

HEPiX Autumn/Fall 2018 Workshop



Report of Contributions

Contribution ID: 1

Type: **not specified**

PIC site report

Monday 8 October 2018 11:05 (15 minutes)

News from PIC since the HEPiX Spring 2018 workshop at Madison, Wisconsin, USA.

Desired length

12

Author: FLIX MOLINA, Jose (Centro de Investigaciones Energéticas Medioambientales y Tecnológicas)

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

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 2

Type: **not specified**

INFN-T1 Site report

Monday 8 October 2018 11:20 (15 minutes)

Updates on INFN-Tier1 centre

Desired length

12

Author: CHIERICI, Andrea (INFN-CNAF)

Presenter: CHIERICI, Andrea (INFN-CNAF)

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 3

Type: **not specified**

Challenges, solutions and lessons learnt in 7 years of Service Management at CERN

Tuesday 9 October 2018 14:00 (25 minutes)

CERN has been using ITIL Service Management methodologies and ServiceNow since early 2011. Initially a joint project between just the Information Technology and the General Services Departments, now most of CERN is using this common methodology and tool, and all departments are represented totally or partially in the CERN Service Catalogue.

We will present a summary of the current situation of Service Management at CERN, as well as its recent evolution. We will talk about service onboarding, user experience, and tool configuration. For this, some topics will be explained in detail, such as:

- The CERN Service Catalogue, the CERN Service Management processes, bringing new services on board, quality management of external contracts, service data privacy management, and some examples of positive outcomes.
- Incident versus request distinction, how to gather feedback about tickets from users, and the Knowledge Base evolution.
- Good practices of tool configuration and integration; data centre and IT services monitoring and status.

The various challenges, adopted solutions and lessons learnt in these topics will be presented. Finally, ongoing and future work will be discussed as well.

Desired length

20 minutes

Author: MARTIN CLAVO, David (CERN)**Co-authors:** BRUGGER, Barbara (CERN); DELAMARE, Catherine (CERN); LIENARD, Eric (CERN); FERNANDEZ GONZALEZ, Isabel (CERN); MOLLER, Mats (CERN); TOTEVA, Zhechka (CERN); SALTER, Wayne (CERN); MARTENS, Reinoud (CERN); VAN DER VOSSEN, Olaf (CERN); CREMEL, Nicole (CERN)**Presenter:** MARTIN CLAVO, David (CERN)**Session Classification:** End-User IT Services & Operating Systems**Track Classification:** End-User IT Services & Operating Systems

Contribution ID: 4

Type: **not specified**

CERN Site Report

Monday 8 October 2018 11:35 (15 minutes)

News from CERN since the HEPiX Spring 2018 workshop at University of Wisconsin-Madison, Madison, USA.

Desired length

20

Author: DUMITRU, Andrei (CERN)

Presenter: DUMITRU, Andrei (CERN)

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 5

Type: **not specified**

CERN Linux services status update

Tuesday 9 October 2018 14:25 (25 minutes)

An update on CERN Linux support distributions and services.

An update on the CentOS community and CERN involvement will be given.

We will discuss software collections, virtualization and openstack SIGs update.

We will present our anaconda plugin and evolution of the locmap tool.

A brief status on alternative arches (aarch64, ppc64le, etc...) work done by the community will be given.

Desired length

20

Authors: OULEVEY, Thomas (CERN); POLOK, Jarek (CERN)

Presenter: OULEVEY, Thomas (CERN)

Session Classification: End-User IT Services & Operating Systems

Track Classification: End-User IT Services & Operating Systems

Contribution ID: 6

Type: **not specified**

Future-Looking Data Storage for Peak Performance in HPC Environments

Thursday 11 October 2018 16:35 (25 minutes)

High performance computing (HPC) environments continually test the limits of technology and require peak performance from their equipment—including storage. Slow overall writing of data and long seek times between file reads due to non-consecutive files, library partitioning, or laborious loading mechanisms often plague tape library efficiencies for large tape users managing massive sets of data. As the growth of data in HPC is constantly expanding, optimizing performance, increasing uptime and lowering costs become vital components to any future-looking HPC storage architecture.

Pertinent storage innovations prioritize improvements in overall access speed and minimize wear to tape media and drives for optimal system reliability at minimal cost. Recent enhancements enable LTO tape-drive-based systems to thrive in large scale storage environments and demonstrate a continual commitment to the development of tape archive technology to serve the HPC space. From minimizing robotic contention without partitions, to increasing mount performance, to speeding the time to repair and lowering overall support costs, new tape features allow users to scale storage performance like never before. To maximize performance of hardware for long-term storage, front-end software must be met with an equally powerful back-end solution.

Join us to learn about the cutting-edge developments in tape library software that help supercomputing environments push the boundaries of their operational objectives, providing cost-effective storage that meets their performance, growth and environmental needs. Hear about the most recent developments to enable cost-effective solutions with unbeatable density.

Desired length

20 minutes

Author: Mr STARR, Matt (CTO at Spectra Logic)**Presenter:** Mr STARR, Matt (CTO at Spectra Logic)**Session Classification:** Storage & Filesystems**Track Classification:** Storage & Filesystems

Contribution ID: 7

Type: **not specified**

LHC Long Shutdown 2 and database changes

Thursday 11 October 2018 11:35 (25 minutes)

In 2019 and 2020 there will be the long shutdown of the LHC.

Besides technical interventions on the accelerator itself, we plan to review our practices and upgrade our Oracle databases.

In this session I will show what goals we have on the Oracle Database side and how we intend to achieve them.

Desired length

20

Author: Mr CALDARA, Ludovico (CERN)

Presenter: Mr CALDARA, Ludovico (CERN)

Session Classification: Storage & Filesystems

Track Classification: Storage & Filesystems

Contribution ID: 8

Type: **not specified**

Indico - Event management in HEP

Tuesday 9 October 2018 14:50 (25 minutes)

In this presentation, we will report on Indico's usage within the High Energy Physics Community, with a particular focus in the adoption of Indico 2.x.

We will also go over the most recent developments in the project, including the new Room Booking interface in version 2.2, as well as plans for the future.

Desired length

20

Author: FERREIRA, Pedro (CERN)**Presenter:** FERREIRA, Pedro (CERN)**Session Classification:** End-User IT Services & Operating Systems**Track Classification:** End-User IT Services & Operating Systems

Contribution ID: 9

Type: **not specified**

Extending computing facilities from local to the cloud: hybrid solutions for computing clusters

Wednesday 10 October 2018 09:15 (20 minutes)

Private and public cloud infrastructures had become a reality in the recent years. Science is looking at this solutions to extend the amount of computing and storage facilities for the research projects that are becoming bigger and bigger. This is the aim of Helix Nebula Science Cloud (HNSci-Cloud) a project lead by CERN in which we submitted two use cases for the astrophysics projects MAGIC and CTA. Both use cases had the purpose to create a hybrid environment to provide a powerful tool based on the Data Analysis as a service (DAaaS) paradigm.

In order to fulfill different requirements, we considered three scenarios: the first one, in the context of MAGIC, we implemented an analysis orchestrator by using the native cloud APIs, in order to send jobs to both cloud providers. The second one was by configuring DIRAC-VM in order to extend DIRAC computing environment for CTA to the cloud. And the last one was to extend the computing facilities at PIC by adding HTCondor cloud nodes transparently to the local farm, using PIC native tools like network, puppet and htcondor tuning.

