

CMS input to the RCS-IT Engagement Process



Marc Dobson, James Letts, Danilo Piparo for the CMS Collaboration

Disclaimer

- Thanks for organising this meeting: it's a pleasure to be here
- CMS is very satisfied with the overall performance of the CERN IT service portfolio. To be concise, this input 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.

This Input

- These slides summarise a **rather complete document** prepared in consultation with many areas of **CMS**, inclusively (attached to the agenda of this meeting)
 - Sections of the document: **IT-CMS activities, Infrastructure and Resources, Physics and General Services, Services and Tools for Communication**
- We wanted to highlight **two categories of items, activities and services**, important for us
 - Criticality level (1,2,3), timelines (short, medium and long term) and description of the impact if absent provided
 - Started from what exists, but also proposed new ones
- CMS is satisfied with the overall level of IT services and collaboration. We nevertheless tried to **highlight the opportunities for improvement**
- The engagement process is new to us and it is not yet clear how it works
 - **We want to make it work for CERN and CMS.** Providing input has a non-negligible cost.
- This input is just a first step, which does not replace other inputs previously provided, such as WLCG Critical Services.
 - **If something is not mentioned by CMS yet, it does not mean it's not relevant**
- To facilitate engagement in the future we strongly suggest to become (even more) familiar also with the documents listed at the end of these slides

Criticality Level

Explanation of the criticality levels (1,2, or 3)

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.

Contents

1 IT-CMS Activities

- 1.1 SAM/ETF
- 1.2 HammerCloud
- 1.3 Proposed activity: Study of the premixing workflow of CMS
- 1.4 Rucio
- 1.5 XRootD

2 Infrastructure and Resources

- 2.1 Network
- 2.2 Online Oracle Database
- 2.3 Network monitoring
- 2.4 Non-x86 and Heterogeneous platforms
- 2.5 Linux

3 Physics and General Services needed by CMS

4 Services and Tools for communication

Highlights: Items with Criticality 1

- **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.
 - Impact of not having: Inability to produce Monte Carlo efficiently and at scales needed for Phase 2.
- **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.
 - Impact of not having: Inability to take control the detector, take data and transfer detector data at sufficient rates during Phase 2.
- **Network monitoring** For CMS it is very important to have coherent and reliable network monitoring (e.g. Grafana dashboards, and Spectrum), allowing notification of issues and identification of bottlenecks in the experimental site, and further out into the CERN network 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.
 - 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 Oracle Database** This database is needed to take data, and its support is required.
 - Impact of not having: Inability to take data.

Highlights: Items with Criticality 2

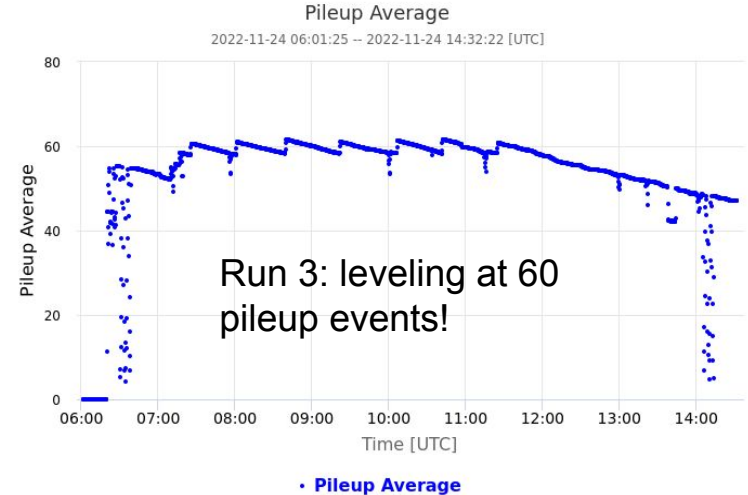
- **SAM/ETF** and **HammerClud** (2 items): SAM/ETF is as important as HammerCloud for CMS. Both are essential common software for accounting, availability and reliability metrics of WLCG sites.
- **Rucio** This is not yet an IT service, nor a tool developed in IT. However we note and support the recommendation 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.
 - Impact of not having: Impacts of risks mentioned in the LHCC recommendation above
- **Non-x86 and Heterogeneous platforms:** needed for CMS CI (non-x86 and accelerators). Needed not to become *“Legacy Computing”*
- **Linux:** CMS online has interest in having the support for CC7 until 6 months after end of Run 3 like ATS sector. 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

The Proposal

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)
 - 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.

Pileup: uninteresting p-p collisions happening during a bunch crossing on top of the interesting hard-scatter collision. They generate occupancy in the detector and have to be simulated.



Useful Documents

- The Phase-2 Upgrade of the CMS Data Acquisition and High Level Trigger [\[link\]](#)
- CMS Phase-2 Computing Model: Update Document [\[link\]](#)
 - ... And recommendations by the LHCC [\[link\]](#)
- Last IT R&D Advisory Group Meeting [\[link\]](#)