



B2SAFE

About

General introduction to the B2SAFE service

Modified: 29 January 2018

Synopsis

B2SAFE is EUDAT's service for the secure long-term preservation of research data. The safety of the data is ensured by means of B2SAFE's replication mechanism which automatically replicates the data to one or several backup sites and which maintains all information on replicas in an additional system to guarantee the findability of data.

This document outlines B2SAFE's functionality. For more insight into the technical details and testing we refer to an extensive modular hands-on tutorial available at the [EUDAT training repository on GitHub](#) which covers:

- Installation
- Usage
- Integration with other EUDAT services

What is B2SAFE?

B2SAFE is EUDAT's service for secure long-term preservation of research data. Data in B2SAFE is kept safe by replicating them to one or several other EUDAT sites, i.e. creating redundant copies of data and maintaining those by different administrative units.

Additionally to the replication workflow, the B2SAFE technology offers the framework to implement community-specific data policies. B2SAFE can store and replicate large amounts of data. It is meant to be used by repositories to preserve and backup their data collections. Moreover, B2SAFE can replicate reference datasets to various compute sites which are usually co-located with the B2SAFE endpoints.

Managing copies (replicas) of data across different sites requires a mechanism to verify data integrity and to manage the different data endpoints. To this end B2SAFE employs Persistent Identifiers (PIDs). PIDs are usually guarantee the identity of data and provide the means to cite data. B2SAFE employs them in a slightly different way. Each data object, original and replica, is assigned with a PID. The PID record itself contains all necessary information to link a data object to its parents if it is a replica, or it lists all direct children if the data object is a original data object. EUDAT designed a specific PID profile which structures this metadata.

B2SAFE can be employed in two modes. 1) A community centre can join the B2SAFE network which requires to deploy the full B2SAFE software stack or 2) a community can use B2SAFE, please see the Section "[More Information](#)" at the bottom of this page.

Technical Details

In this section we will explain how B2SAFE works in detail, which software it is based on and which policies and workflows are supported and can be configured.

EUDAT receives funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 654065.



B2SAFE is based on the data management software iRODS and is implemented as a specific set of policies in iRODS called rules. The B2SAFE rule set integrates data handling in iRODS with tracking data across iRODS instances by means of Persistent Identifiers.

While iRODS itself offers to create workflows and rules with which users can directly work on data, data in iRODS that is subject to the B2SAFE ruleset is not meant to be accessed directly by the user (scientist), i.e. users should not be allowed to change data; data stewards should be careful with changing data and if doing so make sure that all necessary replicas and information on the data are updated and propagated through the replication chain. It is advised to have an iRODS Expert on Site when running the service in joining mode.

Short introduction to iRODS

iRODS is a Data management framework. It consists of storage which can be configured individually per iRODS instance; metadata database called iCAT to keep information on users, access rights, and additional information on files and folders; and a rule engine to implement and execute data management policies.

An iRODS instance is called an iRODS zone and is defined by a metadata database called the iCAT. This database contains all metadata on users, data objects (files), data collections (folders) and storage systems in the iRODS zone. Metadata which the system creates automatically are data size, checksums, last accession date and creation date. Moreover, users and systems can add own metadata structured as key-value-unit triples. This feature is used by B2SAFE to create links to replicas in other iRODS zones.

iRODS abstracts from actual storage system and location and provides a so-called logical path. This feature allows to replicate data in a unique way between iRODS zones without knowing anything about the configured storage media and is one of the concepts B2SAFE relies on.

The rule engine executes iRODS rules - iRODS implementations of low-level data policies. Rules can be called by command line (client side calls), they can be automatically invoked by a certain action in the iRODS system or can be executed on a regular basis (server side calls). B2SAFE implements its data policies as iRODS rules which can be combined to achieve the appropriate behaviour for a community (see Section [Example workflows](#)).

B2SAFE's replication mechanism

B2SAFE replicates data from one iRODS zone to another, i.e. copying data to another administrative domain. This lowers the risk of complete data loss, however it increases the need for proper management of replicas across zones.

