

CERN site report

HEPiX Autumn 2022

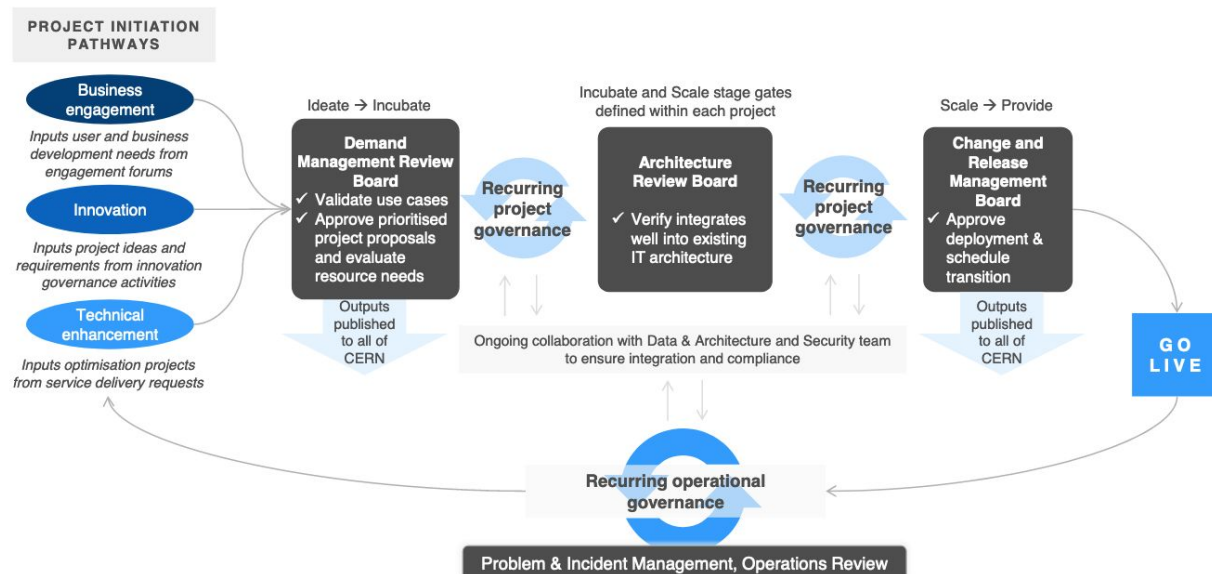
31 October - 03 November, Umeå, Sweden

Jarek Polok <jaroslaw.polok@cern.ch>

On behalf of CERN IT

CERN IT department reorganised

The IT Department's vision is to reinforce its position as a trusted and efficient IT service-provider and technology partner for the CERN community, and to be recognised as a reference catalyst driving collaboration and innovation in the scientific computing environment in the member states and beyond.



Structure, Vision, Strategy, Projects

CERN Data Centre(s)

Prévessin Data Centre (PCC):

- Construction started in March of this year
- Internal works have started, e.g. lighting, painting
- Installation of ventilation ducts and piping starting now
- Building expected to be sealed (doors installed) by the end of the year
- Commissioning due to start in the summer of 2023
- End of commissioning foreseen for October 2023



Hardware deliveries:

- Quite challenging in 2022 due to overall supply chain situation
- Average dispatch time between order and delivery between 16 and 25 weeks
- Some supplies still not delivered after more than 12 months

talk this week: [CERN IT Hardware Procurement activities in 2020-2022](#)

Other:

- Preparing for (un)scheduled power cuts this winter
- OpenDCIM extensively used for planning and de(commissioning)

