

# Development of the Dataset Searcher Webapp for finding data on the Belle II computing grid

Kim Smith<sup>1,\*</sup>, David Dossett<sup>1,\*\*</sup> and Martin Seviar<sup>1,\*\*\*</sup>

<sup>1</sup>University of Melbourne, Australia

**Abstract.** In any large scale scientific experiment involving enormous quantities of data it is crucial that everyone involved has quick and easy access to all the relevant datasets for their research. By the end of the run time of the Belle II experiment there will be a projected  $50 \text{ ab}^{-1}$  of integrated luminosity making it no exception. Until now the only method for locating data of interest was by looking up hand written tables that needed to be regularly updated. In this paper, a new webapp built on the DIRAC software framework will be presented which aims to be the new standard for not only locating data but also storing all its associated metadata.

## 1 Introduction

As Belle II data production continues throughout this and the following years, the number of analysts requiring access to both real and Monte Carlo data will increase dramatically. With this in mind, it becomes essential that physicists are able to locate data of interest in an intuitive and efficient manner. Otherwise, what was once an inconvenience will escalate into a serious hindrance. In the past, the Belle I experiment made use of a file search engine that stored file locations on a remote server that could then be queried by members of the Belle collaboration. A similar system has been employed here but is instead built upon the DIRAC software interface and has the ability to provide not just the locations of files and datasets but also all other metadata that has been stored.


## 2 The Current System

Previous attempts at organising information for users to find have involved storing the locations of files and datasets in lookup tables found on the internal wiki hosted by Confluence. This came with a large amount of overhead in the form of constant and manual updates by the data production team for every new production run. From the analyst's point of view, trying to find the right files was a slow and tedious process which often resulted in frustration. Additionally these tables had a very limited amount of relevant metadata listed along side its location. In terms of a technical framework description, Belle II data is stored on a computing grid across many storage elements (SE) in many institutions around the world [2]. The locations of datasets are known as logical path names (LPN) and their metadata

\*e-mail: [khsmith@student.unimelb.edu.au](mailto:khsmith@student.unimelb.edu.au)

\*\*e-mail: [david.dossett@unimelb.edu.au](mailto:david.dossett@unimelb.edu.au)

\*\*\*e-mail: [martines@unimelb.edu.au](mailto:martines@unimelb.edu.au)



**Figure 1:** Projected peak and integrated luminosity for the Belle II experiment until the end of its run time. 50  $\text{ab}^{-1}$  of integrated luminosity is planned by 2027 [1].

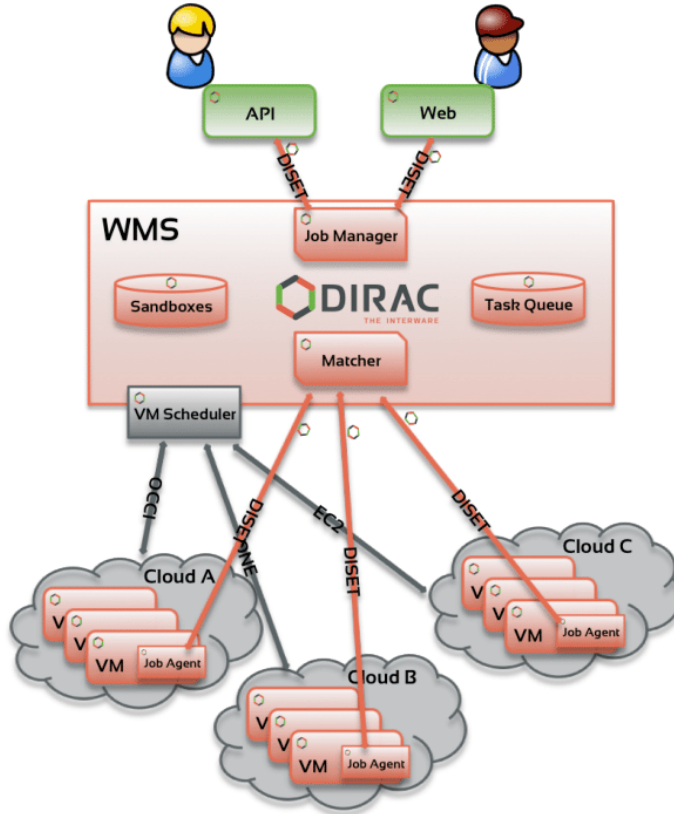
are stored in the AMGA database while their associated SE is tracked by the logical file catalogue (LFC). Software tools for working on the grid are hosted by the DIRAC software framework with a custom subsystem created for the special needs of the Belle II experiment called BelleDIRAC. Finally, file replication and transference is handled by the distributed data management system (DDM).

## 2.1 The Problems

This system has several issues which make it non-ideal for managing and accessing data.

1. The AMGA database is primarily designed for attaching metadata to LPNs and retrieving it back again given the LPN. The task of querying metadata for matching LPNs is not well supported and results in poor retrieval times and overall performance.
2. Due to the difficulty of finding LPNs from metadata, much of it has been encoded into the LPN itself. This has the side effect of making them much longer, which in turn makes storing them in DIRAC and MySQL tables problematic.
3. As mentioned previously, the only way for users to find the locations of data on the grid is by manually searching through large human generated lookup tables.