The replication sites can be configured in iRODS itself as federated iRODS zones (reference to the [iRODS manual for Federations](#))

In a long replication chain not all iRODS zones are directly federated. That means that from the original site one does not have the means to check the integrity of replicas by using iRODS mechanisms (see Figure 1). Hence, there is the need for an external system to log the replication chain and provide some minimal information to ensure data integrity across sites.

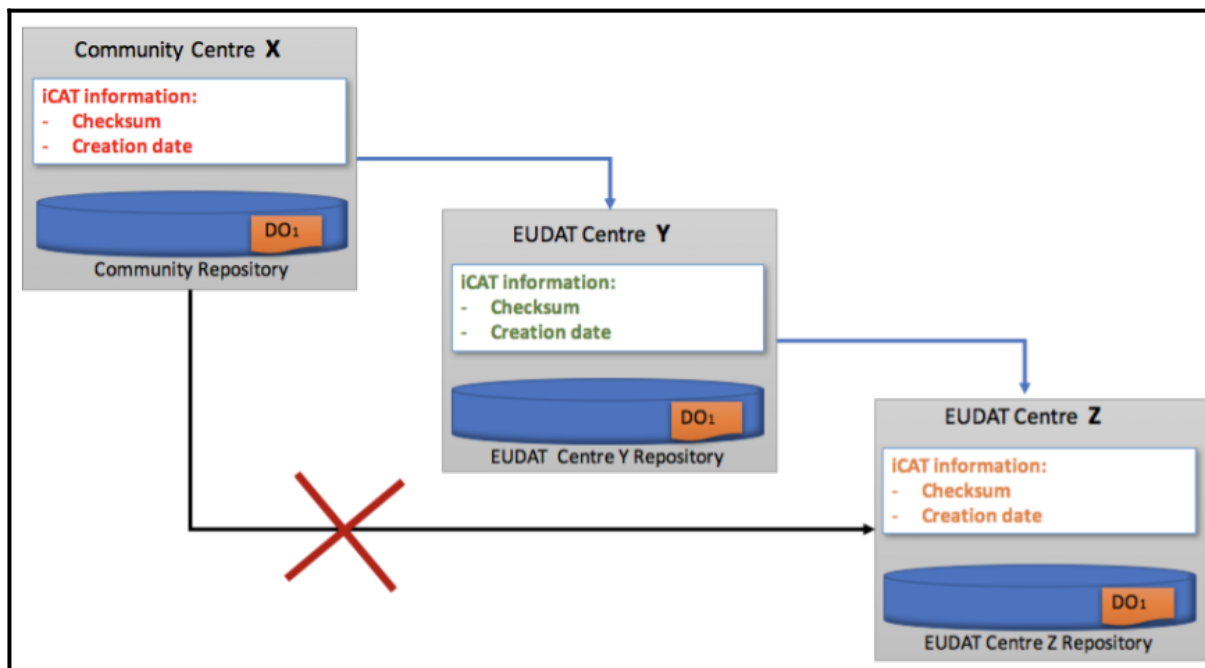


Figure 1: Replication across three iRODS zones; one community centre X replicates its repository data to an EUDAT centre Y, which in turn replicates the data to another EUDAT centre Z. By this three independent copies of the data are made. The blue arrows indicate direct access to the information in iRODS. That is to say the community centre can ensure the data's integrity in the EUDAT centre Y, while the EUDAT centre Y can ensure the data integrity with EUDAT centre Z. However, the community centre cannot verify the data's integrity with EUDAT centre Z by means of iRODS's mechanisms (black arrow).

Tracking replicas across sites

To track replicas and record the whole replication chain of data in B2SAFE an external service is needed. B2SAFE uses the [EUDAT persistent identifier service B2HANDLE](#).

In general PIDs are used to reliably identify and cite data objects throughout their lifecycle and they are thus a vital part of long-term data management. More specifically, B2SAFE employs PIDs and designed specific PID profile to reliably find and identify replicas.

A persistent identifier is an opaque string which usually is resolvable by the HTTP protocol and thus contains the mapping from the opaque string to a URL. Upon creation one can add more information to the PID. In the case of EUDAT's B2SAFE service PIDs of direct replicas, the direct parent's PID and a link (usually also a PID) to the very first data repository is added.

