

HEPiX Spring 2025 Workshop

Monday 31 March 2025 - Friday 4 April 2025

Hotel De La Paix



Book of Abstracts

Contents

Summary of the 2024 autumns european HTC workshop	1
Keeping the LHC colliding: Providing Extended Lifecycle support for EL7	1
Improving CERN's security with an Endpoint Detection and Response Solution	1
Managing Microsoft SQL Infrastrucuture at CERN	2
Design and production experience with a multi-petabyte file-system backup service at CERN	2
How CERN's New Datacenter Enhances Cloud Infrastructure and Data Resilience with Ceph	3
A Distributed Storage Odyssey: from CentOS7 to ALMA9	3
EOS latest developments and operational experience	4
CERN site report	4
Refurbishing the Meyrin Data Centre: Storage Juggling and Operations	5
Single Sign-On Evolution at CERN	5
Continuous calibration and monitoring of WLCG site corepower with HEPsScore23	6
Grafana dashboards as code with Jsonnet	6
The CERN Tape Archive Beyond CERN	7
Evolution of Continuous Integration for the CERN Tape Archive (CTA)	8
Natural job drainage and power reduction studies in PIC Tier-1 using HTCondor	8
CERN Preveessin Datacentre network - Overview and feedback after one year in production.	9
CERN update on tape technology	9
HEPiX Benchmarking Working Group Report	10
Roadmap to LS3: CERN's Linux Strategy	10
Transforming the Disaster Recovery of the Cloud Service	11
Windows device management at CERN: A new era	11

NDGF Site Report	12
Wrap-up	12
Welcome talk	13
Logistics talk	13
Network design and implementation status of HEPS	13
Machine learning for developers and administrators	13
HEPiX Technology Watch Working Group Report	14
IDAF @ DESY: Interdisciplinary Data and Analysis Facility: Status and Plans	14
Updates on CPUs, GPUs and AI accelerators	15
Hyper-converged cloud infrastructure at CSCS	15
Trends of Energy Efficiency for computing and data centre	16
RAL use case of XRootD Managers	16
Update on Energy Efficiency: AmpereOne and Turin	17
NCG-INGRID-PT site report	17
Smart HPC-QC: flexible approaches for Quantum workloads integration	17
Update from the HEPiX IPv6 Working Group	18
IHEP site report	18
Introduction of LHCb Tier2 Site at Lanzhou University	19
Kubernetes and Cloud Native at the SKA Regional Centres	19
Efficiency of job processing in many-core Grid and HPC environments	20

Computing and Batch Services / 1**Summary of the 2024 autumns european HTC workshop****Author:** Christoph Beyer^{None}**Corresponding Author:** christoph.beyer@desy.de

The tenth european HTCCondor workshop took place at NIKHEF Amsterdam autumn last year and as always covered most if not all aspects of up-to-date high throughput computing.

Here comes a short summary of the parts of general interest if you like :)

Desired slot length:

20 minutes

Speaker release:

Yes

Cloud Technologies, Virtualization & Orchestration, Operating Systems / 4**Keeping the LHC colliding: Providing Extended Lifecycle support for EL7****Author:** Ben Morrice¹¹ CERN**Corresponding Author:** ben.morrice@cern.ch

The operation of the Large Hadron Collider (LHC) is critically dependent on several hundred Front-End Computers (FECs), that manage all facets of its internals. These custom systems were not able to be upgraded during the long shutdown (LS2), and with the coinciding end-of-life of EL7 of 30.06.2024, this posed a significant challenge to the successful operation of Run 3.

This presentation will focus on how CERN IT is providing the Red Hat “Extended Lifecycle Support” (ELS) product across the CERN accelerator sector. We will discuss how this solution ensures operational continuity by maintaining software support for legacy hardware, bridging the gap between aging infrastructure and current security requirements. Technical details on how this is achieved, as well as shortcomings and lessons learned will be shared with the audience.

Desired slot length:**Speaker release:**

Yes

Network & Security / 5**Improving CERN’s security with an Endpoint Detection and Response Solution**

Author: Alexandros Petridis^{None}

Co-author: Ricardo Martins Goncalves

Corresponding Authors: alexandros.petridis@cern.ch, ricardo.martins.goncalves@cern.ch

The deployment of an Endpoint Detection & Response (EDR) solution at CERN has been a project aimed at enhancing the security posture of endpoint devices. In this presentation we'll share our infrastructure's architecture and how we rolled out the solution. We will also see how we addressed and overcome challenges on multiple fronts from administrator's fears to fine-tuning detections and balancing performance with security requirements. Finally, we will showcase the capabilities delivered by an EDR solution including real-time threat detection, incident response, and enhanced visibility across endpoints.

