

1 The Rucio File Catalog in DIRAC implemented for Belle II

2 *Cédric Serfon*^{1,*}, *John Steven De Stefano Jr*¹, *Michel Hernández Villanueva*², *Hironori Ito*¹, *Yuji Kato*³, *Paul Laycock*¹, *Ruslan Mashinistov*¹, *Hideki Miyake*⁴, and *Ikuo Ueda*⁴

5 ¹Brookhaven National Laboratory, Upton, NY, USA

6 ²University of Mississippi, MS, USA

7 ³KMI - Nagoya University, Nagoya, Japan

8 ⁴High Energy Accelerator Research Organization (KEK), Japan

9 **Abstract.** DIRAC and Rucio are two standard pieces of software widely used in
10 the HEP domain. DIRAC provides Workload and Data Management function-
11 alities, among other things, while Rucio is a dedicated, advanced Distributed
12 Data Management system. Many communities that already use DIRAC have
13 expressed their interest in using DIRAC for Workload Management in combi-
14 nation with Rucio for Data Management. In this paper, we describe the integra-
15 tion of the Rucio File Catalog into DIRAC that was initially developed for the
16 Belle II collaboration.

17 1 Introduction

18 DIRAC (Distributed Infrastructure with Remote Agent Control) [1] is a piece of software
19 initially developed for the LHCb collaboration [2]. It is designed as "interware" as it provides
20 a complete solution for managing distributed resources. Among the various functionalities of
21 DIRAC, one can find a Workload Management System and a Distributed Data Management
22 system. On the other hand, Rucio [3] is a Distributed Data Management system. Initially
23 developed by the ATLAS collaboration [4], Rucio has quickly gained popularity outside AT-
24 LAS due to its advanced features. Both DIRAC and Rucio [5] are now used by a large
25 community beyond their initial collaboration. In recent years, some communities have ex-
26 pressed interest in using DIRAC for Workload Management in combination with Rucio for
27 Distributed Data Management. The Belle II collaboration [6] developed a Rucio File Cata-
28 log (RFC) plugin for its extension of DIRAC called BelleDirac [7]. In section 2, this article
29 details some of the differences between the RFC plugin and the file catalog plugins already
30 implemented in DIRAC. Section 3 details all the methods that have been implemented and
31 presents differences with respect to the other catalogs.

32 2 DIRAC and File Catalogs

33 For experiments using distributed computing resources it is often the case that there will be
34 multiple copies, or replicas, of files across the different computing sites. An obvious example

*e-mail: cedric.serfon@cern.ch

35 is that all experiments ensure that there is a copy of the precious raw detector data, but there
36 may be multiple copies of processed outputs to allow those data to be used at multiple sites at
37 the same time. File catalogs provide a means to provide coherent access to file replicas. Each
38 file in the catalog has a Logical File Name (LFN) and each LFN can have a list of associated
39 Physical File Names (PFN) that correspond to the physical copies of the files at different
40 sites. If an application or a user wants to locate a particular LFN, then they simply query the
41 catalog to get the list of associated file replicas.

42 **2.1 File catalogs supported by DIRAC**

43 DIRAC provides a File Catalog interface to the actual file catalog service, more details about
44 that interface are described in section 3. To interact with a particular file catalog service, a
45 plugin that implements the methods defined in the File Catalog interface is needed. Before
46 the implementation of the RFC plugin, two file catalogs were supported by DIRAC:

- 47 • The first one is an external catalog called the LCG (LHC Computing Grid) File Catalog
48 (LFC) [8]. It is a hierarchical catalog that allows one to organise the files into a directory
49 structure. The LFC only contains the file replica information and very limited metadata
50 including the checksum and the file size.
- 51 • The DIRAC File Catalog (DFC) is another catalog internal to DIRAC. In contrast to the
52 LFC, the DIRAC File Catalog combines both replica and metadata functionality [9].

