

Repurposing of the Run 2 CMS High Level Trigger Infrastructure as a Cloud Resource for Offline Computing

Marco Mascheroni^{1*}, Antonio Pérez-Calero Yzquierdo^{2,3**}, Edita Kizinevic⁴, Farrukh Aftab Khan⁵, Hyunwoo Kim⁵, Maria Acosta Flechas⁵, Nikos Tsipinakis⁴, Saqib Haleem⁶, Damiele Spiga⁷, Christoph Wissing⁸, and Frank Würthwein¹ on behalf of the CMS Collaboration

¹University of California San Diego, La Jolla, CA, USA

²Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas (CIEMAT), Madrid, Spain

³Port d'Informació Científica (PIC), Barcelona, Spain

⁴European Organization for Nuclear Research (CERN), Geneva, Switzerland

⁵Fermi National Accelerator Laboratory, Batavia, IL, USA

⁶National Centre for Physics, Islamabad, Pakistan

⁷INFN Sezione di Perugia, Perugia, Italy

⁸DESY, Hamburg, Germany

Abstract. The former CMS Run 2 High Level Trigger (HLT) farm is one of the largest contributors to CMS compute resources, providing about 25k job slots for offline computing. This CPU farm was initially employed as an opportunistic resource, exploited during inter-fill periods, in the LHC Run 2. Since then, it has become a nearly transparent extension of the CMS capacity at CERN, being located on-site at the LHC interaction point 5 (P5), where the CMS detector is installed. This resource has been configured to support the execution of critical CMS tasks, such as prompt detector data reconstruction. It can therefore be used in combination with the dedicated Tier 0 capacity at CERN, in order to process and absorb peaks in the stream of data coming from the CMS detector. The initial configuration for this resource, based on statically configured VMs, provided the required level of functionality. However, regular operations of this cluster revealed certain limitations compared to the resource provisioning and use model employed in the case of WLCG sites. A new configuration, based on a vacuum-like model, has been implemented for this resource in order to solve the detected shortcomings. This paper reports about this redeployment work on the permanent cloud for an enhanced support to CMS offline computing, comparing the former and new models' respective functionalities, along with the commissioning effort for the new setup.

1 Introduction

The High-Level Trigger (HLT) system within the CMS experiment [1] at CERN played a pivotal role in the LHC Run 2 by assisting in the selection of collision events, thereby advancing our comprehension of particle physics. As we transition from Run 2 to Run 3, the CMS experiment is adapting its resources for more efficient data processing.

*e-mail: marco.mascheroni@cern.ch

**e-mail: aperez@pic.es

The CMS Online Cloud (deployed on the HLT farm resources) was commissioned during Run 2 to dedicate HLT resources for offline data processing when not needed to support online data taking [2]. Offline computing takes place in the Global pool [3], one of the federated HTCondor [4] pools managed by the Submission Infrastructure team in CMS Computing. This set of federated pools are created using GlideinWMS [5] and collectively form a distributed and dynamic computing environment, capable of efficiently processing the ever-expanding volume of data generated by the CMS experiment at CERN, as shown in Figure 1.

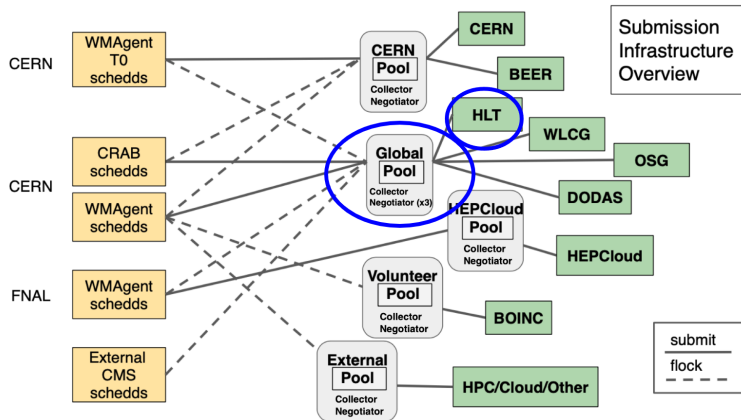


Figure 1. CMS SI current configuration, including multiple federated pools of resources allocated from diverse origins (green boxes) and sets of distributed job schedulers (schedds), handling the Tier 0 and centralized production workloads (WMAgent), as well as analysis job submission (CRAB).