Desired slot length:

15

Speaker release:

Yes

Miscellaneous / 6

Managing Microsoft SQL Infrastructure at CERN

Author: Ricardo Martins Goncalves^{None}

Co-author: Mario Rey Regulez¹

¹ CERN

Corresponding Authors: mario.rey@cern.ch, ricardo.martins.goncalves@cern.ch

For its operations, CERN depends on an extensive range of applications, achievable only through the use of diverse technologies, including more than one relational database management system (RDBMS). This presentation provides an overview of CERN's Microsoft SQL Server (MSSQL) infrastructure, highlighting how we manage servers and design solutions for large-scale databases with different compatibilities and criticality. We will explore the advantages and challenges of using MSSQL and how we provide high availability for critical databases, tailored solutions for applications, and different environments for production and testing. Furthermore, we will discuss our backup infrastructure and data integrity tests to ensure disaster recovery.

Desired slot length:

15

Speaker release:

Yes

Storage & data management / 7

Design and production experience with a multi-petabyte file-system backup service at CERN

Author: Roberto Valverde Cameselle¹

¹ CERN**Corresponding Author:** roberto.valverde.cameselle@cern.ch

We report on our experience with the production backup orchestration via “cback”, a tool developed at CERN and used to back up our primary mounted filesystem offerings: EOS (eosxd) and Ceph (CephFS). In a storage system that handles non-reproducible data, a robust backup and restore system is essential for effective disaster recovery and business continuity. When designing a backup solution, it is crucial to consider the same factors that apply to the production service itself: scalability, performance, security, and operational costs. In this contribution, we will discuss the challenges we encountered, the decisions we made, and the innovative strategies we implemented while designing cback. Many of these insights can be applied to other backup strategies as well.

Desired slot length:

15-20

Speaker release:

Yes

Storage & data management / 8

How CERN’s New Datacenter Enhances Cloud Infrastructure and Data Resilience with Ceph

Author: Roberto Valverde Cameselle¹¹ CERN**Corresponding Author:** roberto.valverde.cameselle@cern.ch

The storage needs of CERN’s OpenStack cloud infrastructure are fulfilled by Ceph, which provides diverse storage solutions including volumes with Ceph RBD, file sharing through CephFS, and S3 object storage via Ceph RadosGW. The integration between storage and compute resources is possible thanks to a close collaboration between OpenStack and Ceph teams. In this talk we review the architecture of our production deployment and how it evolved with the arrival of the new datacenter in Preveessin (France) in the context of supporting BC/DR scenarios

Desired slot length:

15

Speaker release:

Yes

Storage & data management / 9

A Distributed Storage Odyssey: from CentOS7 to ALMA9

Author: Cedric Caffy¹¹ CERN

Corresponding Author: cedric.caffy@cern.ch

On the 30th of June 2024, the end of CentOS 7 support marked a new era for the operation of the multi-petabytes distributed disk storage system used by CERN physics experiments. The EOS infrastructure at CERN is composed of approximately 1000 disk servers and 50 metadata management nodes. Their transition from CentOS 7 to Alma 9 was not as straightforward as anticipated.

This presentation will be all about explaining this transition. From the change of supported certificate and kerberos key signature lengths and algorithms, to openssl library hiccups and Linux kernel crashes, the EOS operation team had to take on different challenges to ensure a seamless operating system transition of the infrastructure while maintaining uninterrupted CERN experiments' data transfers.

Desired slot length:

Speaker release:

Yes

Storage & data management / 10

EOS latest developments and operational experience

Author: Elvin Alin Sindrilaru¹

¹ CERN

Corresponding Author: elvin.alin.sindrilaru@cern.ch

EOS is an open-source storage system developed at CERN that is used as the main platform to store LHC data. The architecture of the EOS system has evolved over the years to accommodate ever more diverse use-cases and performance requirements coming both from the LHC experiments as well as from the user community running their analysis workflows on top of EOS. In this presentation, we discuss the performance of the EOS service during the 2025 Run and also outline the latest developments targeting diverse areas like file system consistency checks improvements, deployment of a high-availability setup for the metadata service, namespace locking and performance optimizations as well as other topics. Apart from new developments, we also discuss changes in the deployment model, especially moving to a native HTTP approach and the commissioning of the GRPC interface used by the CernBox service. To conclude the presentation, we outline some of the key achievements from 2025 that make us confident our system is fully prepared for the challenges that wait ahead as we approach the end of Run 3 and the preparation for High Luminosity LHC.

