



RCS-ICT Technical Committee

The RCS-ICT Technical Committee is an inter-departmental and inter-organizational governance body. Establishes and maintains a roadmap of ICT services and projects jointly engaging the communities in the RCS sector.

Notes of the 15th RCS-ICT Technical Committee held on Friday 21st June 2024

Date: 21st June 2024

Main topic(s): Summary from the 4th RCS-ICT Steering Committee, Discussion on IT-Provided Critical Services for RCS: IT Storage and Data Management Group IT-SD, PSO and Initiatives Status, Action List Status

Agenda (<https://indico.cern.ch/event/1318717/>)

- Summary from the 4th RCS-ICT Steering Committee
- Discussion on IT-Provided Critical Services for RCS: IT Storage and Data Management Group IT-SD
- PSO and Initiatives Status
- Action List Status

Attendance

Present:

Andreas Morsch, Andreas Peters, Armin Nairz, Ben Couturier, Benjamin Bergia, David South, Dirk Duellman, Dmytro Kovalskyi, Enrico Gamberini, Eric Grancher, Jakub Moscicki, Jan Van Eldik, Latchezar Betev, Lorenzo Moneta, Luca Mascetti, Micha Moskovic, Niko Neufeld, Oliver Keeble, Pere Mato, Simone Campana, Wainer Vandelli, Wolfgang Adam, Xavier Espinal, Zach Marshall, Zhechka Toteva

Adoption of Meeting Minutes

Andreas thanked the contributors to the minutes' review. **Minutes were approved.**



Summary from the 4th RCS-ICT Steering Committee (Simone Campana)

Presentation

Simone summarised the presentations, the discussions and the outcome of the last RCS-ICT Steering Committee. The following topics had been addressed:

Scientific Information Landscape

The project consists of the analysis and publication lifecycle management layer, and the repositories layer. The implementation steps consists of:

1. Adopt both Glance for LHC experiments and CAP for SMEs.
2. Enhance CDS for versioning drafts and commenting.
3. Establish a central publication approval workflow.
4. Ensure integration between Glance/CAP and CDS for information consistency.
5. Enable Glance/CAP to transmit the same set of metadata to publication services like arXiv.

3 GRADs and 2 LDs were the requested resources for high and medium priority tasks. The LD in RCS-SIS will support CAP long-term, splitting efforts between the Analysis Preservation Portal (60%) and a lightweight publication lifecycle management system (40%), with the Steering Committee stressing the need to differentiate the roles to avoid confusion. The Scientific Information Landscape task force concluded that customising Glance for SMEs would require significant effort, a point that was partially challenged at the RCS-ICT Technical Committee. While an ad-interim coordinator for Glance is in place, a more sustainable model needs to be put in place for supporting the four large LHC experiments. Although a full LD for this role is considered excessive, large LHC experiments also questioned the relative importance of supporting CAP over Glance as a publication tool for SMEs.

The migration to the new CDS is crucial, as the experiments have not yet been exposed to its functionalities or presented a detailed migration plan. Initial discussions took place at the RCS-ICT Technical Committee in April, and the experiments are eager to be involved early in the process to proactively identify potential issues with the new implementation and to ensure effective migration of the large collections.

Decisions and actions from the SC minutes:

Decision: *the migration to the new CDS and implications for the large LHC experiments should already be evaluated by the existing Scientific Information Landscape working group, which contains the right experts from the experiments. The membership can be adjusted as needed, if people change roles.*

Action: *Alex should write a new request to the working group to define this new task and circulate it to the RCS-ICT Steering Committee by July 15th. Approval will be done via email.*

Action: *Alex should report on the progress of the working group with its new focus at the next Steering Committee in Q4 2024.*

Decision: *there are still diverging opinions on extending CAP vs making a Glance-lite for the SMEs. Glance experts should make a study of what it would imply to develop the needed functionalities for Glance-lite and run it for the SMEs.*



Action: Following the discussion after the meeting, Pippa will try to identify a set of Glance experts to evaluate the feasibility and effort needed for SMEs to use Glance.

Decision: The RCS-ICT Steering Committee considers that it is premature to define resource requests for the 2025 MTP. This should be discussed again at the next Steering Committee.

Document Management Systems update

A document detailing the capabilities currently available at CERN with regard to Document Management Systems was produced by IT in collaboration with other departments. It is currently being reviewed by the STEPS Technical Committee and should be endorsed soon by the STEPS Steering Committee. The document clarifies existing DMS capabilities, emphasises functionality without merging systems, and outlines system lifespans.