53 **2.2 Differences between Rucio and the LFC**

54 The Belle II collaboration previously used the LFC and wanted to maintain the same func-
55 tionality and behaviour. The RFC plugin was designed to meet that goal with some slight
56 differences listed below:

- 57 • Whereas the LFC is inherently hierarchical, Rucio uses by default a flat namespace with
58 files contained in datasets which are a collection of files, but the datasets are not connected
59 to one another. Rucio has another type of data structure called containers which are a
60 collection of datasets and/or containers. Containers cannot contain files. Using containers,
61 it is then possible to reproduce the directory structure of the LFC (see figure 1) with the
62 constraint that directories cannot contain a mixture of files and directories. This same
63 constraint was introduced into Belle II Software to prevent users from encountering this
64 feature.
- 65 • Rucio has the concept of scope : The scope is a way to partition the Rucio namespace and
66 to apply different policies, permissions, etc. Every Data IDentifier (DID) which represents
67 a file, dataset, or a container, has a scope and the DID name must be unique within the
68 scope. Since there is no concept of scope in the LFC, it needs to be hidden from end-users
69 and applications. To achieve this, a deterministic function uses the LFN to associate each
70 LFN to a unique scope in a transparent way.
- 71 • Another big difference is related to the Rucio concept of replication rules. The replication
72 rules are a way to describe how a DID must be replicated on a list of Rucio Storage Ele-
73 ments (RSE). If a rule is created for a particular DID on certain RSE, Rucio will ensure
74 that the rule is fulfilled either by locking the DID at the specified RSE if it is already there,
75 or by transferring and then locking it to the specified RSE. This difference has important
76 consequences for deletion as explained in subsection 3.3.
- 77 • Finally there is no concept of fine-grained Access Control List (ACL) on files or directories
78 in Rucio, contrary to the LFC. In Rucio, permissions are managed at the scope level, and
79 all of the files in the same scope have the same permissions.

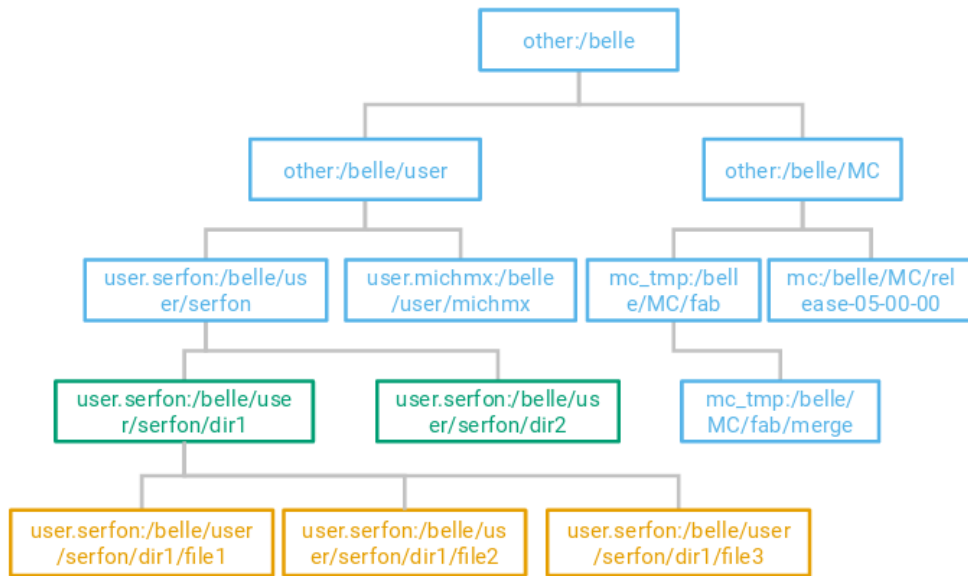


Figure 1. Schema showing how the data are structured in Rucio to reproduce the LFC hierarchy. The orange boxes represent files, bluish green boxes represent datasets that can only contain files, and sky blue boxes are containers that can only contain datasets and other containers. The first part of the name before the colon represents the scope and is associated uniquely to the LFN which follows the colon.

