

CERN site report

HEPiX Spring 2023

27-31 March, Taipei, Taiwan

Jarek Polok <jaroslaw.polok@cern.ch>

On behalf of CERN IT

CERN Data Centre(s)

Prévessin Data Centre (PDC):

- The progress with the construction of the PDC continues well.
- We are on target for it being operational before the end of the year.
- Currently working on the façade and installing technical equipment.

Hardware deliveries:

- All supplies ordered in 2022 finally arrived by the annual closure.
- Acceptance testing completed in Q1 2023 without any issue.
- New contracts awarded at December FC. First round of purchases:
 - 280x quads → **1'120** compute nodes → 143'350 SMT cores (AMD Milan) → 2'380MHS06.
(To be deployed in the Preveessin Data Centre.)
 - 360x JBODs → **8'640** hard drives → 155'520 TB.
(To be deployed in the Meyrin Data Centre.)

Other:

- **Energy saving: one of the containers has been shut down for the winter period (1200 batch nodes, ~ 0.5MW), restarted last week.**
- **Integration of PDC monitoring solution with CERN platforms in progress.**
- **Proposing reduction of 24x7 data center operator coverage (night shift suppression).**



IT Computing Facilities
Data Center 0775-2-0007

0775-2-0007				Showing: Overview	Power/Temp/Hum	
Infrastructure	Occupied	Allocated	Available			
Total U	4230	0	0	4,230		
Percentage	0.0 %	0.0 %	0.0 %	100.0 %		
Delivery Address	CYDC022					
Administrator	IT					
Computed Wattage	0 kW					
Measured Wattage	0 kW					
Design Maximum	1,734 kW					
Data Center Size	344 m ²					
Watts per Square Meter	0 W					
Total Cabinets	90					
Average Temperature	N/A					
Average Humidity	N/A					

Storage Technologies

CERNBox: demonstrated the possibility to share storage resources and applications across organisations and institutions.



- Based on an open interoperable protocol called Open Cloud Mesh, which has been adopted by CERNBox and other popular cloud storage services.
- Presented at the [Cloud Services for Synchronisation and Sharing](#) conference, part of [CS3 conference series](#).



Backup and Restore service: re-negotiating 15 PB license for IBM Spectrum Protect for another 3 years.

CERN TapeArchive:

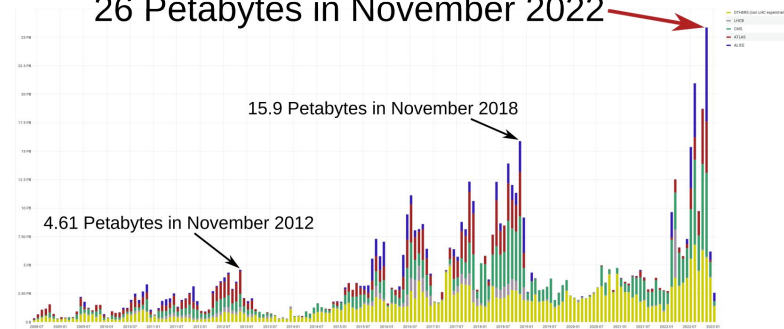
- **Tape infrastructure:** while waiting for new enterprise tape technology, we are purchasing more LTO9 media.
- **HTTP TAPE REST API** deployed in CTA LHC instances.
- **CTA set a new CERN record in monthly archival volume to tape.**

Talk this week: [Status of CERN Tape Archive operations during Run3](#)

New CERN Record

Data archived to tape storage each month since 2008

26 Petabytes in November 2022



Storage Technologies

EOS:

- **2022 statistics**

- 21.8 Billions files read for more than 4EB of data
- 5.73 Billions files recorded counting for around 631 PB of data
- In November 2023 the EOS Physics clusters achieved a new record, we exceeded 500GB/s of data delivered to the experiments

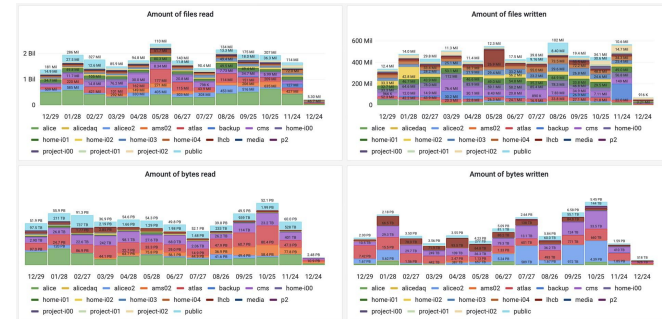

- **50 PB raw will be installed to provide the storage increase for 2023 pledges**