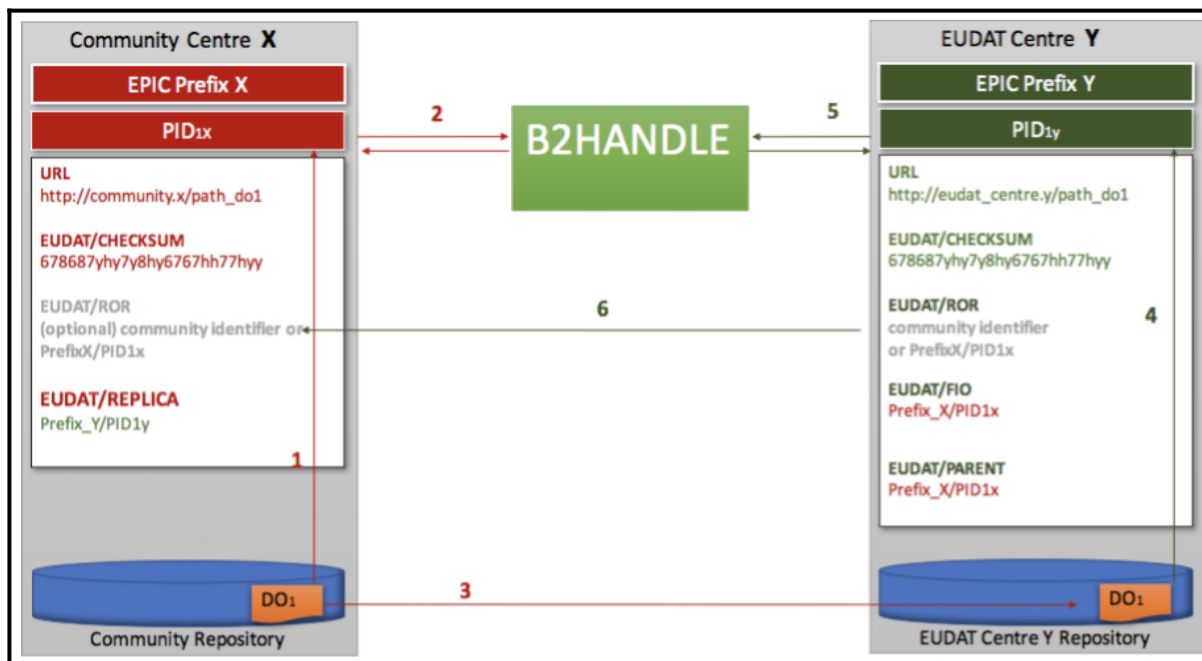


Figure 2: B2SAFE replication with creation of PIDs. Assume we are replicating data from a community centre to one other EUDAT centre. The figure shows the creation of PIDs and their additional information. 1 - The community calls the B2SAFE rule which creates a PID (opaque and unique string) for the data object (DO₁) and 2 - creates an entry in the PID system using B2HANDLE containing the additional information EUDAT/CHECKSUM and the identifier for the data in the community repository EUDAT/ROR. 3 - Subsequently the B2SAFE rule for replication is called which creates a copy of DO₁ at EUDAT Centre Y. 4 - The same rule registers the new data copy with a new PID and 5 - creates an entry in the PID system with the following information: EUDAT/CHECKSUM, the PID of the direct parent EUDAT/PARENT, the EUDAT/ROR and the PID to first EUDAT centre that holds the data (EUDAT/FIO). 6 - Finally the PID of the direct parent at Community Centre X is updated with the location of its replica (EUDAT/REPLICA).

The PID generated by B2SAFE and all information stored in the PID system is publicly accessible. Figure 2 describes how the B2SAFE module integrates the replication of data with the PID system.

With the integration of PIDs a data centre can now track all replicas across the replication chain as shown in Figure 3.

