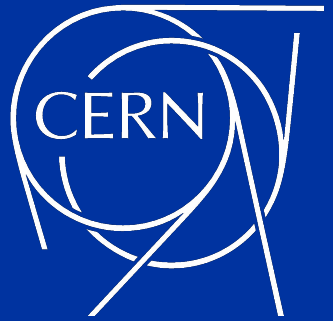

Deploying Apache Ranger in the Big Data ecosystem at CERN

Emil Kleszcz @ CERN

09 Oct 2023

Presentation content

1. About CERN
2. Big Data ecosystem @ CERN
3. Overview of Apache Ranger project
4. Apache Ranger @ CERN
5. Summary

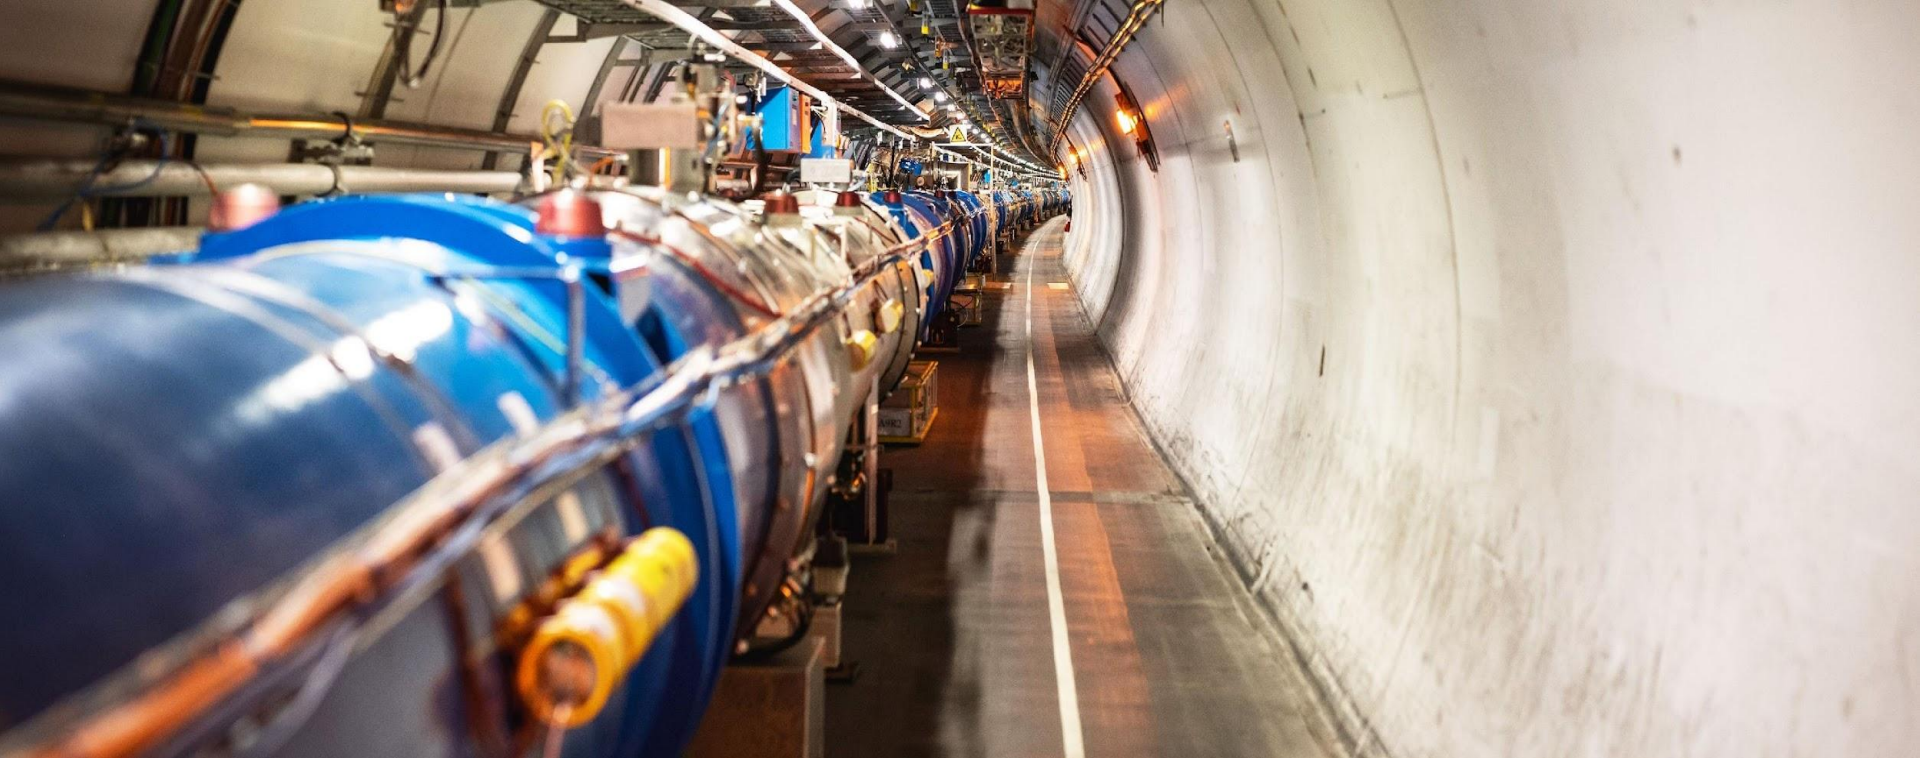

About CERN

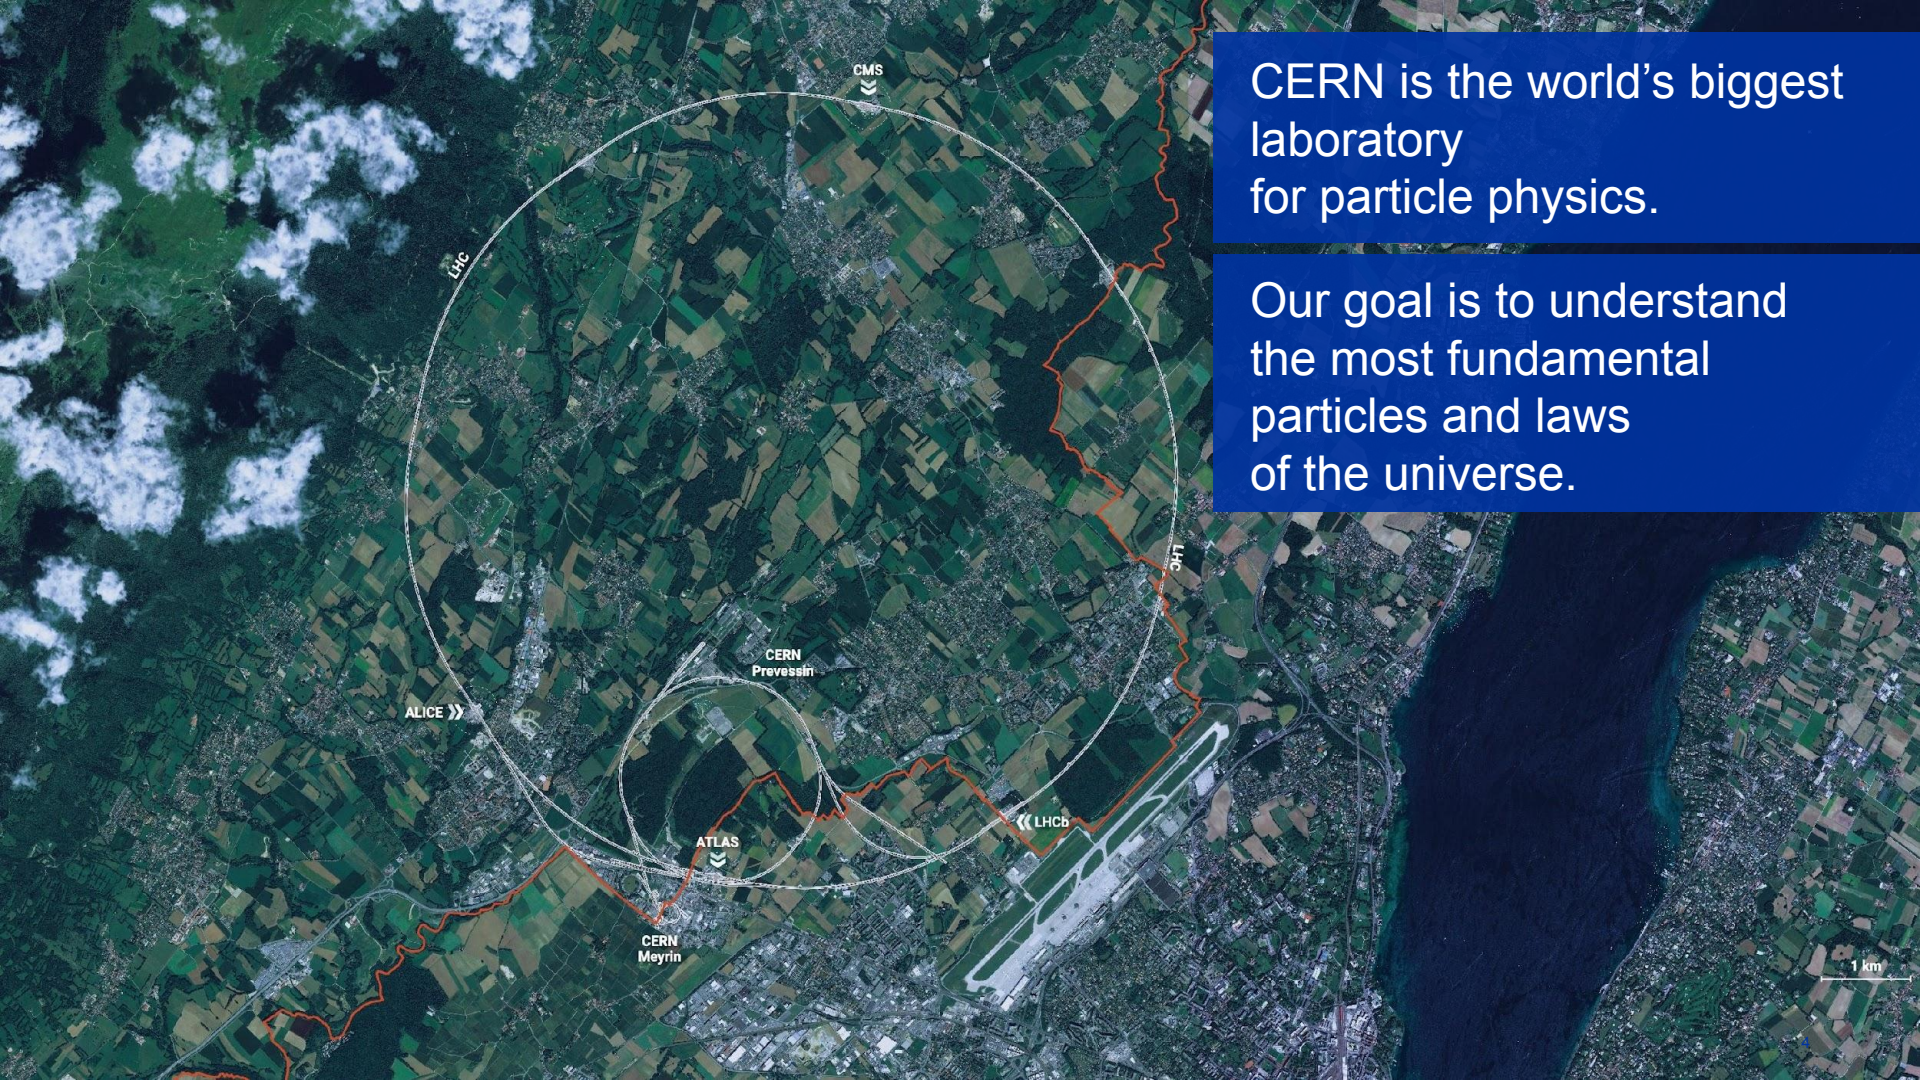

CERN is the world's biggest laboratory
for particle physics.

Our goal is to understand the most fundamental particles and laws of the universe.

COLLABORATION

CERN was founded in 1954 with 12 European Member States

23 Member States

Austria – Belgium – Bulgaria – Czech Republic
Denmark – Finland – France – Germany – Greece
Hungary – Israel – Italy – Netherlands – Norway
Poland – Portugal – Romania – Serbia – Slovakia
Spain – Sweden – Switzerland – United Kingdom

3 Associate Member States in the pre-stage to membership

Cyprus – Estonia – Slovenia

7 Associate Member States

Croatia – India – Latvia – Lithuania – Pakistan
Türkiye – Ukraine

6 Observers

Japan – Russia (suspended) – USA
European Union – JINR (suspended) – UNESCO

Brazil signed Associate Membership Agreement on 3 March 2022. Ratification process in Brazil is underway.

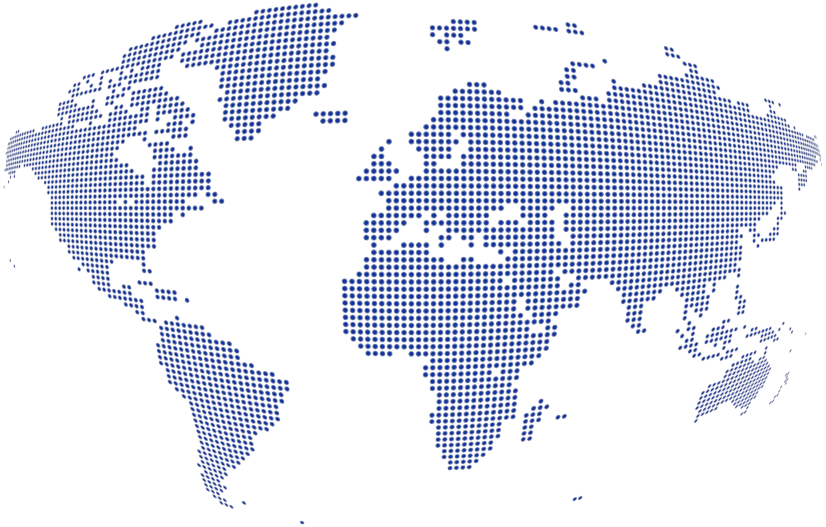

Around 50 Cooperation Agreements with non-Member States and Territories

CERN's annual budget is 1200 MCHF (equivalent to a medium-sized European university)