Year	Quarter	Item	Quantity	Status	Delivery Date	Location
2023	Q3	ARM (ampere) servers	4 + 1	Delivery Accepted	Jan-22	10 October 2022 at openstack COP
N/A		US-ATLAS setup	1	Delivery pending		
2022	Q4	All-flash servers	10	Delivery pending	Delayed - Unknown	CEPH
2022	Q2	IT DB servers (ORACLE)	10	Rail kits delivered	Delayed - Unknown	IT-DB
2022	Q3 / 2023	Q3				
IT-4614/IT		EOS/AliceO2 front-ends	25	Delivery pending	Q4 2022	EOS/AliceO2_pibatch
IT-4615/IT		EOS/AliceO2 JBODs	50	Installed	Q3 2022	
IT-4615/IT		EOS/AliceO2 JBODs	12	Installed (R/150)	Q3 2022 / Q1 2023	
IT-4614/IT		Ceph front-ends	5	Delivery pending	Q4 2022	
IT-4614/IT		Ceph front-ends	11	Delivery pending	Q4 2022	Ceph
IT-4615/IT		Ceph JBODs	46	Installed		
IT-4615/IT		Ceph JBODs	8	Delivery pending	Q1 2023	
IT-4614/IT		DBDD servers	11	Delivery pending	Q4 2022	
IT-4614/IT		HADOOP front-ends	21	Delivery pending	Q4 2022	HADOOP
IT-4615/IT		HADOOP JBODs	6	Installed	Q4 2022	
IT-4614/IT		NXCALS front-ends	14	Delivery pending	Q4 2022	
IT-4615/IT		NXCALS JBODs	9	Installed	Q4 2022	NXCALS
IT-4614/IT		GESMAR servers	2	Delivery pending	Q4 2022 / Q1 2023	EN-EL
IT-4614/IT		Windows servers	4	Delivery pending	Q4 2022 / Q1 2023	IT-G3
IT-4614/IT		DFS servers	16	Delivery pending	Q4 2022 / Q1 2023	IT-SD
IT-4615/IT		DFS storage	8	Delivery pending	Q4 2022 / Q1 2023	IT-SD

talk this week: [Managing CERN data centers with OpenDCIM](#)

Cloud Infrastructure Migration Campaign

What: Move from deprecated to supported network control plane in OpenStack


Why: Unblock upgrades and enable evolution of network features



When: Since November 2021, just finished!

4'400 VMs moved!

talk this week: [Cloud Infrastructure Update: Operations, Campaigns and Evolution](#)



CentOS Stream Status



Re-evaluation of initial recommendation to use CentOS Stream for production:

- Experiences after ~1 year of running services on CentOS Stream
- State of the Red Hat Enterprise Linux rebuilds



Continue to work with the HEP community and encourage feedback from the community!


talk this week:

[CERN Linux Updates and Outlook](#)

&

BoF session:

[Recent experiences with CentOS Stream](#)



SLC6	329	1.2%
CC7	20,680	76.9%
CL8	1,155	4.3%
CS8	3,373	12.5%
CS9	834	3.1%
RHEL7	143	0.5%
RHEL8	147	0.5%
RHEL9	220	0.8%

Storage services

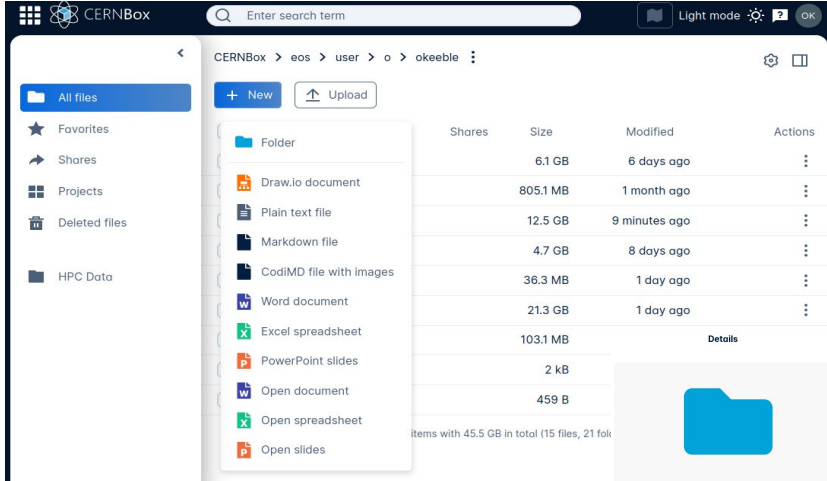
CERNBox:

- **New web interface** with major improvements
- **Integration with MS Office 365 and Markdown (CodiMD) for collaborative editing**
- **New backup orchestrator deployed (based on RESTIC) for EOS home directories and projects, writing to S3, now extended to also back up CephFS shares**

talk this week: [CERNBox: sync, share and science](#)

Ceph:

- Major production clusters upgraded to **Ceph Pacific**
- Capacity expansion of 2nd region Object Storage cluster to 27 PB
- Consolidated storage offer and introduction of Storage Availability zones for block storage
- Ongoing investigation of Ceph replication features for Business Continuity / Disaster Recovery plans



The screenshot shows the CERNBox web interface. The top navigation bar includes the CERNBox logo, a search bar, and a light mode toggle. The main content area displays a file explorer view for the path 'CERNBox > eos > user > o > okeoble'. A 'New' dropdown menu is open, showing options like 'Folder', 'Draw.io document', 'Plain text file', 'Markdown file', 'CodiMD file with images', 'Word document', 'Excel spreadsheet', 'PowerPoint slides', 'Open document', 'Open spreadsheet', and 'Open slides'. A table of files is visible with columns for Shares, Size, Modified, and Actions. The table lists several files with their respective sizes and modification times. A 'Details' panel is open for a folder, showing its last modified date (Oct 24, 2022, 10:58 AM), owner (me), size (54.7 GB), FUSE Path (/eos/user/), Windows Path (\\cernbox-smb(eos)user/), and Direct link (https://cernbox.cern.ch/files/spaces/eos/...). The CERNBox logo is also visible below the screenshot.

Shares	Size	Modified	Actions
	6.1 GB	6 days ago	⋮
	805.1 MB	1 month ago	⋮
	12.5 GB	9 minutes ago	⋮
	4.7 GB	8 days ago	⋮
	36.3 MB	1 day ago	⋮
	21.3 GB	1 day ago	⋮
	103.1 MB		
	2 kB		
	459 B		



Storage services



DFS:

- Review of use-cases, migration of user areas to EOS+SAMBA complete: almost 46000 accounts



CTA:

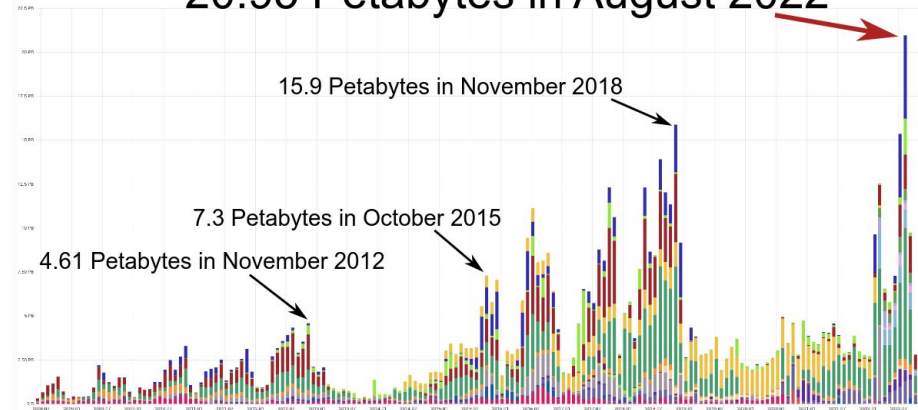
- Nominal performance/efficiency as designed
 - Archival at 10 GB/s per LHC experiment consistently reached
- CASTOR has been decommissioned
 - official EOL 31/10/2022



New CERN Record

Data written to CTA from 1/08/2022 to 31/08/2022

20.95 Petabytes in August 2022



talk this week: [Storage for the LHC: operations during Run 3](#)

Database and Analytics

Start of Next Oracle LTS (long term support) version beta test - Oracle 23c

- 23c will be the version supported for the HL-LHC run. Gradually we want to test the upgrade impact on all critical databases in order to prepare for HL-LHC run

ORACLE®

More than 600 [DBoD](#) databases (MySQL, pgSQL, InfluxDB) upgraded to recent supported releases



- Thanks to automation, large part of Database on Demand landscape was upgraded to higher versions without significant impact to the service



