Data management for the SoLid experiment

Yamiel Abreu[1,*] *Daniela* Bauer[2,**], *Simon* Fayer[2], and *Janusz* Martyniak[2]

[1]Department of Physics, Universiteit Antwerpen, Antwerp, Belgium
[2]Department of Physics, Imperial College, London, London, United Kingdom

Abstract. The SoLid experiment is a short-baseline neutrino project located at the BR2 research reactor in Mol, Belgium. It started data taking in November 2017. Data management, including long term storage will be handled in close collaboration by Vrije Universiteit Brussels, Imperial College London and Rutherford Appleton Laboratory. We describe the SoLid data management model with an emphasis on the software developed for the file distribution on the Grid, the data archiving and the initial data transfer from the experiment to a Grid storage element. We present results from the first six months of data taking, showing how the system performed in a production setting.

1 Introduction

The SoLid experiment aims to measure active-to-sterile anti-neutrino oscillations at short baselines from the BR2 research reactor at SCK•CEN, Mol, Belgium. The SoLid collaboration has developed a novel highly-segmented and modular detector to measure the anti-neutrino flux from the reactor, through the Inverse Beta Decay process [1]. In 2017 a Phase 1 detector with 1.6 ton active mass was deployed at BR2; this detector is constructed out of 16×16×50 hybrid scintillation cubes and read out by 3200 optical fibres coupled to silicon photomultipliers. The experiment generates around 160 Mbit/s of raw data that are recorded to disk for off-line processing and analysis.

2 Data Flow Model

An overview of the SoLid data flow model is given in Fig. 1. The raw data are transferred via a 1 Gbit/s dedicated network from the BR2 reactor to the Grid enabled storage element at the Interuniversity Institute for High Energies (IIHE) in Brussels. From there the data files are replicated to tape at Rutherford Appleton Lab (RAL) for archiving and to disk at Imperial College London.

2.1 Integration with existing services

The SoLid experiment uses the GridPP DIRAC instance [2, 3] as a resource broker and file catalogue. Therefore there is an additional requirement that any transfer system must integrate with the DIRAC services.
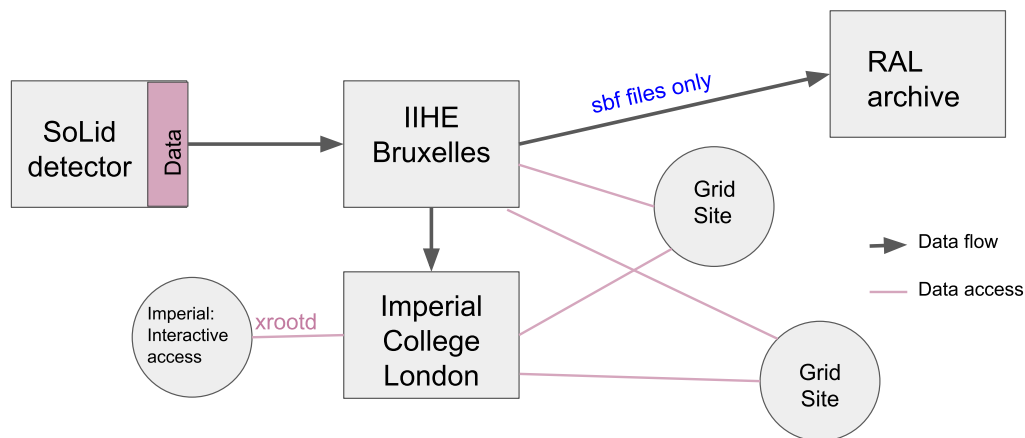
---

*e-mail: yamiel.abreu@uantwerpen.be (for the SoLid collaboration)
**e-mail: lcg-site-admin@imperial.ac.uk

2.2 File access

Imperial College and VUB offer direct data access for analysis via XRootD [4] from any Grid site, and both IIHE and Imperial College offer Grid resources for data processing.


Figure 1. The SoLid data flow model. SBF ('SoLid Binary Files') refers to raw data. Only raw data files are archived.

3 The Data Archiver

The SoLid raw data archiving system is shown in Fig. 2. The system consists of two independent processes running on a dedicated machine at Imperial College. The Collector process queries the DIRAC file catalogue and stores transfer candidates in a dedicated PostgreSQL database. The database holds raw files' metadata, in particular the DIRAC file catalogue registration timestamp which allows the Archiver to perform incremental backups. The Dispatcher process triggers an FTS3 [5] third party transfer between the IIHE storage element and the CASTOR [6] tape storage at RAL. The FTS3 service is fine-tuned to perform parallel file transfers at a speed sufficient to avoid a backlog. The archiving system has proven to work reliably during the 2017-2018 data taking period, uploading 200 TB of SoLid raw data files to permanent storage.