2.3 Differences between Rucio and the DFC

All the differences listed in section 2.2 also apply to the RFC compared to the DFC. In addition, similarly to the DFC, Rucio supports generic metadata. In the DFC, subdirectories inherit the metadata of their parent directories and files inherit the metadata of their directories. In Rucio it is possible to set any metadata to any DID and children do not inherit metadata from their parents.

3 Rucio file catalog plugin functionality

DIRAC provides a File Catalog interface that allows users or any other (Belle)DIRAC component to interact with the file catalog service. It provides several different methods which can be categorised as:

- Read methods: To list the content of a directory, to get stats about files or directories, to get the list of PFN (replicas) associated with one LFN, and many other methods.
- Write methods: To create new files, and to create new replicas.
- Delete methods: To delete file replicas (remove a PFN associated to an LFN), or delete a file completely, i.e. remove an LFN from the catalog.

To implement a new catalog plugin, all these methods need to be implemented. One potential obstacle to implementing a Rucio file catalog plugin is due to the important concept of scope. As mentioned in the previous section, the scope is an unknown concept both for the LFC and DFC. Therefore the scope cannot be passed to the catalog method and needs to be extracted

99 directly from the LFN. This is done with the help of a deterministic function that maps each
100 LFN to one and only one scope.

101 The current implementation of the RFC plugin only contains the replica methods that
102 are available in both the LFC and the DFC, but not the metadata methods that are unique
103 to the DFC. In order to use the RFC plugin, the Rucio clients need to be installed on the
104 DIRAC server. This can be done using the Python package manager pip. In the future, it
105 is foreseen to include the Rucio clients into DIRACOS [10]. To setup the RFC plugin, a
106 few environment variables need to be defined in a RucioFileCatalog section of the DIRAC
107 Configuration system as shown in figure 2.

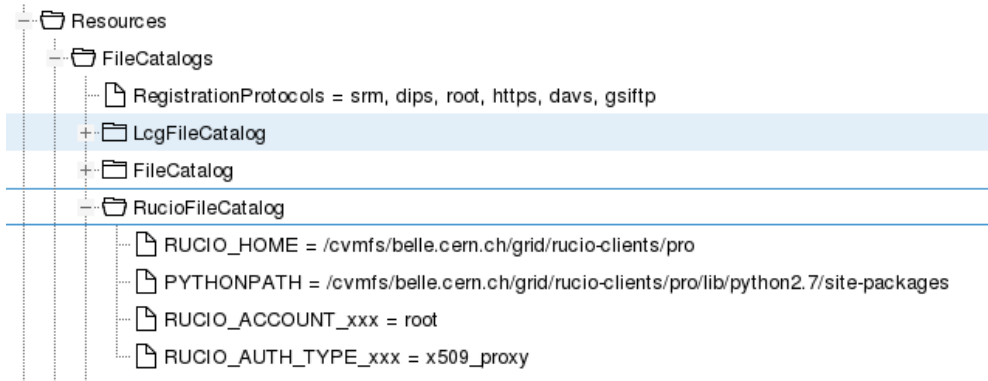


Figure 2. Snapshot of the DIRAC Configuration Service used by Belle II. The RucioFileCatalog section is a subsection of Resources/FileCatalogs that contains a few environment variables that are used to setup the Rucio client.

108 Before starting to use the RFC plugin, the Storage Elements registered in the Resources
109 section of the DIRAC Configuration, as well as the users registered in the Registry section,
110 need to be created on the Rucio server. In the Belle II case this is done automatically by a
111 new DIRAC agent called the RucioSynchronizer that creates the RSE and their associated
112 protocols, as well as the user accounts.

