EPJ Web of Conferences **337**, 01214 (2025) https://doi.org/10.1051/epjconf/202533701214
CHEP 2024

The Cherenkov Telescope Array Observatory Production System Status and Development

Natthan Pigoux[1,*] and *Luisa* Arrabito[1,**]

[1]Laboratoire Univers et Particules de Montpellier, CNRS/IN2P3, Montpellier, France

Abstract. The Cherenkov Telescope Array Observatory (CTAO) is the next-generation instrument in the domain of very-high-energy gamma-ray astronomy. It will consist of tens of Cherenkov telescopes deployed across two CTAO array sites at La Palma (Spain) and Paranal (ESO, Chile). Currently under construction, CTAO will begin operations in the coming years for approximately 30 years. During operations, CTAO is expected to generate about 2 PB of raw data annually, in addition to 5-20 PB of Monte Carlo simulation data. Processing such a large volume of data will require hundreds of millions of CPU HS06 hours per year. To meet these high computing and storage requirements, CTAO will rely on resources distributed across four Data Centers. To manage simulation and data processing workflows in such a distributed environment, we have developed a production system prototype based on the DIRAC interware. In this paper, we present the current status of this prototype, the underlying infrastructure, the technologies in use, as well as recent developments regarding user interfaces for workflow description and automatic recovery of incomplete workflows, along with future perspectives.

1 Introduction

CTAO [1, 2] processes telescope data and performs simulations by executing a series of pipeline steps organised in complex workflows. Each pipeline step consists of multiple computing tasks applied to a specific dataset. The goal of the CTAO production system is to efficiently manage the execution of such complex workflows across the four Data Centers.

In the CTAO software architecture, low-level data processing is managed by the so-called *Data Processing and Preservation System* (DPPS) [3], responsible for processing telescope raw data and generating intermediate-level scientific products. DPPS consists of several subsystems: *Pipelines*, i.e. software applications that form data processing workflows, such as Calibration, Data Processing, Simulation and Data Quality pipelines; the *Workload Management System* (WMS) responsible for managing workflow execution on the computing infrastructure and the *Bulk Data Management System* (BDMS). The *Assembly Integration and Verification* (AIV) system integrates all subsystems into a common software environment, handling integration tests, software quality assurance and the deployment procedures for the entire DPPS. The DPPS software will be developed incrementally through releases, each incorporating new functionalities to support additional use cases.

*e-mail: natthan.pigoux@lupm.in2p3.fr
**e-mail: arrabito@in2p3.fr

In this software architecture, the WMS subsystem corresponds to what we refer to as the CTAO production system. The prototype of the CTAO production system presented in this paper is based on the DIRAC interware [4] and is implemented as a software extension, CTA-DIRAC. We have operated this prototype for over 10 years, running Monte Carlo simulations during the design phase of CTAO. Since there were no dedicated computing resources available to CTAO during this period, we configured the prototype to use opportunistic resources from about 10 EGI grid sites. Over the past two years, we have transitioned to a model that uses exclusively the pledged resources of the four CTAO Data Centers.

In section 2, we describe the server infrastructure that supports the production system, as well as the testbed infrastructure used for release certification before deployment. Section 3 discusses the functionalities of the production system, focusing on key features, recent and future developments and deployment strategy. Finally, in section 4, we highlight the use of the production system prototype for recent Monte Carlo simulations campaigns across the four CTAO Data Centers.

2 CTAO Computing Infrastructure

CTAO has adopted a distributed computing model for simulations and low-level data processing. The entire computational load and storage will be evenly distributed among four Data Centers: CSCS (Lugano, Switzerland), DESY-Zeuthen (Zeuthen, Germany), INAF/INFN (Frascati, Italy) and PIC (Barcelona, Spain). These centers, or a subset of them, will also host the production deployment of DPPS and in particular WMS and BDMS subsystems.