OpenStack virtual machines (VMs) equipped with HTCondor Startd daemons are initiated to utilize available CPU resources during both Interfill and Fill modes, see Figure 2. When the High-Level Trigger (HLT) requires resources, these VMs are gracefully suspended. The *MaxHibernationTime* HTCondor knob is used to allow a delay of 24 hours at both the Central Control Board (CCB) and Workload Management Agents (WMAgents) schedds. This delay ensures that suspended jobs can be seamlessly resumed within the CMS Global pool when the VMs are reintegrated, a process triggered by diminished online processing demands

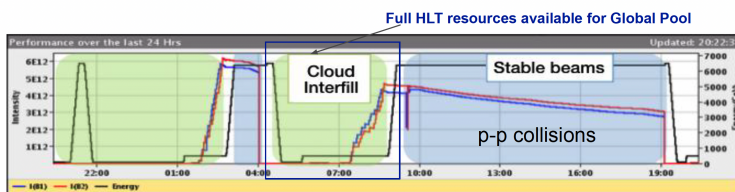


Figure 2. LHC operation cycles and interfill operations of the HLT farm for CMS offline computing tasks.

Figure 3 shows the usage of the HLT farm during Run 2. The farm was intensively used as opportunistic resource for offline processing.

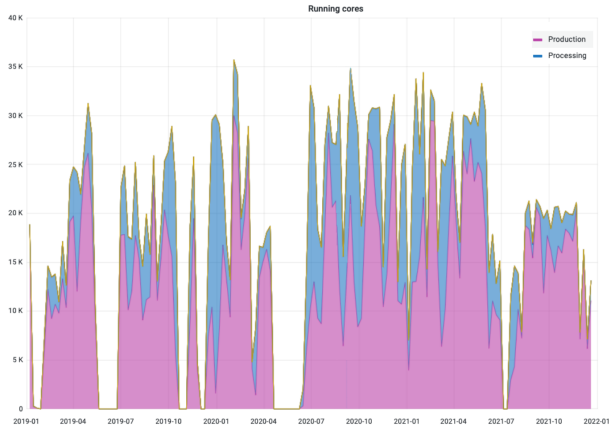


Figure 3. Use of HLT farm resources in support of CMS offline computing tasks during the LS2 years.

2 The HLT farm Hardware status for Run 3

As Run 3 approached, a new computing farm was commissioned for online usage. Currently, CMS benefits from the availability of two cloud-based High-Level Trigger (HLT) farms:

- Run 2 HLT farm (aka Permanent cloud): This farm, boasting approximately 25,000 CPU cores, has been fully repurposed to support offline computing. It operates under the new CMS site name, T2_CH_CERN_P5. While the hardware is out of warranty, it remains within the constraints of the current Data Acquisition (DAQ) and HLT power budget at the data center, with ample space available. As such, it is slated for use until the conclusion of Run 3, though server repairs are subject to best-effort and resources may diminish over time. The decommissioning of this farm is scheduled for the end of Run 3.
- This opportunistic cloud comprises approximately 33,000 CPU cores and will be utilized for offline computing during periods when the LHC is not actively taking data. The CMS site name T2_CH_CERN_HLT has been reserved for this new HLT farm. Starting from Q1 2022, the new HLT nodes in this farm are equipped with dual NVidia T4 GPUs. These resources are expected to remain in use for trigger purposes until the conclusion of Run 3. Subsequently, they are projected to transition into a permanent cloud resource dedicated to offline computing.

3 Original configuration of Run 2 HLT farm

In the old configuration the HLT resources were dynamically repurposed for offline usage through the utilization of *cloud auto* and *HLTd* daemons, as illustrated in Figure 4. Virtual machine (VM) images are crafted with a static HTCondor configuration provided by the Submission Infrastructure (SI) team. These configurations enable HTCondor Startd daemons to maintain persistent connections with the CMS pool, allowing them to operate for extended durations, often spanning weeks, in stark contrast to the typical 48-hour lifespan of regular glideins. Any adjustments, optimizations, or modifications to these configurations necessitate coordination with the SI team and the creation of updated VM images.

However, there are some notable issues and limitations associated with this approach. The prolonged existence of HTCondor Startd daemons raises concerns regarding configuration

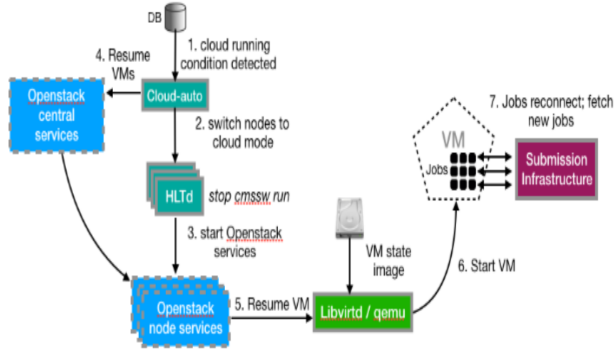


Figure 4. Schematic view of the new HLT farm deployment model, based on vacuum-like pilots.

synchronization with the broader SI setup. Instances of outdated singularity wrapper scripts pointing to obsolete or missing EL8 images have been observed. Additionally, there is a lack of resource validation prior to these daemons connecting with the pool, resulting in a notable number of job failures attributable to cvmfs [6] issues, for example.