- **EOS ALICE Online/Offline (O2) Storage cluster, latest results:**

- the cluster is able to deliver 500 GB/s
- we benchmark the ingress at 280 GB/s
- in mixed mode we achieved 250 GB/s writes while reading 200 GB/s concurrently
- due to ALICE evolving requirement for the Heavy-Ion Run, the cluster will be expanded with an additional 50 PB

Ceph:

As part of the department Business Continuity/Disaster Recovery activities, two clusters, one dedicated to CephFS and one to RBD, were moved entirely onto generator-backed critical power.



Talk this week: [CERN Storage Services Status and Outlook for 2023](#)

Database and Analytics

DataBase on Demand (DBOD) service


- Migrating all hosts from CentOS 7/8 to RHEL 8
- Mandatory MAJOR upgrades for MySQL and PostgreSQL (EOL)
 - MySQL 5.7.x (EOL Oct 2023, ~300 instances)
 - PostgreSQL 11.x (EOL Nov 2023, ~50 instances), PostgreSQL 12.x (EOL Nov 2024, ~260 instances)

State	Type	Version	Port	Category
Running	MYSQL	5.7.37	5503	PROD
Running	MYSQL	8.0.28	5501	PROD
Running	MYSQL	8.0.28	5508	TEST
Running	MYSQL	5.7.37	5508	PROD



OpenSearch migration status

- Out of ElasticSearch 6.X: done!
- Out of ElasticSearch 7.X: Q1 2023
- Out of OpenDistro: Scheduled Q3 2023



Hadoop service

- Dedicated HBase deployment for NXCALS low-latency data pipeline to improve performance/robustness.



Web hosting infrastructure

Consolidated, simplified Short URL service

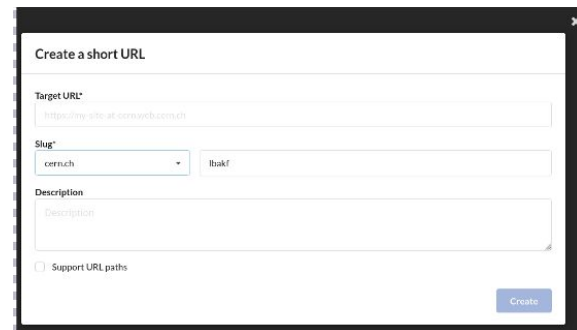
- For all short URLs - `http://cern.ch/something`
- Management from the new web portal

OpenShift PaaS container logs

- Move from ElasticSearch to OpenSearch enabled fine-grained permissions
- All PaaS application logs can be sent to central OpenSearch instance, and each application owner sees only content for their applications

Jenkins CI hosting in OpenShift PaaS

- Jenkins CI lost traction to GitLab CI, GitHub actions
- Discontinuing custom CERN image and template in favor of upstream Jenkins image and template provided by OpenShift



Create a short URL

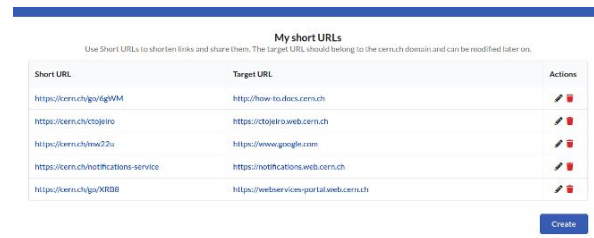
Target URL*
https://my.site.at-cernweb.com.ch

Slug*
cern.ch

Description

Support URL paths

Create



My short URLs

Use Short URLs to shorten links and share them. The target URL should belong to the cern.ch domain and can be modified later on.

Short URL	Target URL	Actions
https://cern.ch/go/gW4M	https://how-to.docx.cern.ch	✎ 🗑
https://cern.ch/ctojeiro	https://ctojeiro.web.cern.ch	✎ 🗑
https://cern.ch/w22a	https://www.google.com	✎ 🗑
https://cern.ch/notifications-service	https://notifications.web.cern.ch	✎ 🗑
https://cern.ch/gu/XR00B	https://webservice-portal.web.cern.ch	✎ 🗑

Create



OPENSHIFT

Identity & Access Management

Single Sign-On

- Migration from ADFS to Keycloak in progress.
- Keycloak upgraded in January, significant performance improvements seen.
- Interested in connecting with other Keycloak users.