Future Requests: Any future requests for additional capabilities must go through the STEPS process to manage potential extra costs.

Public Cloud Strategy

The progress in defining a Public Cloud Strategy, which will encompass a Cloud Policy, Architecture Usage Framework, and establishment of cloud and quantum contracts, was presented. Key milestones include signing the contracts by Q4 2024 and completing the Cloud Operating Model after summer 2024.

Action: the RCS-ICT Steering Committee expects an update once the milestones mentioned

above are met. Simone should schedule it.

Discussion at the Technical Committee: A clarification was made that there would be two separate contracts for Cloud and for Quantum.

Technical Committee: status of ongoing initiatives and new proposals

- Commissioning of RNTuple - fully on track
- Support for Heterogeneous Architectures - Offering new resources will be simpler once the cloud operating model will be available.
- Scientific Data Management - very good progress and collaboration between IT and the experiments. If the RCS' ambition is to go beyond the pilot phase, on which ATLAS and CMS strongly rely on, the current level of effort needs to be maintained both in IT and the experiments.
 - **Action:** a more detailed presentation should be foreseen at the next RCS-ICT Steering Committee, in order to assess the progress.
- Micro-Electronics Cluster (MIC): no-go as the level of effort required for the consolidation is too high with respect to the expected benefits (MIC management).
 - **Decision:** the MIC cluster initiative should be considered closed. The Technical Committee will re-open it if the conditions that led to the decision change in the next few years.
- Collaborative tools: progressing with several successes, pending issue - access for external contributors
 - **Action:** RCS-ICT Technical Committee to present a (set of) proposal(s) at the next RCS-ICT Steering Committee.

- Centralised logbook for the SMEs: finish exploratory phase. EP requested IT to run the service.
 - **Decision**: IT agrees to support this service as part of the IT service portfolio and the provisioning could be as quick as the next six months. Once the service is ready, it should also be advertised more widely, for example to the PS/SPS coordinator and the SPSC. If the scope of the service in future should go beyond the SMEs, the conditions will need to be re-discussed.
- Traffic Marking: It is important in light of a future possible network sharing. Can be carried over with the existing effort in IT. Efforts from the experiments may be required in the future.
- Review of Critical Services: Ongoing
 - **Action**: The RCS-ICT Technical Committee should share the Critical Services review document with the Steering Committee by the end of 2024.”
- Hosting hardware in the CERN IT datacenter infrastructure:
 - **Decision**: IT will consider only requests communicated by the spokespersons of the experiments.”
- Text Transcripts as a Service:
 - **Decision**: CERN needs to support the capability to produce text transcripts online (live) for meetings being attended by the CERN community. The current TTaaS pilot is not considered adequate for this need and will be discontinued in its current implementation and the concerned communities informed. ZOOM text transcripts can be used as an ad-interim solution. CERN IT will investigate different solutions on the market. The evaluation of these solutions should be supported by the CERN community, at the least the large LHC experiments and the Diversity and Inclusion office at CERN.
 - **Action**: In collaboration with relevant communities, IT to start the process of a technical and financial evaluation of different solutions on the market for the online text transcript use case mentioned above. The report will be prepared by the end of the year.
 - **Decision**: TTaaS will continue being supported for the “offline” use case. The engine will however be reevaluated considering more modern technologies.”

Discussion on IT-Provided Critical Services for RCS: IT Storage and Data Management Group IT-SD (Jakub Moscicki)

Presentation

Jakub presented the general and specific uses by the RCS communities of the following IT/SD services:

- General Storage Services
 - CERNBox/EOSHMPM, DFS, AFS, Ceph, CVMFS
- Physics Data Services
 - EOS Physics, FTS, XRootD
 - Recently involved in Rucio devops for ATLAS, CMS and SMEs
- Tapes and Backup
 - CTA Physics/Backup, IBM Spectrum Protect



Jakub demonstrated how different communities used some of the IT-SD group's critical services. The presentation focused on corner cases of reported high criticality or low criticality in comparison with the average criticality. In addition, Jakub explained how misunderstandings of some terms in the criticality matrix can lead to incorrect criticality numbers.

Discussion

Jakub suggested that EOS Physics should be considered separately from EOSHOME/CERNBox. ALICE commented that the term EOS for ALICE is used for EOS Physics.

EOS workspaces, pre-dating EOS projects, are hosted on EOSPUBLIC instance. They should be reviewed by their owners and migrated to EOS Home/Project/Media wherever applicable

CMS is using EOS Home for storing the last stage of the monitoring (action needed).