As of 31 December 2021
Employees:
2676 staff, **783** fellows

Associates:
11175 users, **1556** others

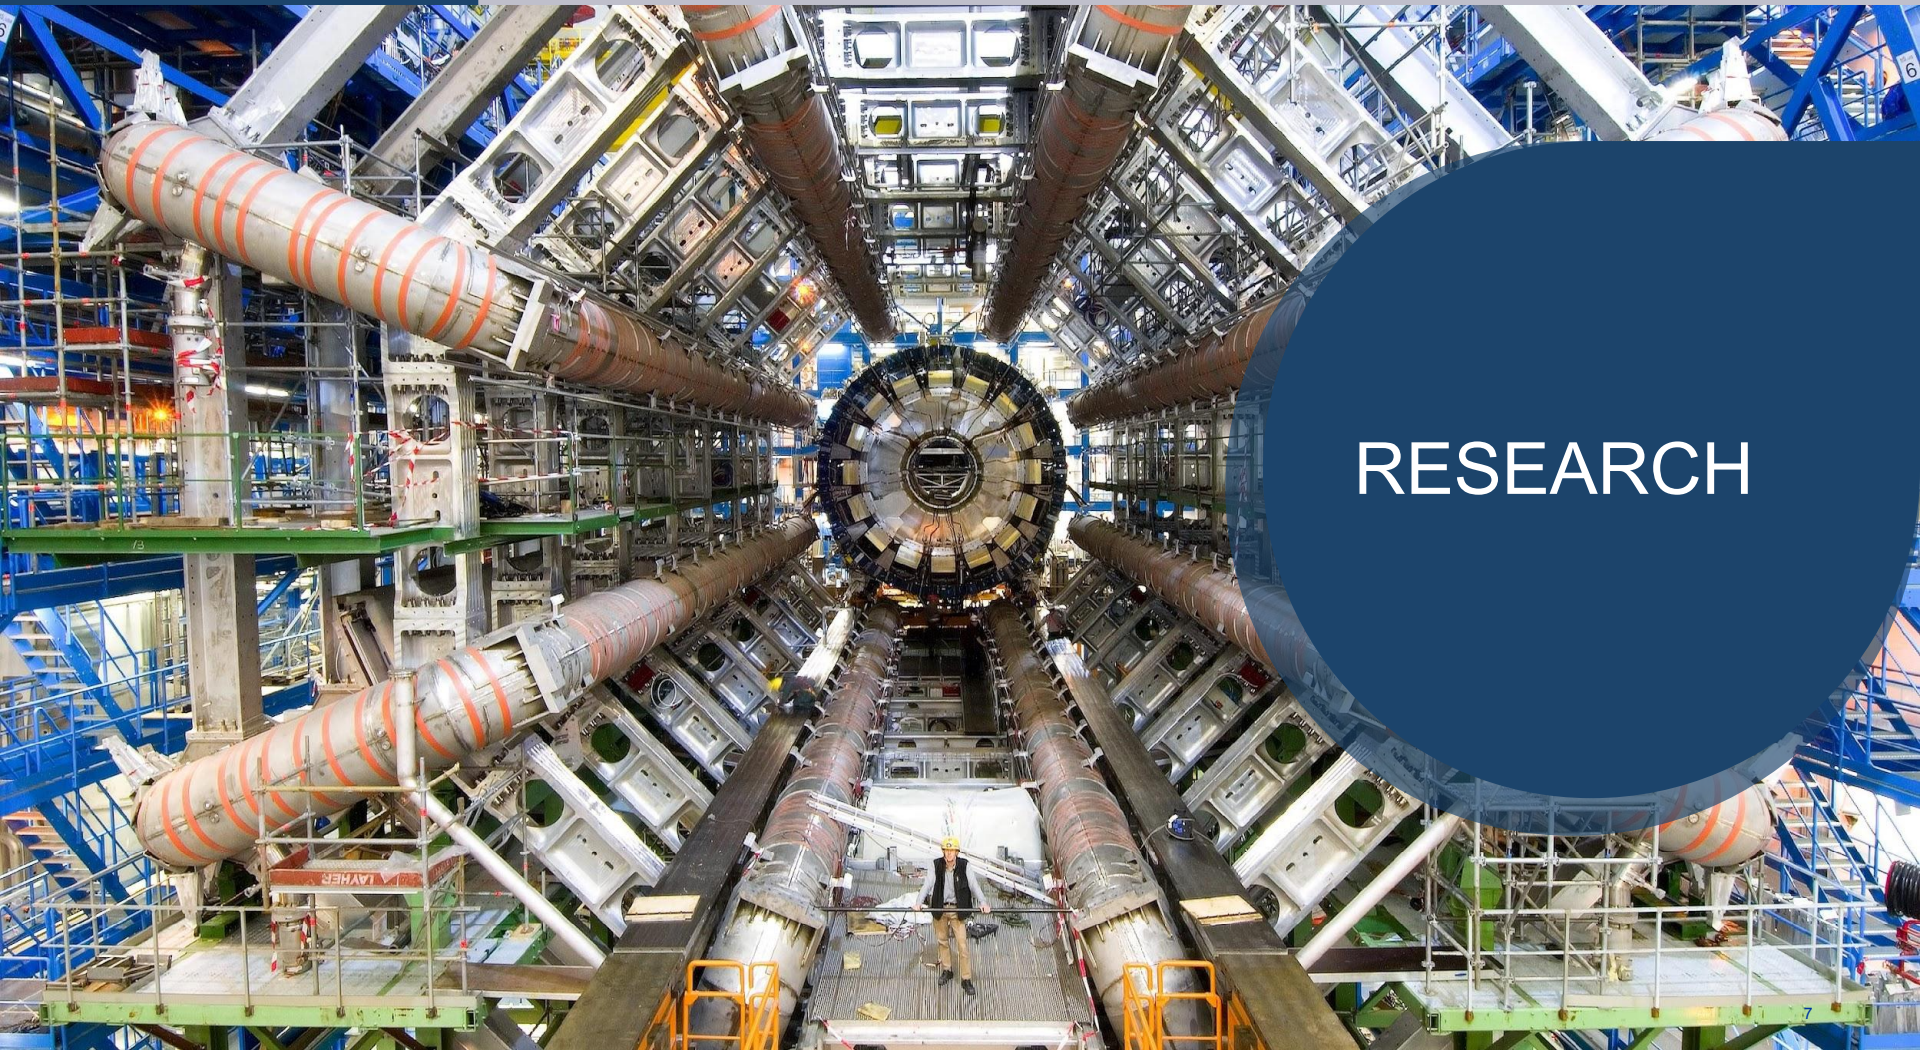

RESEARCH

What is the universe made of?

We study the elementary building blocks of matter and the forces that control their behavior

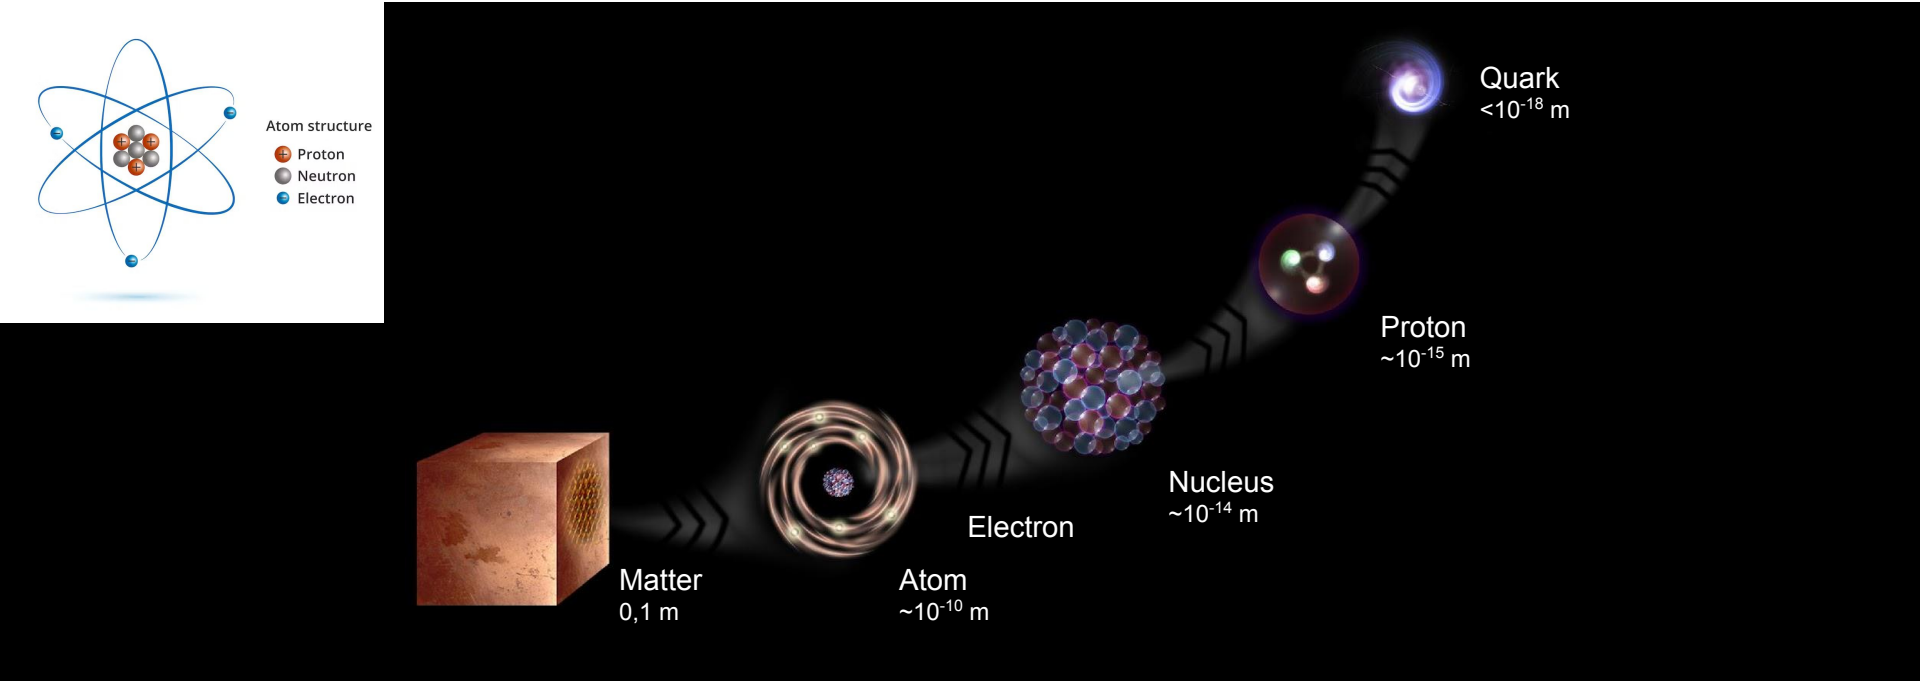

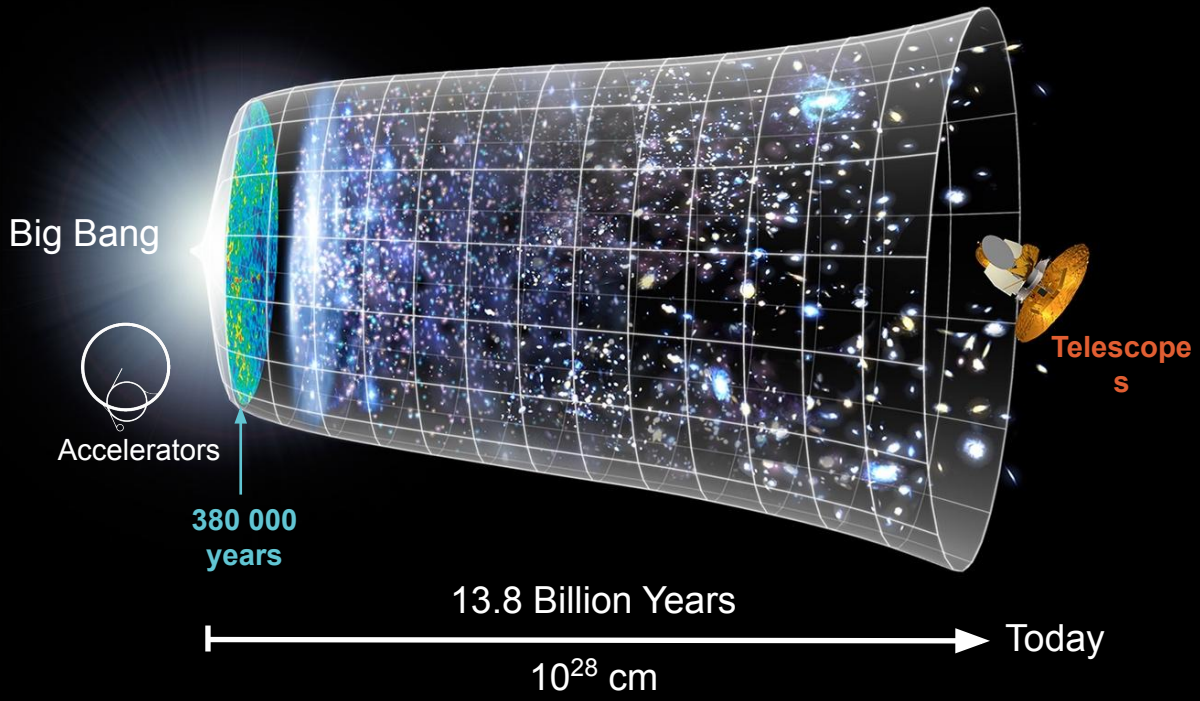

How did the universe begin?

We reproduce the conditions a fraction of a second after the Big Bang, to gain insight into the structure and evolution of the universe.

Large Hadron Collider (LHC)