Desired length

15

Author: Mr CASALS HERNANDEZ, Jordi (Port d'Informació Científica)

Co-authors: ACIN PORTELLA, Vanessa (Institut de Física d'Altes Energies); DELGADO MENGUAL, Jordi (Port d'Informació Científica); ACOSTA SILVA, Carlos (The Barcelona Institute of Science and Technology (BIST) (ES)); Mr LOPEZ MUNOZ, Fernando (PIC)

Presenter: Mr CASALS HERNANDEZ, Jordi (Port d'Informació Científica)

Session Classification: Grid, Cloud & Virtualisation

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: **10**Type: **not specified**

Computing & Software for Big Science Journal

Monday 8 October 2018 09:30 (20 minutes)

The journal Computing and Software for Big Science has been starting 18 months ago as a peer-reviewed journal, with a strong participation of the HEP community. This presentation will give the journal status after this initial period.

Desired length

15

Author: JOUVIN, Michel (Centre National de la Recherche Scientifique (FR))

Presenter: JOUVIN, Michel (Centre National de la Recherche Scientifique (FR))

Session Classification: Miscellaneous

Track Classification: Miscellaneous

Contribution ID: 11

Type: **not specified**

MAlt Project

Tuesday 9 October 2018 15:15 (25 minutes)

Over the years, CERN activities and services have been increasingly relying on commercial software and solutions to deliver core services, often leveraged by interesting financial conditions based on recognizing CERN statuses like academic, non-profit, research, etc. Once installed, well spread, heavily used, the leverage used to attract CERN service managers to the commercial solutions tends to disappear and is replaced by real business models and licensing schemes, that are considered to be unaffordable on the long term.

As a matter of fact, the CERN nature of openness targets delivering the same service to every type of CERN user, from Staff people to Users of the scientific resources. As a result, a high number of licenses is required to deliver services to everyone, and when traditional business models on a per user basis are applied, the costs per product can be huge.

More importantly, the trendy Cloud approach enterprises are now focusing on are introducing a new risk of lock-in on the data this time. It is always easy to import data in a cloud system, but always difficult to get data out. In addition, software vendors deliver easy and integrated cloud access for all their offered services, facilitating the life of end users when choosing storage.

The MAlt project objective is to target if possible Open Source products, to deliver services inclusive to all CERN community. The project's principles of engagement are to deliver the same service to every type of CERN user, avoid vendor lock-in to decrease risk and dependency, keep hands on the data and address bulk use cases.

Desired length

20

Author: ORMANCEY, Emmanuel (CERN)**Presenter:** ORMANCEY, Emmanuel (CERN)**Session Classification:** End-User IT Services & Operating Systems**Track Classification:** End-User IT Services & Operating Systems

Contribution ID: 12

Type: **not specified**

Why would a python developer learn rust, when there are no jobs in it?

Friday 12 October 2018 09:25 (25 minutes)

The language Rust has many potential benefits for the physics community. This talk explains why I think this is true in general and for HEP.

This will be a general introduction to the rust programming language, as a replacement for C/C++ and way to extend python.

Desired length

20 mins

Author: SYNGE, Owen

Presenter: SYNGE, Owen

Session Classification: End-User IT Services & Operating Systems

Track Classification: End-User IT Services & Operating Systems

Contribution ID: 13

Type: **not specified**

Improving OpenMP scaling using openssl

Wednesday 10 October 2018 16:10 (20 minutes)

Scaling an OpenMP or MPI application on modern TurboBoost-enabled CPUs is getting harder and harder. Using some simple ‘openssl’ commands, however, it is possible to adjust OpenMP benchmarking results to correct for the TurboBoost frequencies of modern Intel and AMD CPUs. In this talk I will explain how to achieve better OpenMP scaling numbers and will show how a non-root user can determine other performance characteristics of a CPU.

Desired length

15 min

Author: Mr KEIJSER, Jan Just (Nikhef)**Presenter:** Mr KEIJSER, Jan Just (Nikhef)**Session Classification:** Computing & Batch Services**Track Classification:** Computing & Batch Services

Contribution ID: 14

Type: **not specified**

Update about Wi-Fi service enhancement at CERN

Monday 8 October 2018 14:25 (25 minutes)

During recent times, we can observe a situation where Wi-Fi service starts to be primary Internet connection method in a campus networking. Wireless access is no longer an additional service or just interesting technology for conference and meeting rooms - now support for mobility is expected. In 2016, CERN launched global Wi-Fi renewal project across its campus. The subject of the project is to cover more than 200 office buildings with modern, highly reliable wireless solution. Project itself was presented during previous HEPiX workshop. This year, being in the final phase of deployment, the presentation will focus on implemented features, faced issues and an interaction with the manufacturers. Moreover, project extension for outdoor coverage of key areas across the CERN campus will be discussed.

Desired length

20 minutes

Author: Mr SOSNOWSKI, Adam**Presenter:** Mr SOSNOWSKI, Adam**Session Classification:** Networking & Security**Track Classification:** Networking & Security

Contribution ID: 15

Type: **not specified**

Challenges for connecting Internet of Thing devices at CERN

Monday 8 October 2018 16:55 (25 minutes)

During last years we observed a significant increase of the interest in Internet of Things (IoT) devices. Such equipment of different kind is more and more often considered as valuable and useful tool in the industry. Therefore, IoT comprises a wide variety of devices, user cases and can be connected to the network, using many different access methods.

As CERN network team, we would like to meet users' expectations and be prepared for coming request. The LPWAN network to provide connectivity to non-critical sensors with very low throughput requirements, is currently being deployed across the campus. Nevertheless, there is a potential need for the network, which will be able to interconnect other types of the IoT equipment. Nowadays, we receive some questions and requests considering devices, which work with different access methods, demand more bandwidth and communicate with the network more often than LPWAN equipment. To meet mentioned requirements, CERN has to study different possibilities to offer wireless infrastructure, which will be able to provide connectivity for a new IoT hardware.

Desired length

20 minutes

Author: Mr SOSNOWSKI, Adam**Presenter:** Mr SOSNOWSKI, Adam**Session Classification:** Networking & Security**Track Classification:** Networking & Security

Contribution ID: 16

Type: **not specified**

CERN procurement update

Thursday 11 October 2018 09:50 (25 minutes)

This is an overview of status and plans for the procurements of compute/storage for CERN data centres and some recent adaptations to better benefit from technology trends. The talk will also cover our workflow for hardware repairs as well as status and plans of our ongoing efforts in establishing inventories for deployed assets and spare parts. It will also cover some recent hardware issues and how we are dealing with them.