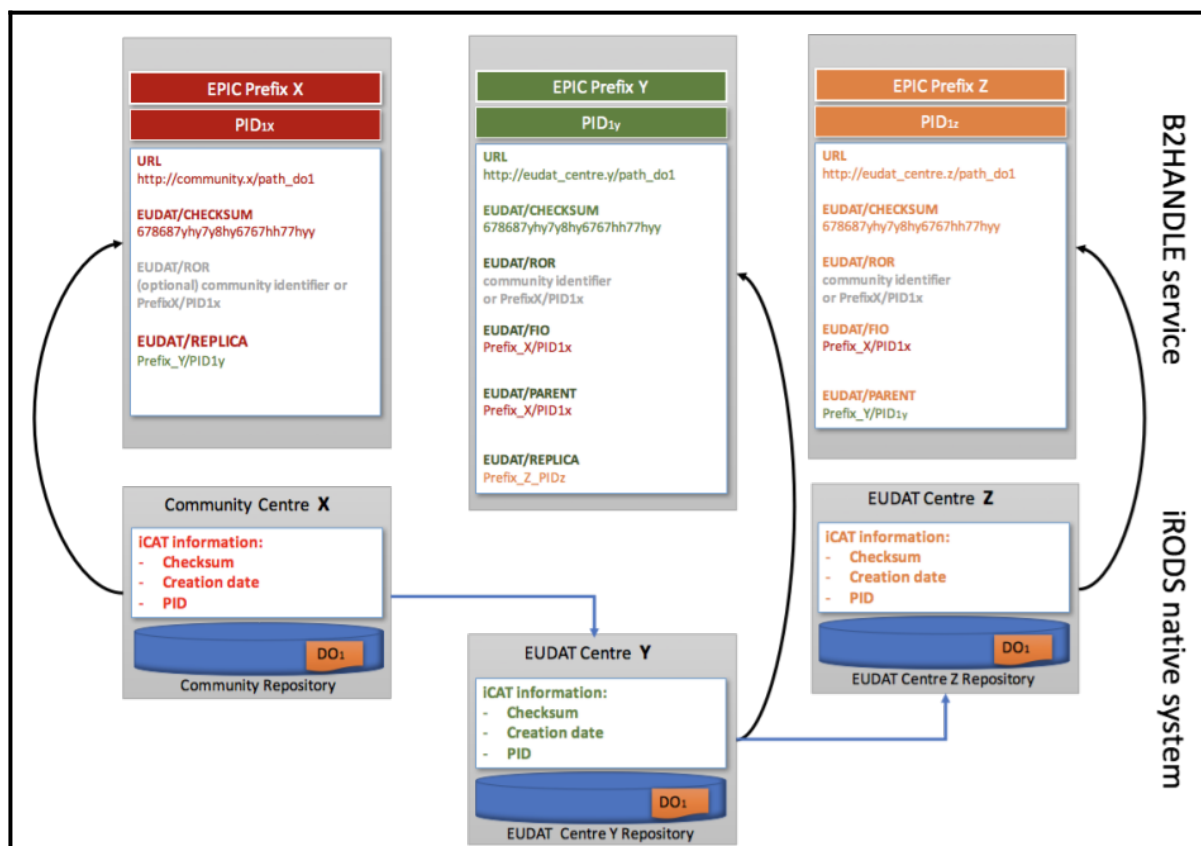


Figure 3: Integration of the PID system and iRODS. Blue arrows indicate data replication between iRODS zones, black arrows indicate PID registration. With the fields EUDAT/PARENT and EUDAT/REPLICA one can follow the full replication chain in the PID system and retrieve the actual location of data (URL field in the PID entry). In addition to the automatic metadata which is created in iRODS, B2SAFE also creates entries for the PID of the data itself, its parent and its replicas. Thus, having access to one data replica in the replication chain one can enter the PID system and retrieve all information to follow the replication chain.

Example workflows

The B2SAFE module is a set of iRODS rules which can be put together in workflows enabling data replication and PID management. In this section we describe several typical B2SAFE workflows. The full documentation of workflows and their respective code examples can be found [on the service's wiki](#).

PID creation

Creating PIDs for files and folders is essential to track replicas across different administrative domains. Figure 4 shows how the B2SAFE rule attaches a PID to either a data object or a collection. The rule *EUDATCreatePID* from the B2SAFE ruleset takes as input the iRODS logical path to the collection or data object which should receive a PID. The rule will then connect to the PID service (B2HANDLE) create a PID, create the PID and store it in the iRODS metadata catalogue. By this the link between iRODS and the PID is established.