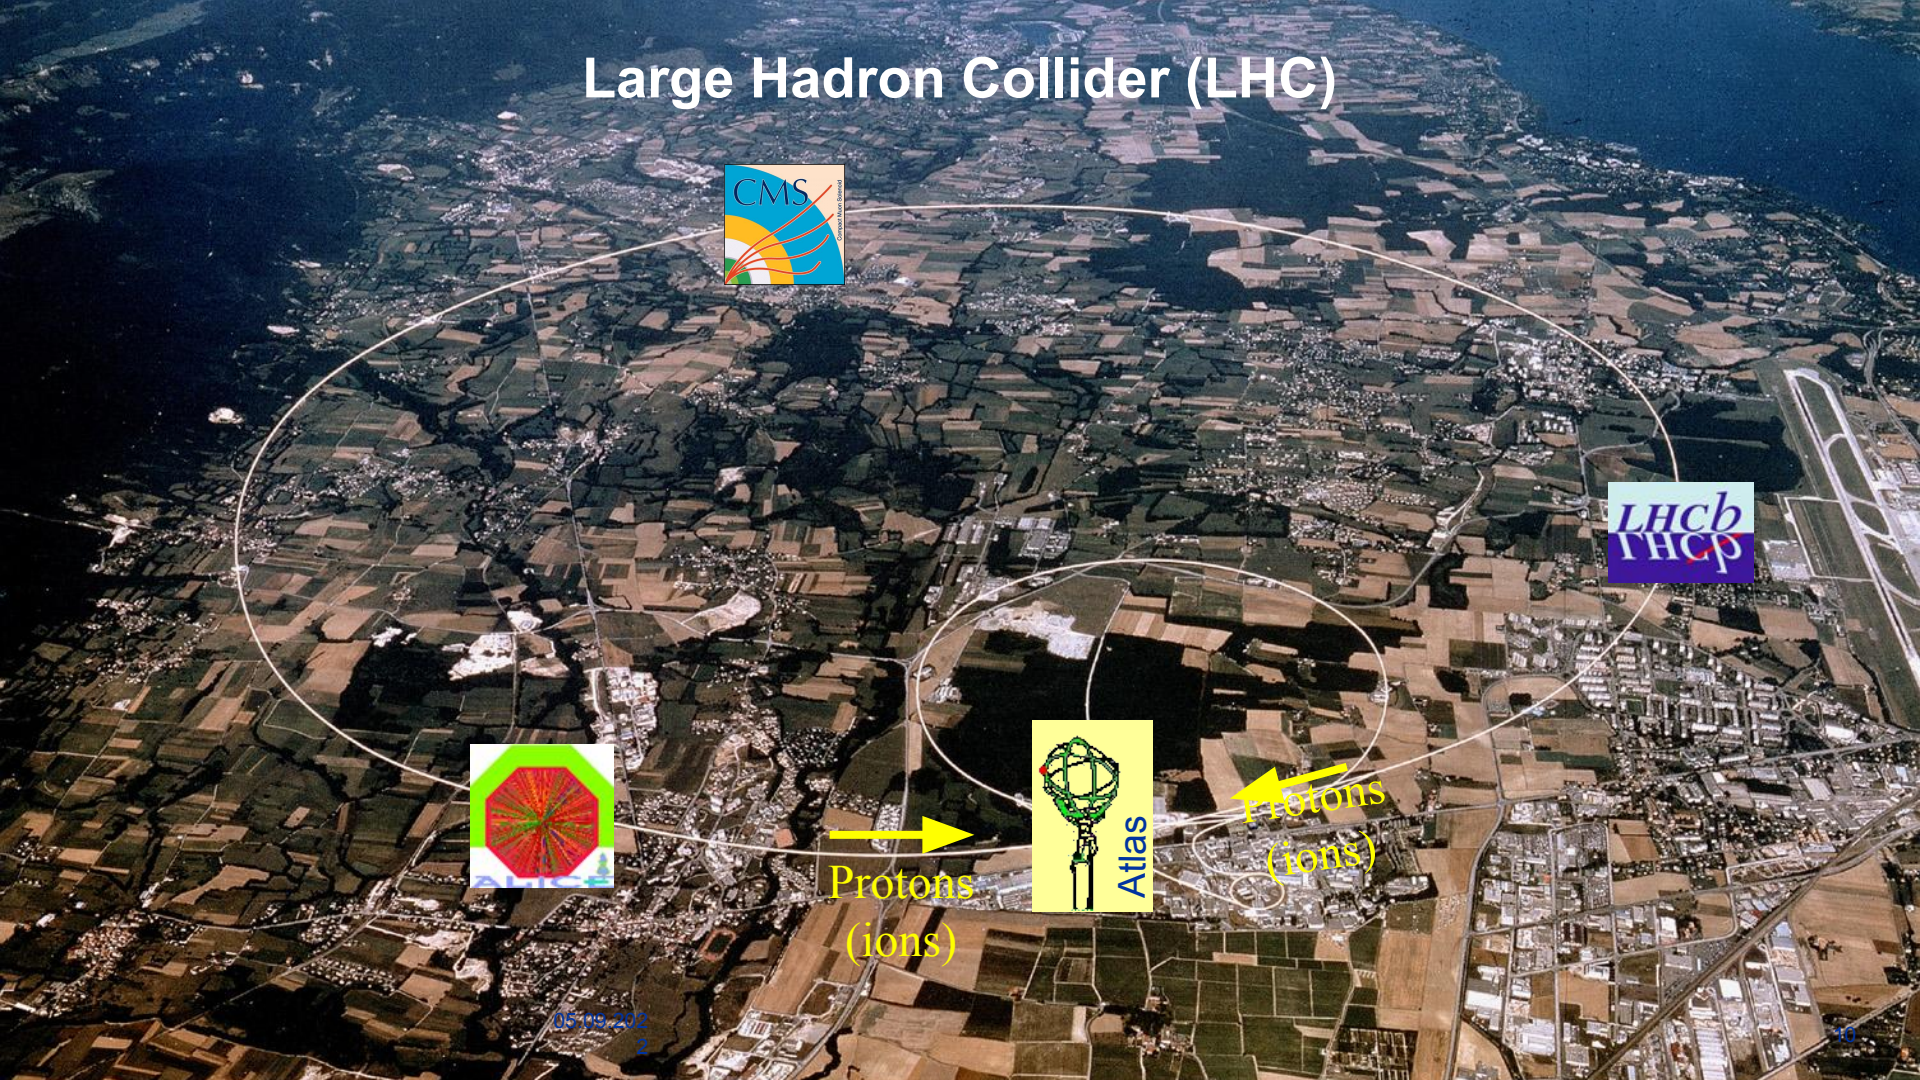

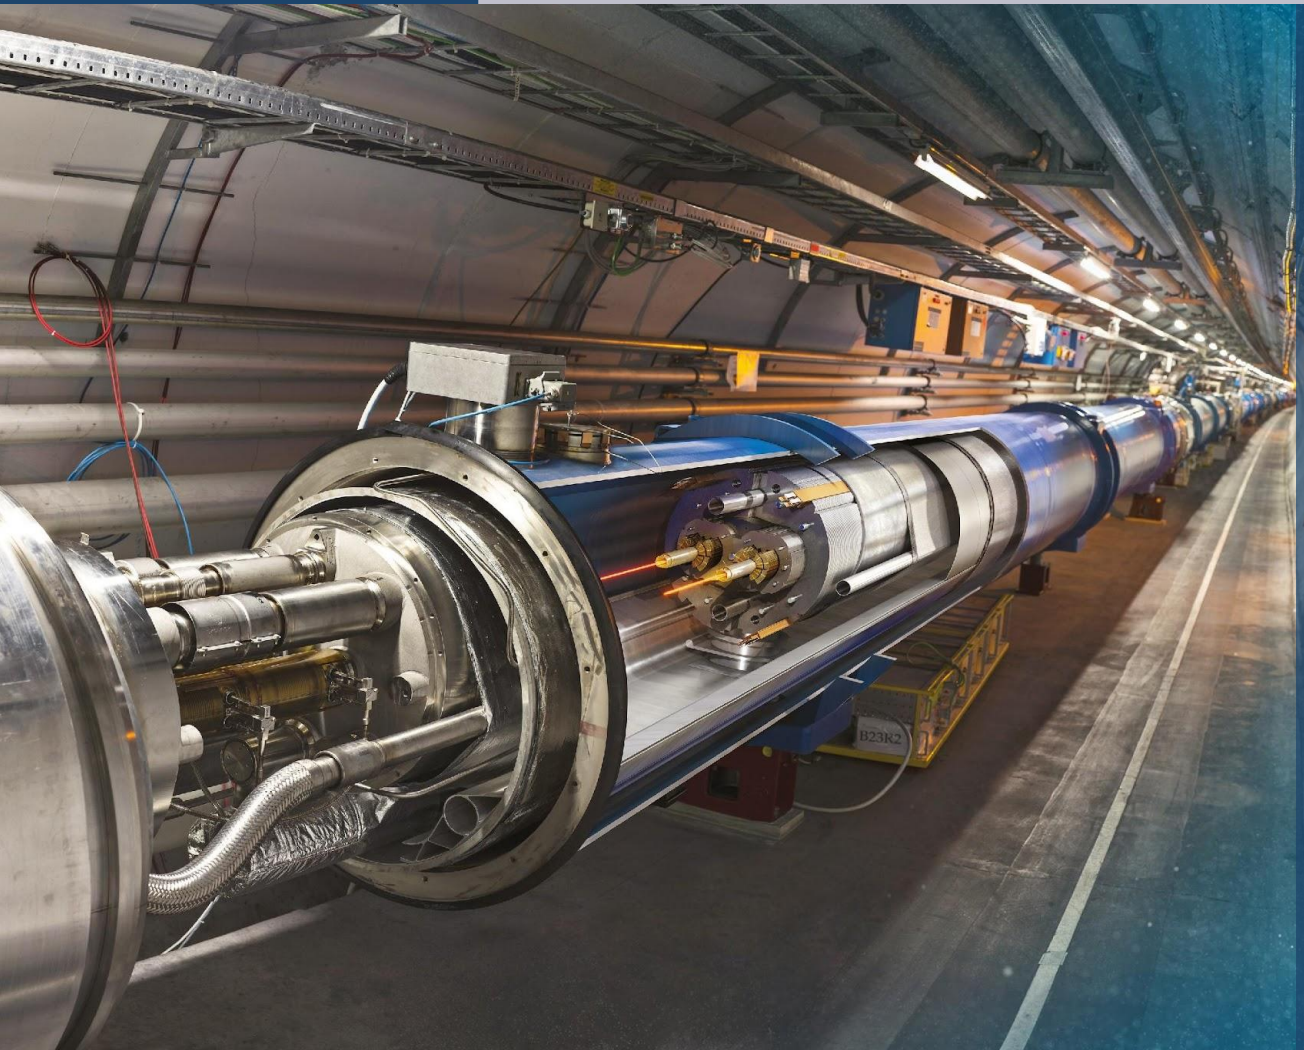

Large Hadron Collider (LHC)

- 27 km in circumference
- About 100 m underground
- Superconducting magnets steer the particles around the ring
- Particles are accelerated close to the speed of light

CERN Accelerators

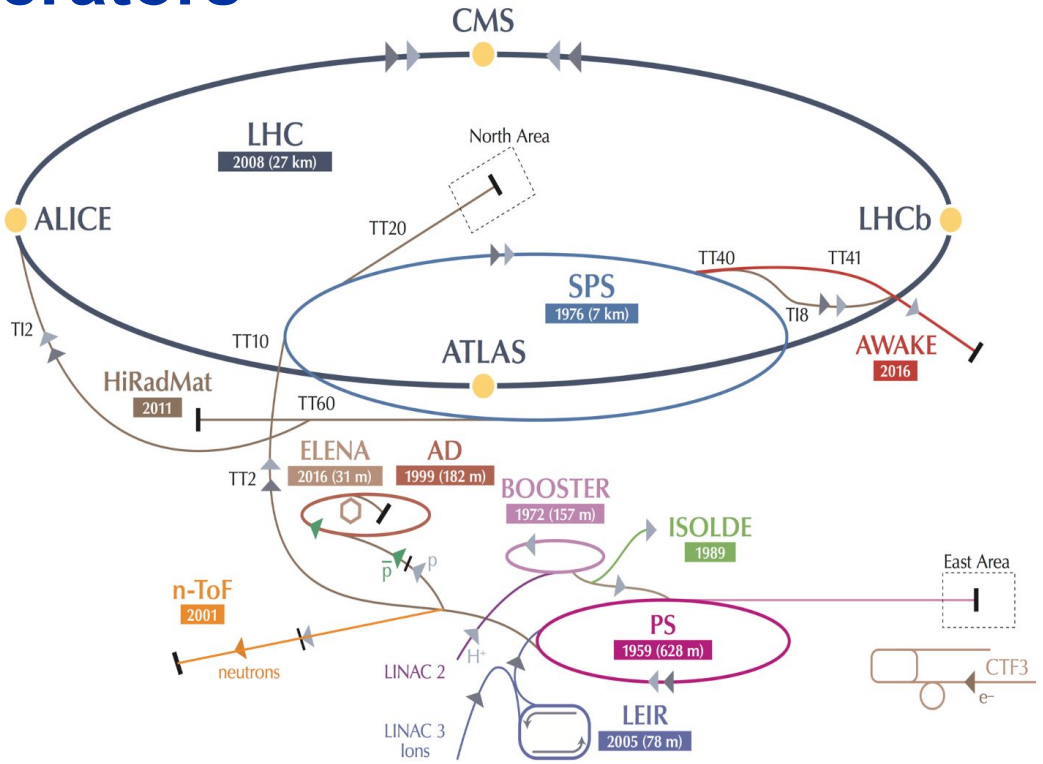

~20 projects other than LHC with > 1200 physicists


DATA
CENTRE

The LHC produces more than 1 billion particle collisions per second

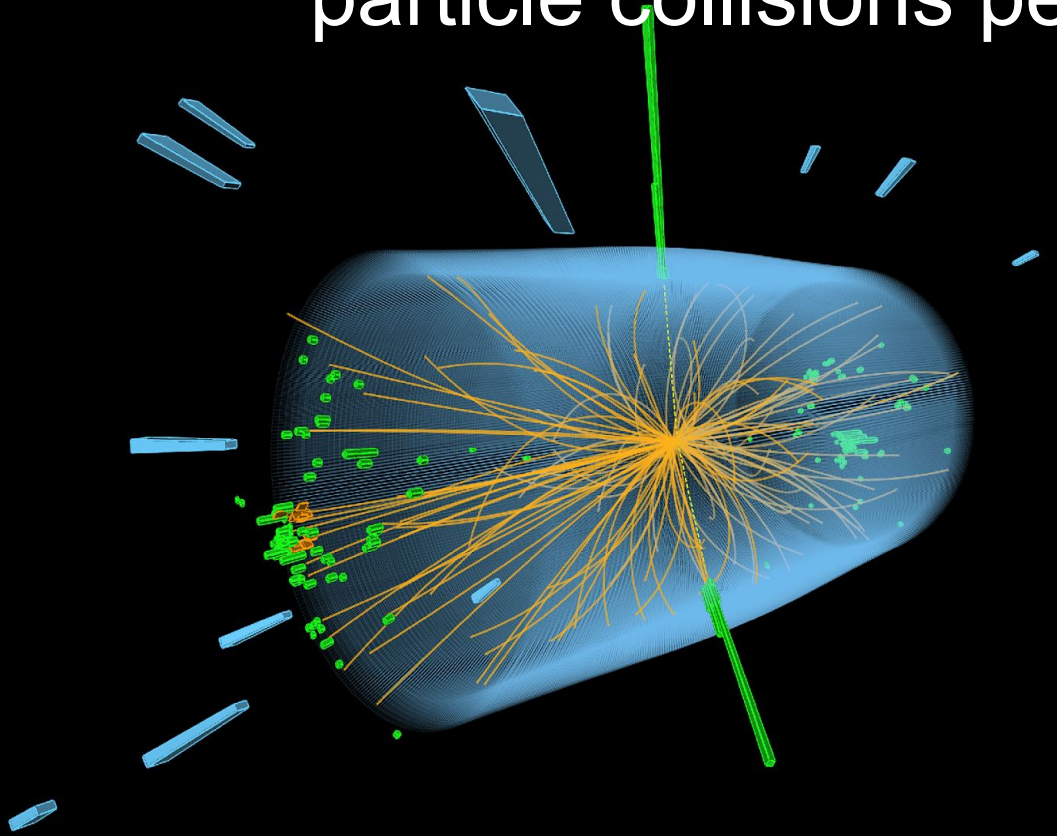

The energy of the particles in collision is converted into new particles.

CERN DC and WLCG (Worldwide LHC Computing Grid)

CERN Science is Data Intensive Science

- Run 1 (2009-2013) we stored 65 PB
- Run 2 (2015-2018) we stored 209 PB
- Run 3 (2022-2026) we expect to store 600 PB

- **CERN Data Centre** in Meyrin
  - **13k servers**, **450k CPU** cores, **320 PB** of storage
  - New data centre being built in Prevessin site

Solution is to combine the power of all our computing centres

Used to store, distribute, process and analyse data.

1 million processing cores in about 170 data centres and 42 countries.

More than 1 Exabyte of CERN data stored world-wide.

Service Availability

15:03:09

Available | Degraded | Unavailable | Unknown

ACRON	AFS	AFS Web Hosting	Alerter	BOINC	Backup and Restore	Batch	CDS Videos Application Support	CERN Document Server (CDS) Application Support	CERN Search	CERNBox	CIXP
CTA	CVMFS	Campus Network	Centralised Elasticsearch	Ceph	Certificate Authority	Conference Rooms	Configuration Management	DFS	Data integration and reporting / Pentaho	Database on Demand	Datacenter Network
Drupal	E-Mail	EOS for Physics	Eduroam	Electronics Design Software	FILER	File Transfer	Fixed Line Phone	GRID Information	Git	HADOOP	HPC
IIS Web Hosting	Indico Event Application Support	JIRA	LXPLUS	Linux Operating System	Load Balancing	Mathematics Software	Mechanical Design Software	Messaging	Monitoring	MyProxy	Network Database and Registration
Network for Projects and Experiments	Oracle Database	PaaS Web Application Hosting	Printing and Copying	Public Information Display	Server Provisioning	Sharepoint	Single Sign On and Account Management	Technical Network	Twiki	VOMS	WLCG Network
Weblogic, Tomcat Java application servers and 3rd party packages	Windows	Windows Terminal Servers									