Desired length

Author: BARRING, Olof (CERN)**Presenter:** BARRING, Olof (CERN)**Session Classification:** IT Facilities & Business Continuity**Track Classification:** IT Facilities & Business Continuity

Contribution ID: 17

Type: **not specified**

Computer Security Update

Wednesday 10 October 2018 11:20 (25 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

Author: VALSAN, Liviu (CERN)**Presenter:** VALSAN, Liviu (CERN)**Session Classification:** Networking & Security**Track Classification:** Networking & Security

Contribution ID: **18**Type: **not specified**

CERN campus network upgrade

Monday 8 October 2018 16:30 (25 minutes)

After several years operating the same manufacturer in the CERN campus network, it's now the time to renew the equipment. CERN Communication Systems group is preparing the introduction of a new manufacturer with some changes based on requirements received from the users at CERN

Desired length

Author: Mr KISHIMOTO BISBE, Carles (CERN)

Presenter: Mr KISHIMOTO BISBE, Carles (CERN)

Session Classification: Networking & Security

Track Classification: Networking & Security

Contribution ID: 19

Type: **not specified**

FUJIFILM: Development of tape technology and challenges to overcome

Thursday 11 October 2018 14:50 (25 minutes)

Fujifilm is the world's leading manufacturer of magnetic tapes (LTO, 3592 and T10000). More than 80% of the storage capacity delivered on tapes comes from Fujifilm's manufacturing and assembly plants in Odawara (Japan) and Bedford (USA). Fujifilm is working in partnership with IBM on the development of the

next tape generations: a roadmap is established, describing the tape formats that will be used until the 2030s. The Scientific environment, HPC users, the genomic research, the satellite imagery or the biotechnology are among the industry segments that provide the most information to IBM and Fujifilm regarding the development criteria for future tape formats.

From a purely industrial point of view, the news in tape technology is to see how manufacturers can develop higher capacity and higher performance levels than previous generations, without hitting the recording stability, nor the data integrity. For example, we can mention the stakes related to storage capacity: in 2011, higher-capacity users could use 3592JC tapes (4TB of data per cartridge) and small and medium-sized businesses used to purchase LTO5 tapes (1.5 TB capacity). By the end of 2018, IBM will launch the 3592JE tapes, which will offer capacities of 20TB per tape cartridge. However, to make tapes of higher data capacity, it is necessary to be able to use smaller particles on tape. The problem is that the magnetic output of a particle depends, theoretically, on its size. A particle that is too small, with a magnetic resonance that is too low, would prevent the read/write process, or, simply, generate data loss. We call SNR (Signal-to-Noise Ratio), the measurement of the perception by the head of the drive, of the signals emitted by the particles that are coated on the tapes. Only a high SNR level can allow a stable recording. Preserving or even improving the level of SNR in tape technology is one of the key requirements for the development of higher capacity tapes.

Another fundamental risk lies in the fact that an increase in the capacity of a tape involves the use of smaller writing tracks. Reducing track width reduces the overall magnetic field and, consequently, hurts the SNR. Finally, an increased write speed, corollary to a greater storage capacity, is logically the third reduction factor of the SNR, and therefore of data loss. The development of new tape generations goes hand in hand with the development of new technologies that are designed to improve the quality of perception of the signals emitted between the tape and the drive's head. On top of Fujifilm's progress in tape manufacturing we can also evoke the technological innovations that result from IBM's R&D:

☒ The use of a larger number of write channels, which drastically increases both write and read speeds. 16 channel heads on LTO6 drives allow 160MB/s transfer rate, while 32 channel heads can increase the write speed up to 500MB/s.

☒ A Larger number of channel heads combined with thinner tracks pushed IBM to produce new heads that would be thinner and more performant at the same time. However, the new performances required in tape technology rendered the technologies used until now obsolete. Hybrid heads or "dual heads" are those that are traditionally used in tape technology. These heads combine two different functions: the writing process, which works by magnetic field injection to the tape (as a toothpaste) and the read process that captures magnetic signals via sensors. By launching its new generation of Terzetto heads, IBM has created the principle of "specialized head". From now on each head will have a unique function, which considerably increases its properties. The concrete result of the launch of these heads is an unprecedented improvement in the field of data integrity, and the SNR.

☒ The TMR head: the improvement of the system's performance increases the transfer of the current within the head, which generates a phenomenon of warming. Major warming disturbs the magnetic field: the new TMR heads, thanks to a greater resistance can multiply the capacity of the

head to capture signals by 4.

These three new features are part of the improvements we will discuss with users during the Hepix exhibition. We will also discuss other major issues such as the development deadlines for tape technologies offering beyond 50TB, how to increase the write speed of new storage systems, or how best to stabilize the transition of the tape within the drive, despite higher unwinding speeds than in the past.

Finally, we will introduce the new Strontium Ferrite technology, a tape coating method using particles, which, with the same size as Barium Ferrite particles, offer a superior magnetic output. This technology will enable the manufacture of tapes of more than 100TB capacity. It is likely that this technology will be the one most commonly used in the 2020s.

Desired length

20 minutes (if possible 12 extra minutes)

Author: Mrs GAMEIRO, Elisabeth (FUJIFILM Recording Media France)

Presenter: Mrs GAMEIRO, Elisabeth (FUJIFILM Recording Media France)

Session Classification: Storage & Filesystems

Track Classification: Storage & Filesystems

Contribution ID: 20

Type: **not specified**

Network configuration management tools evolution

Tuesday 9 October 2018 12:15 (25 minutes)

This talk will provide an overview of the recent changes in architecture and development procedures in use to manage all CERN networks (campus, experiments and technical networks).

These include:

- the migration from 20 years old PL/SQL code to Java using modern microservices architecture,
- the migration from multiple Git repositories to a single one in order to simplify organization and facilitate cross-project changes,
- the creation of a shared library containing the core business logic for managing CERN networks configuration used by multiple applications.

Desired length

15

Author: PAPOUTSIS, Nikolaos (CERN)

Presenters: PAPOUTSIS, Nikolaos (CERN); GARCIA MARTINEZ, Raul (CERN)

Session Classification: Basic IT Services

Track Classification: Basic IT Services

Contribution ID: 21

Type: **not specified**

Cost and system performance modelling in WLCG and HSF: an update

Thursday 11 October 2018 09:25 (25 minutes)

The HSF/WLCG cost and performance modeling working group was established in November 2017 and has since then achieved considerable progress in our understanding of the performance factors of the LHC applications, the estimation of the computing and storage resources and the cost of the infrastructure and its evolution for the WLCG sites. This contribution provides an update on the recent developments of the working group activities, with a special focus on the implications for computing sites.

Desired length

Authors: Dr SCIABÀ, Andrea (CERN); FLIX MOLINA, Jose (Centro de Investigaciones Energéticas Medioambientales y Tecnológicas); SCHULZ, Markus (CERN)

Presenter: Dr SCIABÀ, Andrea (CERN)

Session Classification: IT Facilities & Business Continuity

Track Classification: IT Facilities & Business Continuity

Contribution ID: 22

Type: **not specified**

Commissioning CERN Tier-0 reconstruction workloads on Piz Daint at CSCS

Wednesday 10 October 2018 16:30 (25 minutes)

Predictions for requirements for the LHC computing for Run 3 and for Run 4 (HL_LHC) over the course of the next 10 years show a considerable gap between required and available resources, assuming budgets will globally remain flat at best. This will require some radical changes to the computing models for the data processing of the LHC experiments. The use of large scale general purpose computational resources at HPC centres worldwide is expected to increase substantially the cost-efficiency of the processing. We report on the technical challenges and solution adopted to commission the reconstruction of RAW data from the LHC experiments on Piz Daint at CSCS. This workload is currently executed exclusively at dedicated clusters at the Tier-0 in CERN.