At the same time metadata such as the checksum is stored in the PID system to enable cross-domain integrity

checks. Optionally one can provide a link to the original data (ROR) or the PID to the direct parent of a data object. This information is stored as iRODS metadata and in the newly created PID and by this establishes the upper connection in the replication chain.

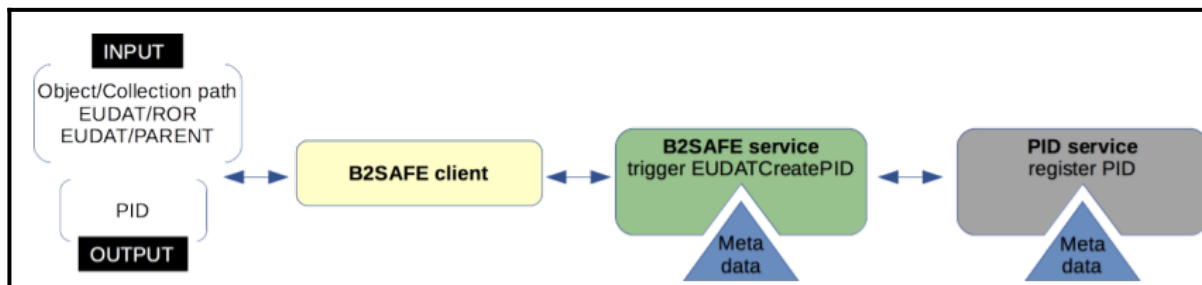


Figure 4: A B2SAFE client rule gathers the iRODS logical path and optionally a PID pointing to the original data (ROR) or the direct parent (PARENT) and propagates these to the B2SAFE EUDATCreatePID rule. In turn this rule establishes the connection to the PID service, creates the PID with respective metadata and stores the created PID as metadata in iRODS.

The rule *EUDATCreatePID* from the B2SAFE rulebase can be called by another rule or an event hook in the iRODS system itself. In both cases the input for the PID creation needs to be propagated to the pid creation rule.

B2SAFE replication

The replication according to the EUDAT policies is triggered by the rule *EUDATReplication* which is part of the B2SAFE ruleset. This rule steers the replication across iRODS zones. It takes as input parameters the iRODS path to the source and the destination object or collection.

One can execute and suppress the creation of PIDs (Figure 5, upper panel) or PIDs can be created synchronously with the replication (Figure 5, lower panel). In the first case there will be no link in the irods metadata database nor in the PID system to build the replication chain. The PID creation and thus the creation of metadata to build the replication chain is triggered by setting the flag *registered* to true.

Since the PID registration costs some time it can be advantageous to decouple the data replication and PID creation when transferring large collections of data. In such a case the replication rule needs to be combined with the *EUDATPIDRegistration* rule from the B2SAFE rulebase. This rule ensures that after data transfer, PIDs are created and the replication chain is built in both the iRODS metadata database and the PID system.

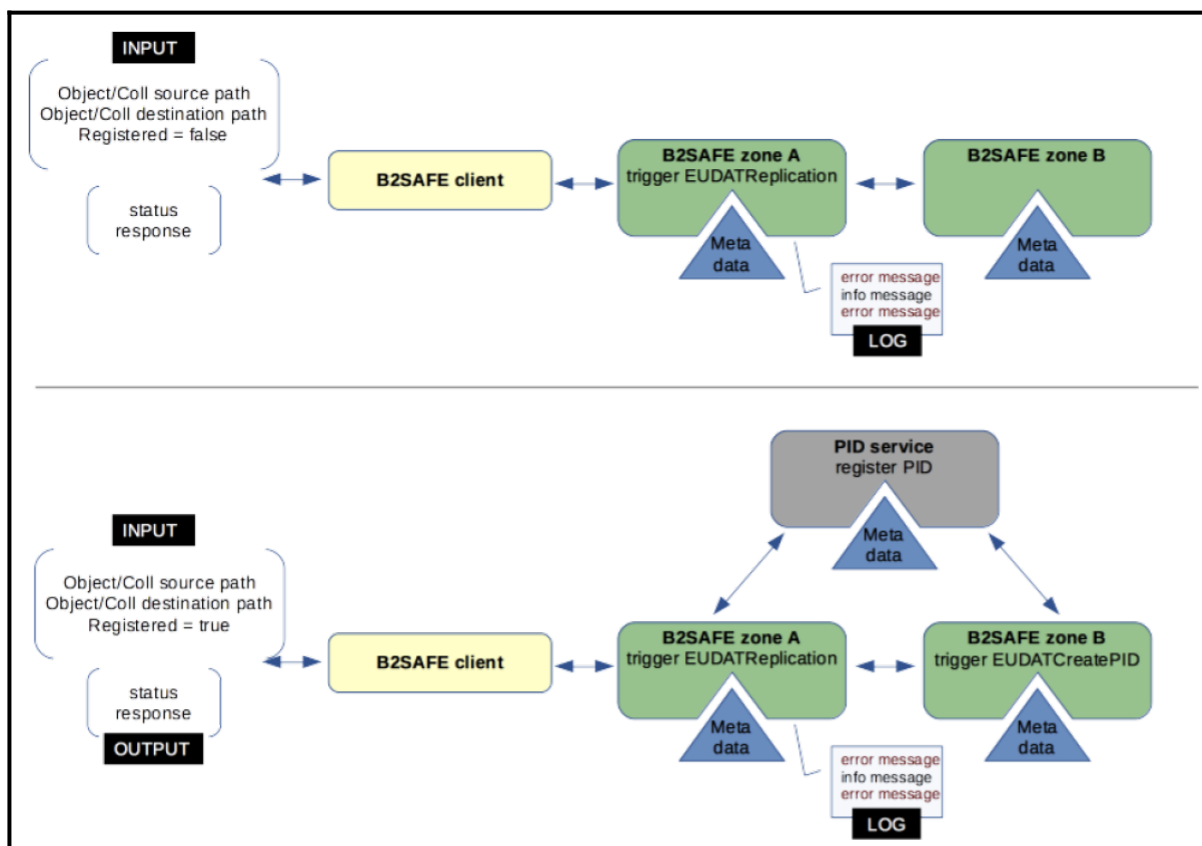


Figure 5: The EUDAT replication. Upper panel: The EUDATReplication rule steers the replication of data across iRODS zones. In this case only minimal information on the data is stored in the iRODS metadata database, the link between the original data and its replica is not introduced. The PID creation upon replication can be triggered by setting the flag 'registered' (see lower panel). Here PIDs are generated as soon as the data is replicated (synchronous PID registration) and the link between the original data and the replica is introduced in the iRODS metadata database and the PID system.

Other rules

The B2SAFE module offers also rules for integrity checks across zones, recovering failed transfers and updating the information on data location in the PID system in case of changing the iRODS path to the data. Furthermore, the ruleset contains experimental features like community metadata handling and messaging.

More information

EUDAT communities can deploy B2SAFE or let an EUDAT site run B2SAFE for them.

- If you want to run B2SAFE, please read the introduction how to [join the B2SAFE replication network](#) and follow our guide to [Configure B2SAFE](#).
- If you prefer that an EUDAT site run B2SAFE for your community, please follow our [Using B2SAFE](#) documentation.

For an extensive, modular, hands-on training course on B2SAFE, please see the [EUDAT training repository on](#)

EUDAT receives funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 654065.



[GitHub](#).

Support

Support for B2SAFE is available via the EUDAT ticketing system through the [webform](#).

If you have comments on this page, please submit them through the [EUDAT ticketing system](#).

Document Data

Version: 1.1

Authors:

Merret Buurman, buurman@dkrz.de

Claudio Cacciari, c.cacciari@cineca.it

Giovanni Morelli, g.morelli@cineca.it

Editors:

Kostas Kavoussanakis, k.kavoussanakis@epcc.ed.ac.uk

Christine Staiger, christine.staiger@surfsara.nl

[Read more](#)