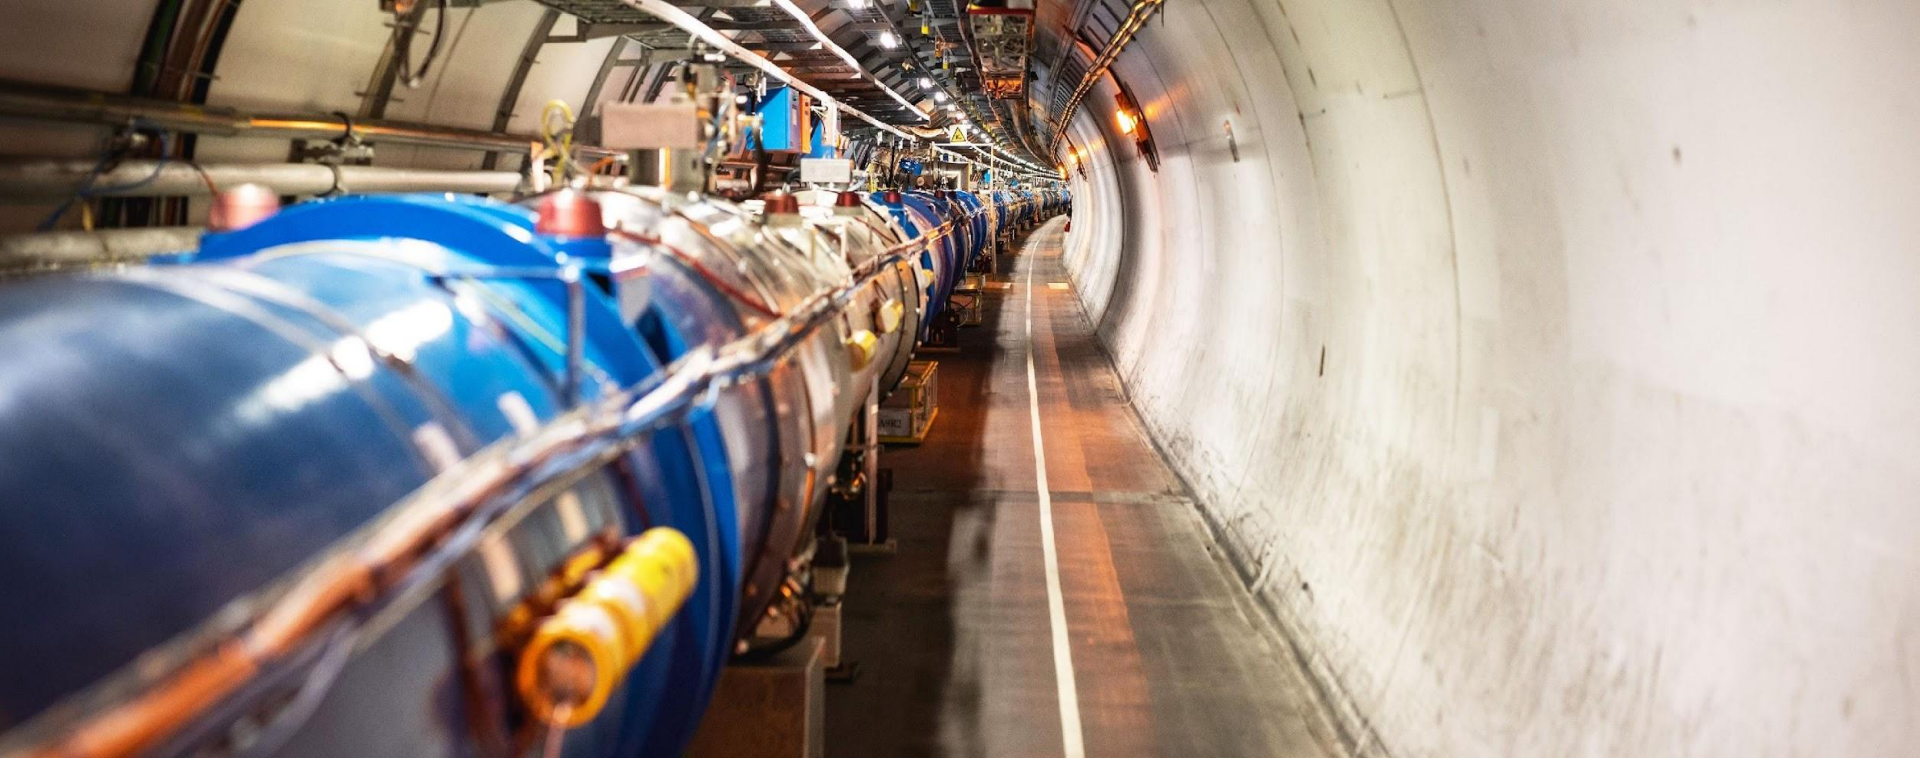

Big Data ecosystem @ CERN

Holistic view on the Big Data components

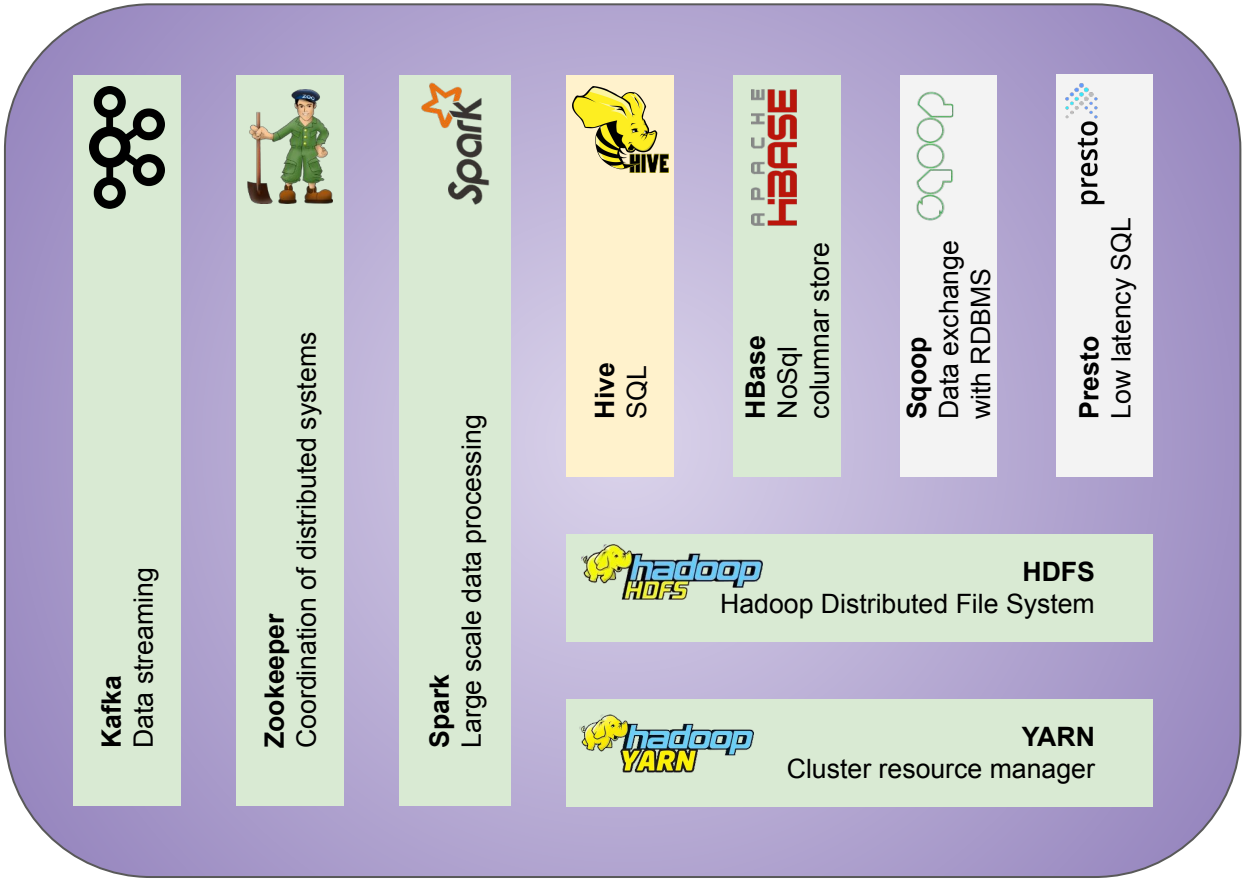

OpenSearch

Apache Ranger

Hadoop ecosystem @ CERN

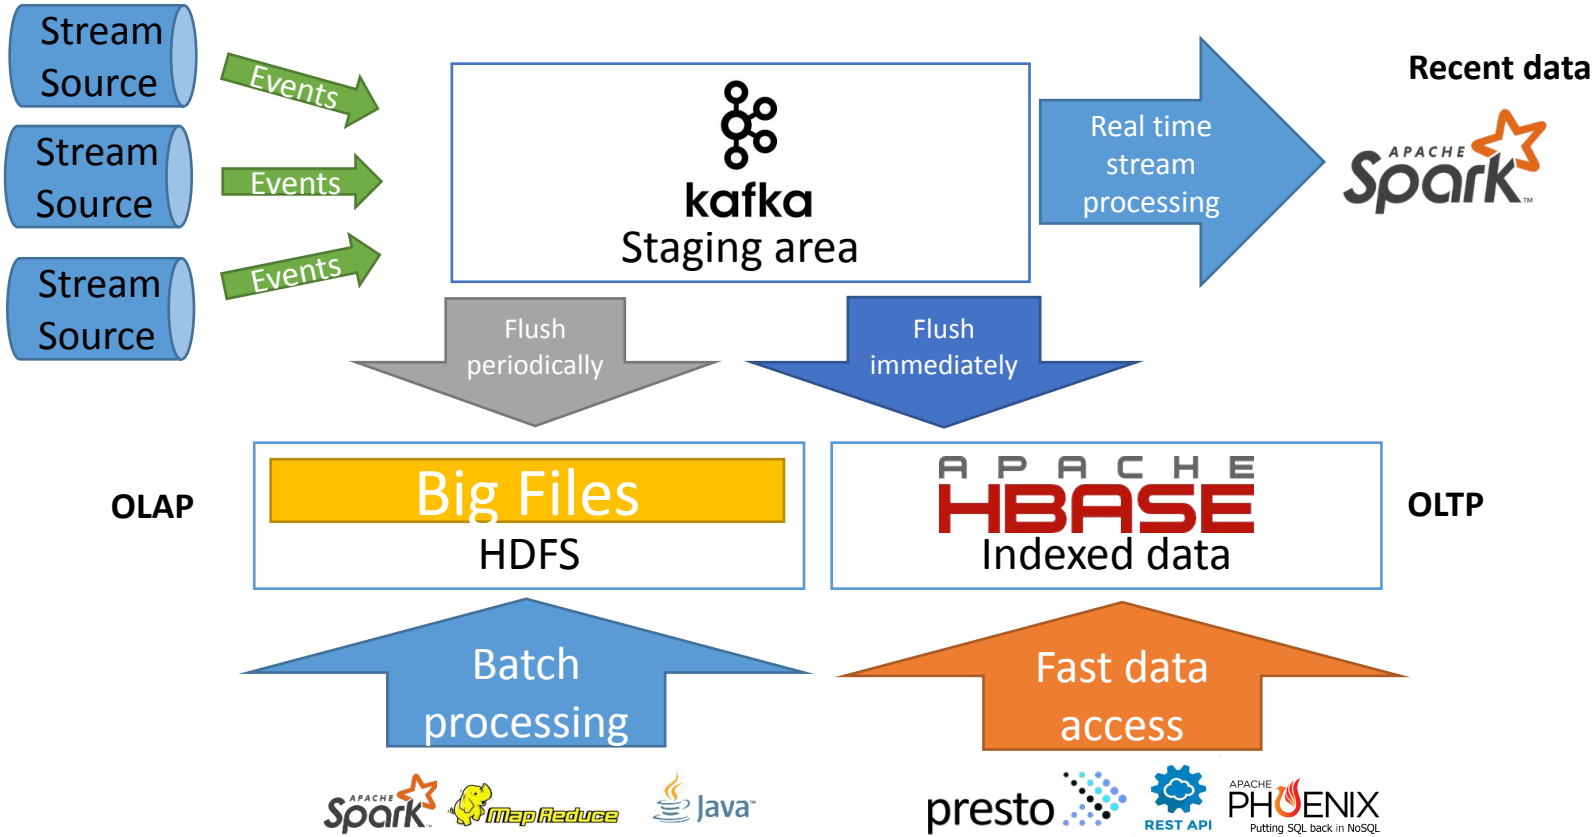

Key users and workload

- Users:
  - IT Security
  - IT Monitoring
  - Experiments
  - Accelerators Monitoring
- HBase and HDFS clusters

analytix_prod

Active Datanodes	HDFS Capacity	HDFS Usage	Number of Files	Active Region Servers	Database Size	Number of Regions
58	24.5 PB	32.0%	39.1 Mil	58	124 TB	13.2 K

nxcals_prod

Active Datanodes	HDFS Capacity	HDFS Usage	Number of Files	Active Region Servers	Database Size	Number of Regions
51	18.5 PB	50.6%	34.5 Mil	No data	No data	No data

nxcalsonline_prod

Active Datanodes	HDFS Capacity	HDFS Usage	Number of Files	Active Region Servers	Database Size	Number of Regions
30	343 TB	33.8%	529 K	30	23.4 TB	23.7 K

Key statistics

- **5 prod clusters** (+ 3QA & multiple DEV)
- Hardware: Intel-based servers, continuous capacity expansion, hyper-threading enabled, HDDs and SSDs, **180+ bare metal, 60+ VMs**
- **Users**: **600+** (general-purpose), **1000+** (accelerators) and growing
- Storage capacity: **40PB+**
- Service started in 2013

Current security state

- Integration with CERN's LDAP (e-groups)
- Kerberos and SSL
- FW rules defined
- Potential improvements:
  - Policies and ACLs defined manually
  - Lack of central management
  - No user-friendly auditing (raw logs)
  - No CERN SSO integration for Web UIs

```
# changing owner
hdfs dfs -chown <owner> <hdfs_path>

# granting new permissions to a user
hdfs dfs -setfacl -m user:<grantee>:<permission> <hdfs_path>

# granting new permissions to a group or an e-group
hdfs dfs -setfacl -m group:<group_name>:<permission> <hdfs_path>
```

Apache Ranger to improve security and ease management?

Apache Ranger

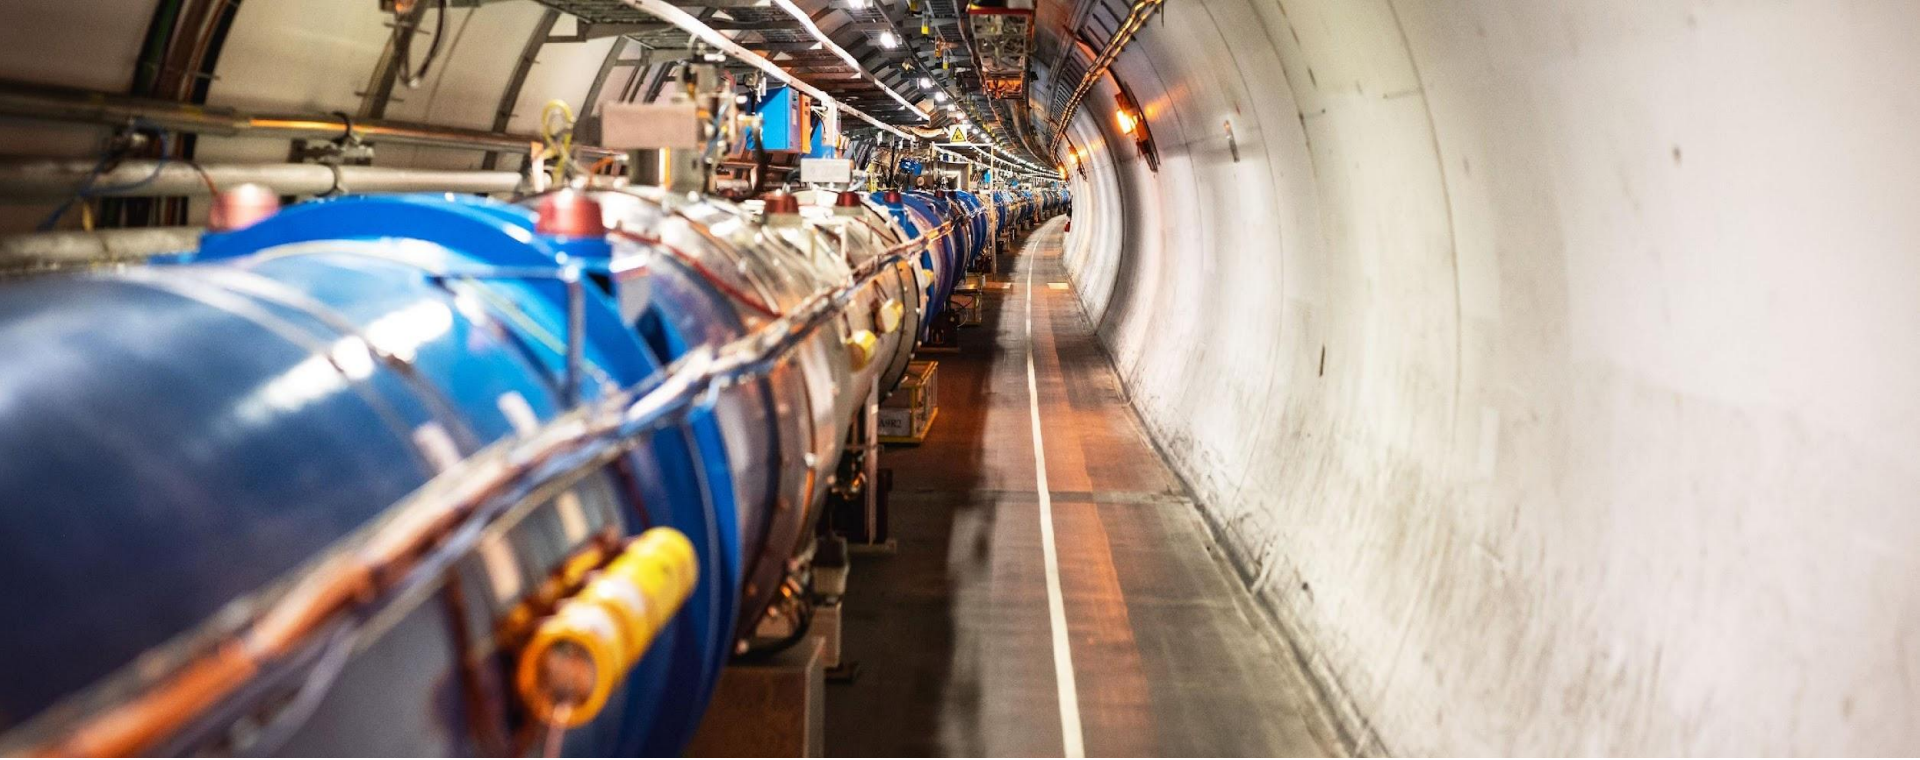

Overview of Apache Ranger project

About the tool

Apache Ranger

- Apache Ranger™ is a framework to enable, monitor and manage comprehensive data security across the Hadoop platform and beyond
- It supports multiple projects such as: Apache Hadoop, Apache HBase, Apache Kafka, YARN, Apache Hive and some more…
- The plugin agent pulls the policy-changes using REST API
- Replaced Apache Sentry

THE APACHE® SOFTWARE FOUNDATION

https://github.com/apache/ranger

https://ranger.apache.org/

What Apache Ranger offers

- Fine-Grained Access Control
- Centralized Policy Management
- Dynamic Policy Enforcement
- Centralized auditing
- Delegate administration of policies to group owners
- Integration with LDAP
- Resource and tag-based Policies
- Extensible Lightweight Plugin Architecture
- No single point of failure (if configured)

