

HEPiX Autumn 2021 online Workshop

Monday 25 October 2021 - Friday 29 October 2021

Book of Abstracts

Contents

HIFIS: VO Federation for EFSS 58	1
ASGC site report 59	1
Zammad: the new ticketing system used at CC-IN2P3 60	2
Thinking About Binary Compatibility and CentOS Stream 61	2
INFN-T1 Site report 62	3
Tape Recall Scheduling Optimization in the Storage System dCache 63	3
AGLT2 Site Report 64	4
Navigating the Stream: automating CentOS releases at CERN 65	4
DESY Site Report 66	5
Full OS image automation: Continuous Integration for CERN's distros 67	5
HEPiX Benchmarking WG update 68	6
WLCG task force on HEP-SCORE deployment 69	6
Diamond Light Source Site Report 70	7
IRFU Site Report 71	7
BNL site report 72	7
KEK site report 73	8
Hardware and batch monitoring at CNAF 74	8
Migration of Archival Tape/Library Technology at SDCC 75	9
Linux OS Plans at The SDCC 76	9
PIC report 77	10
WLCG Network Monitoring and Analytics Update 78	10
Research Networking Technical WG Update 79	11
High-throughput EOS instance for ALICE O2 (online/offline) 80	11

Update from the HEPiX IPv6 working group 81	12
CERN Linux Distro Strategy 82	12
Discussion 85	12
Deployment of a Kubernetes-based Analysis Facility for ATLAS 88	13
IHEP Site Report 89	13
EOS 5 new features and future plans 90	13
Computer Security Update 91	14
RAL Site Report 92	15
Computer Security Update (repeat of 11:00h contrib) 93	15
CERN Site Report 94	15
Report from the HTCondor Workshop September 2021 95	15
Automated Backup and Recovery for the Database on Demand service using EOS 96	16
Migration from CASTOR to the CERN Tape Archive 97	16
Helmholtz IT Services for Science: HIFIS 98	17
Analysis on distributed computing at the Belle II experiment 100	17
Group Photo 101	18
Group Photo 102	18
Migrating from a commercial system to Open Source: Gitlab at DESY 103	18
Survival kit for CERN's organic webscape: Kubernetes and a community of Makers 104	18
From On-Premise to On-Cloud videoconference service 105	19
Running an AV media post-processing system on FOSS 106	20
LoRaWAN at CERN vs pandemics: 1:0 107	20
Threat Intelligence and Security Operations Centres: Collaborative Security 108	20
Resource Trust Evolution: Host Certificates in a Dynamic Landscape 109	21
DESY's user registry 110	22
powerful, cross-disciplinary storage for everyone 111	22
HEPiX closing remarks 112	23
DESY Linux Status & Strategy 113	23
Red Hat's RHEL Academic Model 114	23
The RX Extended SACK Table Proposal - AuriStor RX Prototype Results 115	23

Welcome & Logistics 116	24
Welcome & Logistics 117	24
Discussion on user security awareness 118	24
Linux Distribution BoF (cont'd) 119	25
Alice 120	25
ATLAS 121	25
CMS 122	25
LHCb 123	25
Belle II 124	25
EuXFEL 125	26
General discussion - if needs be into coffee break time 126	26

Grid, Cloud & Virtualisation / 58**HIFIS: VO Federation for EFSS****Authors:** Andreas Klotz¹; Matthias Leander-Knoll²¹ *Helmholtz-Zentrum Berlin*² *KIT***Corresponding Authors:** andreas.klotz@helmholtz-berlin.de, matthias.leander-knoll@kit.edu**HIFIS: VO Federation for EFSS**

Following the first rough ideas on Virtual Organisation (VO; Community AAI ^[1] based group of any size) based Enterprise File Sync&Share (EFSS) Federation ^[2] which were presented by HIFIS ^[3] on CS3 Conference 2021, we have since moved further along working on a first implementation. During Summer 2021, we have clarified the use case and identified the basic technical architecture for this future VO Federation App in Nextcloud.

When users who are distributed across multiple institutes want to collaborate within a Virtual Organisation they currently have two options to do so. One is to use the OCM protocol ^[4] to share files and folders with the individual VO members who are based on remote EFSS instances. This would cause considerable effort on the sharer's side, as they need to keep track of to whom they have shared which content with. The second option is for all VO members to convene on one institution's local EFSS instance, which would cause many redundant accounts and confusion on the user's side. Especially, as they need to know where to log in for working on a specific project and as they have no central entry point for all of their projects on their local EFSS instance.

We want to tackle this issue by enabling users to use federated shares with entire VOs instead of individual users. This way, every user within a VO will receive the share, no matter which EFSS instance they are based on. Updates to VO membership will also be communicated between federation members, resulting in new VO members automatically receiving existing VO shares and former VO members losing access to VO shares. Based on a new interface between EFSS and Community AAI, this whole process is also planned to be GDPR compliant. To ensure that this interface will also work with other Community AAIs, we are collaborating with AARC to create an AARC guideline with the aim of standardizing the interface specifications.

While the initial implementation is set to be done within a Nextcloud environment, the new features will be based on the CS3 APIs ^[5] and consequently be ready to also be implemented by further EFSS vendors.