Desired slot length:

20

Speaker release:

Yes

Site Reports / 11

CERN site report

Author: Elvin Alin Sindrilaru¹

¹ CERN**Corresponding Author:** elvin.alin.sindrilaru@cern.ch

News from CERN since the last HEPiX workshop. This talk gives a general update from services in the CERN IT department.

Desired slot length:**Speaker release:**

Yes

Storage & data management / 12**Refurbishing the Meyrin Data Centre: Storage Juggling and Operations****Author:** Octavian-Mihai Matei^{None}**Corresponding Author:** octavian-mihai.matei@cern.ch

The 50-year-old Meyrin Data Centre (MDC), still remains indispensable due to its strategic geographical location and unique electrical power resilience even if CERN IT recently commissioned the Prévessin Data Centre (PDC), doubling the organization's hosting capacity in terms of electricity and cooling. The Meyrin Data Centre (Building 513) retains an essential role for the CERN Tier-0 Run 4 commitments, notably as primary hosting location for the tape archive and the disk storage. The inevitable investments to the infrastructure (UPS and Cooling) are now triggering the refurbishment of the two main rooms where all the storage equipment is hosted. This presentation will delve into the architectural advancements and operational strategies implemented for and during the Meyrin data centre refurbishment. We will explore how these developments will impact our storage and how the storage operations team will ensure EOS's performance, scalability, and reliability in the coming years.

Desired slot length:

15

Speaker release:

Yes

Network & Security / 13**Single Sign-On Evolution at CERN****Authors:** Hannah Short¹; Paul Van Uytvinck¹; Sebastian Lopienski¹¹ CERN**Corresponding Authors:** paul.van.uytvinck@cern.ch, hannah.short@cern.ch, sebastian.lopienski@cern.ch

The Single Sign-On (SSO) service at CERN has undergone a significant evolution over recent years, transitioning from a Puppet-hosted solution to a Kubernetes-based infrastructure. Since September 2023, the current team has focused on cementing SSO as a stable and reliable cornerstone of CERN's

IT services. Effort was concentrated on implementing best practices in service management - a mid term investment that is already proving worthwhile.

This presentation highlights the strides made in consolidating and modernizing the SSO service. Key achievements include the successful migration from Keycloak 20 to Keycloak 24 and significant improvements in monitoring using Grafana, disaster recovery preparation, and proactive alerting through Telegram and Mattermost.

We also showcase the advantages of Keycloak as a central identity management solution for CERN. Keycloak's extensibility lies in its ability to support custom development through Java-based Service Provider Interfaces (SPIs) to meet specific organizational needs. By implementing these SPIs, the team was able to bridge the gap between modern identity protocols and CERN's diverse legacy systems.

Furthermore, the team has implemented proactive configuration control measures, such as exporting Keycloak realm configurations to GitLab, enabling transparency and traceability for changes made to the SSO configuration.

Desired slot length:

Speaker release:

Yes

Computing and Batch Services / 14

Continuous calibration and monitoring of WLCG site corepower with HEPScore23

Authors: Natalia Diana Szczepanek¹; Domenico Giordano¹

¹ CERN

Corresponding Authors: domenico.giordano@cern.ch, natalia.diana.szczepanek@cern.ch

The performance score per CPU core —corepower— reported annually by WLCG sites is a critical metric for ensuring reliable accounting, transparency, trust, and efficient resource utilization across experiment sites. It is therefore essential to compare the published CPU corepower with the actual runtime corepower observed in production environments. Traditionally, sites have reported annual performance values based on weighted averages of various CPU models, yet until now there was no direct method to validate these figures or to easily retrieve the underlying CPU model weights from each site.

With the official adoption of HEPScore23 as a benchmark in April 2023 by the WLCG, the Benchmarking Working Group introduced new tools, including the HEP Benchmark Suite with Plugins, to address this gap. The new infrastructure is able to continuously monitor and validate the reported performance values by running benchmarks across the grid. This approach ensures the accuracy of annual performance figures, promotes transparency, and enables the timely detection and correction of incorrect values with minimal effort from the sites.

Desired slot length:

20 minutes

Speaker release:

Yes

Software and Services for Operation / 15**Grafana dashboards as code with Jsonnet****Author:** Ewoud Ketele¹¹ CERN**Corresponding Author:** ewoud.rob.l.ketele@cern.ch

Grafana dashboards are easy to make but hard to maintain. Since changes can be made easily, the questions that remain are how to avoid changes that overwrite other work? How to keep track of changes? And how to communicate these to the user? Another question that pops up frequently is how to apply certain changes consistently to multiple visualizations and dashboards. One partial solution is to export Grafana dashboards to their JSON representation and store those in a git repository. However, even simple dashboards can quickly run into the thousands of lines of JSON and version-controlling these is problematic in its own right: the diffs are large and changes do not easily carry over to JSON representations of other dashboards. Instead, we propose to use the Jsonnet configuration language, together with its library Grafonnet, to create so called 'dashboard definitions' that compile into JSON representations of dashboards. We have adopted this solution to manage multiple different dashboards with around 20 visualizations each. This leads to a clear improvement in the maintainability and deployment of these dashboards. In this contribution, we will show how we created complex dashboards with multiple types of visualizations and data sources using functional code and Jsonnet. Additionally, we will show how dashboards-as-code can be integrated in git repositories and their CI/CD pipelines to ensure consistency between the dashboard definitions and the dashboards in a Grafana instance.