These are the problems that need to be addressed in our solution.

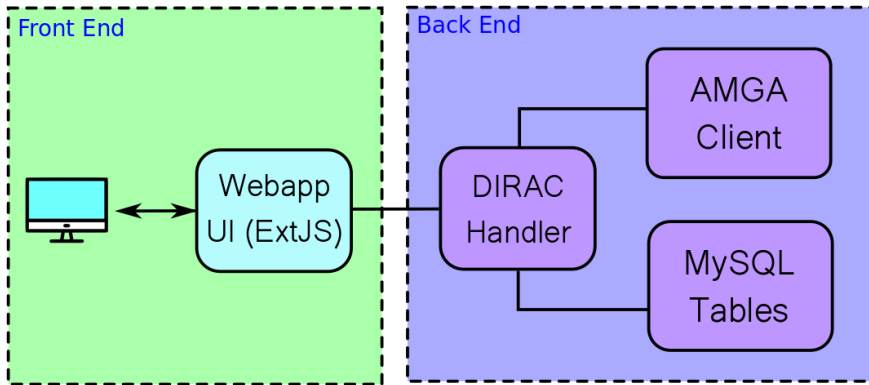


**Figure 2:** The DIRAC Interware software framework [3]. It is an open source solution to the task of managing distributed computing resources via a grid based cloud storage. It provides many communities and organisations the ability to interface with a unified job submission and monitoring system. It also provides the freedom for individual organisations to extend the software to fit their own needs, as is the case with BelleDIRAC.

### 3 Requirements

Before beginning we decided on some design goals that an eventual solution should embody. In order to minimise the time spent locating data by users, it would need to be able to perform the following tasks quickly and seamlessly:

1. Browse the Belle II computing grid via a user friendly interface.
2. Find all metadata associated with a specified LPN provided by the user.
3. Find all LPNs matching a user provided set of metadata.
4. The underlying database must support CRUD functionality, i.e. for those with the correct permissions it should be easily modifiable.
5. It must be able to perform all of these tasks in much less time than it would take to query the AMGA database directly.

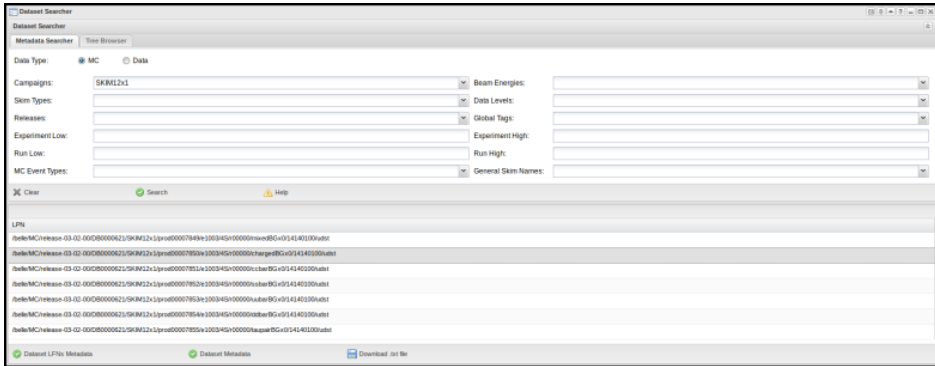


**Figure 3:** End users interact with a GUI written in extended JavaScript which is responsible for passing the search criteria onto the Dirac handler written in Python. The handler then communicates with the AMGA client and MySQL tables before returning the results back to the ExtJS front-end.

## 4 Our Solution

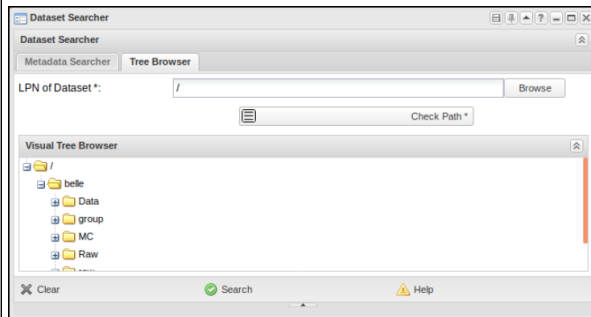
We have created our own much smaller MySQL based database which will only contain databases that are relevant to users. This only needs to query AMGA when new entries are added, resulting in much faster retrievals of metadata. By using a well established relational database management system that is industry standard, we are able to take advantage of a rich feature set that has been highly optimised for performing all the common manipulations one might expect from such a system. One potential pitfall of MySQL is that it has an effective VARCHAR limit of 255 characters when creating a "primary key" in a table. Currently at Belle II our LPNs are coming close to this limit which can result in poor performance during lookups and insertions. Instead we store the LPN with no unique constraints and the SHA1 hash of the LPN as a unique key. Doing this improves the insertion performance and increases the LPN character limit to 1018, which future proofs us in the event that Belle II increases the LPN length.