[1] AARC Blueprint for Community AAIs: <https://aarc-project.eu/architecture/>

[2] 2021 CS3 Contribution: <https://indico.cern.ch/event/970232/contributions/4157924/>

[3] HIFIS Website: <https://hifis.net/>

[4] OCM Project documentation: <https://wiki.geant.org/display/OCM/Open+Cloud+Mesh>

[5] CS3 APIs GitHub page: <https://github.com/cs3org/cs3apis>; CS3 APIS are implemented in the REVA middleware: <https://reva.link/>

Desired slot length:

15

Speaker release:

Yes

Site Reports / 59**ASGC site report**

Authors: Chien-De Li¹; Felix.hung-te Lee¹

¹ *Academia Sinica (TW)*

Corresponding Authors: felix@twgrid.org, chien-de.li@cern.ch

ASGC site report

Desired slot length:

12

Speaker release:

Yes

Basic IT Services / 60

Zammad: the new ticketing system used at CC-IN2P3

Author: Renaud Vernet¹

¹ *CC-IN2P3/CNRS*

Corresponding Author: renaud.vernet@cc.in2p3.fr

In late 2020, CC-IN2P3 decided to look for a service desk ticketing system as an alternative to OTRS. During last winter we carried out a search for a new candidate and chose Zammad. We transitioned to Zammad in summer and are now decommissioning OTRS. This talk will present the features offered by Zammad and the lessons learned from that transition.

Desired slot length:

15' to 20'

Speaker release:

Yes

Birds of a Feather (BoF) session: Linux Discussion / 61

Thinking About Binary Compatibility and CentOS Stream

Author: Patrick Riehecky¹

¹ *Fermi National Accelerator Lab. (US)*

Corresponding Author: jcpunk@gmail.com

CentOS Stream is a great place to develop for what's next in RHEL. But what if I want to have a hybrid infrastructure? How should I think about compatibility between CentOS Stream and released versions of RHEL?

Desired slot length:

20-25

Speaker release:

Yes

Site Reports / 62**INFN-T1 Site report****Author:** Andrea Chierici¹¹ *Universita e INFN, Bologna (IT)***Corresponding Author:** andrea.chierici@cern.ch

A brief update on what's going on at INFN-T1

Desired slot length:

10

Speaker release:

Yes

Storage & File Systems / 63**Tape Recall Scheduling Optimization in the Storage System dCache****Authors:** ALBERT ROSSI¹; Dmitry Litvintsev²; Krishnaveni Chitrapu³; Lea Morschel⁴; Marina Sahakyan^{None}; Paul Millar^{None}; Svenja Meyer⁴; Tigran Mkrtchyan⁵; Vincent Garonne⁶¹ *Fermi National Accelerator Laboratory*² *Fermi National Accelerator Lab. (US)*³ *National Supercomputer Center, Linköping University, Sweden*⁴ *Deutsches Elektronen-Synchrotron DESY*⁵ *DESY*⁶ *Brookhaven National Laboratory (US)***Corresponding Authors:** svenja.meyer@desy.de, arossi@fnal.gov, tigran.mkrtchyan@desy.de, litvinse@fnal.gov, vincent.garonne@cern.ch, marina.sahakyan@desy.de, paul.millar@desy.de, krishnaveni@nsc.liu.se, lea.morschel@desy.de

Given the anticipated increase in the amount of scientific data, it is widely accepted that primarily disk based storage will become prohibitively expensive. Tape based storage, on the other hand, provides a viable and affordable solution for the ever increasing demand for storage space. Coupled with a disk caching layer that temporarily holds a small fraction of the total data volume to allow for low latency access, it turns tape based systems into active archival storage (write once, read many) that imposes additional demands on data flow optimization compared to traditional backup setups (write once, read never). In order to preserve the lifetime of tapes and minimize the inherently higher access latency, different tape usage strategies are being evaluated.

As an important disk storage system for scientific data that transparently handles tape access, dCache is making efforts to contribute to tape recall optimization by introducing a high-level tape recall

request scheduling component within its SRM implementation. This presentation will include first experiences with this component on scientific data in a production environment.

Desired slot length:

Speaker release:

Yes

Site Reports / 64

AGLT2 Site Report

Author: Wenjing Wu¹

Co-authors: Shawn Mc Kee²; Philippe Laurens³

¹ *University of Michigan*

² *University of Michigan (US)*

³ *Michigan State University (US)*

Corresponding Authors: wuwj@umich.edu, shawn.mckee@cern.ch, philippe.laurens@cern.ch

We will present an update on our site since the Fall 2019 report, covering our changes in software, tools and operations. In addition, we will cover significant changes that are underway at both the University of Michigan and Michigan State University sites.

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

Summary

Update on AGLT2 including changes in software, hardware and site configurations and summary of status and future work.

Desired slot length:

Speaker release:

Yes

Birds of a Feather (BoF) session: Linux Discussion / 65

Navigating the Stream: automating CentOS releases at CERN

Author: Alex Iribarren¹

¹ *CERN*

Corresponding Author: alejandro.iribarren@cern.ch

As part of the changing Linux landscape, we now have to support more Linux distributions with higher release cadence than ever before. In order to adapt to these changes, we must automate the entire process to remove the human bottlenecks from the equation.

In this presentation, we will discuss how CERN automates the release of new packages for CentOS Linux 8, CentOS Stream 8 and soon CentOS Stream 9. We produce daily snapshots of upstream content, modify it when necessary, promote packages to production and notify users of changes automatically.