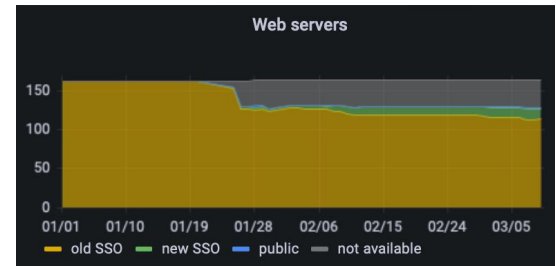
WLCG Token Issuers (INDIGO IAM)

- Now running for 5 experiments.
- Active collaboration between INDIGO IAM Users, participation welcome.


Certificate Authorities

- Investigating GEANT's TCS service for grid user and host certificates.
- Host certificates would be both publicly trusted and IGTF approved.


Talk this week: [Status of CERN Authentication and Authorisation](#)



Remaining standalone webservers connected to old SSO in Jan 2023



Performance improvements following SSO upgrade to Keycloak v 19



WLCG IAM Instances

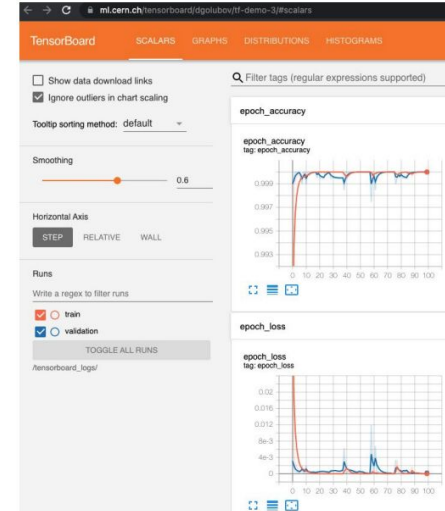

Machine Learning

New CERN IT ML initiative

- Organizing a [series of workshops](#) to better understand and improve the ML/DL infrastructure offering in CERN IT.
- Covering infrastructure, platforms and use cases.

ml.cern.ch

- **Improved monitoring**
 - Add support for TensorBoard, model training monitoring in real time.
 - Visualization of loss and accuracy, model graphs, custom metrics.
 - Integrated in the main kubeflow portal.
- **Dedicated resources**
 - Assignment of resources to specific groups / users.
- **GPU sharing**
 - Sharing of GPUs with time slicing and MIG.
 - Different use cases and SLAs for each concurrency type.
- **Quota enforcement.**



```
apiVersion: v1
kind: ResourceQuota
metadata:
  name: kf-resource-quota
  namespace: dgolubov
status:
  hard:
    limits.cpu: "5"
    limits.memory: 10Gi
    limits.nvidia.com/gpu: "1"
    (soon) limits.nvidia.com/gpu.shared: "1"
    requests.cpu: "5"
    requests.memory: 10Gi
    requests.nvidia.com/gpu: "1"
    (soon) requests.nvidia.com/gpu.shared: "1"
```


Google workspace

Deployed a Google Workspace organization

- Space owned and managed by a domain-verified Google customer
- Integrated with CERN SSO
- 5-year contract

Goal: secure and preserve the work produced by CERN community in the professional context

- CERN fully manages accounts associated with the organization
- CERN fully owns the data - which is contained in CERN tenant located in EU

Basic set of applications for all: Drive, Docs, Sheets, Slides, Forms

- More Google apps can be enabled on-demand

Priority is to onboard *existing* personal Google accounts registered with @cern.ch email address (>2000 accounts)

- Transferred more than 70 personal accounts to CERN Google Workspace.
 - users can choose to migrate professional data.



