

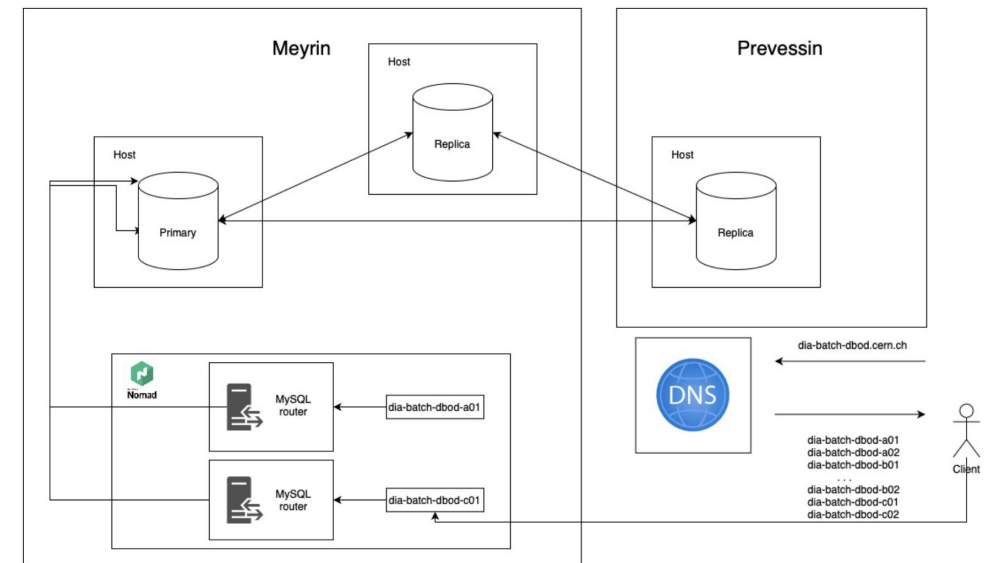
# **CERN Site Report**

## **HEPiX Fall 2024, Oklahoma**

**Elvin Sindrilaru**  
**On behalf of CERN IT**

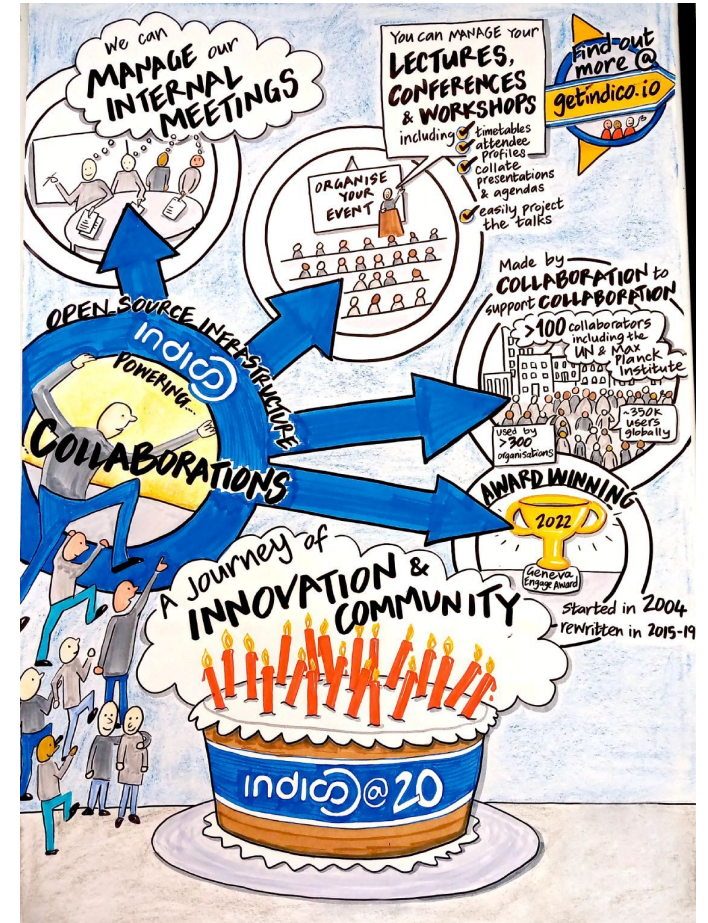
# Databases & Analytics

- **Data Streaming and Messaging:**
  - Progress on **Kafka** integration, including migration efforts and validation of **S3** connectors for archival
  - Integration with **Grafana Mimir - Prometheus long-term storage**
  - Evaluating a centralized **MQTT** (Message Queuing Telemetry Transport) service
- **SWAN and Machine Learning:**
  - Supported multiple events (ATLAS ML workshop, CERN School of Computing sessions) with GPU resources
  - Integration of **SWAN** with **Kubeflow** for enhanced ML capabilities
- **Database Services:**
  - Last legacy **Oracle** databases upgrade to **19c**
  - Introduced **TimescaleDB** to DB-on-Demand service
  - Evolution of **MySQL High Availability** using **InnoDB**
- **OpenSearch and Monitoring:**
  - Migration from OpenDistro to **OpenSearch** continues
  - Expanded monitoring services, including Kubernetes monitoring via **MONIT**
- **Hadoop and Big Data:**
  - Completed upgrade to **Hadoop v3.3** and working on upgrade to **HBase v2.5**
  - Enabling **SSO** integration with **Apache Knox**



# Collaborative Applications

- **CDS (CERN Document Server)** - <https://repository.cern>
  - Migration to a new platform progressing, with all CDS services migrated out of CC7
  - Addressed a **DDoS** attack and improved collaboration with CERN's security team
- **Indico:**
  - Celebrated its **20th** anniversary, with significant events and an ongoing user survey
  - Enhanced integration efforts, including **Zoom Rooms** prototype
