

CERN site report

HEPiX Autumn 2023

16-20 October, Victoria, Canada

Jarek Polok <jaroslaw.polok@cern.ch>

On behalf of CERN IT

CERN Data Centre(s)

Prévessin Data Centre (PDC):


- The project is on time (and budget). Civil engineering works completed.
- Commissioning testing ongoing since few months.
- Integrated Systems Tests of full phase 1 (4 MW) capacity using loadbanks (industrial heaters) in progress.
- Final trial operation with servers in one pod (0.5 MW) planned for December - concluding the acceptance tests: PDC will enter operation phase after.
- 10 years maintenance and (infrastructure) operation contract starts immediately after completing the acceptance.

Hardware deliveries:

- Physics batch processing (PDC) received and ready for deployment: 280x quads → 1'120 compute nodes → 143'350 SMT cores (AMD Milan) → 2'380MHS06
- Data storage: 46x quads → 184 compute nodes → 11'776 SMT cores (AMD Milan)
- 360x JBODs mostly for EOS → 8'640 hard drives → 155'520 TB
- Business Continuity and Disaster Recovery (PDC): ordered - to be deployed in PDC

Other:

- Reduction of 24x7 data center operator coverage (night shift suppression) implemented.



Telephony

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

- **Easy: phone line migration**

- Less than 500 active phone lines - progressive migration of users.
- 4500 inactive phone lines decommissioned 1st of July 2023.

- **Hard(er): replacing the Call Centre functionality**

- Pilot in-house development in operation for reception and CERN guards.
- On track to migrate Service Desk, CERN Control Center and the Fire Brigade.

- **Still expect to complete Alcatel decommissioning by the end of the year.**

The end of a nearly 10-year project to migrate to SIP, Asterisk and softphones!



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 EDI).

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

AI (45%)

CERNphone (1.48 MB) - background (675 KB)

Certificate Installer (172 KB) - background (807 KB)

Chrome (402 MB) - allow background data usage

Allow data usage while Data saver is on

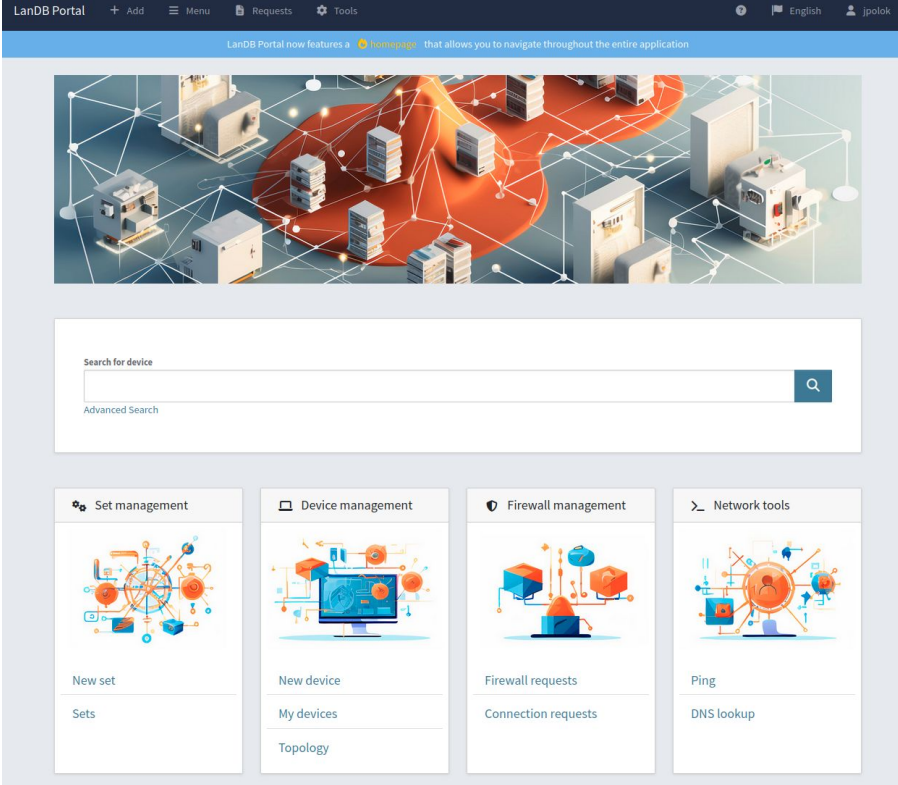
Networking

Changes to the CERN Wi-Fi SSID

- Now restricted to people with CERN accounts
- Visitors must use the CERN-Visitor SSID
- The option for requesting access to the CERN SSID has been removed.