Desired slot length:

20

Speaker release:

Yes

Site Reports / 66

DESY Site Report

Author: Andreas Haupt¹¹ *Deutsches Elektronen-Synchrotron (DESY)***Corresponding Author:** andreas.haupt@desy.de

News from the lab

Desired slot length:**Speaker release:**

Yes

Birds of a Feather (BoF) session: Linux Discussion / 67

Full OS image automation: Continuous Integration for CERN's distros

Author: Daniel Juarez¹¹ *CERN***Corresponding Author:** daniel.juarez.gonzalez@cern.ch

Back in 2019, CERN Linux Support had to run tedious manual procedures to maintain CERN's distro releases: SLC6, CERN CentOS 7, Red Hat 6 and 7. Since then, we have added CentOS 8, CentOS Stream 8, Red Hat 8, and we may be adding other Red Hat rebuilds soon. Given the growing number of supported distros, our team has been increasingly adopting automation and continuous integration in order to deal with all the extra load while reducing the need for human intervention.

Automation can now be found in every part of our process: cloud and Docker image building, baseline testing, CERN specific testing and full-stack functional testing. For this we use a combination of GitLab CI capabilities, Koji, OpenStack Nova, OpenStack Ironic central services and a healthy dose of Python and Bash. Test suites now cover unmanaged, managed, virtual or physical machines so we can certify that our next image release continues to meet the needs of the organization.

Desired slot length:

15-20min

Speaker release:

Yes

Computing & Batch Services / 68

HEPiX Benchmarking WG update

Authors: Domenico Giordano¹; Michele Michelotto²; Manfred Alef³

¹ *CERN*

² *Universita e INFN, Padova (IT)*

³ *Karlsruhe Institute of Technology (KIT)*

Corresponding Authors: domenico.giordano@cern.ch, michele.michelotto@cern.ch, manfred.alef@kit.edu

The Benchmarking WG will present the semestral report about its activity that since few years is focused on the delivery of a new benchmark for HEP, HEPscore, together with a set of software tools that enable the WLCG community to seamlessly run and collect benchmark results (HEP Benchmark Suite) and to maintain the HEP reference applications needed for the benchmarking purposes (HEP Workloads).

This report will update on the recent software developments as well as on the feedback collected from early adopters of the whole toolkit.

In addition, benchmark results obtained running on tens of CPU models will be presented and will compare the HS06 and HEPscore_beta, a prototype of HEPscore based on LHC run2 workloads.

Desired slot length:

20

Speaker release:

Yes

Computing & Batch Services / 69

WLCG task force on HEP-SCORE deployment

Author: Helge Meinhard¹

¹ *CERN*

Corresponding Author: helge.meinhard@cern.ch

Status and plans of the task force

Desired slot length:

Speaker release:

Yes

Site Reports / 70**Diamond Light Source Site Report****Authors:** James Thorne¹; Frederik Ferner^{None}¹ *Diamond Light Source***Corresponding Authors:** frederik.ferner@diamond.ac.uk, j.i.thorne@gmail.com

Diamond Light Source is a Synchrotron Light Source based at the RAL site. This is a summary of what Diamond has been up to in cloud, storage and compute, as well as a few extras.

Desired slot length:**Speaker release:**

Yes

Site Reports / 71**IRFU Site Report****Authors:** Joel Surget¹; Pierre-Francois Honore²; Sophie Catherine Ferry¹¹ *Université Paris-Saclay (FR)*² *CEA/Université Paris-Saclay***Corresponding Authors:** joel.surget@cern.ch, pierre-francois.honore@cea.fr, sophie.catherine.ferry@cern.ch

We present an update of the changes at our site since the last report. Advancements, developments, roadblocks and achievements made concerning various aspects including: WLCG, Unix, Windows, Infrastructure, will be presented.

Desired slot length:

12min

Speaker release:

Yes

Site Reports / 72**BNL site report****Author:** Tejas Rao^{None}**Corresponding Author:** raot@bnl.gov

An update on BNL activities since the Spring 2021 workshop

Desired slot length:

12

Speaker release:

Yes

Site Reports / 73

KEK site report

Author: Tomoaki Nakamura^{None}

Co-authors: Go Iwai ¹; Koichi Murakami ; Soh Suzuki ; Takashi Sasaki ²

¹ *KEK*

² *High Energy Accelerator Research Organization (JP)*

Corresponding Authors: takashi.sasaki@kek.jp, go.iwai@kek.jp, koichi.murakami@kek.jp, tomoaki.nakamura@kek.jp, soh.suzuki@kek.jp

In this talk, we would like to report the recent update and status on the KEK central computing system, Grid services, and international network situation in Japan from the previous HEPiX workshop.

Desired slot length:

12 minutes

Speaker release:

Yes

Computing & Batch Services / 74

Hardware and batch monitoring at CNAF

Author: Federico Versari¹

¹ *INFN*

Corresponding Author: federico.versari@cnaf.infn.it

CNAF Tier-1, composed of almost 1000 worker nodes and nearly 40000 cores, completed its migration to HTCondor more than one year ago. After having adapted existing monitoring tools (built with Sensu, Influx and Grafana) to work with the new batch system, an effort has started to collect a more rich and “condor oriented” set of metrics that are used to provide better insights on the pool status.

Moreover we developed a similar tool with bare metal information collection, enabling sysadmins to have a global view of hardware (IPMI) events on the farm.

