CRIC as a core instrument for WLCG operations

Julia Andreeva[3], *Alexey* Anisenkov[1,2], *Alessandro* Di Girolamo[3], *Panos* Paparrigopoulos[3], and *Natalia Diana* Szczepanek[3]

[1]Budker Institute of Nuclear Physics, Novosibirsk, Russia
[2]Novosibirsk State University, Novosibirsk, Russia
[3]CERN, Geneva, Switzerland

Abstract. WLCG is a large heterogeneous computing infrastructure which provides resources for the LHC experiments. The Computing Resource Information Catalogue (CRIC) has been designed as the source for the WLCG topology information. CRIC aims to describe WLCG distributed sites, their services and how the provided resources are used by the LHC experiments. The CRIC instance dedicated to WLCG has become a key service for WLCG computing operations, playing an important role in monitoring, accounting and organisation of the large-scale data challenges and upgrade campaigns. The contribution describes CRIC functionality widely used for the WLCG central operations as well as development plans helping to decrease operational effort.

1 Introduction

The Worldwide LHC Computing Grid (WLCG) [1] is a global computing infrastructure whose mission is to provide computing resources to store, distribute and process data of the LHC experiments. WLCG represents a global collaboration of more than 170 computing centres in 42 countries, linking up national and international grid infrastructures. The scale of WLCG and its complexity are steadily growing following the needs of the LHC experiments and technology evolution. Computing teams of the LHC experiments, WLCG operations coordination team and support teams at the WLCG sites are facing the challenge of operating reliably and efficiently in such a large-scale steadily evolving heterogeneous infrastructure. Among the most important operational tasks are testing and monitoring the provided resources, upgrading and deployment of the distributed services and detecting, debugging and resolving the eventual problems. All those tasks require information about WLCG sites and services of all types (computing, storage and network) including a very detailed description of service configuration and capabilities. The Computing Resource Information Catalogue (CRIC) is a core service used by WLCG operations to provide required data. Initially, it has been designed as a source for the WLCG topology information. With time the scope of CRIC grew beyond just topology description.

2 LHC requirements and yearly pledges

WLCG sites are organized within a multi-tiered structure, where the tier level of a specific site can vary for different LHC experiments, depending on the set of activities performed

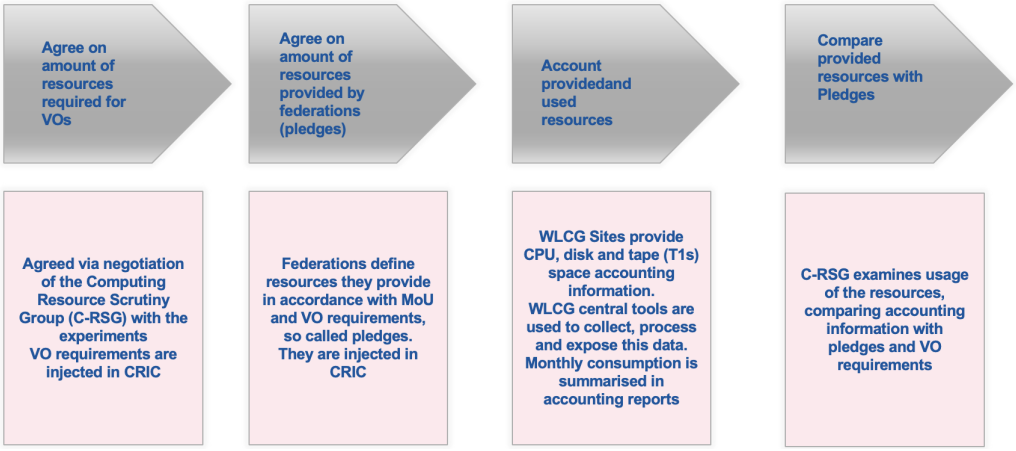

Figure 1. Workflow for defining LHC requirements and pledges

at the site and the criticality of particular services. Consequently, the topology description is specific to each LHC experiment. CRIC not only provides a generic WLCG topology description but also offers a particular experiment-specific topology view.

The WLCG Memorandum of Understanding (MoU) outlines how sites contribute to the WLCG infrastructure, including specifying the minimum levels of service availabilities at Tier-0/1/2 sites and the agreed-upon fraction of storage and compute resources provided by T0, T1, and T2 to meet the resource requirements of the LHC Experiments. The allocation of resources is determined at the WLCG federation level, where federations may comprise one or several computing sites and establish unified technical and operational management to coordinate with the broader WLCG Collaboration.

