CMS Tier-0 data processing during the detector commissioning in Run 3

Jhonatan Amado[1,2], *Jan* Eysermans[3], *German* Giraldo[3], *Dirk* Hufnagel[1], *Dmytro* Kovalskyi[3], and *Antonio* Linares Sancho[1,2], on behalf of the CMS Collaboration

[1]Fermilab, P.O. Box 500, Batavia, IL 60510-5011, USA
[2]University of Wisconsin-Madison, 500 Lincoln Dr, Madison, WI 53706, USA
[3]MIT, 77 Massachusetts Ave, Cambridge, MA 02139, USA
[4]CERN, Esplanade des Particules 1, P.O. Box, 1211 Geneva 23, Switzerland

Abstract. The CMS Tier-0 system is responsible for the prompt processing and distribution of the data collected by the CMS Experiment. A number of upgrades were implemented during the long shutdown 2 of the Large Hadron Collider, which improved the performance and reliability of the system. In this report, these upgrades are discussed and a more detailed description of the Tier-0 system is given. The experience of the data taking during Run 3 detector commissioning as well as the performance of the system are highlighted.

1 Introduction

The Compact Muon Solenoid (CMS) is a general-purpose detector at the CERN Large Hadron Collider (LHC), located 100 meters underground close to the France-Swiss border near Geneva [1]. Data taking started in 2009 with Run 1, continued in 2015 (Run 2) and in 2022 Run 3 started which will last for 4 years. Between the running periods, long shutdowns allow for the accelerator, detector, and computing experts to perform upgrades and maintenance activities that pave the way for collisions with even higher intensity.

Millions of detector readout channels are read out by the Data Acquisition System (DAQ) at a collision rate of 40 MHz. A 2-level trigger system, the Level-1 (L1) and the High-Level Trigger (HLT) filter the events to reduce the output rates to a manageable level of a few kHz. Events that pass the trigger criteria are transferred to Tier-0 for further processing. It is the responsibility of Tier-0 to ensure the data taken is usable by repacking and reconstructing it, safe storage, and providing convenient formats (data tiers) for further offline analysis [2].

An overview of the Tier-0 system is graphically depicted in Fig. 1. Data is continuously injected via the Storage Manager (SM) as streamer files, which are immediately repacked by converting it to the RAW event-level format. Depending on the physics content of the events, the events are split into different Primary Datasets (PD). For safe data storage and eventual secondary reprocessing, the RAW event format is stored at two different tape sites. Simultaneously to repacking, express processing is performed for alignment and calibration (AlCa), and data quality monitoring (DQM) purposes. The outcome of the express workflow is a set of time-dependent detector conditions that are uploaded to the condition database. Finally, the RAW events are (promptly) reconstructed 48 hours after repacking, as it relies on the condition database to incorporate the alignment and calibrations.

The CMS Tier-0 system encompasses various components that are orchestrated by a workflow management system and can be classified into 4 parts, as shown in the yellow boxes of Fig. 1: 1) the input data management consists of the reception of the data through the SM; 2) the heart of Tier-0 is processing infrastructure to provide the necessary computing power and storage to execute the workflows; 3) alignment and calibration and DQM; and 4) the output data management for data storage at the Tier-1 sites. Furthermore, Tier-0 makes use of the CMS software CMSSW for repacking, express, and reconstruction.