Kubernetes Service

Releases 1.24 and 1.25

- Support for ephemeral debug containers, removal of PodSecurityContext and much more - [OTG0075392](#)

Improvements in container storage

- [csi-cvmfs driver](#) now hosted under the official CVMFS github
- New csi-cvmfs release with support for automounts with greatly simplified configuration
- Prototype of OAuth2 support in the EOS integration, including automated renewal of credentials

Sharing of GPUs



- Released support for time-slicing for all Nvidia cards
- Support for A100 cards including Multi-Instance-GPU (MIG) for physical partitioning up to 7 times
- New [blog post series](#) covering setup and benchmarks

Container Webinars

- New [round coming up](#) covering ArgoCD/GitOps, Storage, Workflows and more



GPU "CONCURRENCY" Choices



GitLab

Critical service at CERN

- More than 18k active users
- More than 130k projects
- GitLab Pages - documentation sites
 - More than 500 sites


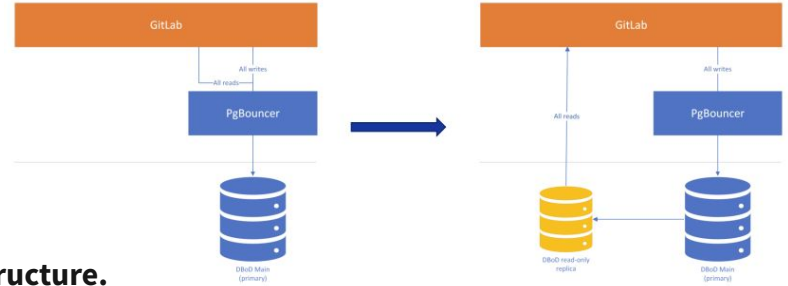
Migration from custom-made to a new Cloud-Hybrid supported infrastructure.

New GitLab Runners infrastructure based on containers ongoing

- AArch64 Runners
- FPGA/SoC
- General Shared Runners

Road to a more scalable and robust infrastructure

- Transaction pooling in place
 - Connection overhead reduction, resources efficiency.
- Database load balancing (primary/secondary)
 - Improving query performance.
- GitLab Cluster (HA) - repositories data
 - Data integrity benefit



Drupal/Atlassian products

Drupal

- **Upgraded Drupal 9 websites to PHP 8.1**
- **Drupal 10 preparations ongoing**
 - Drupal 9 end-of-life is coming in November
 - In Drupal 10 onwards, only CKEditor 5 will be available. (major impact)
 - Ongoing discussions. CERN "Web presence" review taking place: governance, technologies, etc.
 - More news coming soon
- **Continued efforts to decommission legacy infrastructure: Drupal 8 and Drupal 9/PHP 7 websites, Test websites on the legacy VM Infrastructure**



Atlassian

- **CERN procurement is finalising the conditions of the new Atlassian contract for JIRA and Confluence, Data Centre editions**
- **Plan to consolidate multiple instances of JIRA and Confluence into a single, central Data Centre instance**



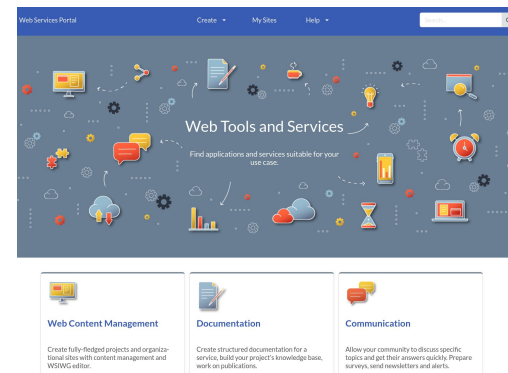
Web Infrastructure & SharePoint Online

Web Infrastructure

- **Windows DFS/WebAFS websites – Phase out campaign**
- **Sites will be migrated to the new WebEOS container based infrastructure**
- **Detailed migration plans will be coordinated with CERN user communities. Complex sites will be followed up case by case**

SharePoint Online Service

- **Move from SharePoint 2013 on-premise to SharePoint Online**
- **Users are either offered to create sites on SharePoint Online or use a suitable alternatives**
- **SharePoint Online training offered by CERN Technical training service**



Communication Systems

Final stages of the Alcatel PABX decommissioning - Two issues to address

- **Easy: phone line migration**
 - Fewer than 500 active phones - progressive migration of users underway.
 - 4,500 inactive lines to be decommissioned on 1st June
- **Hard(er): replacing Call Centre functionality**
 - Pilot in-house development in operation for B55 reception and CERN guards
 - On track to migrate Service Desk, CERN Control Centre and the Fire Brigade by end-September.
- **Still expect to complete Alcatel decommissioning by the end of the year.**



CERNphone Resources Portal

Welcome to the CERNphone Resources Portal!

Manage your numbers

On this webpage you can create, modify or delete your fixed telephony numbers (mobile telephony currently can be managed from EDH).