Desired slot length:

20

Speaker release:

Yes

Storage & data management / 16**The CERN Tape Archive Beyond CERN****Author:** Niels Alexander Buegel^{None}**Co-author:** Michael Davis¹¹ CERN**Corresponding Authors:** niels.alexander.buegel@cern.ch, michael.davis@cern.ch

The CERN Tape Archive (CTA) is CERN's Free and Open Source Software system for data archival to tape. Across the Worldwide LHC Computing Grid (WLCG), the tape software landscape is quite heterogeneous, but we are entering a period of consolidation. A number of sites have reevaluated their options and have chosen CTA for their tape archival storage needs. To facilitate this, the CTA team have added a number of community features, allowing CTA to be used as the tape backend for dCache and to facilitate migrations from other tape systems such as OSM and Enstore. CTA is now packaged and distributed as a public release, free from CERN-specific dependencies, together with a set of operations tools. This contribution presents the latest CTA community features and roadmap.

Desired slot length:

20

Speaker release:

Yes

Storage & data management / 17**Evolution of Continuous Integration for the CERN Tape Archive (CTA)****Author:** Niels Alexander Buegel^{None}**Corresponding Author:** niels.alexander.buegel@cern.ch

The CERN Tape Archive (CTA) software is used for physics archival at CERN and other scientific institutes. CTA's Continuous Integration (CI) system has been around since the inception of the project, but over time several limitations have become apparent. The migration from CERN CentOS 7 to Alma 9 introduced even more challenges. The CTA team took this as an opportunity to make significant improvements in the areas of simplicity, flexibility and robustness. The most impactful change was the migration from plain Kubernetes manifest files to Helm, allowing us to decouple the configuration of CTA from the EOS disk system configuration and opening up opportunities to test other disk buffer systems such as dCache. The new setup allows us to handle complex testing scenarios and perform regression testing on various components independently. We will discuss the challenges we encountered with our CI, the improvements we implemented to address them, and what we hope to do in the future.

Desired slot length:

20

Speaker release:

Yes

Environmental sustainability, business continuity, and Facility improvement / 18**Natural job drainage and power reduction studies in PIC Tier-1 using HTCondor****Author:** Jose Flix Molina¹**Co-authors:** Carles Acosta Silva ²; Gonzalo Merino ³; Jordi Casals Hernandez ⁴; Kevin Fabrega ⁵¹ CIEMAT - Centro de Investigaciones Energéticas Medioambientales y Tec. (ES)² PIC³ IFAE - Institute for High Energy Physics⁴ Port d'Informació Científica (PIC)⁵ UAB**Corresponding Authors:** jcasals@pic.es, cacosta@pic.es, kevin.fabrega@autonoma.cat, merino@pic.es, jose.flix.molina@cern.ch

This study presents analyses of natural job drainage and power reduction patterns in the PIC Tier-1 data center, which uses HTCondor for workload scheduling. By examining historical HTCondor logs from 2023 and 2024, we simulate natural job drainage behaviors, in order to understand natural job drainage patterns: when jobs naturally conclude without external intervention. These findings

provide insights into the center's capability to modulate resource usage according to external factors like green energy availability cycles.

To further validate and extend these observations, simulations were conducted under various load conditions to evaluate the influence of job types and VO-specific durations on drainage cycles. An analysis of power consumption pre- and post-drainage, facilitated by ipmitool, allows for estimating potential power and carbon emission reductions in drainage scenarios. Building on these insights, machine learning models are being developed to predict optimal power scaling adjustments.

We propose a conceptual feedback loop to HTCCondor that could enable real-time power adjustments based on fluctuations in green energy availability. By exploring these ideas, this research aims to contribute to a more sustainable data center model, offering a framework for adapting workload management to dynamic environmental factors.

Desired slot length:

Speaker release:

Yes

Network & Security / 19

CERN Preveessin Datacentre network - Overview and feedback after one year in production.