During the Long Shutdown 2 (LS2), various upgrades were performed on all levels of the Tier-0 infrastructure, which is summarized in Section 2. A more detailed description of the various Tier-0 system is given in Section 3. During operation, frequent changes in the configuration and CMSSW release or (less frequent) infrastructure upgrades are possible, therefore it is important to have extensive test capabilities to test the entire chain, as will be explained in Section 4. Section 5 describes the performance of the Tier-0 system during the Run 3 commissioning and data-taking periods.
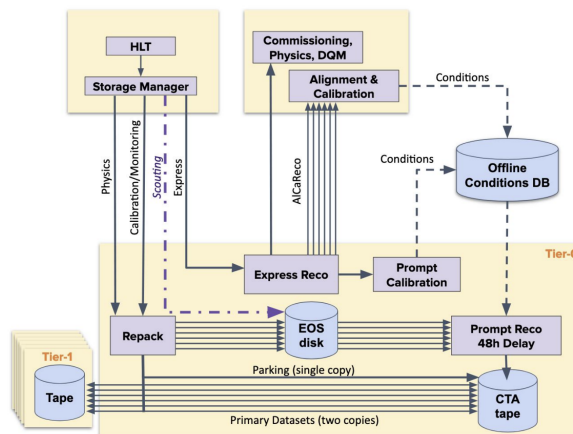

Figure 1. Overview of the CMS Tier-0 system.

2 Upgrades during LS2

2.1 Processing infrastructure

Tier-0 dedicated disk storage area. During Run 2, computing and storage resources for T0_CH_CERN and T2_CH_CERN were merged into T2_CH_CERN. The Tier-0 system started sharing the T2_CH_CERN batch pool with analysis and central production jobs, which allowed for efficient use of all worker nodes, especially during reduced Tier-0 activity (e.g. short technical stops or unexpected issues with the LHC). While sharing disk storage allowed for similar space optimization, it also exposed Tier-0 to the possibility of central production jobs taking over Tier-0 disk space. This was a permanent risk for Tier-0 operations, as running out of disk storage would mean a complete halt of both input and output data operations.

With the purpose of eliminating this risk factor, T0_CH_CERN_Disk site was recreated under the total control of Tier-0. This disk space is used by Tier-0 to store input data received from SM, and output data while it is transferred to its final destination at Tier-1 sites.

T2_CH_CERN batch system use of SSDs. Repacking jobs consist of reading streamer files, sorting their content, and writing them as RAW files. They require high I/O rates and relatively low CPU usage. While the worker nodes during Run 2 could easily handle the processing load of these jobs, running several repack jobs on the same worker node would overload their HDD. To prevent this from happening, Tier-0 limited the number of repack jobs that could run simultaneously in a given worker node.

As part of the upgrades of CERN's batch system during LS2, the HDDs used for scratch storage were gradually replaced by faster SSD devices. This upgrade allowed Tier-0 to remove limitations on concurrent repack jobs, thus achieving faster repacking speeds. During repacking tests performed in June 2021, Tier-0 was able to achieve writing speeds to T0_CH_CERN_Disk of around 30 GB/s, limited by the network link from batch to EOS storage rather than by the I/O capabilities of the batch system itself.

Transition to Python 3. Given that Python 2.7 reached its end of life, during LS2, Tier-0 had to update its code to be compatible with newer Python 3 versions. The transition was completed in October 2021, alongside the release of CMSSW_12_0_0, the first Python 3 exclusive CMSSW release.

2.2 Input/Output data management

Increased in network capacity at CERN. The CERN network capacity was increased during LS2, which both increased the EOS performance and the CERN tape storage performance. Stress tests and tape challenges were performed, resulting in a stable 14 GB/s input transfer to tape and an overall 20-25 GB/s data traffic during collisions with outgoing transfers to CERN disk and tape.

Improvement in Storage Manager. DAQ and SM installed new storage hardware, which made it possible to achieve stable data-taking rates up to 13 GB/s, without producing an instantaneous transfer backlog to Tier-0. Higher rates up to 17 GB/s were tested but a backlog built up, which eventually could be resolved during the interfill. The high rates were generated by dummy data, a feature of the DAQ that was used for several stress tests for DAQ, SM, Tier-0, and others (e.g. tape challenges).

Transition from PhEDEx to Rucio. During Run 2, PhEDEx (Physics Experiment Data Export) was the software in charge of managing the transfer of CMS data among the different sites of the collaboration. Tier-0 made use of the software to make sure its output data was properly transferred to Tier-1 sites. During LS2, PhEDEx was phased out in favor of Rucio, which became the new data management software for CMS [3].