For the Workload and the Bulk Data Management Systems of DPPS, CTAO relies on established and reliable softwares from the High Energy community, adding higher-level functionalities or contributing directly to them. DIRAC interware has been selected for WMS, while Rucio and FTS (File Transfer System) serve as ground softwares for BDMS.

The ecosystem of underlying technologies is based on standards for distributed computing: CernVM-FS (Cern Virtual Machine File System) as the distributed file system, Indigo IAM as the identity provider, OpenSearch for logs and monitoring (coupled with Kibana/Graphana for visualization) and Kubernetes for deploying and running DPPS releases.

2.1 CTA-DIRAC Legacy Prototype

We have developed an initial production system prototype, CTA-DIRAC Legacy, based on the DIRAC interware. This prototype was successfully used throughout the CTAO preparation phase to run large-scale Monte Carlo simulations (*prod2* for site selection [5] and *prod3* for the array layout optimization [6]) on the EGI grid [7]. Currently we continue to operate this prototype to run simulations for the calculation of instrument response functions (IRFs), utilizing the resources of the four CTAO Data Centers. A detailed description of this prototype was provided in a previous report [8]. CTA-DIRAC services and databases were initially deployed at CC-IN2P3 and PIC data centers. We distribute the DPPS *pipelines* softwares via CernVM-FS, with the CTAO repository hosted on Stratum servers located at DESY-Zeuthen and CC-IN2P3. Following the selection of the four CTAO Data Centers, in early 2024, we have migrated all CTA-DIRAC services and databases previously hosted at CC-IN2P3 to CSCS, along with the OpenSearch service, which is now hosted at DESY-Zeuthen.

Finally, we are currently updating CTA-DIRAC services to transition from VOMS to Indigo IAM as the identity provider, using the CTAO IAM instance provided by CNAF/INFN. We already use IAM tokens to submit pilots to Computing Elements and have also successfully tested the use of IAM proxies to access Storage Elements.

2.2 CTAO DPPS Test Cluster

In addition to the production deployment of DPPS subsystems at the four Data Centers, DESY-Zeuthen provides a testbed infrastructure for DPPS developers. This infrastructure offers resources for testing subsystems, their deployment and their integration before producing DPPS releases (see section 3.3). The testbed consists of VMs, a Kubernetes cluster based on Rancher, an OpenSearch and Prometheus servers for monitoring. On this infrastructure we currently run the CTA-DIRAC certification instance based on DIRAC v9 and the DiracX 0.1.0 pre-release running on the Rancher cluster.

3 CTAO Production System

3.1 DIRAC

DIRAC is an interware for managing computing tasks and data across distributed and heterogeneous e-infrastructures. Initially developed by the LHCb collaboration, DIRAC has been providing open-source core functionalities since 2008 [9] that can be extended by other scientific communities. Although DIRAC encompasses both data and workload/workflow management functionalities, DPPS only utilizes the workload and workflow management features, while for data management it relies on Rucio. Below is a brief summary of the key DIRAC components for workload and workflow management:

- Workload Management System (WMS): Manages the entire job lifecycle, including submission, scheduling, execution, status monitoring, and output retrieval. It uses the pilot model for job scheduling and provides plugins for targeting different resource types (e.g., grids, clouds, clusters, HPC data centers). Jobs are described using a custom Job Description Language (JDL).
- Transformation System (TS): Generates and submits jobs based on a transformation description, which includes a job template and input dataset (the result of a query on user-defined metadata). Each transformation consists of multiple identical jobs, each processing a subset of the input dataset. Transformations are linked together to form a complete workflow (referred to as a "production" in DIRAC terminology).
- Production System (PS): Based on a production description, the PS generates and submits transformations to the DIRAC TS.

The DIRAC consortium is currently developing "The neXt Dirac incarnation", DiracX, to modernise the software with today's standards (FastAPI, AutoREST, ...). In particular, DiracX will utilize Helm charts for deployment, and workflows will be described using the Common Workflow Language (CWL) [10].