Author: Vincent Ducret¹

¹ CERN

Corresponding Author: vincent.ducret@cern.ch

This presentation will explain the network design implemented in the CERN Préveessin Datacentre (built in 2022/2023, in production since February 2024). We will show how, starting from an empty building, the current network best practices could be adopted (and partly adapted to match the specific requirements in term of interconnection with the rest of CERN network). We will also provide feedback about some issues encountered during the planning and deployment, give an overview of the network performances after one full year in production, and share with you our current ideas (and questions) regarding CERN Datacentre(s) network possible evolution in the coming years.

Desired slot length:

20 minutes

Speaker release:

Yes

Storage & data management / 20

CERN update on tape technology

Author: Vladimir Bahyl¹

¹ CERN

Corresponding Author: vladimir.bahyl@cern.ch

This presentation will start with the evolution of the tape technology market in the recent years and the expectations from the INSIC roadmap.

From there, with LHC now in the middle of Run 3, we will reflect on the evolution of our capacity planning vs. increasing storage requirements of the experiments. We will then describe our current tape hardware setup and present our experience with the different components of the technology. For example, we will report on performance characteristics of both LTO9 and TS1170 tape drives: RAO, environmental aspects and how the technology evolution is impacting our operations.

Lastly, we will share our thoughts about rack size scale-out tape libraries and consideration to replace FC with SAS.

Desired slot length:

20 minutes

Speaker release:

Yes

Computing and Batch Services / 21

HEPiX Benchmarking Working Group Report

Authors: Domenico Giordano¹; Michele Michelotto²

¹ CERN

² Università e INFN, Padova (IT)

Corresponding Authors: michele.michelotto@cern.ch, domenico.giordano@cern.ch

The Benchmarking Working Group (WG) has been actively advancing the HEP Benchmark Suite to meet the evolving needs of the Worldwide LHC Computing Grid (WLCG). This presentation will provide a comprehensive status report on the WG's activities, highlighting the intense efforts to enhance the suite's capabilities with a focus on performance optimization and sustainability.

In response to community feedback, the WG has developed new modules to measure server utilization metrics, including load, frequency, I/O, and power consumption, during the execution of the HEPsScore benchmark. These advancements enable a more detailed evaluation of power efficiency and computational performance of servers, aligning with WLCG's sustainability goals.

Furthermore, updates on the integration of GPU workloads into the benchmark suite will be presented. This significant development expands the functionality of HEPsScore, increases the catalogue of available workloads, and enhances the suite's applicability to modern and diverse computing environments.

Desired slot length:

Speaker release:

Yes

Cloud Technologies, Virtualization & Orchestration, Operating Systems / 22

Roadmap to LS3: CERN's Linux Strategy

Author: Ben Morrice¹

¹ CERN

Corresponding Author: benjamin.morrice@cern.ch

As CERN prepares for the third Long Shutdown (LS3), its evolving Linux strategy is critical to maintaining the performance and reliability of its infrastructure. This presentation will outline CERN's roadmap for Linux leading up to LS3, highlighting the rollout of RHEL and AlmaLinux 10 to ensure stability and adaptability within the Red Hat ecosystem. In parallel, we will discuss efforts to enhance the adoption of Debian as a robust alternative, bolstering flexibility and long-term sustainability as part of a comprehensive dual-ecosystem approach.

Desired slot length:

Speaker release:

Yes

Environmental sustainability, business continuity, and Facility improvement / 23

Transforming the Disaster Recovery of the Cloud Service

Author: Varsha Bhat^{None}

Corresponding Author: varsha.sharath.bhat@cern.ch

The CERN Cloud Infrastructure Service provides access to large compute and storage resources for the laboratory that includes virtual and physical machines, volumes, fileshares, loadbalancers, etc. across 2 different datacenters. With the recent addition of the Preveessin Data Center, one of the main objectives of the CERN IT Department is to ensure that all services have up-to-date procedures for disaster recovery, including the CERN Private Cloud Service.

To implement the BC/DR policy, the Cloud Team has not only executed the recovery test but also built automation on top of it. This talk will dive into the usage of tools like Terraform to ensure the redeployment and recovery of the control plane in case of a major outage.

Desired slot length:

Speaker release:

Yes

Software and Services for Operation / 24

Windows device management at CERN: A new era

Author: Siavas Firoozbakht¹

Co-authors: Pablo Martin Zamora ¹; Michal Kwiatek ¹; Jonathan Samuel ²

¹ CERN

² CERN - IT-CD-DPP

Corresponding Authors: siavas.firoozbakht@cern.ch, jonathan.rufus.samuel@cern.ch, pablo.martin.zamora@cern.ch, michal.kwiatk@cern.ch

More than 10,000 Windows devices are managed by the Windows team and delegated administrators at CERN. Ranging from workstations on which scientists run heavy simulation software, to security-hardened desktops in the administrative sector and Windows Servers that manage some of the most critical systems in the Organisation –today these systems are managed using a unified MDM solution named CMF (Computer Management Framework), developed at CERN more than 20 years ago.

