*Proceedings of the VIII International Conference "Distributed Computing and Grid-technologies in Science and Education" (GRID 2018), Dubna, Moscow region, Russia, September 10 - 14, 2018*

# CRIC: THE INFORMATION SYSTEM FOR LHC DISTRIBUTED COMPUTING

**A.V. Anisenkov [1,2]**

[1] *Budker Institute of Nuclear Physics, 11 ak. Lavrentieva prospect, Novosibirsk, 630090, Russia*

[2] *Novosibirsk State University, 2 Pirogova street, Novosibirsk, 630090, Russia*

E-mail: Alexey.Anisenkov@cern.ch

The Worldwide LHC Computing Grid infrastructure includes about 200 participating computing centers affiliated with several partner projects over the world. It is built by integrating heterogeneous compute and storage resources in diverse data centers in order to provide CPU and storage capacity to the LHC experiments to perform data processing and physics analysis at petabytes scale data operations. Moreover the experiments extend the capability of WLCG distributed environment by actively connecting opportunistic cloud platforms, HPC and volunteer resources. In order to be effectively used by the LHC experiments, these distributed resources should be well described, which implies easy service discovery and detailed description of service configuration.
CRIC represents the evolution of ATLAS Grid Information System (AGIS) into the common experiment independent high-level information framework which has been evolved in order to serve not just ATLAS Collaboration needs for the description of distributed environment but any other virtual organization relying on large scale distributed infrastructure as well as the WLCG on the global scope. CRIC collects information from various information providers, complements it with experiment-specific configuration required for computing operations, performs data validation and provides coherent view and topology description to the LHC VOs for service discovery and usage configuration.
In this contribution we describe the design and overall architecture of the information system, recent developments and most important aspects of the CRIC framework components implementation and features like flexible definition of information models, built-in collectors, user interfaces, advanced fine-granular authentication/authorization.

Keywords: information system, grid middleware, grid computing facilities, distributed computing

## 1. Introduction

One of the important technical challenges in functioning of a modern high-energy physics experiments is the integration of information and computing resources into a unified system for efficient storing, processing and analyzing of experimental data. For example, the ATLAS experiment [1] at the Large Hadron Collider (LHC) at the European Organization for Nuclear Research (CERN) collects billions of events each data-taking year and produces an even larger amount of events (tens of petabytes data) through the simulation according to several physics and detector models. All these data are stored and processed over heterogeneous distributed environment – the ATLAS Distributed Computing infrastructure (ADC) [2], which is based on Worldwide LHC Computing Grid (WLCG) [3] and extending the capability of the WLCG by including opportunistic cloud platforms, HPC and volunteer resources.

The WLCG is a global collaboration of the computing centers located all over the world (more than 40 countries covered) with a main goal to finally provide a resource to store, distribute, process and analyze unprecedented amount of LHC data, offering transparent access to this data for physics community of more than 10 thousand researches of the world. The variety of involved distributed computing infrastructures and middleware providers used by the LHC Experiments, their software and hardware implementation of computing environments require special attention and approaches to the integration into the unified heterogeneous computing environment of the particular LHC experiment as well as require the development of a central information system to describe the topology of computing resources and theirs specifics. The ATLAS Grid Information System (AGIS) [4] is such middleware framework designed for ATLAS experiment in order to address all these questions, mask the heterogeneity of computing environment, and provide a unified way to operate and configure distributed computing applications.

AGIS collects, structures, validates and exposes the description of computing topology and many others parameters that are needed to effectively make the data transfers, submit the jobs, properly configure high-level services, and monitor coherently the whole ADC services. Being in production since LHC Run-1 data-taking phase as a central information grid middleware system for Distributed Computing in ATLAS, and providing agile, robust and flexible information service, AGIS is continuously evolving to fulfill new requirements, enable enhanced operations and follow the extension of computing environment. The further evolution of the AGIS system beyond the ATLAS Collaboration into an experiment-independent framework (as a unified information system for LHC distributed computing) occurs in the new project called Computing Resource Information Catalog (CRIC), which is aimed to support other communities, in addition to ATLAS, the experiments at LHC by decoupling resource description in common part and isolating all experiment specific logic into dedicated plugins.