- **Zenodo:**
  - Launched new features like fast image zoom and **Software Heritage** integration
  - Launched **EU Open Research Repository** on behalf of European Commission
- **Email Services:**
  - CERN provides anonymous mail gateway on internal network (used by scripts, web sites etc.). Plan to block “domain spoofing” on this service.
  - **Microsoft** announced introduction of [external recipient limits](#) in 2025 (**2k recipients in 24h**), assessing the impact on the CERN mailing lists
  - Implement **DMARC** on **cern.ch** in “**quarantine**” mode; Q1 2025 in “**rejected**” mode
- **Digital Memory and Open Data:**
  - Hosted the first **LHCb Open Data** workshop and executed cold recovery tests
- **Related talk:** [CI4FPGA presentation](#)



HEP X HEPiX Fall 2024 Workshop

4-8 Nov 2024  
US/Central timezone

Overview  
Scientific Program  
Workshop Chairs and Organizers

HEPiX Fall 2024 at the University of Oklahoma, USA

# Communication Systems

- **Networking:**

- Migrated **LHCb** infrastructure to **Preveessin Data Center (PDC)**, upgrading firmware and adopting a new network model
- Addressed **major power cuts impacting Meyrin and LHC networks**; most devices recovered automatically, with some manual interventions.
- Deployed **anycast DNS** server in PDC

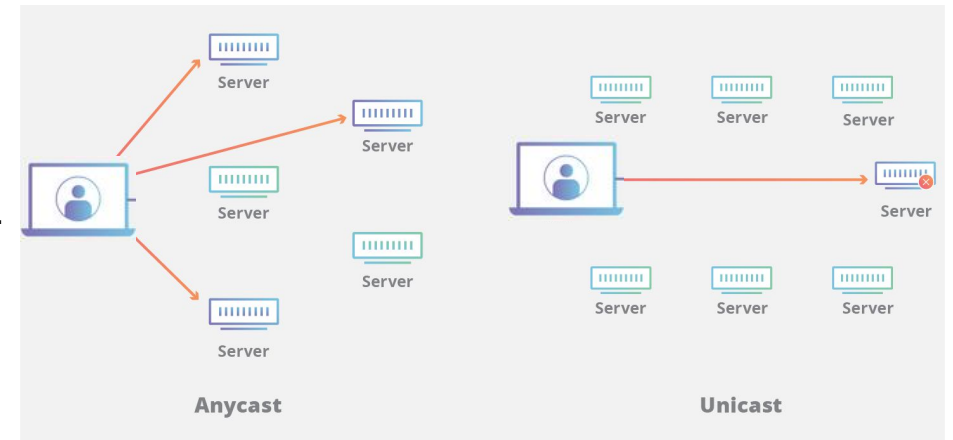
- **Telephony:**

- Finished migration to **RHEL 9** for **TONE/CERNphone** servers
- **Fixed telephony:** Migrate control room phones to a call center-based solution, improving communication reliability
- **Mobile telephony:** Resolved issues with Orange's **VoLTE** service on the Meyrin site and quality improvements