Furthermore, the long-lived HTCondor slots inadvertently contribute to resource allocation inefficiencies, particularly in the context of slot defragmentation. The absence of slot rotation with long-lived slots means that high-priority jobs, such as 8-core tasks like Tier 0 PromptReco, often remain in the queue for extended periods. To address this, an induced defragmentation mechanisms, facilitated by the HTCondor defrag daemon, was introduced and fine-tuned. Such measures are unnecessary for pilots with approximately 48-hour lifespans.

Lastly, it's important to note that the SI team lacks direct control over the tuning of HLT slot parameters in this approach. Any interventions or adjustments in this regard necessitate coordination with the DAQ team.

4 Glideins in a Vacuum

To address the limitations of the old approach, the *Glideins in a Vacuum* method was introduced. This approach enables the direct launching of Glideins from virtual machines (VMs) without the need for any intermediate compute elements or batch systems to manage pilot jobs onto worker nodes. This concept has garnered success within the WLCG context [7]. The configuration of these Glideins is aligned with the standard Global Pool Glideins, resulting in several integrated advantages. Notably, the relatively short-lived pilots provide a continuous and natural renewal of defragmented slots, effectively mitigating resource starvation issues (controlled by `GLIDEIN_Max_Walltime`). Furthermore, these Glideins undergo resource validation before connecting to the pool, contributing to a reduction in operational costs and increased control. Moreover, new configurations are automatically applied to pilots upon Frontend (FE) or factory reconfiguration.

To better reflect the dedicated nature of the old resources for offline computing, and to accommodate the introduction of Glideins in a vacuum, a new site name/endpoint, 'T2_CH_CERN_P5,' was established. A new *Glidein-launcher.service* script has been developed and has seamlessly been integrated into VMs through a contextualization script without necessitating modifications to the original VM image. The Glidein-launcher service, a systemd-based service, retrieves the wrapper script from the GWMS FE whenever

it undergoes updates, subsequently initiating the launch of Glideins. The Submission Infrastructure (SI) team generates and publishes the wrapper for 'T2_CH_CERN_P5' using a post-reconfigure hook of GlideinWMS Frontend.

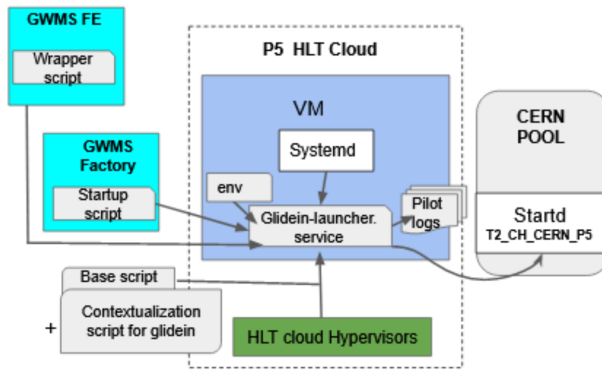


Figure 5. Schematic view of the new HLT farm deployment model, based on vacuum-like pilots.

The transition to the new Glideins-based site started in January 2023, with a complete switch to 'T2_CH_CERN_P5' in February as shown in figure 6.

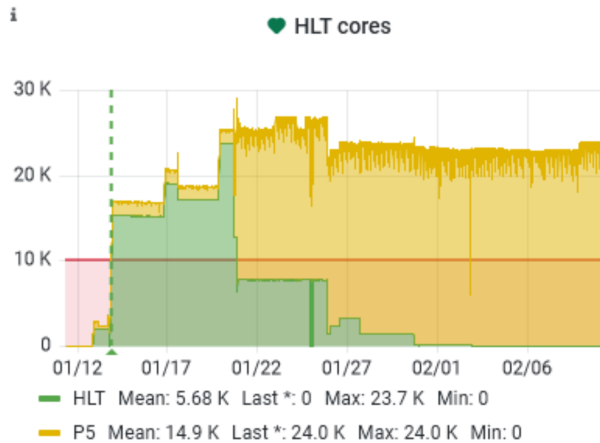


Figure 6. Transition period for the deployment model of the HLT resources, based on vacuum-like pilots.