A pilot for WPA3 userid/password authentication instead of today's MAC-based authentication will start in IT before the end of the year.

- If successful, it will likely be rolled out across CERN during 2024.



The screenshot displays the LanDB Portal interface. At the top, there is a navigation bar with 'LanDB Portal', '+ Add', 'Menu', 'Requests', and 'Tools'. A secondary bar indicates 'LanDB Portal now features a **Homepage** that allows you to navigate throughout the entire application'. Below this is a large banner image depicting a network topology with servers and routers. A search bar for devices is present, with a 'Search for device' input field and a magnifying glass icon, and a link for 'Advanced Search'. The main content area is divided into four columns of network management tools:

- Set management**: Includes 'New set' and 'Sets'.
- Device management**: Includes 'New device', 'My devices', and 'Topology'.
- Firewall management**: Includes 'Firewall requests' and 'Connection requests'.
- Network tools**: Includes 'Ping' and 'DNS lookup'.


Storage Technologies

An Exabyte of disk storage at CERN:

This data capacity is provided using 111 000 devices, predominantly hard disks along with an increasing fraction of flash drives.


Major part of the capacity is managed by EOS.

<https://home.cern/news/news/computing/exabyte-disk-storage-cern>




Tape Storage:

New CERN Record - 35 PB (July 2023)



New CERN Record

Data archived to tape storage each month since 2008



IT Department | Storage and Data Management Group

Database and Analytics

OpenSearch deployment



- Completed the migration of ElasticSearch v6 and v7 clusters to OpenSearch.
- Still missing the migration of OpenDistro clusters (different ES distribution) to OpenSearch.
- Migrated all OpenSearch and OpenDistro clusters to AlmaLinux 9.
- Enabled the k-NN plugin for users interested in VectorDB search use cases.

Hadoop disaster recovery



- Running for several years daily incremental HDFS backups stored on tapes in CTA.
- Missing DR for HBase and Zookeeper. Added hourly snapshots for these tools.
- Snapshots copied to external HDFS clusters for further protection.
- Implementation based on native Hadoop and HBase tools.



Evolution of Monitoring infrastructure at CERN

- Pilot of Prometheus-based metrics service has been open.
- Update to Grafana 10.



Email service

Migration to Exchange Online

- completed as planned end of June 2023

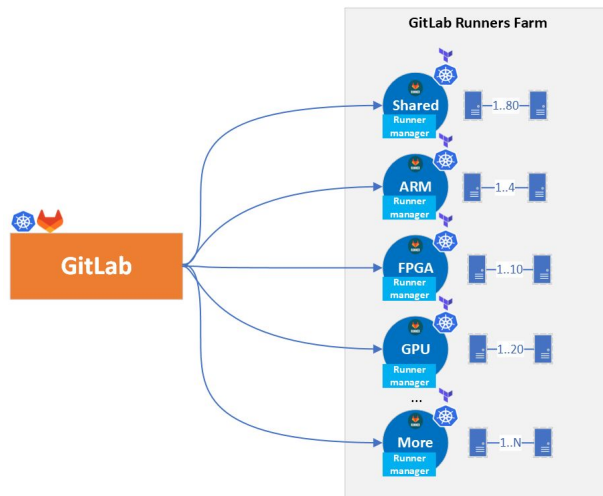
xorlab

xorlab ActiveGuard gradually deployed as a complementary protection measure to Microsoft Exchange Online Protection and Microsoft Defender for Office.





GitLab - New GitLab Runners Architecture

- **New GitLab Runners deployment in place**
- **Docker+machine executor runners cleaning campaign about to finish (ETA October 2023):**
 - Deprecation of the current GitLab Runners based on Docker+machine executor in favor of Kubernetes+executor
 - Taking the opportunity for users to modernize their workflows.
 - Encourage users to use non privileged approaches.
 - Enforce users to use more secured workflows.
 - Security Policies ready to be used.
- **New Architecture for GitLab Runners**
 - GitLab Runners clusters decoupled from the GitLab installation cluster.
 - Several benefits:
 - **Scalability:** GitLab runners can be scaled up independently of the GitLab installation.
 - **Reliability:** If one of the GitLab runners goes down, there are several replicas to get into scene, not impacting the availability of the GitLab installation cluster.
 - **Security:** GitLab runners are isolated from the GitLab installation cluster, and only accessible from the intranet, making them less of a target for attackers.
 - **Flexibility:** Each GitLab Runner serves a specific purpose, and it can be interchanged and scaled up/down upon needs.
- **Automated deployment using Terraform and GitLab**
- **Road to a major GitLab Runners ecosystem**
 - CVMFS, ARM, Spark, TN, etc.