113 3.1 Read methods

114 The following read methods have been implemented: `getReplicas`, `listDirectory`,
115 `getFileMetadata`, `getFileSize`, `isDirectory`, `isFile`, `getDirectorySize`. Using
116 the mapping between datasets/containers and directories described in 2, all these methods
117 have the same behaviour as the ones in the LFC plugin. They use bulk queries to the Rucio
118 server which allows for faster response times in case more than one file is specified. This is
119 particularly important considering that the DIRAC server and the Rucio server can be relatively
120 far away, e.g. for Belle II there is a distance of more than 10000 kilometers between
121 the DIRAC servers and the Rucio servers which represents about 180 ms of Round Trip Time.

122 3.2 Write methods

123 There are only two write methods: `addFile` and `addReplica`. For `addFile` a new atomic
124 bulk method was added to Rucio. Adding new files involves many operations such as creating
125 all of the parent directories if they do not exist, attaching files to the dataset, creating the file

126 replicas, and creating a replication rule for the dataset. The whole workflow is described in
127 figure 3. Regarding `addReplica`, the method simply adds a file replica and a replication rule
128 for this replica.

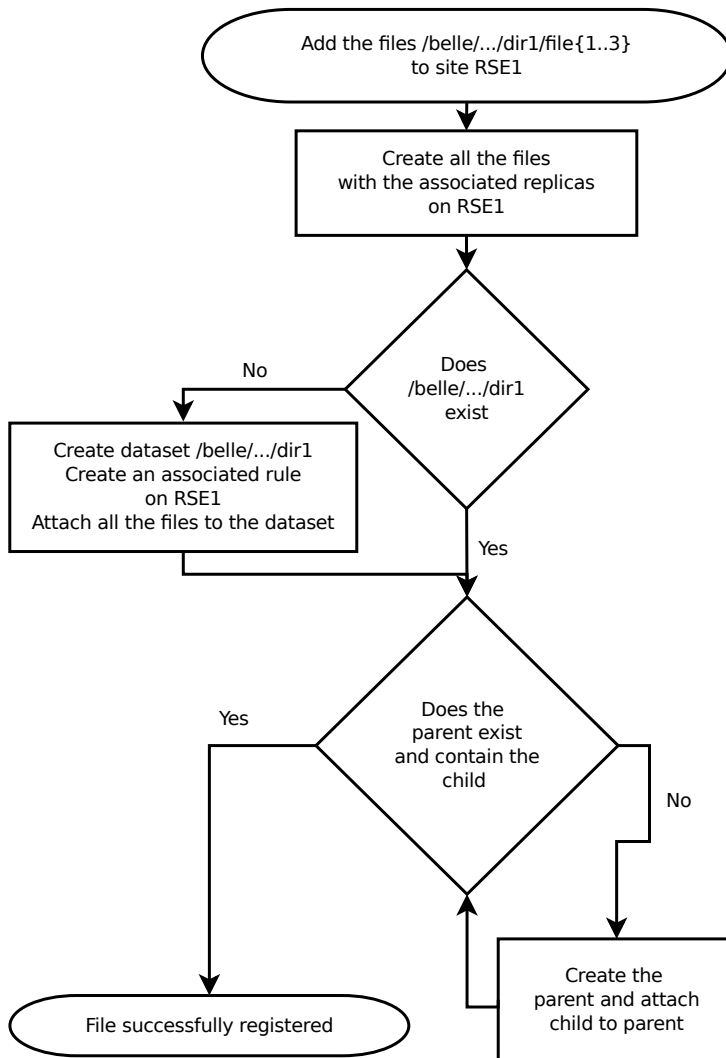


Figure 3. Diagram describing the workflow to add new files at a storage element. The entire procedure is handled in a new atomic bulk operation on the Rucio server side.

