

HEPiX Fall 2024 Workshop



Report of Contributions

Contribution ID: 1

Type: **not specified**

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

Tuesday 5 November 2024 14:00 (30 minutes)

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

Primary authors: FLIX MOLINA, Jose (CIEMAT - Centro de Investigaciones Energéticas Medioambientales y Tec. (ES)); Mr SANTAMARIA RIBA, Marc (PIC)

Co-authors: PLANAS, Elena (PIC); CARRETERO PALACIOS, Jorge (Port d'Informació Científica); Mr TALLADA-CRESPÍ, Pau (PIC-CIEMAT)

Presenters: FLIX MOLINA, Jose (CIEMAT - Centro de Investigaciones Energéticas Medioambientales y Tec. (ES)); Mr SANTAMARIA RIBA, Marc (PIC)

Session Classification: Storage & Filesystems

Track Classification: Storage & Filesystems

Contribution ID: 2

Type: **not specified**

KEK Site Report

Monday 4 November 2024 13:50 (20 minutes)

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

Primary author: IWAI, Go (KEK)**Co-authors:** Dr MURAKAMI, Koichi; SUZUKI, So; NAKAMURA, Tomoaki; KISHIMOTO, Tomoe**Presenter:** IWAI, Go (KEK)**Session Classification:** Site Reports**Track Classification:** Site Reports

Contribution ID: 3

Type: **not specified**

Computer Security Update

Wednesday 6 November 2024 11:25 (35 minutes)

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

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

Contribution ID: 4

Type: **not specified**

One year into the CERN Cyber-Security Audit

Thursday 7 November 2024 13:30 (30 minutes)

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

Primary author: LUEDERS, Stefan (CERN)

Presenter: LUEDERS, Stefan (CERN)

Session Classification: Networking & Security

Track Classification: Networking & Security

Contribution ID: 5

Type: **not specified**

CI4FPGA: Continuous Integration for FPGA/SoC Projects

Friday 8 November 2024 09:30 (30 minutes)

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

Primary authors: MARCOS, Carmen; GENTSOS, Christos (CERN (IT-CA-GES))**Presenter:** MARCOS, Carmen**Session Classification:** Basic and end-user IT services**Track Classification:** Basic and End-User IT Services

Contribution ID: 6

Type: **not specified**

pkcli: A Framework for Scripts to Manage Applications

Friday 8 November 2024 09:00 (30 minutes)

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

Primary author: CARLIN, evan (RadiaSoft LLC)**Presenter:** CARLIN, evan (RadiaSoft LLC)**Session Classification:** Basic and end-user IT services**Track Classification:** Basic and End-User IT Services

Contribution ID: 7

Type: **not specified**

Firewall under attack: operational security rollercoaster

Thursday 7 November 2024 14:00 (30 minutes)

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

Primary authors: STORM, Patrick; WARTEL, Romain (CERN)**Presenters:** STORM, Patrick; WARTEL, Romain (CERN)**Session Classification:** Networking & Security**Track Classification:** Networking & Security

Contribution ID: 8

Type: **not specified**

CERN site report

Tuesday 5 November 2024 16:00 (20 minutes)

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

Primary author: SINDRILARU, Elvin Alin (CERN)

Presenter: SINDRILARU, Elvin Alin (CERN)

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 9

Type: **not specified**

PIC report

Monday 4 November 2024 11:40 (20 minutes)

PIC report to HEPHX Fall 2024.

Desired slot length

Speaker release

Yes

Primary author: FLIX MOLINA, Jose (CIEMAT - Centro de Investigaciones Energéticas Medioambientales y Tec. (ES))

Co-authors: BRUZZESE, Agustin; Dr PACHECO PAGES, Andreu (Institut de Física d'Altes Energies - Barcelona (ES)); ALOU, Antoni (PIC); PEREZ-CALERO YZQUIERDO, Antonio (Centro de Investigaciones Energéticas Medioambientales y Tecnológicas); ACOSTA SILVA, Carles (PIC); PEREZ DENGRA, Carlos (PIC); NEISSNER, Christian (PIC); PLANAS, Elena (PIC); GONZALEZ, Elisabeth (PIC); ACCION GARCIA, Esther (The Barcelona Institute of Science and Technology (BIST) (ES)); TORRADEFLOT, Francesc; MERINO, Gonzalo (IFAE - Institute for High Energy Physics); PRIEGO, Jaime (PIC); CASALS HERNANDEZ, Jordi (Port d'Informació Científica (PIC)); NOT SUPPLIED, Jordi Delgado; CARRETERO PALACIOS, Jorge (Port d'Informació Científica); CABAYOL, Laura (PIC); Mrs PORTO, MC (CIEMAT); Prof. DELFINO REZNICEK, Manuel (The Barcelona Institute of Science and Technology (BIST) (ES)); SANTAMARIA RIBA, Marc (UAB); ERIKSEN, Martin Børstad; Mr TALLADA-CRESPÍ, Pau (PIC-CIEMAT); CRUZ JOSA, Ricard (The Barcelona Institute of Science and Technology (BIST) (ES)); GRAÑA, Ruben Dario; ACIN PORTELLA, Vanessa (Institut de Física d'Altes Energies)

Presenter: FLIX MOLINA, Jose (CIEMAT - Centro de Investigaciones Energéticas Medioambientales y Tec. (ES))

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 11

Type: **not specified**

Cost Comparison of On-Premises Storage with S3 Interfaces

Tuesday 5 November 2024 13:30 (30 minutes)

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

Primary author: Mr THOMPSON, Nathan (Spectra Logic)

Co-authors: Mr NINESLING, Matt (Spectra Logic); STARR, Matthew (Spectra Logic Corporation)

Presenter: Mr THOMPSON, Nathan (Spectra Logic)

Session Classification: Storage & Filesystems

Track Classification: Storage & Filesystems

Contribution ID: 12

Type: **not specified**

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

Tuesday 5 November 2024 14:30 (30 minutes)

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

Primary authors: Mr DELEON, Calos (Stony Brook University); Mr LEONARDI, James (Brookhaven National Laboratory); HUANG, Qiulan (Brookhaven National Laboratory (US))

Co-authors: Dr YOO, Shinjae (Brookhaven National Laboratory); Dr GARONNE, Vincent (Brookhaven National Laboratory)

Presenter: HUANG, Qiulan (Brookhaven National Laboratory (US))

Session Classification: Storage & Filesystems

Track Classification: Storage & Filesystems

Contribution ID: 13

Type: **not specified**

How AI networking Fabrics are different from today's Data Center fabrics

Thursday 7 November 2024 17:00 (30 minutes)

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

Primary author: GILBERT, Paul (Arista Networks)**Presenter:** GILBERT, Paul (Arista Networks)**Session Classification:** Networking & Security**Track Classification:** Networking & Security

Contribution ID: 14

Type: **not specified**

Smart Procurement Utility

Thursday 7 November 2024 10:05 (10 minutes)

The Smart Procurement Utility is a tool that allows the visualisation of HEP Score/Watt vs HEP Score/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

Primary author: BRITTON, David (University of Glasgow (GB))**Co-author:** PROTOPOPESCU, Dan (University of Glasgow (GB))**Presenter:** BRITTON, David (University of Glasgow (GB))**Session Classification:** Topical Session: Carbon & Sustainability in Data Centers

Contribution ID: 15

Type: **not specified**

Exploring the Carbon Compromises

Thursday 7 November 2024 09:00 (25 minutes)

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

Primary authors: BORBELY, Albert Gyorgy (University of Glasgow (GB)); BRITTON, David (University of Glasgow (GB)); SIMILLI, Emanuele (University of Glasgow (GB)); STEWART, Gordon; SKIPSEY, Samuel Cadellin; LLOYD, Steve (University of London (GB))

Presenter: BRITTON, David (University of Glasgow (GB))

Session Classification: Topical Session: Carbon & Sustainability in Data Centers

Track Classification: IT Facilities, Business Continuity and Green IT

Contribution ID: 16

Type: **not specified**

Site report for AGLT2

Monday 4 November 2024 11:20 (20 minutes)

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

Primary authors: LAURENS, Philippe (Michigan State University (US)); MC KEE, Shawn (University of Michigan (US)); Dr WU, Wendy (University of Michigan)

Presenters: LAURENS, Philippe (Michigan State University (US)); MC KEE, Shawn (University of Michigan (US)); Dr WU, Wendy (University of Michigan)

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 17

Type: **not specified**

Network tests at CZ Tier-2

Thursday 7 November 2024 16:30 (30 minutes)

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

Primary authors: CHUDOKA, Jiri (Czech Academy of Sciences (CZ)); VOKAC, Petr (Czech Technical University in Prague (CZ))

Presenter: CHUDOKA, Jiri (Czech Academy of Sciences (CZ))

Session Classification: Networking & Security

Track Classification: Networking & Security

Contribution ID: **18**

Type: **not specified**

RAL Site Report

Monday 4 November 2024 13:30 (20 minutes)

An update on activities at the RAL datacentre.

Desired slot length

Speaker release

Yes

Primary author: BLY, Martin (STFC-RAL)

Presenter: BLY, Martin (STFC-RAL)

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 19

Type: **not specified**

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

Thursday 7 November 2024 16:00 (30 minutes)

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

Primary author: BLY, Martin (STFC-RAL)**Co-author:** KELSEY, David (Science and Technology Facilities Council STFC (GB))**Presenter:** BLY, Martin (STFC-RAL)**Session Classification:** Networking & Security**Track Classification:** Networking & Security

Contribution ID: 20

Type: **not specified**

Purdue CMS Analysis Facility

Thursday 7 November 2024 11:00 (30 minutes)

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 Jupyter-Hub 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

Primary authors: KONDRATYEV, Dmitry (Purdue University (US)); NEUMEISTER, Norbert (Purdue University (US)); PIPEROV, Stefan (Purdue University (US))

Presenter: KONDRATYEV, Dmitry (Purdue University (US))

Session Classification: Operating systems, clouds, virtualisation, grids

Track Classification: Operating Systems, Cloud & Virtualisation, Grids

Contribution ID: 21

Type: **not specified**

S3DF: SLAC Shared Science Data Facility

Tuesday 5 November 2024 09:00 (20 minutes)

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

Primary author: ADESANYA, Adeyemi (SLAC)**Presenter:** ADESANYA, Adeyemi (SLAC)**Session Classification:** Site Reports**Track Classification:** Site Reports

Contribution ID: 22

Type: **not specified**

A Cloud-Native Control Plane for Infrastructure and Platform Management

Tuesday 5 November 2024 09:50 (30 minutes)

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

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

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

Session Classification: Operating systems, clouds, virtualisation, grids

Track Classification: Operating Systems, Cloud & Virtualisation, Grids

Contribution ID: 23

Type: **not specified**

AUDITOR: An Accounting tool for Grid Sites and Opportunistic Resources

Monday 4 November 2024 14:40 (30 minutes)

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 multi-purpose 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

Primary author: SCHNEPF, Matthias Jochen**Presenter:** SCHNEPF, Matthias Jochen**Session Classification:** Computing & Batch Services**Track Classification:** Computing & Batch Services

Contribution ID: 24

Type: **not specified**

Networking Topics for WLCG

Thursday 7 November 2024 15:30 (30 minutes)

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 described 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

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

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

Session Classification: Networking & Security

Track Classification: Networking & Security

Contribution ID: 25

Type: **not specified**

HEPiX Benchmarking WG: Status Report

Monday 4 November 2024 15:45 (30 minutes)

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 HEPsScore 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

Primary author: SCHNEPF, Matthias Jochen

Presenter: SCHNEPF, Matthias Jochen

Session Classification: Computing & Batch Services

Contribution ID: 26

Type: **not specified**

Jefferson Lab Site Report and HPDF Introduction

Tuesday 5 November 2024 16:20 (20 minutes)

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

Primary author: HESS, Bryan

Presenter: HESS, Bryan

Session Classification: Site Reports

Track Classification: Site Reports

Contribution ID: 27

Type: **not specified**

HELP! I have DataCenter Nightmares

Thursday 7 November 2024 14:30 (30 minutes)

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

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

Contribution ID: 28

Type: **not specified**

Shifting Hardware Landscape

Friday 8 November 2024 11:00 (30 minutes)

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

Primary author: MISAWA, Shigeki (Brookhaven National Laboratory (US))**Presenter:** MISAWA, Shigeki (Brookhaven National Laboratory (US))**Session Classification:** Miscellaneous

Contribution ID: 32

Type: **not specified**

Science Cloud based on WLCG Core Technology

Tuesday 5 November 2024 09:20 (30 minutes)

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

Primary author: YEN, Eric (Academia Sinica (TW))

Presenter: YEN, Eric (Academia Sinica (TW))

Session Classification: Operating systems, clouds, virtualisation, grids

Contribution ID: 33

Type: **not specified**

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

Monday 4 November 2024 16:15 (45 minutes)

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

Desired slot length

45

Speaker release

Yes

Primary author: Prof. FAGG, Anderw (University of Oklahoma)

Presenter: Prof. FAGG, Anderw (University of Oklahoma)

Session Classification: Miscellaneous

Contribution ID: 34

Type: **not specified**

dCache on Kubernetes

Tuesday 5 November 2024 11:15 (30 minutes)

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

Primary author: OGGIAN, Elia Luca (ETH Zurich (CH))

Presenter: OGGIAN, Elia Luca (ETH Zurich (CH))

Session Classification: Operating systems, clouds, virtualisation, grids

Track Classification: Operating Systems, Cloud & Virtualisation, Grids

Contribution ID: 35

Type: **not specified**

Operating the 200 Gbps IRIS-HEP Demonstrator for ATLAS

Monday 4 November 2024 14:10 (30 minutes)

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

Primary authors: HELD, Alexander (University of Wisconsin Madison (US)); JORDAN, David (University of Chicago (US)); BENJAMIN, Doug (Brookhaven National Laboratory (US)); GOLNARAGHI, Farnaz (University of Chicago (US)); HU, Fengping (University of Chicago (US)); WATTS, Gordon (University of Washington (US)); VUKOTIC, Ilija (University of Chicago (US)); STEPHEN, Judith Lorraine (University of Chicago (US)); BRYANT, Lincoln (University of Chicago (US)); FEICKERT, Matthew (University of Wisconsin Madison (US)); RIND, Ofer (Brookhaven National Laboratory); GARDNER JR, Robert William (University of Chicago (US)); YANG, Wei (SLAC National Accelerator Laboratory (US))

Presenter: JORDAN, David (University of Chicago (US))

Session Classification: Computing & Batch Services

Contribution ID: 36

Type: **not specified**

Case Study: AI Training Power Demand on a GPU-Accelerated Node

Thursday 7 November 2024 09:45 (20 minutes)

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

Primary author: LATIF, Imran (Brookhaven National Laboratory)**Co-author:** MISAWA, Shigeki (Brookhaven National Laboratory (US))**Presenters:** LATIF, Imran (Brookhaven National Laboratory); MISAWA, Shigeki (Brookhaven National Laboratory (US))**Session Classification:** Topical Session: Carbon & Sustainability in Data Centers**Track Classification:** IT Facilities, Business Continuity and Green IT

Contribution ID: 37

Type: **not specified**

Stories from the TSM to HPSS Migration at KIT

Tuesday 5 November 2024 15:30 (30 minutes)

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

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

Co-authors: GOTTMANN, Artur Il Darovic (KIT - Karlsruhe Institute of Technology (DE)); Mr LOBONTU, Dorin-Daniel; RESSMANN, Doris; MUSHEGHYAN, Haykuhi (Georg August Universitaet Goettingen (DE)); SCHAEFER, Karin; KONSTANTINOV, Preslav (KIT - Karlsruhe Institute of Technology (DE)); AMBROJ PEREZ, Samuel; MOL, Xavier (KIT - Karlsruhe Institute of Technology (DE))

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

Session Classification: Storage & Filesystems

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

Transitioning from RHEV to Openshift

Thursday 7 November 2024 11:30 (30 minutes)

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

Primary author: Mr TANG, Zeyi (Brookhaven National Laboratory)**Co-author:** HANCOCK, Robert**Presenter:** HANCOCK, Robert**Session Classification:** Operating systems, clouds, virtualisation, grids

Contribution ID: 39

Type: **not specified**

Welcome and Workshop Logistics

Monday 4 November 2024 09:15 (10 minutes)

Presenter: SEVERINI, Horst (University of Oklahoma (US))

Session Classification: Welcome and Workshop Logistics

Contribution ID: 40

Type: **not specified**

Research Computing and Storage Strategy at the University of Oklahoma

Monday 4 November 2024 09:25 (45 minutes)

Presenters: NEEMAN, Henry (OU); SEVERINI, Horst (University of Oklahoma (US))

Session Classification: Welcome and Workshop Logistics

Contribution ID: 41

Type: **not specified**

Welcome Address from the Physics Department

Monday 4 November 2024 10:10 (20 minutes)

Presenters: GUTIERREZ, Phillip (University of Oklahoma (US)); SEVERINI, Horst (University of Oklahoma (US))

Session Classification: Welcome and Workshop Logistics

Contribution ID: 42

Type: **not specified**

Summary of the Joint Xrootd and FTS Workshop

Tuesday 5 November 2024 11:45 (30 minutes)

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

Primary author: YANG, Wei (SLAC National Accelerator Laboratory (US))**Co-author:** SEVERINI, Horst (University of Oklahoma (US))**Presenter:** YANG, Wei (SLAC National Accelerator Laboratory (US))**Session Classification:** Operating systems, clouds, virtualisation, grids

Contribution ID: 43

Type: **not specified**

IHEP Site Report

Tuesday 5 November 2024 16:40 (20 minutes)

The progress and status of IHEP site since last Hepix.

Desired slot length

Speaker release

Yes

Primary author: HOU, Siqi

Presenter: HOU, Siqi

Session Classification: Site Reports

Contribution ID: 44

Type: **not specified**

SWT2 site report

Monday 4 November 2024 11:00 (20 minutes)

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

Primary author: BOOTH, Zachary (University of Texas at Arlington)**Co-authors:** VARTAPETIAN, Armen (University of Texas at Arlington); SEVERINI, Horst (University of Oklahoma); SOSEBEE, Mark (University of Texas at Arlington)**Presenter:** BOOTH, Zachary (University of Texas at Arlington)**Session Classification:** Site Reports

Contribution ID: 45

Type: **not specified**

Nebraska Coffea-Casa Analysis Facility Update

Tuesday 5 November 2024 10:45 (30 minutes)

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

Primary author: ATTEBURY, Garhan (University of Nebraska Lincoln (US))**Presenter:** ATTEBURY, Garhan (University of Nebraska Lincoln (US))**Session Classification:** Operating systems, clouds, virtualisation, grids

Contribution ID: 46

Type: **not specified**

CRISP: Collaborative Tools for the ePIC Experiment

Friday 8 November 2024 10:00 (30 minutes)

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

Primary authors: ARKHIPKIN, Dmitry; LANCON, Eric (Brookhaven National Laboratory (US)); Dr LAURET, Jerome (Brookhaven National Laboratory); RIND, Ofer (Brookhaven National Laboratory); STEINBERG, Peter Alan (Brookhaven National Laboratory (US)); HANCOCK, Robert; GANAPATHY, Uma (BNL); GARONNE, Vincent (Brookhaven National Laboratory)

Presenter: RIND, Ofer (Brookhaven National Laboratory)

Session Classification: Basic and end-user IT services

Contribution ID: 47

Type: **not specified**

Carbon costs of storage: a UK perspective

Thursday 7 November 2024 09:25 (20 minutes)

In order to achieve the higher performance year on year required by the 2030s for future LHC up-grades 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

Primary authors: PACKER, Alison; PACKER, Alison (STFC - Science & Technology Facilities Council (GB)); SKIPSEY, Samuel Cadellin

Presenter: SKIPSEY, Samuel Cadellin

Session Classification: Topical Session: Carbon & Sustainability in Data Centers

Contribution ID: 48

Type: **not specified**

Discussion

Thursday 7 November 2024 10:25 (5 minutes)

Session Classification: Topical Session: Carbon & Sustainability in Data Centers

Contribution ID: 49

Type: **not specified**

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

Thursday 7 November 2024 10:15 (10 minutes)

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

Primary author: FLIX MOLINA, Jose (CIEMAT - Centro de Investigaciones Energéticas Medioambientales y Tec. (ES))

Co-authors: ACOSTA SILVA, Carles (PIC); MERINO, Gonzalo (IFAE - Institute for High Energy Physics); CASALS HERNANDEZ, Jordi (Port d'Informació Científica (PIC)); FABREGA, Kevin (UAB)

Presenter: FLIX MOLINA, Jose (CIEMAT - Centro de Investigaciones Energéticas Medioambientales y Tec. (ES))

Session Classification: Topical Session: Carbon & Sustainability in Data Centers

Track Classification: IT Facilities, Business Continuity and Green IT

Contribution ID: 50

Type: **not specified**

Impetus and Drivers

Wednesday 6 November 2024 09:00 (15 minutes)

Presenter: Dr CROOKS, David (UKRI STFC)

Session Classification: Topical Session: Security Operations Center (SOC)

Contribution ID: 51

Type: **not specified**

Mini Intro: Zeek

Wednesday 6 November 2024 09:15 (10 minutes)

Presenter: SHARMA, Aashish (LBNL)

Session Classification: Topical Session: Security Operations Center (SOC)

Contribution ID: 52

Type: **not specified**

Mini-intro: MISP

Wednesday 6 November 2024 09:25 (10 minutes)

Presenter: ACRIS, James (STFC)

Session Classification: Topical Session: Security Operations Center (SOC)

Contribution ID: 53

Type: **not specified**

Mini-Intro: pDNSSOC

Wednesday 6 November 2024 09:35 (10 minutes)

Presenter: WARTEL, Romain (CERN)

Session Classification: Topical Session: Security Operations Center (SOC)

Contribution ID: 54

Type: **not specified**

WLCG SOC at U Chicago

Wednesday 6 November 2024 09:45 (15 minutes)

Presenter: JORDAN, David (University of Chicago (US))

Session Classification: Topical Session: Security Operations Center (SOC)

Contribution ID: 55

Type: **not specified**

Discussion

Wednesday 6 November 2024 10:00 (30 minutes)

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.

Presenters: Dr CROOKS, David (UKRI STFC); VALSAN, Liviu (CERN)

Session Classification: Topical Session: Security Operations Center (SOC)

Contribution ID: 56

Type: **not specified**

Topic 1: Issues capturing ALL traffic with Zeek?

Wednesday 6 November 2024 13:30 (30 minutes)

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?

Presenters: SHARMA, Aashish (LBNL); Dr CROOKS, David (UKRI STFC); JORDAN, David (University of Chicago (US)); MC KEE, Shawn (University of Michigan (US))

Session Classification: SOC Hackathon

Contribution ID: 57

Type: **not specified**

Topic 2: Building a good Zeek alert

Wednesday 6 November 2024 14:00 (30 minutes)

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?

Presenters: SHARMA, Aashish (LBNL); Dr CROOKS, David (UKRI STFC); WARTEL, Romain (CERN)

Session Classification: SOC Hackathon

Contribution ID: 58

Type: **not specified**

Topic 1: Configuring sending alerts for various tools

Wednesday 6 November 2024 15:30 (30 minutes)

How to enable alerts using webhooks and various applications.

Sending to SLACK

Sending to Mattermost

What about Keybase?

Why not email?

Presenters: Dr CROOKS, David (UKRI STFC); ATHERTON, Liam; WARTEL, Romain (CERN)

Session Classification: SOC Hackathon

Contribution ID: 59

Type: **not specified**

Topic 3: How to deploy pDNSSOC (part 1)

Wednesday 6 November 2024 14:30 (30 minutes)

How to deploy pDNSSOC

Example deployment

Working session

Presenter: WARTEL, Romain (CERN)

Session Classification: SOC Hackathon

Contribution ID: **60**

Type: **not specified**

Tool options and choices

Wednesday 6 November 2024 16:00 (30 minutes)

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

Advantages, capabilities, limitations, concerns....

Let's discuss

Presenters: Dr CROOKS, David (UKRI STFC); LUEDERS, Stefan (CERN)

Session Classification: SOC Hackathon