Oracle instances consolidation in preparation for licensing schema change in April 2023:

- As the end of campus license contract nears Database service is consolidating Oracle instances on powerful high frequency/low core count systems to limit the financial impact of Oracle licensing. Consolidation includes (but is not limited to):
 - Review and optimization of existing DataGuard and Active DataGuard configurations
 - Utilization of Oracle RAC One Node as replacement of full Oracle RAC
 - Decommissioning of legacy infrastructure (e.g. GoldenGate replication)



Database and Analytics

ElasticSearch service:



- Migrating all production clusters to [OpenSearch](#) (running v2.2.1)
- New OpenSearch clusters running on physical machines with SSDs and Ceph storage
- Plan is to migrate all legacy ElasticSearch 6/7 clusters by Q1 2023

Hadoop service:



- Testing a dedicated HBase deployment for NXCALS to improve performance/robustness
- Evaluated Apache Ozone as a block storage alternative to HDFS

[SWAN](#) service:



- The interactive data analysis service is gradually being migrated to a Kubernetes deployment. The Kubernetes installation is part of the production service and automatically serving 60% of users regularly. Some GPUs are available to users thanks to this new deployment

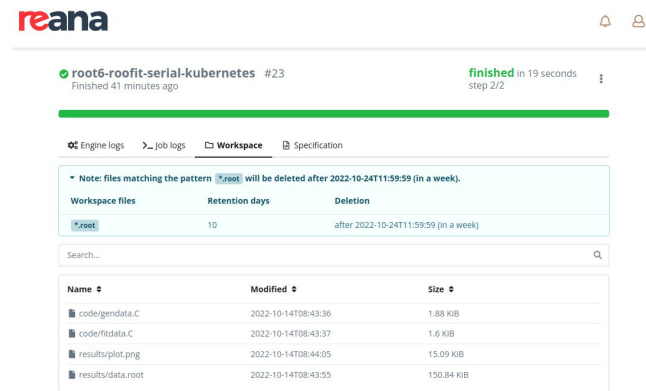
Applications and Reusable Workflows

REANA - Reusable Analyses:

- New user-configurable workspace file retention rules
- Example: retain large temporary "*.root" files for ten days only
- New support for Rucio data access
- Exposure of workflow engine logs and job duration in the web interface
- REANA 0.9.0 official release scheduled for November 2022

JEEDY:

- New cronjob infrastructure based on Argo Workflow
- Wide adoption of GitOps and ArgoCD for internal infrastructure components
- efiles.cern.ch managed fully via GitOps
- More apps and standard approach are coming
- Enabling secret management via teigi



The screenshot shows the REANA web interface for a workflow execution. At the top, the workflow is identified as 'root6-roofit-serial-kubernetes #23', which is 'finished' in 19 seconds, step 2/2, and was completed 41 minutes ago. Below this, there are tabs for 'Engine logs', 'Job logs', 'Workspace', and 'Specification'. A note indicates that files matching the pattern '*root*' will be deleted after 2022-10-24T11:59:59 (in a week). A table lists workspace files with their retention days and deletion dates:

Workspace files	Retention days	Deletion
root	10	after 2022-10-24T11:59:59 (in a week)

Below the table is a search bar and a list of files with columns for Name, Modified, and Size:

Name	Modified	Size
code/gendata.C	2022-10-14T08:43:36	1.88 KiB
code/fitdata.C	2022-10-14T08:43:37	1.6 KiB
results/plot.png	2022-10-14T08:44:05	15.09 KiB
results/data.root	2022-10-14T08:43:55	150.84 KiB