WebUI and REST APIs

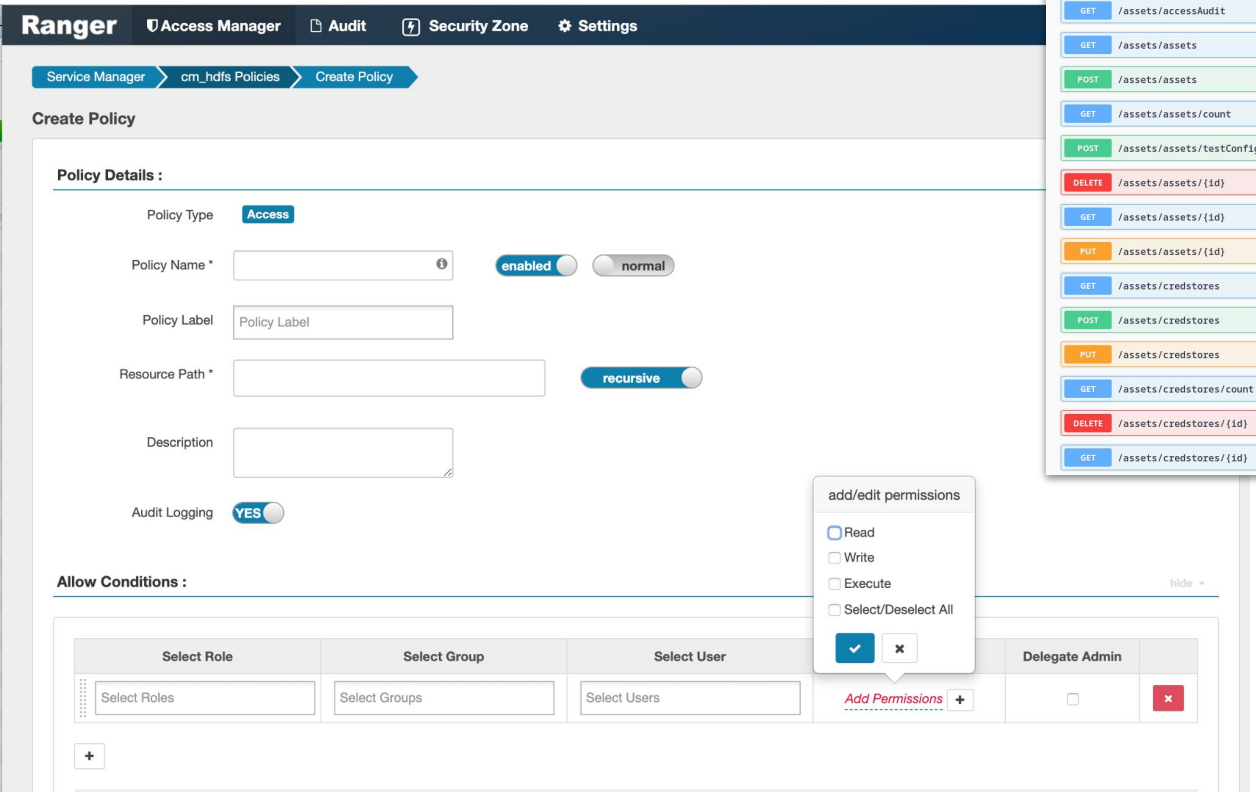

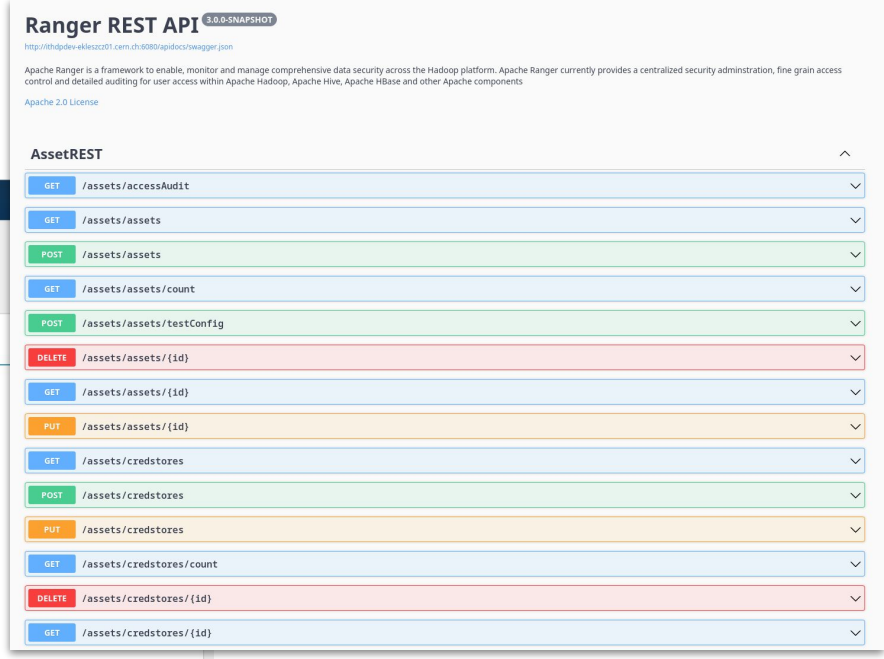

Web interface experience

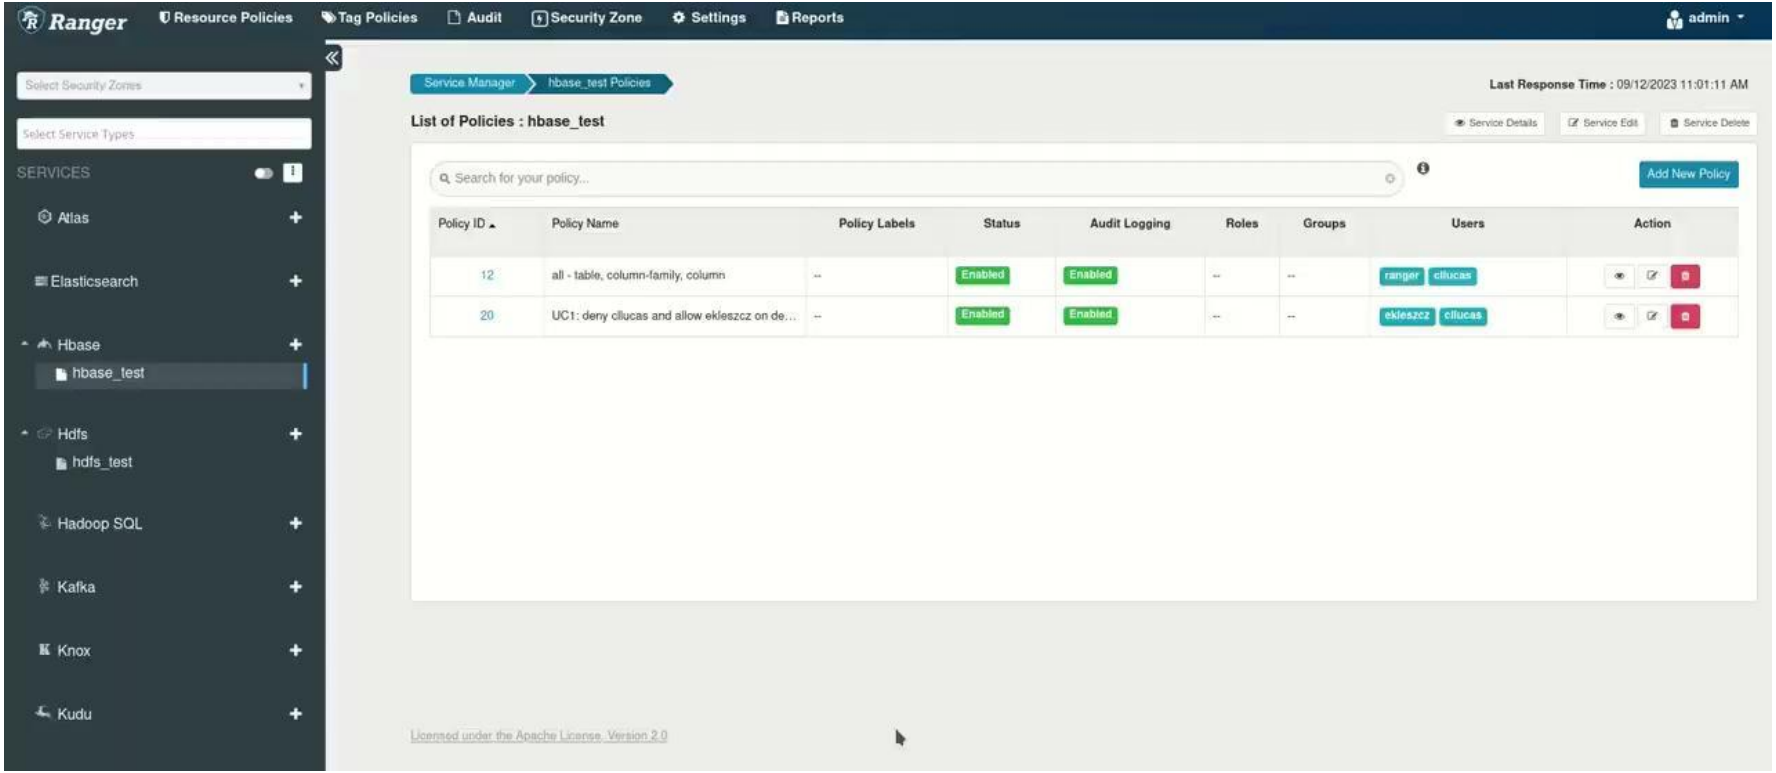

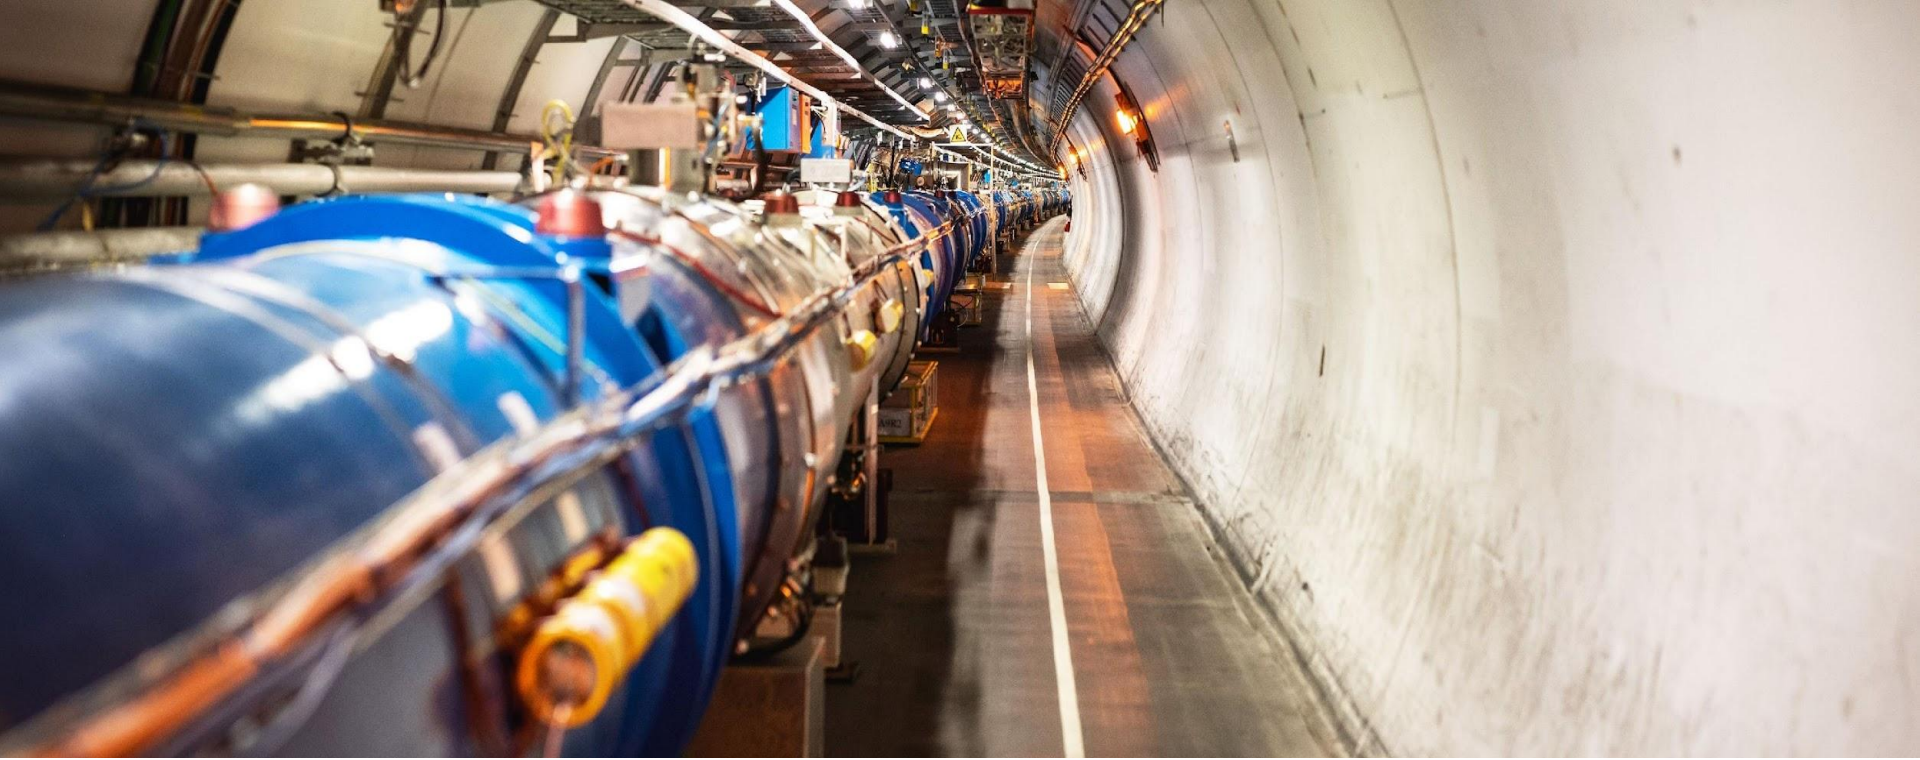

Apache Ranger @ CERN

Building the tool regular way

- Included only plugins for our components in use
- Tested against the master branch and latest release 2.4.0
- Build in CentOS 7 with Java 11

```
mvn clean compile package install -DskipTests=true -Dspotbugs.skip=true
-Dchkstyle.skip=true -Dassembly.plugin.version=3.1.0 -Dhadoop.version=3.3.0
-Dhbase.version=2.3.4 -Dzookeeper.version=3.6.1 -Dhive.version=3.0.0
-Dmysql-connector-java.version=8.0.28
-pl '!plugin-ozone, !plugin-solr, !plugin-nifi, !plugin-nifi-registry, !plugin-kudu,
!plugin-kms, !ranger-ozone-plugin-shim, !storm-agent, !ranger-storm-plugin-shim,
!ranger-solr-plugin-shim, !ranger-atlas-plugin-shim, !plugin-atlas, !plugin-kylin,
!ranger-kylin-plugin-shim'
```

Our setup

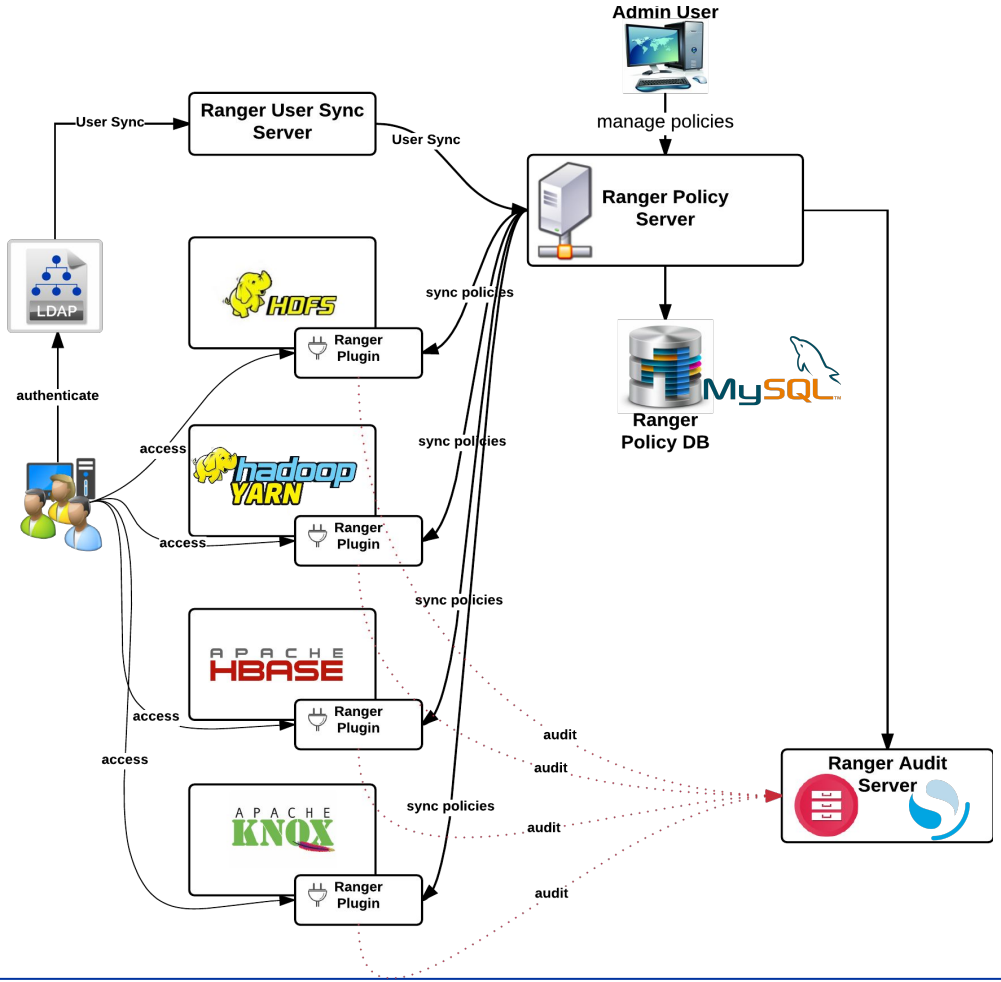
Admin User
manage policies
Ranger User Sync Server
User Sync
User Sync
Ranger Policy Server
LDAP
authenticate
HDFS
Ranger Plugin
sync policies
Ranger Policy DB
MySQL
access
hadoop YARN
Ranger Plugin
sync policies
access
sync policies
APACHE HBASE
Ranger Plugin
access
access
audit
audit
Ranger Audit Server
APACHE KNOX
sync policies
audit
Ranger Plugin
audit