3.2 CTA-DIRAC

DIRAC can be extended to meet the specific needs of user communities. For the needs of CTAO, we have developed CTA-DIRAC, a software extension deployed alongside DIRAC.

One important requirement for CTAO is the ability to execute complex workflows in a fully automated and reproducible manner. To enhance reproducibility and simplify workflows execution across different environments, DPPS workflows will ultimately be described using CWL.

Workflow management in DIRAC is based on the Transformation and Production Systems, described in Section 3.1. The DIRAC Production System was initially developed in the context of CTA-DIRAC before being integrated into the DIRAC core in 2020. In the

CTA-DIRAC extension, we have developed a high-level interface for configuring and submitting specific CTAO workflows [11], based on these two systems. This interface allows users to configure and submit CTAO workflows using a custom YAML description, which is then converted into a DIRAC-compliant format via the DIRAC API. Preliminary support for workflows described in CWL has been developed for a subset of CTAO workflows.

As workflows in a distributed environment rarely complete without manual intervention, we aim to automate failure recovery during workflow execution. To this end, we have recently developed an initial version of a Failover agent in CTA-DIRAC.

This agent monitors the status of jobs and input files within active transformations, marking completed transformations and allowing child transformations to begin. It also ensures that workflows progress, addressing situations where job creation is blocked due to the fact that the number of remaining files to process is lower than the number of input files expected by jobs. The agent triggers job creation for input files associated with failed jobs and generates detailed reports for failures that cannot be automatically resolved.

The Failover agent has been successfully used in the recent simulation campaigns (*prod5b* and *prod6*) and has significantly reduced operational time. We plan to extend it to handle a broader range of errors, including differentiating between pipeline software failures and infrastructure failures. Additionally, the core mechanics of this agent can be shared with other communities using DIRAC. As a result, it can be integrated into the DIRAC core, while CTAO-specific features will remain within the CTA-DIRAC extension.

3.3 Plans on Development and Deployment

A preliminary DPPS release plan has been outlined (see figure 1). As we move forward with releases, we will implement use cases of increasing complexity, from a single job execution to complete workflows producing instrument response functions (IRF).

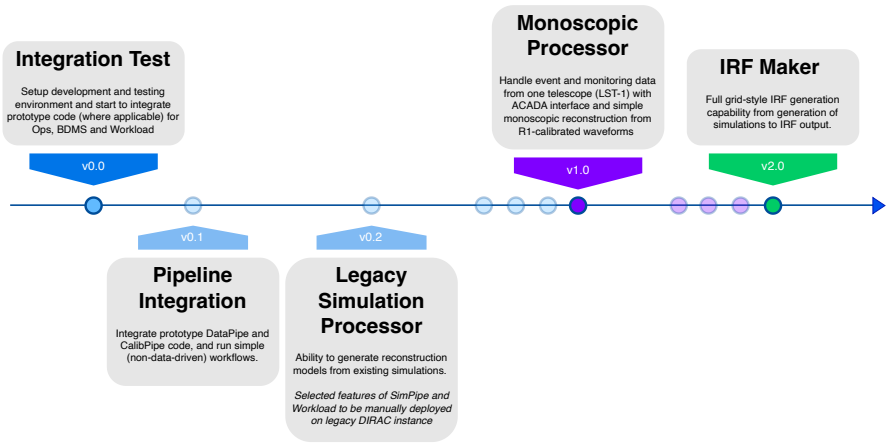

Figure 1. Preliminary overview of the DPPS release plan for the coming years. The first release (v0.0) is scheduled for the end of February 2025.
Source: Karl Kosack, DPPS release plan document (02/2025).