CERNphone numbers

CERNphone is the new softphone solution provided by the IT department created to replace the legacy telephony systems (like Alcatel). From here you can manage your CERNphone numbers:

Number / Description	Type	Service Account	IP Phone	Actions
71834	PERSONAL		TELEPHONE-71834	⚙️ Actions

Create a new number

Do you want to create a new personal or shared CERNphone number? Note that you can only have one active personal number.

Create shared

Query phone number usage statistics

Links

- CERNphone Users Portal
Configure your number settings and mobile devices
- CERNphone Installation Instructions
A guide for configuring your CERNphone number step by step
- Download CERNphone clients
CERNphone is available for Windows, Linux & Mac - Android - IOS
- Contact and Support
For assistance and support please send an email to the CERNphone support team or open a ticket in Service Now. For general questions, suggestions, etc you can contact the CERNphone team.



Optimise battery usage

App	Usage	Background	Allow background data usage
AI (459)	1.48 MB	oreground	675 KB
CERNphone (94.08 MB)	675 KB	ackground	807 KB
Certificate Installer (172 KB)			
Chrome (403 MB)			

Allow background data usage

Allow data usage while Data saver is on

Groupware / Conferencing Technology

E-mail service

- Migration to Exchange Online progressing according to schedule.
- to be completed by early summer



Transcription and Translation as a Service (TTaaS):

- Pilot for offline media transcription is now available and fully integrated within CERN's internal lecture recording workflow.
- The full backlog of CERN lecture videos has now been subtitled and made available to users.

Transcribe
Transcribe your media automatically by just uploading it.

Translate
Translate your media automatically to several languages.

Boost accessibility
Boost the accessibility of your media by adding captions and transcriptions.

Documentation
Documentation of the project and user guides [can be found here](#).

CERN Transcription and Translation Service - 2022 - 2023 - [Credits](#)



CERN Open Data portal

Open Data @ CERN

- CMS released remaining Run1 collision data (490 TB)
👉 <https://opendata.cern.ch/docs/cms-completes-run1-pp-data-release>
- LHCb released first Run1 data stream samples (200 TB)
👉 <https://opendata.cern.ch/docs/lhcb-releases-first-set-of-data>
- New open data releases bring content to over 3 PB
👉 <https://opendata.cern.ch>

Forthcoming

- Started work on integrating LHCb Ntuple Wizard allowing theorists/phenomenologists to ask for custom on-demand data production workflows



REANA Reusable Analyses

REANA @ CERN

- REANA 0.9.0 is released

👉 <https://blog.reana.io/posts/2023/reana-0.9.0/>

- Custom workspace asset expiry rules

👉 <https://blog.reana.io/posts/2022/workspace-file-retention-rules/>

- Support for Rucio external storage system

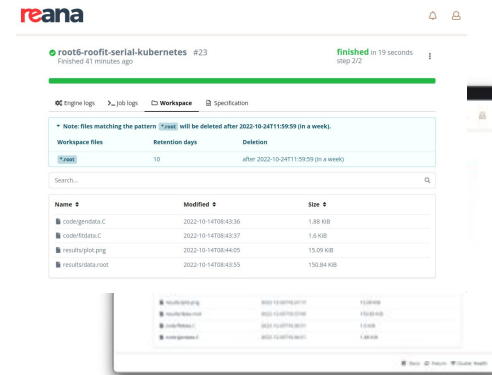
👉 <https://docs.reana.io/advanced-usage/access-control/rucio/>

- Support for Kerberos workflow orchestration

👉 <https://blog.reana.io/posts/2023/kerberos-for-the-whole-workflow/>

Forthcoming

- Support for Keycloak IAM (ESCAPE, PUNCH4NFDI)
- Auto-closure of inactive interactive notebook sessions
- Purge temporary files after workflow execution



```
inputs:
  files:
    - Snakefile
  parameters:
    input: inputs.yaml
workflow:
  type: snakemake
  file: Snakefile
resources:
  kerberos: true
outputs:
  files from snakemake.reana.XRootDExport RucioStProvider as XRootDExportProvider
  - MYO XRootD = XRootDExportProvider(stay_on_reana=True)

file_numbers = XRootDglob_wildcards("root://eosuser.cern.ch/eosuser/jj/shdow/repdata/ing
rule: J11
input:
  expand("mylocaldata/myfile_{i}.csv", nfile_numbers)
rule: Fetch_data
input:
  XRootD-rucioa("root://eosuser.cern.ch/eosuser/jj/shdow/repdata/myfile_{i}.csv")
output:
  "mylocaldata/myfile_{i}.csv"
environment:
  "docker://apptainer54/nginx-w17-va/exasp-20230208-8811"
shell:
  "mkdir -p {input[0]} {output[0]} && wget -O mylocaldata/64_torch_output[0]."
```



Digital Repositories

InvenioRDM

Open Source Research Data Management Platform that will power Zenodo and CDS repositories. Developed in collaboration with >26 international institutions.