Desired length

20

Authors: SCACCA, Francesco Giovanni (Universitaet Bern (CH)); FERNANDEZ FERNANDEZ, Pablo (ETH Zurich (CH)); GILA, Miguel (ETH Zurich); GILA, Miguel (ETH Zurich (CH))

Presenter: SCACCA, Francesco Giovanni (Universitaet Bern (CH))

Session Classification: Computing & Batch Services

Track Classification: Computing & Batch Services

Contribution ID: 23

Type: **not specified**

LTO experiences at CERN

Thursday 11 October 2018 16:10 (25 minutes)

CERN has purchased LTO-8 tape drives that we have started to use with LTO-8 and LTO-7 type M media. While LTO is attractive in terms of cost/TB, it lacks functionality for fast positioning available on enterprise drives, such as a high-resolution tape directory and drive-assisted file access ordering (aka RAO). In this contribution, we will outline our experiences with LTO so far, describe performance measurements and our current investigations on how to improve positioning times.

Desired length

20 mins

Author: CANCIO MELIA, German (CERN)**Presenter:** CANCIO MELIA, German (CERN)**Session Classification:** Storage & Filesystems**Track Classification:** Storage & Filesystems

Contribution ID: 24

Type: **not specified**

Updates on ATLAS Data Carousel R&D

Thursday 11 October 2018 14:00 (25 minutes)

In this talk, I will give an update to the ATLAS data carousel R&D project, focused mainly on the recent tape performance tests, on the T1 sites. Main topics to be covered include :

- 1) the overall throughput delivered by the tape system
- 2) the overall throughput delivered to the end user (rucio)
- 3) any issues/bottlenecks in the various service layers (tape, SRM, FTS, rucio) observed
- 4) recommendations/suggestions to improve the performance in the future

Desired length

15~20m

Author: Dr ZHAO, Xin (Brookhaven National Laboratory)**Presenter:** Dr ZHAO, Xin (Brookhaven National Laboratory)**Session Classification:** Storage & Filesystems**Track Classification:** Storage & Filesystems

Contribution ID: 25

Type: **not specified**

Evaluation of AMD EPYC

Wednesday 10 October 2018 14:00 (25 minutes)

AMD returned to the server CPU market in 2017 with the release of their EPYC line of CPUs, based on the Zen microarchitecture. In this presentation, we'll provide an overview of the AMD EPYC CPU architecture, and how it differs from Intel's Xeon Skylake. We'll also present performance and cost comparisons between EPYC and Skylake, with an emphasis on use in HEP/NP computing environments.

Desired length

Author: HOLLOWELL, Christopher (Brookhaven National Laboratory)

Presenter: HOLLOWELL, Christopher (Brookhaven National Laboratory)

Session Classification: Computing & Batch Services

Track Classification: Computing & Batch Services

Contribution ID: 26

Type: **not specified**

CSCS site report

Monday 8 October 2018 09:50 (15 minutes)

Site report, news and ongoing activities at the Swiss National Supercomputing Centre T2 site (CSCS-LCG2) running ATLAS, CMS and LHCb.

- Brief description about the site, location, size, resource plans until 2020
- Complexities of the collaboration between the 4 parties
- Next steps after LHConCRAY and Tier-0 spillover tests

Desired length

12

Authors: GILA, Miguel (ETH Zurich); Mr CONCIATORE, Dino (CSCS (Swiss National Supercomputing Centre)); FERNANDEZ FERNANDEZ, Pablo (ETH Zurich (CH)); PETRUSIC, Dario (Eidgenoessische Technische Hochschule Zuerich (ETHZ) (CH)); RICCIARDI, Gianni Mario (Eidgenoessische Technische Hochschule Zuerich (CH)); KLJNSMA, Thomas (ETH Zurich (CH)); BERNET, Roland (Universitaet Zuerich (CH)); SCIACCA, Francesco Giovanni (Universitaet Bern (CH)); Mr PASSERINI, Marco (ETHZ - ETH Zurich)

Presenter: Mr PASSERINI, Marco (ETHZ - ETH Zurich)

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 27

Type: **not specified**

Storage at CERN

Thursday 11 October 2018 11:10 (25 minutes)

The CERN IT Storage group operates multiple distributed storage systems and it is responsible for the support of the CERN storage infrastructure, ranging from the physics data of the LHC and non-LHC experiments to users' files.

This talk will summarise our experience and the ongoing work in evolving our infrastructure, focusing on some of the most important areas.

EOS is the high-performance distributed file system from CERN which allows to operate at high incoming data rates for experiments' data taking and running complex production workloads. EOS is also the key component behind CERNBox, the cloud storage service with the goal to provide sync and share files across all major mobile and desktop platforms for EOS data. Due to its popularity and the integration with other CERN services (Batch, SWAN, Microsoft Office, other Office suites) CERNBox is experiencing an exponential usage growth.

CASTOR is the CERN system for the experimental data recording while we are developing CTA, the next-generation tape archival system.

The storage group provides also the storage backbone of the OpenStack infrastructure with a large Ceph deployment, offering also S3 functionality and CephFS for internal storage and to support the CERN HPC facility.

Our group also operates the Stratum 0 and the CERN Stratum 1 for CVMFS (grid data distribution system) and, in connection with the CERN EP-SFT group (in charge with CVMFS development), we are contributing to the CVMFS project in order to cope with new use cases from the experiments.

Desired length

Author: Mr CONTESCU, Cristian (CERN)

Presenter: Mr CONTESCU, Cristian (CERN)

Session Classification: Storage & Filesystems

Track Classification: Storage & Filesystems

Contribution ID: 28

Type: **not specified**

Benchmarking Update

Wednesday 10 October 2018 14:25 (25 minutes)

The HEPiX Benchmarking Working Group is working on a new 'long-running' benchmark to measure installed capacities to replace the currently used HS06. This presentation will show the current status.

Desired length

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

Co-authors: GIORDANO, Domenico (CERN); MICHELOTTO, Michele (Università e INFN, Padova (IT))

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

Session Classification: Computing & Batch Services

Track Classification: Computing & Batch Services

Contribution ID: 29

Type: **not specified**

Network configuration management at CERN: status and outlook

Monday 8 October 2018 14:00 (25 minutes)

CERN's networks comprise approximately 400 routers and 4000 switches from multiple vendors and from different generations, fulfilling various purposes (campus network, datacentre network, and dedicated networks for the LHC accelerator and experiments control).

To ensure the reliability of the networks, the IT Communication Systems group has developed an in-house Perl-based software called "cfmgr" capable of deriving and enforcing the appropriate configuration on all these network devices, based on information from a central network database. While cfmgr has been continuously extended and enhanced over the past 20 years, it has become increasingly challenging to maintain and further expand it due to the decrease in popularity of the technologies it relies upon (programming language and available libraries). Hence, we have evaluated the functionality of various open-source network configuration tools (e.g. NAPALM, Ansible, SaltStack and StackStorm), in order to understand whether we can leverage them and evolve the cfmgr platform.

We will present the result of this evaluation, as well as the plan for evolving CERN's network configuration management by decoupling the configuration generation (CERN specific) from the configuration enforcement (generic problem, partially addressed by vendor or community Python based libraries).

Desired length

20

Authors: SHEVRIKUKO, Arkadiy (CERN); STANCU, Stefan Nicolae (CERN)

Presenter: SHEVRIKUKO, Arkadiy (CERN)

Session Classification: Networking & Security

Track Classification: Networking & Security

Contribution ID: **30**Type: **not specified**

BNL Site Report

Monday 8 October 2018 11:50 (15 minutes)

The Scientific Data & Computing Center (SDCC) at Brookhaven National Laboratory (BNL) serves the computing needs of experiments at RHIC, while also serving as the US ATLAS Tier-1, as well as Belle-2 Tier-1 facility. This presentation provides an overview of the BNL SDCC, highlighting significant developments since the last HEPiX meeting at University of Wisconsin-Madison.

