation of the AFS distributed network filesystem. The OpenAFS Foundation adopted a three-fold mission: to attract and increase the community of OpenAFS users, to foster the OpenAFS community of experts, and to nurture and evolve the OpenAFS technology; each will be explained briefly.

We would like to ask for help from the scientific community and ask its researchers to:

- Contribute to the OpenAFS code, as such contributions are critical for the survival and improvement of the OpenAFS technology. The Foundation has obtained insurance to protect all contributors from potential liability and infringement of IP lawsuits.
- Reviewing code, as their feedback is not only valuable and appreciated in several ways.
- Write documentation for already existing code, which is desperately needed. Communicate the changes to their computing needs, and how they would like OpenAFS to be even more useful to them in the future.
- Help craft code specifications and/or code design.
- Become a guardian and as such, an active, long-term champion shaping the future viability and well-being of OpenAFS technology.
- Donate and/or identify organizations possibly willing and able to contribute funds to sustain a lean operation and/or to fund specific development efforts to be assigned in an open bid process.

The presenter will be either Todd DeSantis or Margarete Ziemer, both Directors and Board members of the Foundation.

Desired length

20 minutes

Primary authors: Mr DESANTIS, Todd (IBM); Dr ZIEMER, Margarete (Sine Nomine Associates)

Presenter: Dr ZIEMER, Margarete (Sine Nomine Associates)

Session Classification: Storage and file systems

Track Classification: Storage & Filesystems

Contribution ID: 158

Type: **not specified**

HPL and HPCG Benchmark on BNL linux farm and SDCC

Wednesday, May 16, 2018 4:50 PM (20 minutes)

HPL and HPCG Benchmark on Brookhaven National Laboratory SDCC clusters and various generations of Linux Farm nodes has been conducted and compared with HS06 results. While HPL results are more aligned with CPU/GPU performance. HPCG results are impacted by memory performances as well.

Desired length

15

Primary author: Dr DONG, Zhihua

Presenter: Dr DONG, Zhihua

Session Classification: Computing and batch systems

Track Classification: Computing & Batch Services

Contribution ID: 159

Type: **not specified**

PDSF Site Report

Monday, May 14, 2018 11:20 AM (20 minutes)

PDSF, the Parallel Distributed Systems Facility, was moved to Lawrence Berkeley National Lab from Oakland CA in 2016. The cluster has been in continuous operation since 1996 serving high energy physics research. The cluster is a tier-1 site for Star, a tier-2 site for Alice and a tier-3 site for Atlas.

This site report will describe lessons learned and challenges met, when migrating from Univa GridEngine to the Slurm scheduler, experiences running containerized software stacks using Shifter, as well as upcoming changes to systems management and the future of PDSF.

Desired length

12

Primary author: RATH, Georg (Lawrence Berkeley National Laboratory)

Co-authors: BOTTS, James (LBNL); PORTER, Jeff (Lawrence Berkeley National Lab. (US)); BALEWSKI, Jan (NERSC); JACOBSEN, Douglas (NERSC); GERHARDT, Lisa (LBNL); DECLERCK, Tina (NERSC); QUAN, Tony (LBL); Mr BASHEER, Ershaad (NERSC)

Presenter: RATH, Georg (Lawrence Berkeley National Laboratory)

Session Classification: Site reports

Track Classification: Site Reports

Contribution ID: **160**Type: **not specified**

OpenAFS Release Team report

Tuesday, May 15, 2018 4:25 PM (20 minutes)

A report from the OpenAFS Release Team on recent OpenAFS releases, including the OpenAFS 1.8.0 release, the first major release in several years. Topics include acknowledgement of contributors, descriptions of issues recently resolved, and a discussion of commits under review for post 1.8.0.

Desired length

20 minutes

Primary authors: Mr MEFFIE, Michael (Sine Nomine Associates); Dr KADUK, Benjamin (MIT (formerly)); Mr WIESAND, Stephan (DESY)

Presenter: Mr MEFFIE, Michael (Sine Nomine Associates)

Session Classification: Storage and file systems

Track Classification: Storage & Filesystems

Contribution ID: 161

Type: **not specified**

IPv6 Deployment Experience at the GridKa Tier-1 at KIT

Monday, May 14, 2018 4:30 PM (20 minutes)

Recently, we've deployed IPv6 for the CMS dCache instance at KIT. We've run into a number of interesting problems with the IPv6 setup we had originally chosen. The presentation will detail the lessons we've learned and the resulting redesign of our IPv6 deployment strategy.

Desired length

10

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

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

Session Classification: Networking and security

Track Classification: Networking & Security