Usersync plugin

- No performance overhead
- Integrates well with thousands of users and groups
- Responsive/robust
- Supports internal and external sync source
- Great for building Roles
- Integrated with the policies’ definitions

Usersync plugin

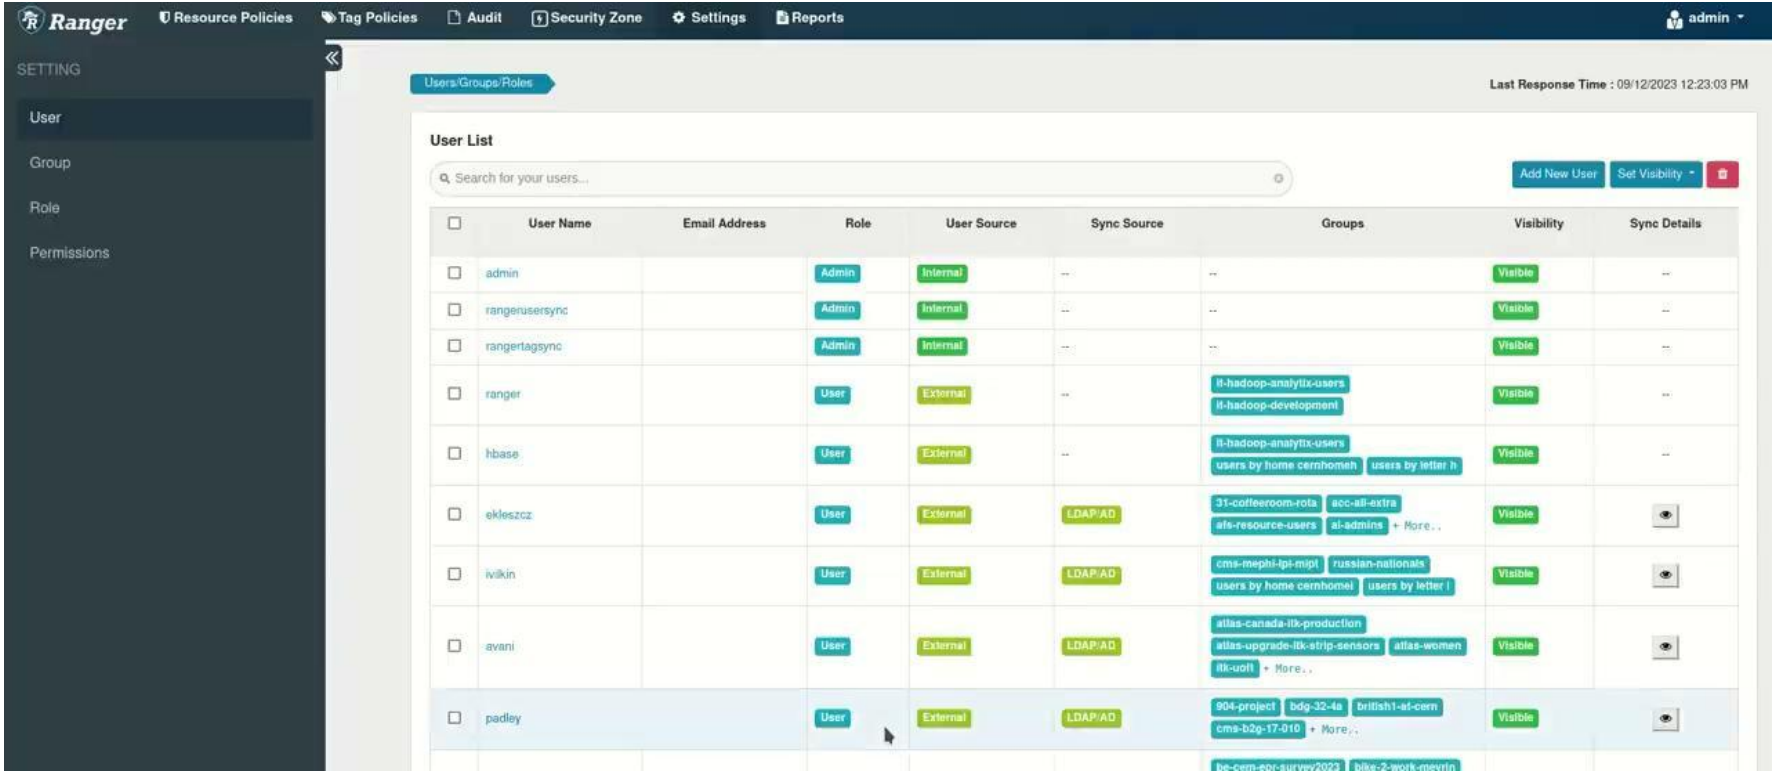

Apache Ranger

Plugin deployment

Demo

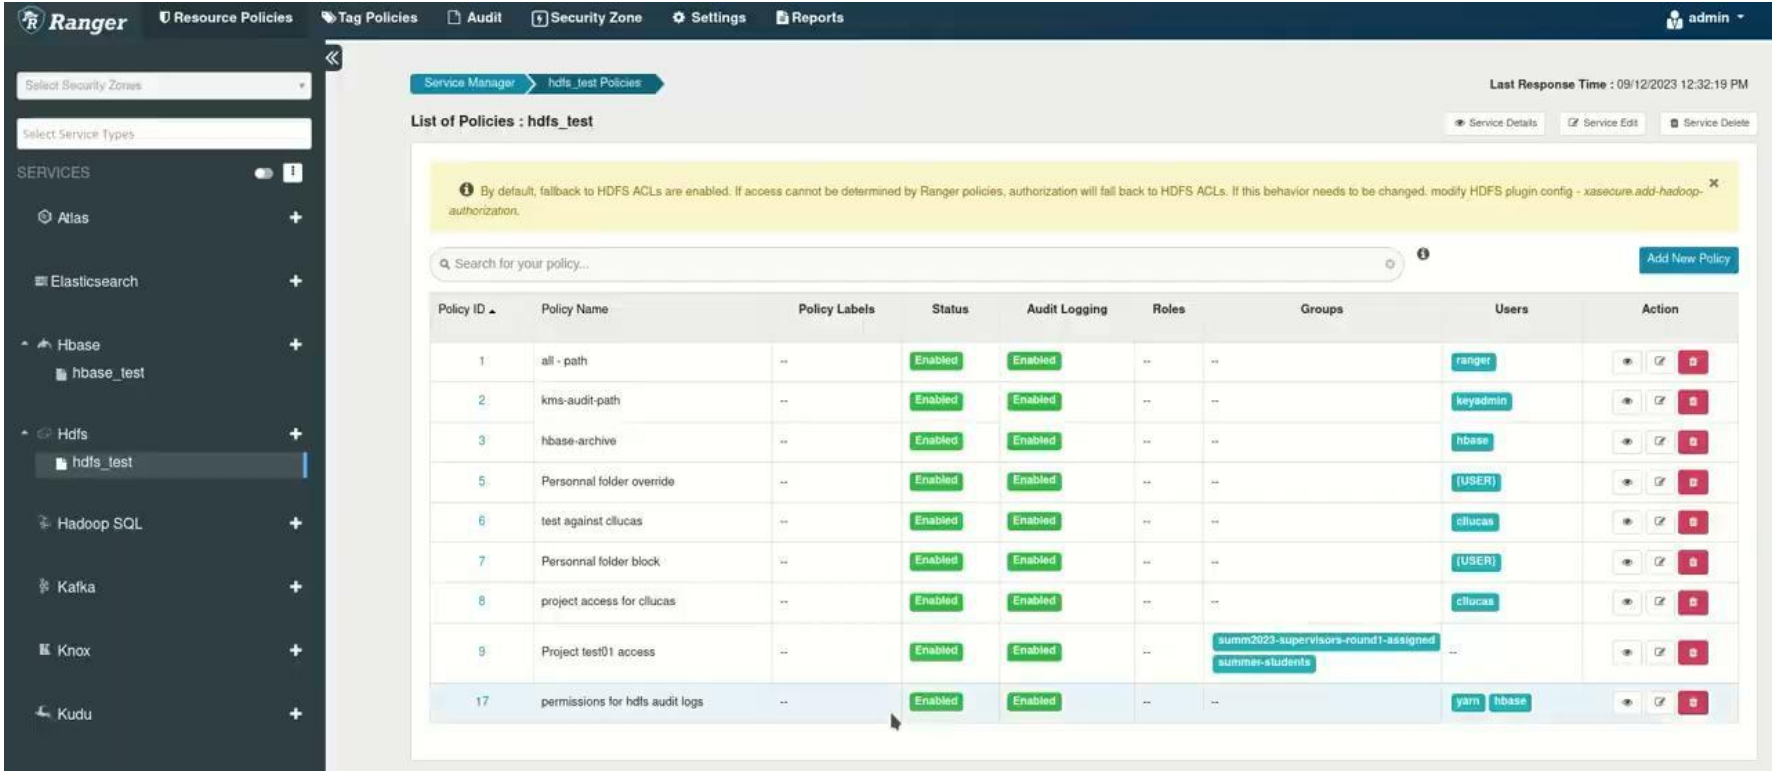
Ranger
Resource Policies
Tag Policies
Audit
Security Zone
Settings
Reports
admin
Select Security Zones
Select Service Types
SERVICES
Atlas
Elasticsearch
Hbase
hbase_test
Hdfs
hdfs_test
Hadoop SQL
Kafka
Knox
Kudu
Service Manager
hdfs_test Policies
Last Response Time : 09/12/2023 12:32:19 PM
List of Policies : hdfs_test
Service Details
Service Edit
Service Delete
By default, fallback to HDFS ACLs are enabled. If access cannot be determined by Ranger policies, authorization will fall back to HDFS ACLs. If this behavior needs to be changed, modify HDFS plugin config - xasecure.add-hadoop-authorization.
Search for your policy...
Add New Policy
Policy ID
Policy Name
Policy Labels
Status
Audit Logging
Roles
Groups
Users
Action
1 all - path -- Enabled Enabled -- -- ranger
2 kms-audit-path -- Enabled Enabled -- -- keyadmin
3 hbase-archive -- Enabled Enabled -- -- hbase
5 Personnal folder override -- Enabled Enabled -- -- {USER}
6 test against cllucas -- Enabled Enabled -- -- cllucas
7 Personnal folder block -- Enabled Enabled -- -- {USER}
8 project access for cllucas -- Enabled Enabled -- -- cllucas
9 Project test01 access -- Enabled Enabled -- summ2023-supervisors-round1-assigned summer-students --
17 permissions for hdfs audit logs -- Enabled Enabled -- -- yarn hbase

Setting resource policies

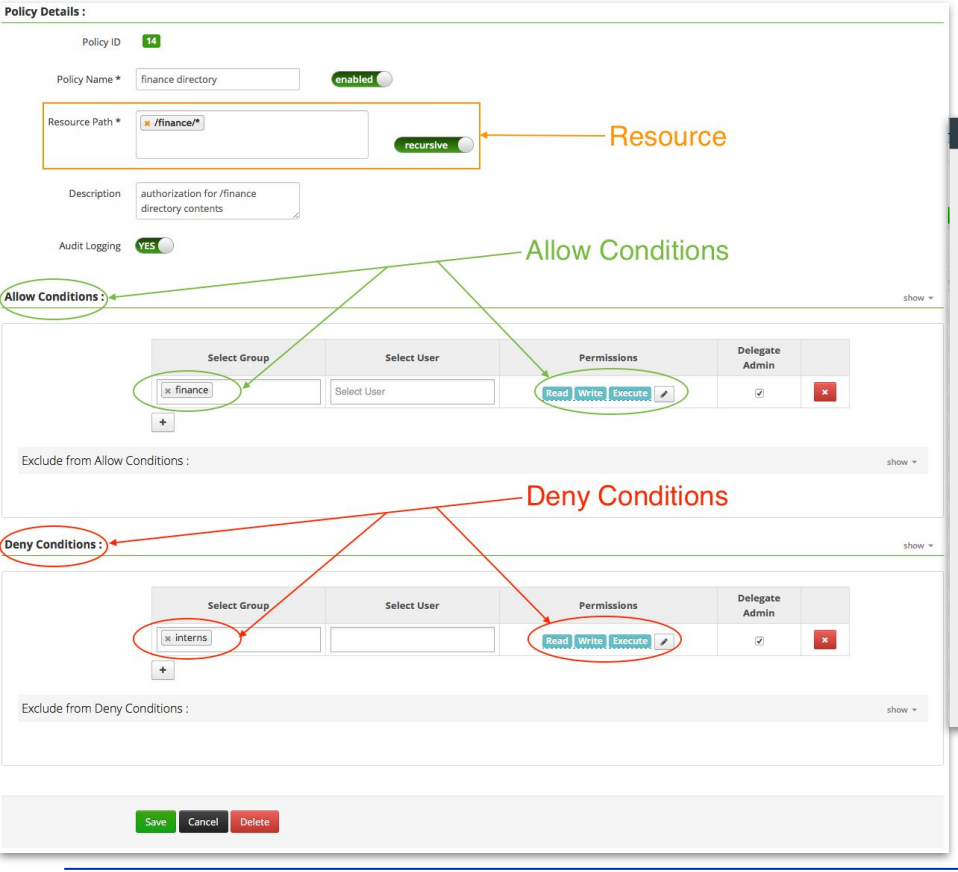

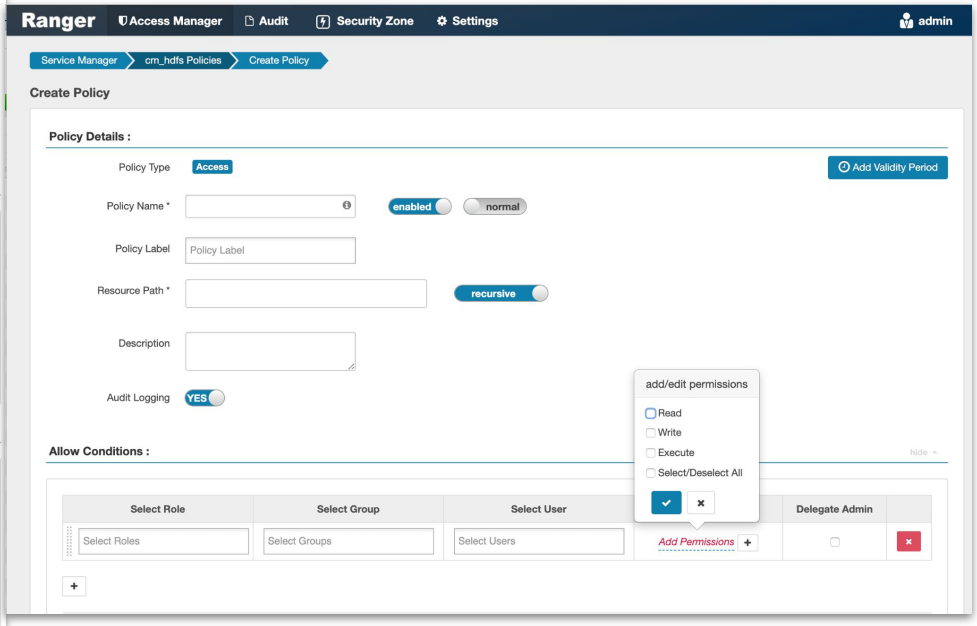

Resource-based policy access flow

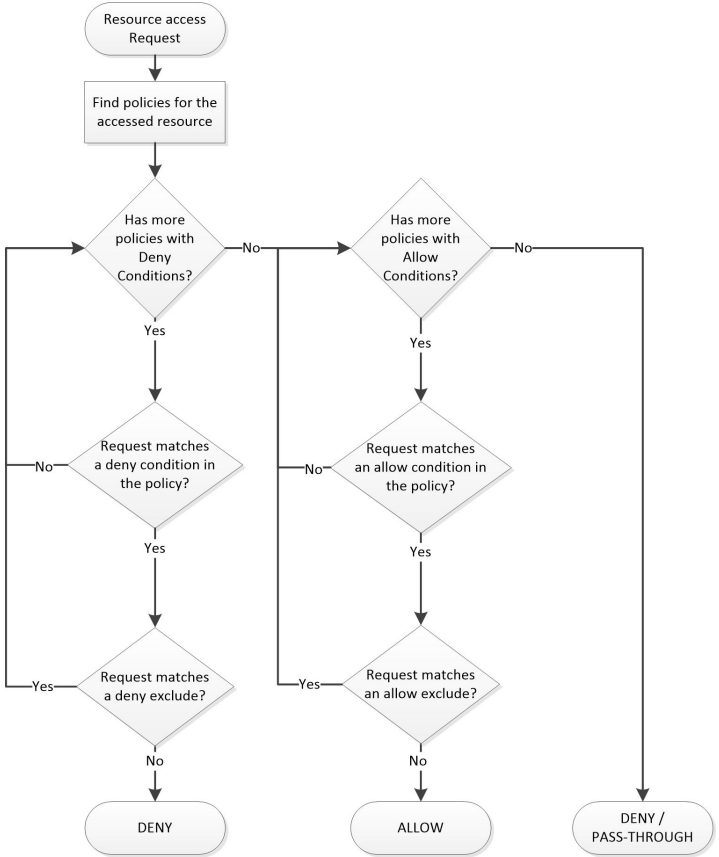

More details

- Support for autocompletion
- User can be set with one of the following roles: **KEYADMIN, ADMIN, USER**
- ROLE is assigned for a specific user by the Administrator
- **Delegated-Admin** permission allows other resource administrators to manage permissions for their managed-resources
- **USER** role gives ability to manage only resources for which the user has been granted with delegated-admin privilege
- Ranger provides **tag-based policies** too