Desired length

12

Author: DE STEFANO JR, John Steven (Brookhaven National Laboratory (US))

Presenter: DE STEFANO JR, John Steven (Brookhaven National Laboratory (US))

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 31

Type: **not specified**

Workflows automation @CERN IT with CERNMegabus

Tuesday 9 October 2018 11:00 (25 minutes)

In the CERN IT agile infrastructure, Puppet is the key asset for automated configuration management of the more than forty thousand machines in the CERN Computer Centre. The large number of virtual and physical machines runs a variety of services that need to interact with each other at run time.

These needs triggered the creation the CERNMegas project, which automates the communication between services on the base of message subscription, message publishing and reaction on message revival.

This presentation focuses on the two recent use cases implemented with the CERNMegabus in the CERN IT. The first is the synchronous management of the application state of machine. The second one is the handling of the Power Cut automation in the CERN Computer Centre.

Desired length

20

Author: TOTEVA, Zhechka (CERN)**Presenter:** TOTEVA, Zhechka (CERN)**Session Classification:** Basic IT Services**Track Classification:** Basic IT Services

Contribution ID: 32

Type: **not specified**

WLCG/OSG Networking Update

Monday 8 October 2018 14:50 (25 minutes)

WLCG relies on the network as a critical part of its infrastructure and therefore needs to guarantee effective network usage and prompt detection and resolution of any network issues, including connection failures, congestion and traffic routing. The 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 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 WLCG/OSG perfSONAR infrastructure as well as cover recent changes in the higher level services that were developed to help bring perfSONAR network to its full potential. This includes new features and highlights from the perfSONAR 4.1 release as well as changes and updates made to the OSG central services. We'll also give details on the two new NSF-funded projects related to the OSG networking area, namely Service Analysis and Network Diagnostics (SAND) and Institute for Research and Innovation in Software for High Energy Physics (IRIS-HEP).

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

Desired length

20

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

Presenter: BABIK, Marian (CERN)

Session Classification: Networking & Security

Track Classification: Networking & Security

Contribution ID: 33

Type: **not specified**

Network Functions Virtualisation Working Group Update

Monday 8 October 2018 16:10 (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 lead 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 Open Virtual Network (OVN), Tungsten, etc., 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 working group activities with details on recent projects, updates from sites on their adoption as well as plans for the near-term future.

Desired length

15

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

Contribution ID: 35

Type: **not specified**

Latest developments of the CERN Data Management Tools

Thursday 11 October 2018 12:00 (25 minutes)

The CERN IT-Storage group Analytics and Development section is responsible for the development of Data Management solutions for Disk Storage and Data Transfer. These solutions include EOS - the high-performance CERN IT distributed storage for High-Energy Physics, DPM - a system for managing disk storage at small and medium sites, and FTS - the service responsible for distributing the majority of LHC data across the WLCG infrastructure.

This presentation will give an overview of the solutions offered by the section, highlighting the latest developments and it will also cover the task of migrating 14000 CERNBox users to a new EOS architecture.

Desired length

Author: Dr KARAVAKIS, Edward (CERN)

Co-authors: BITZES, Georgios (CERN); CONTESCU, Cristian (CERN); DUELLMANN, Dirk (CERN); FURANO, Fabrizio (CERN); GONZALEZ LABRADOR, Hugo (CERN); LAMANNA, Massimo (CERN); MANZI, Andrea (CERN); MASCETTI, Luca (CERN); MOSCICKI, Jakub (CERN); PETERS, Andreas Joachim (CERN); SIMON, Michal Kamil (CERN); SINDRILARU, Elvin Alin (CERN)

Presenter: Dr KARAVAKIS, Edward (CERN)

Session Classification: Storage & Filesystems

Track Classification: Storage & Filesystems

Contribution ID: 37

Type: **not specified**

Good time with data: faas fn

Wednesday 10 October 2018 09:35 (25 minutes)

Presenting what is Function as a service (FaaS) and how it can be used in the data center. Will be showing how NERSC is using FaaS to get added functionality to the data we collect. Also presented will be some basic python on to perform these functions and how to connect to Elastic both natively and via the API.

Desired length

20

Author: WHITNEY, Cary (LBNL)**Presenter:** WHITNEY, Cary (LBNL)**Session Classification:** Grid, Cloud & Virtualisation**Track Classification:** Grid, Cloud & Virtualisation

Contribution ID: **38**Type: **not specified**

Superfacility

Thursday 11 October 2018 10:15 (25 minutes)

This talk will be about the proposed superfacility structure that NERSC is working toward. This is the next model of HPC computing combining all aspects of the data center, infrastructure, WAN and the experiments. We are in the early stages of defining and standing up such a facility.

Desired length

20 or could be 12 if need be

Author: WHITNEY, Cary (LBNL)

Presenter: WHITNEY, Cary (LBNL)

Session Classification: IT Facilities & Business Continuity

Track Classification: IT Facilities & Business Continuity

Contribution ID: 39

Type: **not specified**

PDSF Site Report

Tuesday 9 October 2018 09:00 (15 minutes)

PDSF, the Parallel Distributed Systems Facility 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 running containerized software stacks using Shifter, as well as upcoming changes to systems management and the future of PDSF and the move of the workload into the Cori system.

Desired length

12

Author: RATH, Georg (Lawrence Berkeley National Laboratory)

Co-authors: Dr BALEWSKI, Jan (LBNL); QUAN, Tony (LBL); Dr JACOBSEN, Douglas (LBNL); DECLERCK, Tina (LBNL); BASHEER, Ershaad (LBNL)

Presenter: RATH, Georg (Lawrence Berkeley National Laboratory)

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 40

Type: **not specified**

The new CERN Authentication and Authorization

Tuesday 9 October 2018 11:50 (25 minutes)

The current authentication schemes used at CERN are based on Kerberos for desktop and terminal access, and on Single Sign-On (SSO) tokens for web-based applications.

Authorization schemes are managed through LDAP groups, leading to privacy concerns and requiring a CERN accounts to make possible the mapping to a group.

This scenario is completely separated from WLCG, where authentication is based on X509 certificates, which are mapped to Virtual Organizations for authorization.

While this solution covers the required use cases, it provides a difficult user experience.

Several initiatives, like Indigo-Datacloud platform or the SciTokens projects, are aiming at providing alternative authentication and authorization schemes based on tokens.

The ongoing redesign of the CERN authentication and authorization infrastructures is an occasion to harmonize the different authentication schemes used at CERN, and close the gap between IT's offer and HEP needs, aligning with the token-based authentication schemes that the WLCG is heading towards, and allowing for a full integration between the two worlds.

The new services will provide full support for a federated environment, where users authenticate with their home institute credentials or social account, and more uniform authorization schemes, with builtin privacy.