The objective of the upcoming release v0.0 is to finalize the assembly, integration and verification procedures of DPPS. It will include only the WMS and BDMS subsystems, with

other subsystems integrated in future releases. These two subsystems will be deployed on the DPPS testbed infrastructure on a Kubernetes cluster and in a Gitlab CI pipeline. The focus of Release 0.0 will be on testing integration between WMS and BDMS subsystems, by supporting the use case of uploading and retrieving a file from BDMS using WMS. To test the different use cases, each concerned subsystem provides Helm charts, Docker images and tests for the related use cases, which are then integrated in the DPPS release. The next major DPPS release (v1.0) will also incorporate Data Processing and Calibration pipelines. Its primary goal is to enable data processing from the first Large-Sized Telescope (LST-1), deployed in La Palma (Spain) which is already in commissioning.

Moreover, as we have seen in section 3.1, the DiracX deployment strategy is completely aligned with WMS subsystem requirements. It is a great opportunity for us to get involved in the design and development of new DIRAC features. In particular, the Kubernetes deployment and the introduction of CWL for workflows description, as are both part of the WMS requirements.

4 Operations for Monte Carlo Simulations

Since 2012, large-scale simulation campaigns for CTAO have been conducted using the CTA-DIRAC Legacy prototype, leveraging EGI grid resources. Simulation and processing workflows are submitted through the CTA-DIRAC workflow interface presented in section 3.2. Since operations began, CPU usage has accumulated to approximately 1 billion HS06 CPU hours. The generated simulation data are distributed across seven data centers, totaling more than 7 PB.

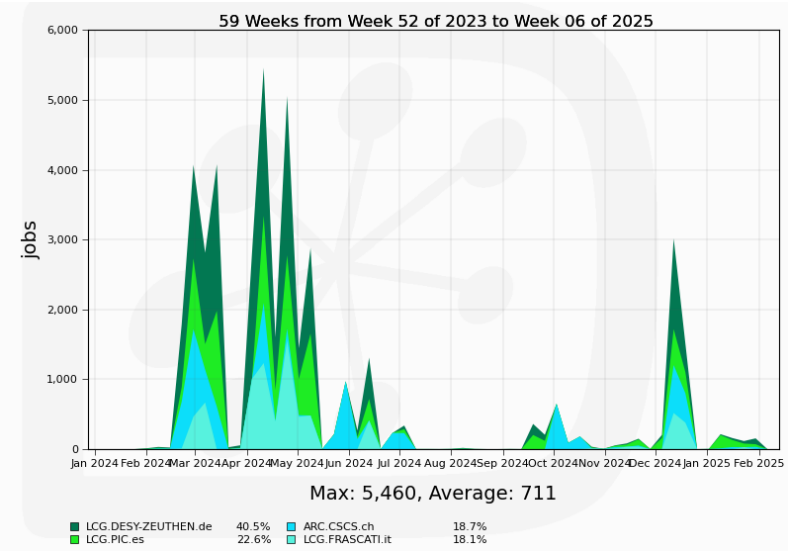

Figure 2. Number of concurrent running jobs by data center, between 01-2024 and 02-2025, using the CTA-DIRAC prototype *legacy* instance. Representing the *prod6* simulations and processing campaign.

From early 2024 to early 2025, we carried out the latest large-scale Monte Carlo campaign (*prod6*), relying exclusively on the four dedicated Data Centers listed in section 2. The CPU load was fairly evenly distributed across the four sites, and we were able to efficiently use the

available slots over several weeks. At peak usage, we reached approximately 5500 concurrent running jobs (figure 2).

At the time of these operations, the first DPPS release, integrating WMS and BDMS, was not yet ready. As a result, we extensively used DIRAC also for data management functionalities, including data transfers, file registration, and metadata queries. For the *prod6* campaign, we produced and permanently stored 1.2 PB of data, distributed across three Data Centers (CSCS, DESY-Zeuthen and PIC). Additionally, simulation and processing jobs collectively transferred approximately 6.9 PB of data.

5 Conclusion and Perspectives

The CTAO production system, based on the DIRAC interware framework, has demonstrated its ability to manage distributed computing resources and handle complex workflows. In the CTA-DIRAC extension, we have implemented a workflow interface to configure and execute CTAO workflows on top of the DIRAC Production System. The system's continued development, including the incorporation of new tools for automation, monitoring and workflow recovery, will be crucial for supporting CTAO's ambitious scientific goals.