The main user interface is built upon the BelleDIRAC software framework and as such is written in ExtJS. Data queries are handled with AJAX requests made to the DIRAC handler which passes information back to the JavaScript in the form of JSON blobs. The main search and display panels are shown in Figures 4a to 4d. Currently the main panel supports searching by MC and real data via the radio buttons at the top as well as the option to directly navigate the Belle II grid file system using the Tree Browser tab. This is of course a much slower method of finding data, but can be useful for finding out what is on the grid in case the user is looking for something not loaded into the MySQL database. One can search by a number of criteria including the beam energy, the skim type, the release version of BASF2 [4] that was used to generate the data and the MC campaign it was generated in. These are all fields that support drop-down selection with some pre-filled options, but also allow for the user to enter their own text. Additionally one has the option to select a search result and show a popup panel that provides all metadata that it has been associated with. Particularly useful for a more thorough search. Finally we have included the ability to download a simple text file with all the search results printed for users to pass into their own workflow offline.




(a) The main search panel. Supports MC and real data using a range of search criteria commonly used by Belle II analysts.

| Name                 | Value   |
|----------------------|---|
| beamEnergy           | 4S  |
| creationDate         | 2019-06-11 18:56:01                               |
| dataLevel            | udst  |
| dataset              | /belle/MC/release-03-02-00/DB0000621/SKIM12x1/... |
| dataType             | mc  |
| dbGlobalTag          | DB0000621   |
| description          | SKIM12x1 BtoCharm:chargedBGx0 MC                  |
| experimentHigh       | 1003  |
| experimentLow        | 1003  |
| generalSkimName      |   |
| lastUpdate           | 2019-06-12 04:41:22                               |
| logLh                |   |
| mc                   | SKIM12x1  |
| mcEventType          | chargedBGx0                                       |
| nFiles               | 1   |
| owner                | g_belle_skim                                      |
| parentDatasets       |   |
| productionId         | 7850  |
| release              | release-03-02-00                                  |
| runHigh              | 0   |
| runLow               | 0   |
| size                 | 474749655   |
| skimDecayMode        | 14140100  |
| sourceCode           |   |
| sourceCodeRevis...   |   |
| status               | good  |
| steeringFile         | MC/SKIM12x1/release-03-02-00/SkimScripts/Bto...   |
| steeringFileRevis... |   |
| stream               |   |
| transformationId     | 28579   |



(c) A lightweight tree browser panel for manual exploration of the Belle II grid file system taken from AMGA.

(b) Show all metadata for a given LPN or all the LFNs contained within by querying AMGA.



(d) Includes the ability to save all search results into a text file for offline reference or for passing into other grid tools.

**Figure 4:** Overview of the main GUI panels that users interact with.

## 5 Summary and Future Plans

In summary, the Belle II data management system has in the past been difficult to navigate for researchers not directly a part of the production team. Due to an increase in total amount of data being generated in the coming years, a new solution to the problem of finding and storing metadata on files and datasets was required. In response to this need, we have created a DIRAC based web application for exactly this purpose. It makes use of the preexisting AMGA database as well as a newly created MySQL database which stores only datasets that are potentially useful to analysts. A user friendly user interface for searching via many commonly used attributes greatly speeds up the process of finding and investigating data stored on the Grid. We have a working build that is currently on the live server and available for use by anyone with the correct permissions and plan on a number of features and milestones coming down the pipeline.

### 5.1 Short Term

In the immediate future we will continue to fill the database with requests from the data production teams and skimming groups. We also have a number of in-development python tools for directly interfacing with the database via a terminal or command line which is useful for management and maintenance. These python scripts could become official Grid tools which would provide an alternative to using the in-browser application for those that prefer it.

### 5.2 Long Term

In the mid to long term, our ultimate goal is to implement a pythonic agent that would constantly monitor for new requests from the data production team. This would then automate the process of adding and updating the to the database without the need for slow human-to-human requests. There are also plans for the current DDM system to migrate over to Rucio [5] along with the LFC during the later stages of deployment. Finally we will continue to receive and act on feedback as it comes in.

## References

- [1] T. Abe et al. (Belle-II) (2010), **1011.0352**
- [2] T. Fifield, *gBasf2 - the distributed analysis client for Belle II*, [http://event.twgrid.org/isgc2011/slides/HEP%20Application/1/ISGC2011\\_gBasf2.pdf](http://event.twgrid.org/isgc2011/slides/HEP%20Application/1/ISGC2011_gBasf2.pdf) (2011), [Online; accessed 30-January-2020]
- [3] N. Brook et al., CoRR **cs.DC/0306060** (2003)
- [4] A. Moll, Journal of Physics: Conference Series **331**, 032024 (2011)
- [5] M. Barisits, T. Beermann, F. Berghaus, B. Bockelman, J. Bogado, D. Cameron, D. Christidis, D. Ciangottini, G. Dimitrov, M. Elsing et al., Computing and Software for Big Science **3**, 11 (2019)