Example with 2 policies

Policies :

1. **Block access for a {USER} to *user/**
2. **Override to allow {USER} access to /user/{USER}**

disabling an HDFS policy doesn't reset the previous state

> kinit cllucas

```
[root@ithdpdev-mmartinm01 ranger-2.4.0-hdfs-plugin]# hdfs dfs -ls /user/cllucas
Found 1 items
drwxr-xr-x   - cllucas supergroup          0 2023-09-11 17:08 /user/cllucas/.sparkStaging
```

```
[root@ithdpdev-mmartinm01 ranger-2.4.0-hdfs-plugin]# hdfs dfs -ls /user/ekleszcz
ls: Permission denied: user=cllucas, access=EXECUTE, inode="/user/ekleszcz"
```

Our experience

- Works with HA clusters (setting namespace)
- No performance overhead
- Extra properties to be set for clusters with Kerb auth in web UI
- Lack of documentation is very painful
- *dfs.namenode.rpc-address.<cluster_namespace>/<hostname>* must be updated to be done when aliases change (hosts upgraded)

```
[root@          ekleszcz]# doAs hdfs hdfs dfs -ls /
Found 13 items
drwxr-xr-x   - hbase    hdfs          0 2017-06-18 19:57 /#
drwxr-xr-x   - hdfs     hdfs          0 2018-04-17 15:59 /Training
drwxr-xr-x   - hdfs     hdfs          0 2019-10-17 11:16 /Trash
drwxr-xr-x   - hdfs     zp            0 2023-04-05 09:33 /atlas
drwxrwxr-x+  - hdfs     zh            0 2023-01-13 18:00 /cms
drwxrwx--x+  - hbase    hdfs          0 2023-05-17 11:51 /hbase
drwxr-xr-x   - ekleszcz hdfs          0 2023-08-18 15:09 /hbase-snapshots-perf-test
drwxr-xr-x   - hdfs     hdfs          0 2016-11-04 14:42 /lost+found
drwxr-xr-x   - hdfs     hdfs          0 2022-07-26 11:42 /project
drwxr-xr-x   - hdfs     hdfs          0 2023-10-05 19:00 /system
drwxrwxrwt   - yarn     hadoop        0 2023-10-05 12:02 /tmp
drwxr-xr-x+  - hdfs     hdfs          0 2023-10-02 18:00 /user
drwxrwxr-x+  - hbase    hadoop        0 2023-10-02 17:11 /var
```

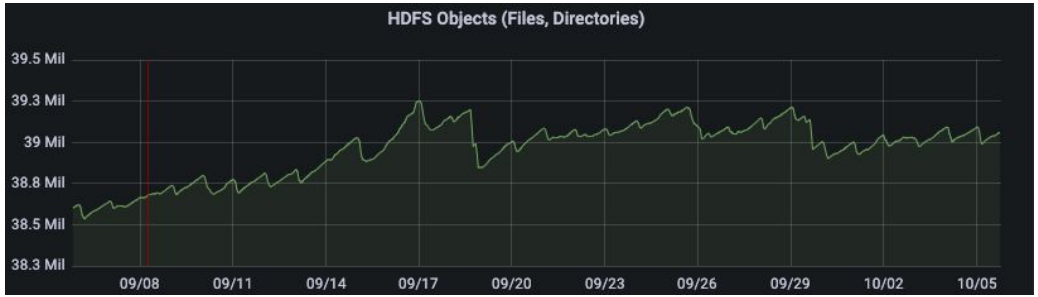

Apache Ranger

Plugin deployment

Yarn resources at CERN

Cluster Metrics

Apps Submitted	Apps Pending	Apps Running	Apps Completed	Containers Running	Memory Used	Memory Total	Me
194399	31	35	194333	400	5.38 TB	30.07 TB	7 GB

Cluster Nodes Metrics

Active Nodes	Decommissioning Nodes	Decommissioned Nodes	Lost Nodes	Unhealthy N
58	0	0	0	0

Scheduler Metrics

Scheduler Type	Scheduling Resource Type	Minimum Allocation	Maximum Allocation
Fair Scheduler	[memory-mb (unit=Mi), vcores]	<memory:2048, vCores:1>	<memory:81920, vCores:32>

Application Queues

Legend: Steady Fair Share | Instantaneous Fair Share | Used | Used (over fair share) | Max Capacity

- root
  - root.certlogcompress
  - root.hconfig
  - root.monitops
  - root.certaenr
  - root.cperezde
  - root.itmonops
  - root.tapeops
  - root.atlevind
  - root.backup
  - root.hmonitor

Yarn Fair scheduler

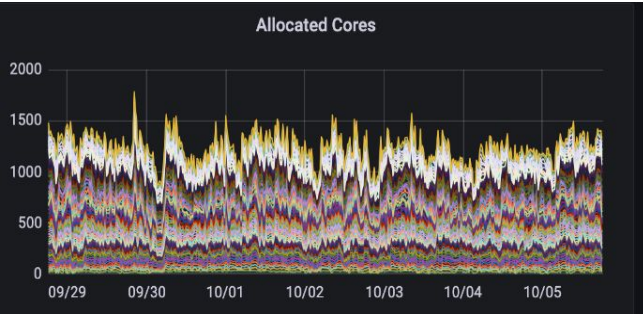

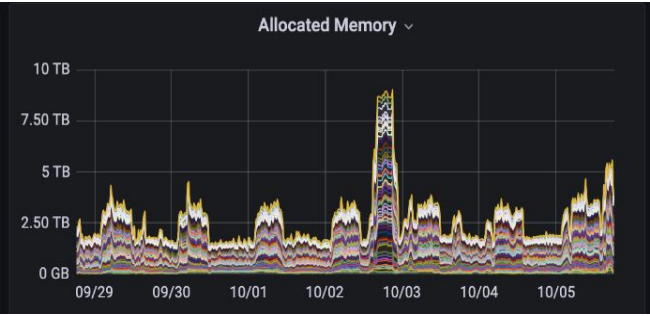

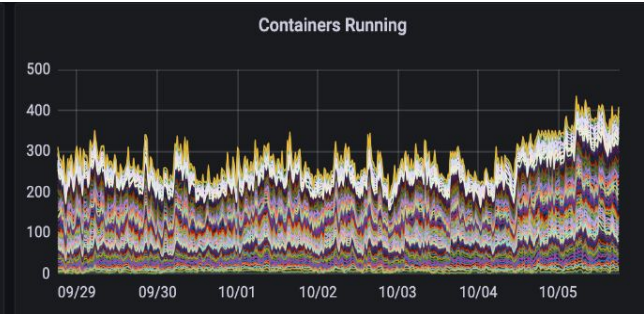

Setting policies for a fair scheduler

queue name is: **root.{group}.{user}**

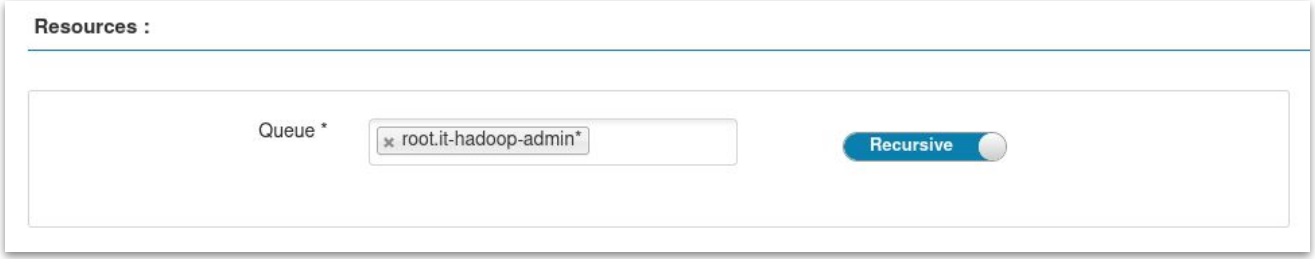

- **submit-app:** submit apps to the queue
- **admin-queue:** administer a queue **(be able to kill an app)**

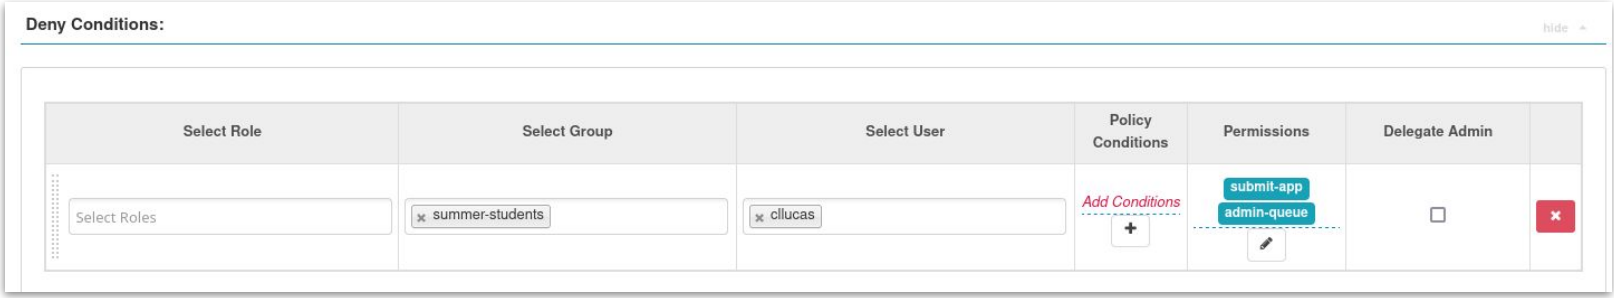

Apache Ranger

APACHE HBASE

Plugin deployment

HBase ACLs management

Around 10k tables in multiple namespaces

```
hbase(main):006:0> user_permission '@ekleszcz'
User                                                    Namespace,Table,Family,Qualifier:Permission
 [redacted]                                             ekleszcz,,,: [Permission: actions=READ,WRITE,EXEC,CREATE,ADMIN]

Took 0.0462 seconds
```

```
hbase(main):002:0> user_permission
User                                                    Namespace,Table,Family,Qualifier:Permission
 [redacted]                                             ,,,: [Permission: actions=READ,WRITE,EXEC,CREATE,ADMIN]
                                                        ,,,: [Permission: actions=CREATE]
                                                        ,,,: [Permission: actions=READ,WRITE,EXEC,CREATE]
                                                        ,,,: [Permission: actions=READ,WRITE,EXEC,CREATE,ADMIN]
                                                        ,,,: [Permission: actions=READ,WRITE,EXEC,CREATE,ADMIN]
[redacted]
Took 0.4352 seconds
```

Setting policies

disabling an HBase policy doesn't reset the previous state

{namespace}:{table}

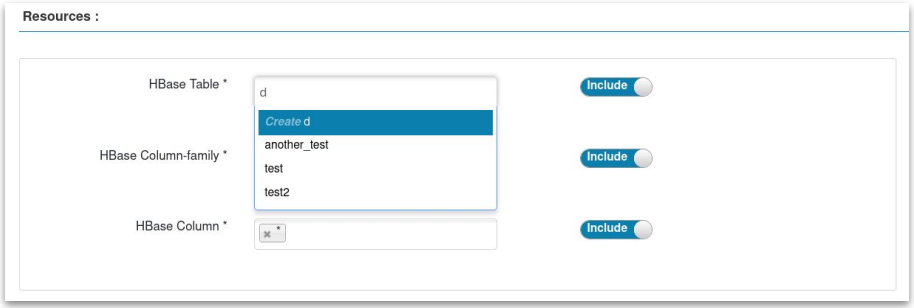

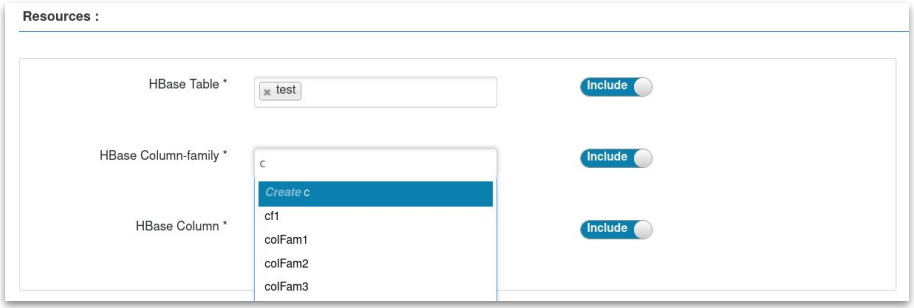

Permissions :