The overall goal of the CRIC information system is focused to provide a consistent description of distributed computing models used by the all LHC Experiments, collect and aggregate data coming from various low-level providers, middleware information services and generic information sources, store and expose different parameters and configuration values which are needed by many distributed computing software components, experiment oriented middleware services and applications.

## 2. The Computing Resource Information Catalog

The Computing Resource Information Catalogue (CRIC) is a high level information system aiming to describe the topology of the WLCG infrastructure as well as experiment-specific configuration required to exploit this infrastructure according to the experiment computing model. It’s also supposed to be the main entry point to consume overall information about WLCG resources required for central WLCG Operations. Based on architecture of the AGIS system, it’s comprised of a core module that is integrated with experiment specific modules (or plugins), as described further in this paper. The core module implements the description of physical resources as they are provided by distributed computing environment, while experiment oriented plugins help to extend the computing model, complement information schema and structure the resources how they are actually used by given experiment. The core part consumes information from different information sources (like GOCDB [5], OIM [6], BDII, REBUS [7]) and it's flexible enough to add or remove low-level information providers and even allow sites to update directly information about their resources using CRIC WebUI. By integrating together all possible information providers required for given experiment in the core part of CRIC implementation the system allows to solve the issue of data inconsistency usually faced by an experiment when data provided by various grid components or information sources sometimes are incomplete or contradictory. Since there is no central place where data can be checked and validated by given experiment due to heterogeneous nature of distributed environment, CRIC plays an essential role as a central experiment-oriented information portal. Finally, such flexibility of decoupling information schema allows LHC experiments to extend own Computing model in a unified way and complement traditional grid resources provided by WLCG computing with other opportunistic resources like commercial cloud platforms, HPC clusters or supercomputer centers, volunteer computing. Moreover, opportunistic sites due to their dynamic nature do not need to be part of GOCDB or OIM systems anymore, nor run grid middleware services (e.g. a BDII for service discovery) to be able to be described in CRIC. This also offers a major advantage to small sites that don't have enough effort to run extra information services. The Figure 1 illustrates how CRIC as a grid information middleware connects together physical resources and LHC VO applications.