EDH



HRT



infor



EDMS



APT



CERN Phonebook



BAAN



Greybook



CET



IMPACT



e-groups

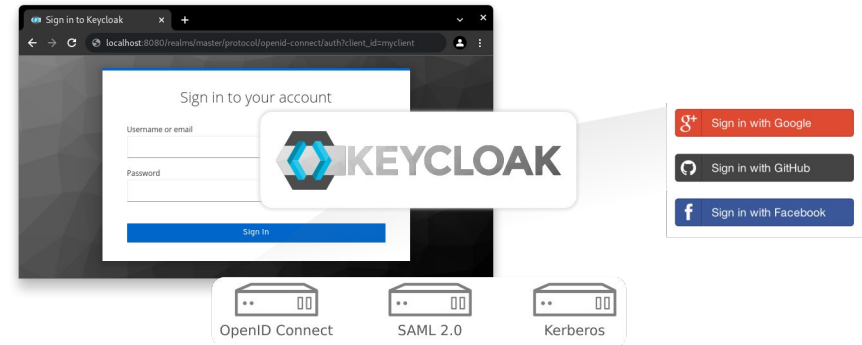
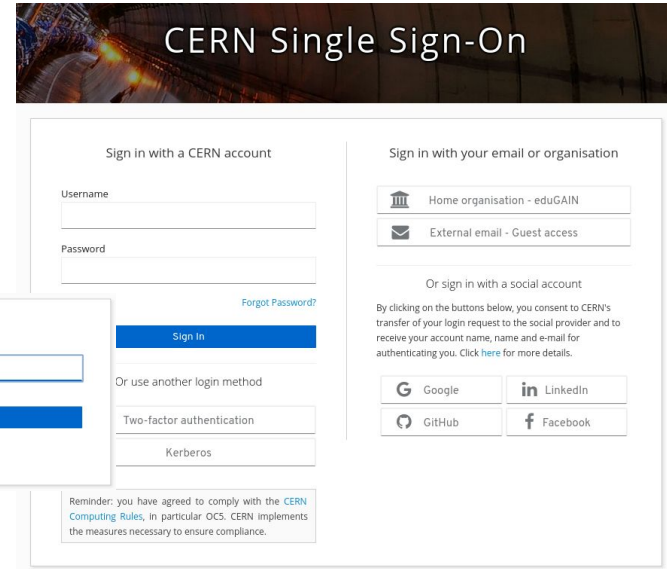
Identity and Access Management

- 2FA rollout in progress
- SSO migration
- Old SSO to be decommissioned summer 2023 (software end of life)
- 8,500 applications on new SSO
- Some service incidents now that high density traffic events protected by new SSO:

architecture being reviewed urgently

- Contact and collaboration with other high traffic Keycloak users would be very welcome!

talk this week: [*Enforcing Two-Factor Authentication at CERN: A Technical Report on Our Experiences with User Migration*](#)



Platforms Infrastructure

PaaS: migration from Openshift3 to OKD4 complete

- **Openshift3 EOL June 2022**
- **Used the Openshift Migration Toolkit for Containers**
- **Documentation sites now served by GitLab Pages: MkDocs, and other static page generators are supported**



Kubernetes / Magnum:

- **Kubernetes 1.24 and 1.25**
- **Support for dual stack networking (ipv4/ipv6)**
- **New CVMFS CSI driver with automount support**
- **Extended support for GPUs, including Multi-Instance-GPU (MIG) for A100**
- **Support for lazy image pulling (estargz) in both builds and deployments**



talk this week: [Kubernetes operators for web hosting at CERN: experience](#)

GitLab:

- **Infrastructure migration from OpenShift 3 to Kubernetes complete**
- **From unsupported custom-made infrastructure to supported Cloud Native Hybrid one**



Collaborative Services

Indico: [Version 3.2](#) (June at CERN, public release in August)

- New batch of privacy-enhancing features
- Updated minutes editor (+ markdown support)
- Event series management interface
- CAPTCHAS in reg. forms and user registration

Transcription / Translation as a Service:

- Ongoing project for a new ASR system based on technology from MLLP (Universitat Politècnica de València)
- Better performance than commercial alternatives, with CERN's vocabulary and accents
- Integration with Zoom, webcast/recording services (both live and offline)
- Web APIs and interface for ad-hoc usage
- Pilot being set up, available before end of year



coming soon!

The screenshot shows the Indico registration form configuration interface. At the top, there is a CAPTCHA section with the text "Confirm that you are not a robot" and a visual CAPTCHA showing the number "9643". Below this, there are several configuration panels:

- Registration Form:** Includes a "Configure" button.
- Privacy settings:** Includes a "Configure" button and the text "Choose who can see the participants of this event."
- List of registrations:** Includes a "Manage" button and the text "Add, remove or modify registrations."