We have operated this prototype since 2012 for large-scale simulation campaigns. These productions required hundreds of millions of HS06 CPU hours per year and the management of Petabytes of data, now exclusively using the four CTAO Data Centers.

For the upcoming DPPS release, we will be able to deploy the WMS alongside the BDMS subsystem on Kubernetes. This aims to carry out integration tests between the two subsystems. Further releases will support more use cases and integrate the different DPPS pipelines software to operate the full workflow for IRFs production and telescope data processing.

Future directions and design choices of the DiracX project are well aligned with CTAO requirements. We will thus continue our commitment in the development of the new incarnation of the software.

References

[1] R. Zanin et al. (CTA Observatory, CTA Consortium, LST), CTA – the World's largest ground-based gamma-ray observatory, PoS **ICRC2021**, 005 (2022). 10.22323/1.395.0005

[2] Science with the Cherenkov Telescope Array (WORLD SCIENTIFIC, 2018), ISBN 9789813270091, http://dx.doi.org/10.1142/10986

[3] CTAO Technical Specifications (2025), www.ctao.org/for-scientists/technical-specifications/

[4] Stagni, Federico, Boyer, Alexandre, Tsaregorodtsev, Andrei, Lytovchenko, Andrii, Sailer, André, Haen, Christophe, Burr, Christopher, Bauer, Daniela, Fayer, Simon, Martyniak, Janusz et al., Dirac current, upcoming and planned capabilities and technologies, EPJ Web of Conf. **295**, 04018 (2024). 10.1051/epjconf/202429504018

[5] T. Hassan, L. Arrabito, K. Bernlöhr, J. Bregeon, J. Cortina, P. Cumani, F. Di Pierro, D. Falceta-Goncalves, R. Lang, J. Hinton et al., Monte carlo performance studies for the site selection of the cherenkov telescope array, Astroparticle Physics **93**, 76 (2017). https://doi.org/10.1016/j.astropartphys.2017.05.001

[6] A. Acharyya, I. Agudo, E. Angüner, R. Alfaro, J. Alfaro, C. Alispach, R. Aloisio, R. Alves Batista, J.P. Amans, L. Amati et al., Monte carlo studies for the optimisation of the cherenkov telescope array layout, Astroparticle Physics **111**, 35 (2019). https://doi.org/10.1016/j.astropartphys.2019.04.001

[7] EGI Use Cases (2022), https://cdn.egi.eu/app/uploads/2022/04/EGI_Use_Cases.pdf

[8] Arrabito, Luisa, Bregeon, Johan, Maeght, Patrick, Sanguillon, Michèle, for the CTA Consortium, Tsaregorodtsev, Andrei, for the DIRAC Consortium, The cherenkov telescope array production system proto-type for large-scale data processing and simulations, EPJ Web Conf. **251**, 02029 (2021). 10.1051/epjconf/202125102029

[9] A. Tsaregorodtsev, M. Bargiotti, N. Brook, A.C. Ramo, G. Castellani, P. Charpentier, C. Cioffi, J. Closier, R.G. Diaz, G. Kuznetsov et al., Dirac: a community grid solution, Journal of Physics: Conference Series **119**, 062048 (2008). 10.1088/1742-6596/119/6/062048

[10] M.R. Crusoe, S. Abeln, A. Iosup, P. Amstutz, J. Chilton, N. Tijanić, H. Ménager, S. Soiland-Reyes, B. Gavrilović, C. Goble et al., Methods included: standardizing computational reuse and portability with the common workflow language, Commun. ACM **65**, 54–63 (2022). 10.1145/3486897

[11] L. Arrabito, J. Bregeon, A. Faure, O. Gueta, N. Pigoux, A. Tsaregorodtsev, The Cherenkov Telescope Array Observatory workflow management system, EPJ Web Conf. **295**, 04044 (2024). 10.1051/epjconf/202429504044