

Input by the CERN ATLAS Team for the Offline Computing and Software

RCS-IT Technical Committee (December 2022)

The CERN ATLAS Team (CAT) is a member institute of the ATLAS Experiment, and as such has host-lab responsibilities for crucial tasks in operation and maintenance, the upgrade, and the physics exploitation of ATLAS. This includes tasks within the ATLAS Software and Computing and Trigger/Data Acquisition domains, and, more concretely, responsibilities in Tier-0 operations, system software development and operations for distributed data management (DDM), and development of tracking, core and analysis offline software.

For fulfilling these tasks, the CAT heavily relies on services provided by CERN IT. CAT also makes heavy use of the products (such as Geant4, ROOT, CVMFS) and services provided by EP-SFT, and greatly profits from the fruitful collaboration. We are very happy and thankful to both for the excellent support. It is of utmost importance for CAT that this level of support by IT and EP-SFT continues also in the future.

The CAT is also engaged in successful open-source projects of scopes beyond ATLAS and even CERN and HEP: Rucio on scientific data management, and ACTS on common tracking solutions. For evolution of these projects, the CAT seeks and suggests even deeper collaboration with IT, and more engagement by IT in the future.

List of main and critical services

Services required by the CAT to fulfil host-lab responsibilities towards ATLAS:

- **Oracle database infrastructure and services**
 - Required for T/DAQ, Tier-0 and Rucio operations
- **CERN storage services**
 - EOS for disk, CTA for tape, associated network
 - Required for T/DAQ, Tier-0 and Rucio operations
 - Required for local CAT physics analysis activities
- **CERN batch services**
 - HTCondor
 - Required for Tier-0 operations
 - Required for local CAT physics analysis activities
- **CERN Cloud services**
 - OpenStack, GitOps, Puppet, containers (Kubernetes), load-balancing
 - Required by all groups and activities

- **GitLab**
 - Used as repository and bug/feature tracking tool
 - Required by all groups and activities
 - Required especially by Rucio for Kubernetes/GitOps

Services required particularly by Rucio and DDM operations:

- **FTS**
 - Required by Rucio for all ATLAS data transfers
- **Data management clients**
 - GFAL, Davix
 - Required for all ATLAS storage interactions
- **MONIT infrastructure**
 - Including Opensearch, Grafana
 - Required for DDM daily operations
 - Including dedicated ATLAS instances of Opensearch
- **ActiveMQ**
 - Required by Rucio and the monitoring infrastructure
- **Hadoop / Spark**
 - Required for DDM operational reporting and overview reports for management
- **Authentication & Authorisation**
 - X509 certificates, OAuth tokens; VOMS, IAM

Evolution of services

- **Support for new hardware architectures and technologies**
 - Testbeds for software development and continuous integration (CI) on GPUs

For our offline software development activities we need access to new development hardware. Testing and deployment on a larger scale exceeds our limited resources, and requires engagement and support by CERN IT, including the implementation of a decision-making process in order to mutually agree on the specifications, scopes and goals.
- **CERN Cloud services: support for containers**
 - Rucio is deployed on Kubernetes on the CERN IT Cloud infrastructure
 - Legacy deployment still exists on Puppet
 - More central technical support by IT for Kubernetes is needed
- **CERN Cloud services: load balancing**
 - Rucio uses implementation in Puppet right now
 - Service needs to be evolved to fully support Kubernetes
- **GitLab**
 - Improved and facilitated access for external (non-CERN) collaborators
 - Hosting of browsable artifact HTML pages

- **FTS**
 - Long-term support
 - Intensified development support
- **Authentication & Authorisation**
 - Long-term support for IAM
 - Support for migration from VOMS to IAM
- **Support for alternative database technologies**
 - Investigate PostgreSQL as a database replacement, with similar level of administrative and technical support as for Oracle

New engagements and extended collaborations

The CERN ATLAS Team would like to see that CERN IT would engage (more) in the **operational support** of the **Rucio data management systems** for ATLAS, CMS and other experiments and sciences using it (or starting to use it).

Over the last years the role of the Rucio project has changed from an ATLAS-internal service to a multi-experiment open-source project. The CERN ATLAS Team is leading the development of the service and is providing operational support for ATLAS and, on best-effort basis, for the other experiments and sciences. We believe that it would not only be beneficial for ATLAS, but for CERN as a whole if IT engaged (more) into the Rucio project and in particular in supporting the experiments in operating the service.