SIS computing needs

Micha Moskovic (RCS-SIS)

RCS-IT Technical Committee #1, 9 December 2022

What we do

In collaboration with international partners, we develop and operate web services related to scientific information for the CERN and global HEP community

- **INSPIRE: main information platform for HEP**
  - 500k-700k visits/month all around the world 24/7, ~100k unique visitors/month
  - Critical to find HEP literature and evaluate job applicants
  - Collaboration of 6 particle physics labs (CERN, DESY, FNAL, IN2P3, IHEP, SLAC)
- **SCOAP3: main Open Access initiative in HEP**
  - Consortium of 3000 institutions worldwide
  - Used for compliance checks in contracts with publishers (60 MCHF since 2014)
  - Partners use the repository as central hub for ingestions into their local repositories
- **HEPData: publication-related HEP data repository**
  - Collaboration with IPPP Durham
  - Data from LHC experiments & more
- **CERN Analysis Preservation: preserve, link & version everything about an analysis**
  - Data, code, documentation, reviews, workflows, gitlab repositories etc.
  - Enabling future use of our unique research output
- **Other smaller tools (CERN Academic Training portal, CERN author guide)**

We are also responsible for data in services provided by IT (library catalogue, CDS)

What we use

For the services we manage, we rely on the following IT resources:

- **Single Kubernetes cluster provisioned through OpenStack Magnum (60 nodes, 1k cores, 1.8 TB RAM)**
- **Multiple ElasticSearch/OpenSearch clusters**
- **Multiple DBOD Postgres databases (but running our own DB for INSPIRE)**
- **File storage:**
  - Ceph block storage (8 TB)
  - CephFS (6 TB)
  - Ceph/S3 (11 TB)
  - EOS (16 TB)
- **LBaaS**
- **CERN DNS**
- **Kubeflow service for ML training**

What we need

- **We don’t expect to have totally new use cases in the coming 5 years, but a slow growth of our datasets and potentially new services using similar technologies**
- **The Scientific Information Landscape project we are steering might impact other services managed by IT (e.g. CDS) and/or require new services if it goes to implementation stage.**
- **Main shortcomings with current IT services (in decreasing order of importance)**
  - **Availability**: high availability requirements as some of our services (e.g. INSPIRE) are critical for the work of a global community that access them 24/7. Affected by many outages in Data Centre. We would welcome SLAs to set clear expectations, also for users and partners.
  - **Business continuity/disaster recovery**: important data currently stored only at CERN (Meyrin + backup copy at Prévessin). It would be good to have other options and be able to run critical services from outside CERN in case of outages (e.g. public cloud provider).
  - **GPU availability:** no need for full GPU, but currently no way to know when a GPU is available. It would be good to have reserved slots (e.g. 1 day/week or month).
  - **DBOD level of service**: not adapted to more complex use cases, no DBA, no access to backups (to replicate data to testing instance). We resorted to running our own DB for INSPIRE but the largest local disk we have access to is barely large enough.