- **Infrastructure and Market Surveys:**

- Launched Market Survey for campus and data center switches
- Upgraded **firewall firmware** to address latency issues, improving connection times

- Service **KPIs** available in Grafana Dashboards





# Fabric - CERN Data Centers

## Plans for evolution of the Data Centers to meet HL-LHC requirements

- **Preveessin Data Center (PDC) - 2024**
  - 4MW (phase I) / 8MW (phase II) high density 25 kW/rack
  - 20% UPS, 5' autonomy
  - **Higher power density** but lower resiliency - **CPU Batch**
- **Meyrin Data Center (MDC) - 1972**
  - 3.8MW, low density 5-10 kW/rack
  - 100% UPS, 10' autonomy (+diesel generator)
  - Lower power density but higher resiliency - **bulk storage**
  - Access to safe power - mission and safety critical services
- **Simplify electrical and cooling infrastructure**
- **Improve power efficiency - update racks + streamline DC operations**



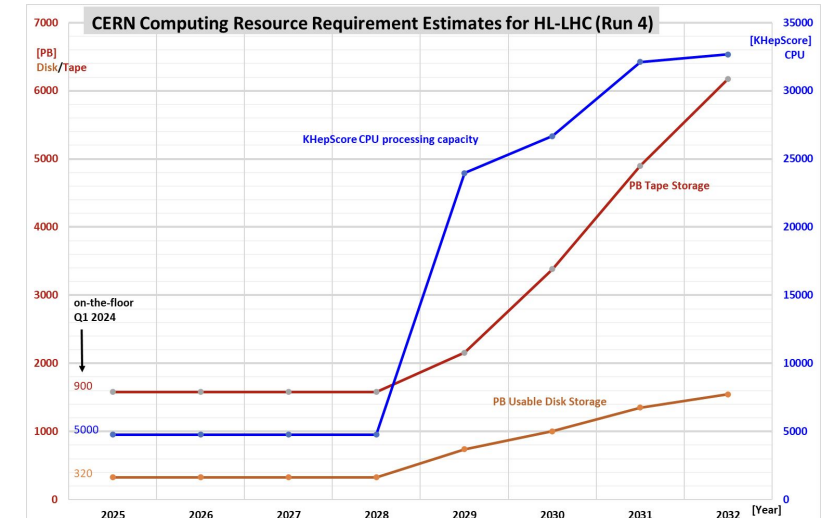
PDC



MDC

## Procurement & Infrastructure:

- New **EOS** capacity procurement (80PB) and **front-end server** acquisitions
- Aging conference room equipment (Codecs) being phased out due to security concerns
- New **HPC cluster** planned for 2025 with ~80 nodes interconnected via **InfiniBand**

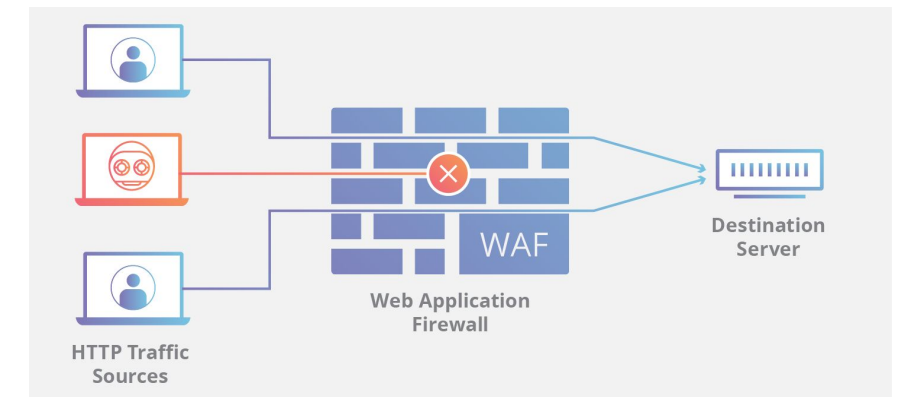
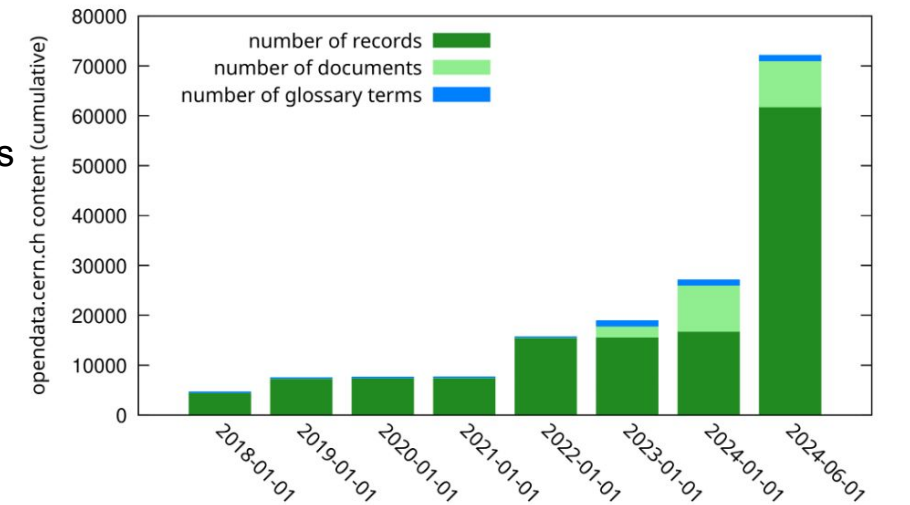