As the next step into the future generation of Windows device management, a new framework to manage Windows devices at CERN has been designed and is gradually being implemented, it leverages the two MDM systems from Microsoft, Intune and Configuration Manager, which allow for co-management of both desktops and servers. The new solution expands the functionality of the former system and aligns CERN's Windows park to industry best practices, facilitating its administration and reinforcing its security posture.

This presentation will describe CERN's path to implement both systems, including the technical challenges, such as the adaptation for interoperability with our open-source SSO, while maintaining compatibility with CERN's well-established Windows infrastructure.

Desired slot length:

30

Speaker release:

Yes

Site Reports / 29

NDGF Site Report

Author: Mattias Wadenstein¹

¹ University of Umeå (SE)

Corresponding Author: mattias.wadenstein@cern.ch

New development in the distributed Nordic Tier-1 and it's participant sites.

Desired slot length:

Speaker release:

Yes

Wrap-up / 31

Wrap-up

Corresponding Author: jose.flix.molina@cern.ch

Welcome / 32

Welcome talk

Welcome / 33

Logistics talk

Corresponding Author: dino.conciatore@cscs.ch**Network & Security / 34**

Network design and implementation status of HEPS

Authors: FAZHI QI¹; MENGYAO QI¹; TAO CUI¹; YANMING WANG¹; 曾珊 zengshan^{None}¹ IHEP**Corresponding Authors:** wym@ihep.ac.cn, cuit@ihep.ac.cn, qmy@ihep.ac.cn, zengshan@ihep.ac.cn

Introduce the network architecture design of HEPS, including the general network, production network and data center network and etc.

The running status for all the network parts will also be described.

Desired slot length:

15

Speaker release:

Yes

Miscellaneous / 35

Machine learning for developers and administrators

Author: Andrey Shevel¹¹ Petersburg Nuclear Physics Institute, University of Information Technology, Mechanics and Optics**Corresponding Author:** shevel.andrey@gmail.com

Developing and managing computing systems is complex due to rapidly changing technology, evolving requirements during development, and ongoing maintenance throughout their lifespan. Significant post-deployment maintenance includes troubleshooting, patching, updating, and modifying components to meet new features or security needs. Investigating unusual events may involve reviewing system descriptions, administrator archives, administrative orders, official recommendations, and system logs. The primary goal is to keep the investigation time within reasonable limits. The machine learning system, Retrieval Augmented Generation (RAG), has been steadily advancing since around 2021. RAG can be regarded as a form of knowledge transfer. In the case studied, large computing systems are recognized as the application point of RAG, which includes a large language

model (LLM) as a collaborator for the development team. This approach offers advantages during the development process of computing systems and in the exploitation phase.

Desired slot length:

15 minutes

Speaker release:

Yes

Environmental sustainability, business continuity, and Facility improvement / 36

HEPiX Technology Watch Working Group Report

Authors: Andrea Chierici¹; Andrea Sciabà²; Shigeki Misawa³

¹ *INFN-CNAF*

² *CERN*

³ *Brookhaven National Laboratory (US)*

Corresponding Authors: chierici@cnafe.infn.it, misawa@bnl.gov, andrea.sciaba@cern.ch

The Technology Watch Working Group, established in 2018 to take a close look at the evolution of the technology relevant to HEP computing, has resumed its activities after a long pause. In this report, we provide an overview of the hardware technology landscape and some recent developments, highlighting the impact on the HEP computing community.

Desired slot length:

20

Speaker release:

Yes

Mid-long term evolution of facilities (Topical Session with WLCG OTF) / 37

IDAF @ DESY: Interdisciplinary Data and Analysis Facility: Status and Plans

Author: Yves Kemp¹

¹ *Deutsches Elektronen-Synchrotron (DE)*

Corresponding Author: yves.kemp@cern.ch

DESY operates the IDAF (Interdisciplinary Data and Analysis Facility) for all science branches: high energy physics, photon science, and accelerator R&D and operations.

The NAF (National Analysis Facility) is an integrated part, and acts as an analysis facility for the German ATLAS and CMS community as well as the global BELLE II community since 2007.

This presentation will show the current status and further plans of the implementation, driven by use cases of the different user communities.