This transition empowers the SI team to swiftly configure, tweak, and optimize slots without the need for coordination with the DAQ team or the complexities of VM image updates. Examples of optimizations include parameters such as `GLIDEIN_Max_Idle` (set to 3600, or 1 hour), `GLIDEIN_Max_Walltime` (approximately 48 hours), `GLIDEIN_CPUS` (configurable for whole node pilots), and `MaxHibernateTime` (set to 86400). Additionally, the system now allows only Production and Tier 0 jobs, a control facilitated through GlideinWMS Frontend. It introduces conditions like:

```

ifthenelse(DESIRED_Sites is not undefined,
stringListMember(GLIDEIN_CMSSite, DESIRED_Sites), undefined)

and

(GLIDEIN_CMSSite != 'T2_CH_CERN_P5' ||
WMAgent_AgentName != UNDEFINED)

```

to further refine job assignments and validations, thereby reducing the likelihood of job failures. Lower idle times for high-priority tasks has also been observed.

With the migration of 'T2_CH_CERN_P5' from the Global to the CERN Pool, several advantages have emerged. This transition brings about a level of similarity with 'T2_CH_CERN' in terms of resource shares and priorities, encompassing Tier 0, production, and analysis jobs (excluding analysis jobs in P5). Furthermore, it effectively isolates 'T2_CH_CERN_P5' from potential issues that may impact the global pool, enhancing resilience and resource stability for tasks primarily associated with Tier 0 operations. This migration also alleviates the pressure on the global pool collector and the global pool T2 negotiator, contributing to the overall robustness and efficiency of the system.

5 Commissioning Tier 0 jobs in the reconfigured HLT farm in Run 3

In the initial year of LHC Run 3 (2022), the Run 2 HLT farm found new opportunities for utilization. This time, it extended its support to Tier 0 operations, specifically in the realm of prompt reconstruction tasks. To ensure a seamless transition, a commissioning phase was initiated in the spring of 2022, aimed at enabling Tier 0 tasks to efficiently access the established HLT farm and secure slots at a rate conducive to their needs, see Figure 7.

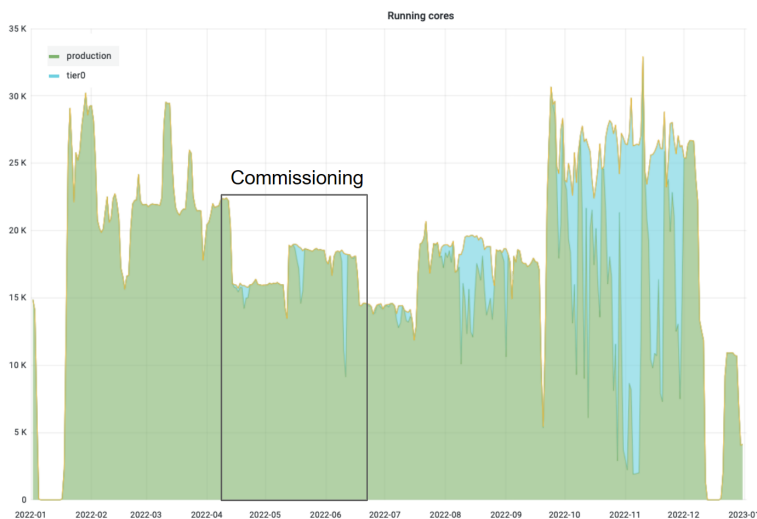


Figure 7. Transition period for the deployment model of the HLT resources, based on vacuum-like pilots.

During this period, a notable challenge emerged: partitionable slots within long-lived virtual machines (VMs) often experienced fragmentation due to low-core-count production

jobs. To address this issue and facilitate the execution of critical PromptReco tasks, which require 8-core slots, an induced defragmentation mechanism for the HLT slots was configured and optimized.

As soon as stable beams were declared, the resource seamlessly extended the Tier 0 capacity, responding to the demands imposed by the data acquisition rates. This adaptive use of the Run 2 HLT farm played a crucial role in supporting the continuous flow of data during LHC Run 3.

In 2023, Tier 0 jobs are efficiently processed on T2_CH_CERN_P5 following the successful transition to the new vacuum model. In 2022, the inclusion of the Run 2 HLT in the resource pool significantly bolstered the capacity for executing Tier 0 prompt reconstruction tasks. The commissioning of the new site name and glidein mechanism, specifically tailored to accommodate Tier 0 jobs, has undergone rigorous testing and validation in anticipation of the 2023 LHC data-taking season. With this new approach it was shown that the allocation of CPU cores for Tier 0, characterized by high-priority jobs, was really fast as shown in Figure 8.