During the discussion on service criticality, RCS communities agreed that only front-end services should be considered for the criticality assessment (direct usage of IT services). It is the responsibility of the IT department to analyse the internal dependency between the different IT services.

EOS Home/Project is part of the CERNBox service and also its main storage backend. Large majority of CERNBox users do use more than one protocol to access the data (EOS fuse, Web, Sync, CIFS). CERNBox, including EOSHOME, has accumulated hundreds of project spaces and is heavily used by the majority of users in the experiments and various user groups for analysis and other activities.

The majority of ATLAS developers use AFS for their core development and deployment workflows. Whether AFS is required as a shared work environment or as a home filesystem for LXPLUS should be clarified.

For SMEs AFS is getting shorter on space and disk. (an action on Karim)

Lxtunnel does not have dependency on AFS as it is running on AFS-homeless LXPLUS configuration.

Generally, AFS future will be discussed based on the outcome of the LXPLUS future working group and the Analysis Facilities activities.

OpenStack relies heavily on CEPH but most of the users do not distinguish between booting an OpenStack instance and attaching a CEPH volume to it, and booting the whole VM from CEPH. In any case, all of this is documented in the OpenStack documentation.

CephFS may become critical in the context of the container registry.



The ALICE experiment builds its software on S3 and then publishes on the CVMFS.

LHCb uses CVMFS for builds. (to be checked with Ben if he meant CVMFS or S3)

For some LHC experiments, CVMFS stratum 0 service has a higher criticality than CVMFS stratum 1 service, which is quite odd. It was clarified that some experiments took into account the redundancy of CVMFS Stratum 1 when evaluating it with a lower criticality than CVMFS Stratum 0, since it is a single point of failure for deploying new versions of software.

Perhaps the criticality of a service should reflect the actual impact and urgency of the service, rather than taking into account the available redundancy, since redundancy is usually implemented because of the service's high criticality. When discussing redundancy, a given IT service is not aware of the redundancy of the other IT services which it relies on.

The number representing criticality in the table represents the impact of an incident and the urgency of its resolution. As a result, it reflects the severity of the service's unavailability.

CTA has been evaluated by LHCb as not being very critical, since it can be down for a few hours, and yet LHCb has an automatic mitigation strategy to write to a buffer for the first X hours.

FTS has the same criticality as CTA.

FILER is used by MIC and a couple of other services in the IT department.
Rucio is not an IT service and it should not be in the table.

In some experiments, the WLCG services were not explicitly reported, so their input should be taken from the WLCG matrix. A table with an explanation for each value in the matrix should be included in the document.

There were several independent suggestions made:

- In order to understand where the criticality value comes from, it is useful to show the impact and urgency values
- Instead of reporting on the already agreed critical services for the WLCG community, we should focus on the services that aren't in the WLCG service matrix or on those that aren't being used in a standard manner.
 - The input from non-LHC communities in RCS about WLCG services should be collected and analysed if this is agreed.
- There was a suggestion that the numbers should not be analysed with such fine granularity, but rather to investigate the exceptional 0s and 100s.

Action: EOS home and AFS should not be used for critical experiment operations. Understand how CMS can remove it from the critical path when storing the last stage of their monitoring operations.

Action: Understand and reduce the AFS critical for SMEs data taking.

Action: Understand the direct S3 usage by the experiments and its criticality.



Action: Review the methodology for the next round of IT Critical Services presentations, and consolidate the list of IT Critical Services.

Action: IT - Check internally in IT how GitLab relies on S3.

PSO and Initiatives Status (Xavier Espinal)

Xavi presented the current status of the PSOs.

Action List Review (Zhechka Toteva)

Zhechka presented the current state of the action list.

AOB

Nothing to add.

Action List

- **2024-15-A1 - Action on CMS and IT/SD:** EOS home and AFS should not be used for critical experiment operations. Understand how CMS can remove it from the critical path when storing the last stage of their monitoring operations.
- **2024-15-A2 - Action on Karim and IT/SD:** Understand and reduce the AFS critical for SMEs data taking.
- **2024-15-A3 - Action on IT/SD:** Understand the direct S3 usage by the experiments and its criticality.
- **2024-15-A4 - Action on RCS chairs:** Review the methodology for the next round of IT Critical Services presentations, and consolidate the list of IT Critical Services.
- **2024-15-A5 - Action on IT/SD:** IT - Check internally in IT how GitLab relies on S3.

Comments/Amendments

- ...