Desired slot length:

Speaker release:

Yes

Environmental sustainability, business continuity, and Facility improvement / 38**Updates on CPUs, GPUs and AI accelerators****Authors:** Andrea Chierici¹; Andrea Sciabà²; Michele Michelotto³; Shigeki Misawa⁴¹ *Universita e INFN, Bologna (IT)*² *CERN*³ *Universita e INFN, Padova (IT)*⁴ *Brookhaven National Laboratory (US)***Corresponding Authors:** misawa@bnl.gov, michele.michelotto@cern.ch, andrea.sciaba@cern.ch, andrea.chierici@cern.ch

In this presentation we try to give an update on CPU, GPU and AI accelerators in the market today.

Desired slot length:

20/30

Speaker release:

Yes

Cloud Technologies, Virtualization & Orchestration, Operating Systems / 39**Hyper-converged cloud infrastructure at CSCS****Authors:** Dino Conciatore¹; Elia Oggian²¹ *CSCS (Swiss National Supercomputing Centre)*² *ETH Zurich (CH)***Corresponding Authors:** elia.oggian@cscs.ch, dino.conciatore@cscs.ch

This presentation provides a detailed overview of the hyper-converged cloud infrastructure implemented at the Swiss National Supercomputing Centre (CSCS). The main objective is to provide a detailed overview of the integration between Kubernetes (RKE2) and ArgoCD, with Rancher acting as a central tool for managing and deploying RKE2 clusters infrastructure-wide.

Rancher is used for direct deployment on MAAS-managed nodes, as well as HPC (High-Performance Computing) nodes designed for high-intensity workloads. In addition, Harvester orchestrates Kubernetes distributions for virtual clusters, improving flexibility and simplifying orchestration on the platform.

ArgoCD plays a key role in automating deployment processes and ensuring consistency between different environments, enabling continuous delivery. The integration of Kubernetes, ArgoCD, Rancher, Harvester and Terraform forms the basis of a hyper-converged, scalable and adaptable cloud infrastructure.

This case study provides information on the architecture, deployment workflows and operational benefits of this approach.

Desired slot length:

Speaker release:

Yes

Environmental sustainability, business continuity, and Facility improvement / 40

Trends of Energy Efficiency for computing and data centre

Author: Eric Yen¹

¹ *Academia Sinica (TW)*

Corresponding Author: han-wei.yen@cern.ch

The objective of this talk is to share the tentative plan of energy efficiency status review by the TechWatch WG. Progress of the primary tasks such as reviewing and understanding of the trends of industry, market and technology, efforts of WLCG and sites, as well as the strategy from measurement/ data collection/ analysis/ modeling/ to estimation will be shared. Through this report, the TechWatch WG would also look forward to advices, case studies, and volunteers for the future work.

Desired slot length:

20

Speaker release:

Yes

Miscellaneous / 41

RAL use case of XRootD Managers

Author: Thomas Jyothish¹

Co-authors: James William Walder²; Thomas Byrne

¹ *STFC*

² *Science and Technology Facilities Council STFC (GB)*

Corresponding Authors: tom.byrne@stfc.ac.uk, jyothish.thomas@stfc.ac.uk, james.william.walder@cern.ch

RAL makes use of the XRootD Cluster Management System to manage our XRootD server frontends for disk based storage (ECHO).

In this session, I'll give an overview of our configuration, custom scripts used and observations on its interaction on different setups.

Desired slot length:

10 minutes

Speaker release:

Yes

Environmental sustainability, business continuity, and Facility improvement / 42**Update on Energy Efficiency: AmpereOne and Turin****Authors:** David Britton¹; Emanuele Simili¹¹ *University of Glasgow (GB)***Corresponding Authors:** david.britton@cern.ch, emanuele.simili@glasgow.ac.uk

Extending the data presented at the last few HEPiX workshops, we present new measurements on the energy efficiency (HEPScore/Watt) of the recently available AmpereOne-ARM and AMD Turin-x86 machines.

Desired slot length:

10

Speaker release:

Yes

Site Reports / 44**NCG-INGRID-PT site report****Authors:** Joao Antonio Tomasio Pina¹; Jorge Gomes²¹ *Laboratory of Instrumentation and Experimental Particle Physics (PT)*² *LIP***Corresponding Authors:** joao.antonio.tomasio.pina@cern.ch, jorge@lip.pt

Evolution of the NCG-INGRID-PT site and future perspectives.