Keycloak - Upgrade and new Architecture

- **Upgrade to version 20.0.5**
 - Sensible improvements in performance
- **Switched from VMs Architecture to Kubernetes one**
 - Why Kubernetes:
 - **Red Hat direction:** Provides a Kubernetes Operator to simplify deployment and replaced Jboss with Quarkus, Java framework designed to run on Kubernetes
 - **Portable:** The only thing you need is a Kubernetes cluster
 - **Reproducible and Immutable:** Speeds up operations, reducing team effort
 - Increased automation and traceability thanks to GitOps approach
 - **Easier to maintain and deploy in long term**
 - Vibrant community supporting Kubernetes
 - Single maintainer for the Keycloak Puppet module
 - Separation of Infinispan cache from Keycloak servers benefits
 - **Monitoring:** Possibility to monitor Keycloak and Infinispan separately
 - **Scalability:** Keycloak and Infinispan scale differently
 - **Flexible deployment model:** keep VMs for Infinispan (but planning to move to K8s as well)
 - **Operations:** User sessions are persisted across Keycloak restarts, helping transparency of future interventions



Kubernetes Service

- **Latest releases: 1.26 and 1.27**

- SeccompDefault now stable, in-place update of Pod resources, provisioning of cross-namespace snapshots, group snapshots

- **Service Highlights**

- [Registry](#): support for artifact signing with cosign, early support for generation and storage of SBOMs (Software Bill Of Materials)
- Runtime: support for ARM nodes, including mixed deployments. Very extensive [series of blog posts](#) on GPU sharing (see also [talk at kubecon](#))
- Security: new [oauth2-refresh-controller](#) for auto refresh of OAuth2 credentials in a cluster - useful for long lived jobs, EOS access
- Storage: [csi-cvmfs](#) with automount support, [csi-eosxd](#) as a replacement of the old eosxd daemonset (deprecated) and including automount and support for OAuth2 based authentication

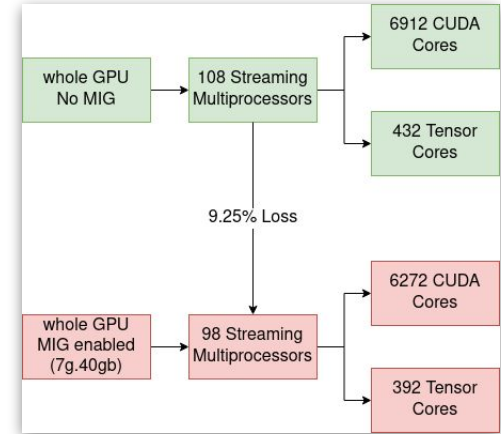
- **Recent Webinars:** <https://indico.cern.ch/category/16678/>

- CVMFS and EOS updates, Full Automation with ArgoCD

- **Major Effort: Kubernetes Service in the TN, Air-Gapped**

- [Collaboration with Accelerator Sector](#), aiming production Q1 2024

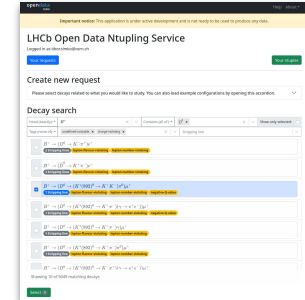
Artifacts	Pull Command	Tags	Signed by Cosign	Size
sha256:4539a049		11.7.1-base-...	✓	337.95MiB
		Accessory		Type
		cosign sha256:58262add		signature.cosign



NAME	KERNEL
continuous-kubernetes-1-26-9-1-rdvxdjocqnia-master-0	6.4.15-200.fc38.aarch64
continuous-kubernetes-1-26-9-1-rdvxdjocqnia-node-0	6.4.15-200.fc38.aarch64

CERN Open Data

- **New open data release**
 - CMS completes Run1 heavy ion open data collection. New release contains about 560 TB of 2013 and 2015 proton-proton reference data, proton-lead collision data and simulations. See the [release announcement](#).
- **LHCb Open Data Ntupling Service**
 - Ongoing integration of LHCb Ntuple Wizard to allow theorists/phenomenologists to ask for custom open data production. Alpha testing phase expected in Q4 2023.
- **“Continuous Reuse”**
 - A new system to periodically execute open data usage examples to ensure future reusability. Testing file presence, log presence, runtime durations. See the [PV2023 talk](#).
- **Forthcoming disk/tape infrastructure updates**
 - Started investigation of a mixed disk/tape storage backend solution to save storage costs for the 4 PB of open data.
 - E.g. 10% data files on disk, 90% on tapes.