Desired length

20

Authors: TEDESCO, Paolo (CERN); ORMANCEY, Emmanuel (CERN); SHORT, Hannah (CERN); SMITH, Tim (CERN)

Presenter: TEDESCO, Paolo (CERN)

Session Classification: Basic IT Services

Track Classification: Basic IT Services

Contribution ID: 42

Type: **not specified**

PDSF - Current Status and Migration to Cori

Wednesday 10 October 2018 16:55 (20 minutes)

PDSF, the Parallel Distributed Systems Facility, has been in continuous operation since 1996 serving high-energy and nuclear 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 the PDSF workload from the existing commodity cluster to the Cori Cray XC40 system. Docker containers enable running the PDSF software stack on different hardware and OS. We will discuss challenges of using highly scalable Cori resources when thousands of user jobs can start within a second and easily saturate IO resources, CVMFS, or external database connectivity.

Desired length

15

Author: BALEWSKI, Jan (Lawrence Berkeley National Lab. (US))

Co-authors: RATH, Georg (Lawrence Berkeley National Laboratory); PORTER, Jeff (Lawrence Berkeley National Lab. (US)); LEE, Rei (Lawrence Berkeley National Laboratory); QUAN, Tony (LBL)

Presenter: BALEWSKI, Jan (Lawrence Berkeley National Lab. (US))

Session Classification: Computing & Batch Services

Track Classification: Computing & Batch Services

Contribution ID: 43

Type: **not specified**

Welcome

Monday 8 October 2018 09:00 (15 minutes)

Introduction to Barcelona and PIC

Desired length

Presenter: MERINO, Gonzalo (IFAE - Institute for High Energy Physics)

Session Classification: Miscellaneous

Track Classification: Miscellaneous

Contribution ID: 44

Type: **not specified**

Workshop logistics

Monday 8 October 2018 09:15 (15 minutes)

Desired length

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

Session Classification: Miscellaneous

Track Classification: Miscellaneous

Contribution ID: 45

Type: **not specified**

Workshop wrap-up

Friday 12 October 2018 11:00 (20 minutes)

Desired length

Presenter: MEINHARD, Helge (CERN)

Session Classification: Miscellaneous

Track Classification: Miscellaneous

Contribution ID: 46

Type: **not specified**

LAL Site Report

Tuesday 9 October 2018 09:15 (15 minutes)

Update on computing at LAL

Desired length

Author: JOUVIN, Michel (Centre National de la Recherche Scientifique (FR))

Presenter: JOUVIN, Michel (Centre National de la Recherche Scientifique (FR))

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 47

Type: **not specified**

The OSiRIS Project: A Multi-institutional Ceph Storage Infrastructure

Thursday 11 October 2018 14:25 (25 minutes)

The NSF-funded OSiRIS project (<http://www.osris.org>), which is creating a multi-institutional storage infrastructure based upon Ceph and SDN is entering its fourth year. We will describe the project, its science domain users and use cases, and the technical and non-technical challenges the project has faced. We will conclude by outlining the plans for the remaining two years of the project and its longer term prospects.

Desired length

20 minutes

Authors: MC KEE, Shawn (University of Michigan (US)); MEEKHOF, Benjeman Jay (University of Michigan (US))

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

Session Classification: Storage & Filesystems

Track Classification: Storage & Filesystems

Contribution ID: 48

Type: **not specified**

The ATLAS Great Lakes Tier-2 Site Update

Monday 8 October 2018 10:05 (15 minutes)

We will present an update on AGLT2, focusing on the changes since the Spring 2018 report. The primary topics to cover include status of IPv6, an update on VMware, dCache and our Bro/MISP deployment. We will also describe the new equipment purchases and personnel changes for our site.

Desired length

12 minutes

Authors: BALL, Robert (University of Michigan (US)); Prof. WU, Wenjing (Computer Center, IHEP, CAS); MC KEE, Shawn (University of Michigan (US)); LAURENS, Philippe (Michigan State University (US)); BROCK, Raymond (Michigan State University)

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

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 49

Type: **not specified**

A Data Lake Prototype for HL-LHC

Wednesday 10 October 2018 10:00 (25 minutes)

Evolving the current data storage, management and access models is the main challenge in WLCG and certainly in scientific computing for the coming years. HL-LHC will be exceeding what funding agencies can provide by an order of magnitude. Forthcoming experiments in particle physics, cosmology and astrophysics also foresee similar magnitudes in data volumes. The concepts of storage federations, data redundancy, transfer protocols and data access need to be revisited with the clear goal to reduce costs while providing the performance required by the experiments in the next decade.

In this contribution we are presenting the experience with a primary Data Lake prototype for WLCG implementing some of the main ideas behind a future concept of storage federations.

Desired length

Authors: ESPINAL, Xavier (CERN); CAMPANA, Simone (CERN); BIRD, Ian (CERN); GIRONE, Maria (CERN); SCHOVANCOVA, Jaroslava (CERN); MCCANCE, Gavin (CERN)

Presenter: ESPINAL, Xavier (CERN)

Session Classification: Grid, Cloud & Virtualisation

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: 50

Type: **not specified**

KEK Site Report

Monday 8 October 2018 10:20 (15 minutes)

All projects hosted at KEK have actively proceeded in 2018. The SuperKEKB/Belle II experiment succeeded in observing the first collisions in April 2018 and accumulating the data until July continuously. The J-PARC accelerator has also provided various kinds of the beam simultaneously. Most of the experimental data are transferred to the KEK central computer system (KEKCC) and stored. In this talk, I would like to report on the status of KEKCC at the Belle II data taking period, the configuration of the private network between J-PARC and KEKCC, and the KEK campus network system which is recently upgraded.

Desired length

12 minutes

Author: NAKAMURA, Tomoaki (High Energy Accelerator Research Organization (JP))**Co-authors:** IWAI, Go (KEK); SUZUKI, Soh; SUZUKI, Jiro (KEK)**Presenter:** NAKAMURA, Tomoaki (High Energy Accelerator Research Organization (JP))**Session Classification:** Site Reports**Track Classification:** Site Reports

Contribution ID: 51

Type: **not specified**

NDGF Site report

Monday 8 October 2018 12:05 (15 minutes)

News and updates from the distributed NDGF Tier1 site.

Desired length

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

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

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 52

Type: **not specified**

RAID is dead

Thursday 11 October 2018 15:15 (25 minutes)

FlexiRemap® Technology

Patented design, built from the ground up for replacing disk- based RAID in flash storage arrays.

Faster performance, longer SSD lifespan, and advanced data protection over traditional RAID technology.

Desired length

15

Author: PATTISON, Alex (AccelStor Ltd)

Presenter: PATTISON, Alex (AccelStor Ltd)

Session Classification: Storage & Filesystems

Track Classification: Storage & Filesystems

Contribution ID: 53

Type: **not specified**

Data Protection at CERN IT

Wednesday 10 October 2018 11:45 (25 minutes)

The processing of personal data is inevitable within a work context and CERN aims to follow and observe best practices as regards the collection, processing and handling of this type of data. This talk aims to give an overview on how CERN and especially the IT department is implementing Data Protection.

Desired length

Author: DUMITRU, Andrei (CERN)

Presenter: DUMITRU, Andrei (CERN)

Session Classification: Networking & Security

Track Classification: Networking & Security

Contribution ID: 54

Type: **not specified**

CERN Tape Archive initial deployments

Friday 12 October 2018 09:50 (25 minutes)

CTA is designed to replace CASTOR as the CERN Tape Archive solution, in order to face scalability and performance challenges arriving with LHC Run-3.

This presentation will give an overview of the initial software deployment on production grade infrastructure. We discuss its performance against various workloads: from artificial stress tests to production condition data transfer sessions with an LHC experiment.

