

HEPiX Autumn 2020 online Workshop



Report of Contributions

Contribution ID: 1

Type: **not specified**

CERN Site Report

Monday, 12 October 2020 09:15 (20 minutes)

News from CERN since the previous HEPiX workshop.

Primary author: DUMITRU, Andrei (CERN)

Presenter: DUMITRU, Andrei (CERN)

Session Classification: Monday morning

Track Classification: Site Reports

Contribution ID: 2

Type: **not specified**

PIC report

Monday, 12 October 2020 17:30 (15 minutes)

This is the PIC report for HEPiX Autumn 2020 Workshop

Primary author: FLIX MOLINA, Jose (Centro de Investigaciones Energéticas Medioambientales y Tecnológico)

Co-authors: Dr MERINO, Gonzalo (IFAE - Institute for High Energy Physics); PORTO FERNANDEZ, Maria Del Carmen (Unknown); ACIN PORTELLA, Vanessa (Institut de Física d'Altes Energies); ACCION GARCIA, Esther (The Barcelona Institute of Science and Technology (BIST) (ES)); ACOSTA SILVA, Carlos (PIC); CRUZ JOSA, Ricard (The Barcelona Institute of Science and Technology (BIST) (ES)); PLANAS, Elena (PIC); RODRÍGUEZ, Bruno; CASALS HERNANDEZ, Jordi (Port d'Informació Científica); PEREZ-CALERO YZQUIERDO, Antonio (Centro de Investigaciones Energéticas Medioambientales y Tecnológico); PACHECO PAGES, Andreu (Institut de Física d'Altes Energies - Barcelona (ES)); PEREZ DENGRA, Carlos (PIC); NEISSNER, Christian (PIC); Prof. DELFINO REZNICEK, Manuel (The Barcelona Institute of Science and Technology (BIST) (ES)); CARRETERO PALACIOS, Jorge (Institute for High Energy Physics (IFAE)); DELGADO, Jordi (PIC); Mr TALLADA CRESPI, Pau (PIC-CIEMAT); TORRADEFLOT, Francesc (PIC); BRUZZESE, Agustín (PIC)

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

Session Classification: Monday evening

Track Classification: Site Reports

Contribution ID: 4

Type: **not specified**

Diamond Light Source Site Report

Monday, 12 October 2020 17:45 (15 minutes)

Diamond Light Source is a synchrotron light source in Oxfordshire, UK. We will present the developments at Diamond since HEPiX Autumn 2019 and our plans for the future.

Primary author: Mr THORNE, James (Diamond Light Source)

Presenter: Mr THORNE, James (Diamond Light Source)

Session Classification: Monday evening

Track Classification: Site Reports

Contribution ID: 5

Type: **not specified**

Welcome & Logistics

Monday, 12 October 2020 09:00 (15 minutes)

Presenter: VAN DER REEST, Peter (DESY)

Session Classification: Monday morning

Track Classification: Miscellaneous

Contribution ID: 14

Type: **not specified**

Welcome & Logistics

Monday, 12 October 2020 17:00 (15 minutes)

Presenter: VAN DER REEST, Peter (DESY)

Session Classification: Monday evening

Track Classification: Miscellaneous

Contribution ID: 23

Type: **not specified**

DESY Site Report

Monday, 12 October 2020 18:00 (15 minutes)

News from the center since Autumn 2019

Primary authors: VAN DER REEST, Peter (DESY); HAUPT, Andreas (Deutsches Elektronen-Synchrotron (DESY))

Presenters: VAN DER REEST, Peter (DESY); HAUPT, Andreas (Deutsches Elektronen-Synchrotron (DESY))

Session Classification: Monday evening

Track Classification: Site Reports

Contribution ID: 24

Type: **not specified**

Learning-based Approaches to Estimate Job Wait Time in HTC Datacenters

Tuesday, 13 October 2020 11:40 (20 minutes)

High Throughput Computing (HTC) datacenters are a cornerstone of scientific discoveries in the fields of High Energy Physics and Astroparticles Physics. These datacenters provide thousands of users from dozens scientific collaborations with tens of thousands computing cores and Petabytes of storage. The scheduling algorithm used in such datacenters to handle the millions of (mostly single-core) jobs submitted every month ensures a fair sharing of the computing resources among user groups, but may also cause unpredictably long job wait times for some users. The time a job will wait can be caused by many entangled factors and configuration parameters and is thus very hard to predict. Moreover, batch systems implementing a fair-share scheduling algorithm cannot provide users with any estimation of the job wait time at submission time.

Therefore, in this talk we investigate how learning-based techniques applied to the logs of the batch scheduling system of a large HTC datacenter can be used to get such an estimation of job wait time. After having illustrated the need for users to get an estimation of the time their jobs will wait, we identify some intuitive causes of this wait time based on the analysis of the information found in the batch system logs. Then, we formally analyze the correlation between these intuitive causes and job wait time and propose learning-based estimators of both job wait time and job wait time ranges. We conclude by presenting the obtained preliminary results and thoughts about how to deploy the proposed estimators in production.

Primary authors: Mr GOMBERT, Luc (CC-IN2P3 / CNRS); Mr SUTER, Frederic (CC-IN2P3 / CNRS)

Presenter: Mr SUTER, Frederic (CC-IN2P3 / CNRS)

Session Classification: Computing and Batch Services

Track Classification: Computing & Batch Services

Contribution ID: 25

Type: **not specified**

Benchmarking Working Group: A status report

Tuesday, 13 October 2020 17:00 (20 minutes)

The HEPiX benchmarking working group has been very active in the past months and will present a report of the activities, concentrating on the new developments after the last report. The new candidate benchmark, designed to overcome the HEP-SPEC06 problems is being published. The HEP Benchmarking Suite has been complemented by GPU and Root benchmark.