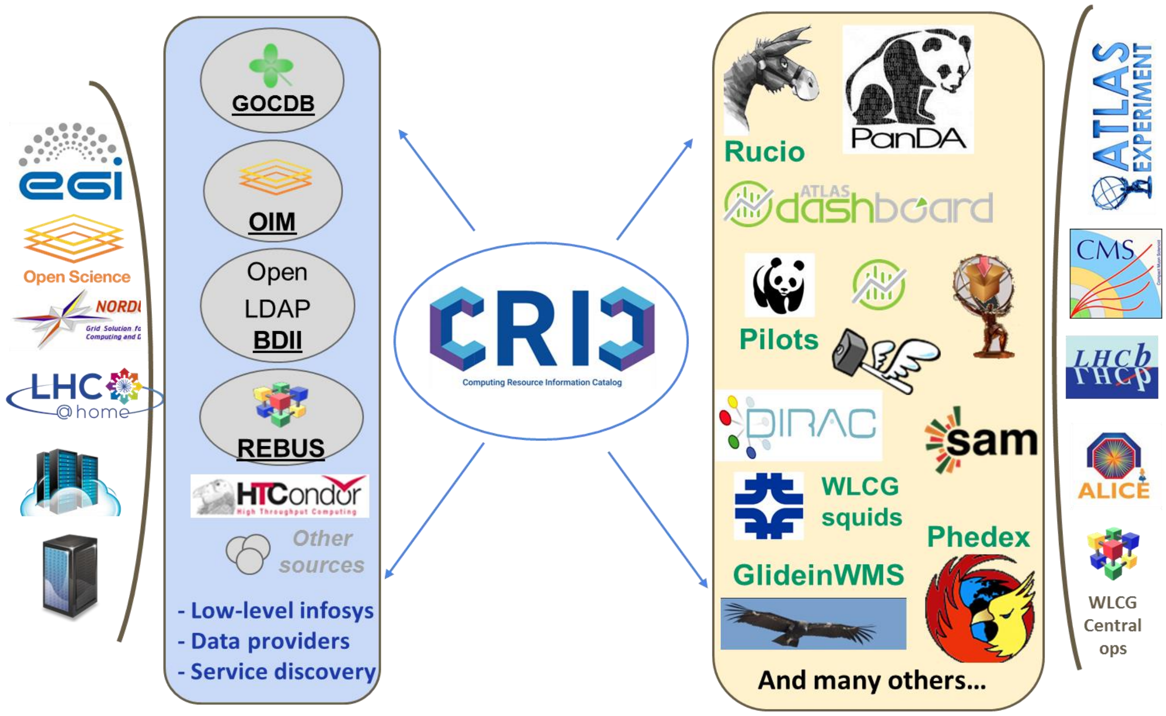


Figure 1. Distributed resource configuration systems and VO applications: gluing them together via high level Information component CRIC

## 3. CRIC architecture overview

The design of CRIC has been driven by the experience gained so far in operating ATLAS Distributed Computing environment using AGIS, as well as by future requests from WLCG operations and upcoming HL LHC challenges. The CRIC system relies on client-server computing technology by providing a set of high-level user web applications (WebUI) and REST-style programming interfaces (API) to access, modify and explore data stored in database backend. Bootstrap, jQuery, web services and many other modern tools and technologies are actively used for the development of the user interfaces.

CRIC is a database oriented system, which is currently implemented using Django framework [8] written in Python. Thanks to the object relation mapping (ORM) technique which is built into the Django, CRIC does not depend on any specific implementation of a database backend. Access to the database tables is applied in terms of high level models, thus avoiding any direct dependence on the relational database system used. This is important functional feature for CRIC, since LHC VOs require to support various database management system, for example ATLAS actively uses Oracle RBMS, while CMS and COMPASS prefer MySQL servers.

One of the key CRIC architecture features is clear separation of the description of the physical computing resources provided by the distributed sites versus the information which is required to use these resources by the experiments. In simple words: "*provided by*" vs "*used by*". By providing an abstraction layer from the physical resources, the system allows the experiment to define their own real organization of resources and complement the information schema with experiment specific structures.

"*Provided by*" information is mostly topology description of the distributed computing infrastructure which implies listing of the service endpoints, their location, implementation, software versions and others. This is generic information which is consumed by all experiments and does not depend on their computing models. It is defined by the site administrators and service providers.

"*Used by*" information defines how experiment uses a particular resource, how it is integrated in the experiment computing system. This information consists of parameters which can be experiment-specific. It is defined by the experts of the workload or data management systems of the experiments, operations teams or experiment-support teams at the sites.

Therefore, the separation of "*provided by*" and "*used by*" objects allows to ensure clean data structures with well-defined responsibilities corresponding to user roles which is translated in corresponding access/update privileges policy in CRIC. This separation makes possible to design a system which can effectively serve various experiments as well as WLCG central operations.

To automatize operations and validate information spread over different external sources available for experiments, CRIC plays essential aggregator role by automatically collecting, caching and correcting data, for example the topology relations and static information about site specifics from various databases and external information sources (gLite BDII, GOCDB, MyOSG and REBUS). It integrates such data with other dynamic details of site resources and services, like site and service status, resource downtimes and blacklisting objects. As an example from ATLAS workflow, CRIC will allows to automatically exclude from the production the resources being temporary blacklisted or currently set in scheduled downtime for Distributed Data Management (Rucio) [9] and ATLAS Production and Distributed Analysis workload management (PanDA) [10] systems.

CRIC follows the approach of modular architecture: modularity implies constructing shared building blocks which allow to optimize development process and to ensure common look and feel of the user interfaces. Such building blocks (for example, table-like visualization, tree-based data export, downtime calendar views, etc.) can be enabled and configured separately by given CRIC instance according to requirements of a LHC experiment.

Flexibility and extensibility of the system are essential to enable description of very dynamic and heterogeneous resources and to follow the technology evolution. The goal is to provide a sort of LEGO constructor and based on it, to be able to re-use implementation of various components, to extend existing functionality or implement a new one.