The screenshot shows a monitoring dashboard with a table of analysis jobs. The table has columns for Name, Last success, Last failure, Last duration, and five status columns (R1, R2, R3, R4, R5). The rows represent different analysis jobs, with some cells in the status columns highlighted in red, indicating failures or issues.

Name	Last success	Last failure	Last duration	R1	R2	R3	R4	R5
alice-lego-brain-test-run	5 hours ago		00:01:35					
alice-pt-analysis	5 hours ago		00:01:10					
alice-mcstat	5 hours ago		00:01:10					
cms-difusion-misss-spectrum	8 days ago		00:01:03					
cms-difusion-spectrum	5 hours ago		00:01:22					
cms-difusion-spectrum-released	5 hours ago		00:01:12					
cms-hf	5 hours ago		00:02:02					
cms-hf-released	5 hours ago	2 days ago	00:02:45					
cms-hfstat-released	5 hours ago		00:01:58					

REANA reusable analyses

- **REANA 0.9.1 was released on September 28th 2023.**

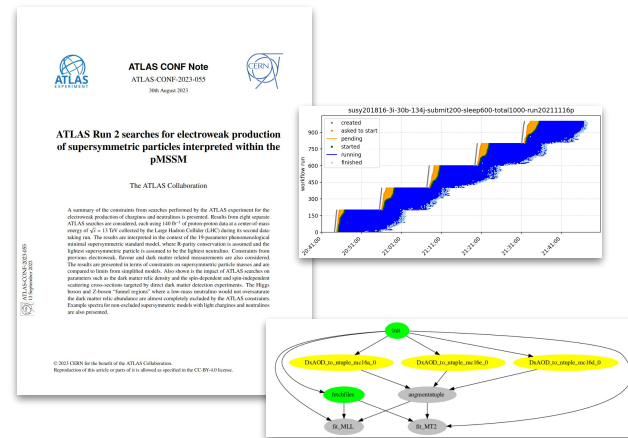
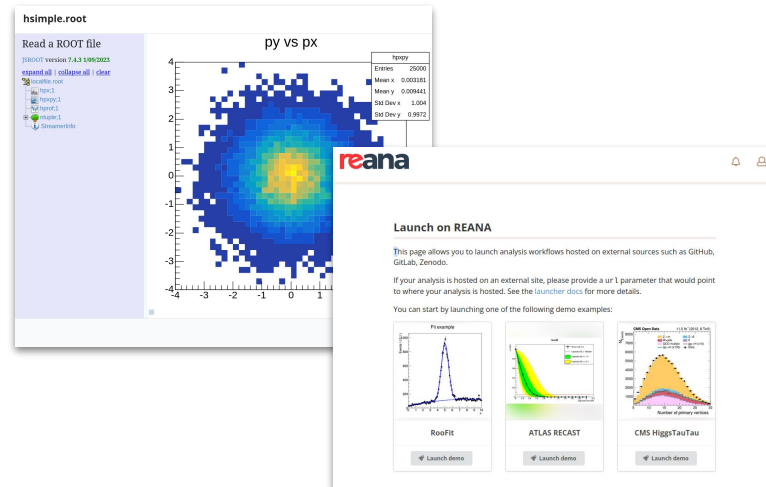
- User-oriented features
 - Visualise PDF and ROOT files in the web interface.
 - Purge temporary files after workflow execution.
 - Improved Snakemake between-run caching of results.
- Admin-oriented features
 - Support any Keycloak IAM. Used by ESCAPE, PUNCH4NFDI.
 - Support Kubernetes clusters from 1.21 to 1.28.
 - Configurable site gallery of launch examples.
 - Auto-closure of inactive interactive notebook sessions.
 - Numerous performance and stability improvements.
- ... and much more. See the [full release notes](#).

- **First results from ATLAS Run2 pMSSM searches using REANA.**

- Running O(10k) workflows to cover enough pMSSM model points.


- **Forthcoming**

- Sharing of workflow runs within research teams.



