


1 December 2022

# CMS Input to the RCS-IT Engagement Process





# Contents

<b>1 IT-CMS Activities</b>	<b>4</b>
1.1 SAM/ETF . . . . .	4
1.2 HammerCloud . . . . .	4
1.3 Proposed activity: Study of the premixing workflow of CMS . . . . .	5
1.4 Rucio . . . . .	5
1.5 XRootD . . . . .	5
<b>2 Infrastructure and Resources</b>	<b>5</b>
2.1 Network . . . . .	5
2.2 Online Oracle Database . . . . .	6
2.3 Network monitoring . . . . .	6
2.4 Non-x86 and Heterogeneous platforms . . . . .	6
2.5 Linux . . . . .	6
<b>3 Physics and General Services needed by CMS</b>	<b>6</b>
<b>4 Services and Tools for communication</b>	<b>7</b>



CMS is very satisfied with the overall performance of the CERN IT service portfolio. To be concise, this document only highlights the opportunities for improvement, needs, and does not describe in detail the good performance of the various services.

This list is just a first step to start the engagement process. For example, a service which is not mentioned is not considered unnecessary by CMS. The WLCG critical services table is still valid and should continue to be considered with priority.

Important elements of the long term scientific software and computing strategy of CMS are summarised in the CMS Phase-2 Computing Model: Update Document, which should be seen as the latest version of the CMS O&C roadmap to HL-LHC, known and understood by all interested parties.

We qualify each item with a level of criticality 1, 2, and 3 and a brief description of the impacts of not having the service.

1. Without, CMS cannot run, and there is no alternative.
2. Without, CMS can continue to run but less efficiently and at a higher cost, not only for CERN, but also other funding countries, and there is no practical alternative.
3. Without, CMS can continue to run, and there are alternatives but perhaps less secure and more costly.

## 1 IT-CMS Activities

In the following we focus on ongoing or proposed activities to be done in collaboration between CMS and IT.

### 1.1 SAM/ETF

Test and monitoring infrastructure: framework development/enhancements/adaptation to new grid and experiment needs/requirements.

- Goal: maintain the ETF framework, adapt it to changes in the relevant middleware, add support for token authentication, ensure full IPv6 support. On the MONIT side, maintain the SAM test results publication mechanism and related user interfaces
- Timeline: currently ongoing, should be preserved
- Criticality: 2
- Impact of not having: Inability to account for fulfilment of pledges to WLCG, no nearly real-time calculation of WLCG site availability and reliability metrics, resulting in less efficient utilisation of resources world-wide, delays in production and analysis, etc.

Note: SAM/ETF is as important as HammerCloud for CMS. Both are essential common software for accounting, availability and reliability metrics of WLCG sites.

### 1.2 HammerCloud

Maintenance and operation of the tool, including updates of the CMSSW release used for the tests, or upgrade of the Python version.

- Goal: maintain the HC framework, migrate the code to Python 3 to ensure long term support, support all current and future versions of CMSSW, improve the flexibility of the framework to support more types of jobs, ensure smooth operations of the HC service
- Timeline: currently ongoing, should be preserved
- Criticality: 2
- Impact of not having: No nearly real-time calculation of WLCG site availability and reliability metrics for job submission, resulting in less efficient utilisation of resources world-wide, delays in production and analysis, etc.

### 1.3 Proposed activity: Study of the premixing workflow of CMS

This activity focuses in particular on the access pattern through ROOT and XRootD of the pileup mixing technique adopted by CMS. On demand, all details about the idea can be added. We underline the value of the activity for other LHC experiments, too.

- Goal: Benchmark and study the CMS premixing workflow, form an optimization strategy of the access pattern of remote data as well as the necessary code for XRootD, ROOT or CMSSW.
- Timeline: short term (the sooner the better)
- Criticality: 3
- Impact of not having: Less efficient utilisation of the computing resources. CMS could potentially do this study, but less effectively or efficiently, lack of sufficient effort.

### 1.4 Rucio

This is not yet an IT service, nor a tool developed in IT. However we note and support the recommendation made by the LHCC in the context of the computing model review of November 2021: *“CMS is congratulated on the recent adoption of Rucio and the currently excellent relationship with the Rucio team is noted. However, we also understand the concern that the core Rucio developers are completely embedded in ATLAS and that, hypothetically, this could lead to tension when setting priorities in the future. This risk would be mitigated by complementing the current Rucio development effort with additional experiment-independent effort and, indeed, by increased CMS contributions to the core Rucio components, leading to a more equitable ownership”*.

- Goal: complement existing experiment specific effort with an activity hosted in IT, where experiment specific investments have the highest probability to land.
- Timeline: depends entirely on CERN IT.
- Criticality: 2
- Impact of not having: Impacts of risks mentioned in the LHCC recommendation above.

### 1.5 XRootD

XRootD is at the heart of the CMS computing model, used for pileup simulation, analysis, and remote processing at opportunistic/storageless sites.

- Goal: The work on the tool should continue in the storage group, aiming to constantly improve the performance, integration with the HEP ecosystem (most notably ROOT), robustness and scalability.
- Criticality: 1
- Impact of not having: Inability to produce Monte Carlo efficiently and at scales needed for Phase 2.

## 2 Infrastructure and Resources

In the following we focus on the aspects closely related to the computing infrastructure and resources.