## 4. CRIC plugins and deployment model

The concept of *provided by* and *used by* resources discussed in the previous chapter is also implemented in decoupling the system services into dedicated plugins, which are deployed at different CRIC instances: shared WLCG CRIC instance [11] and several experiment-related CRIC services. The purpose of the WLCG CRIC instance is to describe physical resources hosted by the distributed computing sites, which are part of the WLCG infrastructure. Experiment CRICs (like CMS CRIC [12]) encapsulate specific configuration for the resources used by the experiments for data storage, data distribution and data processing. It contains all necessary information for organization of the data management and workflow management activities and models described the experiment specific concepts. Therefore, experiment CRIC plugins serve the experiment data management and workload management systems as well as various operational tools, monitoring and accounting systems. It plays a key role in the information flow of the experiment offline computing.

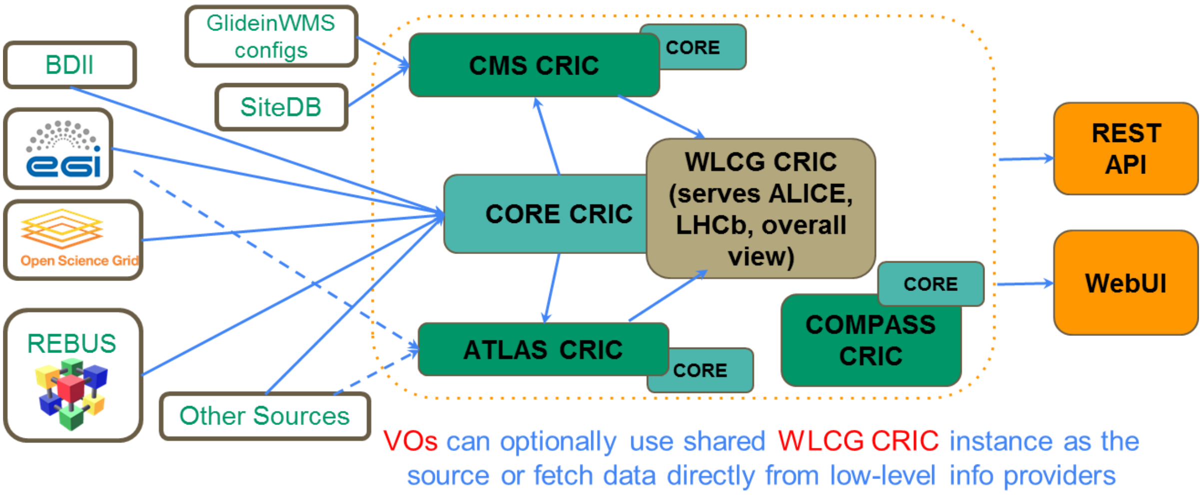


Figure 2. The overview of CRIC plugins and involved information sources

CRIC experiment plugins will be provided for ATLAS and CMS Collaborations, while WLCG CRIC plugin enables tasks required for WLCG central operations and can be used as a primary source of physical resource description for the CRIC experiment plugins. Moreover, it's supposed to implement functionality to manage Site pledges data within WLCG CRIC instance, which is currently performed in REBUS system. REBUS can be retired as soon as WLCG CRIC instance is deployed in production. For LHCb and ALICE experiments which currently intend to rely on their internal systems for "*used by*" type of information, some limited set of experiment-specific configuration (only the part required for central operations) will be provided in the WLCG CRIC instance.

It is foreseen that ATLAS CRIC and CMS CRIC instances could work autonomously without any dependency on the WLCG CRIC instance. The Figure 2 illustrates relations between CRIC plugins and involved information sources.

The COMPASS Experiment at CERN SPS is also evaluating CRIC as information middleware service. COMPASS distributed computing environment relies on similar computing models and uses same with ATLAS high-level frameworks for distributed data management and workload management systems. So that it makes possible to partial share implementations of CRIC modules for ATLAS and COMPASS plugins. The first implementation of COMPASS CRIC plugin is under the testing and planned to be integrated into the COMPASS production infrastructure once the validation step has been finished. Currently implemented COMPASS plugin [13] covers basic functionality required by COMPASS production infrastructure, in particular, the description of compute objects used by the PanDA workload management system.