Rucio organizes files in *Datasets* (equivalent to blocks in CMS), which in turn are grouped in *Containers* (equivalent to datasets in CMS). Instructions to create replicas of a Dataset/Container in specific storage sites are managed by *Rucio rules*. In order to transition Tier-0 from PhEDEx to Rucio, a strategy was defined to fulfill Tier-0's output data management policies, by using the features offered by Rucio. The main objectives for Tier-0 were to secure the output files in Tier-0 disk, to promptly transfer processed data to the final destinations, and to promptly delete data from Tier-0 disk that has already been transferred (see Section 3.5). The final solution required changes in the WMAgent project code to automatically create and delete Rucio rules, as well as the deployment of a dedicated Rucio daemon in T0_CH_CERN_Disk to ensure the prompt deletion of files in the site.

3 Overview of CMS Tier-0

3.1 Workflow Management System (Agent)

Tier-0 has to reliably organize, execute, and supervise a large amount of computational work. While doing this, it has to coordinate with several systems to accomplish tasks like data transfers, bookkeeping, job execution, quality control, etc. The system responsible for managing all these interactions is the Tier-0 agent.

The Tier-0 agent is a workflow management system built as a modification to the central CMS WMAgent system [4]. It consists of several daemons, called *components*, each one of them in charge of a specific task. The different components coordinate their work via the use of a local relational database, called T0AST in the Tier-0 agent. The components also write to a local instance of CouchDB that stores workflow monitoring information and replicates that data to a central monitoring system called *T0 Request Monitor*. Finally, the agent uses databases to communicate with external systems like SM, HLT, and T0 Data Service. Fig. 2 shows the different components and databases that are part of the Tier-0 agent.
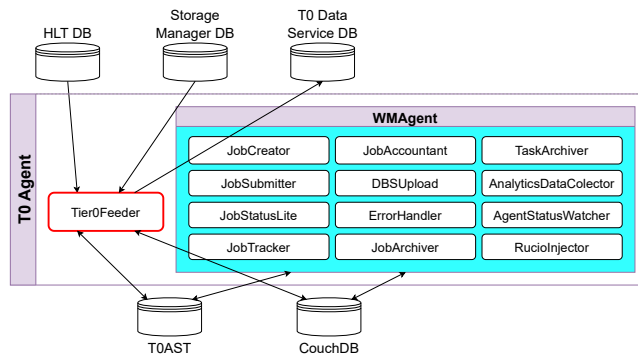

Figure 2. Tier-0 agent components and databases.

The main goal of the Tier-0 agent is to manage the processing of all available input data and handle the placement of the output data in the desired storage sites. To do this, all the computational work is divided into three types of workflows:

- Repack workflows are in charge of repacking the streamer files into the RAW data tier. Each workflow handles data from a single run and a single stream. A single Repack workflow can produce RAW data for several PDs, according to the content of the input stream.

- Express workflows generate data quality and calibration data that is used by the AlCa and DQM systems. Similar to Repack workflows, express workflows use streamer files, however, these streams contain smaller amounts of data so they can be processed more quickly as their output must be generated before the release of PromptReco workflows.

- PromptReco workflow performs the full reconstruction of the physical events. It uses the RAW data as input for a single run and PD combination, generated by Repack workflow. It also requires the culmination of Express workflows for the run that is being processed, to ensure that the proper calibration conditions are used during the reconstruction.

The Tier-0 agent is deployed and runs on a virtual machine from where it executes its tasks. Typically, the agent is deployed when a new data acquisition era is requested due to changing detector/machine configurations or more sophisticated changes in reconstruction (on average once per month).

3.2 Input data management

Per LHC bunch crossing, the DAQ system receives the detector readouts, builds the event, and groups them into successive periods of 23 seconds called lumisections. Events in each lumisection are filtered by the L1 and the HLT trigger systems, resulting in a series of streamer files, grouped by streamers based on the physical content they contain, as a result of the HLT decision. These streamer files are the input for Tier-0 to perform its job.