# Platforms and Workflows

- **Open Data:** CMS released 13 TeV proton-proton data (2016), ~900TB total
  - ATLAS, DELPHI, JADE open data releases expected by year-end
- **REANA:** Significant deployments and upgrades, including integration with Kubernetes and hybrid workflows at various institutions
- **VOMS to WLCG IAM Migration:** Migration completed
  - IAM services moved to High Availability Kubernetes infrastructure
- **Machine Learning (ml.cern.ch)**
  - Support for dynamic GPU allocation and major SW upgrade
  - Collaboration with **UNOSAT** for satellite AI analysis
- **CephFS** instabilities during backup affected PaaS volumes
  - Mitigated by using older cephfs client
  - **Fedora CoreOS (FCOS) -> CentOS Stream CoreOS (SCOS)**
- **SSO** - DB connection instabilities fixed; upgrade to **Keycloak 24** planned
- **GitLab Duo AI-suite** being evaluated
- **WordPress Lite** service available, work in progress on the full central service

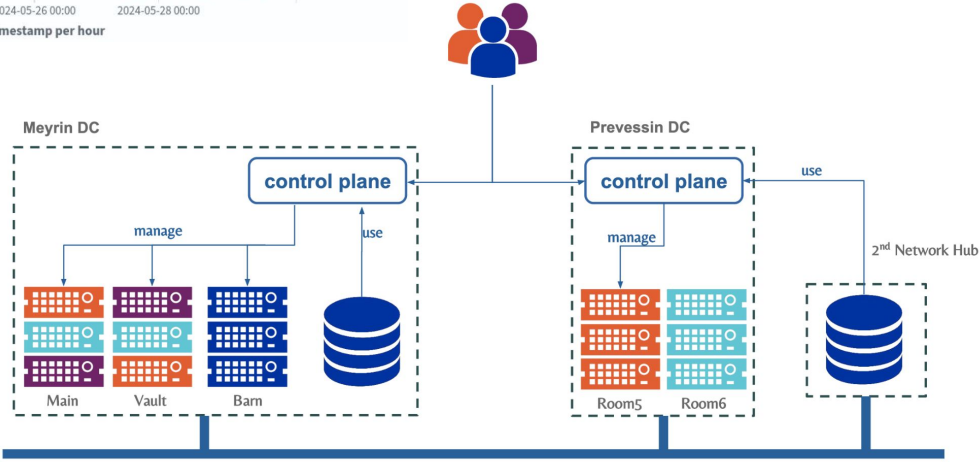
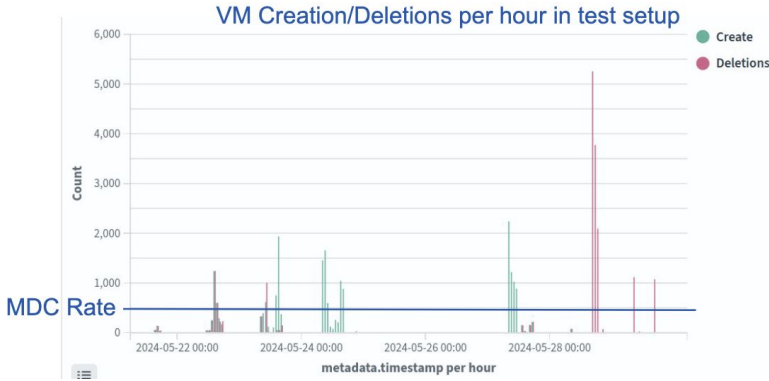
## Plans:

- **Kubernetes and OpenShift 4.16** upgrades planned for late 2024
- **Web Application Firewall (WAF)** to be deployed for **WebEOS**



# Compute and Devices

- **Cloud Services in the Preveessin Data Center**
  - Integrate new data center with **OpenStack**
  - Successful **network scalability tests** 2k hypervisors / 14.5k VMs
- **Linux & Cloud:** hypervisors **EL9** upgrades completed
  - **Debian Linux LTS** support in place
  - New **ARM hypervisors** available
  - **Kubernetes** update to **1.31.1**
- **Batch & Compute**
  - **2000 ARM** cores added to **LxBatch**
  - **HTCondor** update to 23.10
- **Windows Services:** Migration of compatible devices to Windows 11 mostly complete
  - Plan for **Windows 10 Extended Security Updates (ESU)**
- **Openstack APIs** closed from public network and website 2FA from 2nd of December
- **PuppetDB** upgrade from 4 to 8



# Storage and Data Management

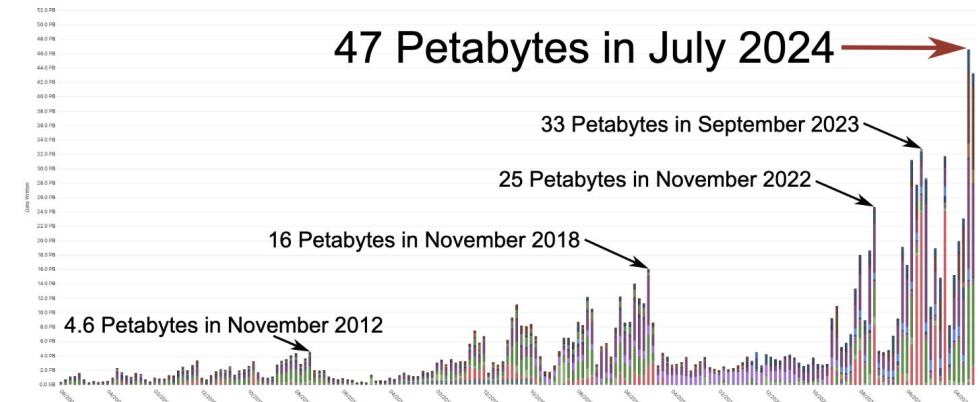
- **Storage Infrastructure:**
  - Significant **tape storage upgrades**
    - Purchased **280 PB** of **LTO-9** and **IBM 3592J** tape cartridges
    - Installed a new **Spectra Logic TFinity+** tape library (up to **270PB LTO-9** capacity)
  - Expanded **Ceph** storage in the **Prévessin Data Centre (PDC)** with 7.8 PB
    - Instabilities traces to kernel changes, mitigation in place
  - Phase-out of **gridFTP** support on **FTS, EOS, and CERNBox**.
  - Hardware issues with older disk nodes and **kernel crashes (XFS)** on EOS
  - **FTS** successful **Archive Metadata** scale test, preparing for Heavy Ion Run
- **Backup and Recovery:**
  - Updated **IBM Storage Protect (TSM)** to the latest release and migrating to **RHEL8** on new server hardware