- **Read**
- **Write**
- **Create**
- **Admin**
- **Execute**

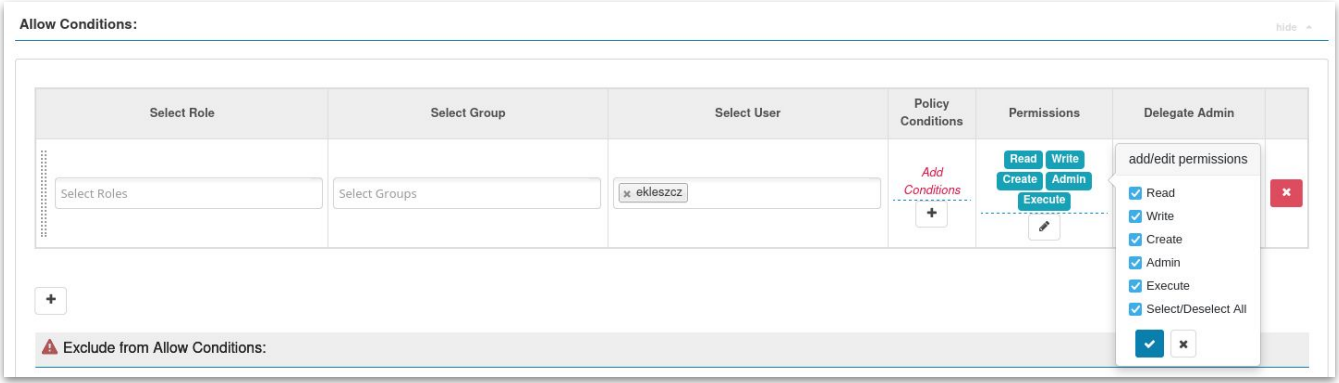

Auditing

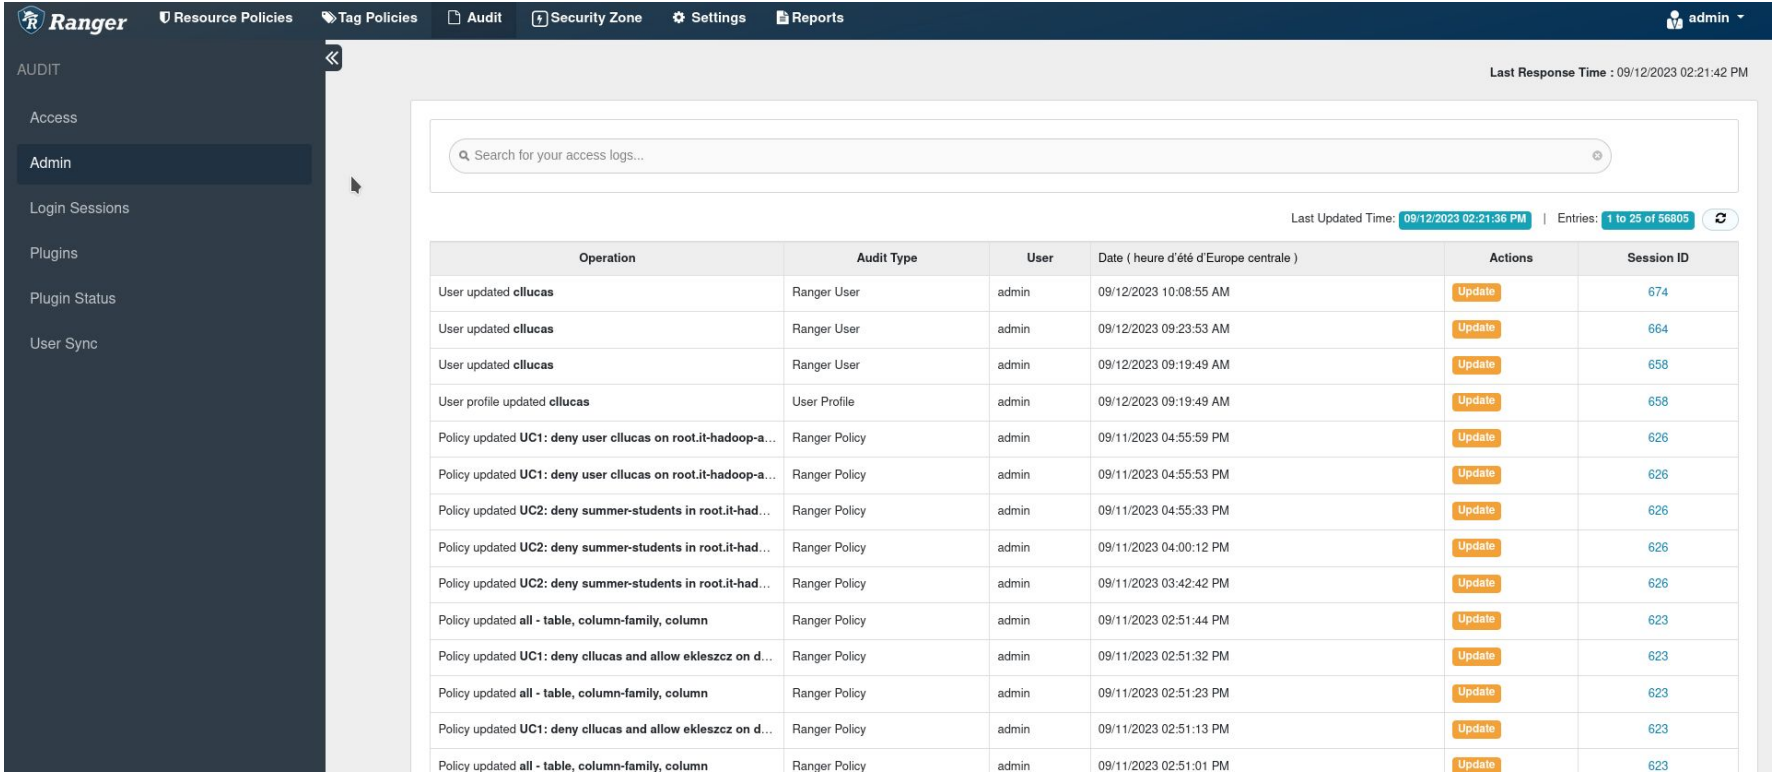
Ranger
Resource Policies
Tag Policies
Audit
Security Zone
Settings
Reports
admin
AUDIT
Access
Admin
Login Sessions
Plugins
Plugin Status
User Sync
Last Response Time : 09/12/2023 02:21:42 PM
Search for your access logs...
Last Updated Time: 09/12/2023 02:21:36 PM | Entries: 1 to 25 of 56805
Operation | Audit Type | User | Date (heure d'été d'Europe centrale) | Actions | Session ID
User updated cllucas | Ranger User | admin | 09/12/2023 10:08:55 AM | Update | 674
User updated cllucas | Ranger User | admin | 09/12/2023 09:23:53 AM | Update | 664
User updated cllucas | Ranger User | admin | 09/12/2023 09:19:49 AM | Update | 658
User profile updated cllucas | User Profile | admin | 09/12/2023 09:19:49 AM | Update | 658
Policy updated UC1: deny user cllucas on root.it-hadoop-a... | Ranger Policy | admin | 09/11/2023 04:55:59 PM | Update | 626
Policy updated UC1: deny user cllucas on root.it-hadoop-a... | Ranger Policy | admin | 09/11/2023 04:55:53 PM | Update | 626
Policy updated UC2: deny summer-students in root.it-had... | Ranger Policy | admin | 09/11/2023 04:55:33 PM | Update | 626
Policy updated UC2: deny summer-students in root.it-had... | Ranger Policy | admin | 09/11/2023 04:00:12 PM | Update | 626
Policy updated UC2: deny summer-students in root.it-had... | Ranger Policy | admin | 09/11/2023 03:42:42 PM | Update | 626
Policy updated all - table, column-family, column | Ranger Policy | admin | 09/11/2023 02:51:44 PM | Update | 623
Policy updated UC1: deny cllucas and allow ekleszcz on d... | Ranger Policy | admin | 09/11/2023 02:51:32 PM | Update | 623
Policy updated all - table, column-family, column | Ranger Policy | admin | 09/11/2023 02:51:23 PM | Update | 623
Policy updated UC1: deny cllucas and allow ekleszcz on d... | Ranger Policy | admin | 09/11/2023 02:51:13 PM | Update | 623
Policy updated all - table, column-family, column | Ranger Policy | admin | 09/11/2023 02:51:01 PM | Update | 623

Auditing: logs kept in HDFS

Stored in /ranger/audit/{plugin}/{date:YYYYMMDD}/{filename}

Example : `hdfs dfs -cat`

```
/ranger/audit/hdfs/20230912/hdfs_ranger_audit_ithdpdev-mmartinm02.cern.ch.log
```


```
{
   "repoType":1,
   "repo":"hdfs_test",
   "reqUser":"ekleszcz",
   "evtTime":"2023-09-12 15:11:43.818",
   "access":"READ_EXECUTE",
   "resource":"/ranger/audit",
   "resType":"path",
   "action":"read",
   "result":1,
   "agent":"hdfs",
   "policy":-1,
   "reason":"/ranger/audit",
   "enforcer":"hadoop-acl",
   "cliIP":"188.184.11.111",
   "agentHost":"ithdpdev-ekleszcz12.cern.ch",
   "logType":"RangerAudit",
   "id":"bdd67338-b9c2-490f-81e2-66304d5c2673-0",
   "seq_num":1,
   "event_count":1,
   "event_dur_ms":0,
   "tags":[],
   "additional_info":"{\"forwarded-ip-addresses\":\"[]\",
\"remote-ip-address\":\"188.184.11.111\",\"accessTypes\":\"[execute, read]\"}",
   "cluster_name":""
}
```


Result of the request

1 is allowed, 0 is denied

Policy ID used to process the request

-1 means that no Ranger policy cover the request

Our experience and concerns

- Documentation
  - Legacy or missing or linked to Ambari or Cloudera but not standalone
  - Little activity for reported issues in GitHub/mailing list
- Compatibility
  - Certain component versions must be respected - no docs
  - Last official release 2.4.0 from 2017
- Integrations
  - Auditing not supported for OpenSearch, only ES 7.10.2
  - Some missing dependencies in yarn plugin (jars missing, eg. apache-commons-lang3)
- Resource intensive - 1G+ of heap needed with certain plugins enabled
- Pull mechanism only - no enforcement of the current state
- Yarn autocompletion didn’t work with fair-scheduler (?)

Upstream contributions

- Clear way to contribute
- **A contribution submitted** to upstream but not followed up
- A few issues reported to the mailing list
  - (dev@ranger.apache.org)

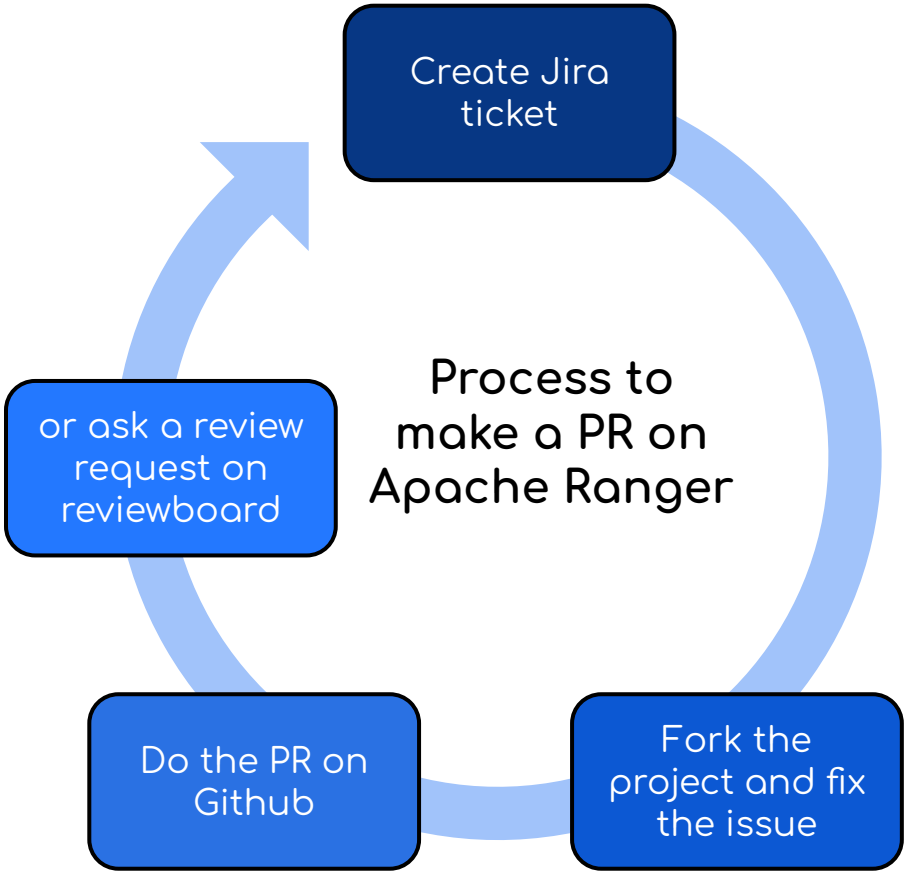

Future plans

- Puppetize and deploy to PROD
- SSO integration (Apache KNOX)
- Auditing with OpenSearch backend
  - We moved away from ElasticSearch/OpenDistro
- Outsource mgmt. to project managers via Roles
- Explore HA deployment
- Fix autocompletion for Apache YARN
- Dynamic management of the configs

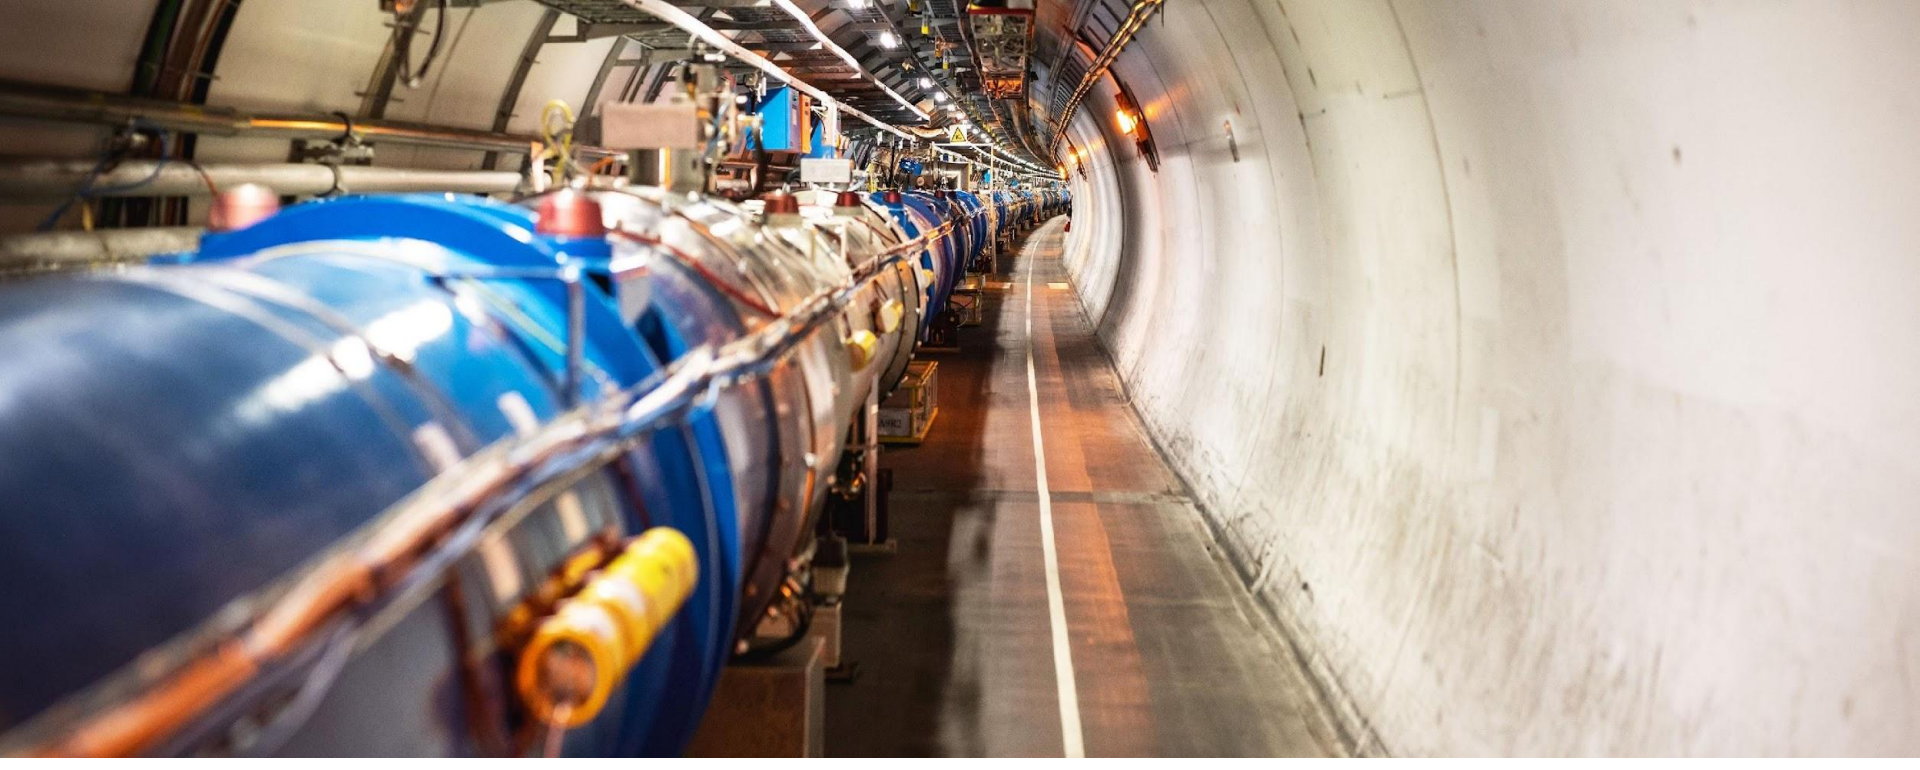

Summary

Summary

- At CERN, **we collect 100s of PBs of data** from various physics experiments
- We have a lot of users who access the data from our Big Data services such as HDFS
- We **need to protect this data and optimize the access rights management**
- Also, we lack a **centralized auditing system**
- **Apache Ranger is a proven-to-work mature project that integrates well** with our main components and provides the solution for our needs
- With Ranger, we faced **a few issues that still need to be tackled**, such as
  - **integration with OpenSearch** for auditing
  - **autocompletion** for the **YARN** plugin
- The tool will be deployed into the prod clusters

Thank you !

Contact: emil.kleszcz@cern.ch

home.cern