Tomcat / Web Java Application Hosting

- **Replace jobs and cronjobs infrastructure based on VMs**
 - New solution based on Argo Workflow
 - Deployed in Kubernetes
 - Users can track/manage jobs definitions in git repositories
 - Possibility to use their own Docker images
- **Finalized migration to new SSO based on Keycloak**
- **Upgrading Kubernetes clusters to 1.25**
 - Replacing deprecate ingress controller with new one
 - Upgrade all internal dependencies (Prometheus, ArgoCD, ArgoWF) to latest version
 - Manage internal infrastructure via GitOps
 - Deploying few new applications completely via GitOps paradigm



Atlassian products

- **CERN obtained a Community License for Jira and Confluence and the required plugins**
- **Preparation for the consolidation of all instances of JIRA (5) and Confluence (4) at CERN into single, centrally provided Data Centre instances**



Web Hosting Infrastructure

- **Preparing for CC7 end of life in June 2024**
 - CC7 still used for a majority of web sites with static or CGI content
 - Central web hosting services will transition directly to Alma9
 - Notable differences between the CC7 and Alma9 environments:
 - PHP 5 replaced by PHP 8
 - Python 2 not available
- **Started migration campaign for web sites on CC7**
 - 4000 web sites with source files on EOS
 - Provided “previews” of the same content served by Alma9
 - Site owners invited to check previews work properly
 - One-click UI for site owners to switch live sites between CC7/Alma9
 - 2300 web sites with source files on AFS
 - Will not update AFS web servers
 - Move site content to EOS instead
 - Special cases will be followed up with the site owners



InvenioRDM and Zenodo

Zenodo migrated to InvenioRDM platform on October 13th

- Major release focused on making collaboration easy for researcher including support for community members and curation, reviews of upload and new possibilities for sharing records with e.g. peer-reviewers
- <https://help.zenodo.org/docs/about/whats-new/>

zenodo Search records... Communities My dashboard lars.holm...

Biodiversity Literature Repository
https://www.biosyslit.org Topic Plazi

Search Requests Members Settings

Back to requests Accept and publish Revoke

Zenodo data and software citation links captured by the Asclepias Broker
Ioannidis, Alex

Conversation Record

Alex Ioannidis commented 1 minute ago: I'd like to have this dataset included in your community.

Lars Holm Nielsen commented 3 seconds ago: Thanks a lot for the submission.

Request type: Draft review

Status: Submitted

Created: 52 seconds ago

zenodo Search records... Communities My dashboard lars.holm...

Biodiversity Literature Repository
https://www.biosyslit.org Topic Plazi

Search Requests Members Settings

Members

Search in members ... Role Visibility Name

0 members selected Member since Visibility Role Invite...

Member since	Visibility	Role
2 minutes ago	Hidden	Curator
5 minutes ago	Public	Owner

2 result(s) found 1 25 results per page

Digital Memory project







A Preserve platform with a web front-end has been implemented and rolled out: <http://preserve.web.cern.ch>

A Digital preservation strategy for IT and for CERN has been drafted, to help setting up a new Operational Circular ruling digital archival practices

Welcome to Preserve, the CERN Digital Memory Platform









This service allows to retrieve your digital assets from CERN digital repositories, request their archival and prepare them for long term preservation, complying to the OAIS specifications.

Features overview

- 1 Submit**

Users can trigger harvesting of records from supported CERN digital repositories.
[SUBMIT](#) 
- 2 Preserve**

The platform will prepare the data for long term preservation, creating Archival and Dissemination packages and saving them on the CERN Tape Archive.
[PRESERVE](#) 
- 3 Access**

Processed data is finally recorded in our Preservation Registry, where it is available to the authorized people.
[ACCESS](#) 





Harvest Data


Here's how to get started and make the platform retrieve records from CERN supported repositories for you.

- 1 Configure**

To be able to fetch your private records you'll need to set up some API tokens first.
[CONFIGURE](#) 
- 2 Search**

You can search for any records on Indico, CodiMD, CDS and Zenodo from the Search page. The platform will fetch the data for you.
[SEARCH](#) 
- 3 Organize**

Once you selected your records, they will be waiting in the 'Staging Area', where you can organize them with custom tags.
[ORGANIZE](#) 
- 4 Download**

Once you confirmed your selection from the staging area, the preservation process will start. Check after a while to get download links.
[DOWNLOAD](#) 

Sources availability

Here's an overview of the available repositories to harvest from. Some of them may need additional [configuration](#).

Source	Configuration Status
CERN Document Server	 Only public records will be available. Configuration is needed for private records.
CodiMD	 Source unavailable. Additional configuration is needed.
Indico	 Source ready
Zenodo	 Source ready



home.cern