The requirements of the LHC experiments are negotiated annually through discussions between the experiments and the Computing Resource Scrutiny Group (C-RSG) [2]. Subsequently, federations define the amount of resources they intend to provide based on the requests of the LHC experiments. WLCG accounting procedures, as described in Section 3, enable tracking the fulfillment of federation obligations and the usage of provided resources by the LHC experiments. The entire workflow is depicted in Figure 1.

CRIC maintains VO (Virtual Organization) requirements and federation pledges, while also offering a user interface for inputting, editing, and viewing this information.

3 Accounting

3.1 Topology source for accounting

In order to assess whether sites fulfill their obligations, the storage and computing resources provided by the sites and consumed by the LHC experiments are accounted for. Accounting metrics are used to generate monthly accounting reports for every federation and for the entire WLCG infrastructure.

For CPU accounting, WLCG relies on the APEL system [3] and the accounting portal of the European Grid Infrastructure (EGI) [4]. EGI provides e-science access to high-throughput computing resources across Europe, and the EGI portal utilizes topology information from

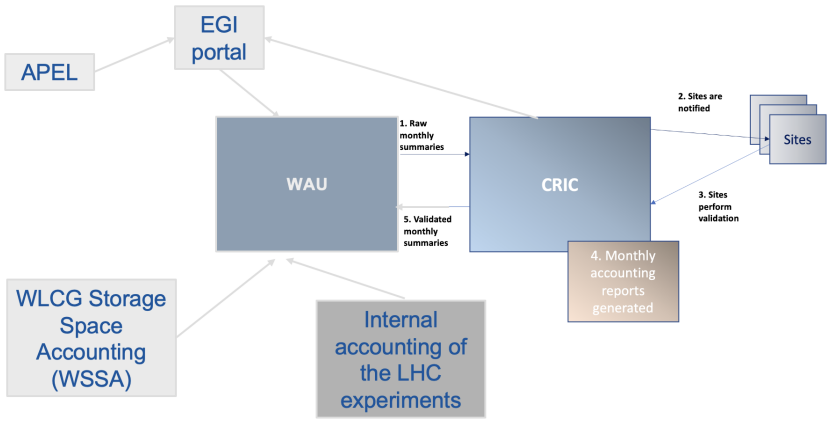

Figure 2. WLCG accounting workflow

CRIC. The CRIC topology describes which sites serve which LHC experiments, the tier level of each site, as well as the WLCG federation structure.

Regarding storage space accounting, which includes accounting for tape and disk storage, the WLCG Storage Space Accounting application (WSSA) [5] has been developed. CRIC is used by this application to store URL links to the files where sites publish their storage accounting metrics every half an hour. WSSA retrieves the set of URLs from CRIC and reads the information published by the sites. The collected data are processed and recorded in the WSSA repository. The accessibility and validity of the published files are assessed every half an hour, and the results of these checks are recorded in CRIC. The WLCG Operations Coordination team regularly reviews the validity and accessibility flags in CRIC and follows up with the sites when problems are detected.

3.2 Accounting validation and report generation

Accounting validation has been introduced to enhance the quality of accounting data. The accounting workflow is depicted in Figure 2. The WLCG Accounting Utility (WAU) imports CPU accounting data from the EGI accounting portal and storage space accounting data from WSSA. The goal is to enable a fully automatic accounting process, starting with data collection from the sites, followed by processing and generation of monthly accounting reports. However, previous experience revealed certain issues in the workflow. In some cases, APEL sensors at the sites were broken or misconfigured, resulting in incorrect accounting information being generated and published in the central accounting repository. Consequently, overall accounting metrics for the infrastructure were affected. Investigating and rectifying site problems can be time-consuming. Such incidents highlighted the need for the ability to correct faulty data promptly while diagnosing and repairing the root cause of the problem. On the other hand, not all sites are currently equipped to publish storage space accounting information, as the instrumentation campaign is ongoing. Therefore, a new functionality was implemented to allow site administrators to review the auto-generated accounting data, validate it, compare it with local accounting metrics, and make corrections if any issues are identified or if data is incomplete. As shown in Figure 2, CRIC plays a pivotal role in the validation process Raw monthly accounting metrics retrieved from the EGI portal and WSSA,