4 Data Replication to Imperial College

SoLid requested a second online copy of the data to be available at Imperial College for analysis work. As an interim solution we implemented a set of simple cron-driven replication scripts based on existing tools. The scripts search the DIRAC file catalogue for any files
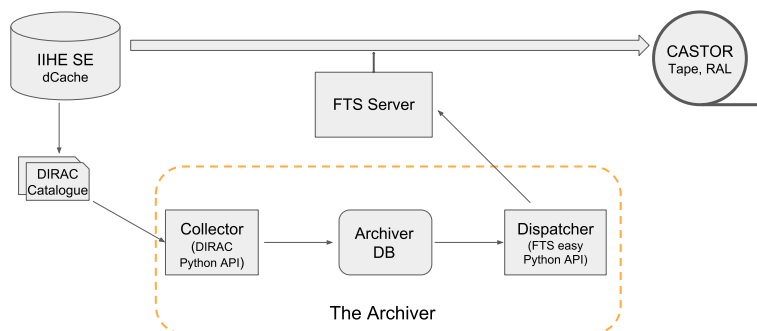
Figure 2. The Solid Archiver. It copies SoLid raw data files from the IIHE storage element to the CASTOR tape storage at RAL.

registered in the previous three days and submit transfer jobs to an FTS3 server. Once the transfer is complete, the additional file replicas are registered in the file catalogue. To mitigate intermittent limitations in bandwidth from IIHE, the FTS3 server was tuned to use a larger number of streams than default. This set-up works well for the bulk of the data, but occasionally requires system administrator intervention for edge cases, such as when files are registered after three days. We have started investigating using the DIRAC transformation system as a possible replacement for the current set-up; this should minimise the need for human intervention.

5 The DIRAC Transformation System

The SoLid experiment uses the DIRAC workload management and file catalogue systems to process their data on the Grid. DIRAC also offers a workflow engine ("Transformation System (TS)") which is capable of executing data management operations.

In the case of SoLid data, the first step is the data being uploaded to the Belgian storage element and registered in the DIRAC file catalogue. There are two ways these files can be processed: Either the TS Client can intercept the file registration requests and then execute the predefined workflow, which in this case is the replication to the Imperial College storage element. In addition the Input Data Agent runs regular metadata queries to scan for files that could not be processed by the TS Client, e.g. because the TS had been temporarily unavailable. Both the TS Client and the Input Data Agent notify the Transformation Agent of the existence of a new file. The Transformation Agent checks any applicable rules (e.g. only certain file types need to be replicated) and creates the replication tasks. A designated module, the Replication Manager, handles the actual replication of the file. An overview of the SoLid specific TS set-up is show in Fig. 3.

So far this has been shown to work in a proof of principle set-up. However a number of issues relating to using a shared multi-VO DIRAC server were identified and it is not possible to move to a production type set-up until these are rectified.
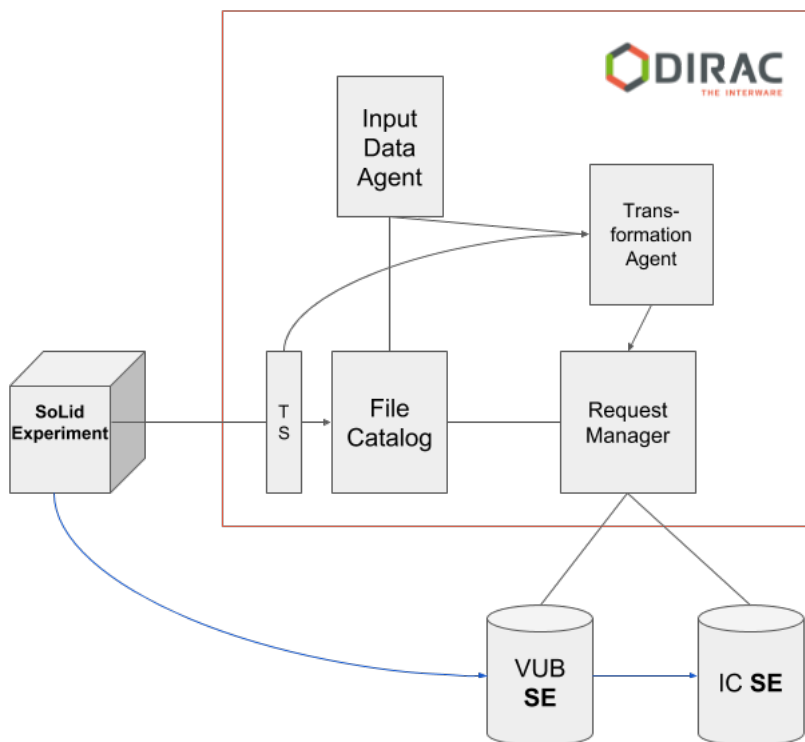
Figure 3. DIRAC Transformation System set-up for automated file replication between two Grid enabled storage elements.

6 Conclusions and Outlook

The SoLid experiment started their commissioning run in November 2017 and has been taking data since January 2018. In accordance with the SoLid experiment data model we developed a reliable archiver to transfer the data to tape for backup. So far we have successfully archived 200 TB of data. While specifically written for the SoLid experiment, this archiver could be readily adapted for other communities.

We have also investigated automated data transfer between the two main SoLid analysis sites. This is currently based on a prototype system, which is sufficient for the current data rate, but will be replaced by a DIRAC based replication system in the near future.

We anticipate that this set-up will then fulfil the requirements of the SoLid experiment for the foreseeable future.

References

[1] SoLid collaboration, *A novel segmented-scintillator antineutrino detector*, JINST **12** P04024 [1703.01683] (2017)
[2] A Tsaregorodtsev et al, *DIRAC: A Community Grid Solution*, Journal of Physics: Conference Series **119** 062048 (2008)
[3] D Bauer, S Fayer, GridPP DIRAC: Supporting non-LHC VOs on LHC centric resources, J. Phys.: Conf. Ser. **898** 052003 (2017)

[4] XRootD [software] http://xrootd.org/ [accessed 2018-09-07]

[5] FTS: File Transfer Service [software] http://fts.web.cern.ch/ [accessed 2018-10-25]

[6] The CERN Advanced Storage manager (CASTOR) [software] http://castor.web.cern.ch [accessed 2018-10-25]