It will also cover CTA's recent participation to the Heavy Ion Data Challenge and a roadmap for future deployments.

Desired length

20 minutes

Author: LEDUC, Julien (CERN)

Presenter: LEDUC, Julien (CERN)

Session Classification: Storage & Filesystems

Track Classification: Storage & Filesystems

Contribution ID: 55

Type: **not specified**

Deployment of compute nodes for the WLCG with Cobbler, Ansible and Salt.

Tuesday 9 October 2018 11:25 (25 minutes)

In Belgium, the ‘Université catholique de Louvain’ (UCLouvain) hosts a Tier-2 WLCG site. The computing infrastructure has recently been merged with the General Purpose cluster of the university. During that merge, the deployment process for the compute nodes has been re-thought, using a combination of three open-source software tools: Cobbler, Ansible and Salt. Those three tools work together to deploy the operating system, install software, configure services, register the new nodes into the inventory, the monitoring system, the resource manager/job scheduler, etc. The setup follows the ‘Infrastructure as Code’ principles, and can adapt to the evolving infrastructure, with nodes being added from times to times depending on funding, and others being decommissioned when they reach a venerable age.

Eventhough Ansible and Salt are often seen as exclusive alternatives in the Devops community, we believe they complement each other very well because they have very different strengths and weaknesses.

In this talk, we will present our setup, how Cobbler, Ansible and Salt interact to go from compute node unboxing to accepting Grid jobs in three smooth operations. We will explain how we decide which tool takes care of which task. We will also present how Salt and Ansible together ease-up the installation of re-compiled software, alongside the usual pre-compiled, CVMFS-distributed, software.

Desired length

20 minutes

Authors: Dr FRANÇOIS, Damien (Catholic University of Louvain); MATTELAER, Olivier (UCLouvain); KEUTGEN, Thomas (Institut de Physique Nucleaire-Departement de Physique-Universit)

Presenter: Dr FRANÇOIS, Damien (Catholic University of Louvain)

Session Classification: Basic IT Services

Track Classification: Basic IT Services

Contribution ID: 56

Type: **not specified**

Backup Infrastructure at CERN

Friday 12 October 2018 10:45 (15 minutes)

CERN's Backup Service hosts around 11 PB of data, in more than 2.1 billion files. We have over 500 clients which back up or restore an average of 80 TB of data each day. At the current growth rate, we expect to have about 13 PB by the end of 2018.

In this contribution we review the impact of the latest changes of the backup infrastructure. We will see how these optimizations helped us to offer a reliable and efficient service, the benefits provided on the day to day operation, and possible future evolutions of the Backup Service.

Desired length

10

Author: GOUNON, Aurelien (CERN)**Presenter:** GOUNON, Aurelien (CERN)**Session Classification:** Storage & Filesystems**Track Classification:** Storage & Filesystems

Contribution ID: 57

Type: **not specified**

Container orchestration as key to the evolution of CERN Web Services

Tuesday 9 October 2018 16:10 (25 minutes)

Introduced in 2016, the “PaaS for Web Applications” service aims at providing CERN users with a modern environment for web application hosting, following the Platform-as-a-Service paradigm. The service leverages the Openshift (now OKD) container orchestrator.

We will provide a quick overview of the project, its use cases and how it evolved over its more than two years of production phase. Then we will focus on integration of Openshift with the CERN environment, and in particular how integration with GitLab allowed to significantly simplify and automate application deployment. We will conclude with the prospects the platform offers for the evolution of CERN Web Services.

Desired length

20

Author: LOSSENT, Alexandre (CERN)**Presenter:** LOSSENT, Alexandre (CERN)**Session Classification:** End-User IT Services & Operating Systems**Track Classification:** End-User IT Services & Operating Systems

Contribution ID: 58

Type: **not specified**

Jefferson Lab Scientific and High Performance Computing

Monday 8 October 2018 12:20 (15 minutes)

JLab high performance and experimental physics computing environment updates since the Fall 2016 meeting, with recent and upcoming hardware procurements for compute nodes, including Skylake, Volta, and/or Intel KNL accelerators; our Supermicro storage; Lustre status; 12GeV computing status; integrating offsite resources; Data Center modernization.

Desired length

12

Author: PHILPOTT, Sandra**Presenter:** PHILPOTT, Sandra**Session Classification:** Site Reports**Track Classification:** Site Reports

Contribution ID: 59

Type: **not specified**

IPv6 & WLCG - update from the HEPiX IPv6 working group

Monday 8 October 2018 15:15 (25 minutes)

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.

All WLCG Tier 1 data centres have IPv6 connectivity and much of their storage is now accessible over IPv6. The LHC experiments have also requested all WLCG Tier 2 centres to provide dual-stack access to their storage by the end of LHC Run 2 (end 2018). 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. I will present the progress of the transition to IPv6 together with current issues.

Desired length

20 minutes

Author: KELSEY, David (Science and Technology Facilities Council STFC (GB))

Presenter: KELSEY, David (Science and Technology Facilities Council STFC (GB))

Session Classification: Networking & Security

Track Classification: Networking & Security

Contribution ID: **60**

Type: **not specified**

Technology watch working group

Thursday 11 October 2018 09:00 (25 minutes)

Report from kick-off meeting

Desired length

20 min

Author: MEINHARD, Helge (CERN)

Presenter: MEINHARD, Helge (CERN)

Session Classification: IT Facilities & Business Continuity

Track Classification: IT Facilities & Business Continuity

Contribution ID: 61

Type: **not specified**

Too much BEER and the Kubernetes dance

Wednesday 10 October 2018 17:15 (25 minutes)

With the demands of LHC computing, coupled with pressure on the traditional resources available, we need to find new sources of compute power. We have described, at HEPiX and elsewhere how we have started to explore running batch workloads on storage servers at CERN, and on public cloud resources. Since the summer of 2018, ATLAS & LHCb have started to use a pre-production service on storage services (BEER) on 3k cores. We have run cloud services with Helix Nebula, Oracle and others at the scale of 10k cores, till now using fat VMs configured using traditional configuration management but with containers for the payload. The kubernetes ecosystem would seem to promise a simplification of public cloud deployments, as well as internal use cases.

The talk will therefore report on the gained experience of running BEER services at some scale, and at the progress and motivation of putting the condor worker itself in a container, managed with kubernetes.

Desired length

20

Authors: JONES, Ben (CERN); FERNANDEZ ALVAREZ, Luis (CERN)**Presenter:** JONES, Ben (CERN)**Session Classification:** Computing & Batch Services**Track Classification:** Computing & Batch Services

Contribution ID: **62**

Type: **not specified**

Surfsara site report

Tuesday 9 October 2018 09:30 (15 minutes)

Site report for Surfsara, formerly known as SARA, part of the Dutch Tier1 site.

Desired length

12

Author: ZWEERS, Onno (SURFsara)

Presenter: ZWEERS, Onno (SURFsara)

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 64

Type: **not specified**

Tokyo Tier-2 Site Report

Tuesday 9 October 2018 10:00 (15 minutes)

The Tokyo regional analysis center located at International Center for Elementary Particle Physics in the University of Tokyo supports ATLAS VO as one of the WLCG Tier2 sites.

It provides 10,000 CPU cores and 10 PB disk storage including the local resource dedicated for the ATLAS-Japan member.

All hardware devices are supplied by the three years rental, and current contract will finish at the end of this year.