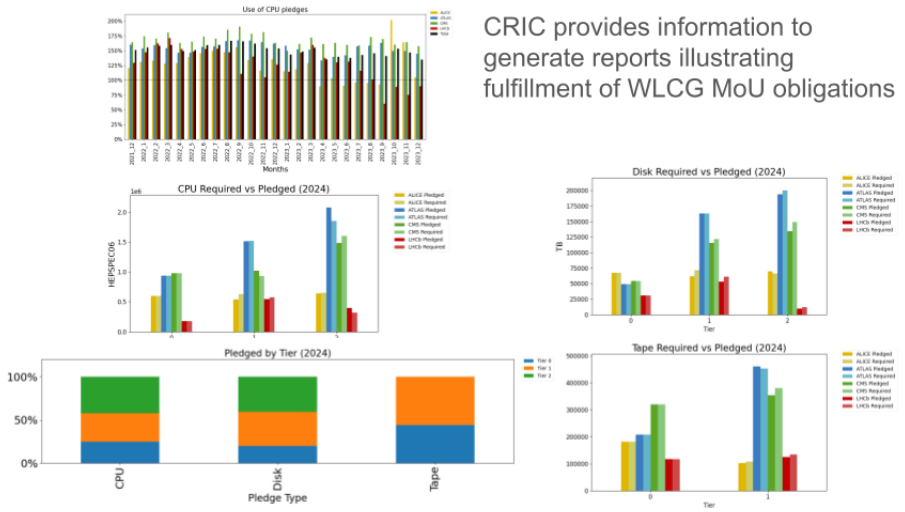

Figure 3. CRIC plots

and imported into WAU, are stored in CRIC. CRIC offers a user interface for site administrators to review and correct accounting information as necessary. Validated monthly summaries are then imported into WAU, which tracks all manually corrected metrics and provides distributions showing differences between raw and validated monthly summaries. The WLCG Operations Coordination team investigates inconsistencies and raises tickets against sites with detected discrepancies. At the end of each month, the WLCG project office generates official monthly accounting reports for the previous month, using the CRIC application for this purpose.

CRIC also provides information that allows for demonstrating the fulfillment of MoU obligations by comparing resource requirements with pledges and actual accounting data. See Figure 3 for further details. In the summer of 2023, WLCG Operations Coordination organized a review of the pledge management and accounting tools used within the WLCG. CRIC is one of these tools. The review gathered input from the experiments, sites, and the WLCG project office. The results of this review will identify potential areas for improvement and, consequently, will shape the CRIC development plans related to WLCG accounting and pledge management functionality.

4 CRIC as a topology source for WLCG testing and monitoring

The ability of CRIC to provide experiment-specific topology description is widely used by WLCG operations tools and activities. The role of CRIC as a topology provider for various WLCG operational tools is described in Figure 4.

4.1 Remote testing and service availability

To ensure the health of distributed sites and their ability to provide the required functionality, various components of the WLCG infrastructure undergo continuous testing. The WLCG

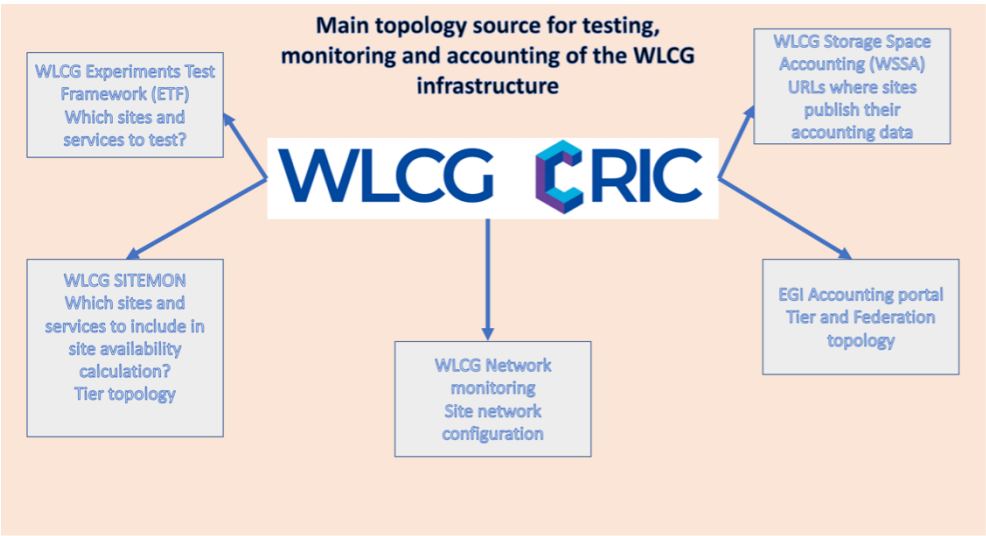

Figure 4. CRIC as a main topology source for WLCG testing, monitoring and accounting.

Experiment Test Framework (ETF) submits tests to the WLCG services, which can be either generic or experiment-specific. CRIC provides the list of sites and services targeted for ETF testing, taking into account experiment-specific topology information.

The results of ETF tests are published to the central repository of the WLCG SITEMON application [6], where they are processed to assess the level of availability delivered by all site services. The WLCG Memorandum of Understanding (MoU) defines minimum service availability levels at the Tier-0/1/2 sites. To perform availability calculations, the ETF test results are enriched with CRIC topology information. Since service and site availability calculations are performed individually for each LHC experiment, experiment-specific topology information provided by CRIC is used for these calculations."