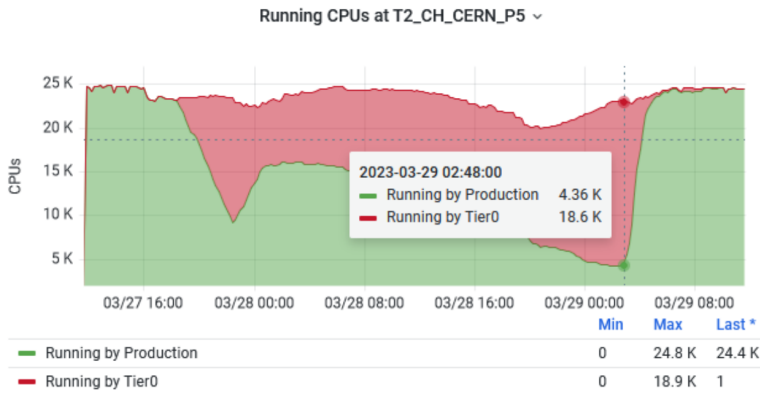


Figure 8. Transition period for the deployment model of the HLT resources, based on vacuum-like pilots.

6 Conclusions and future work

This paper has detailed the successful repurposing of the Run 2 CMS HLT infrastructure into a dedicated farm designed specifically for offline computing tasks. This transformation became feasible due to the deployment of a new HLT farm tailored for Run 3. We have explored the implementation of the Vacuum Model Glideins, an approach that provides enhanced flexibility for offline computing.

Our transition to a Glidein-based acquisition and management approach for P5 resources has yielded several identified advantages, aligning with our initial motivations for this shift. Since its full implementation in February, this approach has performed as anticipated, and we have already realized benefits such as automated pilot reconfiguration. We are now planning to extend this pilot-based methodology to the new Run 3 HLT farm, configuring Glideins with the necessary parameters to accommodate the intermittent nature of these resources. We will continue to utilize the reserved sitename 'T2_CH_CERN_HLT' for this purpose.

Presently, the opportunistic use of the new farm does not encompass GPU nodes due to pending configurations for GPU virtualization.

Additionally, we have examined the potential of leveraging the new HLT infrastructure in both Interfill and Fill modes to meet diverse offline usage requirements in the future. The Run 3 HLT farm will be opportunistically utilized during specific periods, notably during LHC technical stops and in what we refer to as the 'Interfill mode'. This mode is active during brief intervals, typically spanning only a few hours, when the LHC temporarily ceases the generation of physics data. The prerequisites for enabling Interfill mode for offline computing are already in place, including the capabilities for VM suspension and disk state preservation on the DAQ side. Similarly, the SI team has implemented a job suspension and resumption feature to facilitate Interfill mode utilization.

As part of ongoing efforts, a subset of nodes from the Run 3 HLT farm is designated for testing and seamless integration into the broader CMS offline resource pool. This approach ensures the smooth incorporation of these resources into our computational infrastructure, further enhancing our capacity for efficient offline data processing.

Acknowledgements

This work was partially supported by the Spanish Ministry of Science and Innovation under grants PID2019-110942RB-C21, PID2019-110942RB-C22 and PID2020-113807RA-I00, which include FEDER funds from the European Union, and by the US National Science Foundation under Grant No. 2121686.

References

- [1] CMS Collaboration, The CMS experiment at the CERN LHC, JINST 3 (2008) S08004, doi:10.1088/1748-0221/3/08/S08004.
- [2] J.-M. Andre et al, "Experience with dynamic resource provisioning of the CMS online cluster using a cloud overlay", EPJ Web of Conferences 214, 07017 (2019)
- [3] J. Balcas et al. "Using the glideinWMS System as a Common Resource Provisioning Layer in CMS", J. Phys.: Conf. Ser. **664** 062031 (2015).
- [4] The HTCondor Software Suite public web site, <https://research.cs.wisc.edu/htcondor/index.html>
- [5] The Glidein-based Workflow Management System, <https://glideinwms.fnal.gov/doc.prd/index.html>
- [6] D. Weitzel et al, "Accessing Data Federations with CVMFS", EPJ Web of Conferences 898, 062044 (2017)
- [7] A McNab et al, "Running Jobs in the Vacuum", J. Phys.: Conf. Ser. 513 032065 (2014).