In this presentation, we will report the current status and updates on Tokyo Tier-2 operation and the prospect for the new system scheduled to be launched at the beginning of 2019.

Desired length

12

Author: KANEDA, Michiru (ICEPP, the University of Tokyo)

Presenter: KANEDA, Michiru (ICEPP, the University of Tokyo)

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 65

Type: **not specified**

Prague Site Report

Tuesday 9 October 2018 09:45 (15 minutes)

We will give an overview of the site and will share experience with these topics: update to the current release of DPM, HT-Condor configuration, Foreman installation and setup.

Desired length

Authors: ADAM, Martin (Acad. of Sciences of the Czech Rep. (CZ)); CHUDOBA, Jiri (Acad. of Sciences of the Czech Rep. (CZ)); MIKULA, Alexandr (Acad. of Sciences of the Czech Rep. (CZ)); HORAK, Petr (Acad. of Sciences of the Czech Rep. (CZ)); VOKAC, Petr (Czech Technical University (CZ)); SVATOS, Michal (Acad. of Sciences of the Czech Rep. (CZ)); ADAMOVA, Dagmar (Acad. of Sciences of the Czech Rep. (CZ))

Presenter: ADAM, Martin (Acad. of Sciences of the Czech Rep. (CZ))

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 66

Type: **not specified**

KISTI-GSDC site report

Tuesday 9 October 2018 10:15 (15 minutes)

KISTI-GSDC site report

Desired length

Authors: KIM, Jeongheon (Korea Institute of Science and Technology Information); AHN, Sang Un (Korea Institute of Science & Technology Information (KR)); Mr KIM, Jin (Korea Institute of Science and Technology Information); NA, SangHo (KISTI)

Presenter: NA, SangHo (KISTI)

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 67

Type: **not specified**

A Framework for Open Science Cybersecurity Programs

Wednesday 10 October 2018 12:10 (25 minutes)

Trusted CI, the National Science Foundation's Cybersecurity Center of Excellence, is in the process of updating their Guide (published in 2014) and recasting it as a framework for establishing and maintaining a cybersecurity program for open science projects. The framework is envisioned as being appropriate for a wide range of projects, both in terms of scale and life cycle. The Framework is built around the four pillars of Mission Alignment; Governance; Resources; and Controls and then addresses the additional requirements of day-to-day Operations. The talk will cover both the content of the Framework and current thoughts about the presentation design in hopes of getting feedback from the community. The framework should provide a valuable starting point even for science projects facing more severe compliance requirements.

Desired length

20 minutes

Author: COWLES, Bob (BrightLite Information Security)**Presenter:** COWLES, Bob (BrightLite Information Security)**Session Classification:** Networking & Security**Track Classification:** Networking & Security

Contribution ID: 68

Type: **not specified**

Report on the Workshop on Central Computing Facilities Support for Photon Sciences

Wednesday 10 October 2018 15:15 (25 minutes)

This is a report on the recently held workshop at BNL on Central Computing Facilities support for Photon Sciences, with participation from various Light Source facilities from Europe and the US.

Desired length

Author: WONG, Tony (Brookhaven National Laboratory)

Presenter: WONG, Tony (Brookhaven National Laboratory)

Session Classification: Computing & Batch Services

Track Classification: Computing & Batch Services

Contribution ID: 69

Type: **not specified**

Data Analysis as a service for MIRAS beamline users on ALBA synchrotron

Wednesday 10 October 2018 14:50 (25 minutes)

The classic workflow of a experiment in a synchrotron facility starts with the users coming physically to the facility with their samples, they analyze those samples with the beamline equipment and finally they get back to their institution with a huge amount of data in a portable hard disk. The data reduction and analysis is done majorly on the scientific institution of the user. As data reduction and analysis often requires specialized software and a huge amount of computing resources the classical approach leaves no chance to small scientific organisms that cannot afford either the licenses or the computing resources.

Lately some synchrotron facilities are starting to offer data analysis as a service in order to extend the services offered to the visitor users, minimizing the needed data transfer, extending the analysis periods and reinforcing their source scientific organisms capacities.

ALBA synchrotron offers Data Analysis as a service (DaaS) to the users of MIRAS beamline. We are going to explain how ALBA is using Citrix Virtual Desktop Windows based Infrastructure in order to provide that DaaS.

Furthermore, we will include also some insights about other related technologies we used in similar projects such as Linux Virtual Desktops Infrastructure and Virtual GPUs.

Desired length

12

Authors: Mr PUSO, Sergi (ALBA synchrotron); SALABERT, Jordi (ALBA Synchrotron)

Co-authors: Mr PIJUAN, Joan (ALBA synchrotron); Mr VICENTE, Sergio (ALBA synchrotron); Mr PEREZ, Toni (ALBA synchrotron); Ms ROSAS, Gemma (ALBA synchrotron)

Presenter: SALABERT, Jordi (ALBA Synchrotron)

Session Classification: Computing & Batch Services

Track Classification: Computing & Batch Services

Contribution ID: 70

Type: **not specified**

Update on the BNL SDCC Notebook Based Analysis Portal

Tuesday 9 October 2018 16:35 (25 minutes)

The BNL Scientific Data and Computing Center (SDCC) has been developing a user analysis portal based on Jupyterhub and leveraging the large scale computing resources available at SDCC. We present the current status of the portal and issues of growing and integrating the user base.

Desired length

Author: RIND, Ofer**Presenter:** RIND, Ofer**Session Classification:** End-User IT Services & Operating Systems**Track Classification:** End-User IT Services & Operating Systems

Contribution ID: 71

Type: **not specified**

Report from HTCondor workshop

Friday 12 October 2018 09:00 (25 minutes)

Short report on the workshop held at RAL in September, and an outlook to the next workshop

Desired length

20

Author: MEINHARD, Helge (CERN)

Presenter: MEINHARD, Helge (CERN)

Session Classification: Computing & Batch Services

Track Classification: Standby

Contribution ID: 72

Type: **not specified**

IHEP site report

Wednesday 10 October 2018 09:00 (15 minutes)

IHEP computing platform has been running a Tier2 WLCG site and a 12,000 local HTCluster with ~9PB storage. The talk will talk about the optimization of HTCondor local cluster, the next plan for HPC cluster, the progress has been done with IHEP campus network and new functions provided to users.

Desired length

12

Author: ZOU, Jiaheng (IHEP)**Presenter:** ZOU, Jiaheng (IHEP)**Session Classification:** Site Reports**Track Classification:** Standby

Contribution ID: 73

Type: **not specified**

IHEP Campus Network design based on SDN Technology

Wednesday 10 October 2018 10:55 (25 minutes)

The report introduces the design of the IHEP SDN campus network which aims to the separation of the control plane and the data forwarding plane. The new designed SDN network, integrated with IHEP authentication system, has achieved 802.1X-based user access control. It can obtain the network information in the data forwarding plane through the controller and provide more network management functions. The test-bed has been built at the vendor's research center and their research team is cooperating with us to complete functional verification and API development.

Desired length

20

Author: CUI, Tao (IHEP(Institute of High Energy Physics, CAS,China))**Presenter:** CUI, Tao (IHEP(Institute of High Energy Physics, CAS,China))**Session Classification:** Networking & Security**Track Classification:** Standby

Contribution ID: 74

Type: **not specified**

Workshop Photographs

Friday 12 October 2018 11:20 (10 minutes)

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

Session Classification: Miscellaneous