Several demo and production sites already deployed on latest version (e.g. [Caltech Data Repository](#), [WHO Research Data Registry](#) prototype, [TU Wien](#), [TU Graz](#), [GEO Knowledge Graph](#), etc)

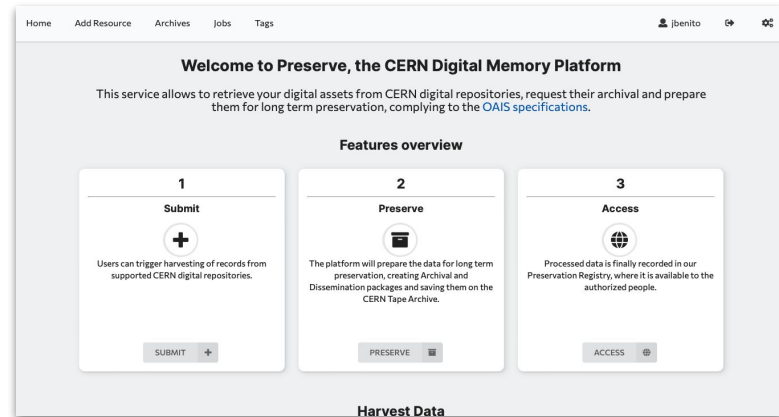
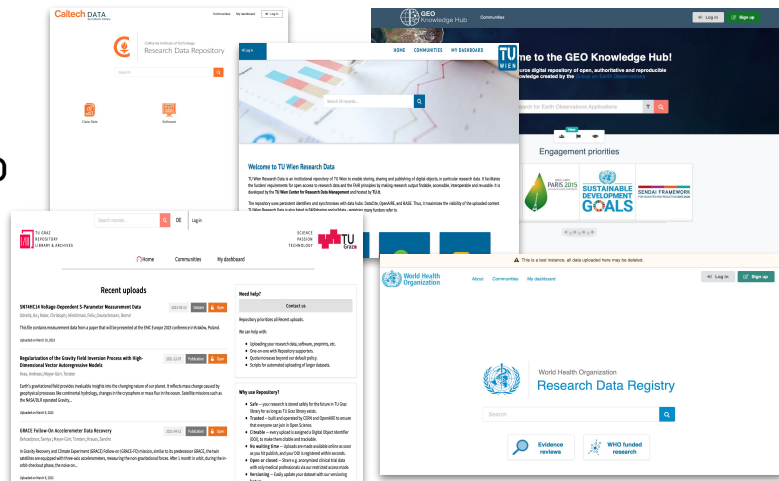
See docs for more info: <https://inveniordm.docs.cern.ch>

Digital Memory

Long-term digital preservation platform launched. Makes sure that content is usable for future generations independently of the changes in technology.

<https://preserve.web.cern.ch> (CERN only)


Contact: digitalmemory@cern.ch



CERN presentations this HEPiX

- [Monitoring Windows Server Infrastructure using Open Source products](#) - Pablo Martin Zamora
- [Status of CERN Authentication and Authorisation](#) - Asier Aguado Corman
- [Cloud Infrastructure update from CERN](#) - Maryna Savchenko
- [ARM Support in the CERN Private Cloud Infrastructure](#) - Maryna Savchenko
- [Fully automated: Updates on the Continuous Integration for supported Linux distributions at CERN](#) - Ben Morrice
- [Update on the Linux Strategy for CERN \(and WLCG\)](#) - Ben Morrice
- [CERN Computer Centre\(s\) Network evolution: achievements during the last 3 years and expected next steps](#) - Vincent Ducret
- [Computer Security Landscape Update](#) - Liviu Valsan
- [Status of CERN Tape Archive operations during Run3](#) - Julien Leduc
- [CERN Storage Services Status and Outlook for 2023](#) - Julien Leduc





home.cern