Primary authors: Dr MICHELOTTO, Michele (Universita e INFN, Padova (IT)); GIORDANO, Domenico (CERN); ALEF, Manfred (Karlsruhe Institute of Technology (KIT))

Presenter: GIORDANO, Domenico (CERN)

Session Classification: Tuesday evening

Track Classification: Computing & Batch Services

Contribution ID: 26

Type: **not specified**

Computer Security Update

Tuesday, 13 October 2020 09:00 (20 minutes)

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

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

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.

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

Primary author: FILIPPAKIS, Nikolaos (CERN)

Presenter: FILIPPAKIS, Nikolaos (CERN)

Session Classification: Tuesday morning

Track Classification: Networking & Security

Contribution ID: 27

Type: **not specified**

Anomaly detection for the Centralized Elasticsearch service at CERN

Thursday, 15 October 2020 09:20 (20 minutes)

For several years CERN offers a centralized service for Elasticsearch. This dynamic infrastructure consists of currently about 30 independent Elasticsearch clusters, covering more than 180 different use cases. Using internal monitoring data, a real time anomaly detection system has been implemented, and is now used in production. This presentation describes how the system works, the experiences gained so far, and the lessons learned from this exercise.

Primary author: Dr SCHWICKERATH, Ulrich (CERN)

Presenter: Dr SCHWICKERATH, Ulrich (CERN)

Session Classification: Thursday morning

Track Classification: IT Facilities & Business Continuity

Contribution ID: 28

Type: **not specified**

IHEP Site Report

Monday, 12 October 2020 09:35 (15 minutes)

News and updates from IHEP since the last HEPiX Workshop. In this talk we would like to present the status of IHEP site including computing farm, HPC, Grid, data storage ,network and so on.

Primary author: LI, Haibo (Institute of High Energy Physics Chinese Academy of Science)

Presenter: LI, Haibo (Institute of High Energy Physics Chinese Academy of Science)

Session Classification: Monday morning

Track Classification: Site Reports

Contribution ID: 30

Type: **not specified**

Simulated phishing campaigns at CERN

Tuesday, 13 October 2020 09:20 (20 minutes)

Since years, e-mail is one of the main attack vectors that organisations and individuals face. Malicious actors use e-mail messages to run phishing attacks, to distribute malware, and to send around various types of scams. While technical solutions exist to filter out most of such messages, no mechanism can guarantee 100% efficiency. E-mail recipients themselves are the next, crucial layer of protection - but unfortunately, they fall for the various tricks used by attackers way too often.

In order to raise awareness and to educate CERN community, CERN Computer Security Team runs regular simulated phishing campaigns. This talk will discuss the motivation behind this activity, various techniques used, as well as the results and lessons learnt. Finally, CERN campaigns will be compared to those run by other organisations, and to available commercial solutions.

Primary author: LOPIENSKI, Sebastian (CERN)

Presenter: LOPIENSKI, Sebastian (CERN)

Session Classification: Tuesday morning

Track Classification: Networking & Security

Contribution ID: 31

Type: **not specified**

WLCG/OSG Network Activities, Status and Plans

Tuesday, 13 October 2020 18:05 (20 minutes)

WLCG relies on the network as a critical part of its infrastructure and therefore needs to guarantee effective network usage and prompt detection and resolution of any network issues, including connection failures, congestion and traffic routing. The OSG Networking Area is a partner of the WLCG effort and is focused on being the primary source of networking information for its partners and constituents. We will report on the changes and updates that have occurred since the last HEPiX meeting. The primary areas to cover include the status of and plans for the WLCG/OSG perfSONAR infrastructure, the WLCG Network Throughput Working Group and the activities in the IRIS-HEP and SAND projects.

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

Presenter: BABIK, Marian (CERN)

Session Classification: Networking and Security

Track Classification: Networking & Security

Contribution ID: 32

Type: **not specified**

Network Functions Virtualisation Working Group Update

Tuesday, 13 October 2020 18:25 (20 minutes)

As the scale and complexity of the current HEP network grows rapidly, new technologies and platforms are being introduced that greatly extend the capabilities of today's networks. With many of these technologies becoming available, it's important to understand how we can design, test and develop systems that could enter existing production workflows while at the same time changing something as fundamental as the network that all sites and experiments rely upon. In this talk we'll give an update on the working group's recent activities, updates from sites and R&E network providers as well as plans for the near-term future.

In particular we'll focus on the white paper surveying number of existing technologies and tools that was written by NFW WG and we are going to discuss possible future directions as well as give an overview of the already existing efforts that are part of the Research Networking Technical WG.

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

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

Session Classification: Networking and Security

Track Classification: Networking & Security

Contribution ID: 33

Type: **not specified**

CERN Cloud Infrastructure status update

Thursday, 15 October 2020 10:10 (20 minutes)

CERN's private OpenStack cloud offers more than 300,000 cores to over 3,500 users with services for compute, multiple storage types, baremetal, container clusters, and more.

The cloud supports CERN's web and administration services and is the tier 0 site in the WLCG.

This update will cover the evolution of the cloud over the past year, and the plans for the upcoming year.

Primary author: Mr HARTLAND, Thomas George (CERN)

Presenter: Mr HARTLAND, Thomas George (CERN)

Session Classification: Grids, Clouds and Virtualisation

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: 34

Type: **not specified**

Kubernetes in the CERN Cloud

Thursday, 15 October 2020 10:50 (20 minutes)

The Magnum component of OpenStack is used to provision container orchestration clusters in the CERN cloud, with Kubernetes being by far the most popular cluster type.