The SM is in charge of sending the streamer files to Tier-0, which belongs to the DAQ group. SM organizes the transfer of data from the Lustre storage area located at Point 5 to the EOS storage area located at CERN's data center [5]. New files are registered into the SM database, and once the files reach EOS, they are available to be processed by Tier-0.

Tier-0 communicates with SM through the SM database. It continuously checks the SM database for new files to process. Upon the arrival of new files, the relevant information about the run and stream of the file is gathered from the metadata available in SM and HLT databases. The streams and runs are injected into the T0AST database, and the files are added and marked as *checked* in the T0 Data Service database. The agent prepares the corresponding Repack workflow jobs, and once the file has been successfully repacked, it is marked in T0 Data Service database as *repacked*. SM continuously checks the status of files in T0 Data Service, to know when it is safe to delete files from the Lustre storage area.

The administration of the EOS streamer area is the responsibility of Tier-0. This is achieved by an automated streamer deletion script that runs every 12 hours to clean up streamer files. Streamer files are deleted if the usage of the streamer area is over a predefined *usage threshold* (oldest streamers are deleted first), and if the streamers are older than a certain *age threshold*. The minimum values for *usage threshold* and *creation threshold* is specified in the Tier-0 policies, and are typically 70 % and 7 days respectively.

These policies ensure that, under normal circumstances, there are 2 copies of every streamer file: one in the Lustre area, and one in the EOS streamer area. Once repacking is finished, the Lustre copy is deleted, leaving only the EOS copy. If the Lustre area is running out of space, SM may prematurely delete their copy of the streamer files, although this should only happen in extreme cases.

3.3 Processing infrastructure

CMS makes use of the Worldwide LHC Computing Grid (WLCG) for its computing and storage resources. These machines are virtually fragmented into various multi-core systems to optimize their use and are accessible by submitting jobs via HTCondor. A job is submitted with specific requirements such as the number of cores or memory availability, and begins to run when these resources become available on the grid.

Tier-0 mainly uses the resources at T2_CH_CERN (Meyrin datacenter) and T2_CH_CERN_P5 (CMS Point 5). In case data is taken at very high rates, or as a fallback option when the T2_CH_CERN_P5 resources are temporarily reduced, more resources are required and Tier-0 can make use of machines at Tier-1 sites.

Jobs submitted by Tier-0 have different core requests, depending on their workflow and task. For example, PromptReco and Express jobs require 8 cores, while Repack jobs are usually submitted as single-core.

Given the importance of Tier-0 computing over other computing operations at CMS, Tier-0 jobs are submitted with the highest priority to obtain the necessary resources quickly. Specifically, Tier-0 will get the resources until it has taken up to 80 % of all available cores.

3.4 Alignment, Calibration, and Data Quality Monitoring

Due to the complexity of the CMS detector and the relatively large periods of data taking, automatic alignment, calibration, and monitoring are crucial in the online data-taking process. It is the task of Tier-0 to provide the necessary information to these external systems within a few hours after data taking, as it is used during the event reconstruction and as monitoring by experts or shifters. Tier-0 distinguishes two types of workflows, the Prompt Calibration Loop (PCL) and DQM, as shown in Fig. 3.
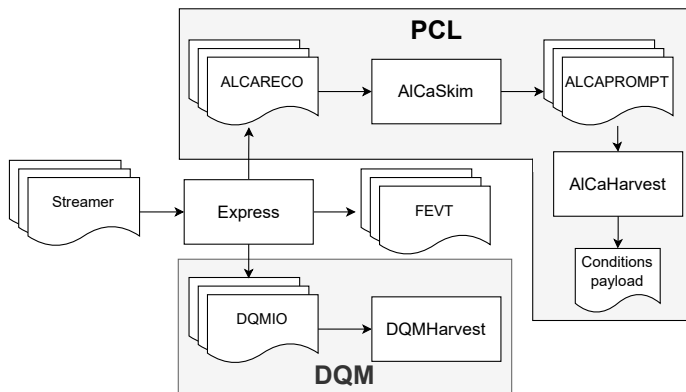

Figure 3. Tasks and outputs of the PCL and DQM Express workflows.

