

HEPiX Fall 2024 Workshop

Monday 4 November 2024 - Friday 8 November 2024



Book of Abstracts

Contents

Optimising Data Access Analytics: Integrating dCache BillingDB with PIC's Scalable Big Data Platform	1
KEK Site Report	1
Computer Security Update	2
One year into the CERN Cyber-Security Audit	2
CI4FPGA: Continuous Integration for FPGA/SoC Projects	3
pkcli: A Framework for Scripts to Manage Applications	3
Firewall under attack: operational security rollercoaster	4
CERN site report	4
PIC report	5
Cost Comparison of On-Premises Storage with S3 Interfaces	5
Using AI/ML for Data Placement Optimization in a Multi-Tiered Storage System within a Data Center	6
How AI networking Fabrics are different from today's Data Center fabrics	6
Smart Procurement Utility	7
Exploring the Carbon Compromises	7
Site report for AGLT2	8
Network tests at CZ Tier-2	8
RAL Site Report	9
Getting closer to an IPv6-only WLCG –update from the HEPiX IPv6 Working Group	9
Purdue CMS Analysis Facility	10
S3DF: SLAC Shared Science Data Facility	10
A Cloud-Native Control Plane for Infrastructure and Platform Management	11
AUDITOR: An Accounting tool for Grid Sites and Opportunistic Resources	11

Networking Topics for WLCG	12
HEPiX Benchmarking WG: Status Report	12
Jefferson Lab Site Report and HPDF Introduction	13
HELP! I have DataCenter Nightmares	13
Shifting Hardware Landscape	13
Science Cloud based on WLCG Core Technology	14
Atmospheric Visibility Estimation From Single Camera Images: A Deep Learning Approach	14
dCache on Kubernetes	15
Operating the 200 Gbps IRIS-HEP Demonstrator for ATLAS	15
Case Study: AI Training Power Demand on a GPU-Accelerated Node	16
Stories from the TSM to HPSS Migration at KIT	16
Transitioning from RHEV to Openshift	17
Welcome and Workshop Logistics	17
Research Computing and Storage Strategy at the University of Oklahoma	17
Welcome Address from the Physics Department	17
Summary of the Joint Xrootd and FTS Workshop	18
IHEP Site Report	18
SWT2 site report	18
Nebraska Coffea-Casa Analysis Facility Update	19
CRISP: Collaborative Tools for the ePIC Experiment	19
Carbon costs of storage: a UK perspective	20
Discussion	20
Natural job drainage and power reduction in PIC Tier-1 using HTCondor	20
Impetus and Drivers	21
Mini Intro: Zeek	21
Mini-intro: MISP	21
Mini-Intro: pDNSSOC	21
WLCG SOC at U Chicago	22
Discussion	22

Topic 1: Issues capturing ALL traffic with Zeek? 22

Topic 2: Building a good Zeek alert 22

Topic 1: Configuring sending alerts for various tools 22

Topic 3: How to deploy pDNSSOC (part 1) 23

Tool options and choices 23

Storage & Filesystems / 1

Optimising Data Access Analytics: Integrating dCache BillingDB with PIC's Scalable Big Data Platform

Authors: Jose Flix Molina¹; Marc Santamaria Riba²

Co-authors: Elena Planas²; Jorge Carretero Palacios³; Pau Tallada-Crespi⁴

¹ CIEMAT - Centro de Investigaciones Energéticas Medioambientales y Tec. (ES)

² PIC

³ Port d'Informació Científica

⁴ PIC-CIEMAT

Corresponding Authors: marc.santamariar@autonoma.cat, carretero@pic.es, eplanas@pic.es, jose.flix.molina@cern.ch, tallada@pic.es

PIC has developed CosmoHub, a scientific platform built on top of Hadoop and Apache Hive, which facilitates scalable reading, writing and managing huge astronomical datasets. This platform supports a global community of scientists, eliminating the need for users to be familiar with Structured Query Language (SQL). CosmoHub officially serves data from major international collaborations, including the Legacy Survey of Space and Time (LSST), the Euclid space mission, the Dark Energy Survey (DES), the Physics of the Accelerating Universe Survey (PAUS), the Gaia ESA Archive, and the Marenstrum Institut de Ciències de l'Espai (MICE) simulations.

This platform is highly scalable and adaptable for various data analytics applications. The recent integration of PIC's dCache billing database records has enabled the exploration of extensive data access logs at PIC, covering roughly eight years. We will share insights from the analysis of CMS data access at PIC, which involved processing approximately 350 million entries using PIC's Hadoop infrastructure. The current system operates with around 1,000 cores, 10 TiB of RAM, 50 TB of NVMe (for caching), and 2 PiB of usable storage. Data is accessed through HiveQL and Jupyter notebooks, with advanced Python scripts enabling efficient interaction.

This framework significantly accelerated data processing, reducing execution times for plot generation to under a minute - a task that previously took several hours using PIC's PostgreSQL databases. This enhanced performance opens up new possibilities for integrating additional data sources, such as job submissions from the local HTCondor batch system, enabling advanced analytics on large datasets.

Desired slot length:

20

Speaker release:

Yes

Site Reports / 2

KEK Site Report

Author: Go Iwai¹

Co-authors: Koichi Murakami ; So Suzuki ; Tomoaki Nakamura ; Tomoe Kishimoto

¹ KEK

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

The KEK Central Computer System (KEKCC) is the KEK's largest-scale computer system and provides several services such as Grid and Cloud computing.

Following the procurement policy for the large-scale computer system requested by the government, we have taken a multiple-year contract and replaced the entire system at the end of every contract year. The new system has been in production since September 2024 and will decommission in August 2028.

In this talk, we would like to review the four-year operation and development of the previous system installed in 2020. In addition, we will then show the difference between before and after September 2024, when the system is in production.

Desired slot length:

Speaker release:

Yes

Networking & Security / 3

Computer Security Update

Author: Stefan Lueders¹

¹ CERN

Corresponding Author: stefan.lueders@cern.ch

This presentation aims to give an update on the global security landscape from the past year. The global political situation has introduced a novel challenge for security teams everywhere. What's more, the worrying trend of data leaks, password dumps, ransomware attacks and new security vulnerabilities does not seem to slow down.

We present some interesting cases that CERN and the wider HEP community dealt with in the last year, mitigations to prevent possible attacks in the future and preparations for when inevitably an attacker breaks in.

Desired slot length:

Speaker release:

Yes

Networking & Security / 4

One year into the CERN Cyber-Security Audit

Author: Stefan Lueders¹

¹ CERN

Corresponding Author: stefan.lueders@cern.ch

This talk presents the findings of the 2023 cybersecurity audit undertaken at CERN, and the resulting plans/progress/accomplishment the Organization experienced in the past 9 months while implementing their recommendations.

Desired slot length:

Speaker release:

No

Basic and end-user IT services / 5

CI4FPGA: Continuous Integration for FPGA/SoC Projects

Authors: Carmen Marcos^{None}; Christos Gentsos¹

¹ CERN (IT-CA-GES)

Corresponding Authors: christos.gentsos@cern.ch, carmen.marcos.sanchez.de.la.blanca@cern.ch

As the complexity of FPGA and SoC development grows, so does the need for efficient and automated processes to streamline testing, building, and collaboration, particularly in large-scale scientific environments such as CERN. This initiative focuses on providing CI infrastructure that is tailored for FPGA development and pre-configured Docker images for essential EDA tools, keeping the learning slope for the more than 100 projected users of the service minimal and using centralized and managed infrastructure that aligns well with CERN's IT services. This centralization facilitates the seamless integration of tools and workflows across diverse experiments, ensuring that development efforts are unified and scalable.

CI4FPGA facilitates testing and building processes by enabling automated pipelines, enhancing collaboration between development teams, and improving overall efficiency. The project frees FPGA designers from the resource-intensive task of maintaining clusters and container images, freeing them up to address key challenges such as automating unit and system-level testing, facilitating shared development of IP cores, among other benefits. One of the features employed is lazy pulling technology, that makes it possible to use scalable VM-based clusters with limited SSD sizes and drastically reduces container image load times from ~15 minutes to ~15 seconds.

Desired slot length:

15 minutes

Speaker release:

Yes

Basic and end-user IT services / 6

pkcli: A Framework for Scripts to Manage Applications

Author: evan carlin¹

¹ RadiaSoft LLC

Corresponding Author: evan@radiasoft.net

System administrators and developers need a way to call application code and other tasks through command line interfaces (CLIs). Some examples include user management (creation, deletion, moderation, etc) or seeding the database for development. We have developed an open source Python framework, `pykern.pkcli`, that simplifies the creation of these application-specific CLIs. In this talk, I will provide an overview of our framework and share examples of how we've used it to administer our systems. I'll discuss the advantages of using `pykern.pkcli` over traditional shell scripts, including improvements in development, testing, modification, and distribution. Additionally, I'll present a case study demonstrating how we use one of these scripts to manage user access control for an application and seamlessly share code between the CLI and a web interface.

Desired slot length:

12

Speaker release:

Yes

Networking & Security / 7

Firewall under attack: operational security rollercoaster

Authors: Patrick Storm^{None}; Romain Wartel¹

¹ *CERN*

Corresponding Authors: storm@lbl.gov, romain.wartel@cern.ch

This talk will walk you through the challenges the ESnet security team faced during an attack against one of its firewalls. It covers the struggle and drama to access the data we needed and, in the end, highlights how nothing quite beats good old-fashioned, down-and-dirty system forensics.

Desired slot length:

20

Speaker release:

Yes

Site Reports / 8

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

Site Reports / 9**PIC report****Author:** Jose Flix Molina¹**Co-authors:** Agustin Bruzzese ; Andreu Pacheco Pages²; Antoni Alou³; Antonio Perez-Calero Yzquierdo⁴; Carles Acosta Silva³; Carlos Perez Dengra³; Christian Neissner³; Elena Planas³; Elisabeth Gonzalez³; Esther Accion Garcia⁵; Francesc Torradeflot ; Gonzalo Merino⁶; Jaime Priego³; Jordi Casals Hernandez⁷; Jordi Delgado Not Supplied ; Jorge Carretero Palacios⁸; Laura Cabayol³; MC Porto⁹; Manuel Delfino Reznicek⁵; Marc Santamaria Riba¹⁰; Martin Børstad Eriksen ; Pau Tallada-Crespí¹¹; Ricard Cruz Josa⁵; Ruben Dario Graña ; Vanessa Acin Portella¹²¹ CIEMAT - Centro de Investigaciones Energéticas Medioambientales y Tec. (ES)² Institut de Física d'Altes Energies - Barcelona (ES)³ PIC⁴ Centro de Investigaciones Energéticas Medioambientales y Tecnológicas⁵ The Barcelona Institute of Science and Technology (BIST) (ES)⁶ IFAE - Institute for High Energy Physics⁷ Port d'Informació Científica (PIC)⁸ Port d'Informació Científica⁹ CIEMAT¹⁰ UAB¹¹ PIC-CIEMAT¹² Institut de Física d'Altes Energies**Corresponding Author:** jose.flix.molina@cern.ch

PIC report to HEPiX Fall 2024.

Desired slot length:**Speaker release:**

Yes

Storage & Filesystems / 11**Cost Comparison of On-Premises Storage with S3 Interfaces****Author:** Nathan Thompson¹**Co-authors:** Matt Ninesling¹; Matthew Starr²¹ Spectra Logic² Spectra Logic Corporation

Corresponding Authors: nathant@spectrallogic.com, mattn@spectrallogic.com, matts@spectrallogic.com

Abstract: To evaluate the cost of various on-premises storage solutions with traditional and S3 interfaces, including flash, disk, and tape.

This presentation compares the costs, factors of flash, disk, and tape-based storage systems, including systems that are compatible with AWS S3. Key metrics to be considered include purchase price, power consumption, cooling requirements, product lifetime and performance characteristics. Additionally, the presentation will explore the long-term implications of each storage type and their impact on the environment.

This presentation will also review current and upcoming technologies that may be leveraged to provide long term exascale storage at a fraction of environment footprint of traditional methods.

Desired slot length:

30

Speaker release:

Yes

Storage & Filesystems / 12

Using AI/ML for Data Placement Optimization in a Multi-Tiered Storage System within a Data Center

Authors: Calos Deleon¹; James Leonardi²; Qiulan Huang³

Co-authors: Shinjae Yoo²; Vincent Garonne²

¹ *Stony Brook University*

² *Brookhaven National Laboratory*

³ *Brookhaven National Laboratory (US)*

Corresponding Authors: qiulan.huang@cern.ch, sjyoo@bnl.gov, codingcarlos23@gmail.com, jleonardi@bnl.gov, vgaronne1@bnl.gov

Scientific experiments and computations, particularly in High Energy Physics (HEP) programs, are generating and accumulating data at an unprecedented rate. Effectively managing this vast volume of data while ensuring efficient data analysis poses a significant challenge for data centers. This paper aims to introduce machine learning algorithms to enhance data storage optimization across various storage media, providing a more intelligent, efficient, and cost-effective approach to data management. We begin by outlining the data collection and preprocessing steps used to explore data access patterns. Next, we describe the design and development of a precise data popularity prediction model using AI/ML techniques. This model forecasts future data popularity based on an analysis of access patterns, enabling optimal data movement and placement. Additionally, the paper evaluates the model's performance using key metrics such as F1 score, accuracy, precision, and recall, alongside a comparison with the Least Recently Used (LRU) strategy. The model achieves an optimal prediction accuracy of up to 92% and an optimal F1 score of 0.47. Finally, we present a prototype use case, leveraging real-world file access data to assess the model's performance.

Desired slot length:

Speaker release:

No

Networking & Security / 13**How AI networking Fabrics are different from today's Data Center fabrics****Author:** Paul Gilbert¹¹ *Arista Networks***Corresponding Author:** pgilbert@arista.com

This presentation looks at what is different about building and deploying AI fabrics. I can if needed remove the Arista logo's from the presentation. I don't see a place to attach the presentation?

Desired slot length:

30minutes

Speaker release:

Yes

Topical Session: Carbon & Sustainability in Data Centers / 14**Smart Procurement Utility****Author:** David Britton¹**Co-author:** Dan Protopopescu¹¹ *University of Glasgow (GB)***Corresponding Authors:** david.britton@cern.ch, dan.protopopescu@cern.ch

The Smart Procurement Utility is a tool that allows the visualisation of HEPScore/Watt vs HEPScore/unit-cost to guide procurement choices and the compromise between cost and carbon. It uses existing benchmarking data and allows the entry of new benchmarking data. Costs can be entered as relative numbers (percentages relative to a chosen baseline) to generate the cost-related plots.

Desired slot length:

10

Speaker release:

Yes

Topical Session: Carbon & Sustainability in Data Centers / 15**Exploring the Carbon Compromises****Authors:** Albert Gyorgy Borbely¹; David Britton¹; Emanuele Simili¹; Gordon Stewart^{None}; Samuel Cadellin Skipsey^{None}; Steve Lloyd²¹ *University of Glasgow (GB)*

² *University of London (GB)*

Corresponding Authors: david.britton@cern.ch, samuel.cadellin.skipsey@cern.ch, emanuele.simili@glasgow.ac.uk, gordon.stewart@glasgow.ac.uk, albert.borbely@cern.ch, s.l.lloyd@qmul.ac.uk

Minimising carbon associated with computing will require compromise. In this presentation I will present the results from simulating a Grid site where the compute is run at reduced frequency when the predicted carbon intensity rises above some threshold. The compromise is a reduction in throughput in exchange for an increased carbon-efficiency for the work that is completed. The presentation will also summarise other, related, work from the Glasgow group.

Desired slot length:

20

Speaker release:

Yes

Site Reports / 16

Site report for AGLT2

Authors: Philippe Laurens¹; Shawn Mc Kee²; Wendy Wu³

¹ *Michigan State University (US)*

² *University of Michigan (US)*

³ *University of Michigan*

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

AGLT2 has a few updates to report since the last HEPix meeting in Spring 2024.

- 1) We transitioned from Cobbler and Satellite plus Capsule server for RHEL provision
- 2) we transitioned from CFengine to Ansible for configuration management for the RHEL9 nodes.
- 3) In order to improve the occupancy of the HTcondor cluster, we started tuning of HTCondor and also new developments of scripts to dynamically adjust the routing rules and monitoring scripts and plots keep track of memory/cpu occupancy.

Desired slot length:

Speaker release:

Yes

Networking & Security / 17

Network tests at CZ Tier-2

Authors: Jiri Chudoba¹; Petr Vokac²

¹ *Czech Academy of Sciences (CZ)*

² *Czech Technical University in Prague (CZ)*

Corresponding Authors: jiri.chudoba@cern.ch, petr.vokac@cern.ch

The CZ Tier-2 in Prague (the Czech Republic) joined the WLCG Data Challenge 24 and managed to receive and sent more than 2 PB during the second week of the DC24. Since then we upgraded our network connection to LHCONE from 100 to 2x100 Gbps. The LHCONE link uses GEANT connection, which was also upgraded to 2x100 Gbps. During July 2024 we executed dedicated network stress tests between Prague and CERN and we observed maxima close to the link capacity - 200 Gbps.

Desired slot length:

10

Speaker release:

Yes

Site Reports / 18

RAL Site Report

Author: Martin Bly¹

¹ *STFC-RAL*

Corresponding Author: martin.bly@stfc.ac.uk

An update on activities at the RAL datacentre.

Desired slot length:

Speaker release:

Yes

Networking & Security / 19

Getting closer to an IPv6-only WLCG –update from the HEPiX IPv6 Working Group

Author: Martin Bly¹

Co-author: David Kelsey²

¹ *STFC-RAL*

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

Corresponding Authors: martin.bly@stfc.ac.uk, david.kelsey@stfc.ac.uk

The HEPiX IPv6 Working Group has been encouraging the deployment of IPv6 in WLCG for many years. At the last HEPiX meeting in Paris we reported that the LHC experiment Tier-2 storage services are now close to 100% IPv6-capable. We had turned our attention to WLCG compute and launched a GGUS ticket campaign for WLCG sites to deploy dual-stack computing elements and worker nodes. At that time 44% of the sites had completed their deployment of dual-stack CPU. The

working group has also continued to monitor the use of IPv4 and IPv6 on the LHCOPN. As before we continue to identify uses of legacy IPv4 data transfers and strive to move these to IPv6. A dual-stack network is not a desirable end-point for all this work; we continue to plan the move from dual-stack to IPv6-only.

This talk will present the activities of the working group since April 2024 and our future plans.

Desired slot length:

20 minutes

Speaker release:

Yes

Operating systems, clouds, virtualisation, grids / 20

Purdue CMS Analysis Facility

Authors: Dmitry Kondratyev¹; Norbert Neumeister¹; Stefan Piperov¹

¹ *Purdue University (US)*

Corresponding Authors: dmitry.kondratyev@cern.ch, norbert.neumeister@cern.ch, stefan.piperov@cern.ch

The Purdue Analysis Facility (Purdue AF) is an advanced computational platform designed to support high energy physics (HEP) research at the CMS experiment. Based on a multi-tenant JupyterHub server deployed on a Kubernetes cluster, Purdue AF leverages the resources of the Purdue CMS Tier-2 computing center to provide scalable, interactive environments for HEP workflows. It supports a full HEP analysis software stack, offers a variety of storage and data access solutions, and integrates modern scale-out tools like Dask Gateway. Since its first deployment in 2023, Purdue AF has been instrumental in numerous published analyses, workshops, and tutorials. We will present the Purdue AF architecture and describe its common use patterns in CMS analyses.

Desired slot length:

15-20

Speaker release:

Yes

Site Reports / 21

S3DF: SLAC Shared Science Data Facility

Author: Adeyemi Adesanya¹

¹ *SLAC*

Corresponding Author: yemi@slac.stanford.edu

A site report on the infrastructure and services that underpin SLAC's data-intensive processing pipelines. The SLAC Shared Science Data Facility hosts the Rubin Observatory DF, LCLS-II and

many other experimental and research workflows. Networking and Storage form the core of S3DF with hardware deployed in a modern Stanford datacenter.

Desired slot length:

Speaker release:

Yes

Operating systems, clouds, virtualisation, grids / 22

A Cloud-Native Control Plane for Infrastructure and Platform Management

Author: Dino Conciatore¹

¹ CSCS (Swiss National Supercomputing Centre)

Corresponding Author: dino.conciatore@cscs.ch

Crossplane is a cloud-native control plane for declarative management of infrastructure and platform resources using Kubernetes-native APIs.

It enables the integration of infrastructure-as-code practices by reusing existing tools such as Ansible and Terraform, while providing flexible, instanceable “compositions” for defining reusable resource configurations. This approach allows organizations to automate, compose, and manage infrastructure alongside application workloads, streamlining operations in a cloud-native ecosystem.

Desired slot length:

20

Speaker release:

Yes

Computing & Batch Services / 23

AUDITOR: An Accounting tool for Grid Sites and Opportunistic Resources

Author: Matthias Jochen Schnepf^{None}

Corresponding Author: matthias.jochen.schnepf@cern.ch

More and more opportunistic resources are provided to the Grid. Often behind one Compute Element several opportunistic computing resource provider exists or are additional to the pledged resources of a Grid site. For such use cases and others, we have developed a most flexible multipurpose accounting ecosystem AUDITOR (AccoUnting Datahandling Toolbox for Opportunistic Resources).

AUDITOR is able to individually collect accounting data from multiple resource providers sharing a CE. The collected information can be used for internal accounting or sent to the European Grid Initiative (EGI) accounting portal. We will show some current use-cases and further plans for AUDITOR.

Desired slot length:

15

Speaker release:

Yes

Networking & Security / 24**Networking Topics for WLCG****Authors:** Marian Babik¹; Shawn Mc Kee²¹ *CERN*² *University of Michigan (US)***Corresponding Authors:** marian.babik@cern.ch, shawn.mckee@cern.ch

We will describe the current activities and plans in WLCG networking, including details about SciTags, the WLCG perfSONAR deployment and the related activities to monitor and analyze our networks. We will also describe the related efforts to plan for the upcoming WLCG Network Data Challenge through a series of mini-challenges that incorporate our tools and metrics.

Desired slot length:

25

Speaker release:

Yes

Computing & Batch Services / 25**HEPiX Benchmarking WG: Status Report****Author:** Matthias Jochen Schnepf^{None}**Corresponding Author:** matthias.jochen.schnepf@cern.ch

HEPSCORE23 has been the official benchmark for WLCG sites since April 2023.

Since then, we have included community feedback and demand. The Benchmarking WG has started a new development effort to expand the Benchmark Suite with modules that can measure server utilization metrics (load, frequency, I/O, power consumption) during the execution of the HEPSCORE benchmark.

This enables a closer look at power efficiency and performance in the Grid environment.

We present the current state of our group and an overview of our current studies.

Desired slot length:

15

Speaker release:

Yes

Site Reports / 26**Jefferson Lab Site Report and HPDF Introduction****Author:** Bryan Hess^{None}**Corresponding Author:** bhess@jlab.org

I will give an report on the Scientific Computing program at Jefferson Lab and a brief introduction to HPDF, the High Performance Data Facility.

Desired slot length:

12

Speaker release:

Yes

Networking & Security / 27**HELP! I have DataCenter Nightmares****Author:** Stefan Lueders¹¹ CERN**Corresponding Author:** stefan.lueders@cern.ch

With the growing complexity of the IT hardware and software stack, with a move from bare-metal to virtual machines & containers, with the prevalent usage of shared central computing resources for Internet-facing services, provisioning of (internal) user services but also the need for serving industrial control systems (OT) in parallel, the design of data centre architectures and in particular its networks can become more and more challenging. This presentation will introduce the dilemma of creating a highly agile and flexible computer center set-up while still trying to maintain security perimeters within. It is bound to fail.

Desired slot length:**Speaker release:**

Yes

Miscellaneous / 28**Shifting Hardware Landscape****Author:** Shigeki Misawa¹¹ Brookhaven National Laboratory (US)**Corresponding Author:** misawa@bnl.gov

Advances in computing hardware are essential for future HEP and NP experiments. These advances are seen as incremental improvements in performance metric over time, i.e. everything works the

same, just better, faster, and cheaper. In reality, hardware advances and changes in requirements can result in the crossing of thresholds that require a re-evaluation of existing practices. The HEPiX Techwatch working group was created to monitor trends in technology that will impact HEP and NP experiments in the future.

Desired slot length:

10 minutes

Speaker release:

Yes

Operating systems, clouds, virtualisation, grids / 32

Science Cloud based on WLCG Core Technology

Author: Eric Yen¹

¹ *Academia Sinica (TW)*

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

This presentation will focus on two topics: 1) status of ATLAS T2 site in Taiwan, and 2) experiences of supporting broader scientific computing over the cloud based on WLCG technology.

Desired slot length:

Speaker release:

Yes

Miscellaneous / 33

Atmospheric Visibility Estimation From Single Camera Images: A Deep Learning Approach

Author: Anderw Fagg¹

¹ *University of Oklahoma*

Corresponding Author: fagg@ou.edu

Atmospheric Visibility Estimation From Single Camera Images: A Deep Learning Approach

Desired slot length:

45

Speaker release:

Yes

Operating systems, clouds, virtualisation, grids / 34**dCache on Kubernetes**

Author: Elia Luca Oggian¹

¹ *ETH Zurich (CH)*

Corresponding Author: elia.oggian@cscs.ch

dCache is composed by a set of components running in Java Virtual Machines (JVM) and a storage backend, Ceph in this case. CSCS moved these JVMs into containers and developed an Helm Chart to deploy them on a Kubernetes cluster. This cloud native approach makes the deployments and management of new dCache instances easier and faster.

Encountered challenges and future developments will be exposed in this presentation.

Desired slot length:

15 minutes

Speaker release:

Yes

Computing & Batch Services / 35**Operating the 200 Gbps IRIS-HEP Demonstrator for ATLAS**

Authors: Alexander Held¹; David Jordan²; Doug Benjamin³; Farnaz Golnaraghi²; Fengping Hu²; Gordon Watts⁴; Ilija Vukotic²; Judith Lorraine Stephen²; Lincoln Bryant²; Matthew Feickert¹; Ofer Rind⁵; Robert William Gardner Jr²; Wei Yang⁶

¹ *University of Wisconsin Madison (US)*

² *University of Chicago (US)*

³ *Brookhaven National Laboratory (US)*

⁴ *University of Washington (US)*

⁵ *Brookhaven National Laboratory*

⁶ *SLAC National Accelerator Laboratory (US)*

Corresponding Authors: fengping@uchicago.edu, ilija.vukotic@cern.ch, matthew.feickert@cern.ch, djordan66@uchicago.edu, robert.w.gardner@cern.ch, ofer.rind@cern.ch, yangw@slac.stanford.edu, judith.lorraine.stephen@cern.ch, farnaz@uchicago.edu, lincoln.bryant@cern.ch, douglas.benjamin@cern.ch, alexander.held@cern.ch, gwatts@uw.edu

The ATLAS experiment is currently developing multiple analysis frameworks which leverage the Python data science ecosystem. We describe the setup and operation of the infrastructure necessary to support demonstrations of these frameworks. One such demonstrator aims to process the compact ATLAS data format PHYSLITE at rates exceeding 200 Gbps. Integral to this study was the analysis of network traffic and bottlenecks, worker node scheduling, disk configurations, and the performance of an S3 object store. The demonstration's performance was measured as the number of processing cores used by the demonstration tests scaled to over 2,000 and as the volume of data accessed in an interactive session approached 200 TB. The presentation will go over the findings and future updates related to the physical infrastructure that supports these demonstrators and what improvements to infrastructure will be made to be better prepared for the future.

Desired slot length:

15 minutes

Speaker release:

Yes

Topical Session: Carbon & Sustainability in Data Centers / 36**Case Study: AI Training Power Demand on a GPU-Accelerated Node****Author:** Imran Latif¹**Co-author:** Shigeki Misawa²¹ Brookhaven National Laboratory² Brookhaven National Laboratory (US)**Corresponding Authors:** misawa@bnl.gov, ilatif@bnl.gov

Data center sustainability, a phenomenon that has grown in focus due to the continuing evolution of Artificial intelligence (AI)/High Performance Computing (HPC) systems; furthermore, the rampant increase in carbon emissions resulted in an unprecedented rise in Thermal Design Power (TDP) of the computer chips at the Scientific Data and Computing Center (SDCC) at Brookhaven National Laboratory (BNL). With the exponential increase of demand towards the usage of such systems, major challenges have surfaced in terms of productivity, Power Usage Effectiveness (PUE), and thermal/scheduling management.

Deploying AI/HPC infrastructure in data centers will require substantial capital investment. This study quantified the energy footprint of this infrastructure by developing models based on the power demands of AI hardware during training. We measured the instantaneous power draw of an 8-GPU NVIDIA H100 HGX node while training open-source models, including the image classifier and the large language model. The peak power draw observed nearly 18% below the manufacturer's rated TDP, even with GPUs near full utilization. For the image classifier, increasing the batch size from 512 to 4096 images reduced total training energy consumption by a factor of four when model architecture remained constant. These insights can aid data center operators in capacity planning and provide researchers with energy use estimates. Future studies will explore the effects of cooling technologies and carbon-aware scheduling on AI workload energy consumption.

Desired slot length:

12

Speaker release:

Yes

Storage & Filesystems / 37**Stories from the TSM to HPSS Migration at KIT****Author:** Andreas Petzold¹**Co-authors:** Artur Il Darovic Gottmann¹; Dorin-Daniel Lobontu; Doris Ressmann; Haykuhi Musheghyan²; Karin Schaefer; Preslav Konstantinov¹; Samuel Ambroj Perez; Xavier Mol¹¹ KIT - Karlsruhe Institute of Technology (DE)² Georg August Universitaet Goettingen (DE)

Corresponding Authors: dorin-daniel.lobontu@kit.edu, doris.ressmann@kit.edu, artur.akhmetshin@cern.ch, preslav.konstantinov@cern.ch, samuel.perez@kit.edu, karin.schaefer@kit.edu, andreas.petzold@cern.ch, haykuhi.musheghyan@cern.ch, xavier.mol@cern.ch

In 2020 we started the migration from our TSM-based tape system to HPSS which was finally finished in the summer of 2024. I'll present lessons learned, pitfalls and also the necessary in-house software developments.

Desired slot length:

Speaker release:

Yes

Operating systems, clouds, virtualisation, grids / 38

Transitioning from RHEV to Openshift

Author: Zeyi Tang¹

Co-author: Robert Hancock

¹ *Brookhaven National Laboratory*

Corresponding Authors: hancock@bnl.gov, zeyi@bnlgov

A description of our experience deploying Openshift both for container orchestration as well as a replacement for Redhat Enterprise Virtualization.

Desired slot length:

20

Speaker release:

Yes

Welcome and Workshop Logistics / 39

Welcome and Workshop Logistics

Corresponding Author: severini@ou.edu

Welcome and Workshop Logistics / 40

Research Computing and Storage Strategy at the University of Oklahoma

Corresponding Authors: hneeman@ou.edu, severini@ou.edu

Welcome and Workshop Logistics / 41**Welcome Address from the Physics Department**

Corresponding Authors: pgutierrez@ou.edu, severini@ou.edu

Operating systems, clouds, virtualisation, grids / 42**Summary of the Joint Xrootd and FTS Workshop**

Author: Wei Yang¹

Co-author: Horst Severini²

¹ *SLAC National Accelerator Laboratory (US)*

² *University of Oklahoma (US)*

Corresponding Authors: hs@nhn.ou.edu, yangw@slac.stanford.edu

The 2nd Joint Xrootd and FTS Workshop at STFC in September 2024 covered many interesting topics. This presentation will summarize the discussion on state of affairs of FTS and Xrootd, plan on FTS4, WLCG token support in FTS, future plan on CERN Data Management Client, The Pelican project and Xrootd/Xcache, Xrootd monitoring, etc. It will cover some of the feedback by experiments, especially with regards to the DC24, and future plan by various experiments.

Desired slot length:

Speaker release:

Yes

Site Reports / 43**IHEP Site Report**

Author: Siqi Hou^{None}

Corresponding Author: housq@ihep.ac.cn

The progress and status of IHEP site since last Hepix.

Desired slot length:

Speaker release:

Yes

Site Reports / 44**SWT2 site report**

Author: Zachary Booth¹

Co-authors: Armen Vartapetian¹; Horst Severini²; Mark Sosebee¹

¹ *University of Texas at Arlington*

² *University of Oklahoma*

Corresponding Authors: zachary.booth@uta.edu, vartap@uta.edu, sosebee@uta.edu, hs@nhn.ou.edu

The Southwest Tier-2 (SWT2) consortium is comprised of two data centers operated at the University of Texas at Arlington (UTA) and at the University of Oklahoma (OU). SWT2 provides distributed computing services in support of the ATLAS experiment at CERN. In this presentation we will describe the resources at each site (CPU cycles and data storage), along with other associated infrastructure required to provide these resources. We will conclude with a discussion of plans for the future evolution of SWT2.

Desired slot length:

Speaker release:

Yes

Operating systems, clouds, virtualisation, grids / 45

Nebraska Coffea-Casa Analysis Facility Update

Author: Garhan Attebury¹

¹ *University of Nebraska Lincoln (US)*

Corresponding Author: garhan.attebury@cern.ch

The CMS Coffea-Casa analysis facility at the University of Nebraska-Lincoln provides researchers with Kubernetes based Jupyter environments and access to CMS data along with both CPU and GPU resources for a more interactive analysis experience than traditional clusters provide. This talk will cover updates to this facility within the past year and recent experiences with the 200 Gbps challenge.

Desired slot length:

15-20

Speaker release:

Yes

Basic and end-user IT services / 46

CRISP: Collaborative Tools for the ePIC Experiment

Authors: Dmitry Arkhipkin^{None}; Eric Lancon¹; Jerome LAURET²; Ofer Rind²; Peter Alan Steinberg¹; Robert Hancock^{None}; Uma Ganapathy³; Vincent Garonne²

¹ *Brookhaven National Laboratory (US)*

² *Brookhaven National Laboratory*

³ *BNL*

Corresponding Authors: jlauret@bnl.gov, ganapathy@bnl.gov, arkipkin@bnl.gov, hancock@bnl.gov, vgaronne1@bnl.gov, peter.steinberg@bnl.gov, elancon@bnl.gov, ofer.rind@cern.ch

This talk describes a project to develop a set of collaborative tools for the upcoming ePIC experiment at the BNL Electron-Ion Collider (EIC). The “Collaborative Research Information Sharing Platform” (CRISP) is built upon an extensible, full-featured membership directory, with CoManage integration and a customized InvenioRDM document repository. The CRISP architecture will be presented, along with plans for future integrations and workflow development.

Desired slot length:

Speaker release:

Yes

Topical Session: Carbon & Sustainability in Data Centers / 47

Carbon costs of storage: a UK perspective

Authors: Alison Packer^{None}; Alison Packer¹; Samuel Cadellin Skipsey^{None}

¹ *STFC - Science & Technology Facilities Council (GB)*

Corresponding Authors: alison.packer@stfc.ac.uk, alison.packer@cern.ch, samuel.cadellin.skipsey@cern.ch

In order to achieve the higher performance year on year required by the 2030s for future LHC upgrades at a sustainable carbon cost to the environment, it is essential to start with accurate measurements of the state of play. Whilst there have been a number of studies of the carbon cost of compute for WLCG workloads published, rather less has been said on the topic of storage, both nearline and archival. We present a study of the embedded and ongoing carbon costs of storage in multiple configurations, from Tape farms through to SSDs, within the UK Tier-1 and Tier-2s and discuss how this directs future policy.

Desired slot length:

Speaker release:

Yes

Topical Session: Carbon & Sustainability in Data Centers / 48

Discussion

Topical Session: Carbon & Sustainability in Data Centers / 49

Natural job drainage and power reduction in PIC Tier-1 using HT-Condor

Author: Jose Flix Molina¹

Co-authors: Carles Acosta Silva ²; Gonzalo Merino ³; Jordi Casals Hernandez ⁴; Kevin Fabrega ⁵

¹ CIEMAT - Centro de Investigaciones Energéticas Medioambientales y Tec. (ES)

² PIC

³ IFAE - Institute for High Energy Physics

⁴ Port d'Informació Científica (PIC)

⁵ UAB

Corresponding Authors: jose.flix.molina@cern.ch, kevin.fabrega@autonoma.cat, merino@pic.es

I will present some preliminary studies and ideas to understand natural job drainage and power reduction in PIC Tier-1, which is using HTCondor. Based on the historical batch system logs, we are simulating natural drainage and understanding how we can modulate the PIC farm without killing jobs.

Desired slot length:

Speaker release:

Yes

Topical Session: Security Operations Center (SOC) / 50

Impetus and Drivers

Corresponding Author: david.crooks@cern.ch

Topical Session: Security Operations Center (SOC) / 51

Mini Intro: Zeek

Corresponding Author: asharma@lbl.gov

Topical Session: Security Operations Center (SOC) / 52

Mini-intro: MISP

Corresponding Author: james.acris@stfc.ac.uk

Topical Session: Security Operations Center (SOC) / 53

Mini-Intro: pDNSSOC

Corresponding Author: romain.wartel@cern.ch

Topical Session: Security Operations Center (SOC) / 54

WLCG SOC at U Chicago

Corresponding Author: djordan66@uchicago.edu

Topical Session: Security Operations Center (SOC) / 55

Discussion

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

We need to have a discussion about what sites and possible users need and expect.

The goal is to both clarify details and get guidance for what we should focus on during the afternoon sessions today.

SOC Hackathon / 56

Topic 1: Issues capturing ALL traffic with Zeek?

Corresponding Authors: david.crooks@cern.ch, shawn.mckee@cern.ch, djordan66@uchicago.edu, asharma@lbl.gov

We have some sites that have question/potential issues concerning the traffic measurements from Zeek vs SNMP.

- Should be expect that the Zeek traffic estimate should be close to the SNMP counters from the corresponding switch ports?
- Is some kind of NIC/hardware offloading hiding traffic from Zeek?
- Do we have best practice recommendations regarding configurations?
- What should sites expect regarding Zeek traffic monitoring and traffic estimations?

SOC Hackathon / 57

Topic 2: Building a good Zeek alert

Corresponding Authors: david.crooks@cern.ch, romain.wartel@cern.ch, asharma@lbl.gov

What does it take to craft a good Zeek alert? Can we work through an example or two? What is the suggested guidance for doing this?

SOC Hackathon / 58

Topic 1: Configuring sending alerts for various tools

Corresponding Authors: liam.atherton@stfc.ac.uk, romain.wartel@cern.ch, david.crooks@cern.ch

How to enable alerts using webhooks and various applications.

Sending to SLACK

Sending to Mattermost

What about Keybase?

Why not email?

SOC Hackathon / 59

Topic 3: How to deploy pDNSSOC (part 1)

Corresponding Author: romain.wartel@cern.ch

How to deploy pDNSSOC

Example deployment

Working session

SOC Hackathon / 60

Tool options and choices

Corresponding Authors: stefan.lueders@cern.ch, david.crooks@cern.ch

Zeek, MISP, pDNSSOC, Elasticsearch, Opensearch, Elastiflow, ElastiAlert, other information sources, other tools?

Advantages, capabilities, limitations, concerns...

Let's discuss