This presentation will look at the new features in Magnum and Kubernetes which make it possible to create highly available Kubernetes clusters that are suitable for hosting critical services.

This includes node groups for splitting nodes across multiple availability zones, software-defined networking loadbalancers for cluster ingress and to enable multi-master clusters, pod and cluster autoscaling for dynamically responding to increased load on a service, and other features.

By combining these features the availability and stability that is possible for a cluster is high enough for us to start running parts of the OpenStack control plane on Kubernetes.

Primary author: Mr HARTLAND, Thomas George (CERN)

Presenter: Mr HARTLAND, Thomas George (CERN)

Session Classification: Grids, Clouds and Virtualisation

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: 35

Type: **not specified**

EOS storage for Alice O2

Wednesday, 14 October 2020 09:00 (20 minutes)

In this contribution we report on the ongoing R&D activity aiming at preparing the EOS ALICE O2 storage cluster for the very demanding requirements of Run 3. After the planned upgrades of LHC and ALICE detectors, the ALICE experiment is expected to increase the data-taking rate handled by the online system and then recorded into permanent storage by one order of magnitude. In order to accommodate the data sent by the ALICE Data Acquisition system (aggregated throughput of 100GB/s), the EOS ALICE O2 cluster will be equipped with one of the latest generation storage servers and configured to apply erasure coding to the registered data.

During this talk we are going to present the latest tests and experiments we conducted on the storage server cluster in the context of ALICE O2 project. In particular we give an overview of the aggregate throughput tests with EOS native erasure coding, as well as with client side erasure coding. In addition, we discuss the tuning of the storage servers including testing different kernel versions and firewall settings that enables utilizing the hardware at its nominal speed.

Primary authors: SIMON, Michal Kamil (CERN); PETERS, Andreas Joachim (CERN); CONTESCU, Cristian (CERN); MASCETTI, Luca (CERN)

Presenter: SIMON, Michal Kamil (CERN)

Session Classification: Wednesday morning

Track Classification: Storage & Filesystems

Contribution ID: 36

Type: **not specified**

The design of networking and computing system for High Energy Photon Source(HEPS)

Tuesday, 13 October 2020 10:40 (20 minutes)

High Energy Photon Source (HEPS) is the first national high-energy synchrotron radiation light source in Beijing China, and will be ready for users and scientists in 2025. According to the estimated data rates, we predict 30 PB raw experimental data will be produced per month from 14 beamlines at the first stage of HEPS, and the data volume will be even greater after over 90 beamlines are completed at the second stage in the near future.

This report will introduce the design of networking, computing and data service system for HEPS.

Primary author: QI, Fazhi (Chinese Academy of Sciences (CN))

Co-authors: Dr ZHAO, Haifeng (IHEP); Ms HU, Hao (IHEP); Dr TIAN, Haolai (IHEP); Mrs ZHANG, Hongmei (IHEP); Dr WANG, Lu (IHEP); Dr HUANG, Qiulan (IHEP); Dr ZENG, Shan (IHEP)

Presenters: QI, Fazhi (Chinese Academy of Sciences (CN)); Dr HUANG, Qiulan (IHEP)

Session Classification: Miscellaneous

Track Classification: Miscellaneous

Contribution ID: 37

Type: **not specified**

BNL Site Report

Monday, 12 October 2020 17:15 (15 minutes)

An update on BNL activities since the Fall 2019 workshop

Primary author: RAO, Tejas

Presenter: RAO, Tejas

Session Classification: Monday evening

Track Classification: Site Reports

Contribution ID: **38**

Type: **not specified**

CERN's Business Continuity Working Group

Thursday, 15 October 2020 09:00 (20 minutes)

In April 2020 CERN has formed a working group on business continuity. In this presentation, we will describe the mandate and the direction of the group, and will discuss some of the tools used. Areas for potential collaboration with other labs will be discussed as well.

Primary author: MEINHARD, Helge (CERN)

Presenter: MEINHARD, Helge (CERN)

Session Classification: Thursday morning

Track Classification: IT Facilities & Business Continuity

Contribution ID: 39

Type: **not specified**

CERNphone –CERN’s upcoming softphone solution

Wednesday, 14 October 2020 19:30 (20 minutes)

CERNphone is the new softphone-based solution that will gradually be deployed across CERN. Based on open-source components and developments, CERNphone provides mobile and desktop clients and back-end services as a replacement for legacy hard phones and commercial PBX systems. In this contribution, we will describe the architecture and main components of CERNphone, discuss the main challenges for delivering reliable mobile and secure SIP clients, and provide an outlook for its deployment at CERN.

Primary author: CANCIO, German (CERN)

Presenter: CANCIO, German (CERN)

Session Classification: Basic IT Services

Track Classification: Basic IT Services

Contribution ID: 40

Type: **not specified**

FTS: Towards tokens, QoS, archive monitoring and beyond

Wednesday, 14 October 2020 10:20 (20 minutes)

The File Transfer Service (FTS) is a fundamental component for the LHC experiments, distributing the majority of the LHC data across the WLCG infrastructure. Tightly integrated with experiment frameworks, it has transferred more than 1 billion files and a total of 950 petabytes of data in 2019 alone. With more than 30 experiments using FTS at CERN and outside, it has steadily gained popularity in data-intensive sciences.

Playing a crucial role in data distribution, FTS is constantly evolving in preparation for LHC RUN-3 and forward. With the 2018 participation in the EU-funded project eXtreme Data Cloud (XDC) and continuous involvement within the WLCG DOMA TPC and QoS working groups, a series of developments have been performed in order to meet the requirements of the LHC experiments and community alike.