Prompt Calibration Loop. The PCL is a series of processes that generate time-dependent information about the conditions of the detector at the moment of the data taking. This information is then used during the PromptReco workflow for optimal event reconstruction.

To generate the necessary detector conditions, several processes in CMSSW are defined in the form of AlCa producers. Within an Express workflow, the initial Express jobs use the defined producers to generate ALCARECO files that contain all the defined products. From these files, skim jobs split the products into different ALCAPROMPT datasets. After the processing is finished, a harvesting job recovers all the AlCa products from the ALCAPROMPT files and generates a single SQLite payload that contains the run's data-taking conditions. Finally, the Tier0Feeder component of the Tier-0 agent (as described in Section 2), uploads the payload to the conditions database.

Data Quality Monitoring. DQM consists of Monitoring Elements (ME), typically exposed as histograms, that provide useful information to determine the quality of the data that was taken for any particular run. Express and PromptReco workflows generate a selection of MEs according to the DQM sequences that are added to their configuration. After reconstruction, the generated MEs are stored in tree-based DQMIO files. Finally, a harvesting job recovers the generated MEs and uploads them to the DQM offline service.

3.5 Output data management

Each of the Repack, Express, and PromptReco workflows produces its own output data, which is temporarily stored at T0_CH_CERN_Disk. The data, at low-level ROOT files, is organized by block (as a group of files), and per PD (as a group of blocks). Once produced, they are injected by the Tier-0 agent (see Fig. 2) into the Dataset Bookkeeping System (DBS)

for bookkeeping purposes and into the Rucio system that is responsible for the output data management.

Tier-0 uses Rucio to transfer output data to the various storage sites (i.e. disk and tape sites). The destination sites for the output data of a particular PD are defined in the Tier-0 configuration. Destinations are chosen by balancing the expected size of the PDs with the storage capacity of the Tier-1 sites. The Tier-0 agent automatically assigns Rucio Container rules to each newly created dataset. Then, Rucio is in charge of replicating the corresponding files to all the desired destination sites.

Moreover, Rucio is also used to manage the deletion of data from T0_CH_CERN_Disk. This is achieved by assigning a Rucio rule to each new block Tier-0 creates, binding the block to T0_CH_CERN_Disk. When these rules are removed by the Tier-0 agent, it signals Rucio that it is safe to delete the blocks. The Tier-0 agent removes a block rule if the following three conditions are met: The block is no longer needed for processing, it has been copied to all its destination sites, and it was created for more time than the deletion delay (typically 7 days).

4 Testing procedures and capabilities

Apart from hardware or software upgrades to the Tier-0 infrastructure and/or the workflow management system, the most frequent (weekly, on average) changes are driven by configuration changes or the usage of a newer CMSSW release. For each such change (unless the risk of failure is proven to be very low), the entire Tier-0 processing chain has to be tested a priori in order to avoid potential failure of the workflows when taking data.

It is the responsibility of Tier-0 to test any changes before deployment in production. This is done using a replay, which consists of reprocessing the streamer files for a given run. Indeed, for dedicated runs that could be useful for potential future testing, the corresponding streamer files are kept on disk, by excluding them from the streamer deletion script (see Section 3.2). A dedicated Tier-0 agent is set up on a virtual machine that performs the full reprocessing (Repack, Express, and PromptReco) for that run. Monitoring of the replay is managed through WMStats, which is a tool that feeds from a dedicated CouchDB instance to provide the status of the workflows performed by the agent. Since data is being reprocessed, it is necessary to differentiate the replay tests from production, to avoid the risk of data duplication. For this, Tier-0 has dedicated T0AST databases for replays, and datasets are uploaded to DBS with a slightly different name. For storage purposes, the output data from replays is not saved to tape and is only kept temporarily on disk for eventual further checking.