The information system is a critical service for any kind of computing operations. Support for shared responsibilities between various teams and individual users who are in charge of certain operations or services has to be enabled in the information system. CRIC provides out of the box the list of enhanced authorization and authentication methods to have customized and fine-grained access control management. CRIC interfaces provide several authentication methods which can be used by end-clients depending of their needs: starting from local password- or SSL certificate-based authentication and ending with unified single-sign-on authentication (SSO, Shibboleth based). Each experiment could configure own data access policies using several types of supported permissions within the system:

- a global permission (used to restrict user actions, for example to allow clients to modify only part of information in WebUI forms);
- an instance specific permission (for example to allow to modify only objects associated with given resource center or site);
- a model specific permission, that affects particular type of objects (for example to allow modification of all Sites).

## 5. Conclusion

The Computing Resources Information Catalog is the evolution of the ATLAS grid information system: it benefits of the many years experience gained within the ATLAS distributed computing, and it is evolved in such a way to offer the possibility to be used by many collaboration, in particular the LHC experiments, which could decide the level of involvement, in terms of deepness of resource description, that they prefer.

All LHC experiments shares common computing infrastructure. CRIC offers a common framework describing this infrastructure with also an advanced functionality to describe all necessary experiment-specific configurations. The way the CRIC system is designed each experiment can independently describe its world and still coexist with the others under the same roof. CRIC architecture allows the collaborations to fully describe their computing model, and to use the system as central entry point for topology description and configuration of their frameworks and applications.

The first CRIC version is mainly focused on the CMS-required functionality. The CMS CRIC plugin has been deployed into production early Autumn 2018. This version enables all functionality which has been provided by the CMS SiteDB service. After integration of all CMS applications with CMS CRIC, SiteDB will be retired. Further extending of functionality aims to facilitate CMS computing operations both for data processing and data transfer. This part of work is performed in close collaboration with the CMS computing community.

WLCG CRIC plugin is a central CRIC instance which satisfies the needs of the WLCG central operations and can be used also as a primary source of topology description for other CRIC plugins. It should provide complete description of the topology and generic configuration of the WLCG resources used by all four LHC experiments.

## References

[1] Aad G. et al. The ATLAS Collaboration, "The ATLAS Experiment at the CERN Large Hadron Collider // Journal of Instrumentation. 2008. Vol. 3. S08003.

[2] S. Campana et al. 2015 ATLAS Distributed Computing in LHC Run2 J. Phys.: Conf.Ser. 664 032004

[3] The Worldwide LHC Computing Grid (WLCG), http://wlcg.web.cern.ch/ (accessed 05.11.2018)

[4] A. Anisenkov et al. 2017 AGIS: Integration of new technologies used in ATLAS Distributed Computing J. Phys.: Conf.Ser. 898 092023

[5] Grid Configuration Database (GOCDB), https://goc.egi.eu/portal/ (accessed 05.11.2018)

[6] Open Science Grid Information service (MyOSG), http://myosg.grid.iu.edu/about (accessed 05.11.2018)

[7] WLCG REsource Balance and USage (REBUS), http://wlcg-rebus.cern.ch (accessed 05.11.2018)

[8] Django project, http://www.djangoproject.com (accessed 05.11.2018)

[9] C. Serfon et al. 2016 Rucio, the next-generation Data Management system in ATLAS Nuclear and Particle Physics Proceedings 273-275 969

[10] Maeno T. on behalf of the ATLAS Collab. Evolution of the ATLAS PanDA Workload Management System for Exascale Computational Science // J. Phys. Conf. Ser. 2014. V. 513. P. 032062; doi:10.1088/1742-6596/513/3/03206

[11] WLCG CRIC instance, https://wlcg-cric.cern.ch (accessed 05.11.2018)

[12] CMS CRIC instance, https://cms-cric.cern.ch (accessed 05.11.2018)

[13] COMPASS CRIC instance, https://compass-cric.cern.ch (accessed 05.11.2018)