Desired slot length:

15

Speaker release:

Yes

Storage & File Systems / 75**Migration of Archival Tape/Library Technology at SDCC****Author:** Yingzi Wu^{None}**Co-authors:** Tim Chou¹; Christopher Henry Hollowell¹; Ognian Novakov²; Alexandr Zaytsev¹¹ *Brookhaven National Laboratory (US)*² *Brookhaven National Lab***Corresponding Authors:** iriswu@bnl.gov, alezayt@bnl.gov, tim.chou@cern.ch, hollowec@bnl.gov

This presentation reports on the ongoing migration of archival library and tape technology at The Scientific Data and Computing Center (SDCC). With the planned transition from Oracle libraries to IBM libraries, we deployed our first IBM TS4500 Library with ~20k tape slots in early 2021. This talk discusses our experience with the IBM TS4500, our continued transition to these libraries, as well as our experience with new tape technologies.

Desired slot length:

15-20minutes

Speaker release:

Yes

Birds of a Feather (BoF) session: Linux Discussion / 76**Linux OS Plans at The SDCC****Author:** Christopher Henry Hollowell¹¹ *Brookhaven National Laboratory (US)***Corresponding Author:** hollowec@bnl.gov

In December 2020, Red Hat and CentOS announced the early end-of-life of CentOS 8 (scheduled for December 2021), and its replacement with CentOS Stream 8. Unlike CentOS, which is a clone of RHEL, CentOS Stream is a forward-distribution of RHEL, containing package updates before they are released, or possibly ever included in RHEL. In this talk, we discuss the history and current status of Linux use at BNL's Scientific Data and Computer Center (SDCC), as well our future plans for this operating system.

Desired slot length:

15 minutes

Speaker release:

Yes

Site Reports / 77**PIC report****Author:** Jose Flix Molina¹**Co-authors:** Gonzalo Merino ²; Maria Del Carmen Porto Fernandez ³; Vanessa Acin ⁴; Esther Accion Garcia ⁵; Carlos Acosta Silva ⁶; Jordi Casals Hernandez ⁷; Ricard Cruz Josa ⁵; Elena Planas ⁶; Agustin Bruzzese ; Andreu Pacheco Pages ⁸; Jorge Carretero Palacios ⁷; Antonio Perez-Calero Yzquierdo ⁹; Jordi Delgado ⁶; Manuel Delfino Reznicek ⁵; Carlos Perez Dengra ⁶; Christian Neissner ⁶; Francesc Torradeflot ⁶; Pau Tallada Crespi ¹⁰; Elisabet Carrasco ⁷¹ *Centro de Investigaciones Energéticas Medioambientales y Tecnológicas (ES)*² *IFAE - Institute for High Energy Physics*³ *Unknown*⁴ *IFAE*⁵ *The Barcelona Institute of Science and Technology (BIST) (ES)*⁶ *PIC*⁷ *Port d'Informació Científica*⁸ *Institut de Física d'Altes Energies - Barcelona (ES)*⁹ *Centro de Investigaciones Energéticas Medioambientales y Tecnológicas*¹⁰ *PIC-CIEMAT***Corresponding Authors:** merino@pic.es, vacin@pic.es, eplanas@pic.es, rcruz@pic.es, cperez@pic.es, bruzzese@pic.es, esthera@pic.es, pacheco@ifae.es, jose.flix.molina@cern.ch, mporto@pic.org.es, neissner@pic.es, jordidem@pic.es, carlos.acosta.silva@cern.ch, delfino@pic.es, tallada@pic.es, antonio.perez.calero.yzquierdo@cern.ch, elisabet@pic.es, carretero@pic.es, torradeflot@pic.es, jcasals@pic.es

This is the PIC report for HEPiX Autumn 2021 Workshop

Desired slot length:**Speaker release:**

Yes

Network & Security / 78**WLCG Network Monitoring and Analytics Update****Authors:** Marian Babik¹; Shawn Mc Kee²¹ *CERN*² *University of Michigan (US)***Corresponding Authors:** marian.babik@cern.ch, shawn.mckee@cern.ch

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 primary areas to cover include the status of and plans for the WLCG/OSG perfSONAR infrastructure, the WLCG Throughput Working Group and the activities in the IRIS-HEP and the now completed SAND projects.

Desired slot length:

Speaker release:

Yes

Network & Security / 79

Research Networking Technical WG Update

Authors: Marian Babik¹; Shawn Mc Kee²

¹ *CERN*

² *University of Michigan (US)*

Corresponding Authors: marian.babik@cern.ch, shawn.mckee@cern.ch

As the scale and complexity of the current HEP network grows rapidly, new technologies and platforms are being introduced that greatly extend the capabilities of today's networks. 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 Research Networking Technical working group activities, challenges and recent updates.

In particular we'll focus on the packet marking technologies (scitags), tools and approaches that have been identified and we are going to discuss status of the implementation and plans for the near-term future.

Desired slot length:

15

Speaker release:

Yes

Storage & File Systems / 80

High-throughput EOS instance for ALICE O2 (online/offline)

Authors: Andreas Joachim Peters¹; Cristian Contescu¹

¹ *CERN*

Corresponding Authors: cristian.contescu@cern.ch, andreas.joachim.peters@cern.ch

In this contribution we are going to report on the latest results of the R&D activity aiming at preparing the EOS ALICE O2 storage cluster for the extremely demanding requirements of LHC Run 3.