A red box highlights the "Retention period" section, which contains the text: "Specify for how many weeks the registration data, including the participant list, should be stored. Retention periods for individual fields can be set in the registration form designer".



The screenshot shows the Indico Privacy dashboard for a talk titled "My awesome talk 31 Mar". The dashboard includes a sidebar with navigation options: Settings, Protection, Privacy, Organization, Room booking, Services, Reports, Customization, and Advanced options. The main content area is titled "Privacy dashboard" and includes sections for "Data controller", "Person/Institution", "Contact email", "Privacy notice", "External page", and "Text". A "Privacy Information" modal is open, showing "Overview", "Timetable", and "Privacy Information" tabs. The "Privacy Information" tab displays the "Data controller" as John Doe with the email address controller@example.com. A "Privacy notice" section is also visible at the bottom of the modal.

Digital Repositories

Zenodo:



- REANA badge to launch analysis directly from Zenodo integrated
- US National Institute of Health have funded Zenodo to enhance access to biomedical research
- New counter measures have significantly reduced spam

Digital Memory:



- Consolidating Long-term Preservation platform into a production service
- Supports lots of different sources of content:
 - repositories (CDS, Zenodo), Indico, Gitlab, CERNBox, CodiMD, etc
- Integrated with Archivematica to support the OAIS model
 - e.g. CASTOR software archived in the platform

InvenioRDM:



- New common platform for Zenodo and CERN Document Server
- New features: Community members support, review workflow, custom fields and administration interface
- Larger production sites launched: TU Graz, Caltech Data

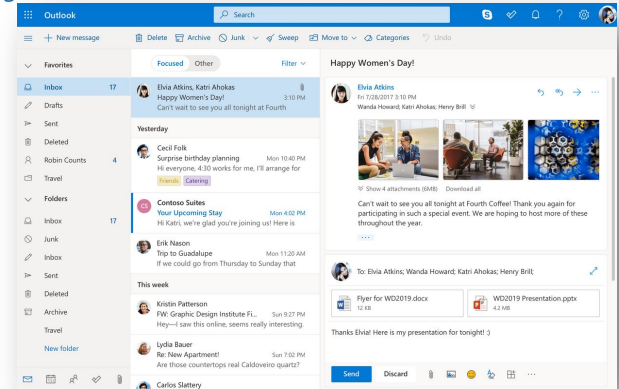
Email and Engineering Software Services

Email Service: Following [MALT closure](#), move to [Exchange Online](#):

- Upgrade of the current Email platform
- Same functionality, same supported clients, renewed Outlook for the web
- Ongoing migration
- Migrating emails, calendar items, contacts and tasks
- From August 2022 to mid-2023

Engineering Software Service:

- New CAD software made available for 2D design: [BricsCAD Pro](#)
- Alternative to AutoCAD



CERN presentations this HEPiX

- [*Kubernetes operators for web hosting at CERN: experience*](#) - Alexandre Lossent
- [*CERN Linux Updates and Outlook*](#) - Alex Iribarren
- [*Scalable Machine Learning at CERN with Kubeflow*](#) - Dejan Golubovic
- [*Automation of \(Remediation\) procedures for Batch Services in CERN with StackStorm*](#) - Ankur Singh
- [*Cloud Infrastructure Update: Operations, Campaigns, and Evolution*](#) - Domingo Rivera Barros
- [*Transferring data from ALICE at 2.4Tbps with DWDM*](#) - Daniele Pomponi
- [*Enforcing Two-Factor Authentication at CERN: A Technical Report on Our Experiences with User Migration*](#) - Adeel Ahmad
- [*Computer Security Landscape Update*](#) - Christos Arvanitis
- [*CERN IT Hardware Procurement activities in 2020-2022*](#) - Luca Atzori
- [*Managing CERN data centers with OpenDCIM*](#) - Jarek Polok
- [*CERNBox: sync, share and science*](#) - Diogo Castro
- [*Storage for the LHC: operations during Run 3*](#) - Cedric Caffy



home.cern