### 2.1 Network

Refurbishment and extension of the network for CMS in P5 for the HL-LHC period. Part of the work is during Run 3 (3562-R), the rest is during LS3 (USC55 and 3562-\*). The latter part is time constrained and on the critical path for the restart of CMS for Run 4. Project already presented to IT Engagement. CMS needs to identify a contact person for this activity in IT for discussing costing, effort and timelines.

- Criticality: 1
- Impact of not having: Inability to take control the detector, take data and transfer detector data at sufficient rates during Phase 2.

## 2.2 Online Oracle Database

This database is needed to take data, and its support is required.

- Criticality: 1
- Impact of not having: Inability to take data.

## 2.3 Network monitoring

For CMS it is very important to have coherent and reliable network monitoring (e.g. Grafana dashboards), where the contributions of individual experiments are visible. The questions to be answered by this monitoring are for example what links or sites are most loaded, by what experiment and why (e.g. scheduled transfers or XRootD traffic). All other activities such as SDN R&Ds, packet marking, and pacing have lower priority and can get more priority once an assessment of the current network usage is available and an overall strategy for an optimized usage of this resource is established.

- Criticality: 1
- Impact of not having: Potentially huge if network becomes a constrained, shared resource now and even more during Phase 2 (offline). Loss of visibility and understanding of issues and constraints affecting data taking performance (online).

## 2.4 Non-x86 and Heterogeneous platforms

CMS needs to have at disposal nodes provisioned with Power and ARM CPU for development, integration and, potentially, physics validation. Moreover, infrastructure and support for Gitlab CI for FPGA/SoC synthesis and simulation is needed by the online community. GPU provisioning for development and integration is also required (no need for resources for validation for the moment). The kind of GPUs should not only include NVIDIA (“small” GPUs such as T4 in first approximation) but also AMD, and eventually Intel, when hardware becomes available.

- Criticality: 2
- Impact of not having: Inability to use heterogeneous platforms at CERN and WLCG sites, resulting in HEP software at CERN effectively becoming legacy code.

## 2.5 Linux

CMS online has interest in having the support for CC7 until 6 months after end of Run 3, like the ATS sector. For offline processing, CMS moved to AlmaLinux during Q2 2022. We would find it very beneficial if CERN converges on a RHEL clone instead of CentOS Stream, and this clone is AlmaLinux.

- Criticality: 2
- Impact of not having: Status quo ante - Having this would potentially reduce costs and effort needed from the experiment.

N.B. Regarding transitioning to a RHEL Clone, we are willing and happy to share the experience and testing infrastructure CMS built up during the evaluation of Linux distributions at the beginning of 2022 and led the experiment to choose AlmaLinux.

## 3 Physics and General Services needed by CMS

In the following, input about opportunities for improvement is provided about important services for CMS.

- **CVMFS**, for software distribution and CMS Offline CI/CD infrastructure `unpacked.cern.ch` repository for container distribution `grid.cern.ch` repository for certificates. We would like to make use of REANA and the CERN Analysis Preservation Portal for that. While the computing resources for the services in general seem to be sufficient, we'd appreciate increased person power in particular on the REANA side to address specific CMS needs.

- **K8S**: hosting CMSWEB services of CMS, which are critical for data taking.
- **CERNBox**: a storage element to collect the output of user Grid jobs (*T3.CH\_CERNBOX*). The way to share sensitive data. Synchronized and shared storage, also accessible from lxbatch/plus.
- **OpenStack**: hosting some CMS internal services critical for data taking.
- **Centralised Elasticsearch Service**, including MONIT
- **Other important services** (not less important, but nothing very concrete to suggest):
  - **Batch**
  - **Configuration Management Service**
  - **IAM** (and VOMS until its retirement)

## 4 Services and Tools for communication

This is a special category of tools and services, not used for core scientific computing, but are crucial for the operations of an experiment that is geographically distributed among several institutes.

- **Authentication and SSO** (critical). This also includes the auth-get-sso-cookie utility, needed for programmatic access to websites/services behind CERN firewall.
- **Discourse**: crucial for CMS. It replaced HyperNews and became a veritable support for the CMS community. Hard to imagine to go back. Need to be sure that backups are available and the history can be restored.
- **Mattermost**: Crucial for secure communication. Need to be sure that backups are available and the history can be restored.
- **Tools for web static pages** (e.g. EOS) as well as dynamic apps (e.g. via OpenShift) and via gitlab CI/CD and mkdocs and hosting services; web analytics
- **CERNPhone**, in particular
  - Preserving the capability of having service numbers linked to egroups. This is a key feature for operations, when an on call person is needed and redirecting calls from the same number are required.
  - We don't receive all the calls (average of 50%), meaning that the phone is not ringing at all.
  - There is also very often some issue with the quality of the sound and some "delay" with the sound
- **Gitlab**: We expect an increase in use of central GitLab CI runners for physics analysis testing.
- **Other important services** (not less important, but nothing very concrete to suggest):
  - **CodiMD** (also in connection to CERNbox)
  - **Indico**
  - **CDS**
  - **Twiki** or equivalent (e.g. FosWiki)
  - **WordPress**
  - **JIRA** or its replacement
  - **Email**
  - **Egroups** or equivalent
  - **Ticketing**
  - **Zoom**
  - **Notification** service (not available, but it would be useful to have)
  - **Newdle** or equivalent
  - **OpenShift**
  - **LimeSurvey**