4.2 Transfer and network monitoring

WLCG transfer monitoring applications also use CRIC as a topology source. Recently CRIC has been enabled with WLCG network topology description. This information includes: - IPv4 and IPv6 prefixes used for LHCONE and LHCOPN - Autonomous System Numbers (ASN used by the BGP routing protocol) - Available Bandwidth - Acknowledge to LHCONE Acceptance Use Policy (AUP) - Contact information of the support teams - Links to network monitoring pages of the WLCG sites

Initially CRIC database has been bootstrapped with information collected by WLCG network experts. As a next step, WLCG Operations Coordination ran a campaign requesting WLCG sites to validate collected network data and provide missing information. All this information is used to debug and solve network issues.

One of the conclusions of the first WLCG Network Data Challenge in fall of 2021 was the need to improve WLCG global network monitoring. In order to achieve this goal the WLCG Monitoring Task Force suggested to instrument sites to publish every minute their inbound and outbound traffic which can be then consumed by the central monitoring system to generate a global WLCG view. Site URL links pointing to json files containing monitoring data are stored in CRIC. Central data collector will use CRIC API to retrieve site URLs. Then it will

iterate over them and record inbound and outbound network traffic in the monitoring repository. This is ongoing work. WLCG Operations Coordination is running a global campaign on site instrumentation for network monitoring.

5 WLCG Upgrade and deployment campaign

Large-scale upgrade and deployment campaigns on the WLCG infrastructure are integral to WLCG operations. CRIC is employed in such campaigns to define the list of sites and services targeted for a particular upgrade or deployment. CRIC also enables progress tracking. Examples of recent campaigns organized by WLCG operations are:

- Enabling the publishing of storage resource reporting (SRR) files used for WLCG storage accounting.
- Validation of site network information by site network experts.
- Enabling monitoring of site inbound and outbound network traffic.
- Migration of DPM (Disk Pool Manager) storage instances to other storage solutions.

DPM has been the most widely deployed solution for storing large data repositories on Grid (or High Throughput Computing) sites for a wide range of applications. Support for DPM is scheduled to end by summer 2024. Therefore, sites were encouraged to migrate to other solutions, and a large-scale migration campaign was initiated by EGI and WLCG operations. CRIC stores information about the implementation and implementation version of the storage services provided by each site. This information is utilized in the DPM migration campaign to identify sites running storage services with DPM implementations. The migration process for a given site is considered complete when CRIC is updated with new storage implementation details.

The experience gained from using CRIC for upgrade and deployment campaigns has contributed to a better understanding of the operational requirements for a specific CRIC application designed to support such activities. This application is currently in development.

Another important missing functionality in CRIC is the generation of notifications to be sent to sites based on the results of selection using specific criteria. For example, notifying sites that have not enabled network description. This functionality has been requested by the WLCG Operations Coordination team and is included in the CRIC development tasks pipeline.

6 Summary

Over the last few years, WLCG CRIC has been effectively utilized for WLCG operations. The service has evolved in response to the requirements of the WLCG Operations Coordination team and the computing teams of the LHC experiments. The new functionality introduced in CRIC has enabled us to save operational effort, reduce the time required for operational tasks, and enhance the overall quality of the WLCG infrastructure. For example, significant improvements in WLCG accounting data would not have been possible without CRIC. Similarly, CRIC plays a pivotal role in the ongoing enhancement work on monitoring the WLCG network infrastructure. Several tasks aimed at further improving WLCG operational efficiency are included in the development plan for CRIC. Among them are site notifications and a dedicated application for massive upgrade campaigns.

References

[1] I. Bird, Computing for the Large Hadron Collider, *Annual Review of Nuclear and Particle Science*, **61**, 99-118 (2011)

[2] CRSG: https://wlcg.web.cern.ch/collaboration/management/computing-resources-scrutiny-group

[3] APEL, Accounting Processor for Event Logs, http://goc.grid.sinica.edu.tw/gocwiki/ApelHome

[4] G. Mathieu, A. Richards, J. Gordon, C.D.C. Novales, P. Colclough, and M. Viljoen. GOCDB, a topology repository for a worldwide grid infrastructure, J. Phys.: Conf. Ser. **219**, 062021 (2010)

[5] J. Andreeva, D. Christidis, A. Di Girolamo, O. Keeble, WLCG space accounting in the SRM-less world, EPJ Web Conf. **214** (2019)

[6] P. Andrade, A. Aimar, S. Brundu, B. G. Bear, G. M. Borge, L. Magnoni, D. L. Nicolau and N. Tsvetkov, WLCG Dashboards with Unified Monitoring, EPJ Web Conf. **245** (2020)