## Future Plans:

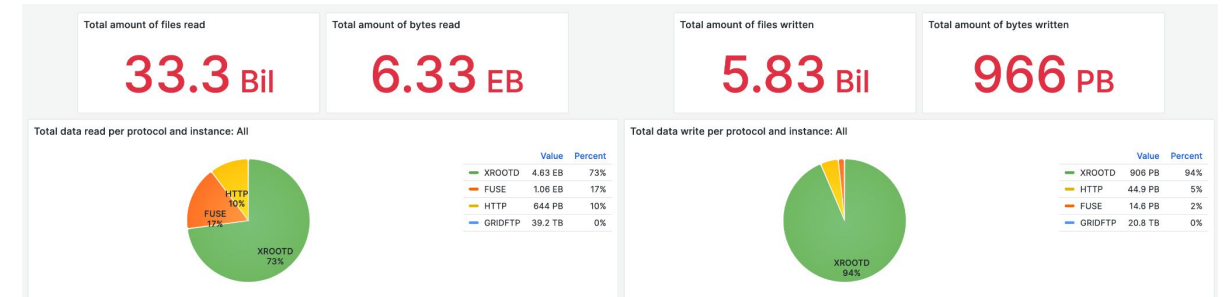
- Conduct **large-scale disaster recovery tests** and resume **Archive Metadata** deployment for data-intensive tasks
- **Disk technology challenges**
  - Currently **cost/TB** optimized, ignoring **IOPS**
- **Tape technology challenges**
  - 20TB -> 50TB, but no backwards compatibility for new drive generations
  - **Consequences:**
    - Older tape drives can not be removed until data is repacked
    - **Repack** on a more aggressive schedule than in the past

## New CERN Record

Data archived to tape storage each month since 2008



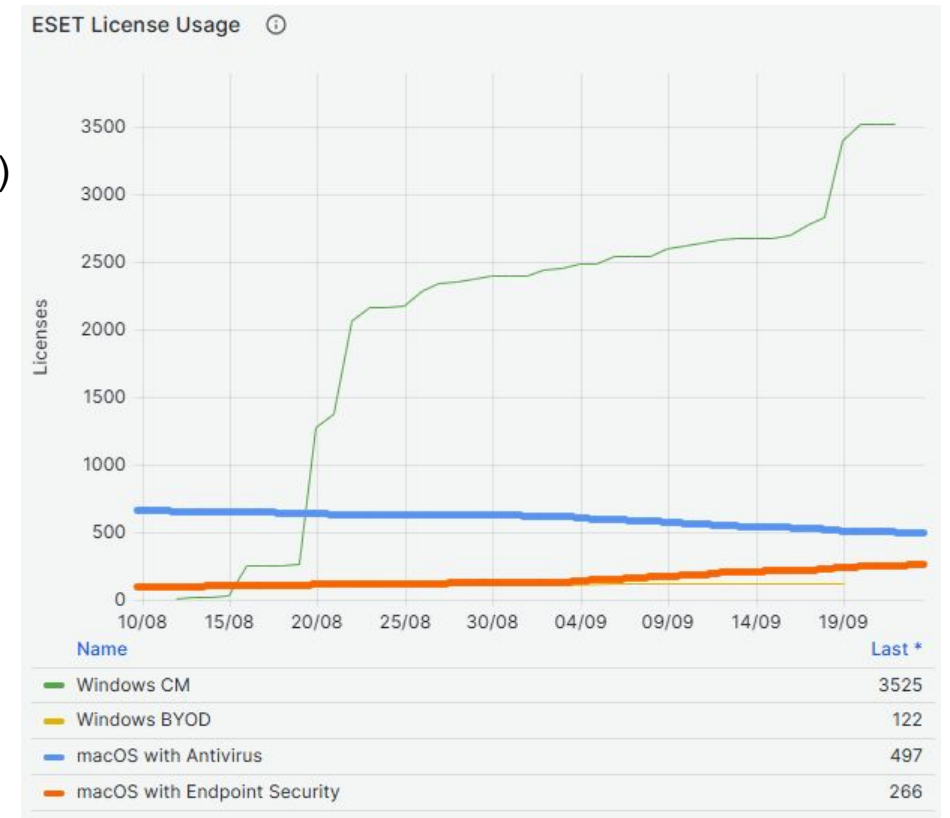
## EOS for Physics 2023





# Security

- Major security actions include **ESET** deployment across departments and vulnerability monitoring for **Microsoft Outlook**
- Completing deployment of **anti-spoofing techniques** (SPF, DKIM, DMARC) with DMARC enforcement on “cern.ch” by Q1 2025
  - **Note: mailing lists** which are **spoofing** the sender address may end up in **quarantine**
- Remote access using **service accounts** disabled on **LxPlus & LxTunnel**
- **Two-factor authentication** expansion on **LxPlus**
- Roll out for **40k SSO users** to be finished by Q1 2025
- **Related talk:** [Computer Security Updates](#)
- **Related talk:** [One year into the CERN Cyber-Security Audit](#)
- **Related talk:** [HELP! I have DataCenter Nightmares](#)





[home.cern](https://home.cern)