To reduce the manual work to be done by the Tier-0operators and to ease replay requests by non-experts, replays can be requested and deployed automatically via the Tier-0 GitHub repository by creating a new pull request with the desired configuration. Before deployment of the replay, a Python syntax check is in place to avoid erroneous configurations.

Replay tests are quite cumbersome and rely on the entire Tier-0 infrastructure. More advanced testing capabilities are under development such as unit tests targeting dedicated workflows or routines. Furthermore, the ability to perform (partial) reprocessing on the local virtual machine can be useful for quick debugging.

5 Early commissioning and performance

Run 3 started in 2022 with a slightly higher center-of-mass energy than Run 2. During the first year, CMS recorded an integrated luminosity of 38.48 fb^{-1}. The average input(output) data volume was about 2(3) GB/s, roughly 1.5 times larger compared to the end of Run 2. Tier-0 managed to repack and reconstruct all the data, including the parking datasets for which it

was initially foreseen to reconstruct only 25 %. The overall performance was good, reaching 100 % availability without any processing delays.

Initially, the T2_CH_CERN computing resources were used for repacking and reconstructing the data, where a nominal usage of 20k cores was achievable (roughly 50 % of the site capacity), distributed over single and multi-core jobs. Towards the end of the year, the T2_CH_CERN_P5 site was gradually commissioned by running prompt reconstruction on part of the parking datasets. Thanks to the site support, T2_CH_CERN_P5 was fully commissioned as a reliable site and Tier-0 was able to get an additional 20k cores. Since then, reconstruction of the data has been performed on both sites.

In terms of output data management, 2022 was a year of large-scale testing of Rucio. Detailed monitoring and close follow-up on the transfers by the Tier-0 operators made it possible to directly spot critical issues, which were addressed and solved by the CMS Rucio development team. The transfer backlog to the Tier-1 tape sites was kept under control at all times, and a stable transfer regime was established, that continued in 2023.

6 Summary

In order to cope with the higher expected luminosity by the LHC and accordingly larger data volumes to process during Run 3, several upgrades to the CMS Tier-0 system were foreseen during LS2. They ranged from hardware upgrades on the infrastructure and resources, utilized for repacking and reconstruction of the data as core tasks of Tier-0, as well as several software upgrades. Notably, the major transition from PhEDEx to Rucio for output data management was crucial.

Before the start of Run 3, dedicated tests were performed to ensure the validity of the entire chain of data processing at Tier-0. During Run 3, which started in 2022, Tier-0 was able to process all the data with good performance and high availability. The outgoing transfers using Rucio were followed up closely without producing a considerable backlog.

Tier-0 will continue its operations in the current configuration till the end of Run 3. It has several handles in dealing with the required resources (processing capacity and disk space), that allow for flexible operation during the course of Run 3.

References

[1] Chatrchyan S. and others [CMS Collaboration] "The CMS experiment at the CERN LHC. The Compact Muon Solenoid experiment." *JINST*. **3** pp. S08004 (2008), https://cds.cern.ch/record/1129810

[2] Hufnagel, D., [CMS Offline and Computing] "The architecture and operation of the CMS Tier-0." *Journal Of Physics: Conference Series*. **331**, 032017 (2011,12), https://dx.doi.org/10.1088/1742-6596/331/3/032017

[3] Vaandering, Eric "Transitioning CMS to Rucio Data Management." *EPJ Web Conf.*. **245** pp. 04033 (2020), https://doi.org/10.1051/epjconf/202024504033

[4] Cinquilli, M. and others "The CMS workload management system." *Journal Of Physics: Conference Series*. **396**, 032113 (2012,12), https://dx.doi.org/10.1088/1742-6596/396/3/032113

[5] Andre, J. and others [CMS DAQ] "Online Data Handling and Storage at the CMS Experiment." *J. Phys. Conf. Ser.*. **664**, 082009 (2015)