Taking into consideration the latest upgrades of the LHC and of the ALICE detectors, the data throughput from the ALICE Data Acquisition system is expected to increase significantly, reaching 100GB/s data rate during Heavy-Ions collisions.

During this talk we are going to display the roadmap to meeting these demands with the EOS ALICE O2 storage instance, including software improvements, the storage nodes hardware setup, operating system tweaks and how we have overcome some of the hurdles in the process.

Desired slot length:

Speaker release:

Yes

Network & Security / 81

Update from the HEPiX IPv6 working group

Author: David Kelsey¹

¹ *Science and Technology Facilities Council STFC (GB)*

Corresponding Author: david.kelsey@stfc.ac.uk

During this year the HEPiX IPv6 working group has continued to encourage the deployment of dual-stack IPv4/IPv6 services. We also recommend dual-stack clients (worker nodes etc). Many data transfers are happening today over IPv6. This talk will present our recent work including the ongoing planning for moving to an IPv6-only core WLCG.

Desired slot length:

Speaker release:

Yes

Birds of a Feather (BoF) session: Linux Discussion / 82

CERN Linux Distro Strategy

Author: Ben Morrice¹

¹ *CERN*

Corresponding Author: benjamin.morrice@cern.ch

Birds of a Feather (BoF) session: Linux Discussion / 85

Discussion

Author: Frank Wurthwein¹

¹ UCSD

Corresponding Author: fkw@ucsd.edu

Grid, Cloud & Virtualisation / 88

Deployment of a Kubernetes-based Analysis Facility for ATLAS

Authors: David Jordan¹; Lincoln Bryant¹; Ilija Vukotic¹; Judith Lorraine Stephen¹; Robert William Gardner Jr¹

¹ University of Chicago (US)

Corresponding Authors: lincoln.bryant@cern.ch, robert.w.gardner@cern.ch, djordan66@uchicago.edu, ilija.vukotic@cern.ch, judith.lorraine.stephen@cern.ch

We present a technical report on the new ATLAS Analysis Facility hosted at the University of Chicago. Designed to support both traditional batch computing and novel analysis frameworks, this facility represents a shift in how future clusters will be deployed, federated, monitored and operated. We will describe the “cloud native” underpinnings of the facility (Kubernetes, GitOps, and Rook), how we leverage federated service deployment via SLATE, and how we incorporate technologies from IRIS-HEP such as ServiceX and Coffea Casa. The flexibility offered by these technologies allows the resource to be configured for the IRIS-HEP Scalable Systems Laboratory, providing a declarative platform for analysis and data delivery systems development.

Desired slot length:

Speaker release:

Yes

Site Reports / 89

IHEP Site Report

Author: Xuantong Zhang^{None}

Corresponding Author: zhangxuantong@ihep.ac.cn

Site news report from IHEP, including status of computing platform construction, grid, network, storage and so on, since last workshop report.

Desired slot length:

Speaker release:

Yes

Storage & File Systems / 90**EOS 5 new features and future plans****Author:** Abhishek Lekshmanan^{None}**Corresponding Author:** abhishek.lekshmanan@cern.ch

EOS is the open source distributed storage technology developed in the CERN IT Department and used at the Large Hadron Collider (LHC). EOS has been operated in production for more than 10 years and it now manages over half an exabyte of disk storage for both LHC & non LHC experiments. Since its first deployment in 2010, the software has evolved a lot catering to the large amounts of data and user requirements. After 4 major version releases, the next version, EOS 5, codename *diopside* is getting ready for deployment. This presentation will give an overview about EOS, the major features this release brings, a roadmap for expected future features and how all of these features contribute towards the next LHC Run.

Desired slot length:

15 minutes

Speaker release:

Yes

Network & Security / 91**Computer Security Update****Author:** Christos Arvanitis¹¹ CERN**Corresponding Author:** christos.arvanitis@cern.ch

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.

The COVID-19 pandemic has introduced a novel challenge for security teams everywhere by expanding the attack surface to include everyone's personal devices / home networks and causing a shift to new, risky software for a remote-first working environment. It was also a chance for attackers to get creative by taking advantage of the fear and confusion to devise new tactics and techniques.

What's more, the worrying trend of data leaks, password dumps, ransomware attacks and new security vulnerabilities does not seem to slow down.

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

Desired slot length:

20

Speaker release:

Yes

Site Reports / 92**RAL Site Report****Author:** Martin Bly¹¹ *STFC-RAL***Corresponding Author:** martin.bly@stfc.ac.uk

Update of news from RAL

Desired slot length:**Speaker release:**

Yes

Network & Security / 93**Computer Security Update (repeat of 11:00h contrib)****Corresponding Author:** christos.arvanitis@cern.ch**Site Reports / 94****CERN Site Report****Author:** Andrei Dumitru¹¹ *CERN***Corresponding Author:** andrei.dumitru@cern.ch

News from CERN since the last HEPiX workshop.

Desired slot length:

20 (17 talk + 3 for q&a)

Speaker release:

Yes

Computing & Batch Services / 95**Report from the HTCondor Workshop September 2021****Author:** Helge Meinhard¹¹ *CERN*

Corresponding Author: helge.meinhard@cern.ch

A summary of the annual “European” HTCondor workshop recently held on-line

Desired slot length:

Speaker release:

Yes

Basic IT Services / 96

Automated Backup and Recovery for the Database on Demand service using EOS

Author: Anirudh Goel¹

¹ CERN

Corresponding Author: anirudh.goel@cern.ch

Last year we deployed a new automated backup and recovery setup for the Database on Demand Service at CERN. This setup comprises of 2 services-

1. Backup service - Stores encrypted and zipped backups of all DBOD production instances to EOS storage.
2. Restore service - Continuously restores and tests connectivity of snapshots of all production instances.