129 3.3 Delete methods

130 There are two delete methods: `removeReplica` and `removeFile`. Due to the very different
131 concepts between Rucio and the LFC regarding deletion, their behaviour is different in the
132 RFC compared to the LFC methods. As explained in section 2, the DID in Rucio are locked

133 at a specific site using replication rules. These rules prevent the deletion of a file replica,
134 whereas this operation is a valid one for the LFC. Therefore if a file that belongs to a dataset
135 has a replica on site A and this dataset has a replication rule on site A, it is impossible to
136 remove the replica. Therefore the `removeReplica` method simply doesn't do anything in
137 the RFC plugin. The `removeFile` method that removes a file from the namespace also has
138 a different behaviour : in the LFC plugin, the command only succeeds if this file has no
139 replicas, whereas in the RFC plugin, the file is removed even if replicas exist. Additionally,
140 whereas in the LFC case the file deletion from storage is done synchronously, in the RFC the
141 file is only logically removed from its parent directory synchronously, while the logical and
142 physical deletion of the file itself and from its replicas is done asynchronously by a separate
143 Rucio daemon.

144 3.4 Future work

145 The current implementation only supports the replica functionality and not the metadata func-
146 tionality supported by the DFC. The RFC plugin could be extended to also support some
147 metadata methods similarly to the DFC as Rucio provides the possibility to store generic
148 metadata (key/value pairs). In addition, the plugin that is currently only part of BelleDirac is
149 foreseen to be integrated to the generic DIRAC.

150 4 Conclusion

151 A new Rucio File Catalog plugin has been developed to interface DIRAC with Rucio. The
152 plugin is being used successfully by the Belle II collaboration in production. This paper sum-
153 marized some of the differences between this new catalog and the other catalogs supported by
154 DIRAC: the LCG File Catalog and the DIRAC File Catalog. Once included into the generic
155 DIRAC, this new interface should help with the adoption of Rucio with DIRAC by more
156 communities.

157 References

- 158 [1] Federico Stagni, Andrei Tsaregorodtsev, André Sailer and Christophe Haen, “The
159 DIRAC interware: current, upcoming and planned capabilities and technologies“, EPJ
160 Web Conf. **245** 03035 (2020). doi: 10.1051/epjconf/202024503035
- 161 [2] A. A. Alves, Jr. *et al.* [LHCb], “The LHCb Detector at the LHC” JINST **3**, S08005
162 (2008) doi:10.1088/1748-0221/3/08/S08005
- 163 [3] Martin Barisits *et al.*, “Rucio - Scientific data management,” Comput. Softw. Big Sci. **3**
164 (2019) no.1, 11 doi:10.1007/s41781-019-0026-3
- 165 [4] G. Aad *et al.* [ATLAS Collaboration], “The ATLAS Experiment at the CERN Large
166 Hadron Collider” JINST **3** S08003 (2008) doi:10.1088/1748-0221/3/08/S08003
- 167 [5] Martin Barisits *et al.*, “Rucio beyond ATLAS: experiences from Belle II, CMS,DUNE,
168 EISCAT3D, LIGO/VIRGO, SKA, XENON”, EPJ Web Conf. **245** 03035 (2020).
169 doi:10.1051/epjconf/202024511006
- 170 [6] T. Abe *et al.*, KEK-REPORT-2010-1, arXiv:1011.0352 (2010)
- 171 [7] H. Miyake *et al.* [Belle-II computing group], “Belle II production system,” J. Phys.
172 Conf. Ser. **664** (2015) no.5, 052028 doi:10.1088/1742-6596/664/5/052028
- 173 [8] J.P. Baud, J. Casey, S. Lemaitre and C. Nicholson, “Performance analysis of a file cat-
174 alog for the LHC computing grid”, HPDC-14. Proceedings. 14th IEEE International
175 Symposium on High Performance Distributed Computing, 2005., Research Triangle
176 Park, NC, 2005, pp. 91-99, doi: 10.1109/HPDC.2005.1520941.

- 177 [9] A. Tsaregorodtsev *et al.* [DIRAC], “DIRAC file replica and metadata catalog”, J. Phys.
178 Conf. Ser. **396** (2012), 032108 doi:10.1088/1742-6596/396/3/032108
- 179 [10] M. Petrič, C. Haen and B. Couturier, “DIRACOS: a cross platform solution for grid
180 tools”, EPJ Web Conf. **245** (2020), 03020 doi:10.1051/epjconf/202024503020