Desired slot length:**Speaker release:**

Yes

Computing and Batch Services / 46**Smart HPC-QC: flexible approaches for Quantum workloads integration**

Author: Simone Rizzo¹

¹ *E4 COMPUTER ENGINEERING Spa*

Corresponding Author: simone.rizzo@e4company.com

Many efforts have tried to combine the HPC and QC fields, proposing integrations between quantum computers and traditional clusters. Despite these efforts, the problem is far from solved, as quantum computers face a continuous evolution. Moreover, nowadays, quantum computers are scarce compared to the traditional resources in the HPC clusters: managing the access from the HPC nodes is non-trivial, as it is easy to turn the accelerator into a bottleneck. Through the SmartHPC-QC project, we design solutions to this integration issue, defining interactions based on the application pattern and depending on the underlying technology of the quantum computer. The project aims to define an integration plan that can satisfy the users' different needs without burdening them with excessive technical complexity. To achieve this goal, we use various approaches, from more typical ones (workflow-based) to more niche solutions (like virtualisation and malleability).

Desired slot length:

Speaker release:

Yes

Network & Security / 47

Update from the HEPiX IPv6 Working Group

Author: Bruno Heinrich Hoeft¹

Co-author: David Kelsey²

¹ *KIT - Karlsruhe Institute of Technology (DE)*

² *Science and Technology Facilities Council STFC (GB)*

Corresponding Authors: bruno.hoeft@kit.edu, david.kelsey@stfc.ac.uk

The HEPiX IPv6 Working Group has been encouraging the deployment of IPv6 in WLCG and elsewhere for many years. At the last HEPiX meeting in November 2024 we reported on the status of our GGUS ticket campaign for WLCG sites to deploy dual-stack computing elements and worker nodes. Work on this has continued. We have also continued to monitor the use of IPv4 and IPv6 on the LHCOPN, with the aim to identify uses of legacy IPv4 data transfers and to remove these. A dual-stack network is not the desirable end-point for all this work; we continue to plan the move from dual-stack to IPv6-only.

Desired slot length:

25

Speaker release:

Yes

Site Reports / 49

IHEP site report

Authors: Chaoqi Guo¹; Fazhi Qi²; Jingyan Shi²; Xiaowei Jiang³

¹ *Institute of High Energy Physics of the Chinese Academy of Sciences*

² *Chinese Academy of Sciences (CN)*

³ *IHEP (中国科学院高能物理研究所)*

Corresponding Authors: jiangxw@ihep.ac.cn, jingyan.shi@cern.ch, fazhi.qi@cern.ch, guocq@ihep.ac.cn

The progress and status of IHEP site since last Hepix.

Desired slot length:

12

Speaker release:

Yes

Site Reports / 50

Introduction of LHCb Tier2 Site at Lanzhou University

Author: Dong Xiao¹

¹ *Lanzhou University*

Corresponding Author: memoryfield0525@gmail.com

This report introduce the LHCb Tier-2 site at Lanzhou University(LZU-T2), which is a major new computing resource designed to support the LHCb experiment. It is part of the Worldwide LHC Computing Grid, which distributes data processing and storage across a network of international computing centers. The LZU-T2 site plays a critical role in processing, analyzing, and storing the vast amounts of data produced by the experiment. Its establishment helps balance the computing load regionally, complementing other Chinese sites such as the Beijing Tier-1 facility. The LZU-T2 site not only supports LHCb's data challenges but also helps integrate regional computing resources.

Desired slot length:

Speaker release:

Yes

Cloud Technologies, Virtualization & Orchestration, Operating Systems / 51

Kubernetes and Cloud Native at the SKA Regional Centres

Author: Rohini Joshi^{None}

Corresponding Author: rohini.joshi@skao.int

The SKA Observatory is expected to be producing up to 600 petabytes of scientific data per year, which would set a new record in data generation within the field of observational astronomy. The SRCNet infrastructure is meant for handling these large volumes of astronomy data, which requires a global network of distributed regional centres for the data- and compute-intensive astronomy use cases. On the Swiss SRCNode, we aim to use Kubernetes as a service management plane which

interacts with external storage and compute services as part of SRCNet to build a science analysis platform.

Desired slot length:

15-20 min

Speaker release:

Yes

Computing and Batch Services / 53**Efficiency of job processing in many-core Grid and HPC environments**

Author: Gianfranco Sciacca¹

¹ *Universitaet Bern (CH)*

Corresponding Author: gsciacca@mail.cern.ch

Developments in microprocessor technology have confirmed the trend towards higher core-counts and decreased amount of memory per core, resulting in major improvements in power efficiency for a given level of performance. Per node core-counts have increased significantly over the past five years for the x86_64 architecture, which is dominating in the LHC computing environment, and the higher core density is not only a feature of large HPC systems, but is also readily available on commodity hardware preferentially used at Grid sites. The baseline multi-core workloads are however still largely based on 8 cores, and the LHC experiments employ different strategies for scheduling their payloads at sites. In this work we investigate possible implications of scaling up core-counts for grid jobs, up to whole node where possible.

Desired slot length:

15 minutes

Speaker release:

Yes