This presentation aims to shed light on the structure and working of this setup.

Desired slot length:

Speaker release:

Yes

Storage & File Systems / 97

Migration from CASTOR to the CERN Tape Archive

Authors: Andreas Joachim Peters¹; Edward Karavakis¹; Giuseppe Lo Presti¹; Julien Leduc¹; Michael Davis¹; Cristian Contescu¹; Elvin Alin Sindrilaru¹; Maria Arsuaga Rios¹; Michal Kamil Simon¹; Mihai Patrascoiu¹; Oliver Keeble¹; Steven Murray¹; Vladimir Bahyl¹; Volodymyr Yurchenko¹

¹ CERN

Corresponding Authors: andreas.joachim.peters@cern.ch, maria.arsuaga.rios@cern.ch, julien.leduc@cern.ch, michael.davis@cern.ch, vladimir.bahyl@cern.ch, volodymyr.yurchenko@cern.ch, edward.karavakis@cern.ch, cristian.contescu@cern.ch, oliver.keeble@cern.ch, elvin.alin.sindrilaru@cern.ch, steven.murray@cern.ch, michal.simon@cern.ch, giuseppe.lopresti@cern.ch, mihai.patrascoiu@cern.ch

CASTOR was used as CERN’s primary archival storage system for the last two decades, including Run-1 and Run-2 of the LHC. For Run-3, CASTOR has been replaced by the CERN Tape Archive (CTA). At the end of Run-2, there were 340 Petabytes of data stored in CASTOR, which had to

be migrated to CTA during Long Shutdown 2. Over 90% of this data is an active archive —the custodial copy of physics data belonging to the four LHC experiments and around a dozen smaller experiments at CERN. The migration and switch from CASTOR to CTA had to be accomplished with minimal interruption to experiment activities; to further complicate the problem, each experiment has a slightly different workflow and data management stack. This presentation will describe our experiences and lessons learned during the two-year period of the migration.

Desired slot length:

20

Speaker release:

Yes

Grid, Cloud & Virtualisation / 98**Helmholtz IT Services for Science: HIFIS****Author:** Uwe Jandt¹¹ DESY**Corresponding Author:** uwe.jandt@desy.de

The Helmholtz-based platform HIFIS builds and sustains an IT infrastructure connecting all Helmholtz research fields and centres.

The services provided by HIFIS include a secure and easy-to-use collaborative environment with efficiently accessible IT services from anywhere. HIFIS further supports Research Software Engineering (RSE) with a high level of quality, visibility and sustainability.

In this talk, we present the scope, current implementation status, as well as exemplary services of HIFIS. The interplay of the multiple modular and decentralized HIFIS components, including the Helmholtz Cloud Portal, the Authentication and Authorization Infrastructure (Helmholtz AAI), intermediary components such as the Helmholtz Cloud Agent and cloud services provided at various sites will be showcased.

Desired slot length:**Speaker release:**

Yes

Miscellaneous / 100**Analysis on distributed computing at the Belle II experiment****Author:** Michel Hernandez Villanueva¹¹ DESY**Corresponding Author:** michel.hernandez.villanueva@desy.de

The Belle II experiment is a detector coupled to the SuperKEKB electron-positron collider, designed to collect 50 times the data produced by the previous generation of B-factories. Process and analyzing

these high volumes of data in a timely fashion requires an efficient interconnection of software and computing resources. In this talk, the analysis model on the Belle II experiment is presented, summarizing the distributed computing technologies adopted and current challenges.

Desired slot length:

Speaker release:

Yes

Miscellaneous / 101

Group Photo

Miscellaneous / 102

Group Photo

Basic IT Services / 103

Migrating from a commercial system to Open Source: Gitlab at DESY

Author: Ingo Ebel¹

Co-authors: Johannes Reppin¹; Elena Gapon¹

¹ *Deutsches Elektronen-Synchrotron DESY*

Corresponding Authors: johannes.reppin@desy.de, ingo.ebel@desy.de, elena.gapon@desy.de

As ATlassian is changing its product offering, DESY has started to evaluate alternatives to the products from the firm's product portfolio.

For the field of repositories and CI/CD software development, a suitable replacement for Jira Software/Bitbucket/Bamboo was found in *gitlab*.

This talk will outline the pilot project, the current production system and the offering we make to users to help them during migration.

Desired slot length:

Speaker release:

Yes

Grid, Cloud & Virtualisation / 104

Survival kit for CERN's organic webscape: Kubernetes and a community of Makers

Author: Konstantinos Samaras-Tsakiris¹

Co-authors: Catharine Noble ¹; Vineet Reddy Rajula ¹; Francisco Borges Aurindo Barros ²

¹ CERN

² Universidade de Lisboa (PT)

Corresponding Authors: francisco.borges.aurindo.barros@cern.ch, rajula.vineet.reddy@cern.ch, catharine.noble@cern.ch, konstantinos.samaras-tsakiris@cern.ch

How can we turn a “chore” into a community of makers, in a scientific Organization that doesn't enforce strict standards? Drupal has a long history at CERN. Site builders of >1k unique websites take advantage of a highly automated, but bespoke infrastructure, that automates website operations: provision, backup, clone, update, delete. Relying on direct support through tickets worked, until Drupal 7 migrations overloaded the support lines, while engineers constantly come and go, taking know-how with them. Technical debt started weighing us down.

Does this sound familiar?

To overcome these service challenges, we came up with a Technical and a Human solution. We built a new, cloud-native infrastructure on Kubernetes and a standard “CERN Drupal Distribution”, fully open source and available. Onboarding newcomers becomes simpler thanks to the standard technologies. Patiently we also grew a Drupal Community at CERN that brought site builders together to absorb knowledge in a social network and rely less on the singular expert.

At HEPiX we'll share with you our journey through these problems and our responses, describing them in details. We hope to give you both a reference frame and concrete solutions.

Desired slot length:

20

Speaker release:

Yes

Basic IT Services / 105

From On-Premise to On-Cloud videoconference service

Author: Ruben Domingo Gaspar Aparicio¹

¹ CERN

Corresponding Author: ruben.gaspar.aparicio@cern.ch

A retrospective view of the transition from Vidyo to Zoom platform.

Desired slot length:

20

Speaker release:

Yes

Basic IT Services / 106**Running an AV media post-processing system on FOSS**

Authors: Miguel Angel Valero Navarro¹; Ruben Domingo Gaspar Aparicio²; Rene Fernandez Sanchez²

¹ *Valencia Polytechnic University (ES)*

² *CERN*

Corresponding Authors: rene.fernandez@cern.ch, ruben.gaspar.aparicio@cern.ch, miguel.angel.valero.navarro@cern.ch

The Weblecture service is in charge of capturing, processing & delivery CERN productions: e-learning, Computing Seminars, Conferences, Outreach events, CERN related communication e.g. DG, HR, Staff association, etc.. Tightly linked to Webcast and Videoconference services from which most AV content is provided. The services delivers AV content on different formats so our community can watch them in different platforms e.g. mobile or a desktop device.

The talk presents the past status running on a legacy stack and the work done to provide same and expanded functionality: new capture agents like Epiphany with improved capabilities, better and faster transcoding which is also provided as a service to other units, unified video player for both conference and web-lectures, keeping backward compatibility, new CERN SSO, etc.

The new architecture has been built around an in-house development and integrations of two FOSS projects Paella player 1 and Opencast 2.

1 <https://paellaplayer.upv.es/>

2 <https://opencast.org/>

Desired slot length:

Speaker release:

Yes

Network & Security / 107**LoRaWAN at CERN vs pandemics: 1:0**

Author: Christoph Merscher¹

¹ *CERN*

Corresponding Author: christoph.merscher@cern.ch

To counteract the spread of the COVID-19 virus as much as possible, a device has been developed at CERN, the Proximenter, which sends information about contact tracing via an IoT network. In this respect, some hurdles had to be overcome in terms of data protection and compromises concerning the network and the actual protocol.

Desired slot length:

Speaker release:

Yes

Network & Security / 108**Threat Intelligence and Security Operations Centres: Collaborative Security**

Authors: David Crooks¹; Liviu Valsan²; Romain Wartel²

¹ UKRI STFC

² CERN

Corresponding Authors: liviu.valsan@cern.ch, david.crooks@cern.ch, romain.wartel@cern.ch

The threat faced by the research and education sector from determined and well-resourced attackers has been growing in recent years and is now acute. We must act together as a community to defend against these attacks. A vital means of achieving this is to share threat intelligence - key indicators of compromise of an ongoing incident including network locations and file hashes - with trusted partners. We must couple this with a robust, fine-grained source of network monitoring. The combination of these elements along with storage, visualisation and alerting is called a Security Operations Centre. The WLCG SOC working group has been pursuing an interconnected network of SOC-equipped sites for several years. We report here on recent progress, including new deployments against multiple 100Gb/s sites, and future plans for the coming year.

Desired slot length:

Speaker release:

Yes

Network & Security / 109**Resource Trust Evolution: Host Certificates in a Dynamic Landscape**

Authors: David Crooks¹; David Kelsey²; Maarten Litmaath³; Romain Wartel³

¹ UKRI STFC

² Science and Technology Facilities Council STFC (GB)

³ CERN

Corresponding Authors: romain.wartel@cern.ch, david.kelsey@stfc.ac.uk, david.crooks@cern.ch, maarten.litmaath@cern.ch

Increasing use of cloud resources, and other developments in new workflows, have posed an important question on which certificate providers are most appropriate for different use cases. Certificate authorities under discussion include Let's Encrypt but also the CAs of commercial cloud providers. We discuss the posing of this question along with the key stakeholders, including sites, operations, experiments, and identity management and security experts, as well as requirements that need to be included in this discussion. We report on the formation of new WLCG task force on this topic as we approach the first meeting of that group.

Desired slot length:

Speaker release:

Yes

Basic IT Services / 110**DESY's user registry**

Author: Dirk Jahnke-Zumbusch¹

¹ *DESY*

Corresponding Author: dirk.jahnke-zumbusch@desy.de

DESY's current user registry is about to retire. During its use boundary conditions have shifted, the user groups have become more heterogeneous and in addition to collaborations with stable memberships the photon science community brought a higher rate of fluctuation of IT users. The successor incorporates those changes and delivers functionality based on Oracle's database (database, APEX, ORDS).

Desired slot length:

Speaker release:

Yes

Storage & File Systems / 111**powerful, cross-disciplinary storage for everyone**

Author: Paul Millar^{None}

Co-authors: Patrick Fuhrmann ; Christian Voss ¹

¹ *DESY-HH*

Corresponding Authors: paul.millar@desy.de, christian.voss@desy.de, patrick.fuhrmann@desy.de

While WLCG may be considered at the vanguard of data-intense scientific research, many other scientific communities are finding their data storage requirements growing beyond their current capabilities. Simultaneously, with science increasingly involving broad collaborations, the ability to support and manage scientists from different institutes is becoming essential.

At DESY, we have developed and deployed a community storage solution. The primary users are the Helmholtz scientific community, which includes a broad spectrum of fundamental research activity. The service is not restricted to Helmholtz users; for example, through the ESCAPE project, we also support other communities such as SKA, CTA, LSST and FAIR.

In this talk will detail the primary use-cases we wish to address with this service, the powerful features that support them, and the technologies that underpin those features. We will also describe the current status of the storage and our plans for the future.

Desired slot length:

Speaker release:

Yes

Miscellaneous / 112

HEPiX closing remarks

Authors: Peter van der Reest^{None}; Tony Wong^{None}

Corresponding Authors: tony@bnl.gov, peter.van.der.reest@desy.de

Birds of a Feather (BoF) session: Linux Discussion / 113

DESY Linux Status & Strategy

Corresponding Author: yves.kemp@desy.de

Speaker release:

Desired slot length:

Birds of a Feather (BoF) session: Linux Discussion / 114

Red Hat's RHEL Academic Model

Author: Brian Exelbierd¹

¹ *Red Hat*

Corresponding Author: bexelbie@redhat.com

In conjunction with our proposal to CERN and the broader HEPiX community, I wanted to share the basis for the proposal and provide an opportunity for questions to be asked.

Desired slot length:

20

Speaker release:

Yes

Storage & File Systems / 115

The RX Extended SACK Table Proposal - AuriStor RX Prototype Results

Author: Jeffrey Altman^{None}

Corresponding Author: jaltman@auristor.com

The RX protocol inherited from IBM AFS is incapable of filling a network pipe with a single RPC when the pipe's bandwidth delay product exceeds 44 1/4 KB. On a 1 Gbit/sec pipe with a 1ms RTT, the maximum theoretical throughput is 360 Mbit/sec with a maximum window size of 44 KB. The RX ACK packet format provides for a theoretical maximum window of 65535 packets but the Selective Acknowledgment (SACK) Table is limited to 255 packets. with 255 packets the maximum window size is 351 KB or maximum throughput of 2.875 Gbit/sec with 1ms RTT. Increase the RTT to 8ms and the maximum throughput is once again reduced to 360 Mbit/sec. The RTT on cross-Atlantic commodity internet pipes often exceed 110ms which reduces the theoretical throughput to 28 MBit/sec.

This presentation will describe AuriStor's proposed RX Extended SACK Table protocol extension, prior efforts at extending the maximum window size, and preliminary results on real-world networks using AuriStor's prototype supporting maximum windows up to 8192 packets or 11MB.

Desired slot length:

20 minutes

Speaker release:

Yes

Miscellaneous / 116

Welcome & Logistics

Author: Peter van der Reest^{None}

Corresponding Author: peter.van.der.reest@desy.de

Miscellaneous / 117

Welcome & Logistics

Author: Peter van der Reest^{None}

Corresponding Author: peter.van.der.reest@desy.de

Network & Security / 118

Discussion on user security awareness

Author: Jean-Michel Barbet¹

¹ *Centre National de la Recherche Scientifique (FR)*

Corresponding Author: jean-michel.barbet@subatech.in2p3.fr

On account of a question after the subject came up in site reports: what do sites have planned in terms of measures to increase users security awareness

Birds of a Feather (BoF) session: Linux Discussion / 119

Linux Distribution BoF (cont'd)

Input and Views from the experiments

Birds of a Feather (BoF) session: Linux Discussion / 120

Alice

Author: Maarten Litmaath¹

¹ *CERN*

Corresponding Author: maarten.litmaath@cern.ch

Birds of a Feather (BoF) session: Linux Discussion / 121

ATLAS

Authors: David Cameron¹; David South²

¹ *University of Oslo (NO)*

² *Deutsches Elektronen-Synchrotron (DE)*

Corresponding Authors: david.south@cern.ch, david.cameron@cern.ch

Birds of a Feather (BoF) session: Linux Discussion / 122

CMS

Birds of a Feather (BoF) session: Linux Discussion / 123

LHCb

Author: Concezio Bozzi¹

¹ *INFN Ferrara*

Corresponding Author: concezio.bozzi@cern.ch

Birds of a Feather (BoF) session: Linux Discussion / 124

Belle II

Author: Takanori Hara¹

¹ *High Energy Accelerator Research Organization (JP)*

Corresponding Author: takanori.hara@kek.jp

Birds of a Feather (BoF) session: Linux Discussion / 125

EuXFEL

Author: Yves Kemp^{None}

Corresponding Author: yves.kemp@desy.de

Birds of a Feather (BoF) session: Linux Discussion / 126

General discussion - if needs be into coffee break time