This presentation will provide a detailed overview of activities carried out for the upcoming 3.10 release, focusing on OpenID Connect (OIDC) token support, QoS functionality, Third Party Copy (TPC) support for XRootD and HTTP protocols, archive monitoring for the new CERN Tape Archive (CTA) system, service scalability improvements and the future direction of FTS.

Primary author: PATRASCOIU, Mihai (CERN)

Co-authors: Dr KARAVAKIS, Edward (CERN); VOINEAG, Cristina (Polytechnic University of Bucharest); GARCIA CABOT, Carles (CERN); MANZI, Andrea

Presenter: PATRASCOIU, Mihai (CERN)

Session Classification: Wednesday morning

Track Classification: Storage & Filesystems

Contribution ID: 41

Type: **not specified**

Roadmap for the DNS Load Balancing Service at CERN

Tuesday, 13 October 2020 09:40 (20 minutes)

This presentation delivers a holistic view of the current and future state of the DNS Load Balancing service at CERN. CERN runs this service to provide the necessary tools for managing the nodes that aliases should present. The service contains three main components:

1. An administrative interface, where users can define aliases and their policies.
2. A client that runs on the nodes to assess their healthiness.
3. An arbiter which, based on the information from the clients, selects the best nodes according to the aliases' policies.

Currently, the service is under migration from Python to Golang. The aim is to improve the parallelism of the service and to decrease the latency of the checks. At the same time, the service is also under migration from a VM-oriented implementation to a Cloud-Native one. The final goal is for this setup to provide a more scalable and reliable system.

Primary author: Mr KOUROS, Kristian (VI Trainee)

Co-authors: SAIZ, Pablo (CERN); REGUERO, Ignacio (CERN)

Presenter: Mr KOUROS, Kristian (VI Trainee)

Session Classification: Tuesday morning

Track Classification: Networking & Security

Contribution ID: 42

Type: **not specified**

Federated Identity Management at BNL

Tuesday, 13 October 2020 19:05 (20 minutes)

We will report on recent activities on integrating Federated Identity Management at the Scientific Data & Computing Center (SDCC).

Primary author: MISAWA, Shigeki (Brookhaven National Laboratory (US))

Presenter: MISAWA, Shigeki (Brookhaven National Laboratory (US))

Session Classification: Networking and Security

Track Classification: Networking & Security

Contribution ID: 43

Type: **not specified**

An Evaluation of Podman

Thursday, 15 October 2020 17:30 (20 minutes)

Podman is the default container execution platform shipped with RHEL/CentOS 8, and is now also available in RHEL/CentOS/SL 7. It provides a user command line interface that is effectively identical to Docker's, and supports rootless container execution. In this talk we'll give an overview of Podman, our experiences with this software, and a comparison with Singularity and Docker.

Primary authors: HOLLOWELL, Christopher Henry (Brookhaven National Laboratory (US)); CARA-MARCU, Costin (Brookhaven National Laboratory (US))

Presenter: HOLLOWELL, Christopher Henry (Brookhaven National Laboratory (US))

Session Classification: Thursday evening

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: 44

Type: **not specified**

HPSS migration to IBM tape technologies

Wednesday, 14 October 2020 17:00 (20 minutes)

The Scientific Data and Computing Center is migrating part of its 200 PB tape-resident data archive to new high-density robotic libraries. The talk will focus on the product/vendor evaluation process and the complex of decision criteria observed. We will also discuss the potential implications of moving from high-performance to high-density robotic systems for hosting active data archives.

Primary author: Mr CHOU, Tim (Brookhaven National Laboratory (US))

Presenter: Mr CHOU, Tim (Brookhaven National Laboratory (US))

Session Classification: Wednesday evening

Track Classification: Storage & Filesystems

Contribution ID: 45

Type: **not specified**

Development of online application portal

Monday, 12 October 2020 11:20 (20 minutes)

The KEK Computing Research Center provides various IT services, such as email, WiFi network, data analysis system, etc. A user submits paper application forms to apply those services. Moreover, some services require endorsement by a KEK staff. From the Computing Research Center point of view, we receive applications from many users every day. These forms are checked, processed, and handed out among the responsible persons to complete applications.

To improve these circumstances, we have developed an online application portal system named “ccPortal.” This portal allows users to apply new services and modify account information online. The portal also supports the center staff’s online workflow, checking validity, creating an account, DB registration, and sending a completion notice. The ccPortal service was rolled out to KEK staff from the middle of September and gradually opened to all others by November.

We will describe the portal service and share our experience of improving work efficiency in IT service operations.

Primary authors: TAKASE, Wataru (High Energy Accelerator Research Organization (JP)); OKADA, Shogo; KANEKO, Sari (KEK); MURAKAMI, Koichi (KEK)

Presenter: TAKASE, Wataru (High Energy Accelerator Research Organization (JP))

Session Classification: End-user Services, Operating Systems

Track Classification: End-User IT Services & Operating Systems

Contribution ID: 46

Type: **not specified**

IGWN submit node: Designing EU-based gateway nodes for distributed computing for Virgo, LIGO, Kagra

Tuesday, 13 October 2020 11:00 (20 minutes)

While there are a handful of International Gravitational-Wave Observatory Network (IGWN) submit nodes deployed in the US for LIGO, Virgo, Kagra (aka IGWN) data-processing pipelines, Nikhef has worked with LIGO and Virgo collaborators to design a submit node that can be deployed at EU sites. More EU-based submit nodes will allow additional points of entry for IGWN computing resources and make the resources more readily accessible for the scientists. This talk will discuss some of the design choices that made the Nikhef IGWN submit node different from the US nodes and how this fits in with the data lifecycle for IGWN.

Primary author: Ms HESTER, Mary (Nikhef)

Presenter: Ms HESTER, Mary (Nikhef)

Session Classification: Miscellaneous

Track Classification: Miscellaneous

Contribution ID: 47

Type: **not specified**

MySQL High Availability in the Database On Demand service

Monday, 12 October 2020 19:20 (20 minutes)

The DBOD Service is a Database as a Service platform that provides MySQL, PostgreSQL, and InfluxDB database instances to CERN users.

During the last few years, more and more critical services have moved to open source database solutions and are now making use of the DBOD service. As a consequence, high availability is expected for some of these services.

This presentation describes how to achieve High Availability for MySQL databases using ProxySQL. I will also present other open source alternatives for MySQL HA.

Primary author: CABEZAS ALONSO, Abel

Presenter: CABEZAS ALONSO, Abel

Session Classification: End-user Services, Operating Systems

Track Classification: End-User IT Services & Operating Systems

Contribution ID: 48

Type: **not specified**

Towards a redundant, robust, secure and reliable IoT network

Wednesday, 14 October 2020 18:50 (20 minutes)

The interest in the Internet of Things (IoT) is growing exponentially so multiple technologies and solutions have emerged to connect mostly everything. A 'thing' can be a car, a thermometer or a robot that, when equipped with a transceiver, will exchange information over the internet with a defined service. Therefore, IoT comprises a wide variety of user cases with very different requirements.

After having studied various Low-Power Wide Area Network (LPWAN) protocols CERN finally selected the Long Range Wide-Area (LoRa) network as base protocol which marked the start of the establishment of an IoT network at CERN. In order to build a functioning basic infrastructure LoRa gateways and a network server have been selected. With these basic components it was possible to build a simple IoT network. Unfortunately this configuration is not suited for a productive environment as it is not reliable enough.

To improve the reliability and security of the network it is necessary to identify the weaknesses of the technology used as well of those of the current infrastructure. Furthermore, it must be determined which of these vulnerabilities can be improved and which cannot.

CERN is currently running two projects that comprise several thousand end devices and thus put the IoT infrastructure to the test.

Primary author: MERSCHER, Christoph (CERN)

Presenter: MERSCHER, Christoph (CERN)

Session Classification: Basic IT Services

Track Classification: Basic IT Services

Contribution ID: 49

Type: **not specified**

KEK Site Report

Monday, 12 October 2020 09:50 (15 minutes)

KEK Central Computer System (KEKCC) has been upgraded by the full-scale hardware replacement and started the operation in September 2020. In this report, we present the specifications of the new KEKCC with the usage achievement of the previous system last year.

Primary author: NAKAMURA, Tomoaki (High Energy Accelerator Research Organization (JP))

Co-authors: IWAI, Go (KEK); MURAKAMI, Koichi; SASAKI, Takashi (High Energy Accelerator Research Organization (JP)); SUZUKI, Soh; TAKASE, Wataru (High Energy Accelerator Research Organization (JP))

Presenter: NAKAMURA, Tomoaki (High Energy Accelerator Research Organization (JP))

Session Classification: Monday morning

Track Classification: Site Reports

Contribution ID: 50

Type: **not specified**

Operating a production HTCondor cluster: Seamlessly automating maintenance, OS and HTCondor updates with (almost) zero downtimes

Tuesday, 13 October 2020 11:20 (20 minutes)

Our HTC cluster using HTCondor has been set up at Bonn University in 2017/2018. All infrastructure is fully puppetised, including the HTCondor configuration.

OS updates are fully automated, and necessary reboots for security patches are scheduled in a staggered fashion backfilling all draining nodes with short jobs to maximize throughput. Additionally, draining can also be scheduled for planned maintenance periods (with optional backfilling) and tasks to be executed before a machine is rebooted or shutdown can be queued. This is combined with a series of automated health checks with large coverage of temporary and long-term machines failures or overloads, and monitoring performed using Zabbix.

In the last year, heterogeneous resources with different I/O capabilities have been integrated and MPI support has been added. All jobs run inside Singularity containers allowing also for interactive, graphical sessions with GPU access.

Combining increasingly heterogeneous resources and different data centre locations in one cluster allows operations with almost zero (full) downtime. During this talk, some examples will be presented on how the automations can be leveraged for different interventions and how the cluster the impact on users and cluster CPU efficiency is minimized.

Primary author: Dr FREYERMUTH, Oliver (University of Bonn (DE))

Co-author: WIENEMANN, Peter (University of Bonn (DE))

Presenter: Dr FREYERMUTH, Oliver (University of Bonn (DE))

Session Classification: Computing and Batch Services

Track Classification: Computing & Batch Services

Contribution ID: 51

Type: **not specified**

Watching File Storage and Transfers with Elastic Stack at BNL

Wednesday, 14 October 2020 19:10 (20 minutes)

The Elastic stack monitors several systems at BNL. At the SDCC two of these are BNLbox and Globus Connect Server. BNLBox- an implementation of nextcloud- is a service for file storing and sharing while Globus is a platform for file transfers. In this talk I will present how we configured Elastic and its components and how they give us insights into client app usage via log ingestion (filebeats) and system monitoring (metricbeats).

Primary author: SNYDER, Matthew (Brookhaven National Laboratory)

Presenter: SNYDER, Matthew (Brookhaven National Laboratory)

Session Classification: Basic IT Services

Track Classification: Basic IT Services

Contribution ID: 52

Type: **not specified**

Running event-driven workflows with dCache storage events

Thursday, 15 October 2020 10:30 (20 minutes)

dCache is an open-source distributed storage system for scientific use cases, actively used by large-scale experiments, including within the WLCG community. In 2018, developers started to introduce the concept of storage events within dCache. We are going to present this concept and see how it can be used to trigger automated workflows, with the example of a proof-of-concept implemented at CC-IN2P3 to help with the LSST image simulation campaign. Finally, we will see how events can be used to transform the way we currently manage infrastructure and workloads when paired with function-as-a-service computing, a new way to deploy highly flexible, automated and smart services.

Primary author: Mr GOUNON, Bastien (CC-IN2P3 / CNRS)

Presenter: Mr GOUNON, Bastien (CC-IN2P3 / CNRS)

Session Classification: Grids, Clouds and Virtualisation

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: 53

Type: **not specified**

A New CMS for a New Decade: Content Management at SDCC

Monday, 12 October 2020 19:00 (20 minutes)

The Scientific Data & Computing Center (SDCC) at BNL is migrating its web content management system from Plone to Drupal. This presentation provides a status update on the project. Several technologies were evaluated and tested according to facility and user requirements and specifications.

Primary authors: DE STEFANO JR, John Steven (Brookhaven National Laboratory (US)); LEPORE, Christian (Brookhaven National Laboratory); PELOSI, Louis (Brookhaven National Lab); SMITH, Jason Alexander (Brookhaven National Laboratory (US))

Presenters: LEPORE, Christian (Brookhaven National Laboratory); PELOSI, Louis (Brookhaven National Lab)

Session Classification: End-user Services, Operating Systems

Track Classification: End-User IT Services & Operating Systems

Contribution ID: 54

Type: **not specified**

The BNLBox service at the SDCC

Monday, 12 October 2020 19:40 (20 minutes)

This presentation covers the BNLBox file sharing service recently introduced at the Scientific Data & Computing Center (SDCC).

Primary author: RIND, Ofer

Presenter: RIND, Ofer

Session Classification: End-user Services, Operating Systems

Track Classification: End-User IT Services & Operating Systems

Contribution ID: 56

Type: **not specified**

Cloud computing to support experiment online computing from the data center.

Thursday, 15 October 2020 17:10 (20 minutes)

Next generation nuclear and high energy physics experiments are moving filtering and processing tasks, previously done with online resources, to the data center. High resolution imaging systems at light sources and electron microscopes require significant amounts of “online” computing resources to rapidly reconstruct images to allow researchers to make “on the fly” adjustments to running experiments, but physical infrastructure prevents these resources from being co-located with the instrument. Financial constraints will also force these resources to be shared. These trends are blurring the line between online and offline computing. Cloud computing technologies, e.g. SDN, VM/Container provisioning, are possible solutions to the myriad of problems that arise from supporting these experiments from the data center.

Primary author: MISAWA, Shigeki (Brookhaven National Laboratory (US))

Presenter: MISAWA, Shigeki (Brookhaven National Laboratory (US))

Session Classification: Thursday evening

Track Classification: Grid, Cloud & Virtualisation

Contribution ID: 57

Type: **not specified**

Scalable High Performance Storage based on Lustre/ZFS over NVMe SSD

Wednesday, 14 October 2020 17:20 (20 minutes)

The presentation shows the Storage test bed set up that I implemented in 2018/2019 when working at SLAC

to meet the 2020 LCLS2 data reduction pipeline requirements reaching 100GB/s.

This set up is primarily oriented on high performance rather than high availability and is based on Lustre/ZFS with NVMe SSD storage and EDR Infiniband connection between Object Storage Servers and clients generating traffic.

Primary author: VERALDI, Riccardo (INFN)

Co-author: FORNARI, Federico

Presenters: VERALDI, Riccardo (INFN); FORNARI, Federico

Session Classification: Wednesday evening

Track Classification: Storage & Filesystems

Contribution ID: 58

Type: **not specified**

NOTED - identification of data transfer in FTS to understand network traffic.

Tuesday, 13 October 2020 10:00 (20 minutes)

NOTED is a project that aims to better exploit WAN bandwidth needed by FTS data transfers. The main component is the Transfer Broker, which interpreters information coming from FTS to identify large data transfer which could benefit of network optimization. The Transfer Broker then enriches the FTS transfers with network information coming from CRIC, the resource database used by ATLAS and CMS.

In this presentation we describe the work done to identify the large transfers in FTS, how these information are enhanced with the data from CRIC, how we try to estimate the duration of the queue, caused by network overloaded. In the nutshell we will show you how we can explain network traffic by using information from FTS.

Primary author: WACZYNSKA, Joanna (Wroclaw University of Science and Technology (PL))

Co-author: MARTELLI, Edoardo (CERN)

Presenter: WACZYNSKA, Joanna (Wroclaw University of Science and Technology (PL))

Session Classification: Tuesday morning

Track Classification: Networking & Security

Contribution ID: 59

Type: **not specified**

ASGC Site Report

Monday, 12 October 2020 10:05 (15 minutes)

Report of development at ASGC recently.

Primary authors: YEN, Han-Wei; Mr LEE, Felix (ASGC)

Presenter: YEN, Han-Wei

Session Classification: Monday morning

Track Classification: Site Reports

Contribution ID: **60**Type: **not specified**

Invenio Based Digital Repositories at BNL

Monday, 12 October 2020 18:40 (20 minutes)

A Research Digital Management (RDM) repository is a web based service that would provide BNL's scientific community a means to share and preserve their scientific results while making them Findable, Accessible, Interoperable, and Reusable (FAIR). Towards this goal Invenio, an open-source software framework for building large-scale digital repositories, is being used to create a research data management platform called InvenioRDM. An overview of the invenio service at BNL and the status of invenioRDM development will be provided in this talk.

Primary author: Dr GAMBOA, Carlos Fernando (Brookhaven National Laboratory (US))

Presenter: Dr GAMBOA, Carlos Fernando (Brookhaven National Laboratory (US))

Session Classification: End-user Services, Operating Systems

Track Classification: Miscellaneous

Contribution ID: 61

Type: **not specified**

IPv6-only on WLCG - update from the IPv6 working group

Tuesday, 13 October 2020 18:45 (20 minutes)

The transition of WLCG central and storage services to dual-stack IPv4/IPv6 has gone well, thus enabling the use of IPv6-only CPU resources as agreed by the WLCG Management Board. More and more WLCG data transfers now take place over IPv6. The dual-stack deployment does however result in a networking environment which is much more complex than when using just IPv6. During the last year the HEPiX IPv6 working group has been investigating the removal of the IPv4 protocol in more places. We will present our recent work and future plans.

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

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

Session Classification: Networking and Security

Track Classification: Networking & Security

Contribution ID: 62

Type: **not specified**

CVMFS service evolution and infrastructure improvements

Wednesday, 14 October 2020 09:40 (20 minutes)

The Cern VM File System (CVMFS) is a service for fast and reliable software distribution on a global scale. It is capable of delivering scientific software onto physical nodes, virtual machines, and HPC clusters by providing POSIX read-only file system access. Files and metadata are downloaded on-demand by means of HTTP requests and take advantage of aggressive caching on intermediate caches and clients. The choice of the HTTP protocol also enables the exploitation of standard web servers and web caches, including commercially-provided content delivery networks.

CVMFS is widely adopted in the HEP community for the distribution of production software, integration builds, auxiliary datasets, and has recently introduced new capabilities to broaden its scope of application. As a prime example, it implements extensive support for container images with DUCC (Daemon that Unpacks Container images into CVMFS), a specialized component that unpacks container images and publishes their extracted form on a repository, and tight integration with container runtimes, making published container images usable by widely-adopted container platforms (e.g., Singularity, Docker, Kubernetes). Such functionality provides an alternative to traditional container registries (e.g., Docker Hub, GitLab Container Registry) and makes the distribution of container images more efficient by leveraging on file-based deduplication and on-demand caching provided by CVMFS.

CVMFS at CERN has also been subject to several infrastructural updates. The repository storage for Stratum Zero servers is now hosted on the Ceph-based S3 service, which provides a relevant performance improvement with respect to block storage provided via Cinder volumes. Also, content distribution to clients has been made more resilient by deploying dedicated caches for sets of repositories, which greatly reduces the problem of interference across repositories and cache thrashing phenomena.

Primary authors: BOCCHI, Enrico (CERN); FURANO, Fabrizio (CERN); VAN DER STER, Dan (CERN); MOSCICKI, Jakub (CERN)

Presenter: BOCCHI, Enrico (CERN)

Session Classification: Wednesday morning

Track Classification: Storage & Filesystems

Contribution ID: 63

Type: **not specified**

HIFIS backbone transfer service: FTS for everyone

Wednesday, 14 October 2020 10:40 (20 minutes)

The German Helmholtz Association (HGF) encompasses 19 research institutes distributed all over Germany, covering a wide variety of research topics ranging from particle and material physics over cancer research to marine biology. In order to stimulate collaborations between different centres, the HGF established so-called incubator platforms. Two of those platforms, relevant for this presentation, are the Helmholtz Artificial Intelligence Cooperation Unit (Helmholtz AI) and Helmholtz Federated IT Services (HIFIS). While Helmholtz AI was established to connect domain scientists and AI experts for a stronger adoption of AI solutions for increasingly complex research tasks, HIFIS targets the exploitation of synergy effects in federated IT services offered by the different HGF centres.

During the ongoing ramp-up phase of both platforms, specific use cases of interdisciplinary research are arising and showing that there is a definitive need to transfer a significant amount of large data sets between centres. This primarily results from the fact that the currently used AI solutions are trained on specific data sets and that the processing of that data is sensitive to network latencies. Consequently, remote data access is less efficient in those cases and consequently data needs to be transferred from the domain scientists'home institutions to the AI experts'location, where the model training is taking place.

In order to cater to those needs, a file transfer service is being established by HIFIS for convenient and automated data transfer between the sites of those interdisciplinary research groups. After evaluating competing solutions like Globus Online and Onedata we agreed to go for FTS3 for reasons we will elaborate on during the presentation. FTS3 is a file transfer service that can commission data transfers between storage endpoints and has been developed at CERN for the transfer of WLCG research data between CERN and several hundred LHC Tier centres. Those endpoints need to be able to communicate with each other as well as with FTS, using a third-party copy (TPC) extension of the HTTP protocol. In order to facilitate an easy installation of endpoints, which are not the known WLCG storage systems (like dCache, DPM and EOS), we provide an Apache web server extension that complies with the needs of FTS3 and can thus act as an storage endpoint for data transfers via HTTP-TPC.

We will present the necessary prerequisites for such an endpoint, it's configuration details as well as the modifications applied to the Apache modules. Adding to that, we will present insights into the performance and reliability of the TPC data transfers measured with test data sets.

Primary authors: WETZEL, Tim (Deutsches Elektronen-Synchrotron DESY); MILLAR, Paul; Dr JANDT, Uwe (Deutsches Elektronensynchrotron (DESY)); FUHRMANN, Patrick

Co-authors: Mr HELMERT, Stefan (Helmholtz-Zentrum Dresden-Rossendorf); Mr HUSTE, Tobias (Helmholtz-Zentrum Dresden-Rossendorf); Dr JUCKELAND, Guido (Helmholtz-Zentrum Dresden-Rossendorf)

Presenter: WETZEL, Tim (Deutsches Elektronen-Synchrotron DESY)

Session Classification: Wednesday morning

Track Classification: Storage & Filesystems

Contribution ID: 64

Type: **not specified**

XCache: exploring cache storage technology

Wednesday, 14 October 2020 09:20 (20 minutes)

Distributed computing involves processing data from far remote site. We explore a new type of cache, XCache, developed within XROOTD, and explore the its capability to improve the access of the jobs to remote storages. The proof-of-concept is realized within the ESCAPE european project and measurement of the performances are done using remote sites. The presentation will develop the configuration of the cache and the different tests done to evaluate the efficiency of the cache for the use cases we have at CC-IN2P3.

Primary author: MUSSET, Paul (IN2P3)

Presenter: MUSSET, Paul (IN2P3)

Session Classification: Wednesday morning

Track Classification: Storage & Filesystems

Contribution ID: 65

Type: **not specified**

MALT Project

Wednesday, 14 October 2020 18:30 (20 minutes)

Presenter: BARROSO LOPEZ, Maite (CERN)

Session Classification: Basic IT Services

Track Classification: Basic IT Services

Contribution ID: 66

Type: **not specified**

CERN Appstore for BYOD devices

Monday, 12 October 2020 11:40 (20 minutes)

The number of BYOD devices at CERN is growing and there is interest in moving from a centrally-managed model to a distributed model where users are responsible for their own devices. Following this strategy, new tools are needed to distribute and - in case of licensed software - also manage licences for applications provided by CERN. Available open source and commercial solutions were previously analyzed and none of them proved to be a good fit for CERN academic environment. Therefore, we started to develop a system that can integrate existing open source solutions and provide desired functionality for multiple platforms, both mobile and desktop.

The functionality of the tool has been growing since the project started. The list of supported platforms has been growing and very recently the integration with Flatpak has been added in order to provide applications not just on Windows, but on Linux as well. The number of features provided is also growing and the UX is enhanced along the way.

This talk will describe the architecture and design decisions made to develop a platform-independent, modern, maintainable and extensible system for future software distribution at CERN.

Presenter: BATO, Tamas (CERN)

Session Classification: End-user Services, Operating Systems

Track Classification: End-User IT Services & Operating Systems

Contribution ID: 67

Type: **not specified**

Workshop Wrap-Up

Thursday, 15 October 2020 17:55 (20 minutes)

Presenter: VAN DER REEST, Peter (DESY)

Session Classification: Miscellaneous

Track Classification: Miscellaneous

Contribution ID: **68**

Type: **not specified**

HEPiX board meeting

Thursday, 15 October 2020 15:30 (1h 25m)

Presenters: VAN DER REEST, Peter (DESY); WONG, Tony (Brookhaven National Laboratory)

Session Classification: Board Meeting

Track Classification: Miscellaneous

Contribution ID: 69

Type: **not specified**

HTCondor 2020 workshop report

Tuesday, 13 October 2020 17:20 (20 minutes)

The “European” HTCondor workshop was held from 21 to 25 September as a purely virtual event. We will report on topics covered and observations made.

Primary author: FLIX MOLINA, Jose (Centro de Investigaciones Energéticas Medioambientales y Tecnológico)

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

Session Classification: Tuesday evening

Track Classification: Computing & Batch Services

Contribution ID: 70

Type: **not specified**

Experiences from online Wokshops and Conferences

Monday, 12 October 2020 11:00 (20 minutes)

The talk will meld the experience of the HSF Workshop, PyHEP and ICHEP and discuss what proved to work best and things that didn't work as well as hoped.

Primary author: STEWART, Graeme A (CERN)

Presenter: STEWART, Graeme A (CERN)

Session Classification: End-user Services, Operating Systems

Track Classification: End-User IT Services & Operating Systems

Contribution ID: 71

Type: **not specified**

dCache in the cloud environment

Wednesday, 14 October 2020 17:40 (20 minutes)

Presenter: Mr MKRTCHYAN, Tigran (DESY)

Session Classification: Wednesday evening

Track Classification: Storage & Filesystems

Contribution ID: 72

Type: **not specified**

Production deployment of the CERN Tape Archive (CTA) for Atlas

Wednesday, 14 October 2020 11:00 (20 minutes)

Presenter: LEDUC, Julien (CERN)

Session Classification: Wednesday morning

Track Classification: Storage & Filesystems

Contribution ID: 73

Type: **not specified**

group photo morning

Tuesday, 13 October 2020 10:20 (5 minutes)

Session Classification: Tuesday morning

Track Classification: Networking & Security

Contribution ID: 74

Type: **not specified**

group photo evening

Tuesday, 13 October 2020 17:40 (5 minutes)

Session Classification: Networking and Security

Track Classification: Networking & Security

Contribution ID: 76

Type: **not specified**

Security Update

Thursday, 15 October 2020 09:40 (10 minutes)

HEPiX might have been hacked - stand by for breaking news.

Presenter: WARTEL, Romain (CERN)

Session Classification: Thursday morning

Contribution ID: 77

Type: **not specified**

Security update

Thursday, 15 October 2020 17:00 (10 minutes)

Presenter: WARTEL